



SCUOLA INTERNAZIONALE SUPERIORE DI STUDI AVANZATI

SISSA Digital Library

Approximation and inference methods for stochastic biochemical kinetics-a tutorial review

Original

Approximation and inference methods for stochastic biochemical kinetics-a tutorial review / Schnoerr, D; Sanguinetti, G; Grima, R. - In: JOURNAL OF PHYSICS. A, MATHEMATICAL AND THEORETICAL. - ISSN 1751-8113. - 50:9(2017), pp. 1-60. [10.1088/1751-8121/aa54d9]

Availability:

This version is available at: 20.500.11767/132451 since: 2023-05-17T16:14:10Z

Publisher:

Published

DOI:10.1088/1751-8121/aa54d9

Terms of use:

Testo definito dall'ateneo relativo alle clausole di concessione d'uso

Publisher copyright

note finali coverpage

(Article begins on next page)

Approximation and inference methods for stochastic biochemical kinetics - a tutorial review

David Schnoerr^{1,2,3}, Guido Sanguinetti^{2,3}, and Ramon Grima^{1,3,*}

¹School of Biological Sciences, University of Edinburgh, Edinburgh EH9 3JH, UK

²School of Informatics, University of Edinburgh, Edinburgh EH8 9AB, UK

³SynthSys, University of Edinburgh, Edinburgh EH9 3JD, UK

*ramon.grima@ed.ac.uk

January 13, 2017

Abstract

Stochastic fluctuations of molecule numbers are ubiquitous in biological systems. Important examples include gene expression and enzymatic processes in living cells. Such systems are typically modelled as chemical reaction networks whose dynamics are governed by the Chemical Master Equation. Despite its simple structure, no analytic solutions to the Chemical Master Equation are known for most systems. Moreover, stochastic simulations are computationally expensive, making systematic analysis and statistical inference a challenging task. Consequently, significant effort has been spent in recent decades on the development of efficient approximation and inference methods. This article gives an introduction to basic modelling concepts as well as an overview of state of the art methods. First, we motivate and introduce deterministic and stochastic methods for modelling chemical networks, and give an overview of simulation and exact solution methods. Next, we discuss several approximation methods, including the chemical Langevin equation, the system size expansion, moment closure approximations, time-scale separation approximations and hybrid methods. We discuss their various properties and review recent advances and remaining challenges for these methods. We present a comparison of several of these methods by means of a numerical case study and highlight some of their respective advantages and disadvantages. Finally, we discuss the problem of inference from experimental data in the Bayesian framework and review recent methods developed the literature. In summary, this review gives a self-contained introduction to modelling, approximations and inference methods for stochastic chemical kinetics.

Contents

1	Introduction	3
2	Stochasticity in biological systems	5
2.1	Emergence of stochasticity	5
2.2	Experimental evidence	7
3	Stochastic chemical kinetics	9
3.1	Chemical reaction networks and deterministic rate equations	9
3.2	Stochastic methods	12
3.2.1	The Chemical Master Equation	12
3.2.2	Moment equations	14
3.2.3	Non-mass-action propensity functions	16
3.3	Stochastic simulations	16
3.3.1	Extrinsic noise	17
3.4	Exact results	18
3.4.1	Finite state space	18
3.4.2	Linear systems	19
3.4.3	Non-linear systems	19
3.4.4	Exact results for special cases	20
4	Approximation methods	22
4.1	The chemical Langevin equation	22
4.1.1	Stochastic simulations	24
4.1.2	Properties and recent developments	24
4.2	The system size expansion	26
4.2.1	Derivation	27
4.2.2	The linear noise approximation	29
4.2.3	Higher order corrections	30
4.2.4	Properties and recent developments	32
4.3	Moment closure approximations	33
4.3.1	General formulation	33
4.3.2	Properties and recent developments	36
4.4	Construction of distributions from moments	38
4.5	Software	39
4.6	Other approximations	40
4.6.1	State space truncation	40
4.6.2	Tau-leaping	40
4.6.3	Time-scale separation	41
4.6.4	Hybrid methods	44
5	Comparison of approximation methods	46
5.1	Enzymatic protein degradation	47
5.2	Bursty protein production	49
5.3	Discussion	51

6 Inference	52
6.1 General problem formulation	53
6.2 The Forward-Backward algorithm	54
6.3 Parameter inference	54
6.4 Computational methods for Bayesian inference in stochastic chemical re- action networks	55
6.4.1 Methods for general networks	55
6.4.2 Inference for gene expression data	56
6.5 Summary	56
7 Conclusions	57

1 Introduction

Understanding the functioning of living cells and biological organisms at the system level has gained increasing attention in recent years and defines a key research programme for the next decades. Experimental techniques are developing at breathtaking speed producing a wealth of data at finer and finer resolutions. However, such experimental data does not by itself reveal the function of such biological systems. The underlying processes typically involve large numbers of interacting components giving rise to highly complex behaviour [1]. Moreover, experimental data are generally corrupted by measurement noise and incomplete, thus posing the mathematical and statistical challenge to infer the relevant biological information from such measurements.

We focus here on mathematical and statistical modelling of chemical reaction networks in biological systems in which *random fluctuations* of molecule numbers play an important role. Recent experiments have shown this to be the case in many biological processes, gene expression being a prominent example [2]. Such random fluctuations or *stochastic effects* in gene expression have been found to lead to dramatically differing behaviours of genetically identical cells. From a modelling point of view, stochastic systems are considerably harder to analyse than their deterministic counterparts.

The *rate equations* give a macroscopic, deterministic description of the dynamics of chemical reaction networks. They consist of a set of ordinary differential equations governing the dynamics of the mean concentrations of the species in a system, thereby ignoring stochastic fluctuations. The rate equations have been successfully applied to various problems and have the advantage of being relatively straightforward to analyse. They typically give an accurate description of systems with large numbers of molecules, which is generally the case in *in vitro* experiments. However, the rate equations do no longer provide a valid description in cases where the effects of stochastic fluctuations become significant. This is typically the case when some species in a system occur at low molecule numbers, a common feature of chemical networks in cells. In this case, the *Chemical Master Equation* (CME) constitutes the accepted probabilistic description of the resulting *stochastic process* [3]. In this framework, the state of a system is given by the molecule numbers of different species, and the CME governs the single-time probability of being in a certain state at a given time.

Despite its simple structure, no analytic solutions to the CME are known for all but the simplest systems. It is however possible to simulate exact sample paths of the underlying stochastic process, by means of the *stochastic simulation algorithm* [4, 5]. The latter allows us to draw exact samples from the process. However, since the stochastic simulation algorithm simulates each and every chemical reaction event in a system, it is computationally expensive and quickly becomes infeasible for larger systems. Accordingly, significant effort has been spent in recent decades on the development of approximation methods of the CME, and a large variety of different methods has emerged. Similarly, due to the growing amount of experimental data becoming available, a considerable amount of work has been devoted to the development of statistical inference methods for such data, i.e., methods that allow to calibrate a model to observational data.

There exists a significant amount of literature on modelling of stochastic chemical kinetics. However, most reviews are mainly concerned with simulation based methods [6, 7, 8, 9, 10, 11, 12]. Others are rather technical and require a high level of pre-existing and/or mathematical knowledge [13, 14, 15, 16, 17]. None of these references gives a thorough introduction into modelling of stochastic chemical kinetics in the biological context that is accessible for non-experts, neither do they give an overview or comparison of state of the art approximation and inference methods.

With this article, we aim at filling these gaps in several ways:

1. We give a self-contained introduction to deterministic and stochastic modelling techniques for chemical reaction networks. First, we introduce and give a historic motivation to the deterministic rate equations. Next, we derive and give a detailed discussion of the CME. We review exact solution methods for the CME as well as simulation methods.
2. We give a detailed derivation and discussion of the following approximation methods of the CME: the chemical Langevin equation and its associated chemical Fokker-Planck equation; the system size expansion; moment closure approximations. Moreover, we give an introduction and overview of other types of approximations, including time-scale separation based methods and hybrid approximations.
3. We perform a numerical case study comparing the various approximation methods mentioned before.
4. We give an introduction to inference methods in the Bayesian framework and review existing methods from the literature.

The presentation is written to be accessible for non-experts that are new to the field of stochastic modelling.

Even though this review is motivated by stochastic effects in systems biology, it is important to stress that many systems in other scientific fields are frequently modelled by means of Master Equations. The methods discussed in this article can therefore readily be applied to such systems. Examples include ecology [18, 19, 20], epidemiology [21, 22, 23, 24, 25, 26], social sciences [27, 28, 29] and neuroscience [30, 31, 32, 33, 34, 35, 36].

This article is structured as follows. We start by discussing the importance of stochasticity in biological systems in Section 2. We describe the underlying mechanisms giving rise to stochastic fluctuations in cells (Section 2.1) and discuss key experimental studies that measured such fluctuations (Section 2.2). Next in Section 3 we discuss deterministic and stochastic modelling methods for chemical kinetics. We start by introducing the concept of chemical reaction networks and deterministic descriptions in terms of macroscopic rate equations in Section 3.1. Next, we introduce stochastic modelling techniques in terms of the CME and stochastic simulation algorithms in Sections 3.2 and 3.3, respectively. We discuss analytic solution methods for certain classes of reaction systems in Section 3.4. Section 4 is devoted to approximation methods of the CME. We give detailed introductions and discussions of the chemical Langevin equation (Section 4.1), the system size expansion (Section 4.2) and moment closure approximations (Section 4.3). Next, we discuss how approximate moment values obtained from these methods can be used to construct probability distributions by means of the maximum entropy principle in Section 4.4. In Section 4.5 we review existing software packages implementing the approximation methods discussed in Sections 4.1-4.3. We give an introduction to other approximation methods in Section 4.6, including the finite state projection algorithm, time-scale separation based approximations and hybrid methods. In Section 5 we perform a numerical case study and compare the chemical Langevin equation, the system size expansion and moment closure approximations. Section 6 deals with the problem of inference for CME type systems from observational data. We introduce the Bayesian approach to this problem and review existing methods from the literature. Finally, we summarise and conclude in Section 7.

2 Stochasticity in biological systems

Stochastic effects play an important role in many chemical reaction networks in living cells. Examples are enzymatic/catalytic processes, transduction of external signals to the interior of cells or the process of gene expression, to name just a few. Here we discuss different sources of stochasticity and illustrate their emergence in biochemical networks using gene expression as an example in Section 2.1. Subsequently in Section 2.2 we explain how the different sources of stochasticity can be measured and distinguished experimentally and highlight the importance of stochasticity for living cells.

2.1 Emergence of stochasticity

The term “gene expression” denotes the process of synthesis of functional gene products such as proteins. The mechanism is illustrated in Figure 1 for prokaryotic cells, namely cells lacking in-membrane-bound organelles (e.g. nucleus, mitochondria), such as bacteria. The process includes two main steps: *transcription* during which mRNA molecules are produced, and *translation* during which protein molecules are synthesised [37]. The transcription process involves the enzyme RNA polymerase. For the mechanism to initiate, an RNA polymerase enzyme must bind to the initial sequence of a gene. It then slides along the gene and produces an mRNA molecule that reproduces the DNA code of the gene. The RNA polymerase molecules move around randomly in the cell, a process

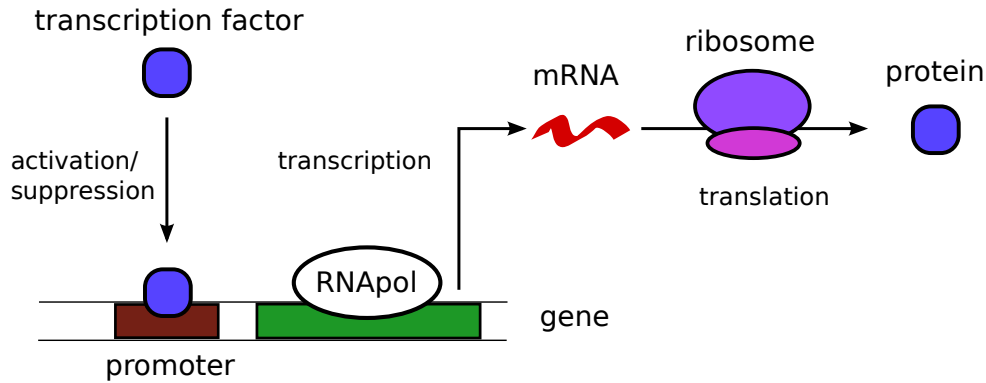


Figure 1: Illustration of gene expression in prokaryotic cells. Transcription is conducted by a RNA polymerase enzyme (RNAPol) that binds to the gene and produces an mRNA molecule from the gene’s DNA. The mRNA is then translated into proteins by ribosomes. Transcription factors may bind to the promoter of the gene and thereby influence the recruitment of RNA polymerase and hence the transcriptional rate of the gene. Gene expression in eukaryotes happens similarly but is compartmentalised in the nucleus (transcription) and the cytosol (translation).

which can be approximately described as Brownian motion [15]. This means that the RNA polymerase binding to the gene is a stochastic event that happens randomly in time. As it turns out, not only the binding of the RNA polymerase to the gene, but also the sliding along the gene happens stochastically. Therefore, the production of mRNA molecules is a stochastic process.

The production of protein molecules from mRNA during translation is conducted by ribosomes, which are RNA and protein complexes. The ribosomes and mRNA diffuse in the cell and hence meet randomly before translation can occur. Translation is thus also a stochastic process. Similarly, the degradation of mRNA molecules and proteins is conducted by certain enzymes and hence happens stochastically.

The transcriptional regulation of gene expression is frequently modulated by regulatory proteins named *transcription factors*. Transcription factors are gene specific and bind to the *promoter region*, which is located upstream of the gene’s encoding region. Upon binding, transcription factors tune the affinity of the RNA polymerase molecules for the promoter, thereby modulating the rate at which transcription initiation events occur and hence the overall transcription rate. A transcription factor can either increase or decrease the binding rate of RNA polymerase and thus either enhance or suppress gene expression.

In eukaryotes, i.e., cells possessing in-membrane-bound organelles such as a nucleus, gene expression happens in a similar but somewhat more complicated way involving more steps. For instance, genes are generally located in the nucleus and transcription hence happens in the nucleus. Processes such as mRNA maturation and diffusion of mRNA to the cytoplasm need to happen before the mature mRNA can be translated into protein.

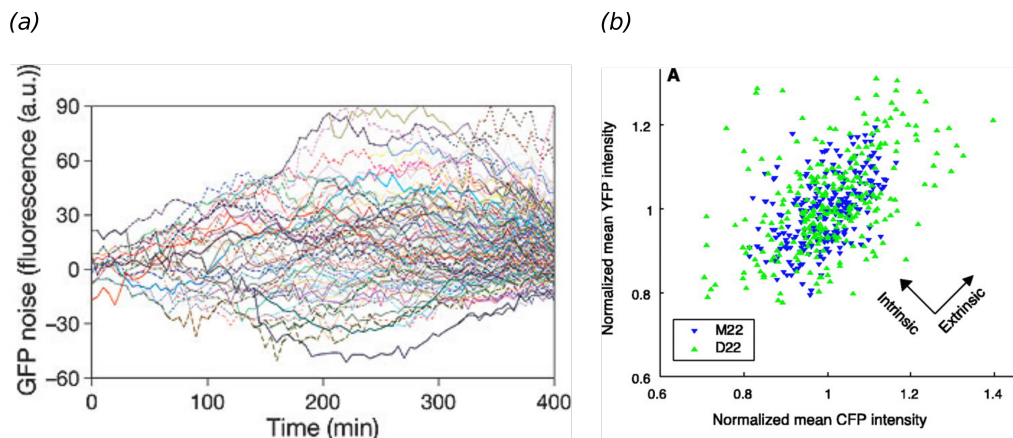


Figure 2: Measurements of intrinsic and extrinsic noise. (a) Fluorescent time series acquired on *Escherichia coli* cells expressing a green fluorescent protein (GFP) encoding gene [39]. The time trajectories correspond to the difference of the measured intensities in individual cells and the population average. We observe that the protein numbers fluctuate strongly over time [39]. Permission kindly granted by the publisher. (b) Fluorescence intensities for the dual reporter technique measured in [2] for the two *Escherichia coli* strains M22 and D22. Permission kindly granted by the publisher. Each triangle corresponds to the measured fluorescence intensity in one cell. Variations along and perpendicular to the diagonal correspond to extrinsic and intrinsic noise, respectively. We observe that in this system both intrinsic and extrinsic noise contribute significantly to the overall fluctuations.

The inherent stochasticity of chemical processes leads to fluctuations of molecule numbers in time. As an example Figure 2 (a) shows fluorescence intensity measurements of a fluorescent protein in individual *Escherichia coli* cells, indicating strong fluctuations of protein numbers in time. Temporal fluctuations that are due to the stochasticity of chemical processes is what we call *intrinsic noise* [2]. Differences in molecule numbers of a certain species, say proteins, between different cells can originate from this type of fluctuations. However, such differences can also stem from other effects, such as physiological differences between cells or differing environmental conditions. For example, the numbers of RNA polymerase or ribosomes may differ between different cells, or different cells may be exposed to varying nutrient concentrations due to environmental fluctuations. Such differences that are not due to the stochasticity of chemical reactions are referred to as *extrinsic noise* [38].

2.2 Experimental evidence

As explained above, stochastic fluctuations are inherent to biochemical processes such as gene expression. The question hence arises: what effects do these fluctuations have on the functioning of cells and which processes dominate the emergence of these fluctuations?

To answer such questions it is crucial to be able to empirically distinguish and quantify intrinsic and extrinsic noise.

One of the first experiments that aimed at explicitly separating the effects of intrinsic and extrinsic noise on gene expression was conducted by Elowitz et al. in 2002 [2] by means of the *dual reporter technique*. In this study, the authors integrated two distinguishable but identically regulated fluorescent reporter genes into the chromosome of *Escherichia coli* cells, one expressing cyan fluorescent protein (CFP) and one yellow fluorescent protein (YFP). The two genes were controlled by identical promoters, which means that they experienced the same external effects, i.e., the same extrinsic noise. Assuming that the two promoters are independent, fluctuations of the two proteins originating from the stochasticity of chemical reactions, i.e., intrinsic noise, should be completely uncorrelated. On the other hand, extrinsic effects should influence both expression systems equally, and the corresponding fluctuations should therefore be strongly correlated. This can be visualised as in Figure 2 (b) which is taken from [2]. The figure shows the measured intensities of individual cells for two different strains, with the YFP intensity on the y -axis and the CFP on the x -axis. The width of the cloud along the diagonal corresponds to correlated fluctuations and hence extrinsic noise, while the width orthogonal to the diagonal corresponds to uncorrelated fluctuations, i.e., intrinsic noise. The figure indicates that for this particular system both intrinsic and extrinsic noise significantly contribute to the overall noise.

Many other experimental studies have been conducted investigating the origins and role of intrinsic and extrinsic noise in cells (for example [40, 41, 42, 43, 39, 44]). For an overview of experimental techniques see [45, 46]. In [43], for example, the authors measured the expression levels of 43 genes in *Escherichia coli* under various different environmental conditions. They observed that the variance of protein numbers scales roughly like the mean protein numbers and that for intermediate abundance proteins the intrinsic fluctuations are comparable or larger than extrinsic fluctuations. In another large scale study on *Escherichia coli* in [44], the authors found that fluctuations are dominated by extrinsic noise at high expression levels.

The reported fluctuations can have significant influence on the functional behaviour of cells [1]. A particularly important example are *stochastic cell fate decisions* [47], where genetically identical cells under the same environmental conditions stochastically differentiate into functionally different states. Such differing cell fate decisions of genetically identical cells are believed to be beneficial for a population of cells experiencing fluctuating environments [47].

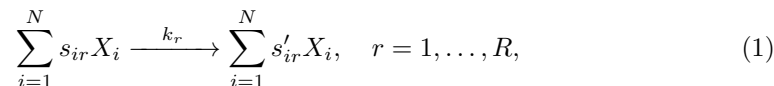
Most of the mentioned experimental studies measure certain statistics of gene expression systems, such as the equilibrium mean and variance of protein numbers. While experimental methods such as the *dual reporter technique* allow to estimate the magnitude of intrinsic and extrinsic noise, they do not explain where the stochasticity originates from and what the underlying processes are. In order to answer such questions, mathematical methods are needed that allow researchers to model the dynamics and stochasticity of such biochemical reaction systems.

3 Stochastic chemical kinetics

After having given a biological motivation for the importance of stochasticity in chemical reaction networks in the previous section, we consider next the question of how such systems can be described mathematically. First, we introduce the concept of chemical reaction networks in Section 3.1 and describe the classical mathematical description of such systems in terms of deterministic rate equations. Next, we discuss the stochastic description of such systems in Section 3.2. We derive the Chemical Master Equation and discuss its validity and solutions. Next, we discuss the stochastic simulation algorithm in Section 3.3. Finally, we review methods for exact analytic solutions of the CME for certain classes of systems in Section 3.4 and give an overview of available analytic solutions of specific systems that do not fall under any of these categories.

3.1 Chemical reaction networks and deterministic rate equations

Biological processes such as gene expression generally consist of complicated mechanisms involving several different types of molecules and physical operations. For a mathematical description of certain processes, one typically does not model all these mechanisms explicitly, but rather replaces them by an effective single chemical reaction event. In the context of gene expression, for example, transcription or translation may be modelled as single chemical reactions. A finite set of chemical species that interact via a set of such chemical reactions constitutes what we call a *chemical reaction network*. Given a set of chemical species X_i , $i = 1, \dots, N$, we define R chemical reactions by the notation



where the *stoichiometric coefficients* s_{ir} and s'_{ir} are non-negative integer numbers denoting numbers of reactant and product molecules, respectively. We say that the r^{th} reaction is “of order m ” if $\sum_{i=1}^N s_{ir} = m$, i.e., if it involves m reactant molecules. We further call a reaction “unimolecular” if $m = 1$, “bimolecular” if $m = 2$ and a system “linear” if $m \leq 1$ for all reactions occurring in the system. We further call a system “open” if it contains a chemical process that generates molecules, and “closed” otherwise. The quantity k_r in Equation (1) is called the *reaction rate constant* of the r^{th} reaction.

Classically, the dynamics of a chemical reaction system as in Equation (1) has been modelled by *the Law of Mass Action* which was developed by Guldberg and Waage in its first version in the 1860s in the context of macroscopic *in vitro* experiments, i.e., macroscopic amounts of chemical molecules in solutions [48, 49, 50, 51]. The Law of Mass Action states that the rate of a reaction is proportional to the product of the concentrations of reactant molecules. Specifically, if we define $\phi = (\phi_1, \dots, \phi_N)$, where ϕ_i denotes the concentration of species X_i , the rate g_r of a reaction as in Equation (1) is given by

$$g_r(\phi) = k_r \prod_{i=1}^N \phi_i^{s_{ir}}. \quad (2)$$




Figure 3: Illustration of a gene expression system with negative feedback loop. When the gene is in the “on” state it produces proteins of a certain type. The protein can decay or bind to the gene’s promoter. In the bound state the gene is in the “off” state, i.e., no protein is produced. The corresponding reactions are given in Equation (5). The protein hence suppresses its own production.

We call g_r the *macroscopic rate function* of the r^{th} reaction. Let us further define the *stoichiometric matrix* \mathbf{S} as

$$S_{ir} = s'_{ir} - s_{ir}, \quad i = 1, \dots, N, \quad r = 1, \dots, R. \quad (3)$$

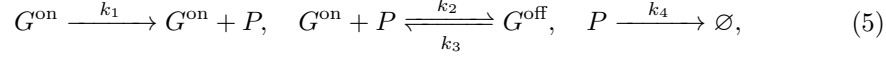
The entry S_{ir} corresponds to the net change of X_i molecules when the r^{th} reaction occurs. Consequently, $S_{ir}g_r(\phi)$ is the rate at which the concentration ϕ_i of species X_i changes due to the r^{th} reaction. Summing up the contributions from all reactions in a given system one obtains [52]

$$\frac{d}{dt}\phi_i = \sum_{r=1}^R S_{ir}g_r(\phi), \quad i = 1, \dots, N. \quad (4)$$

This set of ordinary differential equations is called the *rate equations*. The Law of Mass Action and the associated rate equations assume continuous concentrations, which means that they ignore the discreteness of molecule numbers. Moreover, they constitute a deterministic method, i.e., for a fixed initial condition the state of the system is exactly determined for all times. However, in the previous section we have seen that chemical reactions occur stochastically leading to fluctuations of molecule numbers in time. The reason why the rate equations give an accurate description in experiments such as the ones of Guldberg and Waage is that they studied chemical systems in solutions which typically contain a large number of substrate molecules. For large molecule numbers, it has been found experimentally that the relative fluctuations, i.e., the standard deviation divided by the mean value of molecule concentrations, scales like the inverse square root of the mean concentration and hence becomes small for systems with large molecule numbers [43]. One can therefore expect the rate equations to be an accurate description whenever a system has large molecule numbers for all species.

Example. As an example, consider the gene system in Figure 3. This system can be viewed as a simplified version of the gene expression process described in Figure 1: we replace the process of transcription (gene produces an mRNA molecule) and translation (mRNA molecule produces a protein) by a single effective reaction in which the gene directly produces the protein.

The example in Figure 3 depicts negative autoregulation. The latter is one of the most common regulatory mechanisms: the gene's product downregulates its own expression. In the bound state G^{off} , the gene does not produce any protein. The protein thus suppresses its own production, which means the system is an example of a *negative feedback loop*. The corresponding reactions in the notation of Equation (1) read



where we call the gene in the bound and unbound state G^{on} and G^{off} , respectively and the protein P . In our nomenclature, $G^{\text{on}} + P \rightarrow G^{\text{off}}$ is a second-order or bimolecular reaction, while the other three reactions are first order or linear. By " $P \rightarrow \emptyset$ " we indicate that P leaves the system under consideration. This could for example mean that P becomes converted into different types of chemical species or degraded in its constitutive elements that are not included in our model.

