

Stochastic averaging and sensitivity analysis for two scale reaction networks

Araz Hashemi, Marcel Núñez, Petr Plecháč, and Dionisios G. Vlachos

Citation: [The Journal of Chemical Physics](#) **144**, 074104 (2016); doi: 10.1063/1.4942008

View online: <https://doi.org/10.1063/1.4942008>

View Table of Contents: <http://aip.scitation.org/toc/jcp/144/7>

Published by the [American Institute of Physics](#)

Articles you may be interested in

[Steady state likelihood ratio sensitivity analysis for stiff kinetic Monte Carlo simulations](#)

The Journal of Chemical Physics **142**, 044108 (2015); 10.1063/1.4905957

[A graph-theoretical kinetic Monte Carlo framework for on-lattice chemical kinetics](#)

The Journal of Chemical Physics **134**, 214115 (2011); 10.1063/1.3596751

[Efficient estimators for likelihood ratio sensitivity indices of complex stochastic dynamics](#)

The Journal of Chemical Physics **144**, 104107 (2016); 10.1063/1.4943388

[Parallel kinetic Monte Carlo simulation framework incorporating accurate models of adsorbate lateral interactions](#)

The Journal of Chemical Physics **139**, 224706 (2013); 10.1063/1.4840395

[Parallel replica dynamics method for bistable stochastic reaction networks: Simulation and sensitivity analysis](#)

The Journal of Chemical Physics **147**, 234110 (2017); 10.1063/1.5017955

[Acceleration and sensitivity analysis of lattice kinetic Monte Carlo simulations using parallel processing and rate constant rescaling](#)

The Journal of Chemical Physics **147**, 164103 (2017); 10.1063/1.4998926

PHYSICS TODAY

WHITEPAPERS

ADVANCED LIGHT CURE ADHESIVES

READ NOW

Take a closer look at what these
environmentally friendly adhesive
systems can do

PRESENTED BY

 **MASTERBOND**
ADHESIVES | SEALANTS | COATINGS

Stochastic averaging and sensitivity analysis for two scale reaction networks

Araz Hasemi,^{1,a)} Marcel Núñez,^{2,b)} Petr Plecháč,^{1,c)} and Dionisios G. Vlachos^{2,d)}

¹Department of Mathematical Sciences, University of Delaware, Newark, Delaware 19716, USA

²Department of Chemical and Biomolecular Engineering, University of Delaware, Newark, Delaware 19716, USA

(Received 15 September 2015; accepted 2 February 2016; published online 17 February 2016)

In the presence of multiscale dynamics in a reaction network, direct simulation methods become inefficient as they can only advance the system on the smallest scale. This work presents stochastic averaging techniques to accelerate computations for obtaining estimates of expected values and sensitivities with respect to the steady state distribution. A two-time-scale formulation is used to establish bounds on the bias induced by the averaging method. Further, this formulation provides a framework to create an accelerated “averaged” version of most single-scale sensitivity estimation methods. In particular, we propose the use of a centered ergodic likelihood ratio method for steady state estimation and show how one can adapt it to accelerated simulations of multiscale systems. Finally, we develop an adaptive “batch-means” stopping rule for determining when to terminate the micro-equilibration process. © 2016 AIP Publishing LLC. [<http://dx.doi.org/10.1063/1.4942008>]

I. INTRODUCTION

Stochastic simulations have been an essential tool in analyzing reaction networks encountered in biology, catalysis, and materials growth. However, it is common place for reaction networks to exhibit a large disparity in time scales. These multi-scale stochastic reaction networks can impose an enormous computational burden in order to simulate them exactly. Exact techniques require computation of every reaction at the fastest time-scale, resulting in an exponentially growing load to observe dynamics on the slowest time-scale. Many works have attempted to develop approximate algorithms which allow faster computation with minimal loss of accuracy.^{1–14}

One approach, which we refer to as stochastic averaging, takes its inspiration from classical singular perturbation theory of ordinary differential equations.^{8–12} The idea is that the fast dynamics come to quasi-equilibrium before the slow dynamics take effect, hence one may model the slow time scale dynamics with their averages against the steady-state distribution of the fast dynamics. By estimating the steady-state expectations of the slow propensities, one can then jump the system ahead to the next slow reaction and advance the time clock on the slow scale (skipping over needless computations of fast reactions).

In addition, one often desires the sensitivities of the system $S_f(\theta_i) = \frac{\partial}{\partial \theta_i} \mathbb{E}_{\theta} \{f(X(t))\}$ with respect to the reaction parameters θ_i . The sensitivities give important insight into the system, indicating directions for gradient-descent type optimization as well as determining bounds for quantifying the uncertainty.¹⁵ Current techniques for estimating the

sensitivities have *large variance*, requiring many more samples than those for estimating $\mathbb{E}_{\theta} \{f(X)\}$ alone.^{16–19} Thus, computing sensitivities of multi-scale systems using single-scale techniques is often a computationally intractable problem.

In this work, we use results from Two-Time-Scale (TTS) Markov chains²⁰ to show that the error induced by stochastic averaging algorithms is inverse to the scale disparity in the system. As opposed to the previous approaches of transforming the system variables into auxiliary fast and slow variables,^{11,12} we partition the underlying (discrete) state space and derive a singular perturbation expansion of the corresponding probability measure. The first order term can then be identified from computables of the averaged process, leading to a rigorous theoretical framework for applying singular perturbation averaging for stochastic systems.

Furthermore, this new formulation allows one to identify a macroscopic “averaged” reaction network on a reduced state space whose time-steps are on the macro (slow) time-scale. Thus, it provides a framework for applying single-scale sensitivity analysis techniques to the multi-scale system. Previous works have exploited similar model reduction techniques to estimate sensitivities via finite differences¹⁴ or a “truncated” version of a likelihood ratio estimator.²¹ This work develops an accelerated “Two-Time-Scale” version of the likelihood ratio (Girsanov transform) method^{18,22–24} for estimating system sensitivities of the multiscale system. The TTS-LR method computes sensitivity reweighting coefficients of the macro (reduced-state) process using a representation in terms of the steady-state sensitivities of the micro (fast) process. These micro-level sensitivities can in turn be computed online during the micro-equilibration process.

The outline of the remainder of the paper is as follows. Section II gives the theoretical basis of the paper. The two-time-scale formulation is presented and error bounds are

^{a)}Author to whom correspondence should be addressed. Electronic mail: araz@udel.edu

^{b)}Electronic mail: mpnunez@udel.edu

^{c)}Electronic mail: plechac@udel.edu

^{d)}Electronic mail: vlachos@udel.edu

established. Section III then uses the TTS framework for the purpose of sensitivity analysis. The centered ergodic likelihood ratio estimator is described for single-scale steady-state sensitivity analysis, and is then adapted to the multiscale system. Section IV develops a batch-means stopping rule for determining when the micro-scale system has come to equilibrium. Numerical results are presented in Section V supporting the effectiveness of the methods presented. Concluding remarks are given in Section VI, and proofs of theorems are relegated to the Appendices A and B.

II. FORMULATION

A. Markov chain model of reaction networks

We briefly review the Markov chain model of reaction networks. While our motivation stems from chemical reaction networks, we note that much of the formulation carries over to general Markov chains on integer lattices.

Suppose we have d species described by $X(t) = [X_1(t), X_2(t), \dots, X_d(t)] \in \mathcal{M} \subset \mathbb{Z}^d$ and M reactions r_1, r_2, \dots, r_M . In stochastic reaction networks, one views $X(t)$ as a continuous-time Markov chain (CTMC) in the state space \mathcal{M} . When reaction r fires at time t , the state is updated by the stoichiometric vector ζ_r so that $X(t) = X(t-) + \zeta_r$. Given the set of reaction parameters $\theta = [\theta_1, \theta_2, \dots]$, one characterizes the probabilistic evolution of $X(t)$ by the *propensity functions*

(intensity functions) $\lambda_r(x; \theta)$. The propensity functions are such that, given $X(t) = x$, the probability of one or more firing of reaction r during time $(t, t+h]$ is $\lambda_r(x; \theta)h + o(h)$ as $h \rightarrow 0$; i.e., $\lambda_r(x; \theta)$ is the instantaneous rate/probability of reaction r firing.

A common model for the propensity functions is that of *mass-action kinetics*. Under this assumption, the propensity functions are of the form

$$\lambda_r(x; \theta) = \theta_r \cdot b_r(x) = \theta_r \cdot \prod_{i=1}^d \frac{x_i!}{(x_i - \nu_{r,i})!} \mathbb{I}_{\{x - \nu_{r,i} \geq 0\}}, \quad (1)$$

where $\nu_{r,i}$ is the number of molecules of species i required for reaction r to fire. Mass action kinetics assumes the system is well-mixed, so molecular interactions are proportional to their counts.

From the propensity functions $\lambda_r(x; \theta)$, one can construct the infinitesimal generator $\mathcal{Q} = \mathcal{Q}(\theta)$ of the Markov chain. Viewed as an operator on functions f of the state-space \mathcal{M} , we have

$$(\mathcal{Q}f)(x) = \sum_{r=1}^M \lambda_r(x; \theta) (f(x + \zeta_r) - f(x)).$$

For finite state-space \mathcal{M} (bounded molecule counts), one may also view the generator $\mathcal{Q}(\theta)$ as a matrix. While the state-space is typically intractably large, the generator is sparse with only $M + 1$ non-zero entries in each row,

$$Q(\theta) = x \begin{pmatrix} \dots & x + \zeta_2 & \dots & & & & \\ \vdots & & & & & & \\ \dots & \lambda_2(x; \theta) & \dots & -\lambda_0(x; \theta) & \dots & \lambda_1(x; \theta) & \dots \\ \vdots & & & & & & \end{pmatrix},$$

where $\lambda_0(x; \theta) = \sum_{r=1}^M \lambda_r(x; \theta)$.

Writing $R_r(t)$ to be the counting process representing how many reactions of type r have fired by the time t , we have that $X(t) = X(0) + \sum_{r=1}^M R_r(t) \zeta_r$. Using the *random time change* representation,^{25,26} we write $X(t)$ as

$$X(t) = X(0) + \sum_{r=1}^M Y_r \left(\int_0^t \lambda_r(X(s); \theta) ds \right) \zeta_r, \quad (2)$$

where $Y_r(\cdot)$ are independent unit-rate Poisson processes. This representation is tremendously useful in conducting analysis of the trajectories. In particular, it leads to formulations of the next-reaction method^{27,28} and interpreting simulated trajectories in the path-space to allow for coupling paths^{19,29,30} as well as path-wise differentiation.^{16,17}

When simulating exact trajectories (using any exact method; direct SSA, next-reaction, etc.), the propensity functions $\lambda_r(x; \theta)$ probabilistically determine both the time between reactions Δt as well as the next reaction r^* to fire. The likelihood that reaction r_k is the next to fire is proportional

to its propensity $\lambda_k(x; \theta)$; i.e., $\mathbb{P}_{\theta, x} \{r^* = r_k\} \propto \lambda_k(x; \theta)$. The time between reactions has an exponential distribution with the rate $\lambda_0(x; \theta) = \sum_{r=1}^M \lambda_r(x; \theta)$; i.e., $\Delta t \sim \text{Exp}(\lambda_0(x; \theta))$ with the mean $\mathbb{E}_{x, \theta} \{\Delta t\} = 1/\lambda_0(x, \theta)$.

Multi-scale dynamics occur when the propensity functions have large magnitude disparities. If $\lambda_k(x; \theta) \gg \lambda_j(x; \theta)$ for all $j \neq k$, then $\mathbb{P}\{r^* = r_k\} \approx 1$ and $\Delta t \sim \text{Exp}(\lambda_0(x; \theta)) \approx \text{Exp}(\lambda_k(x; \theta))$. Thus, with a high probability the next reaction in an exact trajectory will be r_k and the time clock will advance on the order of $1/\lambda_k(x; \theta)$. Such multi-scale networks then require an enormous number of computations to sample “slow” reactions and reach the required time horizon for the entire system to relax.

B. Two-time-scale reaction networks

We now consider reaction networks with two scales of dynamics. For further motivation and discussion of reaction networks with multiple time-scales, we refer readers to Refs. 11–14 and references therein. We instead focus on our

formulation for the separation of time-scales and the averaged process via the partitioning of the state space into “fast-classes.” Though analogous to the techniques of transforming the species variables into auxiliary fast/slow variables^{11,12} or projecting to remainder spaces,¹⁴ the direct partitioning of the state space will allow us to construct a singular perturbation expansion of the probability measure and establish the rate of convergence of such averaging methods. In addition, it provides a framework for applying likelihood ratio type sensitivity estimates to the averaged process as we shall see in the sequel.

Suppose a reaction network has species $[X_1, X_2, \dots, X_d]$ and reactions r_1, \dots, r_M . We shall assume that the propensity functions $\lambda_r(x; \theta)$ are of the form (1) (mass-action kinetics, though other forms may also be treated), and that each reaction is indexed by its own reaction parameter θ_r . We assume that there is a scale disparity in the reaction parameters between a set of “fast reactions” and a set of “slow reactions.” Thus we can write $\theta = [\theta_1, \dots, \theta_M] = [\alpha/\varepsilon, \beta]$, where $\beta = [\beta_1, \beta_2, \dots, \beta_{M_s}]$ are the slow reaction parameters, $\varepsilon \ll 1$ is the stiffness parameter, and $\alpha = [\alpha_1, \dots, \alpha_{M_f}]$ are the underlying (rescaled) reaction parameters for the fast reactions.

To ease referencing, we will often index reactions and propensity functions directly by their reaction parameter. For example, r_{β_i} is the reaction with reaction parameter β_i and propensity function $\lambda_{\beta_i}(x; \theta) = \lambda_{\beta_i}(x; \beta) = \beta_i b_{\beta_i}(x)$ (where b_{β_i} is given by (1)). For the fast reactions α_i , we use $\lambda_{\alpha_i}^\varepsilon(x; \theta) = \lambda_{\alpha_i}^\varepsilon(x; \alpha) = (\alpha_i/\varepsilon) b_{\alpha_i}(x)$ to denote the exact propensity function and $\lambda_{\alpha_i}(x; \theta) = \alpha_i b_{\alpha_i}(x)$ to denote the rescaled version.

Let $X^\varepsilon(t)$ denote the Markov chain determined by the exact propensity functions $\lambda_\alpha^\varepsilon(x; \theta)$ and $\lambda_\beta(x; \theta)$. We can write the generator $Q^\varepsilon = Q^\varepsilon(\alpha, \beta)$ of the exact process as before, and observe that

$$Q^\varepsilon(\alpha, \beta) = \frac{1}{\varepsilon} \tilde{Q}(\alpha) + \hat{Q}(\beta), \quad (3)$$

where $\hat{Q}(\beta)$ is the generator of the chain under only the slow dynamics (determined by slow reactions r_β), and $\tilde{Q}(\alpha)$ is the generator of the chain under only the fast dynamics (with the rescaled propensity functions $\lambda_\alpha(x; \alpha)$).

As $\varepsilon \rightarrow 0$ only the fast reactions r_α fire, and so we define an equivalence relation on states $s \in \mathcal{M}$, by $s_i \leftrightarrow s_j$ if they are mutually accessible through only fast reactions. This defines a partition of the state space \mathcal{M} into “fast-classes” \mathcal{M}_k which are by construction the invariant (irreducible) classes of \mathcal{M} under $\tilde{Q}(\alpha)$; e.g.,

$$\mathcal{M} = \bigcup_{k=1}^{N_C} \mathcal{M}_k = \{x_1^{(1)}, x_2^{(1)}, \dots, x_{m_1}^{(1)}, x_1^{(2)}, \dots, x_{m_2}^{(2)}, \dots\},$$

where N_C are the number of invariant “fast-classes,” and $m_k = |\mathcal{M}_k|$ is the number of states inside fast-class \mathcal{M}_k . For ease of presentation, in the present discussion we shall assume the state space is finite.

Assumption 2.1 (finite state space). The state space \mathcal{M} is finite, such that $|\mathcal{M}| = m$. Thus the number of fast classes

$N_C < \infty$ and the number of states in each fast-class $m_k < \infty$, so that $m = m_1 + m_2 + \dots + m_{N_C}$.

