



Data-Driven Computational Methods for Quasi-Stationary Distribution and Sensitivity Analysis

Yao Li¹ · Yaping Yuan¹

Received: 4 March 2021 / Revised: 4 January 2022 / Accepted: 25 January 2022

© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2022

Abstract

This paper studies computational methods for quasi-stationary distributions (QSDs). We first proposed a data-driven solver that solves Fokker–Planck equations for QSDs. Similar to the case of Fokker–Planck equations for invariant probability measures, we set up an optimization problem that minimizes the distance from a low-accuracy reference solution, under the constraint of satisfying the linear relation given by the discretized Fokker–Planck operator. Then we use coupling method to study the sensitivity of a QSD against either the change of boundary condition or the diffusion coefficient. The 1-Wasserstein distance between a QSD and the corresponding invariant probability measure can be quantitatively estimated. Some numerical results about both computation of QSDs and their sensitivity analysis are provided.

Keywords Quasi-stationary distribution · Monte Carlo simulation · Data-driven computation · Coupling method

1 Introduction

Many models in various applications are described by Markov chains with absorbing states. For example, any models with mass-action kinetics, such as ecological models, epidemic models, and chemical reaction models, are subject to the population-level randomness called the demographic stochasticity, which lead to extinction in finite time. There are also many dynamical systems that have interesting short term dynamics but trivial long term dynamics, such as dynamical systems with transient chaos [19]. A common way of capturing asymptotical properties of these transient dynamics is to find the quasi-stationary distribution (QSD),

Yao Li is partially supported by NSF DMS-1813246 and DMS-2108628.

✉ Yao Li
yaoli@math.umass.edu

Yaping Yuan
yuan@math.umass.edu

¹ Department of Mathematics and Statistics, University of Massachusetts Amherst, Amherst, MA 01002, USA

which is the conditional limiting distribution conditioning on not hitting the absorbing set yet. However, most QSDs do not have a closed form. So numerical solutions are necessary in various applications.

Computational methods for QSDs are not very well developed. Although the relation between QSD and the Fokker–Planck equation is well known, it is not easy to use classical PDE solver to solve QSDs because of the following two reasons. First a QSD is the eigenfunction of the Fokker–Planck operator whose eigenvalue is unknown. The cost of solving eigenfunction of a discretized Fokker–Planck operator is considerably high. Secondly the boundary condition of the Fokker–Planck equation is unknown. We usually have a mixture of unbounded domain and unknown boundary value at the absorbing set. There are some literatures about Monte Carlo sampling from a QSD, which usually include regenerating samples from an empirical distribution once hitting the boundary [4–7, 26]. However the efficiency of the Monte Carlo simulation is known to be low. To get the probability density function, one needs to deal with undesired noise associated to the Monte Carlo simulation. Methods like the kernel density estimator and the variable kernel density estimation can smooth the solution but also introduce undesired diffusions to the solution, especially when a QSD is highly concentrated at the vicinity of some low-dimensional sets.

The first goal of this paper is to extend the data-driven Fokker–Planck solver developed by the first author to the case of QSDs [10]. Similar to [10], we need a reference solution $\mathbf{v}$ generated by the Monte Carlo simulation. Then we discretize the Fokker–Planck operator in a numerical domain D without the boundary condition. Because of the lack of boundary conditions, the discretization only gives an underdetermined linear system, denoted by $\mathbf{A}\mathbf{u} = 0$. Then we minimize $\|\mathbf{v} - \mathbf{u}\|$ in the null space of $\mathbf{A}$. As shown in [11], this optimization problem projects the error terms of $\mathbf{v}$ to a low dimensional linear subspace, which significantly reduces its norm. Our numerical simulations show that this data-driven Fokker–Planck solver can tolerate very high level of spatially uncorrelated error, so the accuracy of $\mathbf{v}$ does not have to be high. The main difference between QSD solver and the Fokker–Planck solver is that we need a killing rate to find the QSD, which is obtained by a Monte Carlo simulation. We find that the QSD is not very sensitive against small error in the estimation of the killing rate.

The second goal of this paper is to study the sensitivity of QSDs. Some modifications of either the boundary condition or the model parameters can prevent the Markov process from hitting the absorbing state in finite time. So the modified process would admit an invariant probability measure instead of a QSD. It is important to understand the difference between the QSD of a Markov process and the invariant probability measure of its modification. For example, many ecological models do not consider demographic noise because the population size is large and the QSD is much harder to study. But would the demographic noise completely change the asymptotical dynamics? More generally, a QSD captures the transient dynamics of a stochastic differential equation. If we separate a domain from the global dynamics by imposing a reflecting boundary condition, how would the local dynamics be different from the corresponding transient dynamics? All these require some sensitivity analysis with quantitative bounds.

The way of sensitivity analysis is similar to [11]. We need both finite time error and the rate of contraction of the transition kernel of the Markov process. The finite time error is given by both the killing rate and the change of model parameters (if any). Both cases can be estimated easily. The rate of contraction is estimated by the data-driven method proposed in [21]. We design a suitable coupling scheme for the modified Markov process that admits an invariant probability measure. Because of the coupling inequality, the exponential tail of the coupling time can be used to estimate the rate of contraction. The sensitivity analysis is demonstrated in several numerical examples. We can find that the invariant probability

measure of the modified process is a better approximation of the QSD if (i) the speed of contraction is faster and (ii) the killing rate is lower.

The organization of this paper is as follows. A short preliminary about QSD, Fokker–Planck equation, simulation method, and coupling method is provided in Sect. 2. Section 3 introduces the data-driven QSD solver. The sensitivity analysis of QSD is given in Sect. 4. Section 5 is about numerical examples.

2 Preliminary

In this section, we provide some preliminaries to this paper, which are about the quasi-stationary distribution (QSD), the Fokker–Planck equation, the coupling method, and numerical simulations of the QSD.

2.1 Quasi-Stationary Distribution

We first define the QSD and the exponential killing rate λ of a Markov process with an absorbing state. Let $X = (X_t : t \geq 0)$ be a continuous-time Markov process taking values in a measurable space $(\mathcal{X}, \mathcal{B}(\mathcal{X}))$. Let $P^t(x, \cdot)$ be the transition kernel of X such that $P^t(x, A) = \mathbb{P}[X_t \in A \mid X_0 = x]$ for all $A \in \mathcal{B}$. Now assume that there exists an absorbing set $\partial\mathcal{X} \subset \mathcal{X}$. The complement $\mathcal{X}^a := \mathcal{X} \setminus \partial\mathcal{X}$ is the set of allowed states.

The process X_t is killed when it hits the absorbing set, implying that $X_t \in \partial\mathcal{X}$ for all $t > \tau$, where $\tau = \inf\{t > 0 : X_t \in \partial\mathcal{X}\}$ is the hitting time of set $\partial\mathcal{X}$. Throughout this paper, we assume that the process is almost surely killed in finite time, i.e. $\mathbb{P}[\tau < \infty] = 1$.

For the sake of simplicity let $\mathbb{P}_x$ (resp. $\mathbb{P}_\mu$) be the probability conditioning on the initial condition $x \in \mathcal{X}$ (resp. the initial distribution μ).

Definition 2.1 A probability measure μ on $\mathcal{X}^a$ is called a quasi-stationary distribution(QSD) of the Markov process X_t with an absorbing set $\partial\mathcal{X}$, if for every measurable set $C \subset \mathcal{X}^a$

$$\mathbb{P}_\mu[X_t \in C \mid \tau > t] = \mu(C), \quad t \geq 0, \quad (2.1)$$

If there is a probability measure μ exists such that

$$\lim_{t \rightarrow \infty} \mathbb{P}_x[X_t \in C \mid \tau > t] = \mu(C), \quad \forall x \in \mathcal{X}^a. \quad (2.2)$$

in which case we also say that μ is a quasi-limiting distribution(QLD).

Remark 2.1 When μ satisfies (2.2), it is called a quasi-limiting distribution, or a Yaglom limit. A QLD must be a QSD. Under some mild ergodicity assumptions, a QSD is also a QLD [9].

In the analysis of QSD, we are particularly interested in a parameter λ , called the killing rate of the Markov process. If the distribution of the killing time $\mathbb{P}_x(\tau > t)$ has an exponential tail, then λ is the rate of this exponential tail. The following theorem shows that the killing time is exponentially distributed when the process starts from a QSD[8].

Theorem 2.1 Let μ be a QSD of stochastic process X . Then

$$\exists \lambda = \lambda(\mu) \text{ such that } \mathbb{P}_\mu[\tau > t] = e^{-\lambda t}, \quad \forall t \geq 0,$$

where λ is called the killing rate of X .

Throughout this paper, we assume that X admits a QSD denoted by μ with a strictly positive killing rate λ .

2.2 Stochastic Differential Equations and the Fokker–Planck Equation

Consider a stochastic differential equation (SDE)

$$dX_t = f(X_t)dt + \sigma(X_t)dW_t, \quad (2.3)$$

where $X_t \in \mathbb{R}^d$ and X_t is killed when it hits the absorbing set $\partial\mathcal{X} \subset \mathbb{R}^n$; $f: \mathbb{R}^d \rightarrow \mathbb{R}^d$ is a continuous vector field; σ is an $d \times d$ matrix-valued function; and dW_t is the white noise in $\mathbb{R}^d$. The following well known theorem shows the existence and the uniqueness of the solution of Eq. (2.3).

Theorem 2.2 *Assume that there are two positive constants C_1 and C_2 such that the two functions f and σ in (2.3) satisfy*

(1) *(Lipschitz condition) for all $x, y \in \mathbb{R}^d$ and t*

$$|f(x) - f(y)|^2 + |\sigma(x) - \sigma(y)|^2 \leq C_1|x - y|^2;$$

(2) *(Linear growth condition) for all $x, y \in \mathbb{R}^d$ and t*

$$|f(x)|^2 + |\sigma(x)|^2 \leq C_2(1 + |x|^2).$$

Then there exists a unique solution $X(t)$ to Eq. (2.3).

There are quite a few recent results about the existence and convergence of QSD. Since the theme of this paper is numerical computations, in this paper we directly assume that X_t admits a unique QSD μ on set $\mathcal{X}^a$ that is also the quasi-limit distribution. The detailed conditions are referred to in [12, 17, 22, 24, 25].

Let u be the probability density function of μ . We refer [2] for the fact that u satisfies

$$-\lambda u = \mathcal{L}u = -\sum_{i=1}^d (f_i u)_{x_i} + \frac{1}{2} \sum_{i,j=1}^d (D_{ij}u)_{x_i x_j}, \quad (2.4)$$

where $D = \sigma^T \sigma$, and λ is the killing rate.

2.3 Simulation Algorithm of QSD

It remains to review the simulation algorithm for QSDs. In order to compute the QSD, one needs to numerically simulate a long trajectory of X . Once X_t hits the absorbing state, a new initial value is sampled from the empirical QSD. The re-sampling step can be done in two different ways. We can either use many independent trajectories that form an empirical distribution [3] or re-sample from the history of a long trajectory [5]. In this paper we use the latter approach.

Let $\hat{X} = \{\hat{X}_n^\delta, n \in \mathbb{Z}_+\}$ be a stochastic process that samples the “numerical QSD”. When $\hat{X}_n^\delta \in \mathcal{X}^a$, $\hat{X}$ simply approximates the time- δ sample chain of X_t , i.e., $X_{n\delta}$. The approximation uses either Euler–Maruyama scheme or the Milstein scheme. The Euler–Maruyama numerical scheme is given by

$$\hat{X}_{n+1}^\delta = \hat{X}_n^\delta + f(\hat{X}_n^\delta)\delta + \sigma(\hat{X}_n^\delta)(W_{(n+1)\delta} - W_{n\delta}), \quad (2.5)$$

where $\hat{X}_0^\delta = X_0$, $W_{(n+1)\delta} - W_{n\delta} \sim \mathcal{N}(0, \delta \text{Id}_d)$, $n \in \mathbb{Z}_+$ is a d -dimensional normal random variable. A more accurate scheme is the Milstein scheme, which reads

$$\hat{X}_{n+1}^\delta = \hat{X}_n^\delta + f(\hat{X}_n^\delta)\delta + \sigma(\hat{X}_n^\delta)(W_{(n+1)\delta} - W_{n\delta}) + \sigma(\hat{X}_n^\delta)IL,$$

where I is a $d \times d$ matrix with its (i, j) -th component being the double Itô integral

$$I_{i,j} = \int_{n\delta}^{(n+1)\delta} \int_{n\delta}^{s_2} dW^i(s_1) dW^j(s_2),$$

and $L \in \mathbb{R}^d$ is a vector of left operators with i -th component

$$u\mathcal{L}_i = \sum_{j=1}^d \sigma_{i,j}(\hat{X}_n^\delta) \frac{\partial u}{\partial x_j}.$$

Under suitable assumptions of Lipschitz continuity and linear growth conditions for f and σ , the Euler–Maruyama approximation provides a convergence rate of order $1/2$, while the Milstein scheme is an order 1 strong approximation [18].