Let us order reactions according to the rate constants in Equation (5), and species as G^{on} , P and G^{off} , respectively. Consider the stoichiometric coefficients s_{ir} and s'_{ir} defined in Equation (1), which correspond to the number of reactant and product molecules, respectively, of species X_i in the r^{th} reaction. For the reaction $G^{\text{on}} \rightarrow G^{\text{on}} + P$, for example, we have $s_{11} = 1, s_{21} = 0$ and $s_{31} = 0$, since there is one G^{on} on the l.h.s. of the reaction but no P or G^{off} , and $s'_{11} = 1, s'_{21} = 1$ and $s'_{31} = 0$ since there is one G^{on} , one P and no G^{off} molecules on the r.h.s. Proceeding similarly for the other reactions in Equation (5), we find

$$\mathbf{s} = \begin{pmatrix} 1 & 1 & 0 & 0 \\ 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}, \quad \mathbf{s}' = \begin{pmatrix} 1 & 0 & 1 & 0 \\ 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \end{pmatrix}. \quad (6)$$

Accordingly, the stoichiometric matrix $\mathbf{S} = \mathbf{s}' - \mathbf{s}$ reads

$$\mathbf{S} = \begin{pmatrix} 0 & -1 & 1 & 0 \\ 1 & -1 & 1 & -1 \\ 0 & 1 & -1 & 0 \end{pmatrix}. \quad (7)$$

Let $\boldsymbol{\phi} = (\phi_1, \phi_2, \phi_3)$, where ϕ_1, ϕ_2 and ϕ_3 denote the concentrations of G^{on}, P and G^{off} , respectively. The macroscopic rate vector $\mathbf{g}^{(0)}(\boldsymbol{\phi}) = (g_1^{(0)}(\boldsymbol{\phi}), \dots, g_r(\boldsymbol{\phi}))^T$, with the g_r defined in Equation (2), is obtained by using \mathbf{s} in Equation (6) and reads

$$\mathbf{g}(\boldsymbol{\phi}) = (k_1\phi_1, k_2\phi_1\phi_2, k_3\phi_3, k_4\phi_2)^T. \quad (8)$$

Using Equations (4), (7) and (8) it is easy to write down the corresponding rate equations. However, note that the system has a conservation law in particle numbers which we can use to find a simplified description by reducing the number of variables: the total number of genes in the "on" state and genes in the "off" state is constant, i.e., $\phi_1 + \phi_3 = \text{const.} \equiv g_0$. We can thus reduce the system to a two species system by using $\phi_3 = g_0 - \phi_1$. The matrices \mathbf{s} and \mathbf{s}' for the reduced system are obtained from Equation (6) by dropping the last row,

$$\mathbf{s} = \begin{pmatrix} 1 & 1 & 0 & 0 \\ 0 & 1 & 0 & 1 \end{pmatrix}, \quad \mathbf{s}' = \begin{pmatrix} 1 & 0 & 1 & 0 \\ 1 & 0 & 1 & 0 \end{pmatrix}, \quad (9)$$

and the stoichiometric matrix and propensity vector of the reduced system read accordingly

$$\mathbf{S} = \begin{pmatrix} 0 & -1 & 1 & 0 \\ 1 & -1 & 1 & -1 \end{pmatrix}, \quad \mathbf{g}(\phi) = (k_1\phi_1, k_2\phi_1\phi_2, k_3(g_0 - \phi_1), k_4\phi_2)^T. \quad (10)$$

Using Equation (4) we hence obtain the rate equations

$$\begin{aligned} \frac{\partial}{\partial t}\phi_1 &= -k_2\phi_1\phi_2 + k_3(g_0 - \phi_1), \\ \frac{\partial}{\partial t}\phi_2 &= k_1\phi_1 - k_2\phi_1\phi_2 + k_3(g_0 - \phi_1) - k_4\phi_2. \end{aligned} \quad (11)$$

The gene system in Figure 3 will be used to showcase several different methods in this article, and we will throughout use the reduced system.

3.2 Stochastic methods

The rate equations discussed in the previous section constitute a deterministic description of chemical reaction systems in terms of continuous molecule concentrations. As we have seen in Section 2, however, fluctuations become important in many biological processes such as gene expression. The reason for this is that often some species occur at low molecule counts. A more general, stochastic description would keep the discrete nature of the X_i molecules in Equation (1) intact. One general approach is to explicitly model the spatial positions of molecules and to model their movement as Brownian motion, with chemical reactions happening stochastically under certain rules. Exact analytic results for such systems are generally not known. Simulations, on the other hand, have to keep track of every single particle and quickly become computationally unfeasible. However, under certain conditions, which we discuss in the following section, simplified descriptions can be employed making spatial models and the simulation of single particles unnecessary.

3.2.1 The Chemical Master Equation

We now consider a chemical reaction system as in Equation (1) in a closed compartment of volume Ω . We seek a stochastic description of the system under well-mixed and dilute conditions. By “well-mixed” we mean that the diffusion of particles in the compartment constitutes the fastest time scale of the system, in the sense that the expected distance travelled by each particle between successive reactive collisions is much larger than the length scale of the compartment. This implies that the spatial positions of molecules can be ignored and the dynamics of the system only depends on the total molecule numbers. By “dilute” we mean that the combined volume of all the considered molecules is much smaller than the total volume, which means that the molecules can be considered as point particles.

If these two conditions are met, it can be shown [3] that the state of the system at any time is fully determined by the state vector $\mathbf{n} = (n_1, \dots, n_N)$, where n_i is the molecule number of species X_i in the compartment. In particular, the spatial locations

and diffusion of molecules does not have to be modelled, and the system corresponds to a *continuous-time Markov jump process*. It can further be shown that the probability for the r^{th} reaction to happen in an infinitesimal time step dt is given by $f_r(\mathbf{n})dt$ where $f_r(\mathbf{n})$ is the *propensity function* of the r^{th} reaction and proportional to the number of combinations of reactant molecules in $\mathbf{n} = (n_1, \dots, n_N)$. Consider for example a bimolecular reaction of the form $A + B \rightarrow \emptyset$. The number of pairs with one A and one B molecule is $n_A n_B$, where n_A and n_B are the molecule numbers of A and B , respectively. The corresponding propensity function is hence given by $k_r n_A n_B / \Omega$. The scaling with the volume Ω stems from the fact that the probability for two molecules to collide is proportional to $1/\Omega$. Generalising these arguments to reactions as in Equation (1) leads to

$$f_r(\mathbf{n}) = k_r \Omega \prod_{i=1}^N \frac{n_i!}{(n_i - s_{ir})! \Omega^{s_{ir}}}. \quad (12)$$

Propensity functions of this form are called *mass-action kinetics type* [14]. Although Equation (2) is often, due to historical reasons, stated to be the Law of Mass Action, from a microscopic perspective it is more accurate to state Equation (12) as the Law of Mass Action. Equation (2) can be viewed as the macroscopic version of (12) obtained in the limit of large molecule numbers and small fluctuations. Specifically, for reactions of up to order two Equation (12) becomes

$$\begin{array}{lll} \text{zeroth order} & \emptyset \rightarrow & f_r(\mathbf{n}) = k_r \Omega, \\ \text{first order} & A \rightarrow & f_r(\mathbf{n}) = k_r n_A, \\ \text{second order} & A + B \rightarrow & f_r(\mathbf{n}) = \frac{k_r}{\Omega} n_A n_B, \\ \text{second order} & A + A \rightarrow & f_r(\mathbf{n}) = \frac{k_r}{\Omega} n_A (n_A - 1). \end{array} \quad (13)$$

Let us now consider how we can mathematically describe the dynamics of such a system. To this end, consider the probability distribution $P(\mathbf{n}, t | \mathbf{n}_0, t_0)$ for the system to be in state \mathbf{n} at time t given that it was in state \mathbf{n}_0 at time t_0 . We will use the shorthand $P(\mathbf{n}, t) = P(\mathbf{n}, t | \mathbf{n}_0, t_0)$ in the following and implicitly assume conditioning on an initial state. The probability $P(\mathbf{n}, t + dt)$ after an infinitesimal time step dt is given by $P(\mathbf{n}, t)$ plus the probability to transition into state \mathbf{n} from a different state \mathbf{n}^* minus the probability to leave state \mathbf{n} , which leads us to

$$P(\mathbf{n}, t + dt) = P(\mathbf{n}, t) + dt \left(\sum_{r=1}^R f_r(\mathbf{n} - \mathbf{S}_r) P(\mathbf{n} - \mathbf{S}_r, t) - \sum_{r=1}^R f_r(\mathbf{n}) P(\mathbf{n}, t) \right), \quad (14)$$

where we used the fact that probability of the r^{th} reaction to happen in an infinitesimal time interval dt is given by $f_r(\mathbf{n})dt$ and where \mathbf{S}_r is the r^{th} column of the stoichiometric matrix \mathbf{S} . Subtracting $P(\mathbf{n}, t)$, dividing by dt and taking the limit $dt \rightarrow 0$ gives the *Chemical Master Equation* (CME) [3],

$$\partial_t P(\mathbf{n}, t) = \sum_{r=1}^R f_r(\mathbf{n} - \mathbf{S}_r) P(\mathbf{n} - \mathbf{S}_r, t) - \sum_{r=1}^R f_r(\mathbf{n}) P(\mathbf{n}, t). \quad (15)$$

Since \mathbf{n} is a discrete-valued vector, Equation (15) is a coupled system of linear ordinary differential equations. Note that this system is infinite whenever \mathbf{n} is unbounded. Despite

its simple structure, there are generally no analytic solutions known to the CME. We call a distribution $P(\mathbf{n}, t)$ a *steady-state solution* of the CME if it solves the CME and fulfils $\partial_t P(\mathbf{n}, t) = 0$.

In the context of chemical reactions, one of the first applications of the CME was done by Delbrück in 1940 for describing the dynamics of an autocatalytic reaction [53]. It has later been applied to linear reactions by Bartholomay [54] and McQuarrie [55], and bimolecular reactions by Ishida [56] and McQuarrie et al [57]. Gillespie derived the CME from molecular physics of a dilute and well-mixed gas in 1992 [3], which he extended to liquids in 2009 [58].

Example. Consider again the gene system in Figure 3 with reactions given in Equation (5). The corresponding stoichiometric matrix \mathbf{S} is given in Equation (10). Let n_1 and n_2 be the molecule numbers of the gene in the unbound state G^{on} and the protein P , respectively. Let us assume that there is only one gene in the system, which implies that the number of bound genes G^{off} is $1 - n_1$. Using Equation (12) we find the propensity vector

$$\mathbf{f}(\mathbf{n}) = (k_1 n_1, \frac{k_2}{\Omega} n_1 n_2, k_3(1 - n_1), k_4 n_2)^T. \quad (16)$$

The corresponding CME becomes (c.f. Equation (15))

$$\begin{aligned} \partial_t P(n_1, n_2, t) &= k_1 n_1 P(n_1, n_2 - 1, t) + \frac{k_2}{\Omega} (n_1 + 1)(n_2 + 1) P(n_1 + 1, n_2 + 1, t) \\ &\quad + k_3(2 - n_1) P(n_1 - 1, n_2 - 1, t) + k_4(n_2 + 1) P(n_1, n_2 + 1, t) \quad (17) \\ &\quad - (k_1 n_1 + \frac{k_2}{\Omega} n_1 n_2 + k_3(1 - n_1) + k_4 n_2) P(n_1, n_2, t). \end{aligned}$$

Despite having a relatively simple system here with effectively only two species, no time-dependent solution for its CME in Equation (17) is known to our knowledge. A solution in steady state has been derived in [59], but for most other systems not even a steady state solution is available. Therefore, one generally needs to rely on stochastic simulations or approximations to study the behaviour of such systems.

3.2.2 Moment equations

Suppose we are not interested in the whole distribution solution of the CME but only in its first few moments, say the mean and variance. Starting from the CME in Equation (15) one can derive time evolution equations for the moments of its solution as follows. To obtain the time evolution equation for the moment $\langle n_i \dots n_l \rangle$ we multiply Equation (15) by $n_i \dots n_l$ and sum over all molecule numbers, leading to

$$\partial_t \langle n_i \dots n_l \rangle = \sum_{r=1}^R \langle (n_i + S_{ir}) \dots (n_l + S_{lr}) f_r(\mathbf{n}) \rangle - \sum_{r=1}^R \langle n_i \dots n_l f_r(\mathbf{n}) \rangle. \quad (18)$$

Here, $\langle \cdot \rangle$ denotes the expectation with respect to the solution $P(\mathbf{n}, t)$ of the CME in Equation (15). For moments of up to order two Equation (18) becomes

$$\partial_t \langle n_i \rangle = \sum_{r=1}^R S_{ir} \langle f_r(\mathbf{n}) \rangle, \quad (19)$$

$$\partial_t \langle n_i n_j \rangle = \sum_{r=1}^R [S_{jr} \langle n_i f_r(\mathbf{n}) \rangle + S_{ir} \langle f_r(\mathbf{n}) n_j \rangle + S_{ir} S_{jr} \langle f_r(\mathbf{n}) \rangle]. \quad (20)$$

We see that if all $f_r(\mathbf{n})$ are zeroth or first-order polynomials in \mathbf{n} , i.e., the system is linear without any bimolecular or higher order reactions, the equation of a moment of order m depends only on moments of order m or lower, i.e., the equations are not coupled to higher order equations. The equations up to a certain order hence constitute a finite set of linear ordinary differential equations which can be readily solved numerically or by matrix exponentiation as will be described in the context of the CME in Section 3.4. Note that no approximation has been made here which means that the equations describe the exact moments of the CME's solution. This means that *for linear systems, the exact moments up to a finite order of the process can be obtained by solving a finite set of linear ordinary differential equations.*

If the systems is non-linear, i.e., contains bimolecular or higher order reactions, the equation of a certain moment depends on higher order moments. This means that the moment equations of different orders are coupled to each other, leading to an infinite hierarchy of coupled equations. This can obviously not be solved directly but gives the basis of approximation methods such as *moment closure approximations* which we will discuss in Section 4.3.

Example. Let us again consider the gene system in Figure 3 with reactions given in Equation (5). The corresponding stoichiometric matrix and propensity vector are given in Equation (10). Using these in Equations (19) and (20) one obtains

$$\partial_t y_1 = -\frac{k_2}{\Omega} y_{1,2} + k_3(1 - y_1), \quad (21)$$

$$\partial_t y_2 = k_1 y_1 - \frac{k_2}{\Omega} y_{1,2} + k_3(1 - y_1) - k_4 y_2, \quad (22)$$

$$\partial_t y_{1,1} = \frac{k_2}{\Omega} (-2y_{1,1,2} + y_{1,2}) + k_3(-2y_{1,1} + y_1 + 1), \quad (23)$$

$$\partial_t y_{1,2} = k_1 y_{1,1} + \frac{k_2}{\Omega} (-y_{1,1,2} - y_{1,2,2} + y_{1,2}) + k_3(-y_{1,1} - y_{1,2} + y_2 + 1) - k_4 y_{1,2}, \quad (24)$$

$$\begin{aligned} \partial_t y_{2,2} = & k_1(2y_{1,2} + y_1) + \frac{k_2}{\Omega} (-2y_{1,2,2} + y_{1,2}) \\ & + k_3(-2y_{1,1} - y_1 + 2y_2 + 1) + k_4(-2y_{2,2} + y_2), \end{aligned} \quad (25)$$

where we introduced the shorthand

$$y_{i_1, \dots, i_k} = \langle n_{i_1} \dots n_{i_k} \rangle. \quad (26)$$

Note that the equations for the first order moments in Equations (21) and (22) depend on the second moment $y_{1,2}$, and that the equations for the second moments (23)-(25)

depend on the third moments $y_{1,1,2}$ and $y_{1,2,2}$. Similarly, it is easy to see that moment equations of any order depend on higher order moments, which means that we have an infinite system of coupled equations. Note that all terms in Equations (21)-(25) depending on higher order moments are proportional to the rate constant k_2 of the bimolecular reaction in Equation (5), illustrating that the moment equations decouple in the absence of bimolecular (and higher order) reactions. This could be achieved here by setting $k_2 = 0$ for which the moment equations would decouple and could thus be solved numerically.

3.2.3 Non-mass-action propensity functions

So far we only considered propensity functions of mass-action type defined in Equation (12). However, in the literature propensity functions that are not of mass-action kinetics type are frequently used, such as Michaelis-Menten or Hill functions to model the dependence of the production rate of a product on the substrate concentration in an enzymatic catalysis or the dependence of the expression level of a gene on its transcription factor. Such propensity functions typically arise in reduced models where an effective reaction replaces several microscopic reactions. For a gene that becomes regulated by a transcription factor, for example, the binding of the transcription factor to the promoter of the gene is not modelled explicitly, but the effect of its concentration included in the modified propensity function of the expression reaction. Such non-mass-action type reactions should thus normally be seen as an effective reaction replacing a set of mass-action type reactions.

A possible reduced model of the gene expression system in Figure 3, for example, eliminates the gene from the system and combines the binding, unbinding and protein production reactions into one effective reaction $\emptyset \rightarrow P$ with a Michaelis-Menten type propensity function $k_1 k_3 n_P / (k_2 n_P + k_3)$, which means that the reaction producing protein now depends on the protein number n_P . Such reductions often emerge from separations of time-scales in a system and become exact in certain limits, see Section 4.6 for more details.

Independently of the origin of a non-mass-action propensity function, as long as it is interpreted as the firing rate of the corresponding reaction, the CME description remains valid. Unless explicitly stated otherwise, we will assume mass-action kinetics in this article.

3.3 Stochastic simulations

As mentioned above, there are generally no analytic solutions known to the CME. However, it is possible to directly simulate the underlying process. The stochastic simulation algorithm (SSA) is a popular Monte Carlo method that allows one to simulate exact sample paths of the stochastic process described by the CME.

The stochastic simulation algorithm was first proposed in the context of chemical kinetics by Gillespie [4, 5], and several variants have been proposed in the literature, see [9, 10, 11] for reviews. The basic idea is to simulate reaction events explicitly in time and to update the time and state vector accordingly. The stochastic process described by the CME is a continuous-time Markov jump process, which has the important property

that waiting times, i.e., time intervals between successive reaction events, are exponentially distributed [15]. Since it is easy to sample from exponential distributions, it is straightforward to simulate the occurrences of chemical reactions.

One example is the so-called *direct method* [4], which samples the time step τ for the next reaction to happen and subsequently which of the different reactions occurs. Specifically, let $p(\tau|\mathbf{n}, t)$ be the probability for the next reaction to happen in an infinitesimal time interval dt around $\tau + t$, given that the state of the system is \mathbf{n} at time t , and $p(r|\mathbf{n}, t)$ the probability that the next reaction is a reaction of type r . Using that $f_r(\mathbf{n})dt$ is the probability for the r^{th} reaction to happen in dt , it can be shown that [4]

$$p(\tau|\mathbf{n}, t) = \lambda \exp(-\tau\lambda), \quad \lambda = \sum_{r=1}^R f_r(\mathbf{n}), \quad (27)$$

$$p(r|\mathbf{n}, t) = \frac{f_r(\mathbf{n})}{\lambda}. \quad (28)$$

Samples from Equations (27) and (28) can be respectively obtained as

$$\tau = -\ln(u_1)/\lambda, \quad (29)$$

$$r = \text{smallest integer satisfying } \sum_{i=1}^r f_i(\mathbf{n}) > u_2\lambda, \quad (30)$$

where u_1 and u_2 are uniform random numbers between 0 and 1. The direct SSA method iteratively updates the state vector and time of the process by first sampling the time point of the next reaction event according to Equation (29) and subsequently sampling which reaction happens according to Equation (30).

Unfortunately, the applicability of the SSA is severely limited due to its computational cost. Since each and every reaction is simulated explicitly, the SSA becomes computationally expensive even for systems with few species. This is particularly the case if the molecule numbers have large fluctuations or if many reactions happen per unit time. In the first case a large number of samples have to be simulated to obtain statistically accurate results, whereas in the second case single simulations become expensive since the time between reaction events becomes small.

3.3.1 Extrinsic noise

So far we have only discussed stochastic simulations of systems with propensity functions that depend solely on the system's molecule numbers and the system's volume. This implies that the propensity functions are constant between reactions events, leading to inter-reaction intervals being exponentially distributed. However, as already mentioned in Section 3.2.1, it is frequently of interest to consider stochastically varying propensity functions. This could for example represent the effect of a randomly fluctuating environment on a cell.

To model the effect of extrinsic noise, one typically includes a stochastic variable in one or several propensity functions, whose dynamics may for example be governed by a stochastic differential equation. This means that the corresponding propensities

become explicit (stochastic) functions of time. While the CME in Equation (15) is still valid in this case, stochastic simulation algorithms discussed above are not since the inter-reaction times are now not exponentially distributed anymore.

The simplest approach is to assume the propensity functions to be constant between reaction events, which allows one to use the standard simulation algorithm [4, 60]. However, this is of course an approximation and can lead to inaccuracies if the extrinsic process (and hence the corresponding propensity functions) fluctuates strongly between consecutive reaction times. The most straightforward way to simulate such processes exactly is to integrate the propensity functions step-wise over time until a certain target value is reached [61, 62, 63]. Apart from numerical integration errors this class of algorithms is exact. However, due to the numerical integration it becomes computationally highly expensive. More recently an alternative exact method has been proposed [64]. This method introduces an additional reaction channel in such a way that the total propensity (i.e., λ in Equation (27)) is constant between successive reactions. This in turn allows one to use a standard SSA algorithm on the augmented system. It is hence generally more efficient than the aforementioned integral methods.

3.4 Exact results

As pointed out before, for most systems no analytic solutions of the CME are known. However, for certain classes of systems analytic solutions do exist. For some classes the general time-dependent case can be solved while others can be solved only in steady state. Moreover, analytic solutions have been derived for various simple example systems that do not fall under any of these classes. We give here first an overview of general classes of systems that can be solved analytically, and subsequently list some additional systems for which analytic solutions are known.

3.4.1 Finite state space

Suppose the state space of \mathbf{n} is finite with M elements. Let us associate with each state a probability $p_i(t)$, $i = 1, \dots, M$. In Section 3.2 we have noted that in this case the CME in Equation (15) is a finite system of coupled ordinary differential equations. If we write $\mathbf{p}(t) = (p_1(t), \dots, p_M(t))^T$, the CME can be written in the matrix form as

$$\partial_t \mathbf{p}(t) = \mathbb{E} \mathbf{p}(t), \quad (31)$$

where \mathbb{E} is a $M \times M$ matrix whose elements can be easily derived for a given system using Equation (15). The solution of Equation (31) is simply given by

$$\mathbf{p}(t) = \exp(\mathbb{E}t) \mathbf{p}(0). \quad (32)$$

Therefore, for systems with finite state space a solution of the CME for all times can be in principle computed using Equation (32). However, even for systems with finite state space, the dimension of Equation (32) is often quite large in practice, making matrix exponentiation computationally expensive. Efficient numerical methods for matrix exponentiation have been developed in recent years [65, 66], but for many chemical reaction systems of practical interest it remains an intractable task. For systems with large

or infinite state space Equation (32) is hence not of direct practical use. It forms the basis for certain approximation methods, however, see Section 4.6.

3.4.2 Linear systems

In Section 3.2.2 we found that the moment equations decouple from higher order equations for linear reaction systems, i.e., systems without bimolecular or higher order reactions. The moment equations up to a finite order can hence be directly solved numerically for such systems. For non-linear systems, in contrast, this is not possible since the moment equations couple to higher orders. Similarly, exact solutions for the whole distribution solution of the CME are easier to obtain for linear systems. Often this can be done by means of the *generating function method* which transforms the CME into a partial differential equation [15]. The form of the latter becomes particularly simple for linear systems and can often be solved analytically.

In 1966 Darvey and Staff used the generating function approach to derive a multinomial solution for $P(\mathbf{n}, t)$ for multinomial initial conditions and for systems with $\sum_{i=1}^N s_{ir} = \sum_{i=1}^N s'_{ir} = 1$ for all reactions r , i.e., exactly one reactant and exactly one product molecule for all reactions [67]. Note that these are closed systems. Later Gardiner used the so-called “Poisson representation” to show that for systems with $\sum_{i=1}^N s_{ir}, \sum_{i=1}^N s'_{ir} \leq 1$ for all reactions r and Poisson product initial conditions, the solution of the CME is a Poisson product for all times [68]. The auto- and cross-correlation functions for such systems have more recently been computed in [69]. These results have been generalised to systems with $\sum_{i=1}^N s_{ir}, \sum_{i=1}^N s'_{ir} \leq 1$ for all reactions r , with arbitrary initial conditions in [70]. In this case the solution can be written as a convolution of multinomial and product Poisson distributions. In all these cases, the solutions can be constructed from the solution of the deterministic rate equations discussed in Section 3.1 and can hence be obtained efficiently by (numerically) solving a finite set of ordinary differential equations.

We would like to stress that the methods described here do not apply to all linear reactions, but only to those with less than two product molecules ($\sum_{i=1}^N s'_{ir} \leq 1$). This excludes reactions of the type $A \rightarrow A + B$ which are common in biology, such as translation of a protein B from its mRNA, A .

3.4.3 Non-linear systems

As discussed in the previous section, for linear systems where each reaction has at most one product molecule, an analytic solution of the CME can be derived. In contrast, for non-linear systems or linear systems including reactions with two product molecules no analytic solutions are known in general. However, for certain subclasses of systems it is possible to derive steady-state solutions, i.e., solutions satisfying $\partial_t P(\mathbf{n}, t) = 0$. One example are *reversible systems* that obey *detailed balance*, which is also called *thermodynamic equilibrium*. By “reversible” we mean that each reaction possesses a corresponding reversing reaction. Note that the steady-state condition $\partial_t P(\mathbf{n}, t) = 0$ in the CME in Equation (15) merely means that the probability flow out of a state (second term in the CME) equals the probability flow into a state (first term in the CME). These probabilities consist of sums over all reactions. Detailed balance is a stronger condition, as it

requires the flow of each single reaction in each state to be balanced by its corresponding reversible reaction. Specifically, let r and r' be the indices of two reactions that revert each other and let $P(\mathbf{n})$ be a steady-state solution. The detailed balance condition reads [15]

$$f_r(\mathbf{n})P(\mathbf{n}) = f_{r'}(\mathbf{n} + \mathbf{S}_r)P(\mathbf{n} + \mathbf{S}_r), \quad (33)$$

which needs to be fulfilled for all reactions and all \mathbf{n} . Note that $\mathbf{S}_r = -\mathbf{S}_{r'}$. The left side of Equation (33) is the flow out of state \mathbf{n} due to the r^{th} reaction, while the right side is the flow into state \mathbf{n} due to the backward reaction r' . An analogue detailed balance condition can be formulated for the deterministic rate equations. It can be shown that for a reversible reaction system with mass-action kinetics, the stochastic system is in detailed balance if and only if the deterministic system is in detailed balance [71]. If the detailed balance condition is fulfilled, the solution of the CME is a product of Poisson distributions times a function accounting for conservation laws in molecule numbers, where the mean values of the Poisson distributions are simply given by the detailed balance solutions of the rate equations [72, 15].

As an example, consider a closed, bimolecular system with reactions $A + A \xrightleftharpoons[k_2]{k_1} B$. Let n_A and n_B be the molecule numbers of A and B respectively. The quantity $n_A + 2n_B = C$ is conserved. The corresponding steady-state solution of the CME in detailed balance reads

$$P(n_A, n_B) = \frac{e^{-\alpha_1} \alpha_1^{n_A}}{n_A!} \frac{e^{-\alpha_2} \alpha_2^{n_B}}{n_B!} \delta(n_A + 2n_B - C), \quad (34)$$

where $\alpha_1 = \Omega\phi_1^0$ and $\alpha_2 = \Omega\phi_2^0$, ϕ_1^0 and ϕ_2^0 are the steady-state solutions of the rate equations of A and B , respectively, Ω is the system volume, and δ is the delta function accounting for the conservation law.

These results have more recently been generalised to *weakly reversible* reaction systems [73]. A system is called weakly reversible when for each reaction the change in state space can be reversed by a chain of other reactions. For example, a system with the reactions $X_1 \rightarrow X_2 \rightarrow X_3 \rightarrow X_1$ is not reversible but it is weakly reversible. Weak reversibility is a generalisation of reversibility since each reversible system is also weakly reversible. Correspondingly, the concept of *complex balance* has been employed which is a generalisation of detailed balance [74]. It can be shown that if the rate equations of a system with mass-action kinetics are complex balanced, the CME possesses a steady-state solution given by a product of Poisson distributions times a function accounting for conservation laws [73]. The mean values of the Poisson distributions are given by the complex-balance solutions of the rate equations. Moreover, it can be shown that for any mass-action kinetics reaction system with deficiency zero (see [75] for a definition), the rate equations have a complex-balance solution if and only if the system is weakly reversible [75].