Assumption 2.1 is made only to simplify the discussion. One may also treat the infinite state case with some mild additional conditions on $\tilde{Q}(\alpha)$ and $\hat{Q}(\beta)$ to ensure non-explosiveness and ergodicity of the underlying (rescaled) chain.²⁰ For example, one can require that the reactions which increase the total molecule count have propensities which grow at most linearly with the total molecule count.^{13,14} In addition, we shall impose the following assumption.

Assumption 2.2 (recurrent states). Each state of \mathcal{M} is recurrent, so that there are no absorbing/transient states.

Assumption 2.2 is satisfied if, for example, all reactions are reversible (or often times if only the fast reactions are reversible). One may also treat the case with transient/absorbing classes with some additional stability assumptions to ensure the fast dynamics decay to steady-state; see Section 4.4 of Ref. 20 for more details.

Under Assumption 2.2, we can reorder the state space so that

$$\tilde{Q}(\alpha) = \text{diag}[\tilde{Q}^{(1)}(\alpha), \tilde{Q}^{(2)}(\alpha), \dots, \tilde{Q}^{(N_C)}(\alpha)]$$

is block-diagonal. Here, one can view the generators $\tilde{Q}^{(k)}(\alpha)$ as the restriction of $\tilde{Q}(\alpha)$ to the (irreducible) fast-class \mathcal{M}_k (fast-only dynamics when $X(0) \in \mathcal{M}_k$). In light of the finite state-space and positive recurrence, each $\tilde{Q}^{(k)}(\alpha)$ is ergodic and has a stationary (steady-state) probability measure $\tilde{\pi}^{(k)} = \tilde{\pi}^{(k)}(\alpha)$ such that $\tilde{\pi}^{(k)} \tilde{Q}^{(k)} = \mathbf{0}$ (with $\tilde{\pi}^{(k)}, \mathbf{0}$ interpreted as row vectors).

Using the above formulation, we can restate the averaging principle⁹⁻¹⁴ as follows. For small ε and $X(0) \in \mathcal{M}_k$, $X^\varepsilon(\cdot)$ will relax to its steady-state distribution $\pi^{(k)}$ on the micro time-scale εt before any slow reaction fires (on the macro time scale t). Thus, one can use the stationary average of the slow propensity functions

$$\lambda_{\beta_j}(X(t); \beta, X(0) \in \mathcal{M}_k) \sim \bar{\lambda}_{\beta_j}(\mathcal{M}_k; \theta) \stackrel{\triangle}{=} \mathbb{E}_{\tilde{\pi}^{(k)}(\alpha)} \{\lambda_{\beta_j}(X; \beta)\} \quad (4)$$

to determine the distribution of time until the next slow reaction as well as the probabilities for the next slow reaction being r_{β_j} . These can then be used to simulate a trajectory of the slow (macro-scale) process. We shall further develop this idea more precisely in the remainder.

Write $m = |\mathcal{M}|$, and $m_k = |\mathcal{M}_k|$ as before, so that $Q^\varepsilon(\theta), \tilde{Q}(\alpha), \hat{Q}(\beta) \in \mathbb{R}^{m \times m}$, $\tilde{Q}^{(k)}(\alpha) \in \mathbb{R}^{m_k \times m_k}$, and $\pi^{(k)}(\alpha) \in \mathbb{R}^{1 \times m_k}$. Write $\tilde{\pi}(\alpha) = \text{diag}[\tilde{\pi}^{(1)}, \tilde{\pi}^{(2)}, \dots, \tilde{\pi}^{(N_C)}] \in \mathbb{R}^{N_C \times m}$. Write $\mathbf{1}_{m_k}$ for $[1, 1, \dots, 1]' \in \mathbb{R}^{m_k \times 1}$ and $\tilde{\mathbf{1}} = \text{diag}[\mathbf{1}_{m_1}, \mathbf{1}_{m_2}, \dots, \mathbf{1}_{m_{N_C}}]$.

With $\tilde{\pi}(\alpha)$ describing the limit behavior inside each fast-class on the micro time scale, one can then consider the distribution of the exact system X^ε on the macro time scale. Heuristically, one expects a trajectory to enter a fast-class of states \mathcal{M}_{k_1} and quickly iterate through many fast reactions until the distribution of the trajectory reaches the steady-state $\tilde{\pi}^{(k_1)}$. Eventually, a slow reaction will fire to move the trajectory

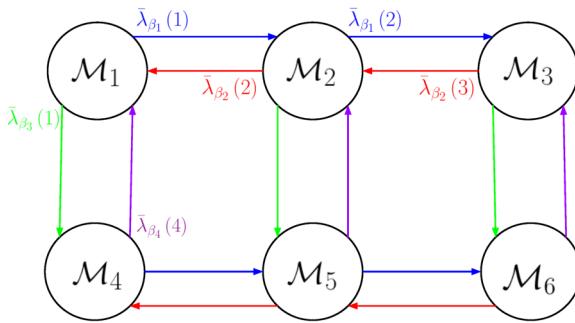


FIG. 1. A depiction of the macro chain of a two-time-scale reaction network. Each fast-class \mathcal{M}_k is an equivalence class of states accessible by fast-reactions r_α . Slow reactions r_β carry each fast-class to a unique new fast-class. The averaged system $\bar{X}(t)$ forms a meta “macro” Markov chain among fast-classes with propensities $\bar{\lambda}_\beta(\bar{X}; \boldsymbol{\beta}, \tilde{\pi}(\boldsymbol{\alpha}))$.

to a new fast-class \mathcal{M}_{k_2} (see Figure 1). Writing

$$\bar{Q} = \bar{Q}(\boldsymbol{\alpha}, \boldsymbol{\beta}) \stackrel{\Delta}{=} [\tilde{\pi}(\boldsymbol{\alpha}) \cdot \bar{Q}(\boldsymbol{\beta}) \cdot \bar{\mathbf{1}}] \in \mathbb{R}^{N_C \times N_C}, \quad (5)$$

we see that \bar{Q} is itself a generator of an “averaged” CTMC reaction network, whose “states” correspond to fast-classes \mathcal{M}_k . Write $\bar{X}(t)$ for the “averaged” process generated by \bar{Q} . Together, \bar{Q} and $\bar{X}(t)$ describe the limit (as $\varepsilon \rightarrow 0$) of the average rate that the exact process $X^\varepsilon(t)$ moves between the fast-classes $\{\mathcal{M}_k\}_{k=1}^{N_C}$ via slow reactions.

Furthermore, we can identify the elements of \bar{Q} from the steady-state averages of the slow propensity functions. First, note that every slow reaction carries each fast class to a unique new fast-class; that is, if $x \leftrightarrow y$ and $\lambda_\beta(x), \lambda_\beta(y) > 0$, then $r_\beta(x) \leftrightarrow r_\beta(y)$. Thus, $r_\beta(\mathcal{M}_k)$ is well-defined. Then, using the form of \bar{Q} together with (5), we have

$$\bar{Q}_{k_1, k_2} = \sum_{\beta \in \boldsymbol{\beta}}_{r_\beta(\mathcal{M}_{k_1}) = \mathcal{M}_{k_2}} \bar{\lambda}_\beta(\mathcal{M}_{k_1}; \boldsymbol{\theta}) \quad (6)$$

for $k_1 \neq k_2$, and similarly we see that

$$\bar{Q}_{k_1, k_1} = - \sum_{\beta \in \boldsymbol{\beta}} \bar{\lambda}_\beta(\mathcal{M}_{k_1}; \boldsymbol{\theta}) \stackrel{\Delta}{=} -\bar{\lambda}_{\beta_0}(\mathcal{M}_{k_1}; \boldsymbol{\theta}). \quad (7)$$

With this formulation, we see that generator \bar{Q} corresponds to a meta “macro” reaction network with the state-space $\bar{\mathcal{M}} = \{\mathcal{M}_1, \mathcal{M}_2, \dots, \mathcal{M}_{N_C}\}$, reactions $\{r_\beta : \beta \in \boldsymbol{\beta}\}$, and propensities $\{\bar{\lambda}_\beta(\bar{X}; \boldsymbol{\theta}) : \beta \in \boldsymbol{\beta}\}$. Figure 1 depicts such a macro chain for the macro-process $\bar{X}(t)$.

If we can estimate the average slow-propensities $\bar{\lambda}_\beta(\mathcal{M}_k; \boldsymbol{\theta})$ within each fast-class (say, through ergodic time averages of the fast-only process), then one can simulate a trajectory of the macro-process $\bar{X}(t)$ from these average propensities using *any* single-scale Monte Carlo simulation (e.g., direct SSA, next-reaction, etc.). Furthermore, if one is ultimately concerned with estimating $\mathbb{E}_\theta \{f(X^\varepsilon(t))\}$ for some observable (quantity of interest) $f : \mathcal{M} \rightarrow \mathbb{R}$, then one can define an augmented functional \bar{f} on $\bar{\mathcal{M}}$ by

$$\bar{f}(\mathcal{M}_k; \boldsymbol{\alpha}) \stackrel{\Delta}{=} \mathbb{E}_{\tilde{\pi}^{(k)}(\boldsymbol{\alpha})} \{f(X)\} = \sum_{x \in \mathcal{M}_k} f(x) \tilde{\pi}_x^{(k)}(\boldsymbol{\alpha}). \quad (8)$$

It shall be shown that $\mathbb{E}_\theta \{\bar{f}(\bar{X}(t))\} \approx \mathbb{E}_\theta \{f(X^\varepsilon(t))\}$ for large enough t and sufficiently small ε .

To illustrate how one can implement the averaging scheme to generate macro-trajectories, we present the following TTS version of the direct SSA (since it is the most succinct to write). In this case, the TTS SSA is essentially the same algorithm as in Refs. 9 and 11. However, we emphasize that the same method can be used to create a TTS version of any exact method defined by the propensity functions. In particular, one can just as easily construct an analogous TTS next-reaction type algorithm^{28–30} for tightly coupled trajectories.

Algorithm 1 (TTS-SSA). To simulate a trajectory of the macro-process $\bar{X}(T)$ until macro time-horizon T_{final} :

- (1) Initialize x at a macro time T ; $x \in \mathcal{M}_k$ for some (unknown) k .
- (2) Simulate the fast-only reaction network $\bar{Q}^{(k)}(\boldsymbol{\alpha})$ until time-averages of observable f and slow propensities $\bar{\lambda}_\beta$ relax to steady-state:

$$\begin{aligned} \frac{1}{t} \int_0^t f(\bar{X}^{(k)}(s)) ds &\rightarrow \mathbb{E}_{\tilde{\pi}^{(k)}(\boldsymbol{\alpha})} \{f(\bar{X})\} \equiv \bar{f}(k), \\ \frac{1}{t} \int_0^t \lambda_\beta(\bar{X}^{(k)}(s); \boldsymbol{\beta}) ds &\rightarrow \mathbb{E}_{\tilde{\pi}^{(k)}(\boldsymbol{\alpha})} \{\lambda_\beta(\bar{X}; \boldsymbol{\beta})\} \\ &\equiv \bar{\lambda}_\beta(k; \boldsymbol{\theta}). \end{aligned}$$
- (3) Observe terminal state $\bar{x}^{(k)} \sim \tilde{\pi}^{(k)}$. Compute $\bar{\lambda}_{\beta_0} = \sum_{j=1}^{M_S} \bar{\lambda}_{\beta_j}$.
- (4) Sample time to next slow reaction: $\Delta T \sim \text{Exp}(\bar{\lambda}_{\beta_0})$.
- (5) Sample next slow reaction to fire $\beta^* \sim 1/\bar{\lambda}_{\beta_0} [\bar{\lambda}_{\beta_1}, \dots, \bar{\lambda}_{\beta_{M_S}}]$.
- (6) Update macro time $T \leftarrow T + \Delta T$ and move to the next fast class by taking $x = \bar{x}^{(k)} + \zeta_{\beta^*}$.
- (7) Return to (1) until macro time horizon T_{final} is reached.

Write $\bar{p}_T = \bar{p}_T(\bar{X}; \boldsymbol{\alpha}, \boldsymbol{\beta})$ for the probability measure (on $\bar{\mathcal{M}}$) induced by the averaged process \bar{X} at time T . At the end of a TTS simulation, one obtains a terminal state $X(T) = x \sim p_T^0 = p_T^0(X; \boldsymbol{\alpha}, \boldsymbol{\beta})$, where p_T^0 is the probability measure on \mathcal{M} induced by the last state observed from the terminal fast-class. Thus, p_T^0 is determined by \bar{p}_T and $\tilde{\pi}(\boldsymbol{\alpha})$. Write $p_T^\varepsilon = p_T^\varepsilon(X^\varepsilon; \boldsymbol{\alpha}, \boldsymbol{\beta})$ for the probability measure on \mathcal{M} induced from the exact process X^ε . Since p_T^ε is the distribution we would see from an exact simulation, and p_T^0 is the distribution from the TTS simulation, the question becomes: What is the error of p_T^0 from p_T^ε ? One can take a singular perturbation expansion of p_T^ε in terms of ε and identify the leading term as p_T^0 to obtain the following result, whose proof we defer to [Appendix A](#).

Theorem 2.3 (Error in Probability). Let $\tilde{\kappa} = -\frac{1}{2} \max \{\text{Re}(\nu) : \nu \text{ is a non-zero eigenvalue of a } \bar{Q}^{(k)}\}$. Then under Assumptions 2.1, 2.2, we have

$$\|p_T^0 - p_T^\varepsilon\| \leq O(\varepsilon + \exp\{-\tilde{\kappa}T/\varepsilon\}), \quad (9)$$

where $\|\cdot\|$ denotes the l_2 norm.

In Theorem 2.3, $\tilde{\kappa}$ is the slowest rate of convergence of $\bar{Q}^{(k)}$ to the steady-state $\tilde{\pi}^{(k)}$ among all fast-classes \mathcal{M}_k .

Thus, as long as the macro time horizon T is large enough to ensure the fast dynamics have relaxed to steady state ($T > -\varepsilon/\tilde{\kappa} \log(\varepsilon)$), then the error becomes $\|p_T^0 - p_T^\varepsilon\| \leq O(\varepsilon)$.

Writing $\pi^0(\alpha, \beta)$ for the stationary distribution corresponding to the TTS probability measure $p_T^0(\alpha, \beta)$, it is not hard to see that $\pi^0(\alpha, \beta) = \bar{\pi} \cdot \tilde{\pi}$, the product of the steady-state distribution between fast-classes $\bar{\pi}(\alpha, \beta)$ and the steady-state distribution within fast-classes $\tilde{\pi}(\alpha)$. Write π^ε for the steady-state distribution corresponding to the exact process generated by Q^ε . Then using Theorem 2.3 and exponential convergence to the steady state, we obtain the following error bounds (see Appendix A 2).

Corollary 2.4 (error in expectation). *Under Assumptions 2.1 and 2.2, $\|\pi^0 - \pi^\varepsilon\| \leq O(\varepsilon)$ and $\|p_T^0 - p_T^\varepsilon\| \leq O(\varepsilon)$ for sufficiently large T . Thus, for all bounded functions f on the state space \mathcal{M} ,*

$$\begin{aligned} \left| \mathbb{E}_{\bar{p}_T} \{ \bar{f}(\bar{X}(T)) \} - \mathbb{E}_{p_T^\varepsilon} \{ f(X^\varepsilon(T)) \} \right| &\leq \|f\|_\infty \|p_T^0 - p_T^\varepsilon\| \\ &\leq O(\varepsilon) \end{aligned} \quad (10)$$

$$\left| \mathbb{E}_{\bar{\pi}} \{ \bar{f}(\bar{X}) \} - \mathbb{E}_{\pi^\varepsilon} \{ f(X^\varepsilon) \} \right| \leq \|f\|_\infty \|\pi^0 - \pi^\varepsilon\| \leq O(\varepsilon).$$

Corollary 2.4 is of great practical use, as it says that the expected value of the macro-process $\bar{X}(T)$ with macro-observable $\bar{f} : \mathcal{M} \rightarrow \mathbb{R}$ provides an $O(\varepsilon)$ estimate of the expected value of the exact system $X^\varepsilon(T)$ with observable $f : \mathcal{M} \rightarrow \mathbb{R}$. Since we can use TTS algorithms (such as Algorithm 1) to quickly generate trajectories of $\bar{X}(T)$ while estimating the macro-observable $\bar{f}(\mathcal{M}_k)$ at each state along the way, this provides a method to very quickly generate estimates of $\mathbb{E}_{p_T^\varepsilon} \{ f(X^\varepsilon(T)) \}$ with at most $O(\varepsilon)$ bias. As $\varepsilon \rightarrow 0$, the bias decreases linearly while the computational savings increase as $O(1/\varepsilon)$.

