

Spectral Methods for Parametric Sensitivity in Stochastic Dynamical Systems

D. Kim, B. J. Debusschere, and H. N. Najm

Sandia National Laboratories, Livermore, California

ABSTRACT Stochastic dynamical systems governed by the chemical master equation find use in the modeling of biological phenomena in cells, where they provide more accurate representations than their deterministic counterparts, particularly when the levels of molecular population are small. The analysis of parametric sensitivity in such systems requires appropriate methods to capture the sensitivity of the system dynamics with respect to variations of the parameters amid the noise from inherent internal stochastic effects. We use spectral polynomial chaos expansions to represent statistics of the system dynamics as polynomial functions of the model parameters. These expansions capture the nonlinear behavior of the system statistics as a result of finite-sized parametric perturbations. We obtain the normalized sensitivity coefficients by taking the derivative of this functional representation with respect to the parameters. We apply this method in two stochastic dynamical systems exhibiting bimodal behavior, including a biologically relevant viral infection model.

INTRODUCTION

In this article, we introduce new methodologies for studying the parametric sensitivity of a stochastic dynamical system (SDS). By parametric sensitivity, we mean the sensitivity of the system dynamics with respect to variations of the parameters that characterize the system. SDSs play an important role as models for biochemical gene regulatory networks, in modeling gene regulation activity in viral kinetics (1,2), in cell signaling pathways (3–5), biochemical switches (6,7), and as models for surface reaction processes (8). The utility of the stochastic representation derives from its ability to capture the true system dynamics, in contrast with deterministic constructions, which neglect the role of random noise in these systems. We focus on SDSs, such as stochastic chemical reaction systems, which exhibit inherent stochastic noise. These systems are governed by the chemical master equation (9,10). The direct solution of this equation is intractable for most practical purposes. Instead, time traces of the system are typically obtained using the stochastic simulation algorithm (SSA) (11,12). Because of inherent internal noise in these systems, the output states become stochastic processes for which the methods of parametric sensitivity that apply to continuous deterministic systems, and that require evaluation of Jacobians (hence derivatives) of the chemical source terms or the state trajectories, are not well defined.

Recently, sensitivity analysis for such systems was performed based on probability density function (PDF) sensitivity, using an analog of the classical sensitivity and the Fisher information matrix (13). Gunawan et al. (13) introduced four measures for parametric sensitivity based on the derivatives of the PDF of the system state with respect to the

parameters of interest. As those derivatives were approximated using centered finite differences between PDFs generated for different values of the parameters, care had to be taken to select step sizes in the parameters that were large enough to accurately see the effect of the parametric changes over the level of internal noise, but at the same time small enough to avoid excessive truncation error in the finite difference scheme. Generally, this balance may be hard to realize with such direct application of finite difference to the PDFs, as systems with large internal noise may require such large parametric changes that the corresponding truncation errors become excessive.

In this work, we present a construction for handling sensitivity in a SDS context that offers the flexibility of a high-order spectral representation of system output statistics in terms of model parameters. The order of the spectral representation is chosen to properly represent the nonlinear changes in system output statistics as a result of larger-magnitude parametric perturbations (as needed to cause changes in the output statistics that are sufficiently large compared to the internal noise in those quantities). This allows an accurate and robust evaluation of sensitivities. Moreover, an additional utility of this construction is that it enables studies of the predictability of the system given uncertainty, variability, or external noise in the model parameters, and allows estimation of corresponding uncertainty of predicted output state statistics. We focus overall on the sensitivity of output state statistics, rather than of the PDF itself, as various statistics can be designed to probe desired aspects of the PDF or, in general, of the observed random response of the system.

The sensitivity/uncertainty construction proposed here introduces presumed/known inherent uncertainties in model parameters, to enable both sensitivity and predictability studies. In the sensitivity context, presumed small parametric uncertainties are chosen, which are nonetheless sufficiently

Submitted March 17, 2006, and accepted for publication September 13, 2006.

Address reprint requests to B. J. Debusschere, Tel.: 925-294-3833; E-mail: bjdebuss@sandia.gov.

D. Kim's present address is Georgetown University, Washington, DC.

© 2007 by the Biophysical Society

0006-3495/07/01/379/15 \$2.00

doi: 10.1529/biophysj.106.085084

large compared to the inherent stochastic noise in the system. On the other hand, in the uncertainty quantification context, parametric uncertainties are known empirically. In either case, we represent uncertain quantities using a stochastic representation, such that each uncertain parameter becomes a random variable in the system model. It is important to differentiate between this randomness, and the inherent random response of the system with deterministic fixed parameters. For a given choice of the parameters, the resulting system statistics (means, or higher moments of the system response) are deterministic quantities. On the other hand, with random (uncertain) model parameters, these same statistics remain as random variables/stochastic processes in terms of the parameters. We model the uncertain model parameters, and the induced random (uncertain) response of the statistics of the system states, using a spectral polynomial-chaos (PC) construction (14–19). In this context, (second-order) random processes are represented using a spectral stochastic expansion in terms of an underlying set of orthogonal polynomials and their associated density, as is further explained below. The strength of the construction stems from its utilization of functional representations, as opposed to much more expensive collocation approaches, to model random quantities. This functional representation also enables direct evaluation of sensitivities of model predictions with respect to the uncertain parameters (20,21). Notably, in the uncertainty quantification context, the PC construction provides not only the sensitivity coefficients, but also a quantification of the confidence in these sensitivities. In general, higher-order PC representations can properly capture the nonlinear behavior of the system dynamics due to significant parametric perturbations. They also allow for inspection of nonlinear sensitivity effects, as variations in multiple parameters can be considered simultaneously. We define the normalized sensitivity coefficient based on the derivative of the PC expansion with respect to the parameter so that we can quantify the parametric sensitivity efficiently.

While this approach for evaluation of sensitivity coefficients has been demonstrated in deterministic chemical systems (21), its use in the SDS context, where internal noise is significant, is novel. Thus, the objective of this work is to outline the requisite development for implementation of the PC-based sensitivity analysis in the SDS context. We discuss means of evaluating the sensitivity coefficients in the presence of significant levels of internal noise. We illustrate the application of the construction in the context of two model stochastic dynamical systems that exhibit nontrivial dynamical behavior, including bifurcation and bimodal statistics, which are relevant in biological systems. The second model, in particular, describes the dynamics of a cellular subsystem response to a viral infection.

The article is arranged as follows: in the next section, we give a brief overview of the mathematical formulation of the PC construction, focusing on Gauss-Hermite PC expansions. Then we provide an algorithmic discussion of PC-based

parametric sensitivity in SDSs. Numerical examples are given in the following section and the main findings are summarized at the end.

POLYNOMIAL CHAOS

Under specific conditions, a random variable can be expressed as a spectral expansion in terms of suitable orthogonal basis functions associated with random variables with a particular density. A well-studied example is the representation of random variables using Hermite polynomials of normally distributed random variables. Depending on the distribution of the random variable to be represented, other combinations of basis functions and associated densities may be more appropriate, such as Charlier polynomials with the Poisson distribution or Laguerre polynomials with the γ -distribution (22). These spectral expansions are generally referred to as polynomial-chaos (PC) expansions, following Wiener (23). In this work, we will mainly focus on Gauss-Hermite PC expansions.

Any second-order random variable X can be represented using a Gauss-Hermite PC expansion as

$$X(\theta) = \sum_{k=0}^{\infty} X_k \Psi_k(\xi_1, \xi_2, \dots), \quad (1)$$

where θ samples a random event space, the Ψ_k values are orthogonal Hermite polynomials of the uncorrelated standard Gaussian random variables $\{\xi_j\}_{j=1}^{\infty}$ (14), and the X_k values are (deterministic) spectral coefficients, here called “PC coefficients”. For practical purposes, this expansion expression, Eq. 1, is generally truncated after the term(s) of order N_{ord} over finite N_{dim} stochastic dimension(s) so that

$$X(\theta) = \sum_{k=0}^P X_k \Psi_k(\xi_1, \xi_2, \dots, \xi_{N_{\text{dim}}}), \quad P = \frac{(N_{\text{dim}} + N_{\text{ord}})!}{(N_{\text{dim}}! N_{\text{ord}}!)} - 1. \quad (2)$$

For simplicity, we refer to this truncated Gauss-Hermite PC expansion as a “PC expansion”.