3.4.4 Exact results for special cases

Most reaction systems of practical interest are neither linear nor do they satisfy the detailed/complex balance conditions. However, the CME has been solved for several

specific systems that do not fall under any of these categories. In particular for systems with only a single reaction it is often straightforward to obtain a time-dependent solution of the CME. Some examples are given below. Most of these solutions are obtained either by means of the aforementioned generating function method or the Poisson representation. While the former transforms the CME into a partial differential equation, the latter maps the discrete process underlying the CME to a continuous diffusion process governed by a Langevin equation defined in a complex-valued space [68]. We would like to stress that both approaches are exact. However the Poisson representation has not been used much in the literature, probably because of its relative complexity compared to other methods and hence we do not cover it in detail here. We refer the reader to the following references for specific applications [76, 77, 78, 79, 80, 81, 82]. For more details on the Poisson representation and the generating function approach we refer the reader to [15].

Single reaction systems.

Time-dependent solutions have been derived for

- $A \rightarrow A + A$ [53]
- $A + B \rightarrow C$ [56]
- $A + A \rightarrow C$ [56]
- $A + B \leftrightarrow C$ [83]

Gene regulatory networks.

- $G^{\text{on}} \rightarrow G^{\text{on}} + P, G^{\text{on}} + P \leftrightarrow G^{\text{off}}, P \rightarrow \emptyset, G^{\text{off}} \rightarrow G^{\text{off}} + P, G^{\text{off}} \rightarrow G^{\text{on}}$: a steady-state solution has been derived if there is one gene in the system [59]. Note that if we ignore the last two reactions this system corresponds to the gene system in Figure 3.
- A system similar to the previous one but with geometrically distributed bursts of protein production has been solved in steady state in [84].

Enzyme systems.

- $E + S \leftrightarrow C \rightarrow E + P$: A time-dependent solution has been derived for the case of a single enzyme in the system [85].
- Adding the reaction $E + P \rightarrow C$ to the previous system makes it reversible and a corresponding steady-state solution has been derived in [86].
- The system with added substrate input reaction $\emptyset \rightarrow S$ has been solved in steady state in [87].
- A similar system but with multiple substrate species competing for the enzyme has been solved in steady state in [88].

Most of these examples only consider steady-state solutions and all of them consider systems with few species. The lack of solutions for more complicated systems makes the development of approximation methods necessary.

4 Approximation methods

As discussed in the previous Section, analytic solutions of the CME are known only for very restrictive classes of systems and few simple special cases. For most systems of practical interest, no analytic solutions are known to date. Stochastic simulation algorithms introduced in Section 3.3 allow to simulate exact samples of the underlying stochastic process, but they quickly become computationally infeasible for larger systems. For these reasons significant effort has been spent in the literature on the development of approximation methods. We give here an introduction to a wide range of such methods. First, we give a detailed introduction to a few approximation methods that can be applied to (almost) arbitrary systems without any pre-knowledge of the system necessary. We derive the approximations, give examples and discuss their properties. Subsequently, we give an overview of other types of approximation methods developed in the literature.

The first method aims at approximating the solution of the CME, namely the *chemical Fokker-Planck equation* (CFPE) and the associated *chemical Langevin equation* (CLE), which we introduce in Section 4.1. The CFPE/CLE define an approximating stochastic process. An alternative method for approximating the solution of the CME is given by the *system size expansion*, which constitutes a systematic expansion of the CME in the inverse volume size. The system size expansion includes the popular *linear noise approximation* (LNA) and also provide an efficient way of approximating the moments of a process. We discuss the system size expansion in Section 4.2. Next, in Section 4.3 we introduce a certain class of *moment closure approximations*, which approximate the moments of a process. In Section 4.4 we show how such approximate moments can be used to construct distributions using the maximum entropy principle. Next, in Section 4.5 we review software packages implementing the discussed approximation methods. Finally, in Section 4.6 we give a brief overview of other approximation methods found in the literature. As we shall see, many of these methods use the CLE, the system size expansion or moment closure approximations as building blocks.

4.1 The chemical Langevin equation

The *chemical Langevin equation* (CLE) and the corresponding *chemical Fokker-Planck equation* (CFPE) constitute a popular diffusion approximation of the CME. Kramers and Moyal derived the CFPE by applying a Taylor expansion to the CME which upon truncation leads to a partial differential equation approximation of the CME [89, 90]: Suppose we let the variables in the CME in Equation (15) become continuous and let $\mathbf{x} = (x_1, \dots, x_N)$, where x_i is the continuous variable denoting the molecule number of species X_i . Performing a Taylor expansion to second order around \mathbf{x} in the first term of the r.h.s. of Equation (15) gives (with \mathbf{n} replaced by \mathbf{x})

$$f_r(\mathbf{x} - \mathbf{S}_r)P(\mathbf{x} - \mathbf{S}_r, t) \approx f_r(\mathbf{x})P(\mathbf{x}, t) - \sum_{i=1}^N S_{ir} \frac{\partial}{\partial x_i} [f_r(\mathbf{x})P(\mathbf{x}, t)] \\ + \sum_{i,j=1}^N S_{ir} S_{jr} \frac{\partial}{\partial x_i} \frac{\partial}{\partial x_j} [f_r(\mathbf{x})P(\mathbf{x}, t)]. \quad (35)$$

Inserting this into the CME in Equation (15), we see that the first term on the r.h.s. of Equation (35) cancels the last term on the r.h.s. of Equation (15) leading to the CFPE:

$$\partial_t P(\mathbf{x}, t) = - \sum_{i=1}^N \frac{\partial}{\partial x_i} [A_i(\mathbf{x})P(\mathbf{x}, t)] + \frac{1}{2} \sum_{i,j=1}^N \frac{\partial}{\partial x_i} \frac{\partial}{\partial x_j} [B_{ij}(\mathbf{x})P(\mathbf{x}, t)], \quad (36)$$

where the drift vector \mathbf{A} and diffusion matrix \mathbf{B} are respectively given by

$$A_i(\mathbf{x}) = \sum_{r=1}^R S_{ir} f_r(\mathbf{x}), \quad (37)$$

$$B_{ij}(\mathbf{x}) = \sum_{r=1}^R S_{ir} S_{jr} f_r(\mathbf{x}). \quad (38)$$

Note that the drift vector $\mathbf{A}(\mathbf{x})$ and diffusion matrix $\mathbf{B}(\mathbf{x})$ do not depend on time. Note also that whereas the state variables denote discrete molecule numbers in the CME, they denote continuous real numbers in the CFPE. The CFPE in Equation (36) is equivalent to the CLE

$$d\mathbf{x} = \mathbf{A}(\mathbf{x})dt + \mathbf{C}(\mathbf{x})d\mathbf{W}, \quad \mathbf{C}(\mathbf{x})\mathbf{C}(\mathbf{x})^T = \mathbf{B}(\mathbf{x}), \quad (39)$$

which is an *Ito stochastic differential equation* [15]. \mathbf{W} in Equation (39) is a multi-dimensional Wiener process. It can be shown [15] that the distribution of a process described by Equation (39) agrees exactly with the solution of the corresponding Fokker-Planck equation in (36). One can thus interpret the CLE in Equation (39) as a generator of realisations of the stochastic process described by the corresponding CFPE. In this sense, the CLE and CFPE are considered to be equivalent to each other. An alternative derivation, which leads directly to the CLE in Equation (39), was given by Gillespie in [91].

Generally there exist different choices for $\mathbf{C}(\mathbf{x})$ in Equation (39) corresponding to different factorisations of the matrix $\mathbf{B}(\mathbf{x})$; these lead to as many different representations of the CLE. One possibility is given by

$$dx_i = \sum_{r=1}^R S_{ir} f_r(\mathbf{x})dt + \sum_{r=1}^R S_{ir} \sqrt{f_r(\mathbf{x})} dW_r, \quad i = 1, \dots, N. \quad (40)$$

This representation of the CLE is the one most commonly used in the literature [91, 92].

Example. Let us come back to the gene expression system in Figure 3 with reactions in Equation (5) and consider the corresponding CFPE and CLE. Using the stoichiometric matrix in Equation (10) and propensity vector in Equation (16) we obtain for the drift vector and diffusion matrix defined in Equation (37) and Equation (38), respectively,

$$\mathbf{A}(\mathbf{x}) = \begin{pmatrix} -\frac{k_2}{\Omega} x_1 x_2 + k_3(1 - x_1) \\ k_1 x_1 - \frac{k_2}{\Omega} x_1 x_2 + k_3(1 - x_1) - k_4 x_2 \end{pmatrix}, \quad (41)$$

$$\mathbf{B}(\mathbf{x}) = \begin{pmatrix} \frac{k_2}{\Omega} x_1 x_2 + k_3(1 - x_1) & \frac{k_2}{\Omega} x_1 x_2 + k_3(1 - x_1) \\ \frac{k_2}{\Omega} x_1 x_2 + k_3(1 - x_1) & k_1 x_1 + \frac{k_2}{\Omega} x_1 x_2 + k_3(1 - x_1) + k_4 x_2 \end{pmatrix}, \quad (42)$$

where x_1 and x_2 are the (continuous) particle numbers of G^{on} and P , respectively. To obtain the CLE in Equation (39), we have to compute \mathbf{C} , which is the square root of \mathbf{B} and thus generally not uniquely defined. One possibility as given in Equation (40) reads

$$C_{ij}(\mathbf{x}) = \begin{pmatrix} 0 & -\sqrt{\frac{k_2}{\Omega}x_1x_2} & \sqrt{k_3(1-x_1)} & 0 \\ \sqrt{k_1x_1} & -\sqrt{\frac{k_2}{\Omega}x_1x_2} & \sqrt{k_3(1-x_1)} & -\sqrt{k_4x_2} \end{pmatrix}. \quad (43)$$

This gives rise to the CLE

$$dx_1 = \left(-\frac{k_2}{\Omega}x_1x_2 + k_3(1-x_1)\right)dt - \sqrt{\frac{k_2}{\Omega}x_1x_2}dW_2 + \sqrt{k_3(1-x_1)}dW_3, \quad (44)$$

$$dx_2 = \left(k_1x_1 - \frac{k_2}{\Omega}x_1x_2 + k_3(1-x_1) - k_4x_2\right)dt \quad (45)$$

$$+ \sqrt{k_1x_1}dW_1 - \sqrt{\frac{k_2}{\Omega}x_1x_2}dW_2 + \sqrt{k_3(1-x_1)}dW_3 - \sqrt{k_4x_2}dW_4,$$

where the W_i are independent Wiener processes. Note that it does not make a difference if one changes the signs in front of the square roots in Equation (43), or equivalently in front of the noise terms in Equation (44), as long as one does so simultaneously for each occurrence of a specific term, i.e., changes the sign of whole columns in Equation (43). To see that such changes are equivalent note that the diffusion matrix $\mathbf{B} = \mathbf{C}\mathbf{C}^T$ of the CFPE is invariant under such changes. This can also be seen directly from the CLE in Equations (44) and (45) since the Wiener processes are symmetric.

4.1.1 Stochastic simulations

As for the CME, there are no analytic solutions known for the CLE for most systems. However, the computational cost of CLE simulations scales with the number of species, rather than with the rate of reaction events as CME simulations. This means that CLE simulations are often more efficient than simulations of the CME, typically if the molecule counts in a system are not too small.

A popular method to simulate the CLE is the *Euler-Maruyama algorithm* which discretises time into intervals dt and simulates the process iteratively as [93]

$$x(t+dt) = x(t) + A(x(t))dt + \sqrt{B(x(t))dt}dw, \quad dw \sim \mathcal{N}(0,1), \quad (46)$$

where $\mathcal{N}(0,1)$ is a normal distribution with mean 0 and variance 1. The smaller dt , the better the true process is approximated by Equation (46) but also the slower the algorithm becomes. The right choice of dt is therefore a tradeoff between accuracy and efficiency. It is important to point out that more efficient simulation methods exist, see for example [93, 94, 95, 96].

4.1.2 Properties and recent developments

The CLE is typically a good approximation whenever the molecule numbers of the system are not too small and is particularly useful if one is interested in time trajectories or

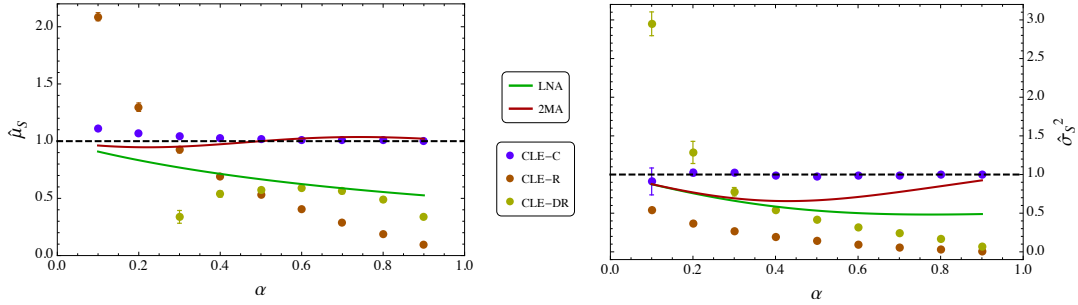


Figure 4: **Comparison of different CLE implementations.** The figures show the steady-state mean (left) and variance (right) of a substrate in a Michaelis-Menten reaction system as a function of the strength of the substrate-to-enzyme binding strength. The results are normalised by the exact CME result, which means that the horizontal dashed line corresponds to the exact value. CLE-C corresponds to the complex-valued CLE, CLE-R to a real-valued implementation with rejecting boundaries at zero molecules, and CLE-DR to the real-valued version proposed in [98]. LNA and 2MA correspond to the linear noise approximation and the second-order normal moment closure, which we will introduce in Sections 4.2 and 4.3, respectively. We observe that both real-valued implementations give large deviations from the CME result, significantly more than the LNA and 2MA. The complex CLE on the other hand is significantly more accurate than all the other methods. The figure is taken from [87] with permission kindly granted by the publisher.

distributions of a process. It can be shown that the differences between the CLE and the CME tend to zero in the limit of large molecule numbers [97]. In other words, the CLE becomes exact in the thermodynamic limit.

By multiplying the CFPE in Equation (36) by $x_i \dots x_l$ and integrating over the whole state space, one obtains ordinary differential equation for the moment $\langle x_i \dots x_l \rangle$ of the process described by the CLE. Importantly, it turns out that the equations for moments of up to order two are *exactly the same as the corresponding equations derived from the CME*. These are given in Equations (19) and (20). Note however that since they are generally coupled to higher order moments for which the evolution equations derived from the CLE and CME do not agree, the first two moments (and higher order moments) of the CLE do not generally agree with the ones of the CME. However, since the moment equations decouple for linear systems as shown in Section 3.2.2, we obtain the important result that *the moments up to order two of the processes described by the CLE and CME agree exactly for linear reaction systems*.

Note that the CME has a natural boundary at zero molecule numbers, i.e., for a sensible initial condition with zero probability to have a negative number of molecules, this probability remains zero for all times. Until recently it has not been clear how this boundary condition behaves when approximating the discrete process underlying the CME by a continuous process using the CLE. As it turns out, this boundary issue leads to an ill-definedness of the CLE due to occurrences of square roots of negative expressions

in finite time with finite probability [99, 100]. Since traditionally the domain of the CLE is (implicitly) assumed to be that of real numbers, the CLE is not well-defined in this case. More recently, it has been shown that this problem is independent of the chosen factorisation of the CFPE’s noise matrix \mathbf{B} (c.f. Equation (39)), which means that the CLE is not well-defined for real variables [87]. In this study it has been shown that the same can be expected for the majority of reaction systems.

Several modified versions of the CLE have been proposed that try to keep the state space real, for example [99, 98]. However, these are *ad hoc* modifications and have been found to introduce high inaccuracies for some non-linear reaction systems [87]. Importantly, they also have been found to violate the CLE’s exactness for the moments up to order two for linear systems. Alternatively, the ill-defined problem can be solved by extending the state space of the CLE to complex variables. It has been proven that this leads to real-valued moments, real-valued autocorrelation functions and real-valued power spectra, and to restore the CLE’s exactness for the moments up to order two for linear systems [87]. This complex CLE was found to be highly accurate for some non-linear systems in comparison to other modified versions. For one example the results from [87] are shown in Figure 4. The complex-valued CLE has other drawbacks, however. For example, it does not directly give approximations of the process or of distributions. Rather, the results have to be projected to real space.

In recent years the CLE has been frequently used for approximating the dynamics of intermediate or high abundance species in so-called *hybrid methods*, see Section 4.6.4. In this case, the probability of negative concentrations becomes very small. If the partitioning into low and high abundance species is done adaptively, the problem of square roots of negative expressions is avoided completely. In this case the complex-valued CLE automatically reduces to the real-valued CLE. The CLE is particularly useful in simulation based hybrid methods since it gives an approximation of the whole process (rather than for example only the process’s moments).

Recently, a tensor-based method has been proposed in [101] for the direct numerical solution of the CFPE. This is a promising approach since it avoids computationally expensive ensemble averaging and has been used for sensitivity and bifurcation analysis [101]. Recall that the CLE becomes exact in the limit of large system sizes. As one may expect, the CLE therefore generally captures the multimodality of the CME’s solution if the multimodality persists for large volumes, since this is the case if the deterministic rate equations are multistable. Surprisingly, however, it has recently been found that the CLE is also able to reproduce *noise-induced multimodality* of the CME’s solution, i.e., multimodality that occurs only if the system volume is decreased below a certain critical value [102]. However, it was found that this is not the case for all reaction systems [103]. In [104], the CLE as a diffusion approximation was applied to Petri nets.

4.2 The system size expansion

Suppose we are not interested in approximating the whole process but only in its distribution or first few moments. Running stochastic simulations of the CLE seem like unnecessary computational effort in this case and the question arises if there exist more efficient approximation methods to achieve this. We next discuss the *system size expansion*

sion which aims at approximating the distribution or the first few moments of a process. The system size expansion is a perturbative expansion of the CME in the inverse system size originally developed by van Kampen [105, 14]. The idea is to separate concentrations into a deterministic part, given by the solution of the deterministic rate equations, and a part describing the fluctuations about the deterministic mean.

4.2.1 Derivation

The system size expansion splits the instantaneous particle numbers n_i of the CME into a deterministic part and a fluctuating part as

$$\frac{n_i}{\Omega} = \phi_i + \Omega^{-1/2}\epsilon_i, \quad i = 1, \dots, N, \quad (47)$$

where Ω is the volume of the system, ϕ_i is the solution of the deterministic rate equations in (4) and we introduced the new variables ϵ_i representing fluctuations about the deterministic mean. Following [106], the system size expansion proceeds by transforming the master equation to these new variables ϵ_i . To this end, we rewrite the CME in Equation (15) as

$$\partial_t P(\mathbf{n}, t) = \sum_{r=1}^R \left(\prod_{i=1}^N E_i^{-S_{ir}} - 1 \right) f_r(\mathbf{n}) P(\mathbf{n}, t), \quad (48)$$

where we introduced the step operators $E_i^{-S_{ir}}$, which are defined in terms of their action on a general function of the state space, $h(\mathbf{n}) = h(n_1, \dots, n_N)$ as

$$E_i^{-S_{ir}} h(n_1, \dots, n_i, \dots, n_N) = h(n_1, \dots, n_i - S_{ir}, \dots, n_N). \quad (49)$$

The system size expansion now proceeds in three steps:

1. **Time derivative.** Due to the change of variables in Equation (47) we need to transform the distribution $P(\mathbf{n}, t)$ into a distribution $\Pi(\boldsymbol{\epsilon}, t)$ of the new variables. Note that the time derivative in the CME in Equation (48) is taken at constant \mathbf{n} , which implies $\partial \epsilon_i / \partial t = -\Omega^{1/2} \partial \phi_i / \partial t$. This in turn leads to

$$\partial_t P(\mathbf{n}, t) = \frac{\partial \Pi(\boldsymbol{\epsilon}, t)}{\partial t} - \Omega^{1/2} \sum_{i=1}^N \frac{\partial \phi_i}{\partial t} \frac{\partial \Pi(\boldsymbol{\epsilon}, t)}{\partial \epsilon_i}. \quad (50)$$

2. **Step operator.** Using Equation (47) in the definition of the step operator in Equation (49), we obtain

$$\prod_{i=1}^N E_i^{-S_{ir}} h(\epsilon_1, \dots, \epsilon_N) = h(\epsilon_1 - \Omega^{-1/2} S_{1r}, \dots, \epsilon_N - \Omega^{-1/2} S_{Nr}). \quad (51)$$

Expanding the r.h.s. in powers of $\Omega^{-1/2}$ we can expand the term of the CME involving the step operators as

$$\begin{aligned} \prod_{i=1}^N E_i^{-S_{ir}} - 1 &= -\Omega^{-1/2} \sum_{i=1}^N S_{ir} \frac{\partial}{\partial \epsilon_i} + \frac{\Omega^{-1}}{2} \sum_{i,j=1}^N S_{ir} S_{jr} \frac{\partial^2}{\partial \epsilon_i \partial \epsilon_j} \\ &\quad - \frac{\Omega^{-3/2}}{6} \sum_{i,j,l=1}^N S_{ir} S_{jr} S_{lr} \frac{\partial^3}{\partial \epsilon_i \partial \epsilon_j \partial \epsilon_l} + O(\Omega^{-2}). \end{aligned} \quad (52)$$

3. Propensity functions. We assume that the propensity functions in the CME can be expanded as

$$f_r(\mathbf{n}) = \Omega \sum_{i=0}^{\infty} \Omega^{-i} f_r^{(i)}\left(\frac{\mathbf{n}}{\Omega}\right). \quad (53)$$

If we assume mass-action kinetics for which the propensity functions take the form in Equation (12), the propensity functions are polynomials and it is easy to see that such an expansion always exists. The same is also true for many propensity functions that are not of mass-action kinetics type, such as Michaelis-Menten and Hill-type propensity functions. Using Equation (47) and expanding the r.h.s. of Equation (53) in $\Omega^{-1/2}$ we obtain

$$\begin{aligned} f_r(\mathbf{n}) &= g_r(\phi) + \Omega^{-1/2} \sum_{i=1}^N \epsilon_i \frac{\partial g_r(\phi)}{\partial \phi_i} \\ &\quad + \Omega^{-1} \left(\frac{1}{2} \sum_{i,j=1}^N \epsilon_i \epsilon_j \frac{\partial^2 g_r(\phi)}{\partial \phi_i \partial \phi_j} - f_r^{(1)} \right) + O(\Omega^{-3/2}), \end{aligned} \quad (54)$$

where we have identified $f_r^{(0)}(\phi) = g_r(\phi)$ and $g_r(\phi)$ is the macroscopic rate function of the r^{th} reaction introduced in Equation (2) in context of the deterministic rate equations.

Applying Equations (50), (52) and (54) to Equation (48) we obtain

$$\partial_t \Pi(\epsilon, t) = [\mathcal{L}^{(0)} + \Omega^{-1/2} \mathcal{L}^{(1)} + \Omega^{-1} \mathcal{L}^{(2)} + O(\Omega^{-3/2})] \Pi(\epsilon, t), \quad (55)$$

where we defined

$$\mathcal{L}^{(0)} = -\partial_i J_{i\epsilon_j}^{(0)j} + \frac{1}{2} D_{ij}^{(0)} \partial_i \partial_j, \quad (56)$$

$$\mathcal{L}^{(1)} = +\Omega^{-1/2} \left(\frac{1}{2} J_{i,j}^{(0)q} \partial_i \partial_j \epsilon_q - \frac{1}{2} J_i^{(0)q,r} \partial_i \epsilon_q \epsilon_r - \frac{1}{6} D_{i,j,k}^{(0)} \partial_i \partial_j \partial_k - D_i^{(1)} \partial_i \right) \Pi, \quad (57)$$

$$\begin{aligned} \mathcal{L}^{(2)} &= +\Omega^{-1} \left(-\frac{1}{6} J_{i,j,k}^{(0)q} \partial_i \partial_j \partial_k \epsilon_q + \frac{1}{4} J_{i,j}^{(0)q,r} \partial_i \partial_j \epsilon_q \epsilon_r - \frac{1}{6} J_i^{(0)q,r,s} \partial_i \epsilon_q \epsilon_r \epsilon_s \right. \\ &\quad \left. - J_i^{(1)q} \partial_i \epsilon_q + \frac{1}{2} D_{i,j}^{(1)} \partial_i \partial_j + \frac{1}{24} D_{i,j,k,l}^{(0)} \partial_i \partial_j \partial_k \partial_l \right). \end{aligned} \quad (58)$$

Here we use the shorthand $\partial_i = \partial/\partial\epsilon_i$ and assume summations over doubly occurring indices for notational simplicity. We note that in Equation (55) also two terms of order $\Omega^{1/2}$ occur. However, these terms cancel each other because they just correspond to the deterministic rate equations. In Equations (56)-(58) we defined

$$D_{i,\dots,k}^{(n)} = \sum_{r=1}^R S_{ir} \dots S_{kr} f_r^{(n)}(\phi), \quad (59)$$

$$J_{i,\dots,k}^{(n)l,\dots,m} = \partial_{\phi_l} \dots \partial_{\phi_m} D_{i,\dots,k}^{(n)}. \quad (60)$$

Note that both the $\mathbf{J}^{(i)}$ and $\mathbf{D}^{(i)}$ matrices depend on the solution ϕ of the rate equations in (4) and are thus generally time dependent. Comparing with Equation (4) we find that $\mathbf{J}^{(0)}$ in Equation (60) is just the Jacobian of the rate equations.

4.2.2 The linear noise approximation

The popular *linear noise approximation* (LNA) [105, 14] is obtained if we truncate the expansion in Equation (55) to zeroth order, leading to

$$\partial_t \Pi(\boldsymbol{\epsilon}, t) = \left[- \sum_{i=1}^N \frac{\partial}{\partial \epsilon_i} \sum_{j=1}^N J^{(0)j}_i \epsilon_j + \frac{1}{2} \sum_{i,j=1}^N D_{ij}^{(0)} \frac{\partial}{\partial \epsilon_i} \frac{\partial}{\partial \epsilon_j} \right] \Pi(\boldsymbol{\epsilon}, t) + O(\Omega^{-1/2}). \quad (61)$$

Equation (61) is a Fokker-Planck equation with drift and diffusion linear and constant in $\boldsymbol{\epsilon}$, respectively, and hence has a multivariate normal solution under appropriate initial conditions. By multiplying Equation (61) with ϵ_i and $\epsilon_i \epsilon_j$ and integrating over all $\boldsymbol{\epsilon}$, one obtains ordinary differential equations for the first and second-order moments, $\langle \epsilon_i \rangle$ and $\langle \epsilon_i \epsilon_j \rangle$, respectively. By doing so one finds that if the mean is initially zero, $\langle \epsilon_i \rangle|_{t=0} = 0$, it remains zero for all times. This is normally the case since otherwise the initial condition of the rate equations would not agree with the initial mean value of the stochastic system. If we additionally assume deterministic initial conditions, we also have $\langle \epsilon_i \epsilon_j \rangle|_{t=0} = 0$. The solution of Equation (61) is thus a multivariate normal distribution with zero mean. Since \mathbf{n} and $\boldsymbol{\epsilon}$ are related by the linear transformation given in Equation (47), the distribution of \mathbf{n} is also given by a multivariate normal distribution. The mean of the latter satisfies the rate equations in Equation (4) and the covariance $\boldsymbol{\Sigma}$ defined as $\Sigma_{ij} = \langle n_i n_j \rangle - \langle n_i \rangle \langle n_j \rangle$ fulfils

$$\partial_t \boldsymbol{\Sigma} = \mathbf{J}^{(0)} \boldsymbol{\Sigma} + \boldsymbol{\Sigma} \mathbf{J}^{(0)T} + \Omega^{-1} \mathbf{D}^{(0)}. \quad (62)$$

The LNA describes the lowest order fluctuations of the system size expansion about the deterministic mean and is valid in the limit of large volumes, i.e., the thermodynamic limit.