III. TWO-TIME-SCALE SENSITIVITY ANALYSIS

Computing the system sensitivities $S_{f,T}(\theta_i) \stackrel{\Delta}{=} \frac{\partial}{\partial \theta_i} \mathbb{E}_\theta \{ f(X(T)) \}$ with respect to reaction parameters $\theta_i \in \theta$ provides great insight into the model. As such, numerous works have constructed and analyzed methods to estimate the sensitivities from sample trajectories of the system.^{16-19,21,23,29,31,32}

Different methods work better for different systems or different criteria, but all methods have higher (sometimes stupendously higher) variance in the estimation of $S_{f,T}(\theta_i)$ compared to the estimation of $\mathbb{E}_\theta \{ f(X(T)) \}$, thus requiring a very large number of samples to estimate the sensitivity precisely. If the system is stiff (as in (3)) so that each exact trajectory $X^\varepsilon(T)$ requires a prohibitively large computational load, then the large number of sample paths required to estimate the sensitivity $S_{f,T}^\varepsilon(\theta_i) \stackrel{\Delta}{=} \partial_{\theta_i} \mathbb{E}_{p_T^\varepsilon(\theta)} \{ f(X^\varepsilon(T)) \}$ makes the problem computationally intractable.

Corollary 2.4 gives that the expectation of the macro “averaged” reaction network $\bar{f}(\bar{X}(T))$ gives an accurate approximation of the expectation of the exact network; $\mathbb{E}_{\bar{p}_T(\theta)} \{ \bar{f}(\bar{X}(T)) \} = \mathbb{E}_{p_T^\varepsilon(\theta)} \{ f(X^\varepsilon(T)) \} + O(\varepsilon)$. A natural question to ask is whether the sensitivities of the exact system converge to the sensitivities of the averaged

system (i.e., if the derivative commutes with the limit). Using the recent result of Ref. 14, one can derive the following (the details are deferred to Appendix A 3).

Proposition 3.1 (convergence of sensitivities).

$$\lim_{\varepsilon \rightarrow 0} S_{f,T}^\varepsilon(\theta_i) = \bar{S}_{f,T}(\theta_i) \stackrel{\Delta}{=} \frac{\partial}{\partial \theta_i} \mathbb{E}_{\bar{p}_T(\theta)} \{ \bar{f}(\bar{X}(T)) \}. \quad (11)$$

Thus, if we can compute the sensitivity of the macro reaction network $\bar{f}(\bar{X}(T))$ (whose sample paths have orders of magnitude less cost to simulate than the exact stiff network $f(X^\varepsilon(T))$), then this provides an accurate estimate of the exact sensitivity. Furthermore, since $\bar{X}(\cdot)$ is formulated as a reaction network with propensities $\{\lambda_\beta(x, \theta) : \beta \in \beta\}$ and observable values $\bar{f}(\mathcal{M}_k)$ (both of which are estimated during a TTS simulation), we can apply most of the existing single-scale sensitivity estimation methods to estimate $\bar{S}_{f,T}(\theta_i)$ and thus $S_{f,T}^\varepsilon(\theta_i)$.

We note that (11) gives that the sensitivity of the exact system converges to the sensitivity of the averaged system, but does not give the rate of convergence. Currently, this is an open question. Since from (10) we have the expectation converges at a rate $O(\varepsilon)$, one might suspect that the sensitivity also converges at rate $O(\varepsilon)$, at least for certain classes of networks (e.g., linear propensities). Computational tests suggest this is the case, and the subject of ongoing work³³ is to establish a precise statement on the rate of convergence. However, the remainder of this work shall focus on the development and practical implementation of a multiscale likelihood ratio estimator of the limit sensitivity $\bar{S}_{f,T}(\theta_i)$.

In what follows, we review the likelihood ratio method for computing system sensitivities for single-scale reaction networks. We address the high variance of the likelihood ratio method by combining centering and ergodic-averaging techniques for a Centered Ergodic Likelihood Ratio (CELR) estimator with dramatically smaller variance for steady-state sensitivity estimation. We then derive a two-time-scale version that allows one to estimate the full gradient of a stiff system using any TTS Monte Carlo method for simulating a macro-trajectory.

A. Likelihood ratio methods

LR methods^{18,21-24,34} (aka the Girsanov transform method) attempt to compute the derivative by reweighting the observed trajectory by its “score” function of the density. Here, one views θ as parametrizing the probability measure on the path-space $P(\cdot, t; \theta)$. If $P(\cdot, t; \theta)$ is differentiable with respect to θ_i , then under mild regularity conditions we have

$$\begin{aligned} S_{f,t}(\theta_i) &\stackrel{\Delta}{=} \frac{\partial}{\partial \theta_i} \mathbb{E}_{\theta^0} \{ f(X(t)) \} \\ &= \int_{\Omega} f(X(t, \omega)) \frac{\frac{\partial}{\partial \theta_i} |_{\theta^0} P(d\omega, t; \theta)}{P(d\omega, t; \theta)} P(d\omega, t; \theta^0) \\ &= \mathbb{E}_{\theta^0} \{ f(X(t)) W_{\theta_i}(t) \}. \end{aligned} \quad (12)$$

Using random-time-change representation (2) it can be shown²² that the reweighting process $W_{\theta_i}(t)$ is a zero-mean

martingale process and can be represented by

$$W_{\theta_i}(t) = \sum_{r=1}^M \int_0^t \frac{\frac{\partial \lambda_r}{\partial \theta_i}(X(s^-), \theta^0)}{\lambda_r(X(s^-), \theta^0)} dR_r(s) - \sum_{r=1}^M \int_0^t \frac{\partial \lambda_r}{\partial \theta_i}(X(s^-), \theta^0) ds, \quad (13)$$

where $dR_r(s)$ is simply the counting measure of reaction r which equals 1 at times s at which reaction r fires and is zero otherwise. Thus, assuming one can compute $\partial_{\theta_i} \lambda_r(x, \theta)$, then $W_{\theta_i}(t)$ has a computationally tractable form as follows.

We write \hat{X}_l for the l th state of the system through a trajectory, and Δ_l for the holding time in the l th state. Write T_l for the time of the l jump, so that $T_l = \sum_{j=0}^{l-1} \Delta_j$. We denote $N(t)$ as the total number of reactions which have fired by time t and r_l^* for the index of the reaction which takes the system from the l th state to the $l+1$ th state. Then W_{θ_i} has the explicit form

$$W_{\theta_i}(t) = \sum_{l=0}^{N(t)-1} \left[\frac{\partial}{\partial \theta_i} \log \lambda_{r_l^*}(\hat{X}_l, \theta^0) - \sum_{r=1}^M \frac{\partial \lambda_r}{\partial \theta}(\hat{X}_l, \theta^0) \Delta_l \right] - \sum_{r=1}^M \frac{\partial \lambda_r}{\partial \theta}(\hat{X}_l, \theta^0) [T - T_{N(t)}].$$

In simulation, the LR estimate is computed via ensemble averages estimated by empirical averaging $S_{f,T}(\theta_i) \approx \widehat{\mathbf{LR}}(N_S, \theta_i)$ with the empirical estimator

$$\widehat{\mathbf{LR}}(N_S, \theta_i) \triangleq \frac{1}{N_S} \sum_{n=1}^{N_S} [\widehat{f(x(T))}]_n [\widehat{W_{\theta_i}(T)}]_n, \quad (14)$$

where N_S is the number sample paths, $[\widehat{f(x(T))}]_n$ is the observable value at terminal time T for the n th sample path, and similarly $[\widehat{W_{\theta_i}(T)}]_n$ is the terminal value of $W_{\theta_i}(T)$ for the n th sample path. While the reweighting process $W_{\theta_i}(t)$ has zero mean, its variance grows with time,^{18,23} making it quite inefficient for large time horizons. The variance can be reduced by using the centered likelihood ratio estimate

$$\widehat{\mathbf{CLR}}(N_S, \theta_i) \triangleq \widehat{\mathbf{LR}}(N_S, \theta_i) - \frac{1}{N_S^2} \left\{ \sum_{n=1}^{N_S} [\widehat{f(x(T))}]_n \right\} \times \left\{ \sum_{n=1}^{N_S} [\widehat{W_{\theta_i}(T)}]_n \right\}. \quad (15)$$

Since $\mathbb{E}_{\theta} \{W_{\theta_i}(T)\} \equiv 0$, the second term does not impose any bias into estimate (12), but is coupled to the first term to reduce the observed variance.¹⁸

Suppose one is interested in the steady-state sensitivities, $S_{f,\infty}(\theta_i) \triangleq \partial_{\theta_i} \mathbb{E}_{\pi(\theta)} \{f(X)\}$. It is well known that $\mathbb{E}_{p_T(\theta)} \{f(X(T))\} = \mathbb{E}_{\pi(\theta)} \{f(X)\} + O(e^{-\kappa T})$ for some mixing rate κ , and thus for large T one can use the terminal distribution of $f(X(T))$ and $W_{\theta_i}(T)$ in (12) to obtain an estimate of the steady-state sensitivity with exponentially small bias.²³ However, the major difficulty in using likelihood ratio estimates is the large variance of the estimator $f(X(T))W_{\theta_i}(T)$, which is proportional to $\text{Var}\{f(X(T))\} \text{Var}\{W_{\theta_i}(T)\}$.^{18,22} It can be seen that $\text{Var}\{W_{\theta_i}(T)\} = O(T)$, so one must balance choosing a terminal time T large enough to ensure sufficient

decay of the bias $\mathbb{E}_{p_T(\theta)} \{f(X(T))\} - \mathbb{E}_{\pi(\theta)} \{f(X)\}$, yet as small as possible to contain the growth of the $\text{Var}\{W_{\theta_i}(T)\}$. While centering as in (15) helps to reduce the variance of the estimator, the variance is usually much larger than comparable finite difference or pathwise derivative methods.¹⁶⁻¹⁸

Instead of using the terminal distribution $f(X(T))$ as an approximation of the steady-state distribution, one could instead use the ergodic-average (time-average) $1/T \int_0^T f(X(s))ds$. The bias of the ergodic-average decays slower than the terminal distribution ($O(1/T)$ compared to $O(e^{-\kappa T})$), but has the advantage that variance decays with time as well; that is, $\text{Var}\{1/T \int_0^T f(X(s))ds\} = O(1/T)$ whereas $\text{Var}\{f(X(T))\} \rightarrow \text{Var}\{f(X(\infty))\} = \sigma^2$ (see Ref. 35, pp. 96 and 228 for more details). While such time-averaging substitutions for steady-state sensitivities are well known in other settings, the technique has not appeared in the recent literature involving steady state sensitivity estimation for stochastic reaction networks^{16,18,21,23} and so we describe it below.

The idea is to simply replace the terminal-state observable $f(X(T))$ with the ergodic average $1/T \int_0^T f(X(s))ds$ in the LR scheme (12)–(15). The philosophy is that by incurring some small amount of additional bias in the mean value, the ergodic steady-state sensitivity estimate has smaller variance than the terminal-state distribution. In fact, using Fubini's theorem and the fact that $\lim_{T \rightarrow \infty} \mathbb{E}_{p_T(\theta)} \{f(X(T))W_{\theta}(T)\} = \partial_{\theta} \mathbb{E}_{\pi(\theta)} \{f(X)\}$,³¹ it is not hard to derive

$$\begin{aligned} \mathbb{E}_{p_T(\theta)} \{1/T \int_0^T f(X(s))ds W_{\theta}(T)\} \\ = 1/T \int_0^T \mathbb{E}_{p_T(\theta)} \{f(X(s))W_{\theta}(T)\} ds \rightarrow \partial_{\theta} \mathbb{E}_{\pi(\theta)} \{f(X)\} \end{aligned}$$

as $T \rightarrow \infty$ and hence the ergodic estimator is asymptotically unbiased. Thus we have the *ergodic likelihood ratio*,

$$\widehat{\mathbf{ELR}}(N_S, \theta_i) \triangleq \frac{1}{N_S} \sum_{n=1}^{N_S} \frac{1}{T} \left[\int_0^T \widehat{f(x(s))} ds \right]_n [\widehat{W_{\theta_i}(T)}]_n. \quad (16)$$

Unfortunately, one usually sees that the variance of the ELR is much the same as the LR because the variance of $W(T)$ dominates in (16) for large T . However, we can make a further attempt to reduce the variance by centering the ELR to derive the *centered ergodic likelihood ratio CELR*,

$$\begin{aligned} \widehat{\mathbf{CELR}}(N_S, \theta_i) \triangleq & \widehat{\mathbf{ELR}}(N_S, \theta_i) \\ & - \frac{1}{N_S^2} \left\{ \sum_{n=1}^{N_S} \frac{1}{T} \left[\int_0^T \widehat{f(x(s))} ds \right]_n \right\} \\ & \times \left\{ \sum_{n=1}^{N_S} [\widehat{W_{\theta_i}(T)}]_n \right\}. \end{aligned} \quad (17)$$

In practice we have observed that the combination of centering and ergodic averaging in the CELR method results in an $O(1)$ estimator, compared to $O(T)$ for the LR and ELR estimators and $\approx O(T^{1/2})$ for the CLR.

As an example, Figures 2 and 3 compare estimators (14)–(17) by sample size for the system

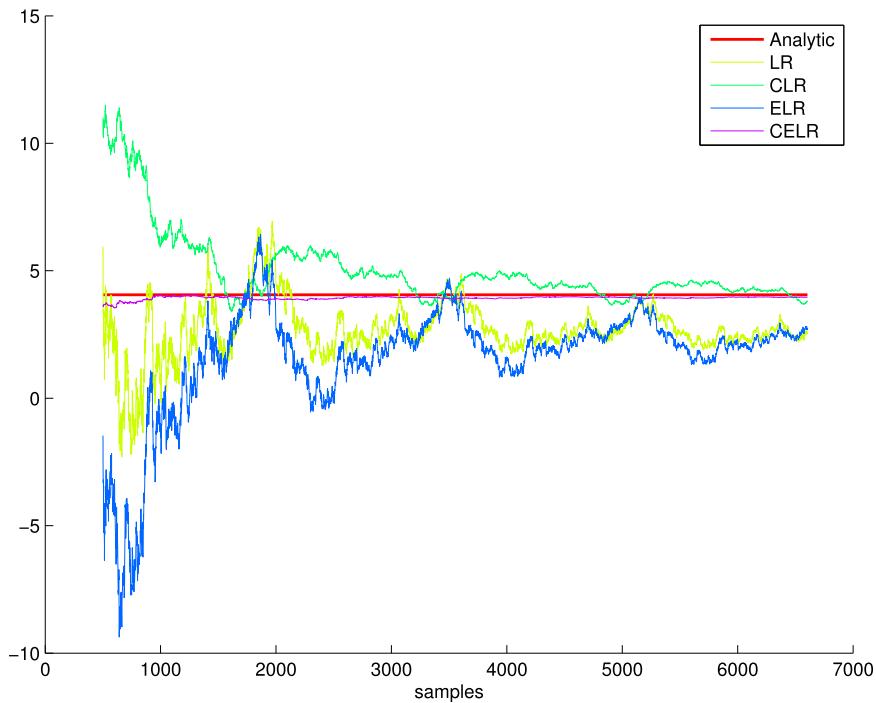
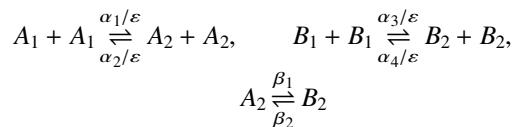


FIG. 2. Comparison of estimators (14)–(17) by sample size for a steady-state sensitivity estimation problem. The oscillations of the estimators by sample size give a depiction of the relative variance of each of the estimators. The CELR method does so well it can barely be discerned from the analytic steady state sensitivity.



at steady-state using observable $f(X) = |B_1|$ and differentiating against parameter $\theta = \beta_1$. Here we use initial point $(A_1, A_2, B_1, B_2) = (10, 15, 15, 10)$ and nominal parameter values $\boldsymbol{\alpha} = [1, 2, 3, 4]$, $\boldsymbol{\beta} = [1.5, 2.5]$, and $\varepsilon = 10^{-3}$. The exact system $X^\varepsilon(t)$ is simulated to time $T = 100$ for which the system is approximately at steady state. Figure 2 plots the exact (analytic) steady state sensitivity with the different estimators

against the number of samples used for ensemble averaging. As the sample size grows, we see that the LR and ELR estimators both perform badly compared to their centered counterparts due to their high variance. In fact, the ELR actually does worse than the LR for this sample path due to its slightly greater bias. The CLR estimator does better at large sample sizes, but nowhere near as well as the CELR estimator which can barely be discerned from the analytic value. Figure 3 shows just the CLR and CELR estimators in order to better see the comparison. The takeaway is that combining centering with ergodic-averaging gives dramatically better variance reduction than either technique applied separately.