To deal with the situation when $\hat{X}$ hits the absorbing set, in addition to $\hat{X}_n^\delta$, we also need to update a temporal occupation measure

$$\mu_n = \frac{1}{n} \sum_{k=0}^{n-1} \delta_{\hat{X}_k^\delta}.$$

If the numerical scheme gives $\hat{X}_{n+1}^\delta \in \partial\mathcal{X}$, we immediately resample $\hat{X}_{n+1}^\delta$ from μ_n . More precisely, let the transition kernel of the numerical scheme of X (without resampling) be $\hat{Q}$. Then $\hat{Q}$ has an absorbing set, i.e., $\hat{Q}(\partial\mathcal{X}, \partial\mathcal{X}) = 1$. The transition kernel of $\hat{X}_n^\delta$ is modified from $\hat{Q}$ such that

$$\mathbb{P}[\hat{X}_{n+1}^\delta \in A \mid \hat{X}_n^\delta = x] = \hat{Q}(x, A) + \hat{Q}(x, \partial\mathcal{X})\mu_n(A).$$

We have the following convergence result from [5].

Proposition 2.1 (Theorem 2.5 in [5]) *Under suitable assumptions about $\hat{X}$, the occupation measure μ_n converges to the QSD μ as $n \rightarrow \infty$.*

The assumption in Proposition 2.1 about $\hat{X}$ is that $\hat{X}$ must be Feller, the absorbing set is accessible, $\hat{X}$ admits a “weak small set”, and that the killing rates from different initial values are uniformly controlled. It is easy to check that these assumptions are satisfied by the numerical scheme of most ergodic SDEs. Since this paper focuses on numerical algorithm, throughout this paper, we assume that all assumptions in Proposition 2.1 are satisfied.

For simplicity, we introduce the algorithm for $n = 2$, specifically, we solve u in Eq. (2.4) numerically on a 2D domain $D = [a_0, b_0] \times [a_1, b_1]$. Firstly, we construct an $N \times M$ grid on D with grid size $h = \frac{b_0 - a_0}{N} = \frac{b_1 - a_1}{M}$. Each small box in the mesh is denoted by $O_{i,j} = [a_0 + (i-1)h, a_0 + ih] \times [a_1 + (j-1)h, a_1 + jh]$. Let $\mathbf{u} = \{u_{i,j}\}_{i=1,j=1}^{i=N,j=M}$ be the numerical solution on D that we are interested in, then $\mathbf{u}$ can be considered as a vector in $\mathbb{R}^{N \times M}$. Each element $u_{i,j}$ approximates the density function u at the center of each $O_{i,j}$, with coordinate $(ih + a_0 - h/2, jh + a_1 - h/2)$. Generally speaking, we count the number of $\{\hat{X}_n^\delta\}$ falling into each box and set the normalized value as the approximate probability density at $O_{i,j}$. The detail of the simulation is shown in Algorithm 1 below.

Sometimes the Euler–Maruyama method underestimates the probability that X moves to the absorbing set, especially when X_t is close to $\partial\mathcal{X}$. This problem can be partially fixed

Algorithm 1 Monte Carlo for QSD

Input: Equation (2.5) and the grid

Output: A Monte Carlo approximation $\mathbf{u} = \{u_{i,j}\}$. Sample size N_s .

Pick any initial value $X_0 \notin \partial\mathcal{X}$ in D

for $n = 1$ to N_s **do**

Use $\hat{X}_n^\delta$ and Eq. (2.5) to compute $\hat{X}_{n+1}^\delta$

Record the coordinates of the small box $O_{i,j}$ where $\hat{X}_{n+1}^\delta$ stands, say i^*, j^*

if $\hat{X}_{n+1}^\delta \notin \partial\mathcal{X}$ **then**

$u_{i^*,j^*} \leftarrow u_{i^*,j^*} + 1$

else

$\hat{X}_{n+1}^\delta = \hat{X}_{[U * n]}^\delta$, where U is a uniformly distributed random variable

end if

end for

Return $u_{i,j}/N_s h^2$ for all i, j as the approximation solution.

by introducing the Brownian bridge correction. We refer to [5] for details. For a sample of $\hat{X}_n^\delta$ falling into a small box which are closed to $\partial\mathcal{X}$, the probability of the continuous-time process falling into the trap $\partial\mathcal{X}$ before the discrete one is relatively high. In fact, this probability is exponentially distributed and the rate is related to the distance from $\partial\mathcal{X}$. Let $B_t^T = W_t - \frac{t}{T} W_T$ be the Brownian Bridge on the interval $[0, T]$. In the 1D case, the law of the infimum and the supremum of the Brownian Bridge can be computed as follows: for every $z \geq \max(x, y)$

$$\mathbb{P}[\sup_{t \in [0, T]} (x + (y - x) \frac{t}{T} + \phi B_t^T) \leq z] = 1 - \exp(-\frac{2}{T\phi^2} (z - x)(z - y)), \quad (2.6)$$

where $x = \hat{X}_n^\delta \in \mathcal{X}^a$, $y = \hat{X}_{n+1}^\delta \in \mathcal{X}^a$, and $\phi = \sigma(\hat{X}_n^\delta)$ is the strength coefficient of Brownian Bridge. This means that at each step n , if $\hat{X}_{n+1}^\delta \in \mathcal{X}^a$, one can compute, with the help of Eq. (2.6), a Bernoulli random variable G with the parameter

$p = \mathbb{P}[\exists t \in (n\delta, (n+1)\delta), \hat{X}_t \in \partial\mathcal{X} | x = \hat{X}_n^\delta, y = \hat{X}_{n+1}^\delta]$ (If $G = 1$, the process is killed).

2.4 Coupling Method

The coupling method is used for the sensitivity analysis of QSDs.

Definition 2.2 (*Coupling of probability measures*) Let μ and ν be two probability measures on a measurable space $(\mathcal{X}, \mathcal{B}(\mathcal{X}))$. A probability measure γ on $(\mathcal{X} \times \mathcal{X}, \mathcal{B}(\mathcal{X}) \times \mathcal{B}(\mathcal{X}))$ is called a coupling of μ and ν , if two marginals of γ coincide with μ and ν respectively.

The definition of coupling can be extended to any two random variables that take value in the same state space. Now consider two Markov processes $X = (X_t : t \geq 0)$ and $Y = (Y_t : t \geq 0)$ with the same transition kernel P . A coupling of X and Y is a stochastic process (X^m, Y^m) on the product state space $\mathcal{X} \times \mathcal{X}$ such that

- (i) The marginal processes X^m and Y^m are Markov processes with the transition kernel P ;
- (ii) If $X_s^m = Y_s^m$, we have $X_t^m = Y_t^m$ for all $t > s$.

The first meeting time of X_t^m and Y_t^m is denoted as $\tau^c := \inf_{t \geq 0} \{X_t^m = Y_t^m\}$, which is called the coupling time. The coupling (X^m, Y^m) is said to be successful if the coupling time is almost surely finite, i.e. $\mathbb{P}[\tau^c < \infty] = 1$. Here the super index m stands for the marginal distribution, which is dropped when it causes no confusion.

In order to give an estimate of the sensitivity of the QSD, we need the following two metrics.

Definition 2.3 (*Wasserstein distance*) Let d be a metric on the metric state space V equipped with distance $d(\cdot, \cdot)$. For probability measures μ and ν on V , the Wasserstein distance between μ and ν for d is given by

$$\begin{aligned} d_w(\mu, \nu) &= \inf\{\mathbb{E}_\gamma[d(x, y)] : \gamma \text{ is a coupling of } \mu \text{ and } \nu.\} \\ &= \inf\left\{\int d(x, y)\gamma(dx, dy) : \gamma \text{ is a coupling of } \mu \text{ and } \nu.\right\} \end{aligned} \quad (2.7)$$

In this paper, without further specification, we assume that the 1-Wasserstein distance is induced by $d(x, y) = \max\{1, \|x - y\|\}$, where $\|x - y\|$ is the Euclidean norm.

Definition 2.4 (*Total variation distance*) Let μ and ν be probability measures on $(\mathcal{X}, \mathcal{B}(\mathcal{X}))$. The total variation distance of μ and ν is

$$\|\mu - \nu\|_{TV} := \sup_{C \in \mathcal{B}(\mathcal{X})} |\mu(C) - \nu(C)|.$$

Lemma 2.1 (*Coupling inequality*) For the coupling given above and the Wasserstein distance induced by the distance given in (2.7), we have

$$\mathbb{P}[\tau^c > t] = \mathbb{P}[X_t^m \neq Y_t^m] \geq d_w(P^t(x, \cdot), P^t(y, \cdot)).$$

Proof By the definition of Wasserstein distance,

$$\begin{aligned} d_w(P^t(x, \cdot), P^t(y, \cdot)) &\leq \int d(x, y)\mathbb{P}[(X_t^m, Y_t^m) \in (dx, dy)] \\ &= \int_{x \neq y} d(x, y)\mathbb{P}[(X_t^m, Y_t^m) \in (dx, dy)] \\ &\leq \int_{x \neq y} \mathbb{P}[(X_t^m, Y_t^m) \in (dx, dy)] \\ &= \mathbb{P}[X_t^m \neq Y_t^m]. \end{aligned}$$

The lemma follows because $\mathbb{P}[\tau^c > t] = \mathbb{P}[X_t^m \neq Y_t^m]$ by definition. $\square$

Consider a Markov coupling $(\hat{X}, \hat{Y})$, where $\hat{X} = \{\hat{X}_n^\delta : n \in \mathbb{N}\}$ and $\hat{Y} = \{\hat{Y}_n^\delta : n \in \mathbb{N}\}$ are two numerical trajectories of the stochastic differential equation described in (2.5). Theoretically, there are many ways to make stochastic differential equations couple. But since numerical computation always has errors, two numerical trajectories may miss each other when the true trajectories couple. Hence we need to apply a mixture of the following coupling methods in practice.

Independent coupling Independent coupling means the noise term in the two marginal processes $\hat{X}$ and $\hat{Y}$ are independent when running the coupling process $(\hat{X}, \hat{Y})$. That is

$$\begin{aligned} \hat{X}_{n+1}^\delta &= \hat{X}_n^\delta + f(\hat{X}_n^\delta)\delta + (W_{(n+1)\delta}^{(1)} - W_{n\delta}^{(1)}) \\ \hat{Y}_{n+1}^\delta &= \hat{Y}_n^\delta + f(\hat{Y}_n^\delta)\delta + (W_{(n+1)\delta}^{(2)} - W_{n\delta}^{(2)}), \end{aligned}$$

where $(W_{(n+1)\delta}^{(1)} - W_{n\delta}^{(1)})$ and $(W_{(n+1)\delta}^{(2)} - W_{n\delta}^{(2)})$ are independent random variables for each n .

Reflection coupling Two Wiener processes meet less often than the 1D case when the state space has higher dimensions. This fact makes the independent coupling less effective. The reflection coupling is introduced to avoid this case. Take the Euler–Maruyama scheme of the SDE

$$dX_t = f(X_t)dt + \sigma dW_t$$

as an example, where σ is a constant matrix. The Euler–Maruyama scheme of X_t reads as

$$\hat{X}_{n+1}^\delta = \hat{X}_n^\delta + f(\hat{X}_n^\delta)\delta + \sigma(W_{(n+1)\delta} - W_{n\delta}),$$

where W is a standard Wiener process. The reflection coupling means that we run $\hat{X}_n^\delta$ as

$$\hat{X}_{n+1}^\delta = \hat{X}_n^\delta + f(\hat{X}_n^\delta)\delta + \sigma(W_{(n+1)\delta} - W_{n\delta}),$$

while run $\hat{Y}_n^\delta$ as

$$\hat{Y}_{n+1}^\delta = \hat{Y}_n^\delta + f(\hat{Y}_n^\delta)\delta + \sigma P(W_{(n+1)\delta} - W_{n\delta}),$$

where $P = I - 2e_n e_n^T$ is a projection matrix with

$$e_n = \frac{\sigma^{-1}(\hat{X}_n^\delta - \hat{Y}_n^\delta)}{\|\sigma^{-1}(\hat{X}_n^\delta - \hat{Y}_n^\delta)\|}.$$

Nontechnically, reflecting coupling means that the noise term is reflected against the hyperplane that orthogonally passes the midpoint of the line segment connecting $\hat{X}_n^\delta$ and $\hat{Y}_n^\delta$. In particular, $e_n = -1$ when the state space is 1D.