An alternative derivation of the LNA is given in [15], where it is derived from the chemical Langevin equation given in Equation (39).

Example. Let us consider the rate equations and LNA for the gene system in Figure 3. The corresponding stoichiometric matrix, the macroscopic rate vector and the rate equations were already derived in Section 3.1 and are given in Equations (10) and

(11), respectively. The LNA defined in Equation (61) requires matrices $\mathbf{J}^{(0)}$ and $\mathbf{D}^{(0)}$ defined in Equations (59) and (60), respectively, for which we obtain

$$\mathbf{D}^{(0)} = \begin{pmatrix} k_2\phi_1\phi_2 + k_3(\frac{1}{\Omega} - \phi_1) & k_2\phi_1\phi_2 + k_3(\frac{1}{\Omega} - \phi_1) \\ k_2\phi_1\phi_2 + k_3(\frac{1}{\Omega} - \phi_1) & k_1\phi_1 + k_2\phi_1\phi_2 + k_3(\frac{1}{\Omega} - \phi_1) + k_4\phi_2 \end{pmatrix}, \quad (63)$$

$$\mathbf{J}^{(0)} = \begin{pmatrix} -k_2\phi_2 - k_3 & -k_2\phi_1 \\ k_1 - k_2\phi_1 - k_3 & -k_2\phi_1 - k_4 \end{pmatrix}. \quad (64)$$

Note that $\mathbf{D}^{(0)}$ and $\mathbf{J}^{(0)}$ are functions of the time-dependent solutions ϕ_1 and ϕ_2 of the rate equations in (11). The solution of the LNA is a normal distribution in $\mathbf{n} = (n_1, n_2)$. Its mean is obtained by (numerically) solving the rate equations in (11), and the covariance by subsequently solving Equation (62) using Equations (63) and (64).

4.2.3 Higher order corrections

Suppose we want to include higher orders in $\Omega^{-1/2}$ in Equation (55) beyond the LNA. In contrast to the LNA, now the PDE in Equation (55) can generally not be solved analytically. However, we can derive ordinary differential equations for the moments of the system accurate to the corresponding order in $\Omega^{-1/2}$, as follows. Assuming an expansion of the distribution $\Pi(\boldsymbol{\epsilon}, t)$ as

$$\Pi(\boldsymbol{\epsilon}, t) = \sum_{i=0}^{\infty} \Omega^{-i/2} \Pi^{(i)}(\boldsymbol{\epsilon}, t), \quad (65)$$

one obtains an expansion of the moments

$$\langle \epsilon_j \dots \epsilon_k \rangle = \sum_{i=0}^{\infty} [\epsilon_j \dots \epsilon_k]^{(i)} \Omega^{-i/2}, \quad [\epsilon_j \dots \epsilon_k]^{(i)} = \int d\boldsymbol{\epsilon} \epsilon_j \dots \epsilon_k \Pi^{(i)}(\boldsymbol{\epsilon}, t). \quad (66)$$

Multiplying Equation (55) with $\epsilon_j \dots \epsilon_k$, using Equations (65) and (66) and integrating over all ϵ , we obtain

$$\partial_t[\epsilon_i]_0 = J^{(0)q}_i[\epsilon_q]_0, \quad (67)$$

$$\partial_t[\epsilon_i]_1 = J^{(0)q}_i[\epsilon_q]_1 + D_i^{(1)} + \frac{1}{2}J^{(0)q,r}_i[\epsilon_q\epsilon_r]_0, \quad (68)$$

$$\partial_t[\epsilon_i]_2 = J^{(0)q}_i[\epsilon_q]_2 + \frac{1}{2}J^{(0)q,r}_i[\epsilon_q\epsilon_r]_1 + \frac{1}{6}J^{(0)q,r,s}_i[\epsilon_q\epsilon_r\epsilon_s]_0 + J^{(1)q}_i[\epsilon_q]_0, \quad (69)$$

$$\partial_t[\epsilon_i\epsilon_j]_0 = J^{(0)q}_i[\epsilon_j\epsilon_q]_0 + J^{(0)q}_j[\epsilon_i\epsilon_q]_0 + D_{ij}^{(0)}, \quad (70)$$

$$\begin{aligned} \partial_t[\epsilon_i\epsilon_j]_1 &= J^{(0)q}_i[\epsilon_j\epsilon_q]_1 + J^{(0)q}_j[\epsilon_i\epsilon_q]_1 + D_i^{(1)}[\epsilon_j]_0 + D_j^{(1)}[\epsilon_i]_0 \\ &\quad + \frac{1}{2}J^{(0)q,r}_i[\epsilon_j\epsilon_q\epsilon_r]_0 + \frac{1}{2}J^{(0)q,r}_j[\epsilon_i\epsilon_q\epsilon_r]_0 + J^{(0)q}_{i,j}[\epsilon_q]_0, \end{aligned} \quad (71)$$

$$\begin{aligned} \partial_t[\epsilon_i\epsilon_j]_2 &= J^{(0)q}_i[\epsilon_j\epsilon_q]_2 + J^{(0)q}_j[\epsilon_i\epsilon_q]_2 + D_i^{(1)}[\epsilon_j]_1 + D_j^{(1)}[\epsilon_i]_1 \\ &\quad + \frac{1}{2}J^{(0)q,r}_i[\epsilon_j\epsilon_q\epsilon_r]_1 + \frac{1}{2}J^{(0)q,r}_j[\epsilon_i\epsilon_q\epsilon_r]_1 + J^{(0)q}_{i,j}[\epsilon_q]_1 \\ &\quad + D_{ij}^{(1)} + \frac{1}{6}J^{(0)q,r,s}_i[\epsilon_j\epsilon_q\epsilon_r\epsilon_s]_0 + \frac{1}{6}J^{(0)q,r,s}_j[\epsilon_i\epsilon_q\epsilon_r\epsilon_s]_0 \\ &\quad + \frac{1}{2}J^{(0)q,r}_{i,j}[\epsilon_q\epsilon_r]_0 + J^{(1)q}_i[\epsilon_j\epsilon_q]_0 + J^{(1)q}_j[\epsilon_i\epsilon_q]_0, \end{aligned} \quad (72)$$

$$\partial_t[\epsilon_i\epsilon_j\epsilon_k]_0 = D_{ij}[\epsilon_k]_0 + (i \leftrightarrow j \leftrightarrow k) + J^{(0)q}_i[\epsilon_j\epsilon_k\epsilon_q]_0 + (i \leftrightarrow j \leftrightarrow k), \quad (73)$$

$$\begin{aligned} \partial_t[\epsilon_i\epsilon_j\epsilon_k]_1 &= D_{ij}[\epsilon_k]_1 + (i \leftrightarrow j \leftrightarrow k) + J^{(0)q}_i[\epsilon_j\epsilon_k\epsilon_q]_1 + (i \leftrightarrow j \leftrightarrow k) \\ &\quad + D_i^{(1)}[\epsilon_j\epsilon_k]_0 + (i \leftrightarrow j \leftrightarrow k) + D_{ijk} + \frac{1}{2}J^{(0)q,r}_i[\epsilon_j\epsilon_k\epsilon_q\epsilon_r]_0 \\ &\quad + J^{(0)q}_{i,j}[\epsilon_k\epsilon_q]_0 + (i \leftrightarrow j \leftrightarrow k), \end{aligned} \quad (74)$$

$$\partial_t[\epsilon_i\epsilon_j\epsilon_k\epsilon_l]_0 = D_{ij}[\epsilon_k\epsilon_l]_0 + (i \leftrightarrow j \leftrightarrow k \leftrightarrow l) + J^{(0)q}_i[\epsilon_j\epsilon_k\epsilon_l\epsilon_q]_0 + (i \leftrightarrow j \leftrightarrow k \leftrightarrow l). \quad (75)$$

Here, $(i \leftrightarrow j \leftrightarrow k)$ denotes a sum over the previous term over all cyclic permutations of the indices i, j and k , and similarly for four indices. For example, we have $D_{ij}[\epsilon_k]_1 + (i \leftrightarrow j \leftrightarrow k) = D_{ij}[\epsilon_k]_1 + D_{ki}[\epsilon_j]_1 + D_{jk}[\epsilon_i]_1$. We find that for deterministic initial conditions $[\epsilon_i]_0 = [\epsilon_i]_2 = [\epsilon_i\epsilon_j]_1 = [\epsilon_i\epsilon_j\epsilon_k]_0 = 0$ holds for all times. The moments in molecule numbers n_i can now be obtained from the moments of the ϵ_i variables by using the ansatz in Equation (47). For the mean and covariance this leads to

$$\left\langle \frac{n_i}{\Omega} \right\rangle = \phi_i + \Omega^{-1}[\epsilon_i]_1 + O(\Omega^{-2}), \quad (76)$$

$$\left\langle \frac{n_i}{\Omega} \frac{n_j}{\Omega} \right\rangle - \left\langle \frac{n_i}{\Omega} \right\rangle \left\langle \frac{n_j}{\Omega} \right\rangle = \Omega^{-1}[\epsilon_i\epsilon_j]_0 + \Omega^{-2}([\epsilon_i\epsilon_j]_2 - [\epsilon_i]_1[\epsilon_j]_1) + O(\Omega^{-3}). \quad (77)$$

To order Ω^0 , Equation (76) corresponds to the macroscopic rate equations, while Equation (77) to order Ω^{-1} corresponds to the LNA estimate for the covariance. Including terms of order Ω^{-1} in Equation (76) gives the leading order corrections to the mean given by the rate equations. To this order, the system size expansion equations have been called ‘‘effective mesoscopic rate equations’’ in the literature [107]. The next leading order corrections to the covariance, which have been called ‘‘Inverse Omega Square’’ in the literature [106], are obtained by keeping terms of order Ω^{-2} in Equations (77).

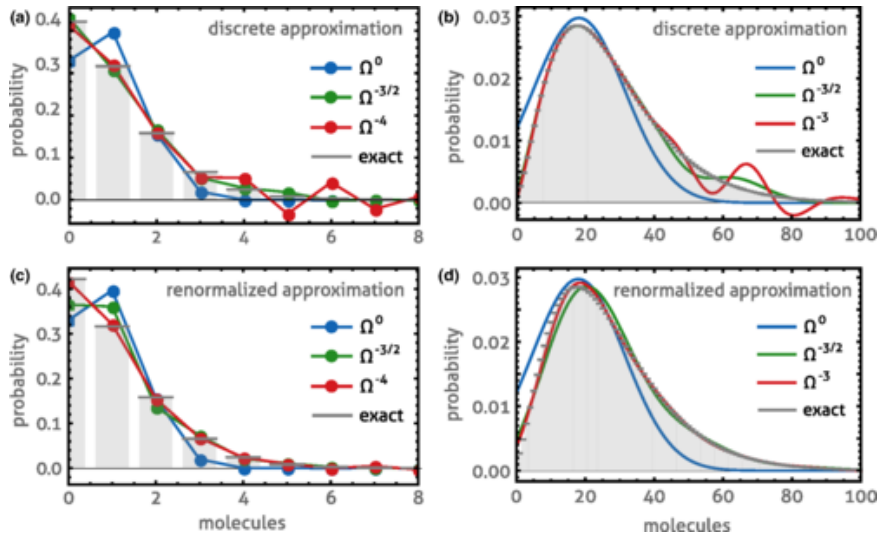


Figure 5: **Distribution of non-linear birth-death process.** The figures show the steady-state distribution for a non-linear birth-death process for two parameter sets, one corresponding to a small mean number of molecules (left panel) and one corresponding to a large mean number of molecules (right panel). Results are shown for different orders of the system size expansion with discrete support (top panel) as well as the renormalised approximation (see main text for a definition). The studied system corresponds to a birth-death process where the propensity of the death process is of Michaelis-Menten type. Such a system is obtained when reducing a Michaelis-Menten system under time-scale separation. This will be demonstrated in Section 4.6.3. The figure is taken from [108] with permission kindly granted by the publisher.

4.2.4 Properties and recent developments

As long as one is only interested in the LNA or higher order corrections to the moments, rather than the distributions of higher order truncations, the system size expansion amounts to the solution of finite sets of ordinary differential equations. It is thus generally significantly more efficient than stochastic simulations of the CME or of the CLE.

Since the system size expansion is an expansion around the deterministic mean, it cannot be used for deterministically multistable systems, i.e., systems whose rate equations have two or more positive stable steady states, unless one is only interested in the short-time behaviour of a process (for example for the purpose of inference from time-series data, c.f. Section 6).

In many cases, the lowest order of the system size expansion, i.e., the LNA, already gives remarkably accurate results and has been used successfully in various applications [109, 110, 111, 112, 113]. However, for some systems it gives rise to significant deviations from the exact result [114]. Higher order approximations beyond the LNA have frequently been found to be highly accurate for such systems. Examples include oscil-

lations in networks of coupled autocatalytic reactions [115, 116], predator-prey systems [117] and circadian oscillators [118].

We note that for truncations of higher orders than zero, the resulting PDEs involve higher order derivatives than two and hence have no probabilistic interpretation due to non positive-definite solutions [119]. Moreover, the solutions to the PDEs are generally not longer known in this case. However, for systems with only one species a general solution to all orders has recently been derived [108]. A similar approach has been developed in [120] for discrete-time models in neuroscience. Note that the variables described by the system size expansion are typically assumed to be continuous. In [108] it has been found that if the support is assumed to be continuous, the first few leading orders of the system size expansion lead to oscillations in the tail of the Poisson distribution for a linear birth-death process. In contrast, if one assumes a discrete support, the distribution is captured accurately. In the same work, a modified system size expansion has been employed. Instead of expanding the variables of the CME around the deterministic mean value given by the rate equations, one expands the variables around the mean value given by the system size expansion itself to the considered order (which is not known a priori). The authors call this method “renormalised approximation”. This approach has been found to give an improved approximation of the distribution for a non-linear process as shown in Figure 5.

4.3 Moment closure approximations

Another popular class of methods that approximate the first few moments of a process are so-called *moment closure approximations*. We present a popular class of such methods here that are particularly easy to derive and implement.

4.3.1 General formulation

In Section 3.2.2 we showed that the CME gives rise to ordinary differential equations for the moments of a process. For linear reaction systems, these equations can be solved numerically up to a certain desired order. For non-linear systems, however, we found that the moment equations of a certain order couple to higher order equations, leading to an infinite hierarchy of equations which can hence not be solved directly. Moment closure approximations truncate this infinite set of equations at a certain order M in some approximate way.

One popular class of moment closure approximations close the moment equations by expressing all moments above a certain order M as functions of lower-order moments. One way to achieve this is for example by assuming the distribution of the system to have a particular functional form, for example a normal distribution. This decouples the equations of the moments up to order M from higher-order moments, and hence leads to a finite set of coupled ordinary differential equations which can then be solved (numerically). We refer to such a moment closure as “ M th order moment closure”.

Let

$$y_{i_1, \dots, i_k} = \langle n_{i_1} \dots n_{i_k} \rangle, \quad (78)$$

$$z_{i_1, \dots, i_k} = \begin{cases} \langle (n_{i_1} - y_{i_1}) \dots (n_{i_k} - y_{i_k}) \rangle & \text{if } k \geq 2, \\ y_{i_1} & \text{if } k = 1, \end{cases} \quad (79)$$

$$c_{i_1, \dots, i_k} = \partial_{s_{i_1}} \dots \partial_{s_{i_k}} g(s_1, \dots, s_N) |_{s_1, \dots, s_N=0}, \quad (80)$$

denote raw moments, central moments and cumulants of order k , respectively. We call y_{i_1, \dots, i_k} a “diagonal moment” if $i_l = i_m$ for all $l, m \in \{1, \dots, k\}$, and a “mixed moment” otherwise, and similarly for central moments and cumulants. The function $g(s)$ in Equation (80) is the *cumulant generating function* and is defined as

$$g(s_1, \dots, s_N) = \log \langle \exp(s_1 n_1 + \dots + s_N n_N) \rangle. \quad (81)$$

We note that all three types of moments are respectively invariant under permutations of their indices. Therefore, only one representative combination of each permutation class has to be considered. Taking this symmetry into account significantly reduces the number of variables and moment equations.

Some popular moment closure methods can be defined as

- “Normal moment closure” [121, 122, 57, 123] (also called “cumulant neglect moment closure” in the literature): all cumulants above order M are set to zero, i.e.,

$$c_{i_1, \dots, i_k} = 0, \quad \text{for } k > M. \quad (82)$$

Note that the cumulants of order higher than two are zero for a normal distribution, hence the name “normal moment closure”. We refer to the normal moment closure approximations to second and third order as “2MA” and “3MA”, respectively.

- “Poisson moment closure” [124]: the cumulants of a one-dimensional Poisson distribution are all equal to the mean value. We assume here a multivariate distribution to be a product of uni-variate Poisson distributions. Accordingly, for the Poisson moment closure approximation of order M we set all diagonal cumulants to the corresponding mean and all mixed cumulants to zero, i.e.,

$$c_{i_1, \dots, i_k} = y_i, \quad \text{for } k > M \quad \text{and} \quad i_1, \dots, i_k = i, \quad \text{for some } i \in \{1, \dots, N\}, \quad (83)$$

$$c_{i_1, \dots, i_k} = 0, \quad \text{for } k > M \quad \text{and} \quad i_m \neq i_n \quad \text{for some } m, n \in \{1, \dots, K\}. \quad (84)$$

- “Log-normal moment closure” [125]: let \mathbf{m} and \mathbf{S} be the mean vector and covariance matrix of a multi-dimensional normal random variable. Then the logarithm of this random variable has a multivariate log-normal distribution whose moments can be expressed in terms of \mathbf{m} and \mathbf{S} as [126]

$$y_{i_1, \dots, i_k} = \exp \left(\mathbf{v}^T \mathbf{m} + \frac{1}{2} \mathbf{v}^T \mathbf{S} \mathbf{v} \right), \quad \text{for } k > M, \quad (85)$$

where $\mathbf{v} = (g_1, \dots, g_N)$, and g_m is the number of i_j 's having the value m . This allows one to express \mathbf{m} and \mathbf{S} in terms of the first two moments y_i and $y_{i,j}$ which then in turn allows one to express higher-order moments in terms of y_i and $y_{i,j}$, too.

- “Central-moment-neglect moment closure” (also called “low dispersion moment closure” in the literature) [127]: all central moments above order M are set to zero:

$$z_{i_1, \dots, i_k} = 0, \quad \text{for } k > M. \quad (86)$$

- “Derivative matching”: the idea of this method is to express moments above order M by lower order moments in such a way that the time derivatives of the moments of the closed system approximate the time derivatives of the moments up to order M of the exact system at some initial time point. In [128] a method is derived that allows one to produce the corresponding expressions.

Each of the closure methods allows one to express all raw moments above a specified order M in terms of lower order moments and hence to close the moment equations. Note that third order central moments and third order cumulants are identical. Therefore, whenever the moment equation of up to order two do not depend on moments of order higher than three, the normal and central-moment-neglect moment closure are identical. Similarly, the normal and Poisson moment closure approximations can be equivalent for certain systems.

Example. As an example, consider again the gene system in Figure 3. The moment equations up to order two for this system are given in Equations (21)-(25). To close these we need to express the third order moments $y_{1,1,2}$ and $y_{1,2,2}$ in terms of lower order moments. We do so by means of the normal moment closure defined in Equation (82) to second order, i.e., we set the third order cumulants $c_{1,1,2}$ and $c_{1,2,2}$ to zero:

$$c_{1,1,2} = y_{1,1,2} - 2y_1 y_{1,2} - y_2 y_{1,1} + 2y_2 y_1^2, \quad (87)$$

$$c_{1,2,2} = y_{1,2,2} - 2y_2 y_{1,2} - y_1 y_{2,2} + 2y_1 y_2^2. \quad (88)$$

Setting these to zero and rearranging gives

$$y_{1,1,2} = 2y_1 y_{1,2} + y_2 y_{1,1} - 2y_2 y_1^2, \quad (89)$$

$$y_{1,2,2} = 2y_2 y_{1,2} + y_1 y_{2,2} - 2y_1 y_2^2. \quad (90)$$

Using these expressions in Equations (21)-(25) the equations decouple from higher order moments. We give here the resulting equations in terms of central rather than raw

moments (c.f. Equation (79))

$$\partial_t z_1 = -\frac{k_2}{\Omega}(z_{1,2} + z_1 z_2) + k_3(1 - z_1), \quad (91)$$

$$\partial_t z_2 = k_1 z_1 - \frac{k_2}{\Omega}(z_{1,2} - z_1 z_2) + k_3(1 - z_1) - k_4 z_2, \quad (92)$$

$$\partial_t z_{1,1} = \frac{k_2}{\Omega}(-2z_2 z_{1,1} - 2z_1 z_{1,2} + z_{1,2} + z_1 z_2) + y_{1,2}) + k_3(-2z_{1,1} - z_1 + 1), \quad (93)$$

$$\begin{aligned} \partial_t z_{1,2} = & k_1 z_{1,1} + \frac{k_2}{\Omega}(-z_2 z_{1,1} - z_1 z_{1,2} - z_2 z_{1,2} - z_1 z_{2,2} + z_{1,2} + z_1 z_2) \\ & + k_3(z_{1,2} - z_{1,1} - z_1 + 1) - k_4 z_{1,2}, \end{aligned} \quad (94)$$

$$\begin{aligned} \partial_t z_{2,2} = & k_1(2z_{1,2} + z_1) + \frac{k_2}{\Omega}(-2z_1 z_{2,2} - 2z_2 z_{1,2} + z_{1,2} + z_1 z_2) \\ & + k_3(-2z_{1,2} - z_1 + 1) + k_4(-2z_{2,2} + z_2). \end{aligned} \quad (95)$$

Note that the equations do not depend on third or higher order moments and hence are closed and can be integrated.

4.3.2 Properties and recent developments

Moment closures are a popular class of approximations of the moments of the CME with many useful properties. First of all, they are easy to derive and implement. Moreover, since they amount to solving a finite set of ordinary differential equations no ensemble averaging is needed, which means they are computationally significantly more efficient than stochastic simulations of the CME or the CLE, and comparable to the system size expansion. However, this computational efficiency comes at a cost: moment closure methods only give approximate moments and not an approximation of the process or distributions. One advantage of moment closures over the system size expansion is that they can be applied to deterministically multistable systems. However, care must be taken for such system since moment closures can lead to unphysical results, see below.

In contrast to the system size expansion, moment closures are not an expansion in any small parameter, and one can therefore generally not expect that increasing the closure order leads to a higher accuracy. Numerical case studies suggest that this is often indeed the case, however [129, 130, 131]. One explanation for this was given in [129] where it has been shown that increasing orders of the normal moment closure agree with increasing orders of the system size expansion in the limit of large volumes. These results can to some extent be generalised to other moment closure methods. For monostable systems, one may therefore expect higher order closures to become more accurate for large enough system sizes. For small system sizes on the other hand, the closures can generally not be expected to converge.

Moreover, since moment closure approximations are an *ad hoc* approximation, it is not even clear if they always give physically meaningful results. As it turns out, this is not always the case, but they sometimes give rise to unphysical behaviour, such as negative mean values, negative variances or negative higher central moments, as well as diverging trajectories [124, 131]. In numerical case studies, it has been found for several non-linear reaction systems that the normal, log-normal, Poisson and central-moment neglect

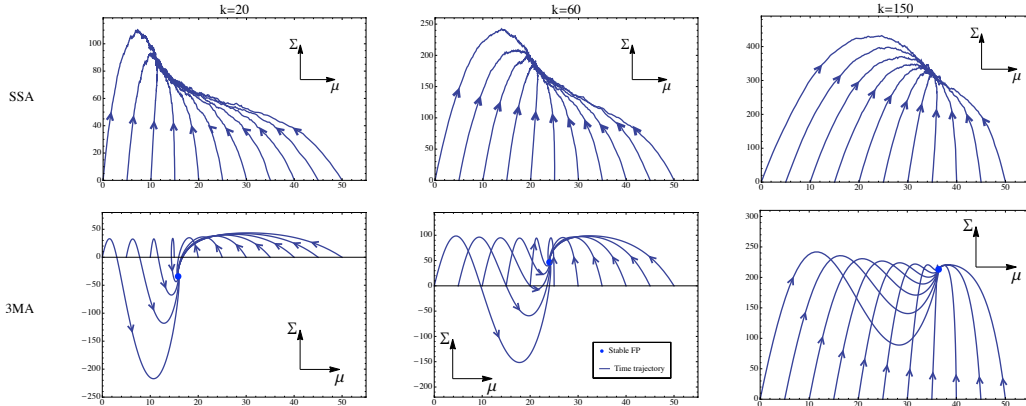


Figure 6: **Time trajectories of moment closure with negative variance.** The figures show time trajectories in the mean-variance ($\mu - \Sigma$) plane of a protein P in a system with bursty protein production. The reactions are $\emptyset \rightarrow mP$, $m \in \mathbb{N} \setminus \{0\}$, $P + P \rightarrow \emptyset$, where the burst size m is a geometrically distributed random number. The top panel shows results obtained from exact stochastic simulations (SSA) and the lower panel results obtained using the normal moment closure to third order (3MA). The parameter k corresponds to the system size. We find that the results from stochastic simulations converge to a positive fixed point for all system sizes and all initial conditions, and that the mean and variance remain non-negative for all times. The 3MA, in contrast, converges to a fixed point with negative variance for the smallest system size $k = 20$. For $k = 60$, the fixed point is now positive, but some initial conditions still lead to trajectories that have a negative variance for some time and are hence physically not meaningful. Only for the largest system size $k = 150$ do all trajectories lead to non-negative trajectories. The figure is taken from [131] with permission kindly granted by the publisher.

moment closure show such unphysical behaviour for system volumes *below a certain critical volume* [131, 132]. An example of such unphysical behaviour is shown in Figure 6. This may be expected, since smaller volumes correspond to stronger non-linearity and hence stronger fluctuations. However, it has been found for deterministically oscillatory and deterministically multistable systems, that the moment closure methods give rise to non-physical oscillations and non-physical multistability, respectively, for system volumes *above a certain critical volume* [131, 132]. In [132], the normal moment closure was found to give physically meaningful results for larger ranges of system volumes than the other three methods. In terms of accuracy, the four methods performed similar to each other. We emphasise that these were numerical case studies, and it remains open to what extent these results can be generalised to arbitrary reaction systems. In [133] and [134], for example, the log-normal moment closure was found to be significantly more accurate than the normal and central-moment-neglect moment closure for several reaction systems.

4.4 Construction of distributions from moments

The system size expansion and moment closure approximations generally only provide approximations of the moments of a process. Suppose we are however interested in approximating distributions of a process. This can be achieved by running Monte Carlo simulations using the SSA or simulations of the chemical Langevin equation. These methods are computationally quite expensive, however. An alternative approach is to first compute approximations of the moments up to a certain order by means of the system size expansion or moment closure methods, and subsequently constructing a distribution from these moments. If the second step can be done efficiently, this can provide a significantly more efficient method than stochastic simulations since one avoids ensemble averaging. One popular method to construct distributions from moments relies on the *maximum entropy principle*.