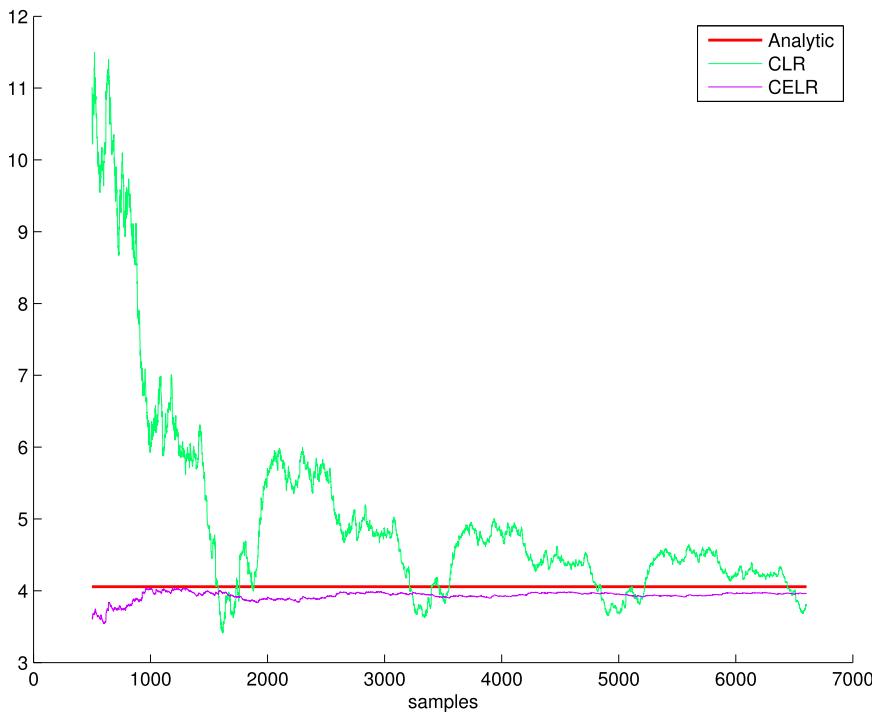


FIG. 3. Comparison of estimators (15) and (17) by sample size for a steady-state sensitivity estimation problem. The variance of the CELR method is drastically less than the CLR method with only a slight increase in the bias.

More examples which contrast the different LR estimators can be seen in the recent comparison study in Ref. 36, which also gives some insight into why $\text{Var}\{\text{ELR}\} \approx \text{Var}\{\text{LR}\} = O(T)$ whereas $\text{Var}\{\text{CELR}\} = O(1)$ as $T \rightarrow \infty$.

B. TTS likelihood ratio

In what follows, we describe how the above single-scale likelihood ratio methods can be adapted to the macro-process $\bar{X}(T)$ for use in (11). Recall that the reaction parameters can be classified as fast or slow, $\boldsymbol{\theta} = [\boldsymbol{\alpha}, \boldsymbol{\beta}]$ with $\boldsymbol{\alpha} = [\alpha_1, \dots, \alpha_{M_f}]$ and $\boldsymbol{\beta} = [\beta_1, \dots, \beta_{M_s}]$. To apply likelihood ratio methods to compute $\partial_{\theta_i} \mathbb{E}_{\boldsymbol{\theta}} \{ \bar{f}(\bar{X}(T)) \}$, we exploit that the macro-process $\bar{X}(T)$ is identified as reaction network with propensities

$$\begin{aligned} \bar{\lambda}_{\beta_r}(\bar{X}; \boldsymbol{\alpha}, \boldsymbol{\beta}) &= \mathbb{E}_{\bar{\pi}(\bar{X})(\boldsymbol{\alpha})} \{ \lambda_{\beta_r}(X; \boldsymbol{\beta}) \} \\ &= \sum_{x \in \mathcal{M}_{\bar{X}}} \lambda_{\beta_r}(x; \boldsymbol{\beta}) \bar{\pi}(\bar{X})(x; \boldsymbol{\alpha}) \end{aligned}$$

(for $\beta_r \in \boldsymbol{\beta}$), and observable

$$\bar{f}(\bar{X}; \boldsymbol{\alpha}) = \mathbb{E}_{\bar{\pi}(\bar{X})(\boldsymbol{\alpha})} \{ f(X) \} = \sum_{x \in \mathcal{M}_{\bar{X}}} f(x) \bar{\pi}(\bar{X})(x; \boldsymbol{\alpha}).$$

Thus the macro-sensitivities can be represented by

$$\begin{aligned} \frac{\partial}{\partial \theta_i} \mathbb{E}_{\bar{p}_T(\boldsymbol{\theta})} \{ \bar{f}(\bar{X}(T); \boldsymbol{\theta}) \} \\ = \mathbb{E}_{\bar{p}_T(\boldsymbol{\theta})} \left\{ \frac{\partial}{\partial \theta_i} \bar{f}(\bar{X}(T); \boldsymbol{\theta}) + \bar{f}(\bar{X}; \boldsymbol{\theta}) \bar{W}_{\theta_i}(T) \right\}, \end{aligned} \quad (18)$$

where the macro-reweighting process $\bar{W}_{\theta_i}(T)$ is given by

$$\begin{aligned} \bar{W}_{\theta_i} &= \sum_{r=1}^{M_s} \int_0^T \frac{\frac{\partial}{\partial \theta_i} \bar{\lambda}_{\beta_r}(\bar{X}(s); \boldsymbol{\theta})}{\bar{\lambda}_{\beta_r}(\bar{X}(s); \boldsymbol{\theta})} dR_{\beta_r}(s) \\ &\quad - \sum_{r=1}^{M_s} \int_0^T \frac{\partial}{\partial \theta_i} \bar{\lambda}_{\beta_r}(\bar{X}(s); \boldsymbol{\theta}) ds. \end{aligned} \quad (19)$$

Therefore, in order to apply (18) we need to be able to compute the derivatives of the averaged observable $\partial_{\theta_i} \bar{f}(\bar{X}(s); \boldsymbol{\theta})$ as well as the derivatives of the averaged propensity functions $\partial_{\theta_i} \bar{\lambda}_{\beta_r}(\bar{X}(s); \boldsymbol{\theta})$.

Suppose $\theta_i = \beta_i \in \boldsymbol{\beta}$ is a slow reaction parameter. If the original observable $f(X)$ has no direct parameter dependence, then $\partial_{\beta_i} \bar{f}(\bar{X}(s), \boldsymbol{\theta}) \equiv 0$. Furthermore, under mass-action kinetics, the averaged propensities have $\partial_{\beta_i} \bar{\lambda}_{\beta_r}(\bar{X}(s); \boldsymbol{\theta}) = \bar{b}_{\beta_r}(\bar{X}; \boldsymbol{\alpha}) \delta_{i,r}$, where $\bar{b}_{\beta_r}(\bar{X}; \boldsymbol{\alpha}) = 1/\beta_r \cdot \bar{\lambda}_{\beta_r}(\bar{X}; \boldsymbol{\alpha})$ is already computed during a TTS simulation and $\delta_{i,r} = 1$ if $i = r$ and 0 otherwise. Thus the slow sensitivities are directly computable from a standard TTS simulation.

Suppose $\theta_i = \alpha_i \in \boldsymbol{\alpha}$ is a fast reaction parameter. Then computing $\partial_{\alpha_i} \bar{f}(\bar{X}(s); \boldsymbol{\alpha})$ and $\partial_{\alpha_i} \bar{\lambda}_{\beta_r}(\bar{X}(s); \boldsymbol{\alpha}, \boldsymbol{\beta})$ is more problematic, as they only depend indirectly on $\boldsymbol{\alpha}$ through the fast-class steady-state measures $\bar{\pi}(\boldsymbol{\alpha})$. Thus explicit computation is often infeasible. However, one may estimate $\partial_{\alpha_i} \mathbb{E}_{\bar{\pi}(\bar{X})(\boldsymbol{\alpha})} \{ f(X) \}$ and $\partial_{\alpha_i} \mathbb{E}_{\bar{\pi}(\bar{X})(\boldsymbol{\alpha})} \{ \lambda_{\beta_r}(X; \boldsymbol{\beta}) \}$ through any sensitivity analysis method from a simulation with only fast reactions. For example, when running the fast-only simulation (under $\bar{Q}^{(\bar{X})}(\boldsymbol{\alpha})$) for equilibration in Algorithm 1, one can compute the corresponding likelihood ratio process $\bar{W}_{\alpha_i}^{(\bar{X})}(t)$ as

in (13) (with t large enough so that $\bar{p}_t^{(\bar{X})}(\boldsymbol{\alpha}) \approx \bar{\pi}^{(\bar{X})}(\boldsymbol{\alpha})$). Then one can estimate the derivatives in (18) and (19) by

$$\begin{aligned} \frac{\partial}{\partial \alpha_i} \bar{f}(\bar{X}; \boldsymbol{\alpha}) &\approx \mathbb{E}_{\bar{p}_t^{(\bar{X})}} \left\{ \frac{1}{t} \left[\int_0^t f(\bar{X}(s)) ds \right] \bar{W}_{\alpha_i}(t) \right\}, \\ \frac{\partial}{\partial \alpha_i} \bar{\lambda}_{\beta_r}(\bar{X}; \boldsymbol{\theta}) &\approx \mathbb{E}_{\bar{p}_t^{(\bar{X})}} \left\{ \frac{1}{t} \left[\int_0^t \lambda_{\beta_r}(\bar{X}(s)) ds \right] \bar{W}_{\alpha_i}(t) \right\}, \end{aligned} \quad (20)$$

using the proposed CELR method (17) during the micro-equilibration computation. Plugging these estimated values into (19) allows one to calculate \bar{W}_{α_i} for each macro-trajectory, which in turn allows for sensitivity estimation with respect to α_i in (18).

We note that our derivation leads to a different form of the multi-scale LR estimator than that of Ref. 21. The latter estimated the reweighting measures for the exact process $W_{\alpha_i}^{\varepsilon}(t)$ by adding together the micro-reweighting measures \bar{W}_{α_i} from within each fast class visited, the idea being that $\bar{W}_{\alpha_i}(t)$ is a zero-mean martingale which adds no new information and only adds in variance once the fast-only process has converged to steady-state. Henceforth, we refer to this approach as the “Truncated Likelihood Ratio,” as it approximates the exact reweighting coefficient $W^{\varepsilon}(t)$ via a truncated observation within each fast-class. Conversely, the TTS likelihood ratio uses the exact representation (19) for the macro-process, and then estimates the terms via (20).

C. Stiff versus rescaled sensitivities

We note that the above procedure will estimate “rescaled” sensitivities with respect to the parameter set $\boldsymbol{\theta} = [\boldsymbol{\alpha}, \boldsymbol{\beta}]$. However, depending on the context of the application, the actual objective may be to estimate the “stiff” sensitivities with respect to $\boldsymbol{\theta}^{\varepsilon} = [\boldsymbol{\alpha}/\varepsilon, \boldsymbol{\beta}] = [\boldsymbol{\alpha}^{\varepsilon}, \boldsymbol{\beta}]$. The sensitivities of the slow parameters $\boldsymbol{\beta}$ are the same, but the TTS sensitivity scheme (as do other multiscale methods^{14,21}) estimates fast sensitivities against the rescaled parameter α_i^{ε} rather than against the stiff parameter $\alpha_i^{\varepsilon} = \alpha_i/\varepsilon$. Naturally, one wonders what the relationship is between the stiff sensitivities $\partial_{\alpha_i^{\varepsilon}}$ and the rescaled sensitivities ∂_{α_i} . Here we show that, against the exact steady state measure π^{ε} , the two sensitivities differ by a factor of ε .

We first present a lemma about general stationary sensitivities of ergodic continuous-time Markov chains, whose proof we defer to [Appendix A 4](#).

Lemma 3.2. *Let $Q = Q(\boldsymbol{\theta})$ be the generator of an ergodic continuous-time Markov chain determined by a parameter vector $\boldsymbol{\theta}$. Let $\pi = \pi(\boldsymbol{\theta})$ be the corresponding stationary distribution. Then we can represent the derivatives $\partial_{\theta} \pi(\boldsymbol{\theta})$ as*

$$\frac{\partial \pi}{\partial \theta} = \pi \left(\frac{\partial Q}{\partial \theta} \right) Q^+ [\mathbf{1}\pi - I], \quad (21)$$

where Q^+ is the pseudo-inverse of the generator Q .

Lemma 3.2 has potential to be interesting in its own right as a method for computing steady-state sensitivities of reaction networks with relatively small (effective) state space \mathcal{M} . If one can construct the generator Q , then (21) gives the

vector $\frac{\partial \pi}{\partial \theta}$ and thus of the steady-state expected value

$$\begin{aligned} \frac{\partial}{\partial \theta} \mathbb{E}_{\pi(\theta)} \{f(X; \theta)\} \\ = \sum_{x \in \mathcal{M}} \frac{\partial f(x; \theta)}{\partial \theta} \pi(x; \theta) + \sum_{x \in \mathcal{M}} f(x; \theta) \frac{\partial \pi(x; \theta)}{\partial \theta}. \end{aligned} \quad (22)$$

For example, for reaction networks with mass-action propensities, the entries of Q shall always be linear in the parameters θ . Therefore, the derivatives $\frac{\partial Q}{\partial \theta}$ are easy to compute (whereas the stationary distribution $\pi(\theta)$ can be quite complex as a function of θ). Thus by computing the pseudo-inverse Q^+ from Q at the nominal value of θ (e.g., via singular value decomposition), one can analytically compute the system sensitivities without need of Monte Carlo simulation. For larger reaction networks, it may be possible to combine the finite state projection method^{37,38} with this pseudo-inverse technique to estimate the exact sensitivities by computing the analytic sensitivities of the reduced system.

Finally, equipped with Lemma 3.2, we can prove the following.

Proposition 3.3 (stiff vs. rescaled sensitivities).

$$\frac{\partial}{\partial \alpha_i^\varepsilon} \mathbb{E}_{\pi^\varepsilon} \{f(X; \theta^\varepsilon)\} = \varepsilon \left[\frac{\partial}{\partial \alpha_i} \mathbb{E}_{\pi^\varepsilon} \{f(X; \theta^\varepsilon)\} \right]. \quad (23)$$

Proof. Write the exact generator Q^ε as $Q^\varepsilon = Q(\theta^\varepsilon) = (1/\varepsilon) \tilde{Q}(\alpha) + \tilde{Q}(\beta) = \tilde{Q}(\alpha^\varepsilon) + \tilde{Q}(\beta)$. For a fast reaction parameter $\alpha_i^\varepsilon = \alpha_i/\varepsilon$ application of the chain rule gives $\partial_{\alpha_i} \lambda_k(x; \alpha^\varepsilon) = [\partial_{\alpha_i^\varepsilon} \lambda_k(x; \alpha^\varepsilon)]/\varepsilon$, from which it follows (using (3)) that $\varepsilon [\partial_{\alpha_i} \tilde{Q}(\alpha^\varepsilon)] = \partial_{\alpha_i} \tilde{Q}(\alpha) = \partial_{\alpha_i^\varepsilon} \tilde{Q}(\alpha^\varepsilon)$. Putting this relation into (21), we then have

$$\frac{\partial}{\partial \alpha_i^\varepsilon} \pi^\varepsilon = \varepsilon \left[\frac{\partial}{\partial \alpha_i} \pi^\varepsilon \right]. \quad (24)$$

Similarly we see that $\partial_{\alpha_i} f(x, \alpha^\varepsilon) = [\partial_{\alpha_i^\varepsilon} f(x, \alpha^\varepsilon)]/\varepsilon$. Finally, putting these relations into (24) we obtain (23). \square

Therefore, by multiplying the TTS sensitivity estimate (against α_i) by a factor of ε , one thus obtains the estimate against the original parameter α_i^ε . Thus one can use the TTS scheme to estimate the full gradient of $\nabla_{\theta^\varepsilon} \mathbb{E}_{\pi^\varepsilon} \{f(X; \theta^\varepsilon)\}$.