Maximal coupling Above coupling schemes can bring $\hat{X}_n^\delta$ moves close to $\hat{Y}_n^\delta$ when running numerical simulations. However, a mechanism is required to make $\hat{X}_{n+1}^\delta = \hat{Y}_{n+1}^\delta$ with certain positive probability. That's why the maximal coupling is involved. One can couple two trajectories whenever the probability distributions of their next step have enough overlap. Denote $p^{(x)}(z)$ and $p^{(y)}(z)$ as the probability density functions of $\hat{X}_{n+1}^\delta$ and $\hat{Y}_{n+1}^\delta$ respectively. The implementation of the maximal coupling is described in the following algorithm that is first proposed by [15].

Algorithm 2 Maximal coupling

Input: $\hat{X}_n^\delta$ and $\hat{Y}_n^\delta$

Output: $\hat{X}_{n+1}^\delta$ and $\hat{Y}_{n+1}^\delta$, and τ^c if coupled

 Compute probability density functions $p^{(x)}(z)$ and $p^{(y)}(z)$

 Sample $\hat{X}_{n+1}^\delta$ and calculate $r = Up^{(x)}(\hat{X}_{n+1}^\delta)$, where U is uniformly distributed on $[0,1]$

if $r < p^{(y)}(\hat{X}_{n+1}^\delta)$ **then**

$\hat{Y}_{n+1}^\delta = \hat{X}_{n+1}^\delta$, $\tau^c = (n+1)\delta$

else

 Sample $\hat{Y}_{n+1}^\delta$ and calculate $r' = Vp^{(y)}(\hat{Y}_{n+1}^\delta)$, where V is uniformly distributed on $[0,1]$

while $r' < p^{(x)}(\hat{Y}_{n+1}^\delta)$ **do**

 Resample $\hat{Y}_{n+1}^\delta$ and V . Recalculate $r' = Vp^{(y)}(\hat{Y}_{n+1}^\delta)$

end while

τ^c is still undetermined

end if

For discrete-time numerical schemes of SDEs, we use reflection coupling when $\hat{X}_n^\delta$ and $\hat{Y}_n^\delta$ are far away from each other, and maximal coupling when they are sufficiently close. The

threshold of changing coupling method is $2\sqrt{\delta}\|\sigma\|$ in our simulation, that is, the maximal coupling is applied when the distance between $\hat{X}_n^\delta$ and $\hat{Y}_n^\delta$ is smaller than the threshold.

3 Data-Driven Solver for QSD

Recall that QSD solves the equation $-\lambda u = \mathcal{L}u$ where $\mathcal{L}$ is the Fokker–Planck operator defined in (2.4). The QSD solver consists of three components: an estimator of the killing rate λ , a Monte Carlo simulator of QSD that produces a reference solution, and an optimization problem similar to in [20].

3.1 Estimation of λ

Let $\hat{X} = \{\hat{X}_n^\delta, n \in \mathbb{Z}_+\}$ be a long numerical trajectory of X_t as described in Algorithm 1. Let $\tau = \{\tau_m\}_{m=0}^M$ be recordings of killing times of the numerical trajectory such that X_t hits $\partial\mathcal{X}$ at $\tau_0, \tau_0 + \tau_1, \tau_0 + \tau_1 + \tau_2, \dots$ when running Algorithm 1. Note that τ is an 1D vector and each element in τ is a sample of the killing time. It is well known that if the QSD μ exists for a Markov process, then there exists a constant $\lambda > 0$ such that

$$\mathbb{P}_\mu[\tau > t] = e^{-\lambda t}.$$

Recall that the killing times τ is exponentially distributed and the rate can be approximated by

$$\lambda = \frac{1}{\text{mean of } \tau}.$$

One pitfall of the previous approach is that Algorithm 1 only gives a QSD when the time approaches to infinity. It is possible that τ has not converged close enough to the desired exponential distribution. So it remains to check whether the limit is achieved. Our approach is to check the exponential tail in a log-linear plot. After having τ , it is easy to choose a sequence of times $t_0, t_1, \dots, t_n$ and calculate $n_i = |\{\tau_m > t_i \mid 0 \leq m \leq M\}|$ for each $i = 0, \dots, n$. Then $p_i = n_i/M$ is an estimator of $\mathbb{P}_\mu[\tau > t_i]$. Now let p_i^u (resp. p_i^l) be the upper (resp. lower) bound of the confidence interval of p_i such that

$$p_i^u = \tilde{p} + z\sqrt{\frac{\tilde{p}}{\tilde{n}_i}(1 - \tilde{p})} \quad (\text{resp. } p_i^l = \tilde{p} - z\sqrt{\frac{\tilde{p}}{\tilde{n}_i}(1 - \tilde{p})}),$$

where $z = 1.96$, $\tilde{n}_i = n_i + z^2$ and $\tilde{p} = \frac{1}{\tilde{n}_i}(n_i + \frac{z^2}{2})$ [1]. If $p_i^l \leq e^{-\lambda t_i} \leq p_i^u$ for each $0 \leq i \leq n$, we accept the estimate λ . Otherwise we need to run Algorithm 1 for longer time to eliminate the initial bias in τ .

3.2 Data Driven QSD Solver

The data driven solver for the Fokker–Planck equation introduced in [20] can be modified to solve the QSD for the stochastic differential Eq. (2.3). We use the same 2D setting in Sect. 2.3 to introduce the algorithm. Let the domain D and the boxes $\{O_{i,j}\}_{i=1,j=1}^{i=N,j=M}$ be the same as defined in Sect. 2.3. Let $\mathbf{u}$ be a vector in $\mathbb{R}^{N \times M}$ such that u_{ij} approximates the probability density function at the center of the box $O_{i,j}$. As introduced in [10], we consider $\mathbf{u}$ as the solution to the following linear system given by the spatial discretization of the

Fokker–Planck Eq. (2.4) with respect to each center point:

$$\mathbf{A}_0 \mathbf{u} = \lambda \mathbf{u}, \quad (3.1)$$

where $\mathbf{A}_0$ is an $(N-2)(M-2) \times (NM)$ matrix, which is called the discretized Fokker–Planck operator, and λ is the killing rate, which can be obtained by the way we mentioned in previous subsection. More precisely, each row in $\mathbf{A}_0$ describes the finite difference scheme of Eq. (2.4) with respect to a non-boundary point in the domain D .

Similar to [20], we need the Monte Carlo simulation to produce a reference solution $\mathbf{v}$, which can be obtained via **Algorithm 1** in Sect. 2. Let $\{\hat{X}_n^\delta\}_{n=1}^N$ be a long trajectory of the numerical process $\hat{X}$ produced by **Algorithm 1**, and let $\mathbf{v} = \{v_{i,j}\}_{i=1,j=1}^{i=N,j=M}$ such that

$$v_{i,j} = \frac{1}{Nh^2} \sum_{n=1}^N \mathbf{1}_{O_{i,j}}(\hat{X}_n^\delta)$$

It follows from the convergence result in Proposition 2.1 that $\mathbf{v}$ is an approximate solution of Eq. (2.4) when the trajectory is sufficiently long. However, as discussed in [20], the trajectory needs to be extremely long to make $\mathbf{v}$ accurate enough. Noting that the error term of $\mathbf{v}$ has little spatial correlation, we use the following optimization problem to improve the accuracy of the solution.

$$\begin{aligned} & \min \|\mathbf{u} - \mathbf{v}\|^2 \\ & \text{subject to } \mathbf{A}_0 \mathbf{u} = \lambda \mathbf{u}. \end{aligned} \quad (3.2)$$

The solution to the optimization problem (3.2) is called the least norm solution, which satisfies $\mathbf{u} = \mathbf{v} - \mathbf{A}^T(\mathbf{A}\mathbf{A}^T)^{-1}(\mathbf{A}\mathbf{v})$, with $\mathbf{A} = \mathbf{A}_0 - \lambda \mathbf{I}$. [20]

An important method called the block data-driven solver is introduced in [10], in order to reduce the scale of numerical linear algebra problem in high dimensional problems. By dividing domain D into $K \times L$ blocks $\{D_{k,l}\}_{k=1,l=1}^{k=K,l=L}$ and discretizing the Fokker–Planck equation, the linear constraint on $D_{k,l}$ is

$$\mathbf{A}_{k,l} \mathbf{u}^{k,l} = \lambda \mathbf{u}^{k,l},$$

where $\mathbf{A}_{k,l}$ is an $(N/K-2)(M/L-2) \times (NM/KL)$ matrix. The optimization problem on $D_{k,l}$ is

$$\mathbf{u}_{k,l} = -\mathbf{A}_{k,l}^T(\mathbf{A}_{k,l}\mathbf{A}_{k,l}^T)^{-1}\mathbf{A}_{k,l}\mathbf{v}_{k,l} + \mathbf{v}_{k,l},$$

where $\mathbf{v}^{k,l}$ is a reference solution obtained from the Monte-Carlo simulation. Then the numerical solution to Fokker–Planck Eq. (3.1) is collage of all $\{u_{k,l}\}_{k=1,l=1}^{k=K,l=L}$ on all blocks. However, the optimization problem “pushes” most error terms to the boundary of domain, which makes the solution less accurate near the boundary of each block. Paper [10] introduced the overlapping block method and the shifting blocks method to reduce the interface error. The overlapping block method enlarges the blocks and set the interior solution restricted to the original block as new solution, while the shifting block method moves the interface to the interior by shifting all blocks and recalculate the solution. We refer [10] for full details of the shifting block method.

Note that in Sect. 3.1, we assume that λ is a pre-determined value given by the Monte Carlo simulation. Theoretically one can also search for the minimum of $\|\mathbf{u} - \mathbf{v}\|^2$ with respect to both λ and $\mathbf{v}$. But empirically we find that the result is not as accurate as using the killing rate λ from the Monte Carlo simulation, possibly because $\mathbf{v}$ has too much error.

One natural question is that how the simulation error in λ would affect the solution $\mathbf{u}$ to the optimization problem (3.2). Some linear algebraic calculation shows that the optimization problem (3.2) is fairly robust against small change of λ .

Theorem 3.1 Let $\mathbf{u}$ and $\mathbf{u}_1$ be the solution to the optimization problem (3.2) with respect to killing rates λ and λ_1 respectively, where $|\lambda - \lambda_1| = \epsilon \ll 1$. Then

$$\|\mathbf{u} - \mathbf{u}_1\| \leq 2s_{\min}^{-1}\epsilon\|\mathbf{v}\| + O(\epsilon^2),$$

where $s_{\min}$ is the smallest singular value of $\mathbf{A}$.