Using the orthogonality of the Ψ_k values, the PC coefficients X_k can be obtained as

$$X_k = \frac{\langle X \Psi_k \rangle}{\langle \Psi_k^2 \rangle}, \quad k = 0, \dots, P, \quad (3)$$

where the angle brackets denote the Gaussian-weighted expectation over the N_{dim} -dimensional stochastic space.

For example, if only one stochastic dimension, i.e., $N_{\text{dim}} = 1$, is considered, $P = N_{\text{ord}}$ in Eq. 2 and the first five one-dimensional Hermite polynomials are

$$\begin{aligned} \Psi_0(\xi) &= 1 \\ \Psi_1(\xi) &= \xi \\ \Psi_2(\xi) &= \xi^2 - 1 \\ \Psi_3(\xi) &= \xi^3 - 3\xi \\ \Psi_4(\xi) &= \xi^4 - 6\xi^2 + 3. \end{aligned} \quad (4)$$

The use of the PC construction for uncertainty quantification has been demonstrated in the literature (14–21). In general, there are two approaches for doing this, an intrusive and a nonintrusive approach. The intrusive approach involves substituting the PC expansions for model parameters and fields in the governing equations, and applying the Galerkin projection, Eq. 3, to arrive at equations governing the evolution of the PC mode strengths of the solution. Given the inherent stochastic noise in the systems at hand, this approach is not directly extensible to the present context. Rather, we apply the nonintrusive spectral projection (NISP) approach (16,19–21) using Gauss-Hermite quadrature (16,24,25), as outlined in the following.

Consider a general SDS model, in which an observable of interest (a state statistic such as its mean or standard deviation) at time t , Y_t , is a function of t , and some system parameter λ , such that $Y_t = F(t; \lambda)$, where $Y_t \in \mathbb{R}$, and where the obvious dependence on initial conditions is omitted for convenience. This model is not available analytically, of course; instead, one would arrive at it, for example, based on many simulations with the stochastic simulation algorithm (SSA), given the initial conditions and the parameter λ . Now, given some known uncertainty in λ , we seek to quantify the uncertainty in Y_t . The NISP approach proceeds first by defining λ as a random variable with known statistics, and constructing its PC expansion accordingly. For illustration, assume a single underlying stochastic dimension ξ , a PC order P , and that λ has a normal distribution with known mean μ and standard deviation σ , such that

$$\lambda = \sum_{k=0}^P \lambda_k \Psi_k(\xi) = \mu + \sigma \xi. \quad (5)$$

The goal then is to evaluate the mode strengths $\{Y_{t,k}\}_{k=0}^P$ of the corresponding PC expansion of Y_t ,

$$Y_t = \sum_{k=0}^P Y_{t,k} \Psi_k(\xi), \quad (6)$$

which provides a complete characterization of the uncertainty in Y_t . These modes are evaluated following Eq. 3,

$$Y_{t,k} = \frac{\langle Y_t \Psi_k \rangle}{\langle \Psi_k^2 \rangle},$$

where the expectations are found by evaluating the equivalent stochastic integrals over ξ -space, e.g.,

$$\langle Y_t \Psi_k \rangle = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} Y_t \Psi_k e^{-\xi^2/2} d\xi, \quad (7)$$

using Gauss-Hermite quadrature (24)

$$\int_{-\infty}^{\infty} f(x) e^{-x^2} dx \simeq \sum_{i=1}^n w_i f(x_i).$$

Note that each quadrature point value for ξ corresponds to an associated value for $Y_t = F(t; \lambda(\xi))$, allowing the evaluation of the integral. As the evaluation of each $Y_t(\xi)$

typically involves extracting a statistical property from a large number of simulated trajectories of the stochastic dynamical system being studied, obtaining realizations of Y_t can be quite expensive computationally. However, those simulations can be sped up significantly by using efficient multiscale algorithms as well as parallelization (26–30). Regardless, for computational efficiency, it is crucial that a minimal number of quadrature points be used. In a one-dimensional (in ξ) setting, it is clear that Gaussian quadrature is optimal, since an n -point quadrature rule provides an exact evaluation of integrals of polynomials up to degree $2n - 1$. As we have shown in previous work (16), Gaussian quadrature is still more efficient than, say, Monte Carlo evaluation of these integrals for up to 3–5 stochastic dimensions. At higher dimensionality, cubature formulae (31) are more feasible. In the present context, we study the sensitivity of the system to one parameter at a time, hence we have a single dimensional stochastic space, such that the choice of Gaussian quadrature is obvious. With $N_{\text{dim}} = 1$, the number of requisite quadrature points depends ultimately on the PC order. Thus, e.g., for second-order chaos ($N_{\text{ord}} = 2$) the product $Y_t(\xi) \Psi_k(\xi)$ is a polynomial of degree 4, requiring $n = 3$ quadrature points only. In general, given P^{th} -order one-dimensional chaos, NISP requires $(P + 1)$ quadrature, or collocation points. It is important to note, however, that this Gauss-Hermite quadrature using $(P + 1)$ quadrature points is exact only if Y_t is well approximated by a polynomial in ξ of degree no more than P .

Thus, in the uncertainty quantification context, the above procedure allows efficient sampling of uncertain parameter distributions (at the Gauss-Hermite quadrature points) to construct the PC representation of the uncertainty in model outputs. Reagan et al. (21) used this approach in the context of deterministic chemical systems, and described a procedure for extracting parametric sensitivity information from the resulting expansions. In the following, we specialize this approach for the case where parametric uncertainties are not, in fact, known. Rather, they will be assumed with the sole goal of probing the resulting PC representation of the model output statistics to evaluate parametric sensitivities.

POLYNOMIAL CHAOS-BASED PARAMETRIC SENSITIVITY

This section outlines in detail the algorithm for polynomial chaos (PC)-based parametric sensitivity analysis in stochastic dynamical systems (SDSs) governed by chemical master equations.

The first step is to determine the statistical properties of the output states and the model parameters of interest for the parametric sensitivity. Moments and conditional moments of the system state distribution, for example, are common statistical properties of interest, and can be obtained from system state realizations generated with the stochastic simulation algorithm (SSA). The sensitivity analysis is then

performed by postulating an artificial, small perturbation to each of the system parameters and analyzing the corresponding changes in the statistical properties of interest. Crucially in the determination of these statistical properties, the number of realizations in SSA simulations should be large enough to minimize the effect of inherent internal stochastic noise on the variations of the statistical properties due to parameter changes.

Let Y_i be a statistical property of the system output state for which we want to analyze the sensitivity with respect to parameter λ_j . Following the above nonintrusive spectral projection (NISP) construction from uncertainty quantification, consider λ_j to be uncertain, specifically a Gaussian random variable whose mean is the nominal parameter value, μ_j , and whose standard deviation σ_j is prescribed as

$$\lambda_j = \mu_j + \sigma_j \xi_j, \quad (8)$$

where the ξ_j values are independent and identically distributed standard Gaussian random variables for each model parameter. The size of the parameter perturbation, σ_j , needs to be chosen large enough so that variations in the system state properties due to this parameter perturbation are significantly larger than the inherent stochastic noise in those properties. However, it must also be kept small enough to avoid highly nonlinear or discontinuous changes in the system response. In general, parameters with small effects on the system response allow larger perturbations (e.g., 10%) without causing discontinuous changes in system response. However, parameters to which the system is very sensitive necessitate smaller associated perturbations (e.g., 0.5%). This postulated distribution in the parameter λ_j can then be propagated through the stochastic dynamical system, following the above NISP procedure, arriving at the P^{th} -order PC expansion for Y_i in ξ_j :

$$Y_i(t, \theta) \approx \sum_{k=0}^P (Y_i)_k(t) \Psi_k(\xi_j). \quad (9)$$

We recall that the important point here is that the statistical properties of interest need to be evaluated only for the values of the model parameters at the Gauss-Hermite quadrature points. This approach is therefore much more efficient than Monte-Carlo based methods, which typically require evaluations of the system properties at many thousands of values of the model parameters.