IV. BATCH-MEANS STOPPING RULE

A crucial question when implementing a TTS simulation is: How long to run the micro-equilibration for? That is, how large a value of t does one use to compute the ergodic averages

$$\begin{aligned} \bar{f}(\bar{X}; \theta) &\approx \frac{1}{t} \int_0^t f(\tilde{X}^{(\bar{X})}(s); \theta) ds, \\ \bar{\lambda}_\beta(\bar{X}; \theta) &\approx \frac{1}{t} \int_0^t \lambda_\beta(\tilde{X}^{(\bar{X})}(s); \beta) ds \end{aligned}$$

for a desired level of accuracy δ ? Taking too small a value for t risks imposing a large bias. However, the $O(1/t)$ rate of convergence for the ergodic average implies almost nothing is gained by integrating t past the relaxation time of the system. Furthermore, when computing the micro-sensitivities one uses the micro-reweighting process $\tilde{W}_{\alpha_i}(t)$ via (20), where

the variance of $\tilde{W}_{\alpha_i}(t)$ increases with the time-horizon t . Thus, we would ideally take the smallest value of t such that $\|\tilde{p}_t^{(\bar{X})} - \tilde{\pi}^{(\bar{X})}\| \leq O(\delta)$. However, different fast-classes \bar{X} can have vastly different sizes. This can result in significantly different relaxation times for each class. It is then ideal to have an *adaptive stopping rule* which terminates the micro (fast-only) simulations when the ergodic averages have converged to the steady state mean.

Current implementations of an “averaged” or “multi-scale” SSA use a constant relaxation time t_f for the micro-averaging step^{11,21} whose choice is motivated by some *a priori* insight into the system. In Refs. 10 and 39 the authors also use a fixed time t , but then exploit algebraic relations of the steady-state means to try to obtain better approximations. In Ref. 9, a stopping rule is developed which determines that equilibrium is reached when the averaged values of the propensities of the forward and backward reactions are approximately equal for each reaction pair. However, experience has shown this “partial-equilibrium” stopping rule can stop prematurely (in the transient regime) with significant probability for systems with relatively few reaction-pairs. Thus, we seek to obtain a robust, adaptive stopping rule for terminating the micro-equilibration simulation.

A. Batch-means for steady-state estimation

The problem at hand is really one about Markov chain mixing-times and the integrated autocorrelation time τ_{int} . Analytically, the mixing and integrated autocorrelation times are related to the spectral gap of the underlying generator.⁴⁰⁻⁴² Unfortunately, for large systems direct computation is usually infeasible. Some common approaches involve estimation autocorrelation function $A(t)$ of the process and then exploiting the relation $\tau_{int} = 2 \int_0^\infty A(t) dt$ to derive estimates of τ_{int} from the estimates of $A(t)$.^{41,43,44} However, if the goal is to terminate the simulation when the ergodic average has converged appropriately, then these methods are indirect and can be computationally intensive. Another common approach is to exploit the regenerative structure of Markov chains⁴⁵ to obtain independent and identically distributed samples of the process and obtain confidence bounds on the ergodic average. However, these methods can be inefficient for complex systems where the return time to the initial state can be quite large.

We instead turn to the method of batch means^{35,46} for determining confidence bounds (and thus a measure of convergence) for the steady-state estimation problem inside each fast-class. The use of batch means is applicable to a wide range of problems (any which satisfy a central limit theorem), and its implementation is very straightforward.

For a general Markov chain $X(s)$ with an observable function f , write $Y(s) = f(X(s))$ and $\bar{Y}(t) = 1/t \int_0^t Y(s) ds$. We denote $f_\pi = \mathbb{E}_\pi \{f(X)\}$ for the steady-state value we wish to estimate. Then under some general conditions⁴² $Y(s)$ satisfies a functional central limit theorem,

$$\frac{t}{\sqrt{\varepsilon}} \left\{ \bar{Y}(t/\varepsilon) - f_\pi \right\}_{t \geq 0} \xrightarrow{\mathcal{D}} \{\sigma B(t)\}_{t \geq 0}$$

in the sense of weak convergence as $\varepsilon \rightarrow 0$.

Suppose that $t \geq N_b \tau_{\text{relax}}$, where τ_{relax} is the relaxation time of the system and N_b is a number of “batches” (bins) to partition the trajectory into. Then the batch means

$$\bar{Y}_k(t) \triangleq \frac{1}{t/N_b} \int_{(k-1)t/N_b}^{kt/N_b} Y(s) ds$$

are approximately (as $t \rightarrow \infty$) independent and identically distributed samples of $\mathcal{N}(f_\pi, \sigma^2 N_b/t)$. Thus

$$\sqrt{N_b} \frac{\bar{Y}(t) - f_\pi}{s_{N_b}(t)} \xrightarrow{\mathcal{D}} T_{N_b-1}$$

as $t \rightarrow \infty$, where T_{N_b-1} is the Student’s t -distribution and $s_{N_b}^2(t)$ is the sample variance among batches,

$$s_{N_b}^2(t) \triangleq \frac{1}{N_b-1} \sum_{k=1}^{N_b} [\bar{Y}_k(t) - \bar{Y}(t)]^2.$$

Thus, for t sufficiently large, a $(1 - \delta_{CI})100\%$ confidence interval for the value of f_π is given by $\bar{Y}(t) \pm \text{MOE}(t, N_b, \delta_{CI})$, where

$$\text{MOE}(t, N_b, \delta_{CI}) \triangleq (t_{\text{quantile}}) \frac{s_{N_b}(t)}{\sqrt{N_b}} \quad (25)$$

and t_{quantile} is the $(1 - \delta_{CI}/2)$ th quantile of the Student’s t -distribution with $N_b - 1$ degrees of freedom.

The usual perspective for applying batch means is that one has a fixed set of data $\{Y(s) : s \in [0, t]\}$ to partition, and then must choose the number of batches N_b appropriately so that each batch length t/N_b is long enough so that the batch mean errors $[\bar{Y}_k(t) - \bar{Y}(t)]$ are approximately independent, identically distributed, and Gaussian. One then often chooses N_b to be relatively small (say, 5–30)^{41,47} to ensure the independent and Gaussian assumptions hold. When viewing the asymptotic structure as the amount of data t grows, then one can ensure that the asymptotic central limit theorem holds if the number of batches grows as $N_b(t) \simeq \sqrt{t}$. In Ref. 46, the authors consider strategies which let the number of batches grow if the correlation between batches is near 0, and otherwise hold $N_b(t)$ fixed until the batch correlation decays to 0.

Since our goal is to simulate only enough (micro-scale) data so as to determine the steady-state values f_π , we instead take the perspective that one has a fixed number of batches N_b desired, and that one should generate data $\{Y(s) : s \in [0, t]\}$ until each of the batch means $\bar{Y}_k(t)$ are (approximately) independent and identically distributed about f_π . For a fixed level of precision δ_{precise} , confidence level δ_{CI} , and the number of batches (independent samples) N_b , the batch-means stopping rule terminates the simulation when $\text{MOE}(t) = \text{MOE}(t, N_b, \delta_{CI}) \leq \delta_{\text{precise}}$, where $\text{MOE}(t)$ is defined by (25). Figure 4 gives a depiction of how the batch-means stopping rule is implemented.

In addition to giving an on-line estimate of the relaxation time of the system, the batch-means stopping rule gives $N_b - 1$ (nearly) independent samples of trajectories with initial distribution approximately equal to the stationary distribution. Furthermore, one can compute the reweighting coefficients $W_{k,\theta}(t)$ in each batch to give N_b (nearly) independent samples of the steady-state reweighting coefficients (in a manner similar to the “Time-Averaged Correlation Function” method of Ref. 23).

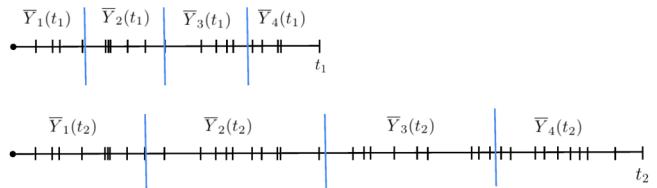


FIG. 4. A sketch of the batch-means stopping rule. The process is simulated for a fixed number of jumps ($N_J = 20$) to a terminal time t_1 , and then the trajectory is partitioned into $N_b = 4$ batches to compute the variance between the batch means $\bar{Y}_k(t_1)$. If the confidence bounds are precise enough ($\text{MOE}(t_1) \leq \delta_{\text{precise}}$), then the simulation is terminated and each batch gives an iid sample of $\mathcal{N}(f_\pi, \sigma^2)$. Otherwise, another $N_J = 20$ jumps are simulated and the process is repeated.

B. Batch-means stopping implementation

Suppose we have a general reaction network with M_r reactions and M_θ reaction parameters. Here we allow the possibility that $M_r \neq M_\theta$ for general propensity functions $\lambda_r(x; \theta)$ (e.g., Michaelis-Menten kinetics), whose parameter derivatives $\partial_{\theta_i} \lambda_r(x; \theta)$ we can compute explicitly for all $i = 1, \dots, M_\theta$ and all $r = 1, \dots, M_r$. Denote by ζ_r the stoichiometric vector for the r th reaction. Our goal is to estimate the gradient $\nabla_{\theta} \mathbb{E}_{\pi(\theta)} \{f(X)\}$ for some observable function f .

We introduce the following notation for the batch-means stopping rule. $\hat{X}(n)$ is the n th state of the reaction network, $T(n)$ is the time of the n th jump, and $\hat{f}(n) = f(\hat{X}(n))$ for the value of the observable at the n th state. $F(n) = \int_0^{T(n)} f(X(s)) ds$ is the time-integrated value of f up to the n th jump, $r^*(n)$ is the reaction which fires at jump n (taking the system from $\hat{X}(n)$ to $\hat{X}(n+1)$), \hat{e}_i is the vector in $\mathbb{R}^{1 \times M_\theta}$ with 1 in the i th component and zeros elsewhere. $R(n, \theta) \in \mathbb{R}^{1 \times M_\theta}$ and $B(n, \theta) \in \mathbb{R}^{1 \times M_\theta}$ are the first and second terms of (13) with respect to each of the parameters $\theta_i \in \theta$ ($i = 1, \dots, M_\theta$). N_b is the number of batches (approximately independent samples) to be used, δ_{CI} is the desired confidence level (for a $(1 - \delta_{CI})100\%$ confidence interval), and δ_{precise} is the maximum allowed radius of the confidence interval at the stopping time. N_J is the number of jumps to simulate before retesting the batches for convergence. Then one can write the batch-means stopping rule as follows.

Algorithm 2 (batch-means stopping rule with sensitivity estimation).

(1) Initialize

- $\hat{X}(0) = x_0$, $\hat{f}(0) = f(\hat{X}(0))$, $T(0) = 0$, $F(0) = 0$, $R(0) = [0, \dots, 0] \in \mathbb{R}^{1 \times M_\theta}$, $B(0) = [0, \dots, 0] \in \mathbb{R}^{1 \times M_\theta}$, $\text{tests} = 0$ (number of times the data has been tested for convergence). Calculate $t_{\text{quantile}} = (1 - \delta_{CI}/2)$ quantile of a Student’s t -distribution with $N_b - 1$ degrees of freedom.

(2) Generate and Record Data

Simulate N_J jumps and record values immediately after each jump. For $n = N_J \cdot \text{tests}, \dots, N_J \cdot (\text{tests} + 1) - 1$,

- compute $\lambda_r(\hat{X}(n); \theta)$, $\frac{\partial}{\partial \theta_i} \lambda_r(\hat{X}(n), \theta)$ for all $r = 1, \dots, M_r$ and all $\theta_i = 1, \dots, M_\theta$. Set $\lambda_0(\hat{X}(n); \theta) = \sum_{r=1}^{M_r} \lambda_r(\hat{X}(n), \theta)$, and $\frac{\partial}{\partial \theta_i} \lambda_0(\hat{X}(n), \theta) = \sum_{r=1}^{M_r} \frac{\partial}{\partial \theta_i} \lambda_r(\hat{X}(n), \theta)$ for all $\theta_i \in \theta$.

- Sample
 $\Delta t(n) \sim \text{Exp}(\lambda_0(\hat{X}(n), \theta))$,
and
 $r^*(n) \sim \frac{1}{\lambda_0(\hat{X}(n), \theta)} \times [\lambda_1(\hat{X}(n); \theta), \dots, \lambda_{M_r}(\hat{X}(n); \theta)]$.
- Update
 $T(n+1) = T(n) + \Delta t(n)$,
 $\hat{X}(n+1) = \hat{X}(n) + \zeta_{r^*(n)}$,
 $\hat{f}(n+1) = f(\hat{X}(n+1))$,
 $F(n+1) = F(n) + \hat{f}(n) \cdot \Delta t(n)$,
 $R(n+1, \theta) = R(n, \theta) + \sum_{i=1}^{M_\theta} \hat{e}_i$
 $\times \left[\frac{\partial}{\partial \theta_i} \lambda_{r^*(n)}(\hat{X}(n), \theta) \right] / \lambda_{r^*(n)}(\hat{X}(n), \theta)$,
 $\hat{b}(n, \theta) = \sum_{i=1}^{M_\theta} \hat{e}_i \cdot \left[\sum_{r=1}^{M_r} \frac{\partial}{\partial \theta_i} \lambda_r(\hat{X}(n), \theta) \right]$,
 $B(n+1, \theta) = B(n, \theta) + \hat{b}(n, \theta) \cdot \Delta t(n)$,
 $W(n+1, \theta) = R(n+1, \theta) - B(n+1, \theta)$.

(3) Test Batches for Convergence

- $N_{end} = N_J \cdot (tests + 1)$ = index of last available data point.
 $\bar{Y} = F(N_{end}) / T(N_{end})$ = total time-averaged value
 $t_{batch} = T(N_{end}) / N_b$ = time length each of batch.
Initialize $F^B = 0$ (total integral through end of previous batch).
- For $k = 1, \dots, N_b$
 - $ind_B(k) = \max \{n : T(n) \leq k \cdot t_{batch}\}$ = index of the last jump in batch k .
 - $F_A^B(k) = F(ind_B(k)) + \hat{f}(ind_B(k))$
 $\times [k \cdot t_{batch} - T(ind_B(k))] - F^B$ = total integrated value of $f(X(s))$ inside batch k .
 - $\bar{Y}_k = F_A^B(k) / t_{batch}$ = k th batch-mean.
 - $F^B \leftarrow F^B + F_A^B(k)$ (update integral to end of previous batch).
- $s_{N_b}^2 = \frac{1}{N_b-1} \sum_{k=1}^{N_b} [\bar{Y}_k - \bar{Y}]^2$ = variance between batches.
- $MOE = t_{quantile} * \sqrt{s_{N_b}^2 / N_b}$ = margin of error for confidence interval.
- If $MOE \leq \delta_{precise}$, then go to (4). Else, $tests \leftarrow tests + 1$ and go back to (2).

(4) Compute LR Weights in each batch.

Initialize $W^B(\theta) = [0, \dots, 0] \in \mathbb{R}^{1 \times M_\theta}$. For $k = 1, \dots, N_b$,

- $W_A^B(k, \theta) = W(ind_B(k), \theta)$
 $- b(ind_B(k), \theta) \cdot [k \cdot t_{batch} - T(ind_B(k))]$
 $- W^B(\theta)$.
- $W^B(\theta) \leftarrow W^B(\theta) + W_A^B(k, \theta)$.