For simplicity, assume we have a one-dimensional problem with discrete-valued variable n and that we have approximate values for the first K moments μ_1, \dots, μ_K , for example obtained by means of the system size expansion. The goal is to construct a distribution $P(n)$ matching these moments. The entropy H of a distribution $P(n)$ is defined as [135]

$$H(P) = - \sum_n P(n) \log(P(n)). \quad (96)$$

The maximum entropy method aims at finding a distribution $P(n)$ in a certain family of distributions that maximises the entropy. In our case, the family of distributions is given by the constraint that the first K moments of the distribution should match μ_1, \dots, μ_K . This is a non-linear constrained optimisation problem, which can be solved by means of Lagrange multipliers as follows. We define the Lagrangian

$$\mathcal{L} = H(P) - \sum_{i=0}^K \lambda_i \left(\sum_n n^i P(n) - \mu_i \right). \quad (97)$$

Optimising \mathcal{L} with respect to P and the λ_i corresponds to solving the constrained optimisation problem. The variation of \mathcal{L} with respect to $P(n)$ reads

$$\frac{\delta \mathcal{L}}{\delta P(n)} = -\log(P(n)) - 1 - \sum_{i=0}^K \lambda_i n^i. \quad (98)$$

Setting this to zero we obtain

$$P(n) = \frac{1}{Z} \exp\left(-\sum_{i=1}^K \lambda_i n^i\right), \quad (99)$$

$$Z = \exp(1 + \lambda_0) = \sum_n \exp\left(-\sum_{i=1}^K \lambda_i n^i\right),$$

with normalisation constant Z . Inserting Equation (99) into Equation (97) one obtains an unconstrained optimisation problem which can be efficiently solved using standard




Figure 7: **Construction of distribution via maximum entropy.** The figure shows the marginal steady-state distribution of a protein in a chemical reaction system with bursty protein production. The grey area shows the exact result computed using the SSA, while the coloured lines show the results obtained using the central-moment-neglect moment closure to order $K = 3$ (red) and $K = 5$ (blue) in combination with maximum entropy. We observe that the approximations accurately capture the skewed distribution, with increasing accuracy at higher order. The figure is taken from [137].

numerical methods. The maximum entropy method for the construction of the marginal distributions of single species has recently been used in combination with moment closure methods and the system size expansion in [136, 137]. Figure 7 shows the result for one example system.

A related approach has been developed in [138], where the maximum entropy principle has been used directly to close the moment equations. For time-dependent approximations this method is computationally expensive because the moment equations have to be solved iteratively in small time steps, and in each time step a multi-variate optimisation problem has to be solved. For steady-state approximations, however, the method was found to be efficient and accurate.

4.5 Software

Several software packages for using the discussed approximation methods, as well as for deterministic rate equations and stochastic simulation algorithms, are freely available. For exact stochastic simulations, for example, available packages include the Java-based program Dizzy [139], stand-alone COPASI [140], stand-alone StochKit [141] and the python package StochPy [142]. The system size expansion for orders beyond the LNA is implemented in the stand-alone package iNA [143], which also allows one to perform exact stochastic simulations. Various moment closure approximations are available in the Matlab toolbox StochDynTools [144], and the normal moment closure in Python package MomentClosure [145]. These packages require programming knowledge and are only applicable to mass-action propensity functions. The mathematica package MOCA [132]

extends the applicability to non-polynomial and time-dependent propensity functions and does not require any programming skills. Similarly, the python package MEANS [146] extends moment closure methods to non-mass-action propensity functions. The Matlab package CERENA [147] implements several of the mentioned methods, including exact stochastic simulations, the system size expansion and moment closure methods. For a detailed overview of available software packages and their capabilities see Figure 1 in [147].

4.6 Other approximations

One reason why the CLE, the system size expansion and moment closure approximations discussed above are so popular is that they are easy to implement, do not require any pre-knowledge of the system (except monostability for the system size expansion), are generally efficient computationally, and often give accurate approximations. The three methods have been frequently applied successfully in the literature. However, there are many scenarios where the three methods give quite inaccurate results. In particular if one or several of the species occur in very low copy numbers, the three methods often perform poorly. A large number of other approximation methods have been developed in the literature, and we give an overview of these methods here. Many of these methods are more sophisticated but only apply to certain classes of systems and require pre-knowledge and/or fine-tuning. As we shall see, the CLE, the system size expansion or moment closure approximations form building blocks of many of these methods.

4.6.1 State space truncation

In Section 3.2 we saw that the CME can be solved exactly by matrix exponentiation whenever the state space is finite (c.f. Equation (32)). However, for many chemical systems of interest the state space is infinite and the solution in Equation (32) can not be computed. The idea of the *finite state projection algorithm* is to truncate the state space to a finite subspace and use matrix exponentiation to obtain an approximation of the distribution on this subspace [148]. It also provides an estimate of the error, which can be systematically reduced by increasing the truncated space.

While this is an efficient method in some cases, often a large truncated space has to be chosen to achieve a reasonable accuracy, making matrix exponentiation intractable, despite recent progress on numerical algorithms, see [66, 16] for overviews. Similarly, sometimes the state space of a system is finite but too large to compute the matrix exponential. Accordingly, several modified versions of the finite state projection algorithm have been developed, see [149] for a review.

4.6.2 Tau-leaping

Tau-leaping methods are approximate ways of simulating a chemical network with the goal of being more efficient than exact stochastic simulations [150]. The basic idea is to “leap” along time in certain steps τ during which several reaction events occur, thereby avoiding the simulation of each single reaction event as exact stochastic simulations do. The time step has to be chosen small enough such that the propensity functions do

not change significantly during one time step. In that case different reactions become independent of each other, and the number of reaction events $N_r(\tau; \mathbf{n}, t)$ of the r^{th} reaction during the time interval τ given that the state at time t is \mathbf{n} becomes a Poisson random variable $\mathcal{P}(f_r(\mathbf{n})\tau)$ with mean $f_r(\mathbf{n})\tau$. Accordingly, the tau-leaping algorithm updates the state $\mathbf{n}(t)$ of the system iteratively as

$$\mathbf{n}(t + \tau) = \mathbf{n}(t) + \sum_{r=1}^R \mathbf{S}_r \mathcal{P}(f_r(\mathbf{n})\tau), \quad (100)$$

where \mathbf{S}_r is the r^{th} row of the stoichiometric matrix. If many reactions are happening per time step the algorithm is more efficient than exact simulations. Increasing τ increases the algorithm’s efficiency but obviously decreases its accuracy. If a system contains species with very low particle numbers τ may have to be chosen smaller than the average inter-reaction time to ensure approximately constant propensity functions, which would lead to the algorithm becoming less efficient than exact simulations. Roughly speaking, tau-leaping therefore works best for systems with not too small average molecule numbers.

Despite its simplicity, implementing the tau-leaping method bears several difficulties. Most importantly, the trade-off between accuracy and efficiency makes the choice of the step size τ non-trivial. A bad choice can lead to highly inefficient or inaccurate results and also to other problems, such as negative molecule numbers. Recent years have therefore seen a wide range of studies developing modified versions of the tau-leaping method that aim at solving these problems, including implicit methods [151, 152], binomial methods [153, 154, 155, 156], multinomial methods [157], K/R-leaping [158, 156] and a post-hoc correcting method [159]. See [8, 9, 160] for overviews of these methods.

4.6.3 Time-scale separation


Many biochemical systems involve processes with highly varying time scales. Given a certain separation of time scales in a system, it is often possible to derive a reduced model of the system which either allows more efficient simulations or more accurate approximations. Such methods have been well-known for deterministic ordinary differential equation models of biochemical kinetics for several decades [161] but have only more recently been developed for stochastic methods [162]. We start here by describing the deterministic setting.

The deterministic case. In the following we describe two popular methods for reducing deterministic rate equations under time-scale separation. The first method assumes that the chemical reactions in a given system can be divided into “slow” and “fast” *reactions*, i.e., reactions that happen very infrequently or very frequently on a certain time scale of interest. In this case it is sometimes possible to derive a reduced model by eliminating the fast reactions, typically by assuming that certain reactions balance each other. Such methods are often called *quasi-equilibrium approximations* (QEA) [163].

The second method separates a system into “slow” and “fast” *species*, rather than reactions. The fast species are assumed to be asymptotically in steady state on the time scale of the slow species. This is known as the *quasi-steady-state approximation*

(QSSA). The idea is to eliminate the fast species from the system and to include their steady-state effect on the slow species.

Example. Let us illustrate the QEA and QSSA in the deterministic setting by means of a simple example, the well-studied Michaelis-Menten system:



A substrate S reversibly binds to an enzyme E to form the substrate-enzyme complex C , from which a product molecule P becomes catalysed. The corresponding rate equations are

$$\begin{aligned} \partial_t \phi_S &= -k_1 \phi_S \phi_E + k_2 \phi_C, \\ \partial_t \phi_E &= -k_1 \phi_S \phi_E + (k_2 + k_3) \phi_C, \\ \partial_t \phi_C &= k_1 \phi_S \phi_E - (k_2 + k_3) \phi_C, \\ \partial_t \phi_P &= k_3 \phi_C, \end{aligned} \quad (102)$$

where ϕ_S, ϕ_E, ϕ_C and ϕ_P denote the concentrations of S, E, C and P , respectively. Let us first make the assumption that the binding and unbinding of the substrate and the enzyme in Equation (101) are fast reactions and that the catalysis reaction is slow. This is the case whenever $k_3 \ll k_2$. On the time scale of the slow reaction, this means that the two fast reactions balance each other, i.e.,

$$k_1 \phi_S \phi_E \approx k_2 \phi_C. \quad (103)$$

This corresponds to the *quasi-equilibrium approximations* (QEA). Note that the total concentration E^0 of enzyme molecules is conserved, i.e., $\phi_E + \phi_C = E^0$. Using this to eliminate ϕ_E in Equation (103) one obtains $\phi_C = (\phi_S E^0) / (\phi_S + K_M)$, where we defined the *Michaelis-Menten constant* $K_M = k_2 / k_1$. Let $v = \partial_t \phi_P = k_3 \phi_C$ be the production rate of the product P . According to Equation (102) this becomes

$$v = \frac{k_3 \phi_S E^0}{\phi_S + K_M}, \quad (104)$$

which is the well-known Michaelis-Menten equation [164].

Next, we illustrate the *quasi-steady-state approximation* (QSSA). To this end, we assume that the complex is approximately in steady state:

$$\partial_t \phi_C \approx 0. \quad (105)$$

It can be shown that this is a reasonable approximation whenever $\phi_S + K'_M \gg \phi_E$ with $K'_M = (k_2 + k_3) / k_1$ [165]. This corresponds to the complex C and free enzyme E being fast species, and the substrate S being a slow species.

Using Equation (105) in Equation (102) and the conservation law $\phi_E + \phi_C = E^0$ we obtain

$$v = \frac{k_3 \phi_S E^0}{\phi_S + K'_M}. \quad (106)$$

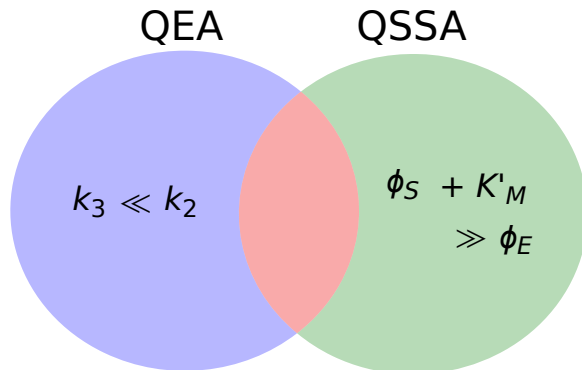


Figure 8: **Visualisation of different time-scale separation regimes.** The figure visualises the two different time-scale separation regimes for the Michaelis-Menten reaction system with reactions in Equation (101). The blue area represents the regime of parameter space where only the quasi-equilibrium assumption holds, the green area where only the quasi-steady-state assumption holds, and the red area where both assumptions hold.

This is also called *Briggs-Haldane kinetics* [166]. Note that Equation (106) has the same form as Equation (104) but with a different constant K'_M , showing that the QEA and QSSA are generally not equivalent. Depending on the parameters, only one of the two conditions (or none) may be satisfied. This is illustrated in Figure 8.

The stochastic case. For stochastic systems described by the CME, reductions based on time-scale separations are not as straightforward as in the deterministic case. Ideally, one would like to derive a reduced CME and/or a stochastic simulation algorithm for a reduced system. As it turns out, under the quasi-equilibrium assumption of fast and slow reactions it is indeed possible to derive a reduced CME [167]. In contrast, the quasi-steady-state assumption of fast and slow species does not necessarily allow to derive a reduced master equation since it leads to a non-Markovian process for the fast species [168]. One can use a reduced master equation with non mass-action propensity functions that corresponds to the reduced deterministic system under the quasi-steady-state assumption. This heuristic master equation has been found to be accurate for some examples [169, 170, 171]. However, even if the quasi-steady-state assumption is valid for the deterministic system, the reduced master equation can be highly inaccurate [172, 173]. The validity of a reduction in the deterministic case does not generally imply a valid reduction in the stochastic case. The relationship between the two descriptions with respect to time-scale separation approximations has for example been studied in [173]. In [174] a reduced description of stochastic models under the quasi-steady-state assumption was derived by means of the linear noise approximation.

It is however possible to derive reduced descriptions based on quasi-steady-state assumptions for a stochastic system if one requires stronger conditions on a system than the conditions needed to reduce the deterministic system. In [175], for example, reactions are split into fast and slow reactions, and slow species are defined as those involved in slow reactions only and fast species as those participating in at least one fast reaction and

any number of slow reactions. As mentioned before, the fast variables conditioned on the slow ones are not Markovian. The authors in [175] solve this problem by approximating the fast process by a virtual fast process that is affected by fast reactions only and therefore Markovian.

A large number of other approximation methods for CME type systems based on time-scale separations have been developed in recent years, see for example [167, 176, 177, 178, 179, 180, 181, 182, 183, 184, 185]. Most of these methods split up the full CME into a CME for the fast variables conditioned on the slow reactions/species, and a CME for the slow reactions/variables with marginalised fast variables. On the time scale of the slow dynamics, the conditional CME of the fast dynamics is assumed to quickly reach steady state. Therefore, the CME of the slow dynamics only depends on the steady-state distribution of the fast dynamics. Most methods rely on SSA simulations of the slow dynamics and assume that the fast dynamics is in steady state [176, 175, 177, 178, 179, 183].

For some systems, the reduced CMEs under time-scale separation have been solved exactly, for example for gene expression [186, 187, 188]. In [189] a perturbative expansion based on time-scale separation for a two-stage gene expression model is derived that allows one to systematically include higher order corrections to the solution.

The studies mentioned so far mainly rely on stochastic simulations of the reduced equations or on analytic solutions for special cases. An alternative approach is to combine time-scale separation reductions with other approximations, such as approximating the fast variables by the chemical Langevin equation [190, 191, 183], deterministic rate equations [190], or moment closure methods [180]. Other methods combine time-scale separation with the finite state projection algorithm [192] (c.f. Section 4.6.1), tau-leaping [193], or the linear noise approximation [194, 174, 195]. An adiabatic approximation derived from a stochastic path integral description has been developed in [196]. Other methods derive reduced models for systems that allow a partition in more than two typical time scales, see for example [197, 179]. In [198] a conditional linear noise approximation for gene expression systems with finite, slow promoter states has been derived.

4.6.4 Hybrid methods

In many biochemical reaction networks of practical interest no time-scale separation assumptions apply. However, often some species occur in low and others in high copy numbers, which motivates the combination of different simulation and/or approximation methods for these two groups of species, similar to the time-scale separation case. One possibility would be to partition species into a discretely modelled group and a continuously modelled group, and accordingly reactions into discrete reactions involving discrete species and continuous reactions that do not involve discrete species. In this case species that are involved in both continuous and discrete reactions describe diffusion-jump processes. Such methods are broadly referred to as *hybrid methods*. Note that many of the methods discussed in Section 4.6.3 fall under this definition. We give here an overview of hybrid methods that do not rely on separation of time scales, although a clear distinction between the two cases is often not possible.

Consider the case where we split the species of a system into a low abundance group which we model discretely using the SSA and a high abundance group which we model

continuously using deterministic rate equations or chemical Langevin equations. Accordingly, we group reactions into discrete reactions involving discrete species and continuous reactions that do not involve discrete species. Between two discrete reaction events, the continuous variables follow rate equations or chemical Langevin equations which can be solved numerically or simulated in a standard way. The propensity functions of the discrete reactions, however, may depend on the continuous variables and hence depend on time in such a hybrid approach. This is akin to the case of extrinsic noise discussed in Section 3.3.1, where a discrete system's propensities were assumed to depend on some external stochastic process. Correspondingly, the discrete system of our hybrid system can not be simulated by a standard SSA. As in the extrinsic noise case, one possibility is to numerically integrate the time-dependent propensity functions over time until a target value is reached and the corresponding reaction occurs. The discrete system is then updated accordingly which may also change the propensity functions of the continuous reactions.

A key challenge of such an approach is to decide which species should be modelled discretely and which continuously, in particular since some species may fluctuate strongly during a simulation. More sophisticated algorithms therefore partition species and reactions adaptively during simulation. Many different methods addressing these and other issues have been developed in the literature, see for example [199, 200, 201, 202, 203, 204, 205, 206, 207, 208, 209, 210, 211]. They differ mainly in how the partitioning is conducted and in the simulation methods for the resulting reduced system. Simulation-based approaches include combination of the SSA for the low abundance species with tau-leaping [204], chemical Langevin equations [202, 206, 208, 209, 212], the chemical Fokker-Planck equation [210] and deterministic rate equation approximations [199, 200, 201, 205, 211] for the abundance species. Other approaches split species into more than two sets [203].

Since many of these methods are heuristic, it is not straightforward to assess their performance or to prove their convergence to the exact system in some limit. In some cases, this is possible, however. For example, error bounds for some hybrid methods have been derived [213, 214, 208]. In [215] the convergence of different types of hybrid methods to the exact system have been studied, and in [184] criteria have been developed for the convergence of a CME to discrete-deterministic hybrid approximations. An error analysis of various methods has recently been conducted in [216].

While simulation-based hybrid methods are often orders of magnitude more efficient than the standard SSA, they all still rely on stochastic simulations for the low abundance species and some of them also for the high abundance species. They therefore can still become computationally expensive depending on the studied system. Some methods aim at circumventing expensive simulations by applying additional approximations to the reduced systems. In [217], for example, the dynamics of the large abundance species is formulated in terms of moment equations conditional on the low abundance (discrete) species. For non-linear systems these can be closed by means of moment closure approximations (c.f. Section 4.3). This method amounts to the solution of a differential algebraic equation which is difficult to solve. The authors in [217] propose several simulation based algorithms. This method has shown to give accurate approximations to the distributions and moments of some gene systems, but its implementation is non-trivial and requires additional approximations and/or simulations. In [218] a simpler

conditional moment closure has been proposed for two-state gene systems. Here, the conditional moment equations are closed by means of the second-order normal closure or derivative matching (c.f. Section 4.3). This method amounts to solving a coupled set of ordinary differential equations. In [219] a hybrid method is developed where the abundance species is described by the (non-conditional) rate equations. This is a simplification of the previously mentioned methods that model the abundance species by rate equations conditional on the low abundance species. However, this simplification allows one to derive analytic solutions of the CME for the low abundance species for certain systems [219].

5 Comparison of approximation methods

In the previous section we gave a detailed introduction to three popular approximation methods of the CME: the chemical Langevin equation (CLE, Section 4.1), the system size expansion (Section 4.2) and moment closure approximations (Section 4.3). These methods have been successfully used in many applications in the literature, but there exist only very few studies comparing their accuracy. It thus remains unclear how the different methods compare to each other.

In Sections 4.1 - 4.3 we discussed advantages and disadvantages of the different methods. Here, we perform a numerical case study to enable the reader to understand the differences between the methods. First, we study an enzyme reaction system of the Michaelis-Menten type in Section 5.1 which can be viewed as a catalysed degradation of a spontaneously produced protein. We then extend the model in Section 5.2 by including transcription of mRNA and translation of the protein from mRNA, which allows for bursts in protein production. Finally in Section 5.3 we summarise the results and give an overview of advantages and disadvantages of the different methods.

Implementation. Implementation details for the different methods are:

- **CLE:** we simulate the CLE using the Euler-Maruyama algorithm introduced in Section 4.1.1. As pointed out in Section 4.1, the CLE is traditionally defined for real variables, but suffers from the occurrences of square roots of negative expressions for which it is not well-defined. We therefore implement two versions of the CLE here: (i) a real-valued implementation which keeps the variables positive by rejecting Euler-Maruyama steps that lead to negative variables. We term this version “CLE-R”; (ii) the recently proposed complex CLE which is defined for complex variables and has been shown to give real-valued moments in the ensemble average [87]. We term this version “CLE-C”.

The Euler-Maruyama algorithm requires a step size dt for time discretisation. For the simulation of steady-state moments or distributions, we take a total number of M samples from a single long trajectory at time steps separated by Δt . The chosen values for these parameters will be given for each of the examples.

- **System size expansion and stochastic simulation algorithm:** we use the stand-alone software package iNA [143] for both the system size expansion and the stochastic simulation algorithm (Section 3.3). We study the zeroth order system

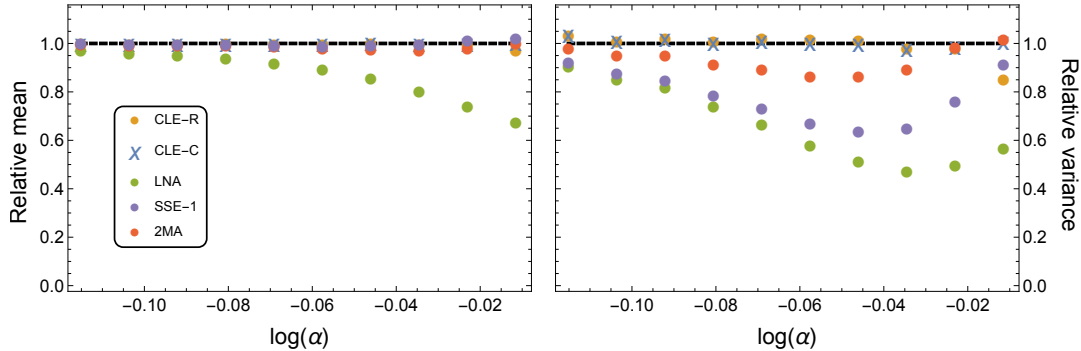


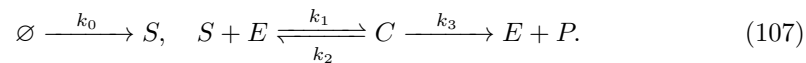
Figure 9: **Steady-state mean and variance of protein for enzyme system in Equation (107)**. The figures show the steady-state mean (left) and variance (right) of the protein S of the enzyme system in Equation (107), as a function of the saturation parameter $\alpha = k_0\Omega/(E^0k_3)$ on logarithmic scale. The values obtained by the approximation methods are normalised by the result obtained from stochastic simulations using the SSA, which means that the dashed black line corresponds to the exact result (up to sampling error). The parameters used are $E^0 = 2500, k_1 = 40, k_2 = 8, k_3 = 60, \Omega = 1$ and we vary k_0 according to the shown values of α . The simulation parameters for the two CLE implementations are $dt = 10^{-5}, \Delta t = 1$ and $M = 10^5$.

size expansion, i.e., the linear noise approximation (termed “LNA”, Section 4.2.2), as well as the first order corrections to the mean and variance (both termed “SSE-1”) given in Equations (76) and (77), respectively.

- **Moment closure approximations:** we study here the second-order normal moment closure (termed “2MA”) introduced in Section 4.3, which is probably the most commonly used one in the literature. For its implementation we use the Mathematica software package MOCA [132].

5.1 Enzymatic protein degradation

Let us consider the well-known Michaelis-Menten system with reactions



We studied this system without the first reaction in Section 4.6.3 in the context of time-scale separations in a deterministic setting. A substrate molecule S is created spontaneously and binds reversibly to a free enzyme molecule E to form a complex C , which catalyses the substrate into a product molecule P . We consider S to be a protein in the following.

Here, we are interested in the accuracy of the different approximation methods as compared to exact stochastic simulations of the corresponding stochastic system. Note

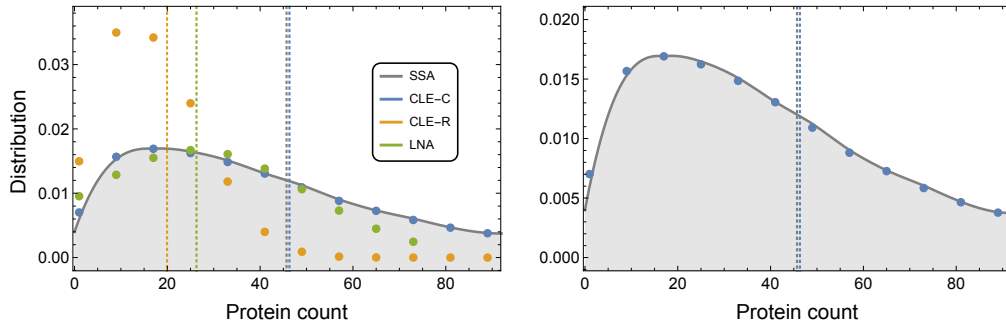


Figure 10: **Steady-state distribution of protein for enzyme system in Equation (107)**. Left: distribution obtained using the CLE-R, CLE-C, LNA and SSA. Right: only CLE-C and SSA. The vertical lines indicate the corresponding steady-state mean values. The mean values of the 2MA and SSE-1 are very close to the SSA result and are not shown here. The used parameters are $E^0 = 60$, $k_0 = 115$, $k_1 = 3.5$, $k_2 = 2$, $k_3 = 2$, $\Omega = 1$. Simulation parameters for the CLE-R and CLE-C are $dt = 10^{-5}$, $\Delta t = 1$ and $M = 10^5$.

that the total number of enzymes E^0 is conserved. The system has a steady state in the protein numbers if and only if $\alpha \equiv k_0\Omega/(E^0k_3) < 1$, which means that the input rate must be smaller than the maximum turnover rate. The parameter α can hence be viewed as a saturation factor.

Figure 9 shows the relative mean and variance of the protein S as a function of α for a system with a large number of total enzymes, $E^0 = 2500$. We find that both the real and complex CLE implementations (CLE-R and CLE-C), the first order system size expansion (SSE-1) and the second-order normal moment closure (2MA) give good approximations for the mean value. Only the LNA, which corresponds to the deterministic rate equations for mean values, shows significant deviations from the exact result. For the variance we observe larger deviations for all methods, with the LNA again being the least accurate. The two CLE implementations show the best performance, being more accurate than the 2MA, which in turn is more accurate than the SSE-1. Note that the CLE-R and the CLE-C give very similar results, meaning that the inclusion of a rejecting boundary has a negligible effect in this case.

Overall, we find that the approximation methods give rise to larger errors for larger values of α , which can be explained as follows. As mentioned before, the system only possesses a steady state for $\alpha < 1$ and becomes unstable for $\alpha > 1$. We therefore expect larger fluctuations for values of α close to unity. Moreover, since most of the enzymes are in the complex state C in this limit, we expect a skewed distribution for the substrate. In the other limit, $\alpha \rightarrow 0$, most enzymes are in the free state E , which reduces the non-linear effect of the bimolecular reaction $S + E \rightarrow C$. We therefore expect the approximation methods to perform well in this limit, and to lead to larger errors for $\alpha \rightarrow 1$. This is exactly what we observe in Figure 9.

