



SCUOLA INTERNAZIONALE SUPERIORE DI STUDI AVANZATI

SISSA Digital Library

Model reduction for the Chemical Master Equation: An information-theoretic approach

Original

Model reduction for the Chemical Master Equation: An information-theoretic approach / Ocal, Kaan; Sanguinetti, Guido; Grima, Ramon. - In: JOURNAL OF CHEMICAL PHYSICS ONLINE. - ISSN 1089-7690. - 158:11(2023), pp. 1-27. [10.1063/5.0131445]

Availability:

This version is available at: 20.500.11767/132230 since: 2023-05-11T12:26:37Z

Publisher:

Published

DOI:10.1063/5.0131445

Terms of use:

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

Publisher copyright

note finali coverpage

(Article begins on next page)

Model Reduction for the Chemical Master Equation: an Information-Theoretic Approach

Kaan Öcal^{1,2}, Guido Sanguinetti³, and Ramon Grima^{2,†}

¹School of Informatics, University of Edinburgh, Edinburgh EH8 9AB, UK

²School of Biological Sciences, University of Edinburgh, Edinburgh EH9 3JH, UK

³Scuola Internazionale Superiore di Studi Avanzati, 34136 Trieste, Italy

[†]Corresponding author: ramon.grima@ed.ac.uk

Abstract

The complexity of mathematical models in biology has rendered model reduction an essential tool in the quantitative biologist's toolkit. For stochastic reaction networks described using the Chemical Master Equation, commonly used methods include time-scale separation, the Linear Mapping Approximation and state-space lumping. Despite the success of these techniques, they appear to be rather disparate and at present no general-purpose approach to model reduction for stochastic reaction networks is known. In this paper we show that most common model reduction approaches for the Chemical Master Equation can be seen as minimising a well-known information-theoretic quantity between the full model and its reduction, the Kullback-Leibler divergence defined on the space of trajectories. This allows us to recast the task of model reduction as a variational problem that can be tackled using standard numerical optimisation approaches. In addition we derive general expressions for the propensities of a reduced system that generalise those found using classical methods. We show that the Kullback-Leibler divergence is a useful metric to assess model discrepancy and to compare different model reduction techniques using three examples from the literature: an autoregulatory feedback loop, the Michaelis-Menten enzyme system and a genetic oscillator.

Keywords: Chemical Master Equation · Model Reduction · Systems Biology

1 Introduction

Stochastic biochemical reaction networks such as those involved in gene expression, immune response or cellular signalling [1–4] are often described using the Chemical Master Equation (CME). The CME describes the dynamics of biochemical processes on a mesoscopic level, viewing them as a discrete collection of molecules interacting and undergoing reactions stochastically; as such it is generally considered more accurate than continuum approximations such as rate equations and the Chemical Langevin Equation [4]. Despite its explanatory power, the CME poses significant analytical and computational difficulties to modellers that have limited its use in practice. Closed-form solutions to the CME are difficult to obtain and are only known for a small number of biologically relevant systems, and solving the CME numerically requires using approximations such as the Finite State Projection (FSP) [5]. Numerical approaches tend to scale poorly with the number of species and reactions present in a system, and as a result there is significant interest in finding ways to simplify a description of a stochastic reaction network that make it easier to analyse and study - this is the goal of model reduction.

Model reduction for deterministic and continuum-limit models in biology is an active research topic [6, 7], but very few existing methods can be applied to the discrete, stochastic setting of the CME. The Quasi-Steady State Approximation (QSSA) is perhaps the best known technique, first considered in the stochastic case in [8]. Here the system is partitioned into 'slow' and 'fast' species such that the fast species evolve very quickly on the timescale of the slow species. On the slow timescale the states of the fast species can therefore be approximated by their steady-state value (conditioned on the slow species), effectively allowing a description of the system in terms of the

slow species only. By its nature the QSSA is only applicable to systems with a clear separation of timescales between species, the existence of which cannot always be established. The QSSA for stochastic systems generally requires more stringent conditions than in the deterministic case, but the exact validity conditions are not well-understood [9–14].

Similar to the QSSA is the Quasiequilibrium Approximation (QEA), which was first considered in [15, 16] for stochastic reaction networks. Here the reaction network is decomposed into ‘slow’ and ‘fast’ *reactions*, and the fast reactions are assumed to equilibrate rapidly on the timescale of the slow reactions. Similar to the QSSA, the QEA can be used to reduce the number of species and reactions in a system, but it relies on the existence of a clear timescale separation between reactions, which is not always present for large systems with many distinct reactions. Much like the QSSA, the validity of the QEA for systems without the appropriate timescale separation has not been generally established, and from the asymptotic nature of the descriptions it is not usually possible to quantify the approximation error. Despite this, both the QSSA and the QEA are by far the most commonly used model-reduction technique for chemical reaction networks owing to their physical interpretability and analytical tractability, most famously in the Michaelis-Menten model of enzyme kinetics.

A distinct approach for model reduction with the Chemical Master Equation is state-space lumping, which originates from the theory of finite Markov chains, see e.g. [17]. Here different states in a system are lumped together such that the coarse-grained system is still Markovian and can be modelled using the CME. For a typical biochemical reaction network it may not be possible to perform any nontrivial lumping while preserving Markovianity, whence approximate lumping methods have been considered e.g. in [18–20]. Here the coarse-grained system is approximated by a Markov process, and the approximation error quantified using the KL divergence between the original model and a lift of the approximation to the state space of the original model. State-space lumping for the CME often occurs in the context of the QEA, as states related by fast reactions are effectively lumped together, or averaged [21–23]. For this reason we will not consider this approach separately, although many of our considerations, such as the optimal form of the lumped propensity functions, extend to state-space lumping.

Finally, a more recent model reduction technique specifically for gene expression systems is the Linear Mapping Approximation (LMA) [24]. The LMA replaces bimolecular reactions of the form $G + P \rightarrow (\dots)$, where G is a binary species such as a gene, by a linear approximation where P is replaced by its mean conditioned on $[G] = 1$. While the LMA does not reduce the number of species or reactions, reaction networks with linear reactions are generally easier to analyse: their moments can be computed exactly, and closed-form solutions are known for many cases [25–28].

At a first glance these approaches - timescale separation, state space lumping and the Linear Mapping Approximation - seem rather disparate, and it is unclear what, if any, relationships there are between these approaches. In this paper we show that all of these methods can be viewed as minimising a natural information theoretic quantity, the Kullback-Leibler (KL) divergence, between the full and reduced models. In particular they can be seen as maximal likelihood approximations of the full model, and one can assess the quality of the approximation in terms of the resulting KL divergence. Based on these results we show how the KL divergence can be estimated and minimised numerically, therefore providing an automated method to choose between different candidate reductions of a given model in situations where none of the above model reduction techniques are classically applicable.

The KL divergences we consider in this paper are computed on the space of trajectories, and as such include both static information and dynamical information, in contrast to purely distribution-matching approaches. The KL divergence and similar information-theoretic measures between continuous-time Markov chains have previously been considered in [29, 30] in the context of variational inference




Figure 1: Model reduction for the Chemical Master Equation. **(a)** Model reduction approximates a high-dimensional model q by a lower-dimensional version. Since the direct projection \tilde{q} of the full model is not easy to describe, we approximate it using a family of tractable candidate models: in this paper, the approximation p is described by the CME. **(b)** Comparison of the full state space of a system, consisting of two species S_1 and S_2 , and a reduced state space containing S_1 only. Species which are not deemed essential can be discarded in the reduction and become unobserved variables. The dynamics of the original system involves all species, whereas the reduced model aims at an effective description only in terms of the reduced species. **(c)** Sample trajectory for a one-dimensional system, defined by the sequence n_0, n_1, \dots of states visited and the corresponding jump times t_1, t_2, \dots (or alternatively the waiting times τ_0, τ_1, \dots).

(with the true model and the approximation reversed compared to our approach), in [31] to obtain approximate non-Markovian reductions, in [32] to analyse information flow for stochastic reaction networks and in [33] in quantifying model discrepancy for Markovian agent-based models.

In Section 2 we introduce the mathematical framework in which we consider model reduction for the Chemical Master Equation, based on KL divergences between continuous-time Markov chains. We show how the KL divergence can be minimised analytically in some important cases, recovering standard results in the literature and providing a mathematical justification for commonly used mean-field arguments as in the QSSA, the QEA and the LMA. We furthermore provide numerical algorithms for estimating as well as minimising the KL divergence in cases where analytical solutions are not available. In Section 3 we illustrate the use of KL divergences as a metric for approximation quality using three biologically relevant examples: an autoregulatory feedback loop exhibiting critical behaviour, Michaelis-Menten enzyme kinematics, where we reanalyse the QSSA and the QEA, and an oscillatory genetic feedback system taken from [11], for which we compare different reductions using our approach. Finally in Section 4 we discuss our observations and how our approach could be used as a stepping-stone towards automated reduction of complex biochemical reaction pathways.

2 Methods

2.1 Stochastic Reaction Networks

The Chemical Master Equation describes a biochemical reaction network as a stochastic process on a discrete state space \mathcal{X} . We will use the letter q to denote such a stochastic process, which for

the purposes of this paper can be seen as a probability distribution over *trajectories* on \mathcal{X} . For a biochemical reaction network the state space will be $\mathcal{X} = \mathbb{N}^s$, where s is the number of species in the system: every state consist of a tuple $\mathbf{n} = (n_1, \dots, n_s)$ of s integers describing the abundances of each species.

Since the model q can consist of many species interacting in complicated ways, we often want to find a reduction that is more tractable, yet approximates q as closely as possible. The reduced model, which we will call p , should be of the same form as q , i.e. described by the CME, but will typically involve fewer species and simpler reactions. In particular p can be defined on a lower-dimensional state space $\tilde{\mathcal{X}}$. A state \mathbf{n} in the original model can then be described by its projection $\tilde{\mathbf{n}}$ onto this lower-dimensional space, together with some unobserved components, which we will denote \mathbf{z} . See Fig. 1a and b for an illustration.