(5) Compute Sensitivity Estimates for $\nabla_\theta \mathbb{E}_{\pi(\theta)} \{f(X)\}$:

- Likelihood Ratio

$$\text{LR} = \frac{1}{N_b} \sum_{k=1}^{N_b} \hat{f}(ind_B(k)) W_A^B(k, \theta).$$

- Centered Likelihood Ratio

$$\text{CLR} = \text{LR} - \left[\frac{1}{N_b} \sum_{k=1}^{N_b} \hat{f}(ind_B(k)) \right] \cdot \left[\frac{1}{N_b} \sum_{k=1}^{N_b} W_A^B(k, \theta) \right].$$

- Ergodic Likelihood Ratio

$$\text{ELR} = \frac{1}{N_b} \sum_{k=1}^{N_b} \bar{Y}_k W_A^B(k, \theta).$$

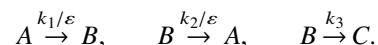
- Centered Ergodic Likelihood Ratio

$$\text{CELR} = \text{ELR} - \left[\frac{1}{N_b} \sum_{k=1}^{N_b} \bar{Y}_k \right] \cdot \left[\frac{1}{N_b} \sum_{k=1}^{N_b} W_A^B(k, \theta) \right].$$

V. SIMULATION RESULTS

Here we present numerical results to display the performance of the proposed algorithms. In what follows, we compare the output of an exact simulation at the single time scale (STS) to the accelerated two time scale (TTS) approximation. From (11), we expect the differences in observable averages and their derivatives to be $O(\varepsilon)$. Because differences are small, we use a simple test system for which many samples can be run to obtain accurate statistics.

Consider a reaction network with species A, B, and C and isomerization reactions given by



For small values of ε , the system becomes stiff as the isomerization between A and B reaches equilibrium much faster than B is converted to C. A TTS approximation assumes that $A \rightarrow B$ and $B \rightarrow A$ are fast and equilibrated.

We first compare the output of the accelerated TTS simulation against the exact STS simulation for varying levels of stiffness ε . For our simulations, initial conditions of $(A_0, B_0, C_0) = (100, 0, 0)$ and the parameters $(k_1, k_2, k_3) = (1, 1.5, 2)$ are chosen. 10 000 replicate (independent) trajectories are run for various values of ε . Statistics are taken at a termination time of $t = 0.5$ s. Species averages are calculated as arithmetic averages over the independent trajectories while sensitivities are computed with the CLR method shown in (15). The error due to statistical averaging is estimated using t -test statistics for averages and a bootstrapping method for sensitivities. Sensitivities with respect to the “slow” parameter k_3 are displayed for each species. As discussed in previous works,²¹ sensitivities with respect to parameters related to fast reactions encounter significant noise, and thus we omit them in order to clearly observe the difference between STS and TTS. Figure 5 shows the disparity between the STS and TTS systems for various values of ε . Errors are normalized by the TTS value such that $\text{Error} = \frac{\text{STS} - \text{TTS}}{|\text{TTS}|}$. Indeed, one observes the difference is proportional to ε , as expected from Corollary 2.4 and (11).

Next, the CLR and CELR methods from Section III A are tested in performing sensitivity analysis in a TTS system. The reaction network described in Section II B is simulated using Algorithm 1. To assess convergence of the microscale distribution, the batch-means stopping criterion described in Section IV is used with a tolerance of $\delta = 0.05$. 1000 replicate trajectories are run to a time horizon of $t_f = 100$ s. The initial conditions used are $(A_0, B_0, *_0) = (30, 60, 10)$ and the parameters used are $(\alpha_1^\varepsilon, \alpha_2^\varepsilon, \beta_1, \beta_2, \beta_3) = (1/\varepsilon, 1.5/\varepsilon, 2, 1, 0.4)$.

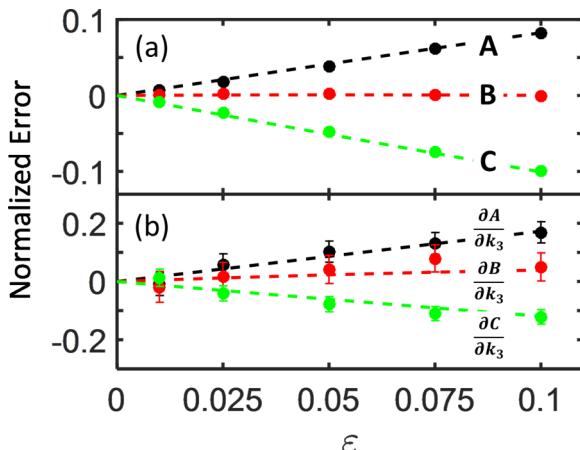


FIG. 5. Normalized error of the two-time-scale (TTS) (accelerated, approximate) simulation from the exact single-time-scale (STS) simulation. The plot (a) shows errors in species averages while the plot (b) shows errors in sensitivities of each species with respect to k_3 . Points indicate simulation statistics while dashed lines confirm the linear trend. The error bars are 95% confidence intervals of statistical noise.

Species populations, time-averaged species populations, and trajectory derivatives are recorded for each run over time. Using these recorded statistics, the sensitivities for all 15 (3 species and 5 parameters) species/parameter combinations are computed at each time point. Figure 6 shows the time evolution of the normalized errors in sensitivity estimates of the species B over time. Estimated values from simulation are referenced to the analytical answer as computed from a differential algebraic equation (see [Appendix B](#)) and normalized by that amount so that $\text{Error} = \frac{\text{estimated}-\text{analytical}}{|\text{analytical}|}$.

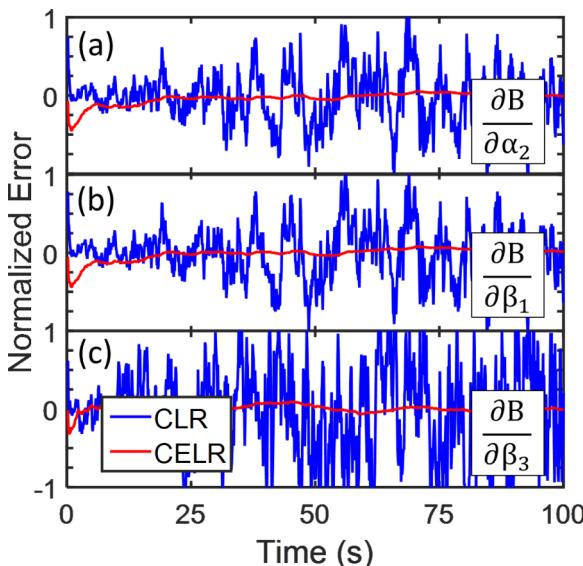


FIG. 6. Time evolution of normalized errors obtained for sensitivities computed by CLR and CELR estimators. Simulation estimates are referenced and normalized by the analytical solution at each time point to compute the error. The graphs (a)–(c) show estimates of different sensitivity indices for the species B. One observes that the variance of the CLR estimates increases with time, making for inefficient estimation. In contrast, the CELR estimates (the solid red line) converge quickly to the steady state sensitivities with variance which is roughly constant with time, allowing for efficient steady-state sensitivity estimation.

TABLE I. Comparison of the error and statistical noise for the CLR and CELR estimators at a short time (during the transient regime) and at a long time (corresponding to steady state). Values in the table refer to the sensitivity of the species B with respect to the parameter given by the row label. Values are reported as a percent of the analytically obtained sensitivity value.

	Percent error			
	$t = 1.3 \text{ s}$		$t = 100 \text{ s}$	
	CLR	CELR	CLR	CELR
α_1	-1.4	-40.0	-83.5	0.0
α_2	-3.0	-40.3	-63.0	-1.1
β_1	0.2	-39.1	-64.1	1.2
β_2	-1.1	-16.7	-22.6	-2.8
β_3	-1.4	-23.3	20.6	-0.1

	Half-length of 95% confidence interval			
	CLR	CELR	CLR	CELR
α_1	12	8	99	11
α_2	12	7	94	11
β_1	12	7	91	11
β_2	13	10	108	12
β_3	21	14	171	18

As expected, CLR estimates are noisy, with variance that grows linearly with time. At short times, the variance is small enough to obtain reasonable estimates. As time increases, the noise becomes significant with respect to the actual values (the magnitude of the normalized error becomes comparable to 1). In contrast, the CELR fails at short times with a noticeable bias. However, the bias, which exists due to a relaxation period, decays as $O(1/t)$ when time increases and the system approaches its steady-state. The variance of the CELR estimates remains constant because the variance of trajectory derivatives increases linearly in time while the variance of ergodic species averages is proportional to $1/t$. At long times (where the CLR is too noisy for efficient estimation), the ergodic likelihood ratio obtains accurate estimates with very low variance. Therefore, it is advisable to use the CLR method for short times (in the transient regime) and the CELR method for long times to obtain steady state values.

Table I shows the error and statistical noise of the CLR and CELR estimations of sensitivities of the species B at short ($t = 1.3 \text{ s}$) and long ($t = 100 \text{ s}$) times. Statistical noise is obtained from bootstrapping the samples used to compute the sensitivity estimates. At $t = 1.3$, the CLR method has low error (theoretically, there is no bias) as well as low variance. The CELR estimator has a similarly low variance, but high error (due to the $O(1/t)$ bias). At $t = 100 \text{ s}$, the CLR estimates have much higher variance which induces large empirical error. In contrast, the bias of the CELR estimate decreases in time while the variance remains low, providing very small empirical errors at large times.

VI. CONCLUSIONS

This work develops a TTS framework for multiscale reaction networks. By decomposing the system into “fast-classes,”

one can approximate the behavior of the multiscale system by a lower-dimensional, single-scale, “macro-averaged” reaction network. By applying a singular perturbation expansion of the underlying probability measures, we have established rigorous bounds on the bias induced by the approximate macro-averaged model. We then proposed a TTS algorithm for simulating the macro reaction network, using an adaptive batch-means stopping rule for determining when the micro-scale dynamics have sufficiently equilibrated.

In addition, we have shown that the sensitivities of the macro-averaged system provide accurate approximations for the multiscale system. Since the macro-averaged system is single-scale, it is possible to incorporate most existing sensitivity estimation methods to the TTS algorithm to obtain estimates of the system sensitivities. We demonstrated the efficiency of a CELR estimator for steady-state sensitivity analysis, and detailed how it can be adapted to the two-time-scale algorithm. Simulations were then used to confirm the analytic error bounds and demonstrate the efficiency of the TTS ergodic likelihood ratio estimator.

ACKNOWLEDGMENTS

The authors are grateful to the referees for their very helpful comments. This work is supported by the U.S. Department of Energy Office of Science, Office of Advanced Scientific Computing Research, and Applied Mathematics program under Award No. DE-SC0010549.

APPENDIX A: PROOFS OF RESULTS

In this section we outline the proofs on the error bounds of the averaged reaction network and convergence of the sensitivities. The error bounds in this work are largely direct applications of the results in Ref. 20 to the two-time-scale reaction networks formulated here. We present an overview of the proofs for insight and completeness. Similarly, the sensitivity convergence result comes from Ref. 14; we shall only show how to fit their result to the two-time-scale framework.

1. Proof of Theorem 2.3

Using the formulation of the exact generator $Q^\varepsilon(\theta) = (1/\varepsilon)\tilde{Q}(\alpha) + \tilde{Q}(\beta)$, the error bound on the induced probability measures of the exact and averaged systems is a direct application of Theorem 4.29 of Ref. 20. We outline the main steps below.

Write $p^\varepsilon(t) \in \mathbb{R}^{1 \times m}$ for the probability measure of the exact system at time t . From the Kolmogorov forward equation (a.k.a. chemical master equation), we have

$$\frac{dp^\varepsilon(t)}{dt} = p^\varepsilon(t) \left[\frac{1}{\varepsilon} \tilde{Q}(\alpha) + \tilde{Q}(\beta) \right]. \quad (\text{A1})$$

Define the differential operator \mathcal{L}^ε on functions with values in $\mathbb{R}^{1 \times m}$ by $\mathcal{L}^\varepsilon f = \varepsilon \frac{df}{dt} - f(\tilde{Q} + \varepsilon \tilde{Q})$. Then $\mathcal{L}^\varepsilon f = 0$ if and only if f solves CME (A1). The form of differential equation (A1) suggests the plausibility of a singular perturbation expansion

of $p^\varepsilon(t)$ by

$$p^\varepsilon(t) = \sum_{i=0}^{\infty} \varepsilon^i \phi_i(t) + \sum_{i=0}^{\infty} \varepsilon^i \psi_i \left(\frac{t}{\varepsilon} \right). \quad (\text{A2})$$

Assuming for the moment that such a representation holds, we proceed to derive the form of the “regular” terms $\phi_i(t)$ and the “boundary layer” terms $\psi_i(t)$. Applying \mathcal{L}^ε to (A2) and equating terms of ε leads to the recursive equations

$$\begin{aligned} \varepsilon^0 : \quad & \phi_0(t)\tilde{Q} = 0 \\ \varepsilon^1 : \quad & \phi_1(t)\tilde{Q} = \frac{d\phi_0(t)}{dt} - \phi_0(t)\tilde{Q} \\ & \vdots \\ \varepsilon^i : \quad & \phi_i(t)\tilde{Q} = \frac{d\phi_{i-1}(t)}{dt} - \phi_{i-1}(t)\tilde{Q} \end{aligned} \quad (\text{A3})$$

and similarly, using the “stretched-time” variable $\tau = t/\varepsilon$ one has equations for $\psi(\tau)$,

$$\begin{aligned} \varepsilon^0 : \quad & \frac{d\psi_0(\tau)}{d\tau} = \psi_0(\tau)\tilde{Q} \\ \varepsilon^1 : \quad & \frac{d\psi_1(\tau)}{d\tau} = \psi_1(\tau)\tilde{Q} - \psi_0(\tau)\tilde{Q} \\ & \vdots \\ \varepsilon^i : \quad & \frac{d\psi_i(\tau)}{d\tau} = \psi_i(\tau)\tilde{Q} - \psi_{i-1}(\tau)\tilde{Q}. \end{aligned} \quad (\text{A4})$$

At $t = 0$,

$$\sum_{i=0}^{\infty} \varepsilon^i (\phi_i(0) + \psi_i(0)) = p^\varepsilon(0), \quad (\text{A5})$$

so $\phi_0(0) + \psi_0(0) = p^\varepsilon(0)$ and $\phi_i(0) + \psi_i(0) = 0$ for all $i \geq 1$. Since $p^\varepsilon(t)$ is a probability measure with $p^\varepsilon(t) \cdot \mathbf{1} = 1$, by sending $\varepsilon \rightarrow 0$ in (A2) it follows that

$$\phi_0(t) \cdot \mathbf{1} = 1 \quad \text{and} \quad \phi_i(t) \cdot \mathbf{1} = 0 \quad (\text{A6})$$

for all $t \in [0, T]$ and all $i \geq 1$.

Turning to the leading regular term $\phi_0(t)$, we note that $\phi_0(t) \cdot \tilde{Q} = 0$ is not uniquely solvable because $\tilde{Q} = \text{diag}[\tilde{Q}^{(1)}, \dots, \tilde{Q}^{(N_C)}]$ has rank $m - N_C$. However, writing $\phi_0^{(k)}(t)$ for the sub-vector of $\phi_0(t)$ corresponding to fast-class \mathcal{M}_k , then we must have $\phi_0^{(k)}(t) \cdot \tilde{Q}^{(k)} = 0$ for all $k = 1, \dots, N_C$. Since each $\tilde{Q}^{(k)}$ is an irreducible generator, we then have $\phi_0^{(k)}(t) = \gamma^{(k)}(t)\tilde{\pi}^{(k)}$ for some scalar multiplier $\gamma^{(k)}(t)$. It can be seen (Proposition 4.24²⁰) that $\gamma^{(k)}(t) = \bar{p}_k(t)$, where $\bar{p}(t)$ is the probability measure among fast-classes $\mathcal{M}_1, \dots, \mathcal{M}_{N_C}$ induced by generator \tilde{Q} (as in (5)) and initial distribution $p^\varepsilon(0) \cdot \mathbf{1} = \mathbb{P}\{X^\varepsilon(0) \in \mathcal{M}_k\}$. This in turn determines a unique solution for $\phi_0^{(k)}(t)$ and therefore $\phi_0(t)$. It follows that $\phi_0(t)$ is exactly the measure p_t^0 in Theorem 2.3 induced by the TTS simulation procedure.