Next, we study the steady-state probability distribution of the protein as predicted by the SSA, CLE-R, CLE-C and the LNA. Note that the CLE-C gives complex-valued

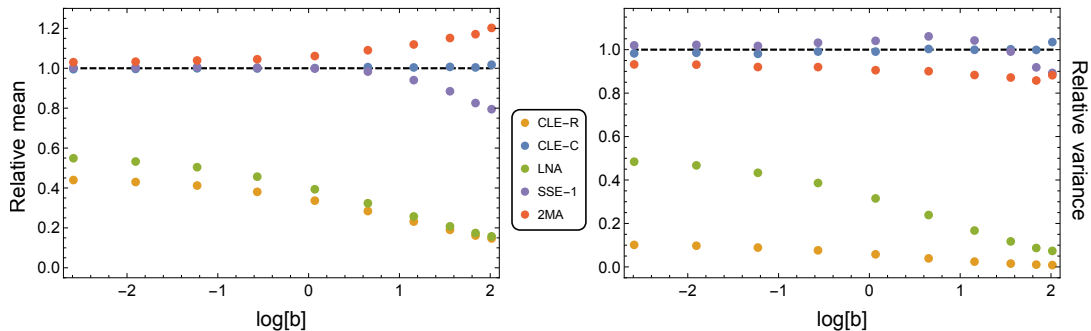
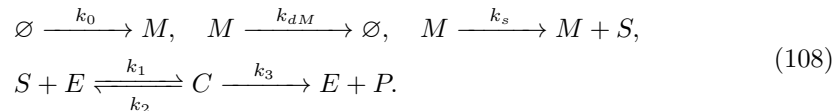


Figure 11: **Steady-state mean and variance of protein for extended system in Equation (108)**. The figures show the steady-state mean (left) and variance (right) of the protein S of the enzyme system in (108), as a function of the burst size $b = k_s/k_{dM}$ on logarithmic scale. The values obtained by the approximation methods are normalised by the result obtained from stochastic simulations (SSA), which means that the dashed black line corresponds to the exact result (up to simulation error). The used parameters are $E^0 = 60, k_s = 1.5, k_1 = 3.5, k_2 = 2, k_3 = 2, \Omega = 1$. k_0 and k_{dM} are varied to vary the burst size $b = k_s/k_{dM}$ such that the average protein production rate $k_0 k_s/k_{dM} = 115$ is held constant. Simulation parameters for the CLE-R and CLE-C are $dt = 10^{-4}, \Delta t = 10$ and $M = 10^4$.

samples. To obtain a distribution in real variables, we take the real parts of these samples. We reduce the total number of enzymes to $E^0 = 60$ to study small molecule number effects. The results are shown in Figure 10. We find that the LNA strongly underestimates the true mean value (obtained using the SSA) and does not reproduce the distribution very accurately. However, surprisingly, the real-valued CLE (CLE-R), does even worse than the LNA, both in terms of mean value and distribution. The complex valued CLE (CLE-C) on the other hand, predicts the true mean value with a negligible error, and even captures the highly skewed distribution very well (right panel of Figure 10). This demonstrates that a naive fixing of the boundary problem of the CLE (CLE-R) can lead to highly inaccurate results.

5.2 Bursty protein production

We now extend the system in Equation (107) by considering the protein S to be produced in a gene expression motif, i.e., via transcription of mRNA M followed by translation to protein:



Depending on the parameters k_{dM} and k_s this system can give rise to bursts in the production of protein S , namely whenever each produced mRNA molecule produces on

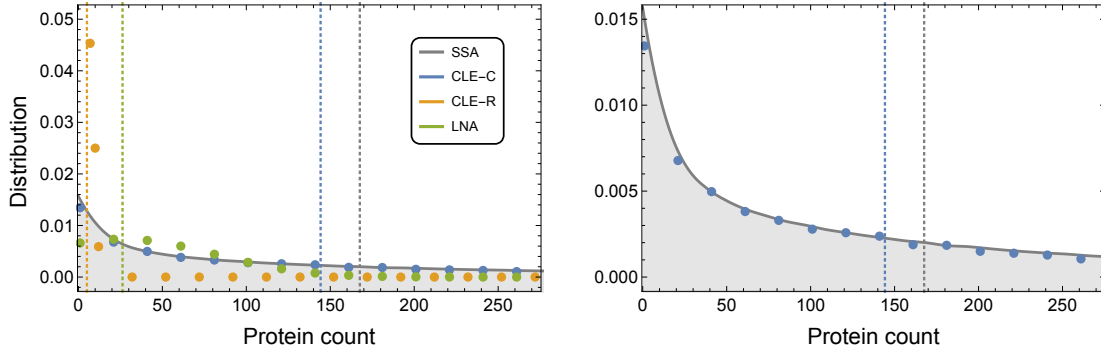


Figure 12: **Steady-state distribution of protein for extended system in Equation (108).** Left: distribution obtained using the CLE-R, CLE-C, LNA and SSA. Right: only CLE-C and SSA. The vertical lines indicate the corresponding steady-state mean values. The used parameters are $E^0 = 60$, $k_s = 1.5$, $k_0 = 15.33$, $k_{dM} = 0.2$, $k_1 = 3.5$, $k_2 = 2$, $k_3 = 2$, $\Omega = 1$, corresponding to a burst size of $b = k_s/k_{dM} = 7.5$. Simulation parameters for the CLE-R and CLE-C are $dt = 10^{-4}$, $\Delta t = 10$ and $M = 10^4$.

average several proteins during its lifetime. The average number of proteins per mRNA molecule is given by $b = k_s/k_{dM}$ and also called “burst size” [220]. Large b correspond to large bursts which in turn lead to large fluctuations.

Here, we are interested in the accuracy of the different approximation methods as a function of the burst size b , since we expect larger errors for larger fluctuations. To isolate the effect of the burst size b , we vary k_0 and k_{dM} such that the average production rate of protein $k_0 k_s/k_{dM} = 115$ in steady-state conditions is held constant and equal to the production rate used in the previous section in Figure 10.

Figure 11 shows the mean and variance of the protein S as a function of b as predicted by the different approximation methods. Similar to the previous section in Figure 9 we find that the LNA performs worse than the other methods for the mean and the variance, with the exception of the real-valued CLE (CLE-R) which performs even worse. Similarly to Figure 10, we find here that the complex-valued CLE (CLE-C) performs surprisingly well, demonstrating again the significance of the boundary problem of the CLE. In fact, the CLE-C performs significantly better than all the other methods, including the 2MA and SSE-1. The latter two perform similarly and both significantly better than the LNA. One should keep in mind, however, that the 2MA, SSE-1 and LNA are computationally significantly more efficient than the CLE-C since they do not rely on sampling.

Next, we consider the steady-state distribution of the protein for a large burst size of $b = 7.5$. Figure 12 shows the results predicted by the stochastic simulation algorithm, the two CLE implementations, and the LNA. We find that the CLE-R and LNA give highly inaccurate results, both for the mean value and the distribution, with even larger deviations than for the simpler system in the previous section in Figure 10. This is to be expected, since the system now leads to large bursts in protein production, while the average production rate of protein and the rate constants of the enzymatic reactions

are the same as in Figure 10. Similar to the previous section, we find here that the complex-valued CLE gives a highly accurate approximation for the distribution, despite the large skewness of the later. Once again we find that the boundary problem of the CLE can lead to highly inaccurate results if not treated carefully.

5.3 Discussion

Numerical results. In this section we gave a numerical comparison of the CLE, the LNA, the normal moment closure to second order (2MA) as well as the next leading order corrections to mean and variance of the system size expansion beyond the LNA (SSE-1). We implemented two versions of the CLE: a real-valued version with rejecting boundary and a complex-valued version. We considered an enzymatic protein degradation system (Section 5.1) and an extension thereof including bursty protein production (Section 5.2).

One important observation we made is the large discrepancy between the two CLE implementations: while we found a negligible difference in the case of large total enzyme numbers (Figure 9), the complex version was significantly more accurate in the case of smaller total enzyme numbers (Figures 10, 11 and 12). Crucially, in the latter cases, the real-valued CLE performed worse than all the other methods, while the complex-valued CLE performed better than all the other methods. This illustrates the boundary problem of the CLE, and how significant inaccuracies can arise when fixing it in a naive way.

Another important observation is that the LNA was throughout found to be less accurate than all the other methods (except the CLE-R). This result is not very surprising, since the LNA corresponds to the deterministic rate equations on the mean level, and gives the zeroth order fluctuations (in terms of the system size expansion) about the deterministic mean for the variance. The SSE-1 includes the leading order corrections to both the mean and variance. Similarly, the 2MA and CLE capture effects of fluctuations on the mean. It is hence not very surprising that these methods perform better than the LNA.

In terms of the mean and variance (Figures 9 and 11), we found that the SSE-1 and 2MA performed similar to each other, except for the variance in Figure 9 where 2MA was more accurate. Surprisingly however, the CLE-C was found to be significantly more accurate than the 2MA and SSE-1 in most cases.

In terms of distributions (Figures 10 and 12) we found the CLE-C to be highly accurate, even though the distributions were highly skewed. The LNA, which predicts a Gaussian distribution, was obviously not able to capture these skewed distributions.

Advantages and disadvantages. The advantage of the CLE over the other methods is that it gives approximations of the process and distributions in contrast to moment closure methods and the system size expansion. Moment closure approximations give only approximations to the first few moments of a process. The system size expansion in principle predicts distributions. However, closed-form solutions for higher orders beyond the LNA have so far only been derived for one-dimensional systems [108]. It is not clear if the same is possible for multi-species systems. The higher orders of the system size expansion have therefore mainly been used to approximate the moments of a process. If one is interested in approximating the whole process or its distributions, the CLE is

therefore a useful method.

Suppose now that we are only interested in the moments of a process. In the numerical case study performed before, we found the CLE to be more accurate than the LNA, the SSE-1 and the 2MA. However, the CLE is computationally significantly more expensive than the other methods. While the CLE requires a large number of stochastic simulations to obtain the moments of a process, the other methods only require the numerical solution of a finite set of ODEs and are hence typically orders of magnitude faster. Moreover, if defined for real-valued variables, the CLE suffers from a boundary problem at zero molecule numbers, and real-valued modifications lead to inaccurate results. The boundary problem is solved by extending the CLE to complex-valued variables, which is however less efficient to simulate [87]. Due to these reasons, it seems preferable to use the system size expansion or moment closure approximations if one is only interested in the moments of a process.

Next, the question arises if the system size expansion or moment closure approximations are preferable. While the system size expansion is a systematic expansion in a small parameter, moment closure approximations are an *ad hoc* approximation. This fact makes the system size expansion more appealing, since it is guaranteed to be accurate for large system volumes. The same cannot generally be expected to be true for moment closure approximations. On the other hand, the system size expansion has the disadvantage that it is not applicable to systems that are deterministically multi-stable, a limitation not shared by moment closure methods. Moreover, higher order corrections of the system size expansion are significantly harder to derive and implement than higher order moment closure methods. Current software packages implementing the system size expansion only implement two orders beyond the LNA for the mean, and one order beyond the LNA for the covariance [143, 147]. Moment closure approximations, on the other hand, are implemented to various orders [132, 146, 147]. Due to these reasons, it depends on the problem at hand as to decide which method is preferable.

6 Inference

So far, we have focussed on the *forward problem* of approximating marginal distributions of a fully specified process. Such distributions depend naturally on the parametrisation of the process: it is not uncommon for e.g. steady-state distributions to exhibit qualitatively different behaviours depending on the specific value of reaction propensities. In many concrete applications, the model parameters may only be known approximately: direct measurements of kinetic reaction parameters are difficult to obtain, and, even in cases when good estimates are available, *in vivo* parameters of a concrete system embedded in a cell may be influenced by a plethora of additional factors, leading to significant uncertainty. It is therefore of considerable interest to also address the *inverse problem*: using (noisy) observations of a system to constrain the uncertainty over model parameters and/ or predictions. This is a well-studied problem in statistics and machine learning: we give here a brief review of recent developments within the *Bayesian* approach to solving this inverse problem, with particular attention to methodologies which have employed the approximation methods described earlier.

6.1 General problem formulation

The setup we will consider is the following: we consider a stochastic process $p(\mathbf{x}_{0:T}|\theta)$ as a measure over the space of trajectories $\mathbf{x}_{0:T}$ of the system in the time interval $[0, T]$, with \mathbf{x}_t representing the value of the state variable at time t . In the case of a CME system, such trajectories will be piecewise constant functions from $[0, T]$ onto a discrete space, while for a continuous approximation (e.g. CLE or LNA) \mathbf{x}_t will be real-valued. The stochastic process is assumed to depend on a set of kinetic parameters θ , whose *a priori* uncertainty is captured by a prior distribution $p(\theta)$. Additionally, we assume the existence of a measurement process which associates each trajectory $\mathbf{x}_{0:T}$ with an observed random variable \mathbf{y} ; in the simplest case, the observed variable \mathbf{y} may just be thought of as the state of the system at a particular set of time points, corrupted by random observation noise. We account for such experimental errors through an *observation model* encoded in a probability distribution $p(\mathbf{y}|\mathbf{x}_{0:T})$, i.e., the likelihood of a measurement given the true state of the system, which may depend on additional parameters (here omitted for notational conciseness). We restrict our interest here to the case where the observations are in the form of a time series of state variable observations with independent and identically distributed noise, i.e., measurements of all or of a subgroup of the species in the system. Hence, the observation vector will take the form $\mathbf{y} = (y_0, \dots, y_T)$ for some discrete time points $0, \dots, T$. More general cases where the observations take other forms, such as continuous-time constraints or penalties over particular areas of the state space, are treated for example in [221, 222]. Also, we consider all observations to come from a single trajectory, or from replicate trajectories with the same (unknown) kinetic parameter values; in other words, we exclude the case where parameters can also be variable, e.g. to accommodate extrinsic noise between cells [223].

The general inference problem is then the problem of computing the joint posterior measure $p(\mathbf{x}_{0:T}, \theta|\mathbf{y})$ over the set of trajectories and parameters of the system. This is formally obtained by applying Bayes' rule

$$p(\mathbf{x}_{0:T}, \theta|\mathbf{y}) \propto p(\mathbf{y}|\mathbf{x}_{0:T})p(\mathbf{x}_{0:T}|\theta)p(\theta). \quad (109)$$

This joint posterior provides information both about the parameters and about the state of the system at all time points during the specific trajectory for which data was collected. In most systems of interest, computation of the normalisation constant of the posterior distribution is analytically intractable, due to the requirement of performing very large sums/ high dimensional integrals. Much of the research in Bayesian statistics and machine learning is therefore focussed on computing efficient approximations to posterior distributions: such approximations can be analytic, usually obtained by variational minimisation of a divergence functional, or sampling based [224]. This latter class has received particular attention in recent years, and is predicated on constructing a Markov chain which has the required posterior as an invariant distribution. This implies that, asymptotically, the Markov chain will sample from the correct posterior distribution, enabling Monte Carlo computation of any desired statistic (*Markov chain Monte Carlo, MCMC*).

6.2 The Forward-Backward algorithm

Inference in dynamical systems is based on the fundamental factorisation of the single time posterior marginal

$$p(\mathbf{x}_t | \mathbf{y}_0, \dots, \mathbf{y}_T) \propto p(\mathbf{x}_t | \mathbf{y}_{i \leq t}) p(\mathbf{y}_{i > t} | \mathbf{x}_t). \quad (110)$$

This factorisation, which is a simple consequence of the Markovian assumption and the product rule of probability [224], states that the posterior probability at time t is a product of the posterior $p(\mathbf{x}_t | \mathbf{y}_{i \leq t})$ based on the data seen so far up to time t (the so called *filtering distribution*) and the likelihood $p(\mathbf{y}_{i > t} | \mathbf{x}_t)$ of future data conditioned on the current state. The factors in equation (110) can be computed iteratively using the celebrated *Forward-Backward algorithm* [225].

The forward part, which is also referred to as *filtering* [226], works as follows. Assume that we know the *posterior* $p(x_{i-1} | y_{i-1}, \dots, y_0)$ at time step $i-1$, and that we can solve the system forward in time (i.e., solve the CME) to obtain the transition probability $p(\mathbf{x}_i | \mathbf{x}_{i-1})$ and hence the *predictive distribution* $p(\mathbf{x}_i | \mathbf{y}_{i-1}, \dots, \mathbf{y}_0) = \int d\mathbf{x}_{i-1} p(\mathbf{x}_i | \mathbf{x}_{i-1}) p(\mathbf{x}_{i-1} | \mathbf{y}_{i-1}, \dots, \mathbf{y}_0)$. The posterior of time step i is then obtained by taking the measurement \mathbf{y}_i at time point i into account by means of the *Bayesian measurement update*

$$\begin{aligned} p(\mathbf{x}_i | \mathbf{y}_i, \dots, \mathbf{y}_0) &= \frac{p(\mathbf{y}_i | \mathbf{x}_i, \mathbf{y}_{i-1}, \dots, \mathbf{y}_0) p(\mathbf{x}_i | \mathbf{y}_{i-1}, \dots, \mathbf{y}_0)}{p(\mathbf{y}_i | \mathbf{y}_{i-1}, \dots, \mathbf{y}_0)} \\ &= \frac{p(\mathbf{y}_i | \mathbf{x}_i) p(\mathbf{x}_i | \mathbf{y}_{i-1}, \dots, \mathbf{y}_0)}{p(\mathbf{y}_i | \mathbf{y}_{i-1}, \dots, \mathbf{y}_0)}, \end{aligned} \quad (111)$$

where we have used the Markov property to obtain the second line. The filtering procedure thus comprises iteratively solving the process between measurements and performing the Bayesian measurement update in Equation (111) and yields the filtering distribution $p(\mathbf{x}_t | \mathbf{y}_{i \leq t})$.

The second term in Equation (110) can be obtained by the Backward algorithm, a recursion procedure similar to filtering (but running backward in time) [225]. Furthermore, a modification of the algorithm, the so called *forward filtering/backward sampling* algorithm [227], can be used to draw sample trajectories from the posterior process, which can then be used in MCMC approaches for joint state and parameter inference.

6.3 Parameter inference

Suppose we are not interested in state inference but only in the parameters of the system. In this case we do not need to compute the posterior marginal in Equation (110). Rather, it is sufficient to either run the forward algorithm (filtering) or the backward algorithm, since each of them independently deliver the marginal likelihood $p(\mathbf{y})$. The backward algorithm computes the likelihood $p(\mathbf{y}_{i > t} | \mathbf{x}_t)$ of future data conditioned on the current state, and $p(\mathbf{y})$ is simply the end result of the recursion algorithm at time 0.

To see that the forward algorithm also allows to compute the likelihood $p(\mathbf{y})$, note

that due to the Markov property the latter can be written as

$$p(\mathbf{y}) = p(\mathbf{y}_0) \prod_{i=1}^T p(\mathbf{y}_i | \mathbf{y}_{i-1}, \dots, \mathbf{y}_0). \quad (112)$$

We find the factors on the r.h.s. of Equation (112) are just the normalisations factors of the Bayesian updates of the filtering procedure in Equation (112).

Optimising $p(\mathbf{y}) = p(\mathbf{y}|\theta)$ with respect to the parameters yields asymptotically consistent parameter estimates, also called *maximum likelihood estimate*. In a Bayesian framework, one would combine the likelihood with a parameter prior $p(\theta)$ to give the posterior over the parameters according to Bayes' law

$$p(\theta|\mathbf{y}) \propto p(\mathbf{y}|\theta)p(\theta), \quad (113)$$

which also allows to quantify uncertainty of the inferred parameter values.

6.4 Computational methods for Bayesian inference in stochastic chemical reaction networks

6.4.1 Methods for general networks

The primary difficulty in applying the forward-backward approach to inference in chemical reaction networks is the requirement for forward integrability of the system dynamics: calculating the transition probabilities $p(\mathbf{x}_i | \mathbf{x}_{i-1})$ requires solving the CME, which is generally not possible analytically. Some approaches resort to numerical integration of the equations, including the variational approach of [228] and the uniformisation sampler of [229]. These approaches can be effective for closed systems with low molecular numbers; however their application to open systems invariably requires an artificial truncation of the state space, introducing a bias which is hard to quantify. Truncations are also used in the recent work of [230]; however, here a random truncation scheme guarantees unbiasedness of the results, as well as leading to substantial computational savings. Other approaches that can handle open systems either introduce additional latent auxiliary variables (such as the number of reactions in the time interval as in [231]), or resort to a sequential Monte-Carlo scheme which relies on multiple simulations from different initial conditions (particle filtering, [232]). Both such schemes incur potentially large computational overheads.

The computationally intensive nature of inference methodologies adopting a microscopic system description constitutes a formidable obstacle to inference in large-scale reaction networks. This has justified a considerable interest in the use of mesoscopic approximations for inference. One of the earliest attempts [233] relied on the chemical Langevin equation approximation to the CME (Section 4.1); this provides a more efficient inference scheme compared to the auxiliary variable approach of [231], however the computational costs remain high due to the need to compute transition probabilities for non-linear diffusion processes. In this light, the LNA provides a more promising avenue, since the sufficient statistics of the (Gaussian) single-time marginals can be efficiently computed by integrating a system of ordinary differential equations. Several

authors have therefore proposed inference schemes which integrate the LNA approximation [234, 235, 236]. Moment closure approximations provide an alternative approximation scheme with similar computational complexity to the LNA, however they do not generally compute a marginal distribution, rather only a few moments of a generally unknown distribution. Their use for time series inference is therefore limited to second-order normal moment closure schemes, where a Gaussian approximation is taken. This approach has been proposed in [237], where it has been shown to yield accurate results with modest computational overheads. In [238] the second-order normal moment closure has been combined with the chemical Langevin equation and integrated into an expectation-propagation algorithm for intractable likelihoods and continuous-time constraints. Another interesting opportunity for moment-based inference is offered by flow-cytometry data: here, simultaneous measurements of millions of cells enable an empirical characterisation of the marginal moments directly (albeit potentially corrupted by extrinsic noise), which can then be fitted to a moment-approximation of the CME [239].

6.4.2 Inference for gene expression data

The inference methods mentioned above do not assume any knowledge about a given system. While this makes them in principle applicable to any type of reaction network, more efficient and/or accurate methods can often be employed by including a priori knowledge in a model. Gene expression systems constitute a particularly important example where this is often the case. Cells typically possess only one or very few copies of a gene. Proteins and mRNA molecules on the other hand often occur at copy numbers that are orders of magnitude larger. Such systems are therefore often suitable for hybrid methods (c.f., Section 4.6.4) that model some of the species as discrete variables (e.g. the genes) and others as continuous variables (e.g. the mRNA and proteins).

While it is often not straightforward to integrate hybrid methods into inference schemes, significant progress has recently been made in this respect. In [240], for instance, the different promoter states of a gene are modelled as a change-point process which drives a linear SDE representing the protein dynamics, and an efficient MCMC inference method is developed. The same model has been integrated into a variational inference method which additionally allows transcriptional feedback in [241, 242]. In [243], a particle MCMC scheme is developed based on a hybrid method that approximates the continuous variables by the linear noise approximation (c.f. Section 4.2.2). More recently a different hybrid approach combining several types of approximations has been integrated into an efficient Bayesian inference scheme in [244].

6.5 Summary

This bird eye survey of inference for stochastic chemical reaction networks highlights the diversity of statistical research in the area. Such diversity can appear baffling to the outsider, and a major problem for the greater diffusion of these ideas is the lack of standard software tools. Inference tools for stochastic chemical reaction networks often form a small subsection of software tools for parameter estimation of deterministic methods [245, 246], and there haven't been systematic comparisons of various inference

schemes that investigate the relative merits of the different algorithms on a number of relevant examples. Furthermore, inference approaches inevitably construct an approximation of a posterior distribution; while much statistics research investigates the convergence properties of these approximations, it remains entirely unclear how inference errors combine with approximation errors when inference schemes are deployed on approximate dynamics. Due to these reasons, we have chosen not to include an explicit numerical comparison between inference methods in this tutorial, as we feel this would deserve a separate review on its own.

7 Conclusions

Recent years have seen an explosion of experimental studies revealing the crucial role that stochastic fluctuations in chemical reaction networks play for living cells. Driven by these discoveries a plethora of methods for the mathematical and statistical analysis of such systems has evolved. The goal of this review is to give a self-contained introduction to the field of modelling for stochastic chemical kinetics. Moreover, it introduces key approximation and inference methods for this field and gives an overview of recent developments.

The Chemical Master Equation (CME) constitutes the accepted non-spatial description of stochastic chemical networks. Recent years have seen a burst of analysis and approximation methods based on the CME. We gave here an introduction to the CME modelling framework and discussed stochastic simulation and analytic solution methods. Next, we introduced various approximation methods with particular focus on the chemical Langevin equation, the system size expansion and moment closure methods. We also gave an introduction to time-scale separation based approximations, as well as hybrid methods and reviewed the existing literature. Finally, we gave an introduction to the problem of statistical inference from experimental data in a Bayesian framework, and reviewed existing methods. The presentation is aimed to be a self-contained introduction for scientists from different disciplines.

In a numerical case study we compared the chemical Langevin equation, the zeroth order system size expansion (linear noise approximation, LNA), the first-order corrections of the system size expansion to mean and covariance (SSE-1) and the second-order normal moment closure (2MA) with exact results obtained using stochastic simulations. In terms of moments, we found that a naive real-valued implementation of the CLE (CLE-R) enforcing positive concentrations was less accurate than all the other methods. A complex version of the CLE (CLE-C), in contrast, was found to be the most accurate of all methods. The SSE-1 and 2MA performed similar to each other and significantly better than the LNA. In terms of steady-state distributions, we compared the CLE-R, the CLE-C and the LNA with exact results obtained using the SSA. We found that the CLE-R and the LNA were not able to accurately capture the distributions, the LNA being more accurate than the CLE-R, however. The CLE-C, in contrast, gave accurate approximations even for highly skewed distributions.

The CME is a valid description of systems that are well-mixed and sufficiently dilute, i.e., the diffusion of particles constitutes the fastest time scale of the system and the total volume of all molecules in the model is much smaller than the system volume. While

these assumptions are valid in some cases, it turns out that they are not met by many biological systems. Whenever this is the case, models need to be employed that take spatial positions and diffusion of particles into account. The main family of such models goes under the name of *stochastic reaction-diffusion processes* (SRDPs), Markovian models where independent particles diffuse in space and react whenever they come in contact (or sufficiently close). The evolution equation for the marginal probabilities of an SRDP, the spatial analogue of the CME, is complicated by the fact that the number of particles varies in time, and needs to be defined on an infinite-dimensional Fock space [247, 248]. Solving, or even approximating, such an equation is essentially impossible, and hence SRDPs are mainly analysed in an algorithmic way: each particle is simulated performing Brownian motion in continuous space and chemical reactions between particles happen stochastically under certain rules [249]. This is computationally extremely expensive and significant effort has been spent in improved simulation methods [250, 251]. An alternative modelling framework is given by the *reaction-diffusion master equation* (RDME) which coarse-grains an SRDP by assuming a compartmentalisation of space and locally homogeneous conditions within each compartment [252]. While simulations in this framework are generally more efficient than in the continuous case, they are typically still expensive [253], generally significantly more expensive than in the non-spatial CME case. More importantly, the RDME is not a systematic discretisation of an SRDP, in the sense that its continuum limit generally does not lead to the original SRDP in two or more spatial dimensions [252], because bimolecular reactions become vanishingly infrequent in the continuum limit.