By substituting ξ_j in Eq. 8 into the P^{th} -order PC expansion of Y_i in Eq. 9 we can rewrite Y_i as a polynomial in λ_j of degree P :

$$Y_i(t, \lambda_j) \approx \sum_{k=0}^P (Y_i)_k(t) \Psi_k\left(\frac{\lambda_j - \mu_j}{\sigma_j}\right). \quad (10)$$

Given the perturbation σ_j , this P^{th} -order polynomial Eq. 10 represents the dependence of Y_i on λ_j in the local neighborhood of μ_j . The order of this polynomial response function needs to be chosen high enough to allow an accurate

representation of the local nonlinear behavior of Y_i over the range of λ_j . This ability to properly represent nonlinear behavior helps to mitigate the effect of the stochastic noise in the evaluations of Y_i at the quadrature points as it allows the selection of larger perturbation sizes that result in changes to Y_i that are significantly larger than the noise in those quantities. This issue will be discussed further below, with particular SDS examples.

Following Reagan et al. (21), we next define the normalized sensitivity coefficient S_{ij}^P and the seminormalized sensitivity coefficient SS_{ij}^P based on the derivative of the P^{th} -order polynomial Eq. 10 with respect to the parameter λ_j . In contrast with Reagan et al. (21), however, we do not have known uncertainties in model parameters. Rather, the uncertainties here are presumed solely for purposes of sensitivity analysis. As a result, we do not address issues of confidence in the sensitivity coefficients as in Reagan et al. (21), but rather focus on the evaluation of the sensitivity coefficient at the nominal value μ_j ,

$$S_{ij}^P(t) := \frac{\partial Y_i / Y_i}{\partial \lambda_j / \lambda_j}(t, \mu_j), \quad (11)$$

which is an estimate of the slope of Y_i in the direction of λ_j at the nominal value μ_j in the parameter space. That is, S_{ij}^P defined in Eq. 11 represents an estimate of the percentage change in Y_i caused by a 1% change in the parameter λ_j from the nominal value μ_j . Therefore, it indicates the effect and relative importance of the various input parameters on the statistical properties of interest at each point in time. The normalization can cause problems, however, when $|Y_i| \approx 0$. In those cases, it is more robust to use the seminormalized formulation for the sensitivity coefficients (21):

$$SS_{ij}^P(t) := \frac{\partial Y_i}{\partial \lambda_j / \lambda_j}(t, \mu_j). \quad (12)$$

Note that the algorithm outlined so far is general for any number of stochastic dimensions and can therefore be used to generate multidimensional, nonlinear response surfaces that characterize the dependency of Y_i on multiple parameters λ_j . These response surfaces can then be used to examine the coupled effect of several parameters. For the purpose of this study, however, we focus only on one-dimensional sensitivity analysis, i.e., the sensitivity will be analyzed with respect to one parameter at a time. For this one-dimensional analysis, Eq. 11 can be further expanded as follows (21). Given the fact that $\Psi'_{k>0}(\xi) = k\Psi_{k-1}(\xi)$ for one-dimensional Hermite polynomials, and by using Eqs. 8 and 9, and the chain rule, we can rewrite Eq. 11 explicitly as

$$S_{ij}^P(t) = \frac{\sum_{k=1}^P (Y_i)_k(t) k \Psi_{k-1}(0) / \sum_{k=0}^P (Y_i)_k(t) \Psi_k(0)}{\sigma_j / \mu_j} \quad (13)$$

and similarly for SS_{ij}^P . Note also that, for $P \geq 2$,

$$S_{ij}^P(t) \neq S_{ij}^1(t) = \frac{(Y_i)_1(t) / (Y_i)_0(t)}{\sigma_j / \mu_j} \quad (14)$$


as the higher-order modes generally have nonzero values at $\xi_j = 0$.

EXAMPLES

We apply the methodologies developed in the previous section for PC-based parametric sensitivity to two SDSs: the Schlögl model (10) and a viral infection model (1). Several statistical properties of interest are examined that reflect the system dynamics based on bimodal behavior. We add an index t to represent the population of the output state at time t , e.g., X_t .

Schlögl model

We use the Schlögl model as a prototype SDS that exhibits bimodality, which is a frequently observed behavior in many biological systems. This model uses the following equations (10):



The species A and B are assumed to be in large excess compared to X so their numbers will be held fixed.

The Schlögl model is bistable for the initial condition $X_0 = 250$ at the nominal parameter set in Table 1, as seen in Fig. 1. The figure shows the probability distribution function of X_t at a number of time instances from $t = 1$ to 20 based on computed system response statistics. Each SSA simulation of the system starts at $t = 0$ with $X_0 = 250$. Starting at this initial condition, each simulation follows a random path in time, eventually clustering around one of two stable states. The figure illustrates the evolution of the X_t PDF in time as this process unfolds. At $t = 0$, the initial condition corresponds to a δ -function (not shown) at $X_t = 250$. At $t = 1$, a unimodal distribution is still evident, including the initial condition at 250, but clearly spanning a significant range of values of X_t , roughly in $[100, 500]$. By $t = 2$, a bimodal structure is already apparent, as most realizations begin to cluster around one or the other limit. This bimodal structure is generally well established, and little changed, after $t = 4$. The distribution of X_t exhibits orbits around two stable states, or basins of attraction, one in the vicinity of $X_t = 100$, and the other around $X_t = 550$. In the subsequent discussion, we refer to these states as the lower and upper states, respectively.

TABLE 1 Nominal parameter set for the Schlögl model

Parameter	Value
$k_1 A$	3×10^{-2}
k_2	10^{-4}
$k_3 B$	2×10^2
k_4	3.5




FIGURE 1 PDFs of X_t at $t = 1, 2, 4, 6$, and 20. The Schlögl model has two stable steady states at $t = 20$ for the initial condition $X_0 = 250$. The number of SSA realizations used to generate this PDF is 40,000.

We note, with reference to Fig. 1, the fast growth in the standard deviation of X_t at early time $[0, 2]$. This growth reflects the random spread of the system state away from its initial condition, and its tendency to migrate toward either of the two basins of attraction as time progresses. Of course, at late time, as the distribution of X_t becomes bimodal, the unconditioned mean and standard deviation of the system population have little practical utility. Instead, we investigate the behavior and sensitivity of the conditional means $E(X_t|\cdot)$ and the conditional standard deviations $\sigma(X_t|\cdot)$ of X_t in both the lower and the upper branches, namely

$$\begin{aligned} Y_1 &:= E[X_t | X_t \leq X_0], & Y_2 &:= \sigma[X_t | X_t \leq X_0], \\ Y_3 &:= E[X_t | X_t > X_0], & Y_4 &:= \sigma[X_t | X_t > X_0]. \end{aligned}$$

These conditional standard deviations initially increase in time, reflecting the spread of the realizations away from the initial condition. However, after $t \approx 4$, most realizations have migrated to either the upper or lower branch, and the conditional standard deviations decrease to a value that reflects the spread of the realizations in the vicinity of those branches.

As discussed in the development of the algorithm for the sensitivity coefficients, several parameters need to be chosen to perform the analysis:

1. The number of SSA realizations used in the averaging of the statistical properties.
2. The size of the applied perturbations.
3. The order of the PC expansion used to represent the system response to these perturbations.

The discussion below shows how these parameters were chosen.

The number of SSA realizations used in the calculation of these statistical properties $\{Y_i\}_{i=1}^4$ was chosen based on the level of internal noise in the system. Fig. 2 shows these