In this paper we assume that the basic structure of p is known *a priori*, i.e. the species and reactions we wish to retain are fixed. Our approach to model reduction therefore consists in finding the optimal propensity functions for the reduced model, and we shall see how this can give rise to various known approximations such as the QSSA, the QEA or the LMA depending on what reductions are performed. We will return to the related problem of choosing the structure of the reduced model p in the discussion.

The original model q defines a probability distribution over trajectories in \mathcal{X} , and projecting each trajectory onto the chosen reduced state space we get the (exact) projection of q onto this space, which we denote \tilde{q} (Fig. 1a). This is a stochastic process on $\tilde{\mathcal{X}}$ that is generally not Markovian and thus cannot be modelled using the CME. We aim to find a tractable approximation p to \tilde{q} that can be described using the CME, and we will do this by minimising the KL divergence $\text{KL}(\tilde{q} \| p)$ between the two models on the space of trajectories. Several well-known examples of model reduction for the CME are illustrated in Fig. 2.

Jumps in q come in two kinds: those that affect the observed species $\tilde{\mathbf{n}}$, which we will call visible jumps, and those that only change \mathbf{z} , which we call hidden. The jumps in \tilde{q} correspond to visible jumps in q . In the context of the CME, jumps are typically grouped into reactions with fixed stoichiometry, often also called reaction channels, and we can similarly distinguish visible and hidden reactions in q . We will always assume that different reactions have different stoichiometries, so that every jump in q and p corresponds to a unique reaction. Reactions with the same stoichiometry can always be combined by summing their propensities.

We introduce some more notation at this point, which is summarised in Table 1 and illustrated in Fig. 1c. A single realisation, or trajectory, of q is defined by the sequence of states $\mathbf{n}_0, \mathbf{n}_1, \mathbf{n}_2, \dots \in \mathcal{X}$ visited and jump times $0 < t_1 < t_2 < \dots$. We will write $\mathbf{n}_{[0,T]} = \{\mathbf{n}(t)\}_{0 \leq t \leq T}$ for a trajectory, where $\mathbf{n}(0) = \mathbf{n}_0$ and $\mathbf{n}(t) = \mathbf{n}_i$ for $t_i \leq t < t_{i+1}$, and denote by $\tau_i = t_{i+1} - t_i$ the waiting times between jumps.

For a continuous-time Markov process p , e.g. one defined using the CME, we denote the transition rate from state \mathbf{n} to $\mathbf{m} \neq \mathbf{n}$ by $p_{\mathbf{m} \leftarrow \mathbf{n}}$. We let $p_{\leftarrow \mathbf{n}} = \sum_{\mathbf{m} \neq \mathbf{n}} p_{\mathbf{m} \leftarrow \mathbf{n}}$ be the total transition rate out of \mathbf{n} . The transition probabilities at \mathbf{n} are then given by $p_{\mathbf{m}|\mathbf{n}} = p_{\mathbf{m} \leftarrow \mathbf{n}} / p_{\leftarrow \mathbf{n}}$. For completeness we define $p_{\mathbf{n}|\mathbf{n}} := 0$.

2.2 Minimising the KL Divergence

The Kullback-Leibler Divergence between two distributions \tilde{q} and p is defined as

$$\text{KL}(\tilde{q} \| p) = \int \tilde{q}(x) \log \frac{\tilde{q}(x)}{p(x)} dx = \mathbb{E}_{\tilde{q}} [\log \tilde{q}(x)] - \mathbb{E}_{\tilde{q}} [\log p(x)] \quad (1)$$

$$= \mathbb{E}_{\tilde{q}} [\log \tilde{q}(x)] + H(\tilde{q}; p) \quad (2)$$




Figure 2: Common model reduction techniques for the Chemical Master Equation. **(a)** The QSSA eliminates intermediate species which evolve on a faster timescale than others, replacing them by their steady-state values. In the case of the pictured Michaelis-Menten enzyme system this is often applied to the enzyme E and the substrate-enzyme complex ES . **(b)** The QEA is an analogue of the QSSA that can be applied when a reaction and its reverse equilibrate rapidly on timescales of interest. This can occur e.g. when a gene switches rapidly between states. **(c)** The LMA replaces protein-gene binding reactions by effective unimolecular reactions, where the concentration of proteins is approximated by its mean conditioned on the gene state. While this does not remove any species from the system, it considerably simplifies the propensities of the reactions.

where the quantity $H(\tilde{q}; p)$ is known as the cross-entropy:

$$H(\tilde{q}; p) = -\mathbb{E}_{\tilde{q}}[\log p(x)] \quad (3)$$

Minimising the KL divergence with respect to p is therefore equivalent to minimising the cross-entropy. Looking at the cross-entropy we see that minimising the KL divergence is a standard inference problem: maximise the (average) log-likelihood of samples drawn from \tilde{q} .

In this paper samples from \tilde{q} and p are trajectories. We therefore need to compute the log-likelihood of an arbitrary trajectory $\mathbf{n}_{[0,T]}$ under p , which by assumption is a stochastic reaction network described by the CME. The log-likelihood of a trajectory can be computed as a product of the transition probabilities as is done in Appendix A. For a trajectory visiting the states $\mathbf{n}_0, \mathbf{n}_1, \dots, \mathbf{n}_k$ with waiting times $\tau_0, \tau_1, \dots, \tau_k$ the log-likelihood is given by

$$\log p(\mathbf{n}_{[0,T]}) = \log p_0(\mathbf{n}_0) - \sum_{i=0}^k \tau_i p_{\leftarrow \mathbf{n}_i} + \sum_{i=1}^k \log p_{\mathbf{n}_i \leftarrow \mathbf{n}_{i-1}} \quad (4)$$

This expression can also be obtained by considering the probability of each transition in the Gillespie

Symbol	Explanation	Symbol	Explanation
q	Full model	$q_{m \leftarrow n}$	Transition rate
\tilde{q}	Projected model	$q_{\leftarrow n}$	Total transition rate
p	Reduced model	$q_{m n}$	Transition probability
\mathbf{n}	Full state vector	$q_0(\mathbf{n})$	Initial distribution
$\tilde{\mathbf{n}}$	Reduced state vector	$q_t(\mathbf{n})$	Single-time marginal
\mathbf{z}	Unobserved state vector		

Table 1: Notation used in this paper.

algorithm. In the presence of time-dependent rates the equation generalises to

$$\log p(\mathbf{n}_{[0,T]}) = \log p_0(\mathbf{n}_0) - \sum_{i=0}^k \int_{t_i}^{t_{i+1}} p_{\leftarrow \mathbf{n}_i}(s) ds + \sum_{i=1}^k \log p_{\mathbf{n}_i \leftarrow \mathbf{n}_{i-1}}(t_i), \quad (5)$$

where t_1, \dots, t_k are the jump times and we set $t_{k+1} = T$ for convenience.

Note that computing the likelihood of a fully observed trajectory is easier than computing the likelihood of a trajectory sampled at discrete time points, which is the form of data typically observed in biological experiments. The difference is that computing the likelihood in the latter case requires integrating over all possible courses of the trajectory between observations; this integral is computed implicitly by the Chemical Master Equation, which is generally hard to solve. For a fully observed trajectory there are no hidden variables to integrate out, which considerably simplifies the task of computing likelihoods.

Returning to the problem of minimizing the KL divergence (1), or equivalently the cross entropy, we need to compute expectations over \tilde{q} that are not in general available in closed form. We can instead approximate them by simulating N trajectories $\tilde{\mathbf{n}}_{[0,T]}^{(i)}$, $i = 1, \dots, N$, from \tilde{q} and computing the estimate

$$\widehat{H}(\tilde{q}; p_\theta) = -\frac{1}{N} \sum_{i=1}^N \log p(\tilde{\mathbf{n}}_{[0,T]}^{(i)}). \quad (6)$$

We can then minimise this with respect to the reduced model parameters by gradient descent. Formulae for the gradients can be computed by differentiating Eqs. (4) or (5) by hand, or by using automatic differentiation software. Minimising the estimated cross-entropy via gradient descent yields a probabilistic estimate for the optimal parameters, which can be expected to converge to the true solution as $N \rightarrow \infty$. Summarising the contents of this section we arrive at Algorithm 1 for automatically fitting the optimally reduced model.

Algorithm 1 Model reduction via stochastic gradient descent.

Inputs: number of simulations N , simulation length T , full model q , model family $\{p_\theta\}$, learning rate η

Output: reduced model parameters $\hat{\theta}$

```

for all  $i = 1, \dots, N$  do
  sample  $\mathbf{n}_{[0,T]}^{(i)}$  from  $q$ 
  project  $\mathbf{n}_{[0,T]}^{(i)}$  to  $\tilde{\mathbf{n}}_{[0,T]}^{(i)}$ 
end
initialize parameters  $\hat{\theta}$ 
while not converged do
  compute  $\widehat{H}(\tilde{q}; p_{\hat{\theta}})$  using (5), (6)
  compute  $\nabla_{\theta} \widehat{H}(\tilde{q}; p_{\hat{\theta}})$ 
   $\hat{\theta} \leftarrow \hat{\theta} - \eta \nabla_{\theta} \widehat{H}(\tilde{q}; p_{\hat{\theta}})$ 
return  $\hat{\theta}$ 

```

Example: Telegraph Model

We will illustrate the above by means of a well-known example, the telegraph model of gene transcription [34], simplified to neglect degradation and illustrated in Fig. 3. The full model q consists

of three species: a gene in the on state (G_{on}) and the off state (G_{off}), together with mRNA (M). The gene switches between both states with constant activation rate σ_{on} and inactivation rate σ_{off} , and in the active state produces mRNA with rate ρ_{on} . The reduced model p consists only of mRNA (M), which is produced at constant rate ρ_{eff} .