Due to these reasons, analysis methodologies for spatial models are much less developed than in the non-spatial CME case. Some studies investigating spatial stochastic phenomena and comparing the SRDP and RDME approaches include [249, 254, 255, 256, 257, 258, 259]. In contrast to the CME case, very few studies have attempted analytical approximations for SRDPs. However, some progress has been made in recent years in this respect. In [260] and [261], for example, the linear noise approximation was extended to spatial systems, and in [262] higher orders of the system size expansion for effective one-species systems. Even fewer studies have addressed inference for SRDPs. Only a handful of studies have approached the issue of inference working directly with the SRDP or RDME framework, using likelihood-free sampling methods [263] or variational approximations [264]. In [265], the linear noise approximation has been integrated into an inference scheme for SRDPs. This method is however limited to certain classes of one-dimensional systems. More recently, in [82] it was shown that SRDPs can be approximated by spatio-temporal point processes, a popular class of models from statistics [266]. This was used to derive an efficient inference algorithm for general SRDPs.

In summary, we have presented an introduction to modelling, approximation and inference methods for stochastic chemical kinetics based on the Chemical Master Equation. We hope that this review will help scientists from other disciplines to dive into this exciting field, and that it will stimulate research in the presented areas.

Acknowledgments

This work was supported by the Biotechnology and Biological Sciences Research Council [BB/F017073/1]; the Leverhulme Trust [RPG-2013-171]; and the European Research Council [MLCS 306999]. The authors thank Elco Bakker, Lucia Bandiera, Lito Pappaxenopoulou and Michael Rule for useful comments on a draft of this article.

References

- [1] Avigdor Eldar and Michael B Elowitz. Functional roles for noise in genetic circuits. *Nature*, 467(7312):167–173, 2010.
- [2] Michael B Elowitz, Arnold J Levine, Eric D Siggia, and Peter S Swain. Stochastic gene expression in a single cell. *Science*, 297(5584):1183–1186, 2002.
- [3] Daniel T Gillespie. A rigorous derivation of the chemical master equation. *Physica A: Statistical Mechanics and its Applications*, 188(1):404–425, 1992.
- [4] Daniel T Gillespie. A general method for numerically simulating the stochastic time evolution of coupled chemical reactions. *Journal of Computational Physics*, 22(4):403–434, 1976.
- [5] Daniel T Gillespie. Exact stochastic simulation of coupled chemical reactions. *The Journal of Physical Chemistry*, 81(25):2340–2361, 1977.
- [6] Thomas E Turner, Santiago Schnell, and Kevin Burrage. Stochastic approaches for modelling in vivo reactions. *Computational biology and chemistry*, 28(3):165–178, 2004.
- [7] Hana El Samad, Mustafa Khammash, Linda Petzold, and Dan Gillespie. Stochastic modelling of gene regulatory networks. *International Journal of Robust and Nonlinear Control*, 15(15):691–711, 2005.
- [8] Hong Li, Yang Cao, Linda R Petzold, and Daniel T Gillespie. Algorithms and software for stochastic simulation of biochemical reacting systems. *Biotechnology Progress*, 24(1):56–61, 2008.
- [9] Jürgen Pahle. Biochemical simulations: stochastic, approximate stochastic and hybrid approaches. *Briefings in Bioinformatics*, 10(1):53–64, 2009.
- [10] Sean Mauch and Mark Stalzer. Efficient formulations for exact stochastic simulation of chemical systems. *IEEE/ACM Transactions on Computational Biology and Bioinformatics (TCBB)*, 8(1):27–35, 2011.
- [11] Daniel T Gillespie, Andreas Hellander, and Linda R Petzold. Perspective: Stochastic algorithms for chemical kinetics. *The Journal of Chemical Physics*, 138(17):170901, 2013.
- [12] Tamas Szekely and Kevin Burrage. Stochastic simulation in systems biology. *Computational and structural biotechnology journal*, 12(20):14–25, 2014.
- [13] Gregoire Nicolis, Ilya Prigogine, et al. *Self-organization in nonequilibrium systems*, volume 191977. Wiley, New York, 1977.
- [14] Nicolaas G van Kampen. *Stochastic processes in physics and chemistry*, volume 1. Elsevier, 1992.
- [15] Crispin W Gardiner. *Handbook of Stochastic Methods*, volume 3. Springer Berlin, 1985.
- [16] John Goutsias and Garrett Jenkinson. Markovian dynamics on complex reaction networks. *Physics Reports*, 529(2):199–264, 2013.
- [17] Markus F Weber and Erwin Frey. Master equations and the theory of stochastic path integrals. *arXiv preprint arXiv:1609.02849*, 2016.

- [18] Richard A Blythe and Alan J McKane. Stochastic models of evolution in genetics, ecology and linguistics. *Journal of Statistical Mechanics: Theory and Experiment*, 2007(07):P07018, 2007.
- [19] Samik Datta, Gustav W Delius, and Richard Law. A jump-growth model for predator–prey dynamics: derivation and application to marine ecosystems. *Bulletin of Mathematical Biology*, 72(6):1361–1382, 2010.
- [20] Andrew J Black and Alan J McKane. Stochastic formulation of ecological models and their applications. *Trends in Ecology & Evolution*, 27(6):337–345, 2012.
- [21] MS Bartlett. Some evolutionary stochastic processes. *Journal of the Royal Statistical Society. Series B (Methodological)*, 11(2):211–229, 1949.
- [22] Norman T J Bailey. A simple stochastic epidemic. *Biometrika*, pages 193–202, 1950.
- [23] Matthew James Keeling and Joshua V Ross. On methods for studying stochastic disease dynamics. *Journal of The Royal Society Interface*, 5(19):171–181, 2008.
- [24] Andrew J Black and Alan J McKane. Stochasticity in staged models of epidemics: quantifying the dynamics of whooping cough. *Journal of The Royal Society Interface*, 7(49):1219–1227, 2010.
- [25] Garrett Jenkinson and John Goutsias. Numerical integration of the master equation in some models of stochastic epidemiology. *PLoS One*, 7(5):e36160, 2012.
- [26] Romualdo Pastor-Satorras, Claudio Castellano, Piet Van Mieghem, and Alessandro Vespignani. Epidemic processes in complex networks. *Reviews of Modern Physics*, 87(3):925, 2015.
- [27] Wolfgang Weidlich. Physics and social science? The approach of synergetics. *Physics Reports*, 204(1):1–163, 1991.
- [28] Wolfgang Weidlich. Sociodynamics – a systematic approach to mathematical modelling in the social sciences. *NONLINEAR PHENOMENA IN COMPLEX SYSTEMS-MINSK-*, 5(4):479–487, 2002.
- [29] Eric Bonabeau. Agent-based modeling: Methods and techniques for simulating human systems. *Proceedings of the National Academy of Sciences*, 99(suppl 3):7280–7287, 2002.
- [30] Toru Ohira and Jack D Cowan. Master-equation approach to stochastic neurodynamics. *Physical Review E*, 48(3):2259, 1993.
- [31] Sami El Boustani and Alain Destexhe. A master equation formalism for macroscopic modeling of asynchronous irregular activity states. *Neural Computation*, 21(1):46–100, 2009.
- [32] Marc Benayoun, Jack D Cowan, Wim van Drongelen, and Edward Wallace. Avalanches in a stochastic model of spiking neurons. *PLoS Computational Biology*, 6(7):e1000846, 2010.
- [33] Michael A Buice, Jack D Cowan, and Carson C Chow. Systematic fluctuation expansion for neural network activity equations. *Neural Computation*, 22(2):377–426, 2010.
- [34] Edward Wallace, Marc Benayoun, Wim Van Drongelen, and Jack D Cowan. Emergent oscillations in networks of stochastic spiking neurons. *Plos One*, 6(5):e14804, 2011.
- [35] Todd K Leen and Robert Friel. Stochastic perturbation methods for spike-timing-dependent plasticity. *Neural computation*, 24(5):1109–1146, 2012.
- [36] Igor Goychuk and Andriy Goychuk. Stochastic Wilson–Cowan models of neuronal network dynamics with memory and delay. *New Journal of Physics*, 17(4):045029, 2015.
- [37] Bruce Alberts, Dennis Bray, Julian Lewis, Martin Raff, Keith Roberts, James D Watson, and AV Grimstone. Molecular Biology of the Cell (3rd edn). *Trends in Biochemical Sciences*, 20(5):210–210, 1995.

- [38] Peter S Swain, Michael B Elowitz, and Eric D Siggia. Intrinsic and extrinsic contributions to stochasticity in gene expression. *Proceedings of the National Academy of Sciences*, 99(20):12795–12800, 2002.
- [39] DW Austin, MS Allen, JM McCollum, RD Dar, JR Wilgus, GS Sayler, NF Samatova, CD Cox, and ML Simpson. Gene network shaping of inherent noise spectra. *Nature*, 439(7076):608–611, 2006.
- [40] Ertugrul M Ozbudak, Mukund Thattai, Iren Kurtser, Alan D Grossman, and Alexander van Oudenaarden. Regulation of noise in the expression of a single gene. *Nature Genetics*, 31(1):69–73, 2002.
- [41] Alejandro Colman-Lerner, Andrew Gordon, Eduard Serra, Tina Chin, Orna Resnekov, Drew Endy, C Gustavo Pesce, and Roger Brent. Regulated cell-to-cell variation in a cell-fate decision system. *Nature*, 437(7059):699–706, 2005.
- [42] Nitzan Rosenfeld, Jonathan W Young, Uri Alon, Peter S Swain, and Michael B Elowitz. Gene regulation at the single-cell level. *Science*, 307(5717):1962–1965, 2005.
- [43] Arren Bar-Even, Johan Paulsson, Narendra Maheshri, Miri Carmi, Erin O’Shea, Yitzhak Pilpel, and Naama Barkai. Noise in protein expression scales with natural protein abundance. *Nature genetics*, 38(6):636–643, 2006.
- [44] Yuichi Taniguchi, Paul J Choi, Gene-Wei Li, Huiyi Chen, Mohan Babu, Jeremy Hearn, Andrew Emili, and X Sunney Xie. Quantifying E. coli proteome and transcriptome with single-molecule sensitivity in single cells. *Science*, 329(5991):533–538, 2010.
- [45] Byron F Brehm-Stecher and Eric A Johnson. Single-cell microbiology: tools, technologies, and applications. *Microbiology and Molecular Biology Reviews*, 68(3):538–559, 2004.
- [46] Narendra Maheshri and Erin K O’Shea. Living with noisy genes: how cells function reliably with inherent variability in gene expression. *Annual Review of Biophysics and Biomolecular Structure*, 36:413–434, 2007.
- [47] Gábor Balázsi, Alexander van Oudenaarden, and James J Collins. Cellular decision making and biological noise: from microbes to mammals. *Cell*, 144(6):910–925, 2011.
- [48] Cato M Guldberg and Peter Waage. Studies concerning affinity. *CM Forhandling: Videnskabs-Selskabet i Christiana*, 35, 1864.
- [49] Peter Waage. Experiments for determining the affinity law. *Forhandling: Videnskabs-Selskabet i Christiana*, 92, 1864.
- [50] Cato Maximilian Guldberg. Concerning the laws of chemical affinity. *CM Forhandling: Videnskabs-Selskabet i Christiana*, 111, 1864.
- [51] Cato Maxilian Guldberg and Peter Waage. Concerning chemical affinity. *Erdmann’s Journal für Praktische Chemie*, 127:69–114, 1879.
- [52] Edda Klipp, Ralf Herwig, Axel Kowald, Christoph Wierling, and Hans Lehrach. *Systems biology in practice: concepts, implementation and application*. John Wiley & Sons, 2008.
- [53] Max Delbrück. Statistical fluctuations in autocatalytic reactions. *The Journal of Chemical Physics*, 8(1):120–124, 1940.
- [54] Anthony F Bartholomay. Stochastic models for chemical reactions: I. Theory of the unimolecular reaction process. *The Bulletin of Mathematical Biophysics*, 20(3):175–190, 1958.
- [55] Donald A McQuarrie. Kinetics of small systems. I. *The Journal of Chemical Physics*, 38(2):433–436, 1963.
- [56] Kenji Ishida. Stochastic model for bimolecular reaction. *The Journal of Chemical Physics*, 41(8):2472–2478, 1964.
- [57] Donald A McQuarrie, CJ Jachimowski, and ME Russell. Kinetics of small systems. II. *The Journal of Chemical Physics*, 40(10):2914–2921, 1964.

- [58] Daniel T Gillespie. A diffusional bimolecular propensity function. *The Journal of Chemical Physics*, 131(16):164109, 2009.
- [59] R Grima, Deena R Schmidt, and Timothy J Newman. Steady-state fluctuations of a genetic feedback loop: An exact solution. *The Journal of Chemical Physics*, 137(3):035104, 2012.
- [60] Maria Luisa Guerriero, Alexandra Pokhilko, Aurora Piñas Fernández, Karen J Halliday, Andrew J Millar, and Jane Hillston. Stochastic properties of the plant circadian clock. *Journal of The Royal Society Interface*, page rsif20110378, 2011.
- [61] David F Anderson. A modified next reaction method for simulating chemical systems with time dependent propensities and delays. *The Journal of Chemical Physics*, 127(21):214107, 2007.
- [62] Vahid Shahrezaei, Julien F Ollivier, and Peter S Swain. Colored extrinsic fluctuations and stochastic gene expression. *Molecular Systems Biology*, 4(1), 2008.
- [63] Giulio Caravagna, Giancarlo Mauri, and Alberto d’Onofrio. The interplay of intrinsic and extrinsic bounded noises in biomolecular networks. *PLoS One*, 8(2):e51174, 2013.
- [64] Margaritis Voliotis, Philipp Thomas, Ramon Grima, and Clive G Bowsher. Stochastic simulation of biomolecular networks in dynamic environments. *PLoS Comput Biol*, 12(6):e1004923, 2016.
- [65] Awad H Al-Mohy and Nicholas J Higham. Computing the action of the matrix exponential, with an application to exponential integrators. *SIAM Journal on Scientific Computing*, 33(2):488–511, 2011.
- [66] Cleve Moler and Charles Van Loan. Nineteen dubious ways to compute the exponential of a matrix, twenty-five years later. *SIAM Review*, 45(1):3–49, 2003.
- [67] IG Darvey and PJ Staff. Stochastic Approach to First-Order Chemical Reaction Kinetics. *The Journal of Chemical Physics*, 44(3):990–997, 1966.
- [68] CW Gardiner and S Chaturvedi. The poisson representation. i. a new technique for chemical master equations. *Journal of Statistical Physics*, 17(6):429–468, 1977.
- [69] William J Heuett and Hong Qian. Grand canonical markov model: a stochastic theory for open nonequilibrium biochemical networks. *The Journal of Chemical Physics*, 124(4):044110, 2006.
- [70] Tobias Jahnke and Wilhelm Huisinga. Solving the chemical master equation for monomolecular reaction systems analytically. *Journal of Mathematical Biology*, 54(1):1–26, 2007.
- [71] Peter Whittle. *Systems in stochastic equilibrium*. John Wiley & Sons, Inc., 1986.
- [72] NG Van Kampen. The equilibrium distribution of a chemical mixture. *Physics Letters A*, 59(5):333–334, 1976.
- [73] David F Anderson, Gheorghe Craciun, and Thomas G Kurtz. Product-form stationary distributions for deficiency zero chemical reaction networks. *Bulletin of Mathematical Biology*, 72(8):1947–1970, 2010.
- [74] Fritz Horn and Roy Jackson. General mass action kinetics. *Archive for Rational Mechanics and Analysis*, 47(2):81–116, 1972.
- [75] Martin Feinberg. The existence and uniqueness of steady states for a class of chemical reaction networks. *Archive for Rational Mechanics and Analysis*, 132(4):311–370, 1995.
- [76] R Voituriez, M Moreau, and G Oshanin. Corrections to the law of mass action and properties of the asymptotic $t = \infty$ state for reversible diffusion-limited reactions. *The Journal of Chemical Physics*, 122:084103, 2005.
- [77] Moritz Gerstung, Jens Timmer, and Christian Fleck. Noisy signaling through promoter logic gates. *Physical Review E*, 79(1):011923, 2009.

- [78] Philipp Thomas, Arthur V Straube, and Ramon Grima. Stochastic theory of large-scale enzyme-reaction networks: Finite copy number corrections to rate equation models. *The Journal of Chemical Physics*, 133(19):195101, 2010.
- [79] KG Petrosyan and Chin-Kun Hu. Nonequilibrium lyapunov function and a fluctuation relation for stochastic systems: Poisson-representation approach. *Physical Review E*, 89(4):042132, 2014.
- [80] Srividya Iyer-Biswas and Ciriya Jayaprakash. Mixed poisson distributions in exact solutions of stochastic autoregulation models. *Physical Review E*, 90(5):052712, 2014.
- [81] István Sugár and István Simon. Self-regulating genes. exact steady state solution by using poisson representation. *Open Physics*, 12(9):615–627, 2014.
- [82] David Schnoerr, Ramon Grima, and Guido Sanguinetti. Cox process representation and inference for stochastic reaction-diffusion processes. *Nature Communications*, 7:11729, 2016.
- [83] Ian J Laurenzi. An analytical solution of the stochastic master equation for reversible bimolecular reaction kinetics. *The Journal of Chemical Physics*, 113(8):3315–3322, 2000.
- [84] Niraj Kumar, Thierry Platini, and Rahul V Kulkarni. Exact distributions for stochastic gene expression models with bursting and feedback. *Physical Review Letters*, 113(26):268105, 2014.
- [85] P Arányi and J Tóth. A full stochastic description of the Michaelis-Menten reaction for small systems. *Acta biochimica et biophysica; Academiae Scientiarum Hungaricae*, 12(4):375–388, 1976.
- [86] PJ Staff. A stochastic development of the reversible Michaelis-Menten mechanism. *Journal of Theoretical Biology*, 27(2):221–232, 1970.
- [87] David Schnoerr, Guido Sanguinetti, and Ramon Grima. The complex chemical Langevin equation. *The Journal of Chemical Physics*, 141(2):024103, 2014.
- [88] William H Mather, Natalie A Cookson, Jeff Hasty, Lev S Tsimring, and Ruth J Williams. Correlation resonance generated by coupled enzymatic processing. *Biophysical Journal*, 99(10):3172–3181, 2010.
- [89] Hendrik A Kramers. Brownian motion in a field of force and the diffusion model of chemical reactions. *Physica*, 7(4):284–304, 1940.
- [90] José E Moyal. Stochastic processes and statistical physics. *Journal of the Royal Statistical Society. Series B (Methodological)*, 11(2):150–210, 1949.
- [91] Daniel T Gillespie. The chemical Langevin equation. *The Journal of Chemical Physics*, 113(1):297–306, 2000.
- [92] Bence Melykúti, Kevin Burrage, and Konstantinos C Zygalakis. Fast stochastic simulation of biochemical reaction systems by alternative formulations of the chemical Langevin equation. *The Journal of Chemical Physics*, 132(16):164109, 2010.
- [93] Peter E Kloeden and Eckhard Platen. *Numerical Solution of Stochastic Differential Equations*, volume 1. Springer Berlin, 1992.
- [94] Kevin Burrage, Pamela Burrage, and Taketomo Mitsui. Numerical solutions of stochastic differential equations—implementation and stability issues. *Journal of Computational and Applied Mathematics*, 125(1):171–182, 2000.
- [95] Desmond J Higham. An algorithmic introduction to numerical simulation of stochastic differential equations. *SIAM Review*, 43(3):525–546, 2001.
- [96] Kevin Burrage, PM Burrage, and Tianhai Tian. Numerical methods for strong solutions of stochastic differential equations: an overview. *Proceedings of The Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 460(2041):373–402, 2004.

- [97] Thomas G Kurtz. Limit theorems and diffusion approximations for density dependent Markov chains. In *Stochastic Systems: Modeling, Identification and Optimization, I*, pages 67–78. Springer, 1976.
- [98] Saswati Dana and Soumyendu Raha. Physically consistent simulation of mesoscale chemical kinetics: The non-negative FIS- α method. *Journal of Computational Physics*, 230(24):8813–8834, 2011.
- [99] Joshua Wilkie and Yin M Wong. Positivity preserving chemical Langevin equations. *Chemical Physics*, 353(1):132–138, 2008.
- [100] Lukasz Szpruch and Desmond J Higham. Comparing hitting time behavior of Markov jump processes and their diffusion approximations. *Multiscale Modeling & Simulation*, 8(2):605–621, 2010.
- [101] Shuohao Liao, Tomáš Vejchodský, and Radek Erban. Tensor methods for parameter estimation and bifurcation analysis of stochastic reaction networks. *Journal of The Royal Society Interface*, 12(108):20150233, 2015.
- [102] Tommaso Biancalani, Louise Dyson, and Alan J McKane. Noise-induced bistable states and their mean switching time in foraging colonies. *Physical review letters*, 112(3):038101, 2014.
- [103] Andrew Duncan, Shuohao Liao, Tomáš Vejchodský, Radek Erban, and Ramon Grima. Noise-induced multistability in chemical systems: Discrete versus continuum modeling. *Physical Review E*, 91(4):042111, 2015.
- [104] Marco Beccuti, Enrico Bibbona, Andras Horvath, Roberta Sirovich, Alessio Angius, and Gianfranco Balbo. Analysis of petri net models through stochastic differential equations. In *International Conference on Applications and Theory of Petri Nets and Concurrency*, pages 273–293. Springer, 2014.
- [105] Nicolaas G van Kampen. The expansion of the master equation. *Advance in Chemical Physics*, 34:245–309, 1976.
- [106] Philipp Thomas, Hannes Matuschek, and Ramon Grima. Computation of biochemical pathway fluctuations beyond the linear noise approximation using iNA. In *IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, pages 1–5, 2012.
- [107] Ramon Grima. An effective rate equation approach to reaction kinetics in small volumes: Theory and application to biochemical reactions in nonequilibrium steady-state conditions. *The Journal of Chemical Physics*, 133(3):035101, 2010.
- [108] Philipp Thomas and Ramon Grima. Approximate probability distributions of the master equation. *Physical Review E*, 92(1):012120, 2015.
- [109] Alan J McKane, James D Nagy, Timothy J Newman, and Marianne O Stefanini. Amplified biochemical oscillations in cellular systems. *Journal of Statistical Physics*, 128(1-2):165–191, 2007.
- [110] Jürgen Pahle, Joseph D Challenger, Pedro Mendes, and Alan J McKane. Biochemical fluctuations, optimisation and the linear noise approximation. *BMC systems biology*, 6(1):1, 2012.
- [111] Joseph D Challenger and Alan J McKane. Synchronization of stochastic oscillators in biochemical systems. *Physical Review E*, 88(1):012107, 2013.
- [112] C Bianca and A Lemarchand. Evaluation of reaction fluxes in stationary and oscillating far-from-equilibrium biological systems. *Physica A: Statistical Mechanics and its Applications*, 438:1–16, 2015.
- [113] Peter G Hufton, Yen Ting Lin, Tobias Galla, and Alan J McKane. Intrinsic noise in systems with switching environments. *Physical Review E*, 93(5):052119, 2016.
- [114] Philipp Thomas, Hannes Matuschek, and Ramon Grima. How reliable is the linear noise approximation of gene regulatory networks? *BMC genomics*, 14(Suppl 4):S5, 2013.

- [115] Thierry Dauxois, Francesca Di Patti, Duccio Fanelli, and Alan J McKane. Enhanced stochastic oscillations in autocatalytic reactions. *Physical Review E*, 79(3):036112, 2009.
- [116] Claudia Cianci, Francesca Di Patti, Duccio Fanelli, and Luigi Barletti. Analytical study of non Gaussian fluctuations in a stochastic scheme of autocatalytic reactions. *The European Physical Journal Special Topics*, 212(1):5–22, 2012.
- [117] M Scott. Non-linear corrections to the time-covariance function derived from a multi-state chemical master equation. *IET Systems Biology*, 6(4):116–124, 2012.
- [118] Philipp Thomas, Arthur V Straube, Jens Timmer, Christian Fleck, and Ramon Grima. Signatures of nonlinearity in single cell noise-induced oscillations. *Journal of Theoretical Biology*, 335:222–234, 2013.
- [119] RF Pawula. Approximation of the linear Boltzmann equation by the Fokker-Planck equation. *Physical Review*, 162(1):186, 1967.
- [120] Todd K Leen and Robert Friel. Perturbation theory for stochastic learning dynamics. In *Neural Networks (IJCNN), The 2011 International Joint Conference on*, pages 2031–2038. IEEE, 2011.
- [121] Leo A Goodman. Population growth of the sexes. *Biometrics*, 9(2):212–225, 1953.
- [122] Peter Whittle. On the use of the normal approximation in the treatment of stochastic processes. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 268–281, 1957.
- [123] Carlos A Gomez-Uribe and George C Verghese. Mass fluctuation kinetics: Capturing stochastic effects in systems of chemical reactions through coupled mean-variance computations. *The Journal of Chemical Physics*, 126(2):024109, 2007.
- [124] Ingemar Nåsell. An extension of the moment closure method. *Theoretical Population Biology*, 64(2):233–239, 2003.
- [125] Matt J Keeling. Multiplicative moments and measures of persistence in ecology. *Journal of Theoretical Biology*, 205(2):269–281, 2000.
- [126] Edwin L Crow and Kunio Shimizu. *Lognormal distributions: Theory and applications*, volume 88. M. Dekker New York, 1988.
- [127] Joao Hespanha. Moment closure for biochemical networks. In *3rd International Symposium on Communications, Control and Signal Processing*, pages 142–147. IEEE, 2008.
- [128] Abhyudai Singh and Joao Pedro Hespanha. Lognormal moment closures for biochemical reactions. In *IEEE Conference on Decision and Control*, pages 2063–2068. IEEE, 2006.
- [129] Ramon Grima. A study of the accuracy of moment-closure approximations for stochastic chemical kinetics. *The Journal of Chemical Physics*, 136(15):154105, 2012.
- [130] Angélique Ale, Paul Kirk, and Michael PH Stumpf. A general moment expansion method for stochastic kinetic models. *The Journal of Chemical Physics*, 138(17):174101, 2013.
- [131] David Schnoerr, Guido Sanguinetti, and Ramon Grima. Validity conditions for moment closure approximations in stochastic chemical kinetics. *The Journal of Chemical Physics*, 141(8):084103, 2014.
- [132] David Schnoerr, Guido Sanguinetti, and Ramon Grima. Comparison of different moment-closure approximations for stochastic chemical kinetics. *The Journal of Chemical Physics*, 143(18):185101, 2015.
- [133] Abhyudai Singh and Joao P Hespanha. Approximate moment dynamics for chemically reacting systems. *IEEE Transactions on Automatic Control*, 56(2):414–418, 2011.
- [134] Eszter Lakatos, Angélique Ale, Paul DW Kirk, and Michael PH Stumpf. Multivariate moment closure techniques for stochastic kinetic models. *The Journal of Chemical Physics*, 143(9):094107, 2015.
- [135] Claude Elwood Shannon. A mathematical theory of communication. *ACM SIGMOBILE Mobile Computing and Communications Review*, 5(1):3–55, 2001.