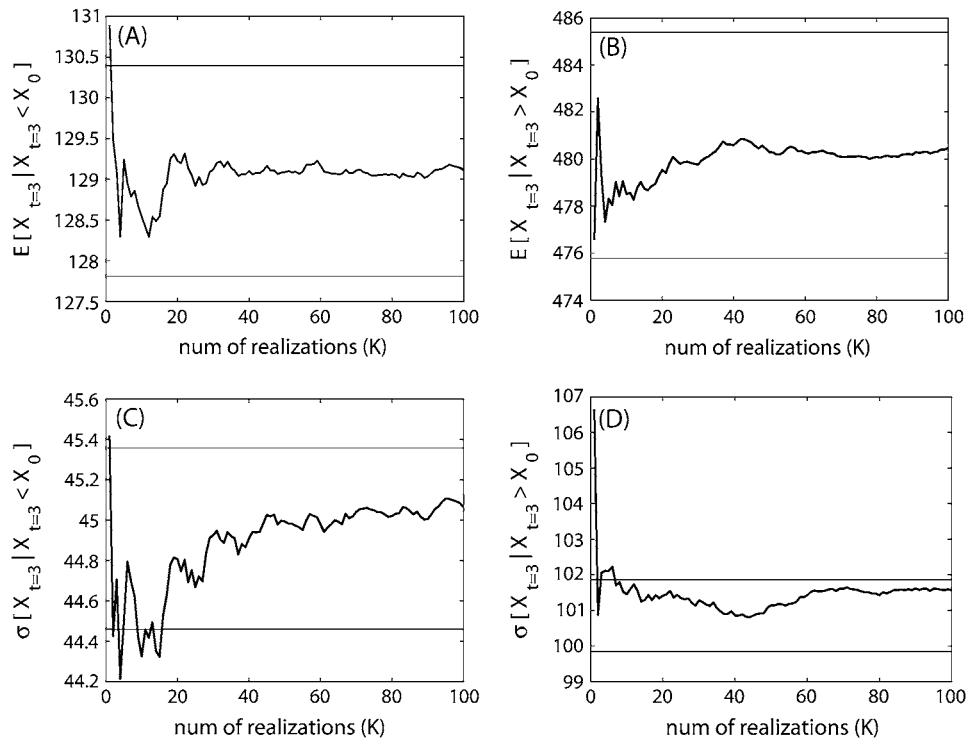


FIGURE 2 Number of realizations for SSA. Statistical properties at $t = 3$ for $X_0 = 250$, as a function of the number of SSA realizations they are averaged over: conditional means (A,B) and conditional standard deviations (C,D). The convergence behavior of these properties is representative of the worst-case scenario over all times. In general, after 40,000 realizations, the effect of the internal stochastic noise on the property estimates is within $\pm 1\%$ of the property value (indicated by the two horizontal lines in each graph).

properties as a function of the number of SSA realizations they were averaged over. The graphs in Fig. 2 were made at $t = 3$, which corresponds to the transient regime where the stochastic fluctuations are the strongest. In general, the stochastic noise is limited to $\pm 1\%$ of the property value after 40,000 SSA realizations. This number of realizations was therefore chosen for the computation of all Schlögl model properties.

As indicated earlier, the size of the applied perturbation should be chosen large enough to accurately capture the variation of the statistical property Y_i due to variations in the model parameters compared to the internal noise level, but small enough to avoid highly nonlinear or discontinuous changes in Y_i . Fig. 3 shows the effect of different sizes of perturbation in the parameter k_1 for the statistical property $Y_4 = \sigma[X_t | X_t > X_0]$ at steady state, $t = 20$. For three different perturbation sizes, this figure shows the values of Y_4 at the five quadrature points that would be needed to project the system response onto a fourth-order PC expansion. The error bars on these values represent the estimated 1% variability left in Y_4 after averaging over 40,000 SSA realizations. As can be seen, a perturbation of 0.3% in k_1 is too small because its effect on Y_4 is overwhelmed by the variability that is due to the inherent internal noise. The nonlinear response in this case comes largely from the internal noise. A 3% perturbation is too large as it causes a discontinuous change in the properties of interest (the standard deviation of the upper branch goes to 0) and is therefore not representative of the local sensitivity of Y_4 . An $\sim 1\%$ perturbation, however, gives a change in Y_4 that is significantly larger than the internal noise, but still only mildly nonlinear so that it can easily be

represented with a PC expansion of a suitable order (see below). This perturbation is therefore appropriate to study the local sensitivity of Y_4 to k_1 at $t = 20$. As the parametric sensitivity in the steady-state regime is of main interest, we checked the proper sizes of perturbation in all the parameters for each property at $t = 20$.

The choice of the order of the PC expansions in the representation of the system response to parametric perturbations is a delicate balance between accuracy and variability. Fig. 4 shows the prediction of the effect of a 3% variation in k_1 on the conditional standard deviation in the upper branch of the solution, Y_4 , using a polynomial chaos approximation that was generated based on a 1% perturbation (σ) in k_1 . Clearly, higher-order polynomial approximations Eq. 10 can adequately capture the local nonlinear behavior over a wider range of the perturbed parameter than lower-order representations over the full time span of the simulation. However, we also observe that a second-order approximation already gives reasonable results.

For the calculation of local sensitivity coefficients, however, we are mainly interested in the accuracy of the representation in the neighborhood of the nominal parameter value, as is studied in Fig. 5, at steady state, $t = 20$. For various PC orders, this figure shows the local behavior of the polynomial response function, which is obtained from the NISP method, for Y_4 , based on a 1% perturbation in k_1 . For each PC order P , a total of $(P + 1)$ quadrature points were used. A first-order response function can obviously not capture any local nonlinear behavior. The second-, third-, and fourth-order representations clearly do show the nonlinear behavior of the system

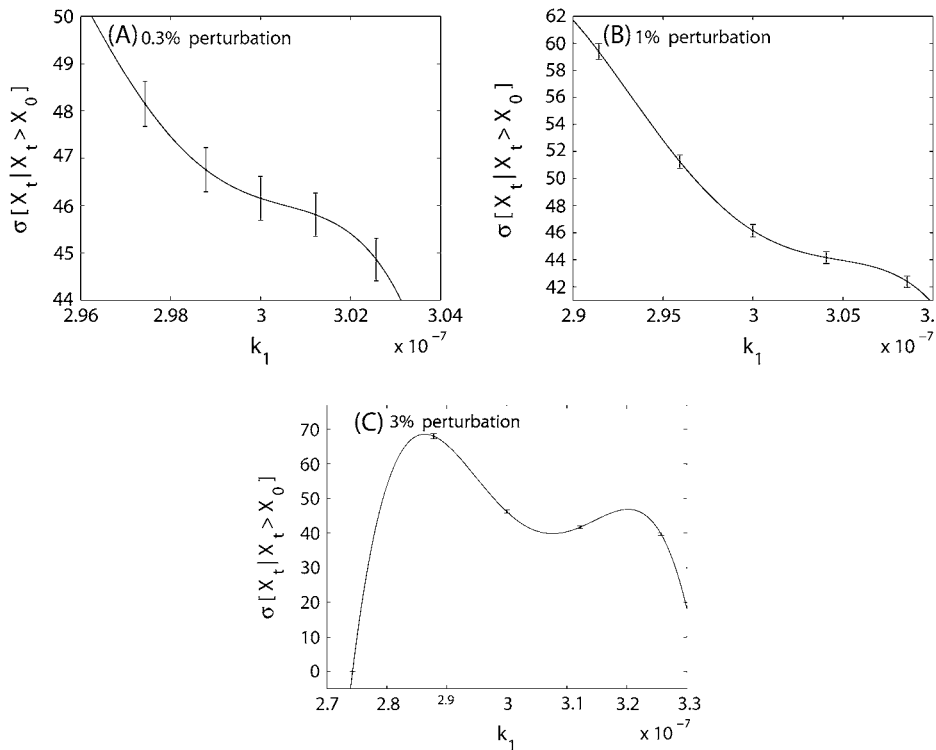


FIGURE 3 Effect of the size of perturbation. Polynomial approximation of $\sigma[X_t | X_t > X_0]$ at steady state, $t = 20$, for $X_0 = 250$ with fourth-order PC expansion as a function of k_1 for three different magnitudes of perturbation. The error bars indicate the $\pm 1\%$ confidence interval on the value of this property at the quadrature points. (B) Appropriate size for sensitivity analysis. The size of perturbation in panel A is too small to generate an appreciable change in the system response compared to the inherent internal noise. The nonlinear behavior comes from the internal noise. The size of the perturbation in panel C is too large to avoid highly nonlinear changes in the properties.

response and qualitatively they all predict the same local behavior in the neighborhood of the nominal values (indicated with an open circle in the graphs). The corresponding sensitivity coefficients S_{41}^P are proportional to the slopes of these representations at the nominal value of k_1 , as indicated with the dashed lines in Fig. 5. Fig. 6 shows those

sensitivity coefficients as a function of the order P of the PC representation and suggests a nonmonotonic convergence with P . To interpret this, it is important to note that these slopes are affected by both the order of the representations as well as the residual internal noise in the values of Y_4 at the quadrature points that are used in the NISP method.




FIGURE 4 Effect of the order of PC expansion on polynomial approximations. Polynomial chaos approximations (solid line) of degrees (A) $P = 1$ through (D) $P = 4$ for $X_0 = 250$ are compared to the true system response (computed directly with SSA) at $+3\%$ perturbed k_1 (dashed line) for the time evolution of the conditional standard deviation in the upper branch. The dotted line is the time evolution at the nominal parameter set. The polynomial approximations were generated based on a 1% perturbation in k_1 .