In this example a state in the full model can be represented as $\mathbf{n} = (g, m)$, where g is the gene state (on or off) and the m is the number of mRNA present, while a reduced state $\tilde{\mathbf{n}} = (m)$ is given by the number of mRNA molecules only, with the unobserved component $\mathbf{z} = (g)$ being the gene state. The only reaction in q that descends to \tilde{q} is mRNA production, whereas the two gene switching reactions are not observed. Note that the projection \tilde{q} of the telegraph model is not Markovian, as the instantaneous mRNA production rate depends on the current gene state g , but we can sample from \tilde{q} by simulating q and discarding the information about the gene state.

The reduced model p is a Poisson process whose only parameter is the production rate ρ_{eff} . The only trajectories possible under p are those starting at $m = 0$ and increasing by one at every jump. The log-likelihood of such a trajectory $\tilde{\mathbf{n}}_{[0,T]}$ under p is given by

$$\log p(\tilde{\mathbf{n}}_{[0,T]}) = \tilde{\mathbf{n}}(T) \log \rho_{\text{eff}} - \rho_{\text{eff}} T, \quad (7)$$

where $\tilde{\mathbf{n}}(T)$ is the total number of mRNA produced up to time T , as can be verified using Eq. (4). The log-likelihood of any other trajectory is $-\infty$.

Note that Eq. (7) describes the likelihood for an entire trajectory, as opposed to the likelihood of observing \mathbf{n}_T molecules at time T , which follows a Poisson distribution with rate $\rho_{\text{eff}} T$ and can be obtained by integrating Eq. (4) over all trajectories that end at the same \mathbf{n}_T . Indeed, conditioned on observing \mathbf{n}_T mRNA molecules at time T their individual production times are uniformly and independently distributed on $[0, T]$ by the properties of the Poisson process, and integrating Eq. (7) over all possible combinations of production times we recover the usual Poisson likelihood, keeping in mind that any permutation of the production times yields the same trajectory.

The reduced model p can model every trajectory obtained from \tilde{q} , so the log-likelihood of any sample from \tilde{q} is finite. The mRNA transcription reactions in both models correspond. If q were to include e.g. mRNA degradation as is usual in the literature, some trajectories from the full model would feature decreases in $\tilde{\mathbf{n}}$ and be impossible under p ; in this case the KL divergence would be infinite, which signals that the reduced model p is not appropriate and needs to be extended to include degradation.

For this simple example the likelihood of a reduced trajectory under p only depends on the total mRNA produced. We can therefore compute the cross-entropy explicitly:

$$H(\tilde{q}; p) = \mathbb{E}_q [\tilde{\mathbf{n}}(T)] \log \rho_{\text{eff}} - \rho_{\text{eff}} T. \quad (8)$$

Minimising this with respect to ρ_{eff} yields the optimum

$$\rho_{\text{eff}}^* = \frac{1}{T} \mathbb{E}_q [\tilde{\mathbf{n}}(T)], \quad (9)$$

which is the average mRNA production rate on the interval $[0, T]$. Thus the optimal approximation to the telegraph model is obtained by setting the mRNA production rate to its mean value. This is the result we obtain by applying the QEA to this system (see 2.5.2), which suggests that the approximation will be better if σ_{on} and σ_{off} are large compared to ρ_{on} . We will see how to evaluate the approximation quality in 2.4.

2.3 Analysing the KL Divergence

Having derived a numerical procedure to minimise the KL divergence (1), or equivalently the cross-entropy (3), we analyse these quantities to understand how information-theoretic model reduction works. As shown in Appendix B, for a Markovian reduced model p the cross-entropy can be written as

$$H(\tilde{q}; p)_{[0, T]} = H(\tilde{q}_0; p_0) + \int_0^T \sum_{\tilde{\mathbf{n}}} \tilde{q}_t(\tilde{\mathbf{n}}) \left(p_{\leftarrow \tilde{\mathbf{n}}}(t) - \sum_{\tilde{\mathbf{m}} \neq \tilde{\mathbf{n}}} \tilde{q}_{\tilde{\mathbf{m}} \leftarrow \tilde{\mathbf{n}}}(t) \log p_{\tilde{\mathbf{m}} \leftarrow \tilde{\mathbf{n}}}(t) \right) dt \quad (10)$$

$$= H(\tilde{q}_0; p_0) + \int_0^T H(\tilde{q}; p)_t dt, \quad (11)$$

with the instantaneous cross-entropy rate at time t defined as

$$H(\tilde{q}; p)_t = - \sum_{\tilde{\mathbf{n}}} \tilde{q}_t(\tilde{\mathbf{n}}) \left(\sum_{\tilde{\mathbf{m}} \neq \tilde{\mathbf{n}}} \tilde{q}_{\tilde{\mathbf{m}} \leftarrow \tilde{\mathbf{n}}}(t) \log p_{\tilde{\mathbf{m}} \leftarrow \tilde{\mathbf{n}}}(t) - p_{\leftarrow \tilde{\mathbf{n}}}(t) \right). \quad (12)$$

Here we use the effective transition rates

$$\tilde{q}_{\tilde{\mathbf{m}} \leftarrow \tilde{\mathbf{n}}}(t) = \lim_{\delta t \rightarrow 0} \frac{1}{\delta t} P(\tilde{q}(t + \delta t) = \tilde{\mathbf{m}} | \tilde{q}(t) = \tilde{\mathbf{n}}, q_0) \quad (\tilde{\mathbf{m}} \neq \tilde{\mathbf{n}}). \quad (13)$$

If \tilde{q} is not Markovian, the transition probability from a state $\tilde{\mathbf{n}}$ at time t will be affected by the history of the process, and hence on the initial distribution q_0 . If \tilde{q} is Markovian, Eq. (13) reduces to the classical transition rate which is independent of q_0 .

Looking at Eq. (10) we make two important observations. As intimated in the discussion of the telegraph example of the previous section, if any trajectory under \tilde{q} has jumps which are not allowed under p , i.e. if $\tilde{q}_{\tilde{\mathbf{m}} \leftarrow \tilde{\mathbf{n}}} \neq 0$ while $p_{\tilde{\mathbf{m}} \leftarrow \tilde{\mathbf{n}}} = 0$, then the cross-entropy (10) will be ∞ : the reduced model p is not flexible enough to model \tilde{q} . On the other hand, if p contains transitions which are impossible under \tilde{q} , i.e. if $p_{\tilde{\mathbf{m}} \leftarrow \tilde{\mathbf{n}}} \neq 0$ while $\tilde{q}_{\tilde{\mathbf{m}} \leftarrow \tilde{\mathbf{n}}} = 0$, then these can only increase the cross-entropy (10). This means that the optimal reduced model p does not contain transitions beyond those in \tilde{q} , and in the context of the CME we can therefore assume that p and \tilde{q} have the same reactions, with different propensities (Markovian for p , not necessarily for \tilde{q}).

If the i -th reaction in p has propensity function $\rho_i(\tilde{\mathbf{n}}; t)$ and net stoichiometry s_i we obtain the following decomposition of the cross-entropy rate at time t :

$$H(\tilde{q}; p)_t = - \sum_i \sum_{\tilde{\mathbf{n}}} \tilde{q}_t(\tilde{\mathbf{n}}) [\tilde{q}_{\tilde{\mathbf{n}} + s_i \leftarrow \tilde{\mathbf{n}}} \log \rho_i(\tilde{\mathbf{n}}; t) - \rho_i(\tilde{\mathbf{n}}; t)], \quad (14)$$

where the first sum is over all reactions in p , or equivalently all visible reactions in q . The total cross-entropy is obtained by integrating Eq. (14) over $[0, T]$, and we can find the optimal p by optimising the cross-entropy for each reaction separately.

We can minimize Eq. (14) analytically if the full model q is Markovian. Assume there is precisely one reaction in q with net stoichiometry s_i in the projection, and let σ_i be its propensity function. Differentiating the above equation with respect to $\rho_i(\tilde{\mathbf{n}}; t)$ and setting the derivative to zero we obtain

$$\rho_i^*(\tilde{\mathbf{n}}; t) = \sum_z q_t(z | \tilde{\mathbf{n}}) \sigma_i(\tilde{\mathbf{n}}, z; t) = \mathbb{E}_z[\sigma_i(\tilde{\mathbf{n}}, z; t) | \tilde{\mathbf{n}}; t]. \quad (15)$$

The optimal propensity for a reaction under p is the expected propensity under the original model conditioned on the observed state $\tilde{\mathbf{n}}$. In particular, if the propensity of the original reaction does not depend on unobserved species it can be taken over directly.




Figure 3: Model reduction illustrated using the telegraph model without degradation. **(a)** Schematic of the model. **(b)** The reduced model is equivalent to a Poisson process with rate ρ_{eff} . **(c)** Example trajectory from the telegraph model, showing mRNA numbers (top) and the gene state (bottom) in time. If we only observe mRNA numbers we can infer the current gene state by filtering, obtaining the probability for the gene being on at a given time (blue line). When mRNA numbers increase the on probability jumps to 1 since transcription only happens in that state. **(d)** KL divergence rate between the telegraph model and its Poisson reduction, for various choices of the switching timescale $\tau = \sigma_{\text{on}} + \sigma_{\text{off}}$ and on probability $p_{\text{on}} = \sigma_{\text{on}} / (\sigma_{\text{on}} + \sigma_{\text{off}})$, assuming a fixed transcription rate $\rho_{\text{on}} = 1$. The Poisson approximation becomes more accurate as the switching timescale increases compared to the transcription timescale.

In practice we often place constraints on the reduced propensities $\rho_i(\tilde{\mathbf{n}}; t)$ such as time-homogeneity and mass-action kinetics, which result in constrained optima. For example, if ρ_i is taken to be independent of time we have to integrate Eq. (14) over $[0, T]$, and as $T \rightarrow \infty$ it can be verified that the total cross-entropy is minimised for

$$\rho_i^*(\tilde{\mathbf{n}}) = \sum_{\mathbf{z}} q_{\infty}(\mathbf{z} | \tilde{\mathbf{n}}) \sigma_i(\tilde{\mathbf{n}}, \mathbf{z}) = \mathbb{E}_{\mathbf{z}}[\sigma_i(\tilde{\mathbf{n}}, \mathbf{z}) | \tilde{\mathbf{n}}; t = \infty], \quad (16)$$

which is the steady-state version of Eq. (15).