Proof Let $\mathbf{E} = [\epsilon \mathbf{I}_{(N-2)(M-2)} | \mathbf{0}]$ be an $(N-2)(M-2) \times (NM)$ perturbation matrix. Let $\tilde{\mathbf{A}} = \mathbf{A} + \mathbf{E}$, $\mathbf{B} = \mathbf{A}^T(\mathbf{A}\mathbf{A}^T)^{-1}\mathbf{A}$ and $\tilde{\mathbf{B}} = \tilde{\mathbf{A}}^T(\tilde{\mathbf{A}}\tilde{\mathbf{A}}^T)^{-1}\tilde{\mathbf{A}}$. Since $\mathbf{u} = \mathbf{v} - \mathbf{B}\mathbf{v}$ and $\mathbf{u}_1 = \mathbf{v} - \tilde{\mathbf{B}}\mathbf{v}$, it is sufficient to prove

$$\|\mathbf{B} - \tilde{\mathbf{B}}\| \leq 2s_{\min}^{-1}\epsilon + O(\epsilon^2).$$

Note that

$$\tilde{\mathbf{A}}\tilde{\mathbf{A}}^T = (\mathbf{A} + \mathbf{E})(\mathbf{A} + \mathbf{E})^T = \mathbf{A}\mathbf{A}^T + \mathbf{A}\mathbf{E}^T + \mathbf{E}\mathbf{A}^T + \mathbf{E}\mathbf{E}^T.$$

Since $\|\mathbf{E}\mathbf{E}^T\|$ is $O(\epsilon^2)$, we can neglect it when we consider the inverse matrix of $\tilde{\mathbf{A}}\tilde{\mathbf{A}}^T$. This means

$$(\tilde{\mathbf{A}}\tilde{\mathbf{A}}^T)^{-1} \approx (\mathbf{A}\mathbf{A}^T)^{-1} - (\mathbf{A}\mathbf{A}^T)^{-1}(\mathbf{A}\mathbf{E}^T + \mathbf{E}\mathbf{A}^T)(\mathbf{A}\mathbf{A}^T)^{-1}.$$

Without considering the high order term $O(\epsilon^2)$, we can see

$$\begin{aligned} \tilde{\mathbf{A}}^T(\tilde{\mathbf{A}}\tilde{\mathbf{A}}^T)^{-1}\tilde{\mathbf{A}} &\approx (\mathbf{A}^T + \mathbf{E}^T)((\mathbf{A}\mathbf{A}^T)^{-1} - (\mathbf{A}\mathbf{A}^T)^{-1}(\mathbf{A}\mathbf{E}^T + \mathbf{E}\mathbf{A}^T)(\mathbf{A}\mathbf{A}^T)^{-1})(\mathbf{A} + \mathbf{E}) \\ &= \mathbf{A}^T(\mathbf{A}\mathbf{A}^T)^{-1}\mathbf{A} + (\mathbf{E}^T(\mathbf{A}\mathbf{A}^T)^{-1}\mathbf{A} - \mathbf{A}^T(\mathbf{A}\mathbf{A}^T)^{-1}(\mathbf{A}\mathbf{E}^T)(\mathbf{A}\mathbf{A}^T)^{-1}\mathbf{A}) \\ &\quad + (\mathbf{A}^T(\mathbf{A}\mathbf{A}^T)^{-1}\mathbf{E} - \mathbf{A}^T(\mathbf{A}\mathbf{A}^T)^{-1}(\mathbf{E}\mathbf{A}^T)(\mathbf{A}\mathbf{A}^T)^{-1}\mathbf{A}) \\ &= \mathbf{B} + [\mathbf{E}^T - \mathbf{A}^T(\mathbf{A}\mathbf{A}^T)^{-1}(\mathbf{A}\mathbf{E}^T)](\mathbf{A}\mathbf{A}^T)^{-1}\mathbf{A} \\ &\quad + \mathbf{A}^T(\mathbf{A}\mathbf{A}^T)^{-1}[\mathbf{E} - (\mathbf{E}\mathbf{A}^T)(\mathbf{A}\mathbf{A}^T)^{-1}\mathbf{A}]. \end{aligned}$$

Consider the singular value decomposition(SVD) of matrix $\mathbf{A}$, i.e. $\mathbf{A} = \mathbf{u}\mathbf{S}\mathbf{v}^T$, wherein

$\mathbf{S} = \begin{bmatrix} s_1 & & & 0 \\ & \ddots & & \vdots \\ & & s_{(N-2)(M-2)} & 0 \\ & & & 0 \end{bmatrix}$ is an $(N-2)(M-2) \times (NM)$ matrix and both $\mathbf{u} \in \mathbb{R}^{(N-2)(M-2) \times (N-2)(M-2)}$, $\mathbf{v} \in \mathbb{R}^{(NM) \times (NM)}$ are orthogonal. Then $\mathbf{A}^T(\mathbf{A}\mathbf{A}^T)^{-1}\mathbf{A} = \mathbf{v}\mathbf{D}_1\mathbf{v}^T$, where $\mathbf{D}_1 = \begin{bmatrix} \mathbf{I}_{(N-2)(M-2)} & 0 \\ 0 & 0 \end{bmatrix}_{(NM) \times (NM)}$, and

$$\begin{aligned} (\mathbf{E}^T - \mathbf{A}^T(\mathbf{A}\mathbf{A}^T)^{-1}\mathbf{A}\mathbf{E}^T) &= (\mathbf{E}^T - \mathbf{v}\mathbf{D}_1\mathbf{v}^T\mathbf{E}^T) \\ &= \mathbf{v}\mathbf{D}_2\mathbf{v}^T\mathbf{E}^T, \text{ where } \mathbf{D}_2 = \mathbf{I} - \mathbf{D}_1. \\ \mathbf{E} - (\mathbf{E}\mathbf{A}^T)(\mathbf{A}\mathbf{A}^T)^{-1}\mathbf{A} &= \mathbf{E} - \mathbf{E}\mathbf{A}^T(\mathbf{A}\mathbf{A}^T)^{-1}\mathbf{A} \\ &= \mathbf{E} - \mathbf{E}\mathbf{v}\mathbf{D}_1\mathbf{v}^T \\ &= \mathbf{E}\mathbf{v}\mathbf{D}_2\mathbf{v}^T. \end{aligned}$$

Note that $\|\mathbf{v}\mathbf{D}_2\mathbf{v}^T\mathbf{E}^T\| \leq \epsilon$, $\|\mathbf{E}\mathbf{v}\mathbf{D}_2\mathbf{v}^T\| \leq \epsilon$. Since $\mathbf{A}^T(\mathbf{A}\mathbf{A}^T)^{-1}$ and $(\mathbf{A}\mathbf{A}^T)^{-1}\mathbf{A}$ are two generalized inverse of $\mathbf{A}$,

$$\mathbf{A}^T(\mathbf{A}\mathbf{A}^T)^{-1} = \mathbf{v}\mathbf{S}^*\mathbf{u}^T, \quad (\mathbf{A}\mathbf{A}^T)^{-1}\mathbf{A} = \mathbf{u}\mathbf{S}^*\mathbf{v},$$

where

$$\mathbf{S}^* = \left[\begin{array}{ccc|c} \frac{1}{s_1} & & & 0 \\ & \ddots & & \vdots \\ & & \frac{1}{s_{(N-2)(M-2)}} & 0 \end{array} \right] \quad (3.3)$$

and $\|\mathbf{A}^T(\mathbf{A}\mathbf{A}^T)^{-1}\| = \|(\mathbf{A}\mathbf{A}^T)^{-1}\mathbf{A}\| = \frac{1}{s_{\min}}$. Hence we conclude that

$$\|\mathbf{B} - \tilde{\mathbf{B}}\| \leq 2s_{\min}^{-1}\epsilon + O(\epsilon^2).$$

□

Remark 3.1 It is very difficult to estimate the minimum singular value of matrix $\mathbf{A}$ analytically, even for the simplest case when the Fokker–Planck equation is just a heat equation. But empirically we find that $s_{\min}^{-1}$ is usually not very large. For example, $s_{\min}^{-1}$ for the gradient flow with a double well potential is Sect. 5.4 is 0.4988, and $s_{\min}^{-1}$ for the “ring example” in Sect. 5.3 is only 0.2225.

4 Sensitivity Analysis of QSD

A stochastic differential equation has a QSD usually because it has a natural absorbing state. For example, in ecological models, this absorbing state is the natural boundary of the domain at which the population of some species is zero. Obviously invariant probability measures are easier to study than QSDs. One interesting question is that if we slightly modify the equation such that it does not have absorbing states any more, how can we quantify the difference between QSD and the invariant probability measure after the modification? This is called the sensitivity analysis of QSDs.

In this section, we focus on the difference between the QSD of a stochastic differential equation $X = \{X_t, t \in \mathbb{R}\}$ and the invariant probability measure of a modification of X , denoted by $Y = \{Y_t, t \in \mathbb{R}\}$. For the sake of simplicity, this paper only compares the QSD (resp. invariant probability measure) of the numerical trajectory of X (resp. Y), denoted by $\hat{X} = \{\hat{X}_n^\delta, n \in \mathbb{Z}_+\}$ (resp. $\hat{Y} = \{\hat{Y}_n^\delta, n \in \mathbb{Z}_+\}$). Denote the QSD (resp. invariant probability measure) of $\hat{X}$ (resp. $\hat{Y}$) by $\hat{\mu}_X$ (resp. $\hat{\mu}_Y$) and the QSD (resp. invariant probability measure) of the original SDE X (resp. Y) by μ_X (resp. μ_Y). The sensitivity of invariant probability measure against time discretization has been addressed in [11]. When the time step size of the time discretization is small enough, the invariant probability measure μ_Y is close to the numerical invariant probability measure $\hat{\mu}_Y$. The case of QSD is analogous. Hence $d(\hat{\mu}_X, \hat{\mu}_Y)$ is usually a good approximation of $d(\mu_X, \mu_Y)$.

We are mainly interested in the following two different modifications of X .

Case 1 Reflection at $\partial\mathcal{X}$ One easy way to modify the numerical trajectory $\hat{X}$ to eliminate the absorbing state is to add a reflecting boundary. This method preserves the local dynamics but changes the boundary condition. More precisely, the trajectory of $\hat{X}$ follows that of $\hat{Y}$ until it hits the boundary $\partial\mathcal{X}$. Without loss of generality assume $\partial\mathcal{X}$ is a smooth manifold embedded in $\mathbb{R}^d$. If $\hat{Y}_n^\delta \in \hat{X}_n^\delta \in \mathcal{X}^a$ but $\hat{X}_{n+1}^\delta \notin \mathcal{X}^a$, then $\hat{Y}_{n+1}^\delta$ is the mirror reflection of $\hat{X}_{n+1}^\delta$ against the boundary of $\mathcal{X}$. Denote the intersection of $\partial\mathcal{X}$ and the line segment from $\hat{X}_n^\delta$ to $\hat{X}_{n+1}^\delta$ by X_a . Denote the unit normal vector of $\partial\mathcal{X}$ at X_a by $\mathbf{n}$. Then some simple calculations give

$$\hat{Y}_{n+1}^\delta = X_a + (\text{Id} - 2\mathbf{n}\mathbf{n}^T)(\hat{X}_{n+1}^\delta - X_a).$$

In most ecological and epidemiological models the natural boundary is $\partial\mathcal{X} = \{(x_1, \dots, x_n) \mid x_i = 0 \text{ for some } 1 \leq i \leq n, x_j \geq 0 \text{ for all } 1 \leq j \leq n\}$. In this setting $\hat{Y}_{n+1}^\delta$ has an easy expression $\hat{Y}_{n+1}^\delta = \text{abs}(\hat{X}_{n+1}^\delta)$, where $\text{abs}(\cdot)$ means element-wise absolute value of a vector.

We remark that this reflection of the numerical trajectory is not consistent with the stochastic differential equation with reflecting boundary conditions, which solves a Skorohod's equation [23]. The problem is that the noise vanishes at $\partial\mathcal{X}$ for most ecological models. Hence $\partial\mathcal{X}$ remains to be absorbing even if the stochastic differential equation has reflecting boundaries. The "reflection" we use here is only for the numerical trajectory. It can be interpreted as a small random number of individuals are artificially added immediately after the extinction of this species happens. The goal here is to investigate the sensitivity of QSDs against a change of the boundary condition. See the summary at the end of Sect. 4.2

Case 2 Demographic noise in ecological models. We would also like to address the case of demographic noise in a class of ecological models, such as population model and epidemic model. For the sake of simplicity consider an 1D birth-death process with some environment noise

$$dY_t = f(Y_t)dt + \sigma Y_t dW_t, \quad (4.1)$$

where $Y_t \in \mathbb{R}$. Note that for the birth-death process model, $\{0\}$ is an absorbing state. It is known that for suitable $f(Y_t)$, the probability of Y_t hitting zero in finite time is zero [13]. However, if we take the demographic noise, i.e., the randomness of birth/death events, into consideration, the birth-death process becomes

$$dX_t = f(X_t)dt + \sigma' X_t dW_t^{(1)} + \epsilon \sqrt{X_t} dW_t^{(2)}, \quad (4.2)$$

where $\epsilon \ll 1$ is a small parameter that is proportional to $-1/2$ -th power of the scale of the population size, σ' is the new parameter that addresses the separation of environment noise and demographic noise. For example, if the steady state of X_t is around 1, we can choose $\sigma' = \sqrt{\sigma^2 - \epsilon^2}$.

Different from Eq. (4.1), the magnitude of random perturbation in Eq. (4.2) near the boundary is much larger. As a result, Eq. (4.2) could hit the boundary in finite time with strictly positive probability. (For example $dX_t = rX_t dt + \epsilon \sqrt{X_t} dW_t$ has strictly positive extinction probability in finite time whenever $\epsilon > 0$. This can be checked by applying the Itô's formula to a test function $1/X_t$ then taking the expectation. Whether Eq. (4.2) has finite time extinction depends on details of $f(x)$. But most ecological model has $f(X_t) \approx O(1) \times X_t$ when X_t is small in order to model the intrinsic growth of the population. That fits the setting of the example above.) Therefore, it is common for Eq. (4.1) to admit an invariant probability measure while Eq. (4.2) has a QSD. One very interesting question is that, if ϵ is sufficiently small, how different is the invariant probability measure of Eq. (4.1) from the QSD of Eq. (4.2)? This is very important in the study of ecological models because theoretically every model is subject a small demographic noise. If the invariant probability measure is dramatically different from the QSD after adding a small demographic noise term, then the Eq. (4.1) is not a good model due to its high sensitivity, and we must study the Eq. (4.2) directly.