FIGURE 5 Local nonlinear behavior for the order of PC expansion. Each solid line/curve shows a local nonlinear behavior of the property $Y_4 = \sigma[X_t | X_t > X_0]$ for the order of PC expansion, from 1 to 4, over the range of the parameter k_1 at steady state, $t = 20$, with $X_0 = 250$. The straight dashed line through Y_4 at the nominal value of $k_1 = 3.0 \times 10^{-7}$ in each graph represents the locally linear response of Y_4 in the direction of k_1 . The error bars indicate the $\pm 1\%$ confidence interval on the value of this property at the quadrature points. Higher-order approximations better represent the system response over the full range of the perturbation.

To examine the effect of the internal noise in the values of Y_4 at the quadrature points, we computed 10 independent realizations of the sensitivity coefficient of Y_4 to k_1 , S_{41}^P . Each of these realizations of S_{41}^P was based on a set of 40,000 SSA realizations of the overall system, with a different random number seed at the start of each set. Fig. 7 shows the time evolution of the ensemble average of those 10 realizations of S_{41}^P for both second- and fourth-order PC expansions. The

standard deviation over those 10 realizations of S_{41}^P is shown in the error bars. The change in the computed sensitivity coefficient in going to higher-order PC is roughly one standard deviation; a statistically significant change. Thus, the difference between the second- and fourth-order results is not




FIGURE 6 Convergence with the PC expansion order. The normalized sensitivity coefficient at $t = 20$, of $Y_4 = \sigma[X_t | X_t > X_0]$ with respect to k_1 , S_{41}^P is shown for increasing PC order P from 1 to 4, with a fixed perturbation size of 1%. The graph suggests a nonmonotonic convergence under these conditions.




FIGURE 7 Effect of the order of PC expansion on the time evolution of the normalized sensitivity coefficient. Average time evolution of the normalized sensitivity coefficient of $Y_4 = \sigma[X_t | X_t > X_0]$ with respect to k_1 , S_{41}^P , based on 10 independent realizations of S_{41}^P that were obtained from 10 sets of 40,000 SSA realizations with a different random number seed for each set. The error bars indicate the standard deviations over the 10 realizations. There is a small difference between the results for second-order PC, $P = 2$ (solid line) and fourth-order PC, $P = 4$ (dashed line). The fourth-order PC results, however, show a larger variability than the second-order results.

solely due to the internal noise in the values of Y_4 at the quadrature points. Another noteworthy point is that the standard deviations in the fourth-order results are larger than in the second-order results. It is natural to expect that, for a given noise amplitude, derivatives based on higher-order response functions will exhibit higher levels of noise. Because of the larger variability in the higher-order PC results, the choice of the PC order is a tradeoff between accuracy and variability. In the examples studied in this work, all quantities vary quite smoothly in the neighborhood of their nominal values. This can be seen in Fig. 5, as the behavior near the nominal parameter values is very similar for all PC orders $P \geq 2$. Note also that of all the properties studied, Y_4 showed the most variability.

As a further comparison, we plot the time evolution of the linear normalized sensitivity coefficients based on both second- and fourth-order PC for the statistical properties $\{Y_i\}_{i=1}^4$ relative to each parameter $\{k_j\}_{j=1}^4$ in Figs. 8 and 9, respectively. Based on the reaction mechanism structure, it is expected to find increased levels of X with increasing values of the forward reaction rates k_1 and k_3 , while one would expect inverse dependence of X on k_2 and k_4 . The results for the sensitivity coefficients of the conditional means with respect to the four rate constants are consistent with this expectation in both figures, and for both the lower and upper

branches of the solution. Moreover, for the lower branch (Fig. 8 A), the conditional mean of X exhibits highest negative sensitivity to k_4 throughout the time span of the solution. On the other hand, the conditional mean of X in the upper branch (Fig. 8 C) exhibits highest positive sensitivity to k_1 . In other words, increasing k_4 has the dominant role in decreasing Y_1 , while increasing k_1 has the dominant role in increasing Y_3 . Moreover, while Y_1 is largely insensitive to k_2 , Y_3 is largely insensitive to k_3 . The dependence of the conditional means in the two branches on the four rate constants clearly exhibits a symmetry. The expected dependence of the conditional standard deviations on the four rate constants is not easily available by inspection of the structure of the mechanism. The two figures generally indicate that the conditional standard deviations are more sensitive to the parameters than the conditional means in each branch. The conditional standard deviation in the lower branch, Y_2 , in frame (Fig. 8 B), exhibits similar positive dependence on k_1 and k_3 , nearly negligible dependence on k_2 , and highest (negative) sensitivity to k_4 ; generally consistent with the observed dependences of the conditional mean in this branch, Y_1 , in frame (Fig. 8 A). Increasing k_4 has the dominant role in reducing both the mean of the solution in the lower branch, and the scatter of the realizations around this mean. On the other hand, the dependence of the conditional standard deviation

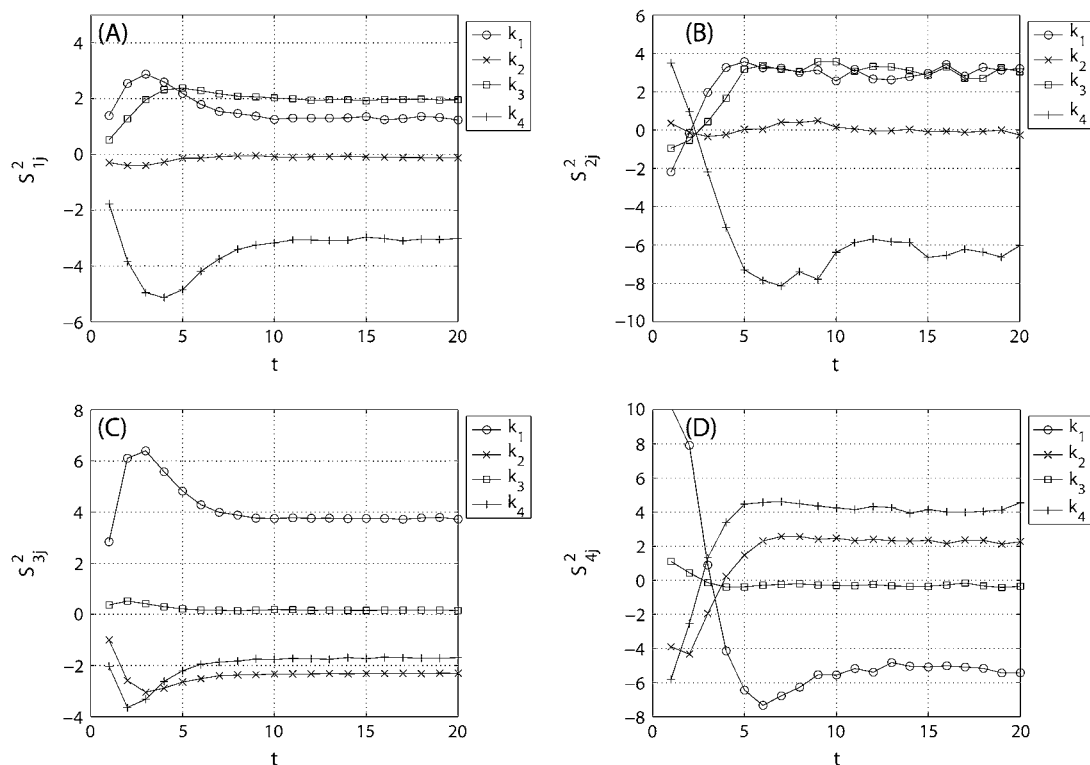


FIGURE 8 Time evolution of the normalized sensitivity coefficients based on second-order PC. The normalized sensitivity coefficients, S^2_{ij} , $i, j = 1, \dots, 4$, based on second-order PC are shown as a function of time until steady state, $t = 20$, in the lower branch (A,B) and the upper branch (C,D). The parameters with respectively the most and the least effect are k_4 and k_2 in the lower branch and k_1 and k_3 in the upper branch. The conditional standard deviations in each branch B and D are in general more sensitive to the parameters than the conditional means in each branch A and C.