For other forms of propensity functions the optimum can generally be obtained by a similar averaging procedure. In all cases the main difficulties lie in estimating the relevant conditional expectations, which can be done numerically using the Monte Carlo approach presented in the previous section. Eq. (15) is perhaps the central result of this paper, and we will see that it coincides with the propensities obtained using the QSSA, the QEA and the LMA in when these are applied. These methods can therefore all be seen as special cases of our variational approach based on minimising the KL divergence.

Example: Telegraph Model, cont.

Consider the telegraph model from the previous section. If we make the reduced model p more flexible by allowing a nonlinear and time-dependent propensity $\rho(m; t) = \rho_{\text{eff}}(m; t)$, then the cross-entropy between \tilde{q} and p can be computed as

$$H(\tilde{q}; p) = \int_0^T \sum_m \tilde{q}_t(m) (\tilde{q}_{m+1 \leftarrow m}(t) \log \rho_{\text{eff}}(m; t) - \rho_{\text{eff}}(m; t)) dt. \quad (17)$$

Note that both \tilde{q} and p have the same initial conditions. The effective propensities of \tilde{q} can be computed as

$$\tilde{q}_{m+1 \leftarrow m}(t) = \rho_{\text{on}} q_t(g = \text{on} | m), \quad (18)$$

which is ρ_{on} weighed by the probability that the gene is on at time t given that there are m mRNA molecules present. This is the optimal propensity for the reduced model p in accordance with Eq. (15), featuring a nonlinear dependence on m , and this can be shown to be the optimal Markovian approximation to \tilde{q} .

If we require mass-action kinetics for the reduced model, i.e. if we let $\rho(m; t) = m \rho_{\text{eff}}(t)$ for an effective rate constant $\rho_{\text{eff}}(t)$, minimising the cross-entropy (17) yields

$$\rho_{\text{eff}}(t) = \rho_{\text{on}} q_t(g = \text{on}). \quad (19)$$

The effective transcription rate learned by p is the expected transcription rate of q at time t , this time averaged over all m .

Finally if we require $\rho(m; t) = m \rho_{\text{eff}}$ to be time-independent we obtain by a similar procedure

$$\rho_{\text{eff}} = \rho_{\text{on}} \cdot \frac{1}{T} \int_0^T q_t(g = \text{on}) dt, \quad (20)$$

which is just the previous result averaged over t . This is the special case considered in the previous section, where p was restricted to be a Poisson process with constant intensity. As $T \rightarrow \infty$, the optimum (20) converges to the expected transcription rate of q at steady state.

2.4 Computing KL Divergences

To this point we have been occupied with minimising the KL divergence (1), or rather the cross-entropy (3), with respect to p . In order to assess the performance of the reduced model obtained this way we will compute the full KL divergence explicitly, which requires access to the entropy of the projected model \tilde{q} :

$$H(\tilde{q})_{[0, T]} = -\mathbb{E}_{\tilde{\mathbf{n}}_{[0, T]}} [\log q(\tilde{\mathbf{n}}_{[0, T]})]. \quad (21)$$

As for the cross-entropy the integral in Eq. (21) is generally intractable, and we will approximate the entropy by simulating N samples $\tilde{\mathbf{n}}_{[0, T]}^{(i)}$ from \tilde{q} and computing

$$\widehat{H}(\tilde{q}) = -\frac{1}{N} \sum_{i=1}^N \log \tilde{q}(\tilde{\mathbf{n}}_{[0, T]}^{(i)}). \quad (22)$$

Here we face another difficulty, for if \tilde{q} is not Markovian we cannot use Eqs. (4) or (5) to compute log-likelihood of a trajectory $\tilde{\mathbf{n}}_{[0, T]}$. Since by assumption q is a Markov process, the projection \tilde{q} is just a partially observed Markov process and we can use the so-called forward algorithm for Hidden Markov Models [35] to compute the log-likelihood of a trajectory.

The forward algorithm computes the joint probability $q(\tilde{\mathbf{n}}_{[0, t]}, \mathbf{z}_t)$ for the observed trajectory up to time t and the current state of the hidden species, \mathbf{z}_t . As shown in Appendix C, the forward algorithm in this case yields the following set of jump ODEs:

$$q(\tilde{\mathbf{n}}_0, \mathbf{z}_0) = q_0(\tilde{\mathbf{n}}_0, \mathbf{z}_0), \quad (23a)$$

$$\frac{d}{dt} q(\tilde{\mathbf{n}}_{[0, t]}, \mathbf{z}_t) = \sum_{\mathbf{z}'} q(\tilde{\mathbf{n}}_{[0, t]}, \mathbf{z}_t) \left(\sum_{i'} \sigma_{i'}(\tilde{\mathbf{n}}_{[0, t]}, \mathbf{z}_t) q(\tilde{\mathbf{n}}_{[0, t]}, \mathbf{z}') - \sum_i \sigma_i(\tilde{\mathbf{n}}_{[0, t]}, \mathbf{z}_t) q(\tilde{\mathbf{n}}_{[0, t]}, \mathbf{z}_t) \right), \quad (23b)$$

$$\lim_{t \searrow t_k} q(\tilde{\mathbf{n}}_{[0, t]}, \mathbf{z}_t) = \sigma_{j_k}(\tilde{\mathbf{n}}_k, \tilde{\mathbf{z}}_{t_k}) \cdot \lim_{t \nearrow t_k} q(\tilde{\mathbf{n}}_{[0, t]}, \mathbf{z}_t), \quad (23c)$$

with jumps at the jump times of the observed trajectory $\tilde{\mathbf{n}}_{[0,T]}$. The first sum in Eq. (23b) represents hidden reactions (those that do not affect $\tilde{\mathbf{n}}$), and the second sum is over all visible reactions in q , with propensity functions given by the σ_i . The marginal likelihood of the observed trajectory can then be computed as

$$\tilde{q}(\tilde{\mathbf{n}}_{[0,T]}) = \sum_{\mathbf{z}} q(\tilde{\mathbf{n}}_{[0,T]}, \mathbf{z}). \quad (24)$$

The above equations take the form of a modified CME where the observed variables are fixed. They can be solved using an adaptation of the Finite State Projection algorithm [5], and in cases where the unobserved state space is finite, exact solutions can be computed numerically. We summarise the procedure to estimate the KL divergence defined in Eq. (1) in Algorithm 2.

Algorithm 2 Computing the projected KL divergence between a full and a reduced model.

Inputs: number of simulations N , simulation length T , full model q , reduced model p

Output: KL divergence estimate $\overline{\text{KL}}(\tilde{q} \| p)$

```

 $\hat{L} := 0$ 
for all  $i = 1, \dots, N$  do
  sample  $n_{[0,T]}$  from  $q$ 
  project  $n_{[0,T]}$  to  $\tilde{\mathbf{n}}_{[0,T]}$ 
  compute  $p(\tilde{\mathbf{n}}_{[0,T]})$  using (5)
  compute  $\tilde{q}(\tilde{\mathbf{n}}_{[0,T]})$  using (23a)–(23c), (24).
   $\hat{L} += \tilde{q}(\tilde{\mathbf{n}}_{[0,T]}) - \tilde{p}(\tilde{\mathbf{n}}_{[0,T]})$ 
return  $\hat{L}/N$ 

```

Rearranging Eqs. (23a)–(23c) we can show that the marginal distribution itself solves the following jump ODE:

$$\log \tilde{q}(\tilde{\mathbf{n}}_0) = \log \tilde{q}_0(\tilde{\mathbf{n}}(0)), \quad (25a)$$

$$\frac{d}{dt} \log \tilde{q}(\tilde{\mathbf{n}}_{[0,t]}) = - \sum_i \mathbb{E}_{z_t} [\sigma_i(\tilde{\mathbf{n}}(t), \mathbf{z}_t) | \tilde{\mathbf{n}}_{[0,t]}, q_0; t], \quad (25b)$$

$$\lim_{t \searrow t_k} \log \tilde{q}(\tilde{\mathbf{n}}_{[0,t]}) = \lim_{t \nearrow t_k} \left(\log \tilde{q}(\tilde{\mathbf{n}}_{[0,t]}) + \log \mathbb{E}_{z_t} [\sigma_{j_k}(\tilde{\mathbf{n}}_k, \tilde{z}_t) | \tilde{\mathbf{n}}_{[0,t]}, q_0; t] \right). \quad (25c)$$

Eqs. (25a)–(25c) generalise Eqs. (4) and (5) to the presence of unobserved species \mathbf{z} .

The equations for the log-likelihood are linear and can be split into a sum of contributions for each reaction in \tilde{q} . Integrating over all reduced trajectories $\tilde{\mathbf{n}}_{[0,T]}$ shows that the entropy $H(\tilde{q})$ can thus be expressed as a sum of contributions from each reaction in \tilde{q} . Eq. (14) shows that the same decomposition holds for the cross-entropy $H(\tilde{q}; p)$, so entire KL divergence $\text{KL}(\tilde{q} \| p)$ can be decomposed into contributions coming from the individual reactions, which allows us to differentially analyse the accuracy of the reduced model p for each reaction separately. We note that all reactions together determine the distribution q on trajectories over which we integrate, so this decomposition is not strict in practice.