We roughly follow [11] to carry out the sensitivity analysis of QSD. Here we slightly modify $\hat{X}_n^\delta$ such that if $\hat{X}_n^\delta \in \partial\mathcal{X}$, instead of sampling from the occupation measure, we immediately re-sample $\hat{X}_n^\delta$ from the QSD $\hat{\mu}_X$. This new process, denoted by $\tilde{X}_n^\delta$, admits an invariant probability measure $\hat{\mu}_X$. Now denote the transition kernel of $\tilde{X}_n^\delta$ and $\hat{Y}_n^\delta$ by P_X and P_Y respectively. The following Proposition is motivated by [16].

Proposition 4.1 For any $T > 0$, if there exists a constant $0 < \alpha < 1$ such that

$$d_w(\hat{\mu}_X P_Y^T, \hat{\mu}_Y P_Y^T) \leq \alpha d_w(\hat{\mu}_X, \hat{\mu}_Y),$$

then we have

$$d_w(\hat{\mu}_X, \hat{\mu}_Y) \leq \frac{d_w(\hat{\mu}_X P_X^T, \hat{\mu}_X P_Y^T)}{1 - \alpha}.$$

Proof Let $d_w(\cdot, \cdot)$ be the 1-Wasserstein distance defined in Sect. 2. We can decompose $d_w(\hat{\mu}_X, \hat{\mu}_Y)$ via the following triangle inequality:

$$d_w(\hat{\mu}_X, \hat{\mu}_Y) \leq d_w(\hat{\mu}_X P_X^T, \hat{\mu}_X P_Y^T) + d_w(\hat{\mu}_X P_Y^T, \hat{\mu}_Y P_Y^T).$$

Since the transition kernel P_Y^T has enough contraction such that

$$d_w(\hat{\mu}_X P_Y^T, \hat{\mu}_Y P_Y^T) \leq \alpha d_w(\hat{\mu}_X, \hat{\mu}_Y)$$

for some $\alpha < 1$, after some simplification, we have

$$d_w(\hat{\mu}_X, \hat{\mu}_Y) \leq \frac{d_w(\hat{\mu}_X P_X^T, \hat{\mu}_X P_Y^T)}{1 - \alpha}. \quad \square$$

Therefore, in order to estimate $d_w(\hat{\mu}_X, \hat{\mu}_Y)$, we need to look for suitable numerical estimators of the finite time error and the speed of contraction of P_Y^T . The finite time error can be easily estimated in both cases. And the speed of contraction α comes from the geometric ergodicity of the Markov process $\hat{Y}$. If our numerical estimation gives

$$d_w(\mu P_Y^T, \nu P_Y^T) \leq C e^{-\gamma T}, \quad (4.3)$$

then we set $\alpha = e^{-\gamma T}$. As discussed in [11], this is a quick way to estimate α . In fact, in all examples that we have tested, when starting from $\hat{\mu}_X$ and $\hat{\mu}_Y$, the prefactor C of the coupling probability in Eq. (4.3) is not far away from 1. Hence in practice it does not differ from the “true upper bound” very much. The “true upper bound” of α in [11] comes from the extreme value theory, which is much more expensive to compute.

4.1 Estimation of Contraction Rate

Similar to in [21], we use the following coupling method to estimate the contraction rate α . Let $\hat{Z} = \hat{Z}_n^\delta = (\hat{Y}_n^1, \hat{Y}_n^2)$ be a Markov process in $\mathbb{R}^{2d}$ such that $\hat{Y}_n^1$ and $\hat{Y}_n^2$ are two copies of $\hat{Y}$, and that $\hat{Z}$ is a coupling of $\hat{Y}_n^1$ and $\hat{Y}_n^2$. Recall that τ^c is the coupling time, which is also the first passage time to the “diagonal” hyperplane $\{(\mathbf{x}, \mathbf{y} \in \mathbb{R}^{2d}) | \mathbf{y} = \mathbf{x}\}$. Then by Lemma 2.1

$$d_w(\hat{\mu}_X P_Y^t, \hat{\mu}_Y P_Y^t) \leq \mathbb{P}[\tau^c > t].$$

As discussed in [21], we need a hybrid coupling scheme to make sure that two numerical trajectories can couple. Some coupling methods such as reflection coupling or synchronous coupling are implemented in the first phase to bring two numerical trajectories together. Then we compare the probability density function for the next step and couple these two numerical trajectories with the maximal possible probability (called the maximal coupling). After doing this for many times, we will have many samples of τ^c denote by τ^c . We use the exponential tail of $\mathbb{P}[\tau^c > t]$ to estimate the contraction rate α . More precisely, we look for a constant $\gamma > 0$ such that

$$-\gamma = \lim_{t \rightarrow \infty} \frac{1}{t} \log(\mathbb{P}[\tau^c > t])$$

if the limit exists. See **Algorithm 3** for the detail of implementation of coupling. Note that we cannot simply compute the contraction rate start from $t = 0$ because only the tail of coupling time can be considered as exponentially distributed. In practice, we apply the same method as we compute the killing rate in Sect. 3.1. After having τ^c , it is easy to choose a sequence of times $t_0, t_1, \dots, t_n$ and calculate $n_i = |\{\tau_m^c > t_i \mid 0 \leq m \leq M\}|$ for each $i = 0, \dots, n$. Then $p_i = n_i/M$ is an estimator of $\mathbb{P}_\mu[\tau^c > t_i]$. Now let p_i^u (resp. p_i^l) be the upper (resp. lower) bound of the confidence interval of p_i such that

$$p_i^u = \tilde{p} + z \sqrt{\frac{\tilde{p}}{\tilde{n}_i} (1 - \tilde{p})} \quad \text{resp.} \quad p_i^l = \tilde{p} - z \sqrt{\frac{\tilde{p}}{\tilde{n}_i} (1 - \tilde{p})},$$

where $z = 1.96$, $\tilde{n}_i = n_i + z^2$ and $\tilde{p} = \frac{1}{\tilde{n}_i} (n_i + \frac{z^2}{2})$. Let t_n be the largest time that we can still collect available samples. If there exist constants C and $i_0 < n$ such that $p_i^l \leq C e^{\gamma t_i} \leq p_i^u$ for each $i_0 \leq i \leq n$, we say that the exponential tail starts at $t = t_{i_0}$. We accept the estimate of the exponential tail with rate $e^{-\gamma t}$ if the confidence interval $p_{i_0}^u - p_{i_0}^l$ is sufficiently small, i.e., the estimate of coupling probability at $t = t_{i_0}$ is sufficiently trustable. Otherwise we need to run Algorithm 3 for longer time to eliminate the initial bias in τ^c .

Algorithm 3 Estimation of contraction rate α

Input: Initial values $x, y \in K$

Output: An estimation of contraction rate α

Choose threshold $d > 0$

for $i = 1$ to N_s **do**

$\tau_i^c = 0, t = 0, (\hat{Y}_t^1, \hat{Y}_t^2) = (x, y)$

Flag = 0

while Flag=0 **do**

if $|\hat{Y}_t^1 - \hat{Y}_t^2| > d$ **then**

 Compute $(\hat{Y}_{t+1}^1, \hat{Y}_{t+1}^2)$ using reflection coupling or independent coupling

$t \leftarrow t + 1$

else

 Compute $(\hat{Y}_{t+1}^1, \hat{Y}_{t+1}^2)$ using maximal coupling

if coupled successfully **then**

 Flag=1

$\tau_i^c = t$

else

$t \leftarrow t + 1$

end if

end if

end while

end for

Use $\tau_1^c, \dots, \tau_{N_s}^c$ to compute $\mathbb{P}(\tau^c > t)$

Fit the tail of $\log \mathbb{P}(\tau^c > t)$ versus t by linear regression. Compute the slope γ .

4.2 Estimator of Error Terms

It remains to estimate the finite time error $d_w(\hat{\mu}_X P_X^T, \hat{\mu}_Y P_Y^T)$. As we mentioned in the beginning of this section, we will consider two different cases and estimate the finite time errors respectively.

4.2.1 Case 1: Reflection at $\partial\mathcal{X}$

Recall that the modified Markov process $\hat{Y}$ reflects at the boundary $\partial\mathcal{X}$ when it hits the boundary. Hence two trajectories

$$\begin{aligned}\tilde{X}_{n+1}^\delta &= \tilde{X}_n^\delta + f(\tilde{X}_n^\delta)\delta + \sigma(\tilde{X}_n^\delta)(W_{(n+1)\delta} - W_{n\delta}) \\ \hat{Y}_{n+1}^\delta &= \hat{Y}_n^\delta + f(\hat{Y}_n^\delta)\delta + \sigma(\hat{Y}_n^\delta)(W_{(n+1)\delta} - W_{n\delta})\end{aligned}$$

are identical if we set the same noise in the simulation process. $\tilde{X}$ only differs from $\hat{Y}$ when $\tilde{X}$ hits the boundary $\partial\mathcal{X}$. When $\tilde{X}$ and $\hat{Y}$ hit the boundary, $\tilde{X}$ is resampled from $\hat{\mu}_X$, and $\hat{Y}$ reflects at the boundary. Hence the finite time error $d_w(\hat{\mu}_X P_X^T, \hat{\mu}_X P_Y^T)$ is bounded from above by the killing probability within the time interval $[0, T]$ when starting from $\hat{\mu}_X$.

In order to sample initial value $\mathbf{x}$ from the numerical invariant measure $\hat{\mu}_X$, we consider a long trajectory $\{\tilde{X}_n^\delta\}$. The distance between $\tilde{X}$ and the modified trajectory $\hat{Y}$ is recorded after time T . Then we let $\hat{Y}_0^\delta = \tilde{X}_0^\delta = \mathbf{x}_{i+1}$ and restart the simulation, where $\mathbf{x}_{i+1} = \tilde{X}_T^\delta$ from the i -th iteration. See the **Algorithm 4** for the detail.

Algorithm 4 Estimate finite time error for Case 1

Input: Initial value $\mathbf{x}_1$

Output: An estimator of $d_w(\hat{\mu}_X P_X^T, \hat{\mu}_X P_Y^T)$

for $i = 1$ to N_s **do**

 Using the same noise, simulate $\tilde{X}$ and $\hat{Y}$ with initial value $\mathbf{x}_i$ up to T

 Set $d_i = 0$

if $\tau < T$ **then**

 Regenerate $\tilde{X}$ as its empirical distribution

$d_i = d(\tilde{X}_T^\delta, \hat{Y}_T^\delta)$

end if

 Let $\mathbf{x}_{i+1} = \tilde{X}_T^\delta$

end for

Return $\frac{1}{N_s} \sum_{i=1}^{N_s} d_i$

When the number of samples is sufficiently large, $\mathbf{x}_1, \dots, \mathbf{x}_{N_s}$ are from a long trajectory of the time- T skeleton of $\tilde{X}_T$. Hence they are approximately sampled from $\hat{\mu}_X$. The error term $d_i = d(\tilde{X}_T^\delta, \hat{Y}_T^\delta)$ for $\tilde{X}_0^\delta = \hat{Y}_0^\delta = \mathbf{x}_i$ estimates $d(\tilde{X}_T, \hat{Y}_T)$. Let $\hat{\mu}_X^2$ be the probability measure on $\mathbb{R}^d \times \mathbb{R}^d$ that is supported by the hyperplane $\{(x, y) \in \mathbb{R}^d \times \mathbb{R}^d \mid x = y\}$ such that $\hat{\mu}_X^2(\{(x, x) \mid x \in A\}) = \hat{\mu}_X(A)$ for any $A \in \mathcal{B}(\mathcal{X})$. Since the pushforward map $\hat{\mu}_X^2(P_X^T \times P_Y^T)$ is a coupling, it is easy to see that the output of **Algorithm 4** gives an upper bound of $d_w(\hat{\mu}_X P_X^T, \hat{\mu}_X P_Y^T)$.

From the analysis above, we have the following lemma, which gives an upper bound of the finite time error $d_w(\hat{\mu}_X P_X^T, \hat{\mu}_X P_Y^T)$.