With $\phi_0(t)$ determined, (A5) then gives the initial condition $\psi_0(0) = p^\varepsilon(0) - \phi_0(0) = p^\varepsilon(0)[I_m - \tilde{\mathbf{1}}\tilde{\pi}]$, from which one can solve (A4) to obtain $\psi_0(\tau) = \psi_0(0) \cdot \exp\{\tilde{Q}\tau\}$. It can be shown (Proposition 4.25²⁰) that

$$\|\psi_0(\tau)\| \leq C \exp\{-\tilde{\kappa}\tau\}, \quad (\text{A7})$$

where C depends on the Jordan-form of \tilde{Q} . Higher order terms can also be solved for recursively (Proposition 4.26²⁰), and it

can be shown (Proposition 4.28²⁰)

$$\sup_{0 \leq t \leq T} \left\| \sum_{i=0}^n \varepsilon^i \phi_i(t) + \sum_{i=0}^n \varepsilon^i \psi_i(t/\varepsilon) - p^\varepsilon(t) \right\| = O(\varepsilon^{n+1}). \quad (\text{A8})$$

In particular, using the 0th-order expansion we have

$$\begin{aligned} \|p_T^0 - p_T^\varepsilon\| &= \|\phi_0(T) - p^\varepsilon(T)\| \\ &\leq \|\phi_0(T) + \psi_0(T/\varepsilon) - p^\varepsilon(T)\| + \|\psi_0(T/\varepsilon)\| \\ &\leq \|O(\varepsilon) + O(\exp\{-\tilde{\kappa}T/\varepsilon\})\|. \end{aligned} \quad (\text{A9})$$

□

2. Proof of Corollary 2.4

With singular perturbation bound (A8), Corollary 2.4 follows immediately. Since the exact process $X^\varepsilon(t)$ is ergodic, there exists a time horizon T^ε such that $\|p^\varepsilon(t) - \pi^\varepsilon\| \leq \varepsilon$ for $t \geq T^\varepsilon$. Similarly, by (A7) there exists \bar{T} such that $\|\psi_0(t)\| \leq \varepsilon$ for $t \geq \bar{T}$, and $\|\bar{p}(t) - \bar{\pi}\| \leq \varepsilon$ for $t \geq \bar{T}$, implying that $\|\phi_0(t) - \bar{\pi}\bar{\pi}\| = \|\bar{p}(t)\bar{\pi} - \bar{\pi}\bar{\pi}\| \leq \varepsilon$. Then taking $T \geq \max\{T^\varepsilon, \bar{T}, \bar{T}\}$ and applying (A8), we have

$$\begin{aligned} \|p^\varepsilon(t) - \bar{\pi}\bar{\pi}\| &\leq \|p^\varepsilon(t) - \phi(t) - \psi(t/\varepsilon)\| \\ &\quad + \|\phi(t) - \bar{\pi}\bar{\pi}\| + \|\psi(t/\varepsilon)\| \\ &\leq O(\varepsilon + \exp\{-\tilde{\kappa}t/\varepsilon\}) \\ \|\pi^\varepsilon - \bar{\pi}\bar{\pi}\| &\leq \|\pi^\varepsilon - p^\varepsilon(T)\| + \|p^\varepsilon(T) - \bar{\pi}\bar{\pi}\| \leq O(\varepsilon) \end{aligned}$$

and the corollary follows. □

3. Proof of Proposition 3.1

Proposition 3.1 is simply the application of Theorem 3.2 of Ref. 14 to the TTS framework. The method and framework for separating time scales in Ref. 14 is slightly different from the TTS framework used here, but it can be seen that the two are equivalent. Here, we briefly review the multiscale framework of Ref. 14 and show how one can translate between their “remainder spaces” and the TTS “fast-classes.”

a. Scaling rates and remainder spaces

As in Refs. 13 and 18, Ref. 14 considers reaction rates of the form $a_k^N(x, \theta) = N^{\rho_k} \lambda_k(x, \theta)$ which scale with the “system size” or “normalization parameter” $N \gg 0$, where ρ_k is the scaling rate for the k th reaction channel. For a given normalizing parameter N , the corresponding system is denoted by $X^N(t)$. One analyzes the system against a reference time scale γ by $X_\gamma^N(t) = X^N(tN^\gamma)$. For a given normalizing parameter $N_0 \gg 0$, the corresponding system is denoted by $X^{N_0}(t)$. One analyzes the system against a reference time scale γ by $X_\gamma^N(t) = X^N(tN^\gamma)$.

The scaling rates ρ_k determine the time scales at which the reaction channels fire. For a system with a single level of stiffness, there are only two scaling rates, $\rho_{fast} > \rho_{slow}$, which partition the reaction channels as either fast or slow. Write $\Gamma_1 = \{k : \rho_k = \rho_{fast}\}$ for the fast reaction index set, and similarly Γ_2 for the slow reaction index set.

Take $\mathbb{S}_2 = \{v \in \mathbb{R}_+^d : \langle v, \zeta_k \rangle = 0 \text{ for all } k \in \Gamma_1\}$ so that $\langle X_\gamma^N(t), v \rangle$ is unchanged by fast reactions. Then take

$\mathbb{L}_2 = \text{span}(\mathbb{S}_2)$ and Π_2 to be the projection map from \mathbb{R}^d to \mathbb{L}_2 , so that $\Pi_2 \zeta_k = \mathbf{0}$ for all $k \in \Gamma_1$.

Let $\mathbb{L}_1 = \text{span}\{(I - \Pi_2)x : x \in \mathcal{M}\}$, and for any $v \in \Pi_2 \mathcal{M}$ let $\mathbb{H}_v = \{y \in \mathbb{L}_1 : y = (I - \Pi_2)x, \Pi_2 x = v, x \in \mathcal{M}\}$, the set of remainders of elements in \mathcal{M} which get projected to v . Then we can define an operator \mathbb{C}^v by

$$\mathbb{C}^v f(z) = \sum_{k \in \Gamma_1} \lambda_k(v + z, \theta) [f(z + \zeta_k) - f(z)]$$

which is a generator of a Markov chain with state space \mathbb{H}_v (note that $y \in \mathbb{H}_v \implies y + \zeta_k \in \mathbb{H}_v$ for all $k \in \Gamma_1$).

Assuming \mathbb{H}_v under \mathbb{C}^v is ergodic, there is a stationary distribution π^v . Then for each slow reaction $k \in \Gamma_2$ one can define the “averaged” propensities $\hat{\lambda}_k(v, \theta) = \int_{\mathbb{H}_v} \lambda_k(v + z, \theta) \pi^v(dz)$ for all $v \in \Pi_2 \mathcal{M}$. Using the random time change representation, define the Markov chain on $\Pi_2 \mathcal{M}$ by

$$\hat{X}_\theta(t) = \Pi_2 x_0 + \sum_{k \in \Gamma_2} Y_k \left(\int_0^t \hat{\lambda}_k(\hat{X}(s), \theta) ds \right) \Pi_2 \zeta_k. \quad (\text{A10})$$

Taking $\gamma_2 = -\rho_{slow}$ as the slow time scale, one has $\Pi_2 X_{\gamma_2, \theta}^N \Rightarrow \hat{X}_\theta$ as $N \rightarrow \infty$ ¹³ under more general conditions than Assumptions 2.1, 2.2. Under this context, Theorem 3.2 of Ref. 14 states that

$$\lim_{N \rightarrow \infty} \frac{\partial}{\partial \theta} \mathbb{E} \{f(X_{\gamma_2, \theta}^N(t))\} = \frac{\partial}{\partial \theta} \mathbb{E} \{f_\theta(\hat{X}_\theta(t))\}, \quad (\text{A11})$$

where $f_\theta(v) = \int_{\mathbb{H}_v} f(v + y) \pi_\theta^v(dy)$.

b. Equivalence of fast-classes and remainder spaces

Here we show how the TTS framework is equivalent to the scaling rate framework. Consider a TTS reaction network as described by (3). Taking $N = 1/\varepsilon$, $\rho_{fast} = 1$, $\rho_{slow} = 0$, it is easy to see that

$$X_{0, \theta}^N(t) = X_\theta^N(t) \stackrel{\mathcal{D}}{=} X^\varepsilon(t)$$

so $\lim_{N \rightarrow \infty} X_{0, \theta}^N(t) = \lim_{\varepsilon \rightarrow 0} X^\varepsilon(t)$. It remains to identify $f_\theta(\hat{X}_\theta(t))$ from (A10), (A11) with $\bar{f}(\bar{X}(t))$ from (8). We do so by showing the equivalence of the fast-classes \mathcal{M}_l and the remainder spaces \mathbb{H}_v .

Lemma 1.1. *The projection map Π_2 is invariant on fast-classes \mathcal{M}_l . The set of remainder spaces $\{\mathbb{H}_v : v \in \Pi_2 \mathcal{M}\}$ is in one-to-one correspondence with the set of fast-classes $\{\mathcal{M}_l\}$. Additionally, each $x \in \mathcal{M}_l$ corresponds to a unique element $y \in \mathbb{H}_v$ for some $v \in \Pi_2 \mathcal{M}$.*

Proof. Define $\eta : \{\mathcal{M}_l\} \rightarrow \{\mathbb{H}_v : v \in \Pi_2 \mathcal{M}\}$ by

$$\eta(\mathcal{M}_l) = \mathbb{H}_{\Pi_2(x)} \quad \text{for any } x \in \mathcal{M}_l.$$

Then η is well-defined, since $x, y \in \mathcal{M}_l$ implies that $y = x + \sum_{k \in \Gamma_1} c_k \zeta_k$ for some $c_k \in \mathbb{N}$, and $\Pi_2(\zeta_k) = \mathbf{0}$ for all $k \in \Gamma_1$ gives $\Pi_2(y) = \Pi_2(x)$. Clearly, η is also onto.

It remains to establish η is injective. It is sufficient to show that if $\Pi_2(x) = v$ and $\Pi_2(x') = v$ for $x, x' \in \mathcal{M}$, then x and x' belong to the same fast-class \mathcal{M}_l . Since Π_2 projects onto the span of the complement of $\text{span}\{\zeta_k : k \in \Gamma_1\}$, we have $v - x = \sum_{k \in \Gamma_1} c_k \zeta_k$ and $v - x' = \sum_{k \in \Gamma_1} c'_k \zeta_k$ for some $c_k, c'_k \in \mathbb{R}$.

Then $x' = x + \sum_{k \in \Gamma_1} (c_k - c'_k) \zeta_k$, with $x', x, \zeta_k \in \mathbb{N}^d$, so it follows that $(c_k - c'_k) \in \mathbb{N}$ for all k . Hence x and x' communicate by fast reactions and thus belong to the same fast-class \mathcal{M}_l . Therefore, Π_2 is invariant on fast-classes and η is injective.

Finally, since $\Pi_2(\cdot)$ is invariant on fast-classes \mathcal{M}_l , it follows that $x \rightarrow x - \Pi_2(x)$ bijectively maps elements of $\cup_l \mathcal{M}_l$ to elements of $\cup_{v \in \Pi_2(\mathcal{M})} \mathbb{H}_v$ such that $x, x' \in \mathcal{M}_l$ implies $(x - \Pi_2(x)), (x' - \Pi_2(x')) \in \mathbb{H}_{\Pi_2(\mathcal{M}_l)}$. \square

Because of the direct correspondence between $\{\Pi_2(x) : x \in \mathcal{M}\}$ and $\{\mathcal{M}_l\}$, we see (upon reordering states) that for $v = \Pi_2(\mathcal{M}_l)$, we have $\pi_\theta^v = \tilde{\pi}_\theta^{(l)}$, so that $\hat{\lambda}_k(v, \theta) = \bar{\lambda}_k(\mathcal{M}_l, \theta)$, and $f_\theta(v) = \bar{f}(\mathcal{M}_l, \theta)$. Thus, $f_\theta(\hat{X}(t))$ has the same distribution as $\bar{f}(\mathcal{M}_l, \theta)$ and so $\frac{\partial}{\partial \theta} \mathbb{E} \{f_\theta(\hat{X}_\theta(t))\} = \frac{\partial}{\partial \theta} \mathbb{E} \{\bar{f}(\bar{X}_\theta(t))\}$. Therefore, using (A11) we have

$$\begin{aligned} \lim_{\varepsilon \rightarrow 0} \partial_\theta \mathbb{E} \{f(X^\varepsilon(t))\} &= \lim_{N \rightarrow \infty} \partial_\theta \mathbb{E} \{f(X_{0,\theta}^N(t))\} \\ &= \partial_\theta \mathbb{E} \{f_\theta(\hat{X}_\theta(t))\} = \partial_\theta \mathbb{E} \{\bar{f}(\bar{X}(t))\} \end{aligned} \quad (\text{A12})$$

and hence Proposition 3.1.

4. Proof of Lemma 3.2

This proof of pseudo-inverse sensitivity representation (21) for continuous-time chains is adapted from an analogous result for discrete-time chains in Ref. 48.

Recall that for an ergodic CTMC with generator $Q(\theta)$, one can compute the steady-state probability vector $\pi(\theta)$ by solving the linear system $\pi \cdot Q = \mathbf{0}$ and $\pi \cdot \mathbf{1} = 1$. By differentiating $\pi(\theta) \cdot Q(\theta) = \mathbf{0}$, we have the relation

$$\frac{\partial \pi}{\partial \theta} Q = -\pi \frac{\partial Q}{\partial \theta}.$$

Then by expanding with the Moore-Penrose pseudo-inverse, we have

$$\left(\frac{\partial \pi}{\partial \theta} \right)' = (Q)' \left(\frac{-\partial Q}{\partial \theta} \right)' \pi' + [I - (Q')^+ Q'] w$$

for some vector w .

Now, we see by the projection property of $(Q')^+ Q' = (QQ')^+ = QQ^+$ that the operator $I - QQ^+$ is the projection operator onto the kernel of $Q' = \text{span of } \pi'$, so that $(I - QQ^+)w = \gamma \pi'$ for some scalar γ . Thus we have the relation

$$\frac{\partial \pi}{\partial \theta} = \pi \left(\frac{-\partial Q}{\partial \theta} \right) Q^+ + \gamma \pi$$

for some $\gamma \in \mathbb{R}$. It remains to determine γ to have a method of relating the sensitivity coefficient to a linear transformation of π .

Now, we can see that $\pi(\theta) \cdot \mathbf{1} = 1$, so that $\frac{\partial \pi}{\partial \theta} \cdot \mathbf{1} = 0$. Thus, we have

$$0 = \frac{\partial \pi}{\partial \theta} \cdot \mathbf{1} = \pi \left(\frac{-\partial Q}{\partial \theta} \right) Q^+ \mathbf{1} + \gamma \pi \mathbf{1}$$

so that

$$\gamma = \gamma \pi \mathbf{1} = \pi \frac{\partial Q}{\partial \theta} Q^+ \mathbf{1}.$$

Putting the above together, we can write $\frac{\partial \pi}{\partial \theta}$ as

$$\frac{\partial \pi}{\partial \theta} = \pi \left(\frac{\partial Q}{\partial \theta} \right) Q^+ [\mathbf{1} \pi - I]$$

as in (21).

APPENDIX B: ANALYTIC SOLUTION OF THE MODEL SYSTEM

In a well-mixed system with linear propensities, the time-evolution of the system can be obtained from a set of ordinary differential equations (ODE). The STS system can be modeled with a system of ODEs. The TTS system imposes algebraic constraints for the fast modes, resulting in an algebraic differential system of equations. In both cases, a set of adjoint ODEs can be used to compute sensitivities alongside species populations.