Example: Telegraph Model (cont. 2)

We can solve the filtering problem for the telegraph model exactly as the unobserved state space is 2-dimensional. Letting $F(i; t) = q(g(t) = i, m_{[0,t]})$, Eq. (23b) reduces to

$$\frac{d}{dt} F(\text{on}; t) = -(\rho_{\text{on}} + \sigma_{\text{off}}) F(\text{on}; t) + \sigma_{\text{on}} F(\text{off}; t), \quad (26)$$

$$\frac{d}{dt} F(\text{off}; t) = \sigma_{\text{off}} F(\text{on}; t) - \sigma_{\text{on}} F(\text{off}; t), \quad (27)$$

where $m_{[0,t]}$ denotes the observed mRNA trajectory. At each mRNA production event we update the probabilities according to Eq. (23c):

$$\lim_{t \searrow t_k} F(\text{on}; t) = \rho_{\text{on}} \cdot \lim_{t \nearrow t_k} F(\text{on}; t), \quad (28)$$

$$\lim_{t \searrow t_k} F(\text{off}; t) = 0. \quad (29)$$

The marginal likelihood of the observed mRNA trajectory is then given by

$$q(m_{[0,t]}) = F(\text{on}; t) + F(\text{off}; t). \quad (30)$$

As a byproduct the above equations yield the conditional distribution over the gene state at each time t given the trajectory prior to that time point. These are the filtered distributions and illustrated in Fig. 3c, but we only need the marginal likelihood for our purposes.

Using the above equations to compute the log-likelihood for a large number of trajectories sampled from the telegraph model q yields a numerical approximation of the entropy $H(\tilde{q})$, and together with Eq. (7) we obtain an estimate of the KL divergence $\text{KL}(\tilde{q} \| p)$. For time-homogeneous propensities the KL divergence will be asymptotically proportional to T , and we call the proportionality factor the (steady-state) KL divergence rate. In Fig. 3d we show how this KL divergence rate between the telegraph model and its Poisson approximation changes for various choices of σ_{on} and σ_{off} . We observe that the KL divergence rate tends to 0 as the switching rates increase; indeed the Poisson approximation can be obtained from the QEA for this example.

Our results on the telegraph model without degradation apply almost verbatim when degradation is included: adding the reaction $M \rightarrow \emptyset$ to the full and reduced models does not affect the KL divergences considered. *A priori* adding degradation changes the probability distribution over trajectories, which can now exhibit decreasing mRNA numbers, but the log-likelihood contributed by the degradation reaction is the same between the telegraph model and its reduction, and therefore the difference cancels in Eqs. (25a)–(25c) and Eq. (14). Since the log-likelihood contributed by the gene switching and transcription reactions does not depend on mRNA numbers, which do change in the presence of degradation, the total KL divergence is unaffected. In the presence of feedback, degradation would indirectly affect the KL divergence via its effect on mRNA numbers; we refer the interested reader to the recent paper [32] for an analysis of this type of information flow in stochastic biochemical reaction networks.

2.5 Relationship with Known Approaches

2.5.1 The Quasi-Steady State Approximation

In deterministic chemical kinetics the Quasi-Steady State Approximation (QSSA) is a model reduction technique that can be applied when a system can be partitioned into so-called slow species \mathbf{n}^s and fast species \mathbf{n}^f such that the fast species \mathbf{n}^f evolve on a much faster time scale than the slow species.

On the timescale on which the slow species evolve, the fast species can therefore be assumed to reach their steady state almost instantaneously. Thus one assumes $\frac{d}{dt}\mathbf{n}^f = 0$, which allows for simplification of the remaining equations for the slow species. The QSSA has famously been applied to Michaelis-Menten enzyme kinetics with the enzyme-substrate complex ES as the fast species, resulting in the classical Michaelis-Menten propensity for product formation (see 3.2 for more information).

The QSSA was extended to the stochastic case in [8]. Here one assumes that conditioned on \mathbf{n}^s , the fast species reach steady state nearly instantaneously compared to the timescale of interest. In our formulation the reduced state space consists of the slow species, $\tilde{\mathbf{n}} = \mathbf{n}^s$, and the unobserved species are the fast species, $\mathbf{z} = \mathbf{n}^f$. The QSSA then yields the following reduced CME for the approximation p (see (11) in [8]):

$$\frac{d}{dt}p_t(\mathbf{n}^s) = \sum_{i=1}^t \left[\tilde{\rho}_i(\mathbf{n}^s - \mathbf{S}_i^s; t) p_t(\mathbf{n}^s - \mathbf{S}_i^s) - \tilde{\rho}_i(\mathbf{n}^s; t) p_t(\mathbf{n}^s) \right], \quad (31)$$

where the reduced propensities $\tilde{\rho}_i$ are defined as

$$\tilde{\rho}_i(\mathbf{n}^s; t) = \mathbb{E}_{\mathbf{n}^f} \left[\rho_i(\mathbf{n}^s, \mathbf{n}^f) \mid \mathbf{n}^s; t \right]. \quad (32)$$

This agrees precisely with Eq. (15), which illustrates how the QSSA can be seen as minimising the KL divergence (50). We can compare (32) with the true propensities of the projection \tilde{q} , which depend on the entire history of the observed trajectory and the initial distribution q_0 :

$$\rho_i^{\text{exact}}(\mathbf{n}; t \mid q_0, \mathbf{n}_{[0,t]}^s) = \mathbb{E}_{\mathbf{n}^f} \left[\rho_i(\mathbf{n}^s, \mathbf{n}^f) \mid \mathbf{n}_{[0,t]}^s, q_0; t \right]. \quad (33)$$

Some intuition for the relationship between timescale separation and our approach can be gained from equations Eqs. (25a)–(25c), which describe the probability of a reduced trajectory $\tilde{\mathbf{n}}_{[0,T]}$ under \tilde{q} . Reaction propensities under q in general depend on the unobserved species, whose distribution is correlated with the history of the current trajectory. If we assume that the timescale of the unobserved species \mathbf{n}^f is very fast, then this correlation will decay very quickly and the time-dependence of the conditional distributions in Eqs. (25a)–(25c) will be negligible, so we can replace these equations by

$$\log \tilde{q}(\tilde{\mathbf{n}}_0) = \log \tilde{q}_0(\tilde{\mathbf{n}}(0)), \quad (34a)$$

$$\frac{d}{dt} \log \tilde{q}(\tilde{\mathbf{n}}_{[0,t]}) = - \sum_i \mathbb{E}_{\mathbf{z}_t} [\sigma_i(\tilde{\mathbf{n}}(t), \mathbf{z}_t) \mid \tilde{\mathbf{n}}_t], \quad (34b)$$

$$\lim_{t \searrow t_k} \log \tilde{q}(\tilde{\mathbf{n}}_{[0,t]}) = \lim_{t \nearrow t_k} \left(\log \tilde{q}(\tilde{\mathbf{n}}_{[0,t]}) + \log \mathbb{E}_{\mathbf{z}_t} [\sigma_{j_k}(\tilde{\mathbf{n}}_k, \mathbf{z}_t) \mid \tilde{\mathbf{n}}_t] \right). \quad (34c)$$

which is another way of showing that \tilde{q} has propensities given by (32).

2.5.2 The Quasiequilibrium Approximation

While the QSSA relies on a reaction network being divisible into slow and fast species that evolve on two different timescales, in practice it is more frequently the case that some *reactions* in a system will be fast and that some will be slow. The Quasiequilibrium Approximation (QEA) developed in [16, 36, 37] modifies the QSSA to model this scenario by reformulating the system using extents $\mathbf{a} = (a_1, \dots, a_r)$, where the extent a_i is defined the number of times reaction i has occurred.

The extents themselves form a Markov chain, and the state of a system can be obtained from its extents as $\mathbf{n}(t) = \mathbf{n}_0 + \mathbf{S}\mathbf{a}(t)$, where \mathbf{S} is the stoichiometry matrix. The CME of this system is

$$\frac{d}{dt}q_t(\mathbf{a}) = \sum_i [\rho_i(\mathbf{S}\mathbf{a} - \mathbf{S}_i) q_t(\mathbf{a} - \mathbf{e}_i) + \rho_i(\mathbf{S}\mathbf{a}) q_t(\mathbf{a})]. \quad (35)$$

Here the sum is over reactions and \mathbf{e}_i is the vector with a 1 in the i -th position and 0 elsewhere. The QEA assumes that the extent vector can be divided into slow and fast components \mathbf{a}^s and \mathbf{a}^f , respectively, corresponding to slow and fast reactions, and we define our reduced system to consist only of the slow reactions. From here we can proceed analogously to the QSSA discussed above and obtain that the reduced propensities are given by

$$\tilde{\rho}_i(\mathbf{a}^s; t) = \mathbb{E}_{\mathbf{a}^f}[\rho_i(\mathbf{S}\mathbf{a}) | \mathbf{a}^s; t], \quad (36)$$

if the fast reactions can be assumed to equilibrate instantaneously.

A subtlety of this argument is the fact that the fast extents \mathbf{a}^f can only increase in time and therefore will not admit a steady-state distribution in general. The conditional means, however, will converge in many cases, e.g. if \mathbf{a}^f consists of both directions of a reversible reaction. If this is the case the QEA yields a well-defined reduction, which moreover agrees with (15) and therefore minimises the KL divergence to the full model on the space of extents.

2.5.3 The Linear Mapping Approximation

The Linear Mapping Approximation [24] replaces a bimolecular reaction of the form $G + X \xrightarrow{\sigma} G^*$, where G is a binary species representing a gene state, by a reaction $G \xrightarrow{\bar{\sigma}} G^*$ with effective propensity $\bar{\sigma}$. Assuming mass action kinetics, the propensity function of the bimolecular reaction is $\rho(\mathbf{n}) = gx\sigma$, where g is the state of G and x the abundance of the species X , and the propensity of the linearised version is $\tilde{\rho}(\mathbf{n}) = \bar{\sigma}g$ for an appropriate choice of $\bar{\sigma}$. Taking the derivative with respect to $\bar{\sigma}$ of the cross-entropy rate (14) at time t yields

$$\frac{\partial H(\tilde{q}; p)_t}{\partial \bar{\sigma}} = \sum_{\mathbf{n}} q_t(\mathbf{n}) \left(g - \frac{\sigma gx}{\bar{\sigma}} \right), \quad (37)$$

which vanishes if and only if

$$\bar{\sigma} = \sigma \cdot \mathbb{E}[x | g = 1; t]. \quad (38)$$

This provides a mathematical justification for the mean-field assumption underlying [24]. The LMA approximates this conditional expectation by imposing a self-consistency condition on the linearised system, thereby deriving an effective approximation to (38).