Lemma 4.1 *For the Wasserstein distance induced by the distance given in (2.7), we have*

$$d_w(\hat{\mu}_X P_X^T, \hat{\mu}_X P_Y^T) \leq \int_{\mathbb{R}^d} \mathbb{P}_x(\tau < T) \hat{\mu}_X(dx),$$

where x is the initial value with distribution μ_Y .

Proof Note that $\mu_X^2(P_X^T \times P_Y^T)$ is a coupling of $\mu_X P_X^T$ and $\mu_X P_Y^T$. From the definition of Wasserstein distance, we have

$$\begin{aligned} d_w(\hat{\mu}_X P_X^T, \hat{\mu}_X P_Y^T) &\leq \int_{\mathbb{R}^d \times \mathbb{R}^d} d(x, y) \hat{\mu}_X^2(P_X^T \times P_Y^T)(dx dy) \\ &= \int_{\mathbb{R}^d} \mathbb{E}_{(x, x)}[d(\tilde{X}_T^\delta, \hat{Y}_T^\delta)] \hat{\mu}_X(dx) \\ &= \int_{\mathbb{R}^d} \mathbb{P}_x(\tau < T) d(\tilde{X}_T^\delta, \hat{Y}_T^\delta) \hat{\mu}_X(dx) \\ &\leq \int_{\mathbb{R}^d} \mathbb{P}_x(\tau < T) \hat{\mu}_X(dx), \end{aligned}$$

the inequality in the last step comes from the definition $d(x, y) = \max(1, \|x - y\|)$. $\square$

4.2.2 Case 2: Impact of a Demographic Noise $\epsilon \sqrt{X_t} dW_t$

Another case that often involves QSD is the demographic noise. Let $\hat{X}$ be the numerical trajectory of the SDE with an additive demographic noise $\epsilon \sqrt{X_t} dW_t$. Let $\tilde{X}$ be the modification of $\hat{X}$ that resample from $\hat{\mu}_X$ whenever hitting $\partial\mathcal{X}$ so that it admits $\hat{\mu}_X$ as an invariant probability measure. Let $\hat{Y}$ be the numerical trajectory of the SDE without demographic noise so that $\hat{Y}$ admits an invariant probability measure. We have trajectories

$$\begin{aligned} \tilde{X}_{n+1}^\delta &= \tilde{X}_n^\delta + f(\tilde{X}_n^\delta) \delta + \sigma(\tilde{X}_n^\delta)(W_{(n+1)\delta} - W_{n\delta}) + \epsilon \sqrt{\tilde{X}_n^\delta}(W'_{(n+1)\delta} - W'_{n\delta}) \\ \hat{Y}_{n+1}^\delta &= \hat{Y}_n^\delta + f(\hat{Y}_n^\delta) \delta + \sigma'(\hat{Y}_n^\delta)(W_{(n+1)\delta} - W_{n\delta}). \end{aligned}$$

Here we assume that $\hat{Y}$ has zero probability to hit the absorbing set $\partial\mathcal{X}$ in finite time. Different from the Case 1, we will need to study the effect of the demographic noise. When estimating the finite time error $d_w(\hat{\mu}_X P_X^T, \hat{\mu}_X P_Y^T)$, we still need to sample the initial value $\mathbf{x}$ from $\hat{\mu}_X$ and record the distance between these two trajectories $\tilde{X}$ and $\hat{Y}$ up to time T . The distance between $\tilde{X}$ and $\hat{Y}$ can be decomposed into two parts: one is from the killing and resampling, the other is from the demographic noise. The first term is the same as in Case 1. The second term is due to the nonzero demographic noise that can separate $\tilde{X}$ and $\hat{Y}$ before the killing. In a population model, this effect is more obvious when one species has small population, because $\sqrt{x} \gg x$ when $0 < x \ll 1$. See the description of **Algorithm 5** for the full detail.

When N_s is sufficiently large, $\mathbf{x}_1, \dots, \mathbf{x}_{N_s}$ are from a long trajectory of the time- T skeleton of $\tilde{X}_T$. Hence they are approximately sampled from $\hat{\mu}_X$. The error term d_i for $\tilde{X}_0^\delta = \hat{Y}_0^\delta = \mathbf{x}_i$ estimates $d(\tilde{X}_T, \hat{Y}_T)$. A similar coupling argument shows that the output of **Algorithm 5** is an upper bound of $d_w(\hat{\mu}_X P_X^T, \hat{\mu}_X P_Y^T)$.

For each initial value $x \in \mathbb{R}^d$, denote

$$\theta_x = \mathbb{E}_x[d(\tilde{X}_T^\delta, \hat{Y}_T^\delta) | \tilde{X}_0^\delta = \hat{Y}_0^\delta = x, \tau > T]. \quad (4.4)$$

Similar to in Case 1, the following lemma gives an upper bound for the finite time error $d_w(\mu_X P_X^T, \mu_X P_Y^T)$.

Lemma 4.2 *For the Wasserstein distance induced by the distance given in (2.7), we have*

$$d_w(\hat{\mu}_X P_X^T, \hat{\mu}_X P_Y^T) \leq \int \mathbb{P}_x(\tau < T) \hat{\mu}_X(dx) + \int \theta_x \hat{\mu}_X(dx),$$

where x is the initial value with distribution $\hat{\mu}_X$ and θ_x .

Algorithm 5 Estimate finite time error for Case 2

Input: Initial value $\mathbf{x}_1$
Output: An estimator of $d_w(\mu_X P_X^T, \mu_Y P_Y^T)$
for $i = 1$ to N_s **do**

 Using the same noise, simulate $\tilde{X}$ and $\hat{Y}$ with initial value $\mathbf{x}_1$ up to T

 Set $\text{Flag} = 0, d_i = 0$
if $\tau < T$ **then**

 Regenerate $\tilde{X}$ as its empirical distribution

 $\eta_i = d(\tilde{X}_T^\delta, \hat{Y}_T^\delta)$
 $\text{Flag} = 1$
else
 $\theta_i = d(\tilde{X}_T^\delta, \hat{Y}_T^\delta)$
end if
 $d_i = \theta_i + \mathbf{1}_{\{\text{Flag}=1\}}(\eta_i - \theta_i)$

 Let $\mathbf{x}_{i+1} = \tilde{X}_T^\delta$
end for

 Return $\frac{1}{N_s} \sum_{i=1}^{N_s} d_i$

Proof Note that $\hat{\mu}_X^2(P_X^T \times P_Y^T)$ is a coupling of $\hat{\mu}_X P_X^T$ and $\hat{\mu}_X P_Y^T$. From the definition of the Wasserstein distance, we have

$$\begin{aligned} d_w(\hat{\mu}_Y P_Y^T, \hat{\mu}_X P_X^T) &\leq \int_{\mathbb{R}^d \times \mathbb{R}^d} d(x, y) \hat{\mu}_X^2(P_X^T \times P_Y^T)(dx dy) \\ &= \int_{\mathbb{R}^d} \mathbb{E}_{(x, x)}[d(\tilde{X}_T^\delta, \hat{Y}_T^\delta)] \hat{\mu}_X(dx) \\ &= \int_{\mathbb{R}^d} \mathbb{P}_x(\tau < T) d(\tilde{X}_T^\delta, \hat{Y}_T^\delta) \hat{\mu}_X(dx) \\ &\quad + \int_{\mathbb{R}^d} \mathbb{P}_x(\tau > T) d(\tilde{X}_T^\delta, \hat{Y}_T^\delta) \hat{\mu}_X(dx) \\ &\leq \int_{\mathbb{R}^d} \mathbb{P}_x(\tau < T) \hat{\mu}_X(dx) + \int \theta_x \hat{\mu}_X(dx) \end{aligned}$$

 according to the definition of θ_x . □

In summary, the sensitivity of QSD depends on both mixing rate of the modified process and the killing probability. Both higher mixing rate and lower killing probability per unit time lead to a more robust QSD that is not sensitive against small change of the dynamics or the boundary condition. More precisely, let T be the constant time we choose and α be the contraction of operator P_Y^T .

(1) If the modified process $\hat{Y}$ only differs from $\hat{X}$ with a reflection, then we have

$$d_w(\hat{\mu}_X, \hat{\mu}_Y) \leq (1 - \alpha)^{-1} \int_{\mathbb{R}^d} \mathbb{P}_x[\tau < T] \hat{\mu}_X(dx).$$

(2) If the demographic noise is removed from $\hat{Y}$, then

$$d_w(\hat{\mu}_X, \hat{\mu}_Y) \leq (1 - \alpha)^{-1} \left\{ \int_{\mathbb{R}^d} \mathbb{P}_x(\tau < T) \hat{\mu}_X(dx) + \int_{\mathbb{R}^d} \theta_x \hat{\mu}_X(dx) \right\},$$

where θ_x is defined in Eq. (4.4).




Fig. 1 Plot of $\mathbb{P}(\tau > t)$ vs. t , confidence interval (upper bound and lower bound) and function $y = e^{-\lambda t}$

5 Numerical Examples

5.1 Ornstein–Uhlenbeck Process

The first SDE example is the Ornstein-Uhlenbeck process :

$$dX_t = \theta(\mu - X_t)dt + \sigma dW_t, \quad (5.1)$$

where $\theta > 0$ and $\sigma > 0$ are parameters, μ is a constant. In addition, W_t is a Wiener process, and σ is the strength of the noise. In our simulation, we set $\theta = 1$, $\mu = 2$, $\sigma = 1$ and the absorbing set $\partial\mathcal{X} = (-\infty, 0] \cup [3, \infty)$. In addition, we apply the Monte Carlo simulation with 512 mesh points on the interval $[0, 3]$.

We first need to use Algorithm 1 to estimate the survival rate λ . Our simulation uses Euler–Maruyama scheme with $\delta = 0.001$ and sample size $N = 10^6$ and $N = 10^8$ depending on the setting. All samples of killing times are recorded to plot the exponential tail. The mean of killing times gives an estimate $\lambda = -0.267176$. The exponential tail of $\mathbb{P}(\tau > t)$ vs. t , the upper and lower bound of the confidence interval, and the plot of $e^{-\lambda t}$ are compared in Fig. 1. We can see that the plot of $e^{-\lambda t}$ falls in the confidence interval for all t . Hence the estimate of λ is accepted.

Furthermore, we would like to show the robustness of our data-driven QSD solver. The QSD is not explicit given so we use very large sample size (10^{10} samples) and very small step size (10^{-4}) to obtain a much more accurate solution, which is served as the ground truth. Then we compare the numerical solutions obtained by the Monte Carlo method and the data-driven method for QSD with $N = 10^6$ and $N = 10^8$ samples, respectively. The result is shown in the first column of Fig. 2. The data-driven solver performs much better than the Monte Carlo approximation for $N = 10^6$ samples. It takes 10^8 samples for the direct Monte Carlo sampler to produce a solution that looks as good as the QSD solver. Similar to the data-driven Fokker–Planck solver, our data-driven QSD solver can tolerate high level spatially uncorrelated error in the Monte Carlo simulation.

It remains to check the effect of Brownian Bridge. We apply different time step sizes $\delta = 0.01$ and $\delta = 0.001$ for each trajectory. We use 10^7 samples for $\delta = 0.001$ and 10^6 samples for $\delta = 0.01$ to make sure that the number of killing events (for estimating the killing rate) are comparable. When $\delta = 0.001$, the error is small with and without Brownian bridge




Fig. 2 Upperpanel: (Left) Monte Carlo estimation vs. data-driven solver estimation for sample size $N = 10^6$. (Right) Effect of Brownian Bridge for sample size $N = 10^7$ and time step size 0.001. Lowerpanel: (Left) Effect of Brownian Bridge for sample size $N = 10^6$ and time step size 0.01. (Right) Monte Carlo estimation vs. data-driven solver estimation for sample size $N = 10^8$

correction. But Brownian bridge correction obviously improves the quality of solution when $\delta = 0.01$. See the lower left panel of Fig. 2. This is expected because, with larger time step size, the probability that the Brownian bridge hits the absorbing set $\partial\mathcal{X}$ gets higher.