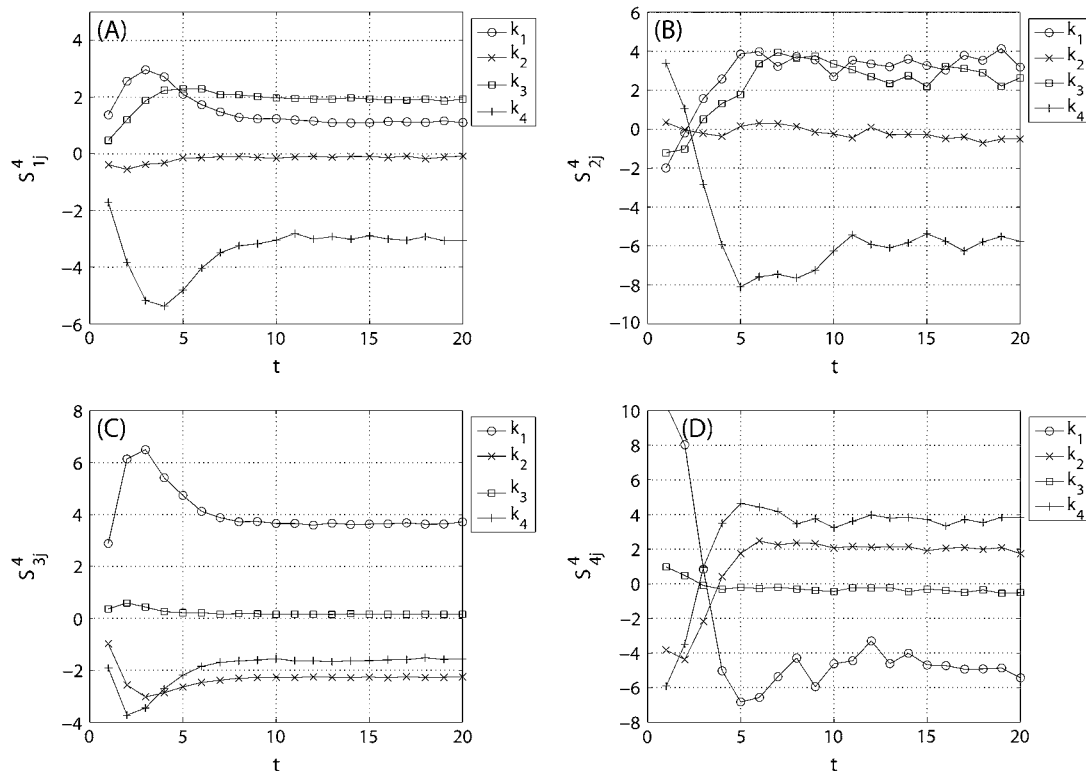


FIGURE 9 Time evolution of the normalized sensitivity coefficients based on fourth-order PC. The normalized sensitivity coefficients, S_{ij}^4 , $i, j = 1, \dots, 4$, based on fourth-order PC are shown in time until steady state, $t = 20$. The results are very similar to the sensitivity coefficients based on second-order PC, shown in Fig. 8.


of the solution in the upper branch, Y_4 , in frame (Fig. 8 D), on the four rate constants is the inverse of the observed dependence of the conditional mean in this branch, Y_3 , in frame (Fig. 8 C). Thus, while increasing k_1 has the dominant role in increasing the conditional mean in this branch, it has the dominant role in decreasing the conditional standard deviation. Evidently, increasing both k_1 and k_4 has the dominant role in moving the conditional means apart, and reducing the overlap of the realizations in the two basins of attraction.

Finally, we note that the Figs. 8 and 9 also show that both second- and fourth-order PC-based parametric sensitivity coefficients are very similar under the same conditions for this system. Thus, despite the challenges with demonstrating convergence in Figs. 5–7 due to the noise in the data and the increased range of the quadrature points, the evident trends in the sensitivity results are largely insensitive to the increase in order from 2 to 4; such that the conclusions are independent of the chosen order for the smoothly behaving observables in this example. For the virus model system below, we will similarly use second-order PC representations to study its behavior.

Virus model

As a more complex and biologically relevant example, we consider a model for intracellular viral kinetics, representative of a generic nonlytic virus, as studied by Srivastava et al.

(1). This model describes the infection of a cell by a virus and considers six reactions between three viral species: template nucleic acid (T), genomic nucleic acid (G), and structural protein (S). The system equations can be written as



where \emptyset represents degradation or secretion and V represents viral progeny. Here, we use the same parameter values as in Srivastava et al. (1), which were determined from the

TABLE 2 Nominal parameter set for the virus model

Parameter	Value
k_1	0.025 day ⁻¹
k_2	0.25 day ⁻¹
k_3	1.0 day ⁻¹
k_4	7.5×10^{-6} molecule ⁻¹ day ⁻¹
k_5	1000 day ⁻¹
k_6	1.99 day ⁻¹




FIGURE 10 PDFs of the virus model at 200 days. PDFs for the low MOI case (*left column*) and the high MOI case (*right column*) at 200 days are shown using 50,000 realizations: T = template, G = genome, S = structural protein. A bimodal population distribution of viral components is observed for the low MOI case, not the high MOI case.

corresponding deterministic system by setting the steady-state values of T , G , and S to 20, 200, and 10,000 molecules, respectively, as listed in Table 2.

In this study, simulations were done for both low and high multiplicity of infection (MOI): as in Srivastava et al. (1), an

infection with one initial molecule of T represents the low MOI case and infection with five molecules of T represents the high MOI case.

We used 50,000 SSA simulations to limit the effect of inherent stochastic noise to $< \pm 1\%$ of the statistical values.

Each SSA simulation represents the infection of an individual cell by the virus. By binning the results of 50,000 runs, PDFs for the numbers of T , G , and S molecules after 200 days for the low and high MOI cases are obtained as shown in Fig. 10. Note that the PDFs for the low MOI case have a bimodal character, where the peak at 0 indicates the probability of a failed infection. In the high MOI case, such failed infections are highly unlikely and the population distributions are unimodal. We are interested in the parametric sensitivity of the probability of such failed infections for the low MOI case as a function of time. This probability is defined as

$$Y_1 := P\{(T, G, S) = (0, 0, 0)\}.$$

In the low MOI case, there is also a subpopulation of low-level infected cells. This subpopulation could act as a viral reservoir, i.e., a group of cells where the infection is latently present for a long time without fully developing. Such a viral reservoir has been suggested as a potential contributing factor to viral persistence (1). In the high MOI case, such a viral reservoir does not appear as most infections typically develop quite rapidly in this case.

To study the parametric sensitivity of this subpopulation, which we will refer to as the viral reservoir in the remainder of this work, we look at the system after the first few days of the infection. This is because most of the infections fail in the first few days, and therefore, most of the cells with a low-level viral population do not represent latent but failing infections at that point in time. After this transient in the failed infections, a cell is considered to be part of the viral reservoir if its level of viral components is no more than 1% of the deterministic steady-state levels, $(T, G, S) = (20, 200, 10,000)$. The probability of a cell being in the viral reservoir then becomes

$$Y_2 := P\{T = 0, G \leq 2, S \leq 100, (T, G, S) \neq (0, 0, 0)\}.$$

We further study the parametric sensitivity of the average virus production rate for both low and high MOI. This production rate is calculated as

$$Y_i := E[k_4 \times G \times S],$$

where $i = 3, 4$ denote the low and high MOI cases, respectively.