As an aside we can compute the KL divergence rate between q and the optimal reduction p given by (38) analytically to obtain

$$\text{KL}(q \| p)_t = \sigma \cdot P(g = 1) \cdot (\mathbb{E}[x \log x | g = 1] - \mathbb{E}[x | g = 1] \mathbb{E}[\log x | g = 1]). \quad (39)$$

Here the dependence of the expectations on the time t are suppressed. This expression for the discrepancy incurred by the LMA, which resembles the definition of the variance, quantifies the intuition that the LMA should be accurate if fluctuations of X in the unbound gene state are small.

The above derivation is not entirely rigorous as the linearisation has a different net stoichiometry than the original reaction and KL divergence between q and p is infinite. To remedy this we can either neglect fluctuations in X due to binding in the original model or consider the KL divergence on a reaction-by-reaction basis as discussed in 2.4, since the two reactions correspond despite their different net stoichiometries.




Figure 4: Reduction of an autoregulatory feedback loop using the Linear Mapping Approximation. **(a)** Schematic of the full reaction network. The number $k \geq 0$ of proteins produced at each reaction is geometrically distributed with mean b . **(b)** Reduced form of the same reaction system, where the protein-gene binding reaction is replaced by a mean-field approximation. **(c)** For the chosen parameter values the system exhibits critical behaviour around $\sigma_b = 2.5$, transitioning from a mostly unbound to a mostly bound state. Near the transition protein fluctuations increase and the mean-field assumption in the LMA breaks down. A comparison of steady-state moments and distributions for the full model, the optimal reduction and the LMA shows that the effective reaction rate computed by the latter is generally close to optimal. **(d)** KL divergence rate at steady state between the full model and the reduced version, computed analytically using Eq. (39) and using Monte Carlo simulations. The peak around the transition matches the observed discrepancy between the full and the reduced model. The remaining model parameters are $\sigma_u = 400$, $\rho_u = 0.3$, $\rho_b = 105$, $d = 1$.

3 Numerical Experiments

3.1 Autoregulatory Feedback Loop

In this section we analyse a simple model of stochastic gene expression featuring positive autoregulation (see Fig. 4a). The system in question consists of a single gene found in two states, bound (G_b) and unbound (G_u), as well as the coded protein P which can bind to the gene to increase its own transcription rate. For simplicity we do not consider mRNA dynamics explicitly, instead modelling protein production as occurring in geometrically distributed bursts (see [38] for a derivation). The linearised version of the system, using the optimal propensity derived in Section 2.5.3, is shown in Fig. 4b. In Fig. 4c we compare the steady-state distributions of the full model with its linearisation, where the effective binding rate is computed numerically using our variational approach, and the LMA, where the binding rate is approximated using the self-consistent approach in [24].

This system exhibits critical behaviour for some parameter values (see Fig. 4c), and it was shown

in [39] that computing the moments of this system is difficult near points of criticality as most moment closure techniques as well as the LMA yield inaccurate results. Our results show that for all parameters the self-consistent equations of the LMA provides results close to the numerically computed optimum. Yet while both reductions reproduce the mean protein numbers of the full model (the LMA incurring a small bias near the critical point), both fail to capture the increased fluctuations near the critical point, significantly underestimating the variance in protein numbers. Comparing the protein distributions for the full and the optimally reduced model shows a large discrepancy near the critical point, compared to parameters far from it.

We compute the KL divergence rate at steady state between the full model and its linearisation using (39) and via Monte Carlo estimation (see Fig. 4d). The steady-state KL divergence rate exhibits a notable peak near the critical point, coinciding with the parameter regime where the linearisation fails to capture the full model. As we move away in either direction from the critical point the KL divergence decreases in accordance with the better approximation of the system by its linearised version. This shows how the KL divergence can be used to assess how well model reduction works for different parameter regimes.

3.2 Michaelis-Menten Kinetics

We next consider one of the most-studied reaction networks in biology, the Michaelis-Menten model of enzyme kinetics, consisting of an enzyme E and a substrate S that reversibly bind to form an enzyme-substrate complex ES which converts the substrate into the product P (see Fig. 5a). This system is often treated as an example application of both the QSSA and the QEA, which remain the standard model reduction techniques applied to this example.

The QSSA classically models the enzyme E and the enzyme-substrate complex ES as the fast species, and replaces the full system by a one-step reaction where the substrate is converted to the product with a Hill-type propensity function (see Fig. 5b). This approximation is known to be valid when

$$E_T \ll [S] + K_m, \quad (40)$$

where E_T is the total number of enzymes and K_m is the Michaelis-Menten constant $(k_{-1} + k_2)/k_1$ [40, 41]. Note that the reduced model is invariant under the scaling

$$E_T \mapsto \varepsilon \cdot E_T \quad k_i \mapsto \varepsilon^{-1} \cdot k_i, \quad i = -1, 1, 2, \quad (41)$$

where the parameter ε can be seen as controlling the accuracy of the reduction: the rescaled QSSA condition (40) reads

$$\varepsilon E_T \ll [S] + K_m, \quad (42)$$

which suggests that small values of ε should result in the full model more closely resembling the QSSA reduction.

The stochastic QEA for this system is analysed in [16] and also results in a one-step reaction where the propensity function is now piecewise linear (see Fig. 5b). As opposed to the QSSA, the QEA is accurate when substrate binding and unbinding is fast, i.e. k_1, k_{-1} are sufficiently large. Note that the QEA reduction is invariant under the scaling

$$k_1 \mapsto c \cdot k_1 \quad k_{-1} \mapsto c \cdot k_{-1}, \quad (43)$$




Figure 5: Comparison of the QEA and the QSSA for Michaelis-Menten kinetics. **(a)** Schematic of the Michaelis-Menten system. We assume mass action kinetics. **(b)** Reduction of the Michaelis-Menten system under the QEA and the QSSA yields a one-step process with effective (non-mass action) propensities that depend on the method used. **(c)** Total KL divergence between the full model and the QEA and QSSA for different values of the timescale parameter c in (43) and the enzyme abundance parameter ε in (41). The QEA generally becomes more accurate as c increases, and the QSSA becomes more accurate as ε decreases. Note the small copy number effects that become apparent in the QSSA for low values of ε and c . **(d)** Comparison of the effective reaction propensities computed according to Eq. (15) with those predicted by the QEA and the QSSA. Here substrates include enzyme-bound substrates. The top figure shows the effective propensities for $\varepsilon = 1$: as $c \rightarrow \infty$ the effective propensities converge to the function predicted by the QEA. The bottom figure shows the effective reaction propensities for $c = 1000$, and as ε decreases the propensities approach the function predicted by the QSSA. The remaining parameters were fixed to $k_1 = k_{-1} = 0.001$, $k_2 = 0.1$, $E_T = 10$ and $S_T = 100$.

where c controls the timescale at which quasiequilibrium is reached, and therefore the accuracy of the QEA. We stress that (41) and (43) are two independent scalings and can be performed simultaneously, as will be the case in Fig. 5c.

In general it may be difficult to predict which of the two approaches is more accurate unless one is clearly in one of the limiting regimes. Based on the KL divergence between the full model and either of the two reactions we can investigate this question for a range of parameters. The reduced model consists of two species, S and P , and since $[S] + [P]$ is conserved we can describe it in terms of either of the two. The correct projection for this example identifies the species S and ES in the full model, ie. we define $[S]_{\text{red}} = [S] + [ES]$ in order to eliminate the binding and unbinding reactions from the projection. This lumping of two rapidly equilibrating species is standard when using the QEA (see e.g. [23]), but it applies equally well to the QSSA in this case.

In Fig. 5b we compute the total KL divergence over $[0, \infty]$ between the full model and both reductions for a fixed number of substrate molecules. The total KL divergence is finite since all trajectories enter the same absorbing state defined by $[S] = 0$ in a finite amount of time. As expected for the QEA its accuracy increases with c , but whereas the QSSA tends to become less accurate for large E_T , in the low enzyme regime we observe a similar decrease in accuracy that is not explained by the deterministic theory.

The observed decrease is a small copy number effect caused by the fact that the waiting time distribution between two productions of P is not exponential: for very small E_T the unbinding of an enzyme immediately after such an event implies that a significant fraction of such productions occur as a two-step process (where the free enzyme binds another substrate and then converts it) [42]. For large enough c the binding step is very fast, and assuming that $k_2 \gg k_{-1}$ so that unbinding is unlikely to occur again, the conversion of substrates to products can be viewed as an effective one-step process. The effect of nonexponential waiting time distributions has previously been analysed in [43]; we see that the KL divergence on trajectories can be sensitive to subtle dynamical effects such as waiting time distribution that are not visible when considering e.g. the moments of a system, which are well predicted by the QSSA for small values of c and ε .

Overall we can see that for the chosen parameter values the QSSA generally performs better than the QEA in the regime of small ε , corresponding to low enzyme numbers, and the QEA is most accurate for c , keeping in mind the scalings in (41) and (43). Neither approximation is satisfactory if the number of total enzymes is larger than the amount of substrates, but the binding and unbinding rates are small. In this scenario other approximations, such as the total QSSA [44], which we do not investigate here, will generally be more accurate.

Figures 5c and 5d compare the effective propensities for the full system, as a function of the unconverted substrate abundance, with the predictions made by the QEA resp. the QSSA. In the case of the QEA we see that the effective propensities slowly converge to their limit as the timescale parameter c increases. In contrast the QSSA provides a good approximation to the effective propensity as long as the number of substrate molecules is larger than the number of enzymes. While the effective propensities are the optimal choice for the reduced model, the actual quality of the approximation is affected both by the size of the fluctuations of the actual propensities around their mean as well as the degree to which the waiting times of the original system follow an exponential distribution.