5.2 Wright–Fisher Diffusion

The second numerical example is the Wright-Fisher diffusion model, which describes the evolution of colony undergoing random mating, possible under the additional actions of mutation and selection with or without dominance [14]. The Wright-Fisher model is an SDE

$$dX_t = -X_t dt + \sqrt{X_t(1-X_t)} dW_t,$$

where W_t is a Wiener process and $\partial\mathcal{X} = \{0\}$ is the absorbing set. By the analysis of [14], the Yaglom limit, i.e., the QSD, satisfies

$$\lim_{t \rightarrow \infty} \mathbb{P}[X_t \in dy | \tau > t] = 2(1-y)dy.$$

The goal of this example is to show the effect of Brownian bridge when the coefficient of noise is vanishing at the boundary. Since the Euler–Maruyama scheme only has an order of accuracy 0.5, in the simulation, we apply the Milstein scheme, which reads

$$\hat{X}_{n+1}^\delta = \hat{X}_n^\delta - \hat{X}_n^\delta \delta + \sqrt{\hat{X}_n^\delta(1-\hat{X}_n^\delta)}(W_{(n+1)\delta} - W_{n\delta}) + \frac{1}{4}(1-2\hat{X}_n^\delta)[(W_{(n+1)\delta} - W_{n\delta})^2 - \delta]$$

One difficulty of using the Brownian bridge correction in this model is that the coefficient of the Brownian motion is vanishing at the boundary. Recall that the strength coefficient of Brownian bridge is denoted by ϕ . Larger ϕ means X_t has higher probability of hitting the boundary. Since the coefficient of the Brownian motion is vanishing at the boundary,




Fig. 3 Effect of Brownian Bridge and a correction of Brownian Bridge. Left: Monte Carlo approximations without Brownian Bridge correction, with original Brownian Bridge correction, and with modified Brownian Bridge correction, in comparison to the analytical QSD. Right: Result from the data-driven QSD solver using the Monte Carlo simulation data from Left panel

the effective strength coefficient ϕ becomes dramatically smaller when X_t gets closer to the boundary. As a result, it is not a good idea to still approximate X_t by a Brownian motion.

And the original strength coefficient $\phi = \sqrt{\hat{X}_n^\delta(1 - \hat{X}_n^\delta)}$ can dramatically overestimate the probability of hitting the boundary. Estimating the hitting probability of this diffusion bridge is a difficult problem that is well beyond the scope of this paper. To the best of our knowledge, it is not possible to explicitly calculate the conditional distribution of the diffusion bridge that starts from $x := \hat{X}_n^\delta$ and ends at $y := \hat{X}_{n+1}^\delta$. Instead, we find that an empirically found strength coefficient $\phi^2 = \frac{1}{3} \min\{x(1-x), y(1-y)\}$ can fix this problem. We note that this is not a simple ad-hoc solution because in a stochastic differential equation, the diffusion plays a dominate role in a very short time interval. Hence a similar modification of the Brownian bridge should work for many stochastic differential equations with $\sqrt{X_t}dB_t$ noise terms. In particular, the vanishing coefficient $\sqrt{X_t}dB_t$ also appears in many ecological models. We will implement this modified Brownian bridge correction when simulating these models.

The effect of Brownian bridge is shown in the left side of Fig. 3. We compare the solutions obtained via Monte Carlo method and the data-driven method with Brownian Bridge by running 10^7 samples on $[0, 1]$ with time step size $\delta = 0.01$. The Monte Carlo approximation is far from the true density function of Beta(1,2) near $x = 0$, while the use of the original Brownian Bridge only makes things worse. The modified Brownian Bridge solves this boundary effect problem reasonably well. The output of the data-driven QSD solver has a similar result (Fig. 3 Right). Let $x = \hat{X}_n^\delta$ and $y = \hat{X}_{n+1}^\delta$. One can see that the numerical QSD is much closer to the true distribution if we replace the strength coefficient of the Brownian bridge $\phi^2 = x(1-x)$ by the modified strength coefficient $\phi^2 = \frac{1}{3} \min\{x(1-x), y(1-y)\}$.

5.3 Ring Density Function

Consider the following stochastic differential equation:

$$\begin{aligned}dX &= (-4X(X^2 + Y^2 - 1) + Y)dt + \sigma dW_t^X \\dY &= (-4X(X^2 + Y^2 - 1) - X)dt + \sigma dW_t^Y,\end{aligned}$$

where W_t^X and W_t^Y are independent Wiener processes. In the simulation, we set the strength of noise $\sigma = 1$.

We first look at the approximation obtained by Monte Carlo method with 256×256 mesh points on the domain $D = [-1.5, 1.5] \times [-1.5, 1.5]$. The simulation uses step size $\delta = 0.001$ and $N = 10^8$ samples. (See upper left panel in Fig. 4). The Monte Carlo approximation has too much noise to be useful. The quality of this solution can be significantly improved by using the data-driven QSD solver. See upper right panel in Fig. 4.

The simulation result shows the estimated rate of killing $\lambda = -0.176302$. We use this example to test the sensitivity of solution $\mathbf{u}$ against small change of the killing rate. We compare the approximations obtained by setting the killing rate be λ , 1.1λ and 0.9λ respectively. Heat maps of the difference between QSDs with “correct” and “wrong” killing rates are shown in two middle panels in Fig. 4. It shows that difference brought by “wrong” rates is only $\approx O(10^{-4})$, which can be neglected. This result coincides with the analysis in Theorem 3.1 in this paper.

Finally, we would like to emphasize that the data-driven QSD solver can tolerate very high level of spatially uncorrelated noise in the reference solution $\mathbf{v}$. For example, if we use the same long trajectory with 10^8 samples that generates the top left panel of Fig. 4, but only select 10^5 samples with intervals of 10^3 steps of the numerical SDE solver, the Monte Carlo data becomes very noisy (Bottom left panel of Fig. 4). However, longer intermittency between samples also reduces the spatial correlation between samples. As a result, the output of the QSD solver has very little change except at the boundary grid points, because the optimization problem (3.2) projects most of error to the boundary of the domain. (See bottom right panel of Fig. 4.) This result highlights the need of high quality samplers. A Monte Carlo sampler with smaller spatial correlation between samples can significantly reduce the number of samples needed in the data-driven QSD solver.

5.4 Sensitivity of QSD: 1D Examples

In this subsection, we use 1D examples to study the sensitivity of QSDs against changes on boundary conditions. Consider an 1D gradient flow of the potential function $V(x)$ with an additive noise perturbation

$$X_t = -V'(X_t)dt + \sigma dW_t. \quad (5.2)$$

Let $(-\infty, 0]$ be the absorbing state of X_t . So if $V(0) < \infty$, X_t admits a QSD, denoted by μ_X . If we let the stochastic process reflect at $x = 0$, the modified stochastic process, denoted by Y_t , admits an invariant probability measure denoted by μ_Y . We will compare the sensitivity of μ_X against the change of boundary condition for two different cases whose speed of mixing are different, namely a single well potential function and a double well potential function.

We choose a single well potential function $V_1(x) = (x - 1)^2$ and a double well potential function $V_2(x) = x^4 - 4\sqrt{2}x^3 + 10x^2 - 4\sqrt{2}x + 1$. The values of minima of both V_1 and V_2 are zero. The values of V_1 and V_2 at the absorbing state are 1. And the height of the barrier between two local minima of V_2 is 1. The strength of noise is $\sigma = 0.7$ in both examples.




Fig. 4 (Ring density) *Upperpanel* :The approximation by Monte Carlo simulation(left) and the algorithm in Sect. 3.2(right) with 256×256 mesh points and 10^8 samples. *Middlepanel*: Sensitivity effect of small change to killing rate λ . Lower panel: The approximation by Monte Carlo simulation with smaller samples(left) and the output of the data-driven QSD solver(right)

See Fig. 5 middle column for plots of these two potential functions. In both cases, the QSD and the invariant probability measure are computed on the domain $D = [0, 3]$. To further distinguish these two cases, we denote the QSD of Eq. (5.2) with absorbing state $x = 0$ and potential function $V_1(x)$ (resp. $V_2(x)$) by μ_X^1 (resp. μ_X^2) and the invariant probability measure of Eq. (5.2) with reflection boundary at $(-\infty, 0]$ and potential function $V_1(x)$ (resp. $V_2(x)$) by μ_Y^1 (resp. μ_Y^2). Probability measures μ_X^1 and μ_Y^1 (resp. μ_X^2 and μ_Y^2) are compared in Fig. 5 right column. We can see that the QSD and the invariant probability measure have small difference for the single well potential function V_1 . But they look very different for




Fig. 5 (Single well vs. Double well potential function) *Leftcolumn*: $\mathbb{P}(\tau^c > t)$ vs. t . *Middlecolumn*: Single well and double well potential functions. *Rightcolumn*: QSD vs. invariant density function

the double well potential function V_2 . With the double well potential function, there is a visible difference between probability density functions of μ_X^2 and μ_Y^2 . The density function of QSD is much smaller than the invariant probability measure around the left local minimum $x = 1 - \sqrt{2}$ because this local minimum is closer to the absorbing set, which makes killing and regeneration more frequent when X_t is near this local minimum. In other words, the QSD of Eq. (5.2) with respect to the double well potential is very sensitive against the change at the boundary.

The reason of the high sensitivity is illustrated by the coupling argument. We first run Algorithm 3 with 8 independent long trajectories with length of 10^6 and collect the coupling times. The slope of exponential tail of the coupling time gives the rate of contraction of P_Y^T . The $\mathbb{P}(\tau^c > t)$ versus t plot is demonstrated in log-linear plot in Fig. 5 left column. The slope of exponential tail is $\gamma = 2.031414$ for the single well potential V_1 , and $\gamma = 0.027521$ for the double well potential case. Then we run Algorithm 4 to estimate the finite time error $d_w(\mu_X P_X^T, \mu_X P_Y^T)$ for both cases. Since the single well potential case has a much faster coupling speed, we can choose $T = 0.5$. The output of Algorithm 4 is $d_w(\mu_X P_X^T, \mu_X P_Y^T) \approx 0.00391083$. This gives an estimate $d_w(\mu_X, \mu_Y) \approx 0.0061$. The double well potential case converges much slower. We choose $T = 20$ to make sure that the denominator $1 - e^{-\gamma T}$ is not too small. As a result, Algorithm 4 gives an approximation $d_w(\mu_X P_X^T, \mu_X P_Y^T) \approx 0.06402$, which means $d_w(\mu_X, \mu_Y) \approx 0.1512$. This is consistent with the right column seen in Fig. 5, the QSD of the double well potential is much more sensitive against a change of the boundary condition than the single well potential case.

5.5 Lotka–Volterra Competitive Dynamics

In this example, we focus on the effect of demographic noise on the classical Lotka–Volterra competitive system. The Lotka–Volterra competitive system with some environmental fluctuation has the form

$$\begin{aligned} dY_1(t) &= Y_1(t)(l_1 - a_{11}Y_1(t) - a_{12}Y_2(t))dt + \sigma'_1 Y_1(t)dW_1(t), \\ dY_2(t) &= Y_2(t)(l_2 - a_{22}Y_1(t) - a_{21}Y_2(t))dt + \sigma'_2 Y_2(t)dW_2(t). \end{aligned} \quad (5.3)$$

Here $l_i > 0$ is the per-capita growth rate of species i and $a_{ij} > 0$ is the per-capita competition rate between species i and j . More details can be found in [13]. Model parameters are chosen to be $l_1 = 2, l_2 = 4, a_{11} = 0.8, a_{12} = 1.6, a_{21} = 1, a_{22} = 5$. Let $\partial\mathcal{X}$ be the union of x-axis and y-axis. For suitable σ'_1 and σ'_2 , Y_1 and Y_2 can coexist such that the probability of (Y_1, Y_2) hits $\partial\mathcal{X}$ in finite time is zero. So Eq. (5.3) admits an invariant probability measure, denoted by μ_Y .

As a modification, we add a small demographic noise term to Eq. (5.4). The equation becomes

$$\begin{aligned} dX_1(t) &= X_1(t)(l_1 - a_{11}X_1(t) - a_{12}X_2(t))dt + \sigma_1 X_1(t)dW_1(t) + \epsilon\sqrt{X_1(t)}dW'_1(t), \\ dX_2(t) &= X_2(t)(l_2 - a_{22}X_1(t) - a_{21}X_2(t))dt + \sigma_2 X_2(t)dW_2(t) + \epsilon\sqrt{X_2(t)}dW'_2(t). \end{aligned} \quad (5.4)$$

It is easy to see that Eq. (5.4) can exit to the boundary $\partial\mathcal{X}$. It admits a QSD, denoted by μ_X .