To study the parametric sensitivity of the properties $\{Y_i\}_{i=1}^4$ with respect to the six parameters $\{k_j\}_{j=1}^6$, we use second-order PC expansions. We begin by examining the overall response of the system, and commenting on the fidelity of the second-order PC representation, as shown in Fig. 11. The dotted line in Fig. 11 shows the time evolution of the probability of failed infections and of being in the viral reservoir at the nominal parameter set given in Table 2, as computed directly from SSA realizations. For this nominal parameter set, 20.6% of infections with an MOI of 1 failed within 10 days compared to the total 25.1% failed infections over a 200-day time span. As explained above, the probabil-




FIGURE 11 Dynamics of the failed infections and the viral reservoir for the low MOI case. Each curve shows the percentage of the failed infections and the viral reservoir as a function of time. The viral reservoir is computed after the first 10 days. The dotted and dashed lines show the time evolution at the nominal parameter set and at +10% perturbed k_2 , respectively. Polynomial approximations (solid lines), which were generated based on a 5% perturbation in k_2 using second-order PC representation, are compared to the true system data at +10% perturbed k_2 . Both approximations are very good, as the solid and dashed lines are nearly indistinguishable in the graph.

ity of being in the viral reservoir is therefore only computed after the first 10 days. As can be seen in the figure, the overall response of the system statistics $\{Y_1, Y_2\}$ is characterized by a very fast initial transient. After the associated fast timescales are exhausted, slower timescales prevail, governing the subsequent evolution of the system. We next evaluate the fidelity of the second-order PC approximation. To do this,

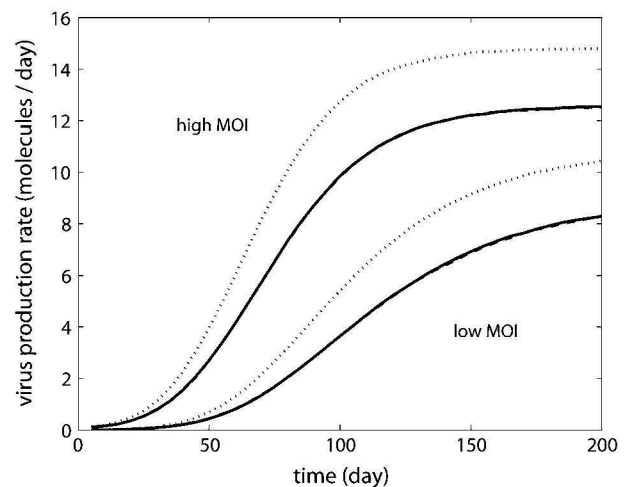


FIGURE 12 Dynamics of the virus production rate. Polynomial approximations (solid line) are compared to the true system data at +10% perturbed k_2 (dashed line) based on a second-order PC representation for the virus production rate in the both low and high MOI cases. The approximation is very good, as the solid and dashed lines are nearly indistinguishable. The dotted lines show the time evolution at the nominal parameter set. The polynomial approximations were generated based on a 5% perturbation in k_2 .

we perturb k_2 by 5%, and construct the corresponding PC expansions for $\{Y_1, Y_2\}$ as in Eq. 10. We then evaluate, based on both direct SSA realizations and the constructed PC expansions, the time evolution of the $\{Y_1, Y_2\}$ for a 10% perturbation in k_2 from the nominal value. The comparison between the PC-based prediction and the true (SSA-based) system response is shown in Fig. 11. As can be seen, the approximations based on second-order PC can accurately reproduce the time evolution of both $\{Y_1, Y_2\}$ for a 10% perturbation in k_2 . Moreover, we note that increasing k_2 , which is the rate of decay of T , is expected to inhibit the rate of progress of the infection by making less T available for the creation of G and S . This can be seen in Fig. 11 since a +10% perturbation in k_2 evidently causes respective increases in the percentages of the failed infections and the viral reservoir of 9.5% and 21.8% on average.

Fig. 12 shows the time evolution of the viral production rate for both the low and high MOI cases, with the rates at the nominal parameter set shown with dotted lines. For these properties as well, a second-order polynomial approximation accurately predicts the effect of a +10% perturbation in k_2 . This perturbation reduces the virus production rate as much as 27.9% for the low MOI case and 22.0% for the high MOI case on average.

The sensitivity coefficients for the probability of failed infections and the viral reservoir in the low MOI case are shown in Fig. 13. Note that because the number of failed infections is near zero at the early time, and because the viral reservoir goes to zero at the late time, the evaluation of the fully normalized sensitivity coefficients is numerically not well conditioned for these properties. In this case, the sensitivities are better represented with the seminormalized sensitivity coefficients, as defined in Eq. 12.

The time evolution of the seminormalized sensitivity coefficients, based on second-order PC, for the probability of failed infections is shown in Fig. 13 A, with a detailed view of the first 20 days in Fig. 13 B. After the initial transient, this probability of failed infections is only sensitive to k_2 and k_3 . Increasing k_2 speeds up the decay of T , which therefore increases the probability of failed infections. Increasing k_3 , on the other hand, reduces the number of failed infections as it catalytically creates more copies of G . In the transient regime, the decay rate of the structural protein S has a strong effect; increasing k_6 increases the number of failed infections.

The seminormalized sensitivity coefficients for the viral reservoir in the low MOI case are shown in Fig. 13 C. An increase in k_2 , the decay rate of T , increases the number of cells in the viral reservoir as it slows down the infection rate

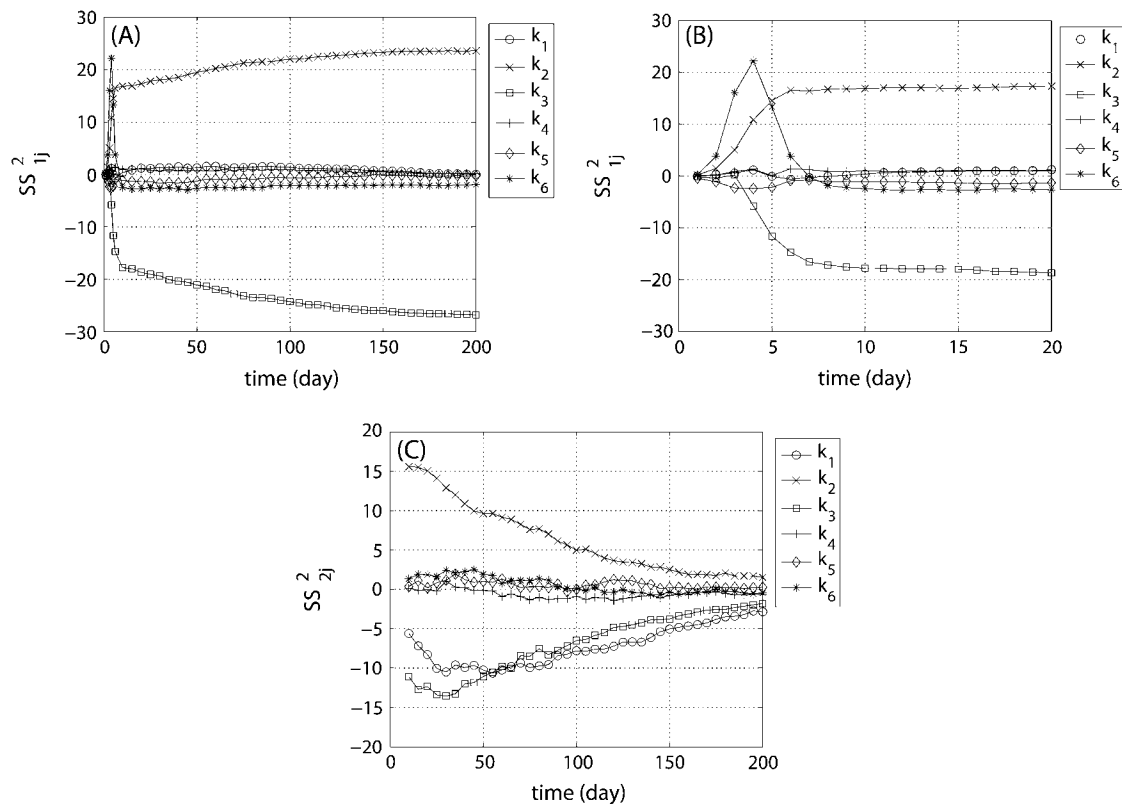


FIGURE 13 Time evolution of the seminormalized sensitivity coefficients based on second-order PC for the failed infections and the viral reservoir in the low MOI case. The sensitivity coefficients for the failed infections are shown in panel A as a function of time, with a detail of the first 20 days in panel B. Panel C shows the seminormalized sensitivity coefficients for the viral reservoir. In general, the parameters k_2 and k_3 have a strong effect on both properties. The reaction rate k_1 has the dominant negative effect on the viral reservoir.

and therefore tends to keep the population of the viral components low. The reactions associated with k_1 and k_3 combine to increase both the G and T population in cells, which speeds up the infection, and the viral reservoir therefore has a strong negative sensitivity to them.