In our model system, gaseous species A adsorbs onto a catalyst surface, isomerizes to species B, and then desorbs. A diagram of the reaction network is shown in Figure 7. The reactions along with their rate laws are shown in Table II. N_A , N_B , and N_* denote the surface coverages of species A, B, and empty sites, respectively. The adsorption/desorption of species A is assumed to be much faster than the others. The separation of time scales is captured with the dimensionless parameter $\varepsilon \ll 1$. The system contains $M = 3$ species and $R = 6$ reactions. Mathematically, we use the $M \times 1$ column vector N to specify the species populations, where $N_1 = N_A$, $N_2 = N_B$, and $N_3 = N_*$.

The linear dependence of the reaction rates is written as

$$r(N) = \begin{bmatrix} \varepsilon^{-1} \alpha_1 N_3 \\ \varepsilon^{-1} \alpha_2 N_1 \\ \beta_1 N_1 \\ \beta_2 N_2 \\ \beta_3 N_2 \end{bmatrix}. \quad (\text{B1})$$

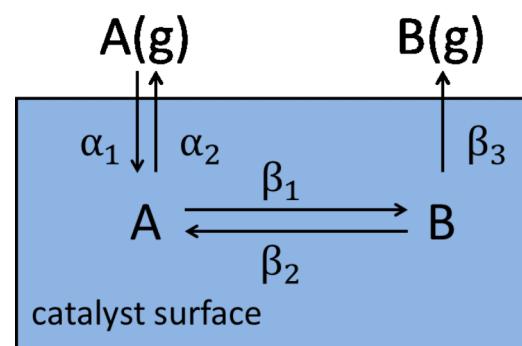


FIG. 7. Visual diagram of reaction network.

TABLE II. Reactions for the benchmark network.

Index	Reaction	Rate
1	$* \rightarrow A$	$\varepsilon^{-1} \alpha_1 N_*$
2	$A \rightarrow *$	$\varepsilon^{-1} \alpha_2 N_A$
3	$A \rightarrow B$	$\beta_1 N_A$
4	$B \rightarrow A$	$\beta_2 N_B$
5	$B \rightarrow *$	$\beta_3 N_B$

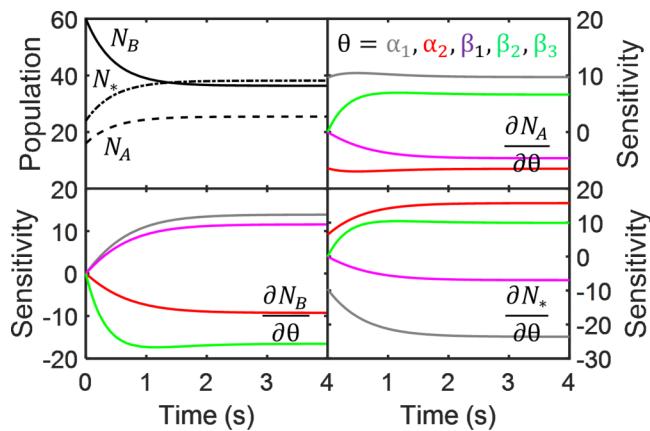


FIG. 8. Network results for the two time-scale system. Graphs show population counts (top left), derivatives of species A (top right), derivatives of species B (bottom left), and derivatives of empty sites (bottom right).

Each row of the stoichiometric matrix corresponds to a different species, which are N_1 , N_2 , and N_3 , respectively. The columns correspond to each of reactions 1-5 in order. Extracting the information from Table II and putting it in mathematical form gives the $M \times R$ stoichiometric matrix

$$S = \begin{bmatrix} 1 & -1 & -1 & 1 & 0 \\ 0 & 0 & 1 & -1 & 1 \\ -1 & 1 & 0 & 0 & -1 \end{bmatrix}. \quad (\text{B2})$$

The transformation matrix

$$T = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 1 & 1 & 1 \end{bmatrix} \quad (\text{B3})$$

yields $y = T \cdot N$ and T can be decomposed into $T_f = [1 \ 0 \ 0]$ and $T_s = \begin{bmatrix} 1 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$ for the slow modes by looking at the 0 rows of S_f^T . This gives us the transformed variables as $y_f = [N_1]$ and $y_s = \begin{bmatrix} N_2 \\ N_1 + N_2 + N_3 \end{bmatrix}$.

TABLE III. Comparison of the CLR and CELR estimators with the ODE solution. Values show actual values rather than errors. Values in the table refer to the sensitivity of the species B with respect to the parameter given by the row label. 95% confidence intervals are based on statistical noise.

	$t = 1.3$ s		
	ODE	CLR	CELR
α_1	11.9	11.8 ± 1.5	7.2 ± 0.9
α_2	-7.9	-7.7 ± 1.0	-4.7 ± 0.6
β_1	9.9	10.0 ± 1.2	6.0 ± 0.8
β_2	-17.4	-17.2 ± 2.4	-14.5 ± 1.7
β_3	-17.4	-17.1 ± 3.4	-13.3 ± 2.4
	$t = 100$ s		
	ODE	CLR	CELR
α_1	13.9	2.3 ± 13.7	13.9 ± 1.4
α_2	-9.3	-3.4 ± 8.9	-9.2 ± 1.0
β_1	11.6	4.1 ± 10.6	11.7 ± 1.2
β_2	-16.5	-12.8 ± 18.1	-16.1 ± 2.0
β_3	-16.5	-19.9 ± 29.0	-16.5 ± 3.2

In the context of our example problem, we can assign physical meaning to the transformation: The variable $y_1 = N_A$ is affected by both slow and fast reactions. For a given set of slow variables, we can solve for y_1 to specify the equilibrium constraint of $r_1 = r_2$. The variable $y_2 = N_B$ is unaffected by the fast adsorption/desorption of A, but is affected by the slow reactions. Finally, the variable $y_3 = N_A + N_B + N_*$ is a second “slow” variable. In this example $y_3 = 1$ applies at all times due to stoichiometric constraints. However, it is still identified as a “slow mode” because this constraint is not a consequence of disparities in reaction time scales.

The system is simulated with the choice of parameters $N_0 = \begin{bmatrix} 30 \\ 60 \\ 10 \end{bmatrix}$, $\alpha_1 = 1$, $\alpha_2 = 1.5$, $\beta_1 = 2$, $\beta_2 = 1$, $\beta_3 = 0.4$, $\varepsilon = 0.01$. The simulation results are shown in Figure 8. Table III shows values at $t = 1.3$ s and $t = 100$ s along with CLR and CELR estimates with statistical confidence intervals. Derivatives with respect to β_2 and β_3 overlap because both affect system properties through the independent parameter $\beta_2 + \beta_3$.

In general, the system parameters need not be the rate constants themselves. A different parameterization would involve a transformation of the rate constants. Sensitivities could be obtained through a chain rule.

- ¹A. Chatterjee and D. G. Vlachos, “An overview of spatial microscopic and accelerated kinetic Monte Carlo methods,” *J. Comput.-Aided Mater. Des.* **14**(2), 253–308 (2007).
- ²D. T. Gillespie, “Approximate accelerated stochastic simulation of chemically reacting systems,” *J. Chem. Phys.* **115**(4), 1716–1733 (2001).
- ³Y. Cao, D. T. Gillespie, and L. R. Petzold, “Avoiding negative populations in explicit Poisson tau-leaping,” *J. Chem. Phys.* **123**(5), 054104 (2005).
- ⁴Y. Cao, D. T. Gillespie, and L. R. Petzold, “Efficient step size selection for the tau-leaping simulation method,” *J. Chem. Phys.* **124**(4), 044109 (2006).
- ⁵M. Rathinam, L. R. Petzold, Y. Cao, and D. T. Gillespie, “Stiffness in stochastic chemically reacting systems: The implicit tau-leaping method,” *J. Chem. Phys.* **119**(24), 12784–12794 (2003).
- ⁶A. Chatterjee, D. G. Vlachos, and M. A. Katsoulakis, “Binomial distribution based-leap accelerated stochastic simulation,” *J. Chem. Phys.* **122**(2), 024112 (2005).
- ⁷T. Tian and K. Burrage, “Binomial leap methods for simulating stochastic chemical kinetics,” *J. Chem. Phys.* **121**(21), 10356–10364 (2004).
- ⁸H. Salis and Y. Kaznessis, “Accurate hybrid stochastic simulation of a system of coupled chemical or biochemical reactions,” *J. Chem. Phys.* **122**(5), 054103 (2005).
- ⁹A. Samani and D. G. Vlachos, “Overcoming stiffness in stochastic simulation stemming from partial equilibrium: A multiscale Monte Carlo algorithm,” *J. Chem. Phys.* **123**(14), 144114 (2005).
- ¹⁰Y. Cao, D. T. Gillespie, and L. R. Petzold, “The slow-scale stochastic simulation algorithm,” *J. Chem. Phys.* **122**(1), 014116 (2005).
- ¹¹E. Weinan, D. Liu, and E. Vanden-Eijnden, “Nested stochastic simulation algorithms for chemical kinetic systems with multiple time scales,” *J. Comput. Phys.* **221**(1), 158–180 (2007).
- ¹²C. Huang and D. Liu, “Strong convergence and speed up of nested stochastic simulation algorithm,” *Commun. Comput. Phys.* **15**, 1207–1236 (2014).
- ¹³H.-W. Kang and T. G. Kurtz, “Separation of time-scales and model reduction for stochastic reaction networks,” *Ann. Appl. Probab.* **23**(2), 529–583 (2013).
- ¹⁴A. Gupta and M. Khammash, “Sensitivity analysis for stochastic chemical reaction networks with multiple time-scales,” *Electron. J. Probab.* **19**(59), 1–53 (2014).
- ¹⁵P. Dupuis, M. A. Katsoulakis, Y. Pantazis, and P. Plechac, “Path-space information bounds for uncertainty quantification and sensitivity analysis of stochastic dynamics,” e-print [arXiv:1503.05136](https://arxiv.org/abs/1503.05136) [math] (2015).

¹⁶E. S. Wolf and D. F. Anderson, "Hybrid pathwise sensitivity methods for discrete stochastic models of chemical reaction systems," *J. Chem. Phys.* **142**(3), 034103 (2015).

¹⁷P. W. Sheppard, M. Rathinam, and M. Khammash, "A pathwise derivative approach to the computation of parameter sensitivities in discrete stochastic chemical systems," *J. Chem. Phys.* **136**(3), 034115 (2012).

¹⁸T. Wang and M. Rathinam, "Efficiency of the Girsanov transformation approach for parametric sensitivity analysis of stochastic chemical kinetics," e-print [arXiv:1412.1005](https://arxiv.org/abs/1412.1005) [math] (2014).

¹⁹A. Gupta and M. Khammash, "An efficient and unbiased method for sensitivity analysis of stochastic reaction networks," *J. R. Soc. Interface* **11**(101), 20140979 (2014).

²⁰G. Yin and Q. Zhang, *Continuous-Time Markov Chains and Applications: A Two-Time-Scale Approach* (Springer, New York, 2013).

²¹M. Núñez and D. G. Vlachos, "Steady state likelihood ratio sensitivity analysis for stiff kinetic Monte Carlo simulations," *J. Chem. Phys.* **142**(4), 044108 (2015).

²²S. Plyasunov and A. P. Arkin, "Efficient stochastic sensitivity analysis of discrete event systems," *J. Comput. Phys.* **221**(2), 724–738 (2007).

²³P. B. Warren and R. J. Allen, "Steady-state parameter sensitivity in stochastic modeling via trajectory reweighting," *J. Chem. Phys.* **136**(10), 104106 (2012).

²⁴P. W. Glynn, "Likelihood ratio gradient estimation for stochastic systems," *Commun. ACM* **33**(10), 75–84 (1990).

²⁵D. F. Anderson and T. G. Kurtz, "Continuous time Markov chain models for chemical reaction networks," in *Design and Analysis of Biomolecular Circuits*, edited by H. Koepll, G. Setti, M. di Bernardo, and D. Densmore (Springer, New York, 2011), pp. 3–42.

²⁶S. N. Ethier and T. G. Kurtz, *Markov Processes: Characterization and Convergence* (John Wiley & Sons, 1986).

²⁷M. A. Gibson and J. Bruck, "Efficient exact stochastic simulation of chemical systems with many species and many channels," *J. Phys. Chem. A* **104**(9), 1876–1889 (2000).

²⁸D. F. Anderson, "A modified next reaction method for simulating chemical systems with time dependent propensities and delays," *J. Chem. Phys.* **127**(21), 214107 (2007).

²⁹M. Rathinam, P. W. Sheppard, and M. Khammash, "Efficient computation of parameter sensitivities of discrete stochastic chemical reaction networks," *J. Chem. Phys.* **132**(3), 034103 (2010).

³⁰D. Anderson, "An efficient finite difference method for parameter sensitivities of continuous time Markov chains," *SIAM J. Numer. Anal.* **50**(5), 2237–2258 (2012).

³¹A. Gupta and M. Khammash, "Sensitivity analysis for stochastic chemical reaction networks with multiple time-scales," *Electron. J. Probab.* **19**(59), 1–53 (2014).

³²Y. Pantazis, M. A. Katsoulakis, and D. G. Vlachos, "Parametric sensitivity analysis for biochemical reaction networks based on pathwise information theory," *BMC Bioinf.* **14**, 311 (2013).

³³A. Hashemi, P. Plechac, and T. Wang, "Rate of convergence of averaged sensitivities in multiscale systems" (unpublished).

³⁴J. A. McGill, B. A. Ogunnaike, and D. G. Vlachos, "Efficient gradient estimation using finite differencing and likelihood ratios for kinetic Monte Carlo simulations," *J. Comput. Phys.* **231**(21), 7170–7186 (2012).

³⁵S. Asmussen and P. W. Glynn, *Stochastic Simulation: Algorithms and Analysis* (Springer Science & Business Media, 2007).

³⁶G. Arampatzis, M. A. Katsoulakis, and L. Rey-Bellet, "Efficient estimators for likelihood ratio sensitivity indices of complex stochastic dynamics," e-print [arXiv:1510.05439](https://arxiv.org/abs/1510.05439) [math, stat] (2015).

³⁷B. Munsky and M. Khammash, "The finite state projection algorithm for the solution of the chemical master equation," *J. Chem. Phys.* **124**(4), 044104 (2006).

³⁸V. Sunkara and M. Hegland, "An optimal finite state projection method," *Procedia Comput. Sci.* **1**(1), 1579–1586 (2010).

³⁹Y. Cao, D. Gillespie, and L. Petzold, "Multiscale stochastic simulation algorithm with stochastic partial equilibrium assumption for chemically reacting systems," *J. Comput. Phys.* **206**(2), 395–411 (2005).

⁴⁰D. A. Levin, Y. Peres, and E. L. Wilmer, *Markov Chains and Mixing Times* (American Mathematical Society, 2009).

⁴¹C. J. Geyer, "Practical Markov chain Monte Carlo," *Stat. Sci.* **7**(4), 473–483 (1992).

⁴²C. Kipnis and S. R. S. Varadhan, "Central limit theorem for additive functionals of reversible Markov processes and applications to simple exclusions," *Commun. Math. Phys.* **104**(1), 1–19 (1986).

⁴³B. A. Berg, "Introduction to Markov chain Monte Carlo simulations and their statistical analysis," in *Markov Chain Monte Carlo: Innovations and Applications*, edited by W. S. Kendall, F. Liang, and J.-S. Wang (World Scientific, 2005).

⁴⁴A. Sokal, "Monte Carlo methods in statistical mechanics: Foundations and new algorithms," in *Functional Integration* (Springer, 1997), pp. 131–192.

⁴⁵S. P. Meyn and R. L. Tweedie, *Markov Chains and Stochastic Stability* (Cambridge University Press, 2009).

⁴⁶C. Alexopoulos and A. F. Seila, "Implementing the batch means method in simulation experiments," in *Proceedings of the 28th Conference on Winter Simulation, WSC '96* (IEEE Computer Society, Washington, DC, USA, 1996), pp. 214–221.

⁴⁷P. W. Glynn and D. L. Iglehart, "Simulation output analysis using standardized time series," *Math. Oper. Res.* **15**(1), 1–16 (1990).

⁴⁸R. E. Funderlic and C. D. Meyer, Jr., "Sensitivity of the stationary distribution vector for an ergodic Markov chain," *Linear Algebra Appl.* **76**, 1–17 (1986).