- [136] Alexander Andreychenko, Linar Mikeev, and Verena Wolf. Model reconstruction for moment-based stochastic chemical kinetics. *ACM Transactions on Modeling and Computer Simulation (TOMACS)*, 25(2):12, 2015.
- [137] Alexander Andreychenko, Luca Bortolussi, Ramon Grima, Philipp Thomas, and Verena Wolf. Distribution approximations for the chemical master equation: comparison of the method of moments and the system size expansion. *arXiv preprint arXiv:1509.09104*, 2015.
- [138] Patrick Smadbeck and Yiannis N Kaznessis. A closure scheme for chemical master equations. *Proceedings of the National Academy of Sciences*, 110(35):14261–14265, 2013.
- [139] Stephen Ramsey, David Orrell, and Hamid Bolouri. Dizzy: stochastic simulation of large-scale genetic regulatory networks. *Journal of Bioinformatics and Computational Biology*, 3(02):415–436, 2005.
- [140] Stefan Hoops, Sven Sahle, Ralph Gauges, Christine Lee, Jürgen Pahle, Natalia Simus, Mudita Singhal, Liang Xu, Pedro Mendes, and Ursula Kummer. COPASI—a complex pathway simulator. *Bioinformatics*, 22(24):3067–3074, 2006.
- [141] Kevin R Sanft, Sheng Wu, Min Roh, Jin Fu, Rone Kwei Lim, and Linda R Petzold. StochKit2: software for discrete stochastic simulation of biochemical systems with events. *Bioinformatics*, 27(17):2457–2458, 2011.
- [142] Timo R Maarleveld, Brett G Olivier, and Frank J Bruggeman. StochPy: a comprehensive, user-friendly tool for simulating stochastic biological processes. *PLoS One*, 8(11):e79345, 2013.
- [143] Philipp Thomas, Hannes Matuschek, and Ramon Grima. Intrinsic noise analyzer: a software package for the exploration of stochastic biochemical kinetics using the system size expansion. *PLoS One*, 7(6):e38518, 2012.
- [144] Joao P Hespanha. StochDynTools - a MATLAB toolbox to compute moment dynamics for stochastic networks of bio-chemical reactions, 2006.
- [145] Colin S Gillespie. Moment-closure approximations for mass-action models. *Systems Biology, IET*, 3(1):52–58, 2009.
- [146] Sisi Fan, Quentin Geissmann, Eszter Lakatos, Saulius Lukauskas, Angelique Ale, Ann C Babbie, Paul DW Kirk, and Michael PH Stumpf. MEANS: python package for Moment Expansion Approximation, inference and Simulation. *Bioinformatics*, page btw229, 2016.
- [147] Atefeh Kazerooni, Fabian Fröhlich, Andreas Raue, Fabian J Theis, and Jan Hasenauer. CERENA: ChEmical REaction Network Analyzer—A Toolbox for the Simulation and Analysis of Stochastic Chemical Kinetics. *PLoS One*, 11(1):e0146732, 2016.
- [148] Brian Munsky and Mustafa Khammash. The finite state projection algorithm for the solution of the chemical master equation. *The Journal of Chemical Physics*, 124(4):044104, 2006.
- [149] Khanh N Dinh and Roger B Sidje. Understanding the finite state projection and related methods for solving the chemical master equation. *Physical biology*, 13(3):035003, 2016.
- [150] Daniel T Gillespie. Approximate accelerated stochastic simulation of chemically reacting systems. *The Journal of Chemical Physics*, 115(4):1716–1733, 2001.
- [151] Muruhan Rathinam, Linda R Petzold, Yang Cao, and Daniel T Gillespie. Stiffness in stochastic chemically reacting systems: The implicit tau-leaping method. *The Journal of Chemical Physics*, 119(24):12784–12794, 2003.
- [152] Yang Cao, Daniel T Gillespie, and Linda R Petzold. Adaptive explicit-implicit tau-leaping method with automatic tau selection. *The Journal of Chemical Physics*, 126(22):224101, 2007.
- [153] Tianhai Tian and Kevin Burrage. Binomial leap methods for simulating stochastic chemical kinetics. *The Journal of Chemical Physics*, 121(21):10356–10364, 2004.

- [154] Abhijit Chatterjee, Dionisios G Vlachos, and Markos A Katsoulakis. Binomial distribution based τ -leap accelerated stochastic simulation. *The Journal of Chemical Physics*, 122(2):024112, 2005.
- [155] Xinjun Peng, Wen Zhou, and Yifei Wang. Efficient binomial leap method for simulating chemical kinetics. *The Journal of Chemical Physics*, 126(22):224109, 2007.
- [156] André Leier, Tatiana T Marquez-Lago, and Kevin Burrage. Generalized binomial τ -leap method for biochemical kinetics incorporating both delay and intrinsic noise. *The Journal of Chemical Physics*, 128(20):205107, 2008.
- [157] Michel F Pettigrew and Haluk Resat. Multinomial tau-leaping method for stochastic kinetic simulations. *The Journal of Chemical Physics*, 126(8):084101, 2007.
- [158] Xiaodong Cai and Zhouyi Xu. K-leap method for accelerating stochastic simulation of coupled chemical reactions. *The Journal of Chemical Physics*, 126(7):74102–74102, 2007.
- [159] David F Anderson. Incorporating postleap checks in tau-leaping. *The Journal of Chemical Physics*, 128(5):054103, 2008.
- [160] Daniel T Gillespie. Simulation methods in systems biology. In *Formal Methods for Computational Systems Biology*, volume 5016, pages 125–167. Springer, 2008.
- [161] JR Bowen, Andreas Acrivos, and AK Oppenheim. Singular perturbation refinement to quasi-steady state approximation in chemical kinetics. *Chemical Engineering Science*, 18(3):177–188, 1963.
- [162] JAM Janssen. The elimination of fast variables in complex chemical reactions. II. Mesoscopic level (reducible case). *Journal of Statistical Physics*, 57(1-2):171–185, 1989.
- [163] Reinhart Heinrich and Stefan Schuster. *The regulation of cellular systems*. Springer Science & Business Media, 2012.
- [164] Leonor Michaelis and Maud L Menten. Die Kinetik der Invertinwirkung. *Biochemische Zeitschrift*, 49(333-369):352, 1913.
- [165] Lee A Segel and Marshall Slemrod. The quasi-steady-state assumption: a case study in perturbation. *SIAM Review*, 31(3):446–477, 1989.
- [166] George Edward Briggs and John Burdon Sanderson Haldane. A note on the kinetics of enzyme action. *Biochemical Journal*, 19(2):338, 1925.
- [167] John Goutsias. Quasiequilibrium approximation of fast reaction kinetics in stochastic biochemical systems. *The Journal of Chemical Physics*, 122(18):184102, 2005.
- [168] JAM Janssen. The elimination of fast variables in complex chemical reactions. III. Mesoscopic level (irreducible case). *Journal of Statistical Physics*, 57(1-2):187–198, 1989.
- [169] Didier Gonze, José Halloy, and Albert Goldbeter. Deterministic versus stochastic models for circadian rhythms. *Journal of Biological Physics*, 28(4):637–653, 2002.
- [170] Christopher V Rao and Adam P Arkin. Stochastic chemical kinetics and the quasi-steady-state assumption: application to the gillespie algorithm. *The Journal of Chemical Physics*, 118(11):4999–5010, 2003.
- [171] Kevin R Sanft, Daniel T Gillespie, and Linda R Petzold. Legitimacy of the stochastic Michaelis-Menten approximation. *IET Systems Biology*, 5(1):58–69, 2011.
- [172] Philipp Thomas, Arthur V Straube, and Ramon Grima. Communication: limitations of the stochastic quasi-steady-state approximation in open biochemical reaction networks. *The Journal of Chemical Physics*, 135(18):181103, 2011.
- [173] Jae K Kim, Krešimir Josić, and Matthew R Bennett. The relationship between stochastic and deterministic quasi-steady state approximations. *BMC Systems Biology*, 9(1):87, 2015.
- [174] Philipp Thomas, Arthur V Straube, and Ramon Grima. The slow-scale linear noise approximation: an accurate, reduced stochastic description of biochemical networks under timescale separation conditions. *BMC Systems Biology*, 6(1):39, 2012.

- [175] Yang Cao, Daniel T Gillespie, and Linda R Petzold. The slow-scale stochastic simulation algorithm. *The Journal of Chemical Physics*, 122(1):014116, 2005.
- [176] Howard Salis and Yiannis N Kaznessis. An equation-free probabilistic steady-state approximation: dynamic application to the stochastic simulation of biochemical reaction networks. *The Journal of Chemical Physics*, 123(21):214106, 2005.
- [177] E Weinan, Di Liu, and Eric Vanden-Eijnden. Nested stochastic simulation algorithm for chemical kinetic systems with disparate rates. *The Journal of Chemical Physics*, 123(19):194107, 2005.
- [178] A Samant and DG Vlachos. Overcoming stiffness in stochastic simulation stemming from partial equilibrium: a multiscale Monte Carlo algorithm. *The Journal of Chemical Physics*, 123(14):144114, 2005.
- [179] E Weinan, Di Liu, and Eric Vanden-Eijnden. Nested stochastic simulation algorithms for chemical kinetic systems with multiple time scales. *Journal of Computational Physics*, 221(1):158–180, 2007.
- [180] Carlos A Gómez-Uribe, George C Verghese, and Abraham R Tzafri. Enhanced identification and exploitation of time scales for model reduction in stochastic chemical kinetics. *The Journal of Chemical Physics*, 129(24):244112, 2008.
- [181] Simone Pigolotti and Angelo Vulpiani. Coarse graining of master equations with fast and slow states. *The Journal of Chemical Physics*, 128(15):154114, 2008.
- [182] Michael W Chevalier and Hana El-Samad. A rigorous framework for multiscale simulation of stochastic cellular networks. *The Journal of Chemical Physics*, 131(5):054102, 2009.
- [183] Simon L Cotter, Konstantinos C Zygalakis, Ioannis G Kevrekidis, and Radek Erban. A constrained approach to multiscale stochastic simulation of chemically reacting systems. *The Journal of Chemical Physics*, 135(9):094102, 2011.
- [184] Hye-Won Kang, Thomas G Kurtz, et al. Separation of time-scales and model reduction for stochastic reaction networks. *The Annals of Applied Probability*, 23(2):529–583, 2013.
- [185] Luca Bortolussi, Dimitrios Milios, and Guido Sanguinetti. Efficient stochastic simulation of systems with multiple time scales via statistical abstraction. In *International Conference on Computational Methods in Systems Biology*, pages 40–51. Springer, 2015.
- [186] Vahid Shahrezaei and Peter S Swain. Analytical distributions for stochastic gene expression. *Proceedings of the National Academy of Sciences*, 105(45):17256–17261, 2008.
- [187] Pavol Bokes, John R King, Andrew TA Wood, and Matthew Loose. Multiscale stochastic modelling of gene expression. *Journal of Mathematical Biology*, 65(3):493–520, 2012.
- [188] Bence Mélykúti, Joao P Hespanha, and Mustafa Khammash. Equilibrium distributions of simple biochemical reaction systems for time-scale separation in stochastic reaction networks. *Journal of The Royal Society Interface*, 11(97):20140054, 2014.
- [189] Nikola Popović, Carsten Marr, and Peter S Swain. A geometric analysis of fast-slow models for stochastic gene expression. *Journal of Mathematical Biology*, 72(1-2):87–122, 2016.
- [190] Eric L Haseltine and James B Rawlings. Approximate simulation of coupled fast and slow reactions for stochastic chemical kinetics. *The Journal of Chemical Physics*, 117(15):6959–6969, 2002.
- [191] Howard Salis and Yiannis Kaznessis. Accurate hybrid stochastic simulation of a system of coupled chemical or biochemical reactions. *The Journal of Chemical Physics*, 122(5):054103, 2005.
- [192] Slaven Peleš, Brian Munsky, and Mustafa Khammash. Reduction and solution of the chemical master equation using time scale separation and finite state projection. *The Journal of Chemical Physics*, 125(20):204104, 2006.

- [193] Jacek Puchalka and Andrzej M Kierzek. Bridging the gap between stochastic and deterministic regimes in the kinetic simulations of the biochemical reaction networks. *Biophysical Journal*, 86(3):1357–1372, 2004.
- [194] Chetan D Pahlajani, Paul J Atzberger, and Mustafa Khammash. Stochastic reduction method for biological chemical kinetics using time-scale separation. *Journal of Theoretical Biology*, 272(1):96–112, 2011.
- [195] Philipp Thomas, Ramon Grima, and Arthur V Straube. Rigorous elimination of fast stochastic variables from the linear noise approximation using projection operators. *Physical Review E*, 86(4):041110, 2012.
- [196] NA Sinitzyn, Nicolas Hengartner, and Ilya Nemenman. Adiabatic coarse-graining and simulations of stochastic biochemical networks. *Proceedings of the National Academy of Sciences*, 106(26):10546–10551, 2009.
- [197] Kevin Burrage, Tianhai Tian, and Pamela Burrage. A multi-scaled approach for simulating chemical reaction systems. *Progress in Biophysics and Molecular Biology*, 85(2):217–234, 2004.
- [198] Philipp Thomas, Nikola Popović, and Ramon Grima. Phenotypic switching in gene regulatory networks. *Proceedings of the National Academy of Sciences*, 111(19):6994–6999, 2014.
- [199] Karan Vasudeva and Upinder S Bhalla. Adaptive stochastic-deterministic chemical kinetic simulations. *Bioinformatics*, 20(1):78–84, 2004.
- [200] Aurélien Alfonsi, Eric Cances, Gabriel Turinici, Barbara Di Ventura, and Wilhelm Huisinga. Exact simulation of hybrid stochastic and deterministic models for biochemical systems. *Rr-5435, INRIA-Rocquencourt*, 2004.
- [201] Kouichi Takahashi, Kazunari Kaizu, Bin Hu, and Masaru Tomita. A multi-algorithm, multi-timescale method for cell simulation. *Bioinformatics*, 20(4):538–546, 2004.
- [202] Natasha A Neogi. Dynamic partitioning of large discrete event biological systems for hybrid simulation and analysis. In *Hybrid Systems: Computation and Control*, pages 463–476. Springer, 2004.
- [203] Martin Bentele and Roland Eils. General stochastic hybrid method for the simulation of chemical reaction processes in cells. In *Computational Methods in Systems Biology*, pages 248–251. Springer, 2004.
- [204] Leonard A Harris and Paulette Clancy. A “partitioned leaping” approach for multiscale modeling of chemical reaction dynamics. *The Journal of Chemical Physics*, 125(14):144107, 2006.
- [205] Andreas Hellander and Per Lötstedt. Hybrid method for the chemical master equation. *Journal of Computational Physics*, 227(1):100–122, 2007.
- [206] Alina Crudu, Arnaud Debussche, and Ovidiu Radulescu. Hybrid stochastic simplifications for multiscale gene networks. *BMC Systems Biology*, 3(1):1, 2009.
- [207] Stephan Menz, Juan C Latorre, Christof Schutte, and Wilhelm Huisinga. Hybrid Stochastic–Deterministic Solution of the Chemical Master Equation. *Multiscale Modeling & Simulation*, 10(4):1232–1262, 2012.
- [208] Arnab Ganguly, Derya Altintan, and Heinz Koepl. Jump-diffusion approximation of stochastic reaction dynamics: error bounds and algorithms. *Multiscale Modeling & Simulation*, 13(4):1390–1419, 2015.
- [209] Andrew Duncan, Radek Erban, and Konstantinos Zygalakis. Hybrid framework for the simulation of stochastic chemical kinetics. *arXiv preprint arXiv:1512.03988*, 2015.
- [210] Cosmin Safta, Khachik Sargsyan, Bert Debusschere, and Habib N Najm. Hybrid discrete/continuum algorithms for stochastic reaction networks. *Journal of Computational Physics*, 281:177–198, 2015.

- [211] Benjamin Hepp, Ankit Gupta, and Mustafa Khammash. Adaptive hybrid simulations for multiscale stochastic reaction networks. *The Journal of Chemical Physics*, 142(3):034118, 2015.
- [212] Alessio Angius, Gianfranco Balbo, Marco Beccuti, Enrico Bibbona, Andras Horvath, and Roberta Sirovich. Approximate analysis of biological systems by hybrid switching jump diffusion. *Theoretical Computer Science*, 587:49–72, 2015.
- [213] Tobias Jahnke. On reduced models for the chemical master equation. *Multiscale Modeling & Simulation*, 9(4):1646–1676, 2011.
- [214] Tobias Jahnke and Michael Kreim. Error bound for piecewise deterministic processes modeling stochastic reaction systems. *Multiscale Modeling & Simulation*, 10(4):1119–1147, 2012.
- [215] Alina Crudu, Arnaud Debussche, Aurélie Muller, Ovidiu Radulescu, et al. Convergence of stochastic gene networks to hybrid piecewise deterministic processes. *The Annals of Applied Probability*, 22(5):1822–1859, 2012.
- [216] Simon L Cotter and Radek Erban. Error analysis of diffusion approximation methods for multiscale systems in reaction kinetics. *SIAM Journal on Scientific Computing*, 38(1):B144–B163, 2016.
- [217] J Hasenauer, V Wolf, A Kazeroonian, and FJ Theis. Method of conditional moments (MCM) for the chemical master equation. *Journal of Mathematical Biology*, 69(3):687–735, 2014.
- [218] Mohammad Soltani, Cesar Augusto Vargas-Garcia, and Abhyudai Singh. Conditional Moment Closure Schemes for Studying Stochastic Dynamics of Genetic Circuits. *IEEE Transactions on Biomedical Circuits and Systems*, 9(4):518–526, 2015.
- [219] Stephen Smith, Claudia Cianci, and Ramon Grima. Model reduction for stochastic chemical systems with abundant species. *The Journal of Chemical Physics*, 143(21):214105, 2015.
- [220] Mukund Thattai and Alexander Van Oudenaarden. Intrinsic noise in gene regulatory networks. *Proceedings of the National Academy of Sciences*, 98(15):8614–8619, 2001.
- [221] Luca Bortolussi and Guido Sanguinetti. Learning and designing stochastic processes from logical constraints. In *QEST*, volume 8054, pages 89–105. Springer, 2013.
- [222] Botond Cseke, Manfred Opper, and Guido Sanguinetti. Approximate inference in latent Gaussian-Markov models from continuous time observations. In *Advances in Neural Information Processing Systems*, pages 971–979, 2013.
- [223] Christoph Zechner, Michael Unger, Serge Pelet, Matthias Peter, and Heinz Koepl. Scalable inference of heterogeneous reaction kinetics from pooled single-cell recordings. *Nature Methods*, 11(2):197–202, 2014.
- [224] Christopher M Bishop. Pattern recognition and machine learning. *Machine Learning*, 2006.
- [225] Lawrence Rabiner and B Juang. An introduction to hidden Markov models. *ieee assp magazine*, 3(1):4–16, 1986.
- [226] Rudolph Emil Kalman. A new approach to linear filtering and prediction problems. *Journal of basic Engineering*, 82(1):35–45, 1960.
- [227] Sylvia Frühwirth-Schnatter. Data augmentation and dynamic linear models. *Journal of Time Series Analysis*, 15(2):183–202, 1994.
- [228] Manfred Opper and Guido Sanguinetti. Variational inference for Markov jump processes. In *Advances in Neural Information Processing Systems*, pages 1105–1112, 2008.
- [229] Vinayak Rao and Yee Whye Teh. Fast MCMC sampling for Markov jump processes and extensions. *The Journal of Machine Learning Research*, 14(1):3295–3320, 2013.

- [230] Anastasis Georgoulas, Jane Hillston, and Guido Sanguinetti. Unbiased Bayesian inference for population Markov jump processes via random truncations. *Statistics and Computing*, 2016.
- [231] Richard J Boys, Darren J Wilkinson, and Thomas BL Kirkwood. Bayesian inference for a discretely observed stochastic kinetic model. *Statistics and Computing*, 18(2):125–135, 2008.
- [232] Andrew Golightly and Darren J Wilkinson. Bayesian parameter inference for stochastic biochemical network models using particle Markov chain Monte Carlo. *Interface Focus*, page rsfs20110047, 2011.
- [233] Andrew Golightly and Darren J Wilkinson. Bayesian inference for stochastic kinetic models using a diffusion approximation. *Biometrics*, 61(3):781–788, 2005.
- [234] Andreas Rutter and Manfred Opper. Efficient statistical inference for stochastic reaction processes. *Physical Review Letters*, 103(23):230601, 2009.
- [235] Vassilios Stathopoulos and Mark A Girolami. Markov chain Monte Carlo inference for Markov jump processes via the linear noise approximation. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 371(1984):20110541, 2013.
- [236] Paul Fearnhead, Vasileios Giagos, and Chris Sherlock. Inference for reaction networks using the linear noise approximation. *Biometrics*, 70(2):457–466, 2014.
- [237] Peter Milner, Colin S Gillespie, and Darren J Wilkinson. Moment closure based parameter inference of stochastic kinetic models. *Statistics and Computing*, 23(2):287–295, 2013.
- [238] Botond Cseke, David Schnoerr, Manfred Opper, and Guido Sanguinetti. Expectation propagation for continuous time stochastic processes. *Journal of Physics A: Mathematical and Theoretical*, 49(49):494002, 2016.
- [239] Christoph Zechner, Jakob Ruess, Peter Krenn, Serge Pelet, Matthias Peter, John Lygeros, and Heinz Koepl. Moment-based inference predicts bimodality in transient gene expression. *Proceedings of the National Academy of Sciences*, 109(21):8340–8345, 2012.
- [240] Florian Stimberg, Manfred Opper, Guido Sanguinetti, and Andreas Rutter. Inference in continuous-time change-point models. In *Advances in Neural Information Processing Systems*, pages 2717–2725, 2011.
- [241] Andrea Ocone and Guido Sanguinetti. Reconstructing transcription factor activities in hierarchical transcription network motifs. *Bioinformatics*, 27(20):2873–2879, 2011.
- [242] Andrea Ocone, Andrew J Millar, and Guido Sanguinetti. Hybrid regulatory models: a statistically tractable approach to model regulatory network dynamics. *Bioinformatics*, 29(7):910–916, 2013.
- [243] Chris Sherlock, Andrew Golightly, and Colin S Gillespie. Bayesian inference for hybrid discrete-continuous stochastic kinetic models. *Inverse Problems*, 30(11):114005, 2014.
- [244] Kirsty L Hey, Hiroshi Momiji, Karen Featherstone, Julian RE Davis, Michael RH White, David A Rand, and Bärbel Finkenstädt. A stochastic transcriptional switch model for single cell imaging data. *Biostatistics*, 16(4):655–669, 2015.
- [245] Juliane Liepe, Chris Barnes, Erika Cule, Kamil Erguler, Paul Kirk, Tina Toni, and Michael PH Stumpf. ABC-SysBio—approximate bayesian computation in python with GPU support. *Bioinformatics*, 26(14):1797–1799, 2010.
- [246] Yuhui Chen, Conor Lawless, Colin S Gillespie, Jake Wu, Richard J Boys, and Darren J Wilkinson. CaliBayes and BASIS: integrated tools for the calibration, simulation and storage of biological simulation models. *Briefings in Bioinformatics*, 11(3):278–289, 2010.
- [247] Masao Doi. Second quantization representation for classical many-particle system. *Journal of Physics A: Mathematical and General*, 9(9):1465, 1976.

- [248] Masao Doi. Stochastic theory of diffusion-controlled reaction. *Journal of Physics A: Mathematical and General*, 9(9):1479, 1976.
- [249] Radek Erban and S Jonathan Chapman. Stochastic modelling of reaction-diffusion processes: algorithms for bimolecular reactions. *Physical Biology*, 6(4):046001, 2009.
- [250] Jeroen S van Zon and Pieter R Ten Wolde. Simulating biochemical networks at the particle level and in time and space: Green’s function reaction dynamics. *Physical Review Letters*, 94(12):128103, 2005.
- [251] Aleksandar Donev, Vasily V Bulatov, Tomas Oppelstrup, George H Gilmer, Babak Sadigh, and Malvin H Kalos. A first-passage kinetic Monte Carlo algorithm for complex diffusion-reaction systems. *Journal of Computational Physics*, 229(9):3214–3236, 2010.
- [252] Samuel A Isaacson. Relationship between the reaction-diffusion master equation and particle tracking models. *Journal of Physics A: Mathematical and Theoretical*, 41(6):065003, 2008.
- [253] Jin Fu, Sheng Wu, Hong Li, and Linda R Petzold. The time dependent propensity function for acceleration of spatial stochastic simulation of reaction-diffusion systems. *Journal of Computational Physics*, 274:524–549, 2014.
- [254] Samuel A Isaacson. The reaction-diffusion master equation as an asymptotic approximation of diffusion to a small target. *SIAM Journal on Applied Mathematics*, 70(1):77–111, 2009.
- [255] Jana Lipková, Konstantinos C Zygalakis, S Jonathan Chapman, and Radek Erban. Analysis of brownian dynamics simulations of reversible bimolecular reactions. *SIAM Journal on Applied Mathematics*, 71(3):714–730, 2011.
- [256] Stefan Hellander, Andreas Hellander, and Linda Petzold. Reaction-diffusion master equation in the microscopic limit. *Physical Review E*, 85(4):042901, 2012.
- [257] Anel Mahmutovic, David Fange, Otto G Berg, and Johan Elf. Lost in presumption: stochastic reactions in spatial models. *Nature methods*, 9(12):1163–1166, 2012.
- [258] Samuel A Isaacson. A convergent reaction-diffusion master equation. *The Journal of Chemical Physics*, 139(5):054101, 2013.
- [259] PR Taylor, RE Baker, and CA Yates. Deriving appropriate boundary conditions, and accelerating position-jump simulations, of diffusion using non-local jumping. *Physical biology*, 12(1):016006, 2014.
- [260] Matthew Scott, Francis J Poulin, and Herbert Tang. Approximating intrinsic noise in continuous multispecies models. In *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, page rspa20100275. The Royal Society, 2010.
- [261] Malbor Asllani, Tommaso Biancalani, Duccio Fanelli, and Alan J McKane. The linear noise approximation for reaction-diffusion systems on networks. *The European Physical Journal B*, 86(11):1–10, 2013.
- [262] Stephen Smith, Claudia Cianci, and Ramon Grima. Analytical approximations for spatial stochastic gene expression in single cells and tissues. *Journal of The Royal Society Interface*, 13(118):20151051, 2016.
- [263] Geoffrey R Holmes, Sean R Anderson, Giles Dixon, Anne L Robertson, Constantino C Reyes-Aldasoro, Stephen A Billings, Stephen A Renshaw, and Visakan Kadirkamanathan. Repelled from the wound, or randomly dispersed? Reverse migration behaviour of neutrophils characterized by dynamic modelling. *Journal of The Royal Society Interface*, page rsif20120542, 2012.
- [264] Michael A Dewar, Visakan Kadirkamanathan, Manfred Opper, and Guido Sanguinetti. Parameter estimation and inference for stochastic reaction-diffusion systems: application to morphogenesis in *D. melanogaster*. *BMC Systems Biology*, 4(1):1, 2010.

- [265] Andreas Ruttor and Manfred Opper. Approximate parameter inference in a stochastic reaction-diffusion model. In *International Conference on Artificial Intelligence and Statistics*, volume 9, pages 669–676, 2010.
- [266] Noel Cressie and Christopher K Wikle. *Statistics for spatio-temporal data*. John Wiley & Sons, 2015.