In order to study the effect of demographic noise, we compare μ_Y , the numerical invariant measure of Eq. (5.3), and μ_X , the QSD of Eq. (5.4). In our simulation, we fix the strength of demographic noise as $\epsilon = 0.05$ and compare μ_X and μ_Y at two different levels of the environment noise $\sigma_1 = \sigma_2 = 0.75$ and $\sigma_1 = \sigma_2 = 1.1$ respectively. The coefficient σ'_1 and σ'_2 in Eq. (5.3) satisfies $\sigma'_i = \sqrt{\sigma_i^2 + \epsilon^2}$ for $i = 1, 2$ to match the effect of the additional demographic noise. Compare Figs. 6 and 7, one can see that μ_Y has significant concentration at the boundary when $\sigma_1 = \sigma_2 = 1.1$.

The result for $\sigma_1 = \sigma_2 = 0.75$ is shown in Fig. 6. Left bottom of Fig. 6 shows the invariant measure. The QSD is shown on right top of Fig. 6. The total variation distance between these two measures are shown at the bottom of Fig. 6. The difference is very small and it just appears around boundary. This is reasonable because with high probability, the trajectories of Eq. (5.4) moves far from the absorbing set $\partial\mathcal{X}$ in both cases, which makes the regeneration events happen less often. This is consistent with the result of Lemma 4.2. We compute the distribution of the coupling time. The coupling time distribution and its exponential tail are shown in Fig. 6 Top Left. Then we use Algorithm 5 to compute the finite time error. To better match two trajectories given by Eqs. (5.3) and (5.4), we separate the noise term in Eq. (5.3) into $\sigma'_i Y_i(t)dW_i(t) = \sigma_i Y_i(t)dW_i^{(1)}(t) + \epsilon Y_i(t)dW_i^{(2)}(t)$ for $i = 1, 2$, where $W_i^{(1)}(t)$ and $W_i^{(2)}(t)$ use the same Wiener process trajectory as $W_i(t)$ and $W'_i(t)$ in Eq. (5.4) for $i = 1, 2$. Let $T = 4$. The finite time error caused by the demographic noise is 0.01773. As a result, the upper bound given in Lemma 4.2 is 0.02835. Note that as seen in Fig. 6, this upper bound actually significantly overestimates the distance between the invariant probability measure and the QSD. The empirical total variation distance is much smaller than our theoretical upper bound. This is because the way of matching $\sigma'_i Y_i(t)dW_i(t)$ and $\sigma_i X_i(t)dW_i(t) + \epsilon\sqrt{X_i(t)}dW'_i(t)$ is very rough. A better approach of matching those noise terms will likely lead to a more accurate estimation of the upper bound of the error.

The results for $\sigma_1 = \sigma_2 = 1.1$ are shown in Fig. 7. The total variation distance between these two measures are shown at the bottom of Fig. 7. It is not hard to see the difference is significantly larger than case $\sigma = 0.75$. The reason is that trajectories of Eq. (5.4) have high probability moving along the boundary in this parameter setting. This makes the probability of falling into the absorbing set $\partial\mathcal{X}$ much higher. Same as above, we compute the distribution of the coupling time and demonstrate the coupling time distribution (as well as the exponential tail) in Fig. 7 Top Left. The coupling in this example is slower so we choose $T = 12$ to




Fig. 6 (Case $\sigma_1 = \sigma_2 = 0.75$) *Upperpanel*: (Left) $\mathbb{P}(\tau^c > t)$ vs. t . (Right) QSD with demographic noise coefficient $\epsilon = 0.05$. *Lowerpanel*: (Left) Invariant density function for $\sigma = 0.75$. (Right) Total variation of QSD and invariant density function




Fig. 7 (Case $\sigma_1 = \sigma_2 = 1.1$) *Upperpanel*: (Left) $\mathbb{P}(\tau^c > t)$ vs. t . (Right) QSD with demographic noise coefficient $\epsilon = 0.05$. *Lowerpanel*: (Left) Invariant density function for $\sigma = 1.1$. (Right) Total variation of QSD and invariant density function

run Algorithm 5. The probability of killing before T is approximately 0.11186 and the total finite time error caused by the demographic noise is 0.06230. As a result, the error estimate of $d_w(\hat{\mu}_X, \hat{\mu}_Y)$ given in Lemma 4.2 is 0.1356. This is consistent with the numerical finding shown in Fig. 7 Bottom Right.

5.6 Chaotic Attractor

In this subsection, we consider a non-trivial 3D example that has interactions between chaos and random perturbations, called the Rossler oscillator. The random perturbation of the Rossler oscillator is

$$\begin{cases} dx = (-y - z)dt + \sigma dW_t^x \\ dy = (x + ay)dt + \sigma dW_t^y \\ dz = (b + z(x - c))dt + \sigma dW_t^z, \end{cases} \quad (5.5)$$

where $a = 0.2$, $b = 0.2$, $c = 5.7$, and W_t^x , W_t^y and W_t^z are independent Wiener processes. The strength of noise is chosen to be $\epsilon = 0.1$. This system is a representative example of chaotic ODE systems appearing in many applications of physics, biology and engineering. We consider Eq. (5.5) restricted to the box $D = [-15, 15] \times [-15, 15] \times [-1.5, 1.5]$. Therefore, it admits a QSD supported by D . In this example, a grid with $1024 \times 1024 \times 128$ mesh points is constructed on D .

It is difficult to use traditional PDE solver to compute a large scale 3D problem. To analyze the QSD of this chaotic system, we apply the blocked version of the Fokker–Planck solver studied in [10]. More precisely, a big mesh is divided into many “blocks”. Then we solve the optimization problem (3.2) in parallel. The collaged solution is then processed by the “shifting block” technique to reduce the interface error, which means the blocks are reconstructed such that the center of new blocks cover the boundary of old blocks. Then we let the solution from the first found serve as the reference data, and solve optimization problem (3.2) again based on new blocks. See [10] for the full details of implementation. In this example, the grid is further divided into $32 \times 32 \times 4$ blocks. We run the “shifting block” solver for 3 repeats to eliminate the interface error. The reference solution is generated by a Monte Carlo simulation with 10^9 samples. The killing rate is $\lambda = -0.473011$. Two “slices” of the solution, as seen




Fig. 8 (Rossler) Projections of 2 “slices” of the QSD of the Rossler system to the xy -plane. z -coordinates of 2 slices are $[-0.0352, 0.2578]$, $[1.1367, 1.4297]$. The solution is obtained by a half-block shift solver on $[-15, 15] \times [-15, 15] \times [-1.5, 1.5]$ with $1024 \times 1024 \times 128$ mesh points, $32 \times 32 \times 4$ blocks, and 10^9 samples

in Fig. 8, are then projected to the xy -plane for the sake of easier demonstration. See the caption of Fig. 8 for the coordinates of these two “slices”. The left picture in Fig. 8 shows the projection of the solution has both dense and sparse parts that are clearly divided. An outer “ring” with high density appears and the density decays quickly outside this “ring.” The right picture in Fig. 8 demonstrates the solution has much lower density when z -coordinate is larger than 1.

6 Conclusion

In this paper we provide some data-driven methods for the computation of quasi-stationary distributions (QSDs) and the sensitivity analysis of QSDs. Both of them are extended from the first author’s earlier work about invariant probability measures. When using the Fokker–Planck equation to solve the QSD, we find that the idea of using a reference solution with low accuracy to set up an optimization problem still works well for QSDs. And the QSD is not very sensitively dependent on the killing rate, which is given by the Monte Carlo simulation when producing the reference solution. The data-driven Fokker–Planck solver studied in this paper is still based on discretization. But we expect the mesh-free Fokker–Planck solver proposed in [10] to work for solving QSDs. In the sensitivity analysis part, the focus is on the relation between a QSD and the invariant probability measure of a “modified process”, because many interesting problems in applications fall into this category. The sensitivity analysis needs both a finite time truncation error and a contraction rate of the Markov transition kernel. The approach of estimating the finite time truncation error is standard. The contraction rate is estimated by using the novel numerical coupling approach developed in [21]. The sensitivity analysis of QSDs can be extended to other settings, such as the sensitivity against small perturbation of parameters, or the sensitivity of a chemical reaction process against its diffusion approximation. We will continue to study sensitivity analysis related to QSDs in our subsequent work.

References

1. Agresti, A., Coull, B.A.: Approximate is better than “exact” for interval estimation of binomial proportions. *Am. Stat.* **52**(2), 119–126 (1998)
2. Anderson, W.J.: *Continuous-Time Markov Chains: An Applications-Oriented Approach*. Springer (2012)
3. Barton, R. R., Schruben, L. W.: Uniform and bootstrap resampling of empirical distributions. In: *Proceedings of the 25th Conference on Winter Simulation*, pp. 503–508 (1993)
4. Benaïm, M., Champagnat, N., Villemonais, D.: Stochastic approximation of quasi-stationary distributions for diffusion processes in a bounded domain. In: *Annales de l’Institut Henri Poincaré, Probabilités et Statistiques*, vol. 57, pp. 726–739. Institut Henri Poincaré (2021)
5. Benaïm, M., Cloez, B., Panloup, F., et al.: Stochastic approximation of quasi-stationary distributions on compact spaces and applications. *Ann. Appl. Probab.* **28**(4), 2370–2416 (2018)
6. Burdzy, K., Hołyst, R., March, P.: A Fleming–viot particle representation of the dirichlet laplacian. *Commun. Math. Phys.* **214**(3), 679–703 (2000)
7. Cloez, B., Thai, M.-N.: Quantitative results for the Fleming–viot particle system and quasi-stationary distributions in discrete space. *Stoch. Process. Appl.* **126**(3), 680–702 (2016)
8. Collet, P., Martínez, S., San Martín, J.: *Quasi-Stationary Distributions: Markov Chains, Diffusions and Dynamical Systems*. Springer (2012)
9. Darroch, J.N., Seneta, E.: On quasi-stationary distributions in absorbing discrete-time finite markov chains. *J. Appl. Probab.* **2**(1), 88–100 (1965)
10. Dobson, M., Li, Y., Zhai, J.: An efficient data-driven solver for fokker-planck equations: algorithm and analysis. *Commun. Math. Sci.* accepted (2021)

11. Dobson, M., Li, Y., Zhai, J.: Using coupling methods to estimate sample quality of stochastic differential equations. *SIAM/ASA J. Uncertain. Quantif.* **9**(1), 135–162 (2021)
12. Ferrari, P.A., Martínez, S., Picco, P.: Existence of non-trivial quasi-stationary distributions in the birth-death chain. *Adv. Appl. Probab.* 795–813 (1992)
13. Hening, A., Li, Y.: Stationary distributions of persistent ecological systems. *J. Math. Biol.* **82**(7), 1–53 (2021)
14. Huillet, T.: On wright-fisher diffusion and its relatives. *J. Stat. Mech. Theory Exp.* **2007**(11), P11006 (2007)
15. Jacob, P.E., O’Leary, J., Atchadé, Y.F.: Unbiased markov chain monte carlo methods with couplings. *J. R. Stat. Soc. Ser. B (Stat. Methodol.)* **82**(3), 543–600 (2020)
16. Johndrow, J.E., Mattingly, J.C.: Error bounds for approximations of markov chains used in bayesian sampling. Preprint [arXiv:1711.05382](https://arxiv.org/abs/1711.05382) (2017)
17. Karatzas, I., Shreve, S.: *Brownian motion and stochastic calculus*, vol. 113. Springer (2014)
18. Kloeden, P.E., Platen, E.: *Stochastic differential equations*. In: *Numerical Solution of Stochastic Differential Equations*, pp. 103–160. Springer (1992)
19. Lai, Y.-C., Tél, T.: *Transient Chaos: Complex Dynamics on Finite Time Scales*, vol. 173. Springer (2011)
20. Li, Y.: A data-driven method for the steady state of randomly perturbed dynamics. *Commun. Math. Sci.* **17**(4) (2019)
21. Li, Y., Wang, S.: Numerical computations of geometric ergodicity for stochastic dynamics. *Nonlinearity* **33**(12), 6935 (2020)
22. Øksendal, B.: *Stochastic differential equations*. In: *Stochastic Differential Equations*, pp. 65–84. Springer (2003)
23. Pilipenko, A.: *An introduction to stochastic differential equations with reflection*, vol. 1. Universitätsverlag Potsdam (2014)
24. Van Doorn, E.A.: Quasi-stationary distributions and convergence to quasi-stationarity of birth-death processes. *Adv. Appl. Probab.* 683–700 (1991)
25. van Doorn, E.A., Schrijner, P.: Geometric ergodicity and quasi-stationarity in discrete-time birth-death processes. *ANZIAM J.* **37**(2), 121–144 (1995)
26. Villemonais, D.: Interacting particle systems and yaglom limit approximation of diffusions with unbounded drift. *Electron. J. Probab.* **16**, 1663–1692 (2011)

Publisher’s Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.