3.3 Genetic Oscillator

Our final example is the gene expression system shown in Fig. 6a. This system consists of a gene, mRNA and protein, as well as a Michaelis-Menten type protein degradation mechanism. Up to two protein molecules can bind the gene, and in the twice bound state the gene pauses transcription. This model has been considered in [11] as an example system where naive reduction of the CME using




Figure 6: Comparison of different reductions for the oscillatory gene network. **(a)** Schematic of the reaction system. We assume mass action kinetics for all reactions in the full model. **(b)** Example trajectories for the full model and its reductions. Note the oscillatory behaviour in the original model that is not preserved in reductions IV and V. **(c)** Power spectra of mRNA and protein concentrations (defined as copy numbers normalised by the system size Ω). Models I, II and I+II closely reproduce the oscillatory behaviour of the full model while IV and V do not show sustained oscillations. **(d)** KL divergence rates between the full model and all reductions. **(e)** Steady-state means and standard deviations for mRNA and protein numbers. While all reductions closely approximate the mean, Models IV and V do not match the variance of the full model. The parameters for the full model, taken from [11], are $\Omega = 1000$, $E_T = 10$, $\rho_m = 50$, $\rho_p = 0.0045$, $d_m = 0.01$, $k_1 = 0.1$, $k_{-1} = 10$, $k_2 = 10$, $\sigma_{b,P} = 0.001$, $\sigma_{u,P} = 100$, $\sigma_{b,PP} = 1000$, $\sigma_{u,PP} = 1$.

Hill-type effective propensities is unable to accurately capture the noise in the original. The true system can exhibit oscillations in mRNA numbers that are caused by the two-step negative feedback and are not present in the reduced version. In this section we want to analyse what reductions can be performed on the full model while still keeping oscillatory behaviour.

We consider five different reductions for this model which are listed in Table 2. Each reduction removes or combines several species such as gene states, thereby reducing the dimensionality of the system. The rightmost column describes the correct projection for each model, derived using an analogous argument as in the Michaelis-Menten example. We fit the unknown (effective) parameters for each model numerically by minimising the KL divergence from the full model. Typical mRNA trajectories for all models can be seen in Fig. 6b. The oscillations in the full model are clearly visible, and they are inherited by reductions I, II and I+II. In contrast, simplifying the first binding step as in IV and V results in a non-oscillatory system (see Fig. 6c).

In Fig. 6d we compare the KL divergence rates between the full model and reductions I-V (note that the different reductions are defined on different state spaces). While models I, II and I+II have comparatively low KL divergences from the original there is a sharp increase with models IV and V. From this alone one could expect that the latter perform significantly worse, which is indeed the case as shown in Fig. 6b. Fig. 6e presents a comparison of the mean and standard deviations for mRNA and protein abundances; while all models approximate the means very closely, the predicted standard deviations for model IV and V are very far from their true values consistent with the lack oscillatory behaviour. This resembles the results in Section 3.1 which also show excellent agreement of the full and reduced model on the mean level, independent of the total approximation quality.

This comparison of various reductions shows how variational model reduction can be employed on a more sophisticated scale, where multiple reductions are possible. Given a list of possible simplifications we can automatically find optimal parameters for each reduction and compute the KL divergence from the true model. Reductions which do not significantly affect the output will generally lead to very low KL divergences compared to those which do, as in this example where a large gap between KL divergence rates separates reductions I, II and I+II from IV and V. The latter do not retain much of the moments or oscillatory dynamics of the full model compared to the other reductions.

4 Discussion

In this work we presented an information-theoretic approach to model reduction for the CME based on minimising the KL divergence from the full model and the proposed reduction. Based on this variational principle we derive Eq. (15), showing that the optimal (effective) propensity for a reduction of the Chemical Master Equation equals the mean propensity of the original system conditioned on the state of the reduced system. As a consequence we show that some of the most common approaches to model reduction for the CME - the QSSA [8–10], the QEA [15, 16] and the LMA [24] - generally try to minimise the KL divergence considered in this paper, providing a general justification for the mean-field arguments proposed in the literature on these methods and connecting them with information theory and likelihood-based approaches. We provide a numerical algorithm for automated fitting of a reduced model based on minimising the KL divergence via stochastic gradient descent, and a numerical algorithm for estimating this KL divergence in order to assess model fit. While the KL divergence between Markov chains has been considered before in e.g. [18–20, 29–32, 45], to our knowledge it has not been studied in connection with standard model reduction methods as done in this paper. Our numerical results show how the KL divergence can be computed in practice and used in the context of model reduction.

Model	Full reactions	Reduced reactions	Projection map
I	$E + P \rightleftharpoons EP \longrightarrow E$	$P \xrightarrow{d_{\text{eff}}} \emptyset$	$[P]_{\text{red}} := [P] + [EP]$
II	$G + P \rightleftharpoons G_P$ $G_P + P \rightleftharpoons G_{PP}$ $G_P \longrightarrow G_P + P$	$G + P \xrightleftharpoons[\sigma_u / (K_u + [P])]{\sigma_{b,P}} G_{P/PP}$ $G_{P/PP} \xrightarrow{\rho_{\text{eff}}} G_{P/PP} + P$	$[P]_{\text{red}} := [P] + [G_{PP}]$ $[G_{P/PP}]_{\text{red}} := [G_P] + [G_{PP}]$
I + II	(I & II above)	(I & II above)	$[P]_{\text{red}} := [P] + [EP] + [G_{PP}]$ $[G_{P/PP}]_{\text{red}} := [G_P] + [G_{PP}]$
IV	$G + P \rightleftharpoons G_P$ $G_P + P \rightleftharpoons G_{PP}$ $G \longrightarrow G + P$ $G_P \longrightarrow G_P + P$	$G_{-}/P + P \xrightleftharpoons[\sigma_{u,PP}]{\sigma_b[P]/(K_b + [P])} G_{PP}$ $G_{-}/P \xrightarrow{\rho_m} G_{-}/P + P$	$[P]_{\text{red}} := [P] + [G_P] + [G_{PP}]$ $[G_{-}/P]_{\text{red}} := [G] + [G_P]$
V	$G + P \rightleftharpoons G_P$ $G_P + P \rightleftharpoons G_{PP}$ $G \longrightarrow G + P$ $G_P \longrightarrow G_P + P$	$\emptyset \xrightarrow{\frac{\alpha + \beta[P]}{1 + \gamma[P] + \delta[P]^2}} P$	$[P]_{\text{red}} := [P] + [G_P] + 2[G_{PP}]$

Table 2: Five possible reductions for the oscillator model. Model I+II is the combination of the reductions in Models I and II. All propensities are shown without their mass action terms. Species that are lumped or removed in the reduced model are shown in gray, and species that do not exist in the full model in blue. Parameters that are introduced in the reduced model are shown in red and found by minimisation of the KL divergence between the models; all remaining parameters are taken over from the full model. The functional forms of the propensity functions for Models II, IV and V were derived using the QEA.

Using three biologically relevant examples we illustrated how these model reduction techniques can be analysed from this variational perspective. For the autoregulatory feedback loop we showed that the KL divergence provides a useful metric of approximation quality and can detect parameter regimes where the mean-field approximation behind the LMA fails. Using the Michaelis-Menten system we demonstrated how our approach can be used to decide between possible reductions of a model, particularly in the non-asymptotic regime where neither the QSSA nor the QEA are strictly valid. Finally we used the genetic oscillator in [11] as an example to show how different reductions of a given model can be fit automatically, separately and in combination, and assesses which steps in a putative model reduction procedure would impact approximation quality more than others.

One important aspect of model reduction we have not addressed is that of choosing an appropriate architecture for a reduced model. As we investigated in the case of the genetic oscillator, while we can always optimise the parameters given an architecture, the quality of the approximation can greatly depend on which reductions are performed. An algorithm that performs automated model reduction for the CME, computing the maximally reduced version of a given model that still approximates the original within a given threshold, is still missing from the literature. In combination with other machine learning approaches to model reduction such as [46], however, we hope that our approach may be one of several steps towards such a procedure and.

While the KL divergence is a well-studied quantity in statistics, it lacks some desirable properties in the context of model reduction. There is in general no simple relationship connecting the KL divergence between two multivariate distributions with those of their marginals, and as a consequence we cannot compute e.g. the KL divergence between single-time marginals of the full and reduced models from our knowledge of the KL divergence on the trajectory level. Similarly, rigorously establishing a relationship between the KL divergence and dynamical information such as power spectra is difficult. Despite this, our empirical results suggest that the KL divergence on the space of trajectories is a useful quantity that relates to both the discrepancy between single-time marginals and power spectra and can be a useful metric of model fit despite a lack of sharp results. We are optimistic that further investigations in this area will demarcate more clearly which properties of a model are well captured by the KL divergence and which are not.

Code availability

Code implementing the methods introduced in this paper can be found at <https://github.com/kaandocal/modred>.

Acknowledgments

This work was supported by the EPSRC Centre for Doctoral Training in Data Science (EPSRC grant EP/L016427/1) and the University of Edinburgh for K. Ö. and a Leverhulme Trust grant (Grant No. RPG-2020-327) for R. G.

References

- [1] M. Kærn, T. C. Elston, W. J. Blake, and J. J. Collins, “Stochasticity in gene expression: from theories to phenotypes,” *Nat. Rev. Genet.*, **6**(6): 451–464, 2005.
- [2] R. Satija and A. K. Shalek, “Heterogeneity in immune responses: from populations to single cells,” *Trends Immunol.*, **35**(5): 219–229, 2014.