The time evolution of the normalized sensitivity coefficients of the viral production rates is shown in Fig. 14 for both the low and high MOI cases. The viral production rates are most increased by increases in k_1 and k_3 , as those reactions act to increase the G population. They are most negatively affected by increases in k_2 as this reduces the amount of T that is available to form G and S . The effects of k_4 and k_5 are almost identical, and are opposite to the effect of k_6 . All three of these reactions have a different effect in the early time of the infection versus the late time. In the early time, increases in k_4 and k_5 increase the viral production rate as k_5 increases the production of S and k_4 increases the rate of

production of V from G and S . However, this increased viral production in the early time removes more G , resulting in a lower G population in the late time of the infection, with a reduced viral production rate at the late time as a consequence. Similarly, increasing k_6 reduces the viral production rate in the early time as it removes S from the system. However, this leaves more G in the system, resulting in a larger G population and a larger viral production rate in the late time. Note also that the virus production rates in the low MOI case exhibit generally higher sensitivities than in the high MOI case.

CONCLUSIONS

We presented a polynomial chaos (PC)-based method for parametric sensitivity analysis in stochastic dynamical systems governed by the chemical master equation. Observables of the system dynamics are represented as polynomial functions of the model parameters using spectral PC representations. While many stochastic realizations of the system are generally required to evaluate the observables, a quadrature-based sampling scheme is used to efficiently obtain this response function. Sensitivity coefficients were defined using derivatives of these polynomial functions with respect to the model parameters. Higher-order polynomial approximations are able to capture the local nonlinear relationship between the system dynamics and finite-sized perturbations of the model parameters. The resulting sensitivity coefficients indicate the effect as well as the relative importance of each reaction on/to the statistical properties of interest.

Numerical examples show that the PC-based sensitivity analysis is a viable tool for efficiently providing valuable information about the system dynamics. For example, in the virus model, the sensitivity analysis shows very clearly that the probability of failed infections, the viral reservoir, and the viral production rate are all very sensitive to the catalytic generation of the genome nucleic acid (parameter k_3) as well as the rate of template nucleic acid degradation (parameter k_2). These two reactions are therefore the dominant reactions in the model. This methodology can readily be extended to analyze the coupled sensitivity with respect to multiple parameters simultaneously, which is the subject of ongoing work.

The authors thank Dr. Olivier Le Maître and Dr. Roger Ghanem for their helpful discussions.

This work was funded by the U.S. Department of Energy Office of Science, MICS program, under award No. 05-012556-547. Sandia National Laboratories is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the U.S. Department of Energy under contract No. DE-AC04-94-AL85000.

REFERENCES

1. Srivastava, R., L. You, J. Summers, and J. Yin. 2002. Stochastic vs. deterministic modeling of intracellular viral kinetics. *J. Theor. Biol.* 218:309–321.




FIGURE 14 Time evolution of the normalized sensitivity coefficients based on second-order PC for the virus production rate. The normalized sensitivity coefficients S_{ij}^2 for the virus production rate are shown as a function of time for the low MOI case (top) and the high MOI case (bottom). Both cases are most sensitive to the parameters k_2 and k_3 , with the low MOI case being more sensitive to parametric perturbations.

2. Haseltine, E. L., and J. B. Rawlings. 2002. Approximate simulation of coupled fast and slow reactions for stochastic chemical kinetics. *J. Chem. Phys.* 117:6959–6969.
3. Resat, H., H. S. Wiley, and D. A. Dixon. 2001. Probability-weighted dynamic Monte Carlo method for reaction kinetics simulations. *J. Phys. Chem. B.* 105:11026–11034.
4. Schoeberl, B., C. Eichler-Jonsson, E. D. Gilles, and G. Muller. 2002. Computational modeling of the dynamics of the map kinase cascade activated by surface and internalized EGF receptors. *Nat. Biotechnol.* 20:370–375.
5. Goldstein, B., J. R. Faeder, and W. S. Hlavacek. 2004. Mathematical and computational models of immune-receptor signalling. *Nature Rev. Immunol.* 4:445–456.
6. Arkin, A., J. Ross, and H. McAdams. 1998. Stochastic kinetic analysis of developmental pathway bifurcation in phage λ -infected *Escherichia coli* cells. *Genetics.* 149:1633–1648.
7. Allen, R. J., P. B. Warren, and P. R. ten Wolde. 2005. Sampling rare switching events in biochemical networks. *Phys. Rev. Lett.* 94:018104.
8. Makeev, A. G., D. Maroudas, and I. G. Kevrekidis. 2002. coarse stability and bifurcation analysis using stochastic simulators: kinetic Monte Carlo examples. *J. Chem. Phys.* 116:10083–10091.
9. Gillespie, D. 1992. A rigorous derivation of the chemical master equation. *Phys. A.* 188:404–425.
10. Gillespie, D. 1992. Markov Processes: An Introduction for Physical Scientists. Academic Press, San Diego, CA.
11. Gillespie, D. 1977. Exact stochastic simulation of coupled chemical reactions. *J. Phys. Chem.* 81:2340–2361.
12. Gillespie, D. T. 2002. The chemical Langevin and Fokker-Planck equations for the reversible isomerization reaction. *J. Phys. Chem. A.* 106:5063–5071.
13. Gunawan, R., Y. Cao, L. Petzold, and F. J. Doyle. 2005. Sensitivity analysis of discrete stochastic systems. *Biophys. J.* 88:2530–2540.
14. Ghanem, R., and P. Spanos. 1991. Stochastic Finite Elements: A Spectral Approach. Springer Verlag, New York.
15. Le Maître, O., O. Knio, H. Najm, and R. Ghanem. 2001. A stochastic projection method for fluid flow. I. Basic formulation. *J. Comput. Phys.* 173:481–511.
16. Le Maître, O., M. Reagan, H. Najm, R. Ghanem, and O. Knio. 2002. A stochastic projection method for fluid flow. II. Random process. *J. Comput. Phys.* 181:9–44.
17. Debusschere, B., H. Najm, A. Matta, T. Shu, O. Knio, R. Ghanem, and O. Le Maître. 2002. Uncertainty quantification in a reacting electrochemical microchannel flow model. In Proc. 5th Int. Conf. on Modeling and Simulation of Microsystems. 384–387.
18. Debusschere, B., H. Najm, A. Matta, O. Knio, R. Ghanem, and O. Le Maître. 2003. Protein labeling reactions in electrochemical microchannel flow: numerical simulation and uncertainty propagation. *Phys. Fluids.* 15:2238–2250.
19. Reagan, M., H. Najm, R. Ghanem, and O. Knio. 2003. Uncertainty quantification in reacting flow simulations through nonintrusive spectral projection. *Combust. Flame.* 132:545–555.
20. Reagan, M., H. Najm, B. Debusschere, O. L. Maître, O. Knio, and R. Ghanem. 2004. Spectral stochastic uncertainty quantification in chemical systems. *Combust. Theory Model.* 8:607–632.
21. Reagan, M., H. Najm, P. Pébay, O. Knio, and R. Ghanem. 2005. Quantifying uncertainty in chemical systems modeling. *Int. J. Chem. Kin.* 37:368–382.
22. Xiu, D., and G. Karniadakis. 2002. The Wiener-Askey polynomial chaos for stochastic differential equations. *SIAM J. Sci. Comput.* 24: 619–644.
23. Wiener, N. 1938. The homogenous chaos. *Am. J. Math.* 60:897–936.
24. Abramowitz, M., and I. Stegun. 1970. Handbook of Mathematical Functions. Dover, New York.
25. Debusschere, B., H. Najm, P. Pébay, O. Knio, R. Ghanem, and O. Le Maître. 2004. Numerical challenges in the use of polynomial chaos representations for stochastic processes. *SIAM J. Sci. Comput.* 26:698–719.
26. Rathinam, M., L. R. Petzold, Y. Cao, and D. T. Gillespie. 2003. Stiffness in stochastic chemically reacting systems: the implicit τ -leaping method. *J. Chem. Phys.* 119:12784–12794.
27. Gillespie, D. T. 2001. Approximate accelerated stochastic simulation of chemically reacting systems. *J. Chem. Phys.* 115:1716–1733.
28. Cao, Y., D. T. Gillespie, and L. R. Petzold. 2005. The slow-scale stochastic simulation algorithm. *J. Chem. Phys.* 122:014116.
29. Cao, Y., D. Gillespie, and L. Petzold. 2005. Multiscale stochastic simulation algorithm with stochastic partial equilibrium assumption for chemically reacting systems. *J. Comput. Phys.* 206:395–411.
30. Petzold, L. 2006. StochKit: stochastic simulation kit web page. <http://www.engineering.ucsb.edu/~cse/StochKit/index.html>.
31. Petras, K. 2001. Fast calculation of coefficients in the Smolyak algorithm. *Numer. Algorithms.* 26:93–109.