- [3] T. Lipniacki, B. Hat, J. R. Faeder, and W. S. Hlavacek, “Stochastic effects and bistability in T cell receptor signaling,” *J. Theor. Biol.*, **254**(1): 110–122, 2008.
- [4] D. Schnoerr, G. Sanguinetti, and R. Grima, “Approximation and inference methods for stochastic biochemical kinetics - a tutorial review,” *J. Phys. A*, **50**(9): 093001, 2017.
- [5] B. Munsky and M. Khammash, “The Finite State Projection algorithm for the solution of the Chemical Master Equation,” *J. Chem. Phys.*, **124**(4): 044104, 2006.
- [6] N. Ali Eshtewy and L. Scholz, “Model reduction for kinetic models of biological systems,” *Symmetry*, **12**(5): 863, 2020.
- [7] T. J. Snowden, P. H. van der Graaf, and M. J. Tindall, “Methods of model reduction for large-scale biological systems: a survey of current methods and trends,” *Bull. Math. Biol.*, **79**(7): 1449–1486, 2017.
- [8] C. V. Rao and A. P. Arkin, “Stochastic chemical kinetics and the quasi-steady-state assumption: application to the Gillespie algorithm,” *J. Chem. Phys.*, **118**(11): 4999–5010, 2003.
- [9] J. K. Kim, K. Josić, and M. R. Bennett, “The relationship between stochastic and deterministic quasi-steady state approximations,” *BMC Syst. Biol.*, **9**(1): 87, 2015.
- [10] —, “The validity of quasi-steady-state approximations in discrete stochastic simulations,” *Biophys. J.*, **107**(3): 783–793, 2014.
- [11] P. Thomas, A. V. Straube, and R. Grima, “The slow-scale Linear Noise approximation: An accurate, reduced stochastic description of biochemical networks under timescale separation conditions,” *BMC Syst. Biol.*, **6**(1): 39, 2012.
- [12] —, “Communication: Limitations of the stochastic quasi-steady-state approximation in open biochemical reaction networks,” *J. Chem. Phys.*, **135**(18): 181103, 2011.
- [13] H.-W. Kang, W. R. KhudaBukhsh, H. Koepl, and G. A. Rempała, “Quasi-steady-state approximations derived from the stochastic model of enzyme kinetics,” *Bull. Math. Biol.*, **81**(5): 1303–1336, 2019.
- [14] J. Eilertsen, K. Srivastava, and S. Schnell, “Stochastic enzyme kinetics and the quasi-steady-state reductions: Application of the slow-scale Linear Noise Approximation à la Fenichel,” *J. Math. Biol.*, **85**(1): 3, 2022.
- [15] E. L. Haseltine and J. B. Rawlings, “Approximate simulation of coupled fast and slow reactions for stochastic chemical kinetics,” *J. Chem. Phys.*, **117**(15): 6959–6969, 2002.
- [16] J. Goutsias, “Quasiequilibrium approximation of fast reaction kinetics in stochastic biochemical systems,” *The Journal of Chemical Physics*, **122**(18): 184102, 2005.
- [17] L. Cardelli, I. Pérez Verona, M. Tribastone, *et al.*, “Exact maximal reduction of stochastic reaction networks by species lumping,” *Bioinform.*, **37**: 2021.
- [18] R. A. Amjad, C. Blöchl, and B. C. Geiger, “A generalized framework for Kullback-Leibler Markov aggregation,” *IEEE Trans. Automat. Contr.*, **65**(7): 3068–3075, 2020.
- [19] K. Deng, P. G. Mehta, and S. P. Meyn, “Optimal Kullback-Leibler aggregation via spectral theory of Markov chains,” *IEEE Trans. Automat. Contr.*, **56**(12): 2793–2808, 2011.
- [20] B. C. Geiger, T. Petrov, G. Kubin, and H. Koepl, “Optimal Kullback-Leibler aggregation via information bottleneck,” *IEEE Trans. Automat. Contr.*, **60**(4): 1010–1022, 2015.
- [21] S. Bo and A. Celani, “Multiple-scale stochastic processes: decimation, averaging and beyond,” *Phys. Rep.*, **670**: 1–59, 2017.

- [22] J. Holehouse, A. Sukys, and R. Grima, “Stochastic time-dependent enzyme kinetics: closed-form solution and transient bimodality,” *J. Chem. Phys.*, **153**(16): 164113, 2020.
- [23] C. Jia and R. Grima, “Dynamical phase diagram of an auto-regulating gene in fast switching conditions,” *J. Chem. Phys.*, **152**(17): 174110, 2020.
- [24] Z. Cao and R. Grima, “Linear Mapping Approximation of gene regulatory networks with stochastic dynamics,” *Nat. Commun.*, **9**(1): 3305, 2018.
- [25] T. Jahnke and W. Huisinga, “Solving the Chemical Master Equation for monomolecular reaction systems analytically,” *J. Math. Biol.*, **54**(1): 1–26, 2007.
- [26] C. Gadgil, C. H. Lee, and H. G. Othmer, “A stochastic analysis of first-order reaction networks,” *Bull. Mathem. Biol.*, **67**(5): 901–946, 2005.
- [27] T. Zhou and J. Zhang, “Analytical results for a multistate gene model,” *SIAM J. Appl. Math.*, **72**(3): 789–818, 2012.
- [28] Y. Li, D.-Q. Jiang, and C. Jia, “Steady-state joint distribution for first-order stochastic reaction kinetics,” *Phys. Rev. E*, **104**(2): 024408, 2021.
- [29] M. Opper and G. Sanguinetti, “Variational inference for Markov jump processes,” in *Advances in Neural Information Processing Systems*, vol. 20, 2007.
- [30] C. Wildner and H. Koepl, “Moment-based variational inference for Markov jump processes,” in *Proceedings of the 36th International Conference on Machine Learning*, 2019, 6766–6775.
- [31] L. Bronstein and H. Koepl, “Marginal process framework: A model reduction tool for Markov jump processes,” *Phys. Rev. E*, **97**(6): 062147, 2018.
- [32] A.-L. Moor and C. Zechner, *Dynamic information transfer in stochastic biochemical networks*, 2022. arXiv: [2208.04162 \[q-bio.MN\]](https://arxiv.org/abs/2208.04162).
- [33] W. R. KhudaBukhsh, A. Auddy, Y. Disser, and H. Koepl, “Approximate lumpability for Markovian agent-based models using local symmetries,” *J. Appl. Probab.*, **56**(3): 647–671, 2019.
- [34] J. Peccoud and B. Ycart, “Markovian modeling of gene-product synthesis,” *Theor. Popul. Biol.*, **48**(2): 222–234, 1995.
- [35] L. Rabiner, “A tutorial on Hidden Markov Models and selected applications in speech recognition,” *Proc. IEEE*, **77**(2): 257–286, 1989.
- [36] Y. Cao, D. T. Gillespie, and L. R. Petzold, “The slow-scale Stochastic Simulation Algorithm,” *J. Chem. Phys.*, **122**(1): 014116, 2005.
- [37] E. L. Haseltine and J. B. Rawlings, “On the origins of approximations for stochastic chemical kinetics,” *J. Chem. Phys.*, **123**(16): 164115, 2005.
- [38] C. Jia and R. Grima, “Small protein number effects in stochastic models of autoregulated bursty gene expression,” *J. Chem. Phys.*, **152**(8): 084115, 2020.
- [39] K. Öcal, R. Grima, and G. Sanguinetti, “Parameter estimation for biochemical reaction networks using Wasserstein distances,” *J. Phys. A*, **53**(3): 034002, 2019.
- [40] L. A. Segel, “On the validity of the steady state assumption of enzyme kinetics,” *Bull. Math. Biol.*, **50**(6): 579–593, 1988.
- [41] D. Gillespie, K. Sanft, and L. Petzold, “Legitimacy of the stochastic Michaelis–Menten approximation,” *IET Syst. Biol.*, **5**(1): 58–69, 2011.

- [42] R. Grima and A. Leier, “Exact product formation rates for stochastic enzyme kinetics,” *J. Phys. Chem. B*, **121**(1): 13–23, 2017.
- [43] D. T. Gillespie, Y. Cao, K. R. Sanft, and L. R. Petzold, “The subtle business of model reduction for stochastic chemical kinetics,” *J. Chem. Phys.*, **130**(6): 064103, 2009.
- [44] S. MacNamara, A. M. Bersani, K. Burrage, and R. B. Sidje, “Stochastic chemical kinetics and the total quasi-steady-state assumption: application to the Stochastic Simulation Algorithm and Chemical Master Equation,” *J. Chem. Phys.*, **129**(9): 095105, 2008.
- [45] Z. Rached, F. Alajaji, and L. Campbell, “The Kullback-Leibler divergence rate between Markov sources,” *IEEE Trans. Automat. Contr.*, **50**(5): 917–921, 2004.
- [46] G. Caravagna, L. Bortolussi, and G. Sanguinetti, “Matching models across abstraction levels with gaussian processes,” in *Computational Methods in Systems Biology*, E. Bartocci, P. Lio, and N. Paoletti, Eds., ser. Lecture Notes in Computer Science, Cham: Springer International Publishing, 2016, 49–66.
- [47] D. T. Gillespie, “A general method for numerically simulating the stochastic time evolution of coupled chemical reactions,” *J. Comput. Phys.*, **22**(4): 403–434, 1976.

A Likelihoods of Fully Observed Trajectories

Let p be a continuous-time Markov chain defined on the space \mathcal{X} with initial distribution p_0 . If p has time-independent transition rates the Stochastic Simulation Algorithm [47], reproduced in Algorithm 3, returns exact samples from p . From the description we deduce that the log probability of sampling a given trajectory $\mathbf{n}_{[0,T]}$ equals

$$\log p(\mathbf{n}_{[0,T]}) = \log p_0(\mathbf{n}(0)) - \int_0^T p_{\leftarrow \mathbf{n}(s)} ds + \sum_{i=0}^{k-1} \log(p_{\mathbf{n}_{i+1} \leftarrow \mathbf{n}_i}) \quad (44)$$

where $\mathbf{n}_0, \mathbf{n}_1, \dots$ is the sequence of states visited. If the rates of p are time-dependent one can verify that the log probability is given by the following generalisation of (44):

$$\log p(\mathbf{n}_{[0,T]}) = \log p_0(\mathbf{n}(0)) - \int_0^T p_{\leftarrow \mathbf{n}(s)}(s) ds + \sum_{i=0}^{k-1} \log(p_{\mathbf{n}_{i+1} \leftarrow \mathbf{n}_i}(t_{i+1})), \quad (45)$$

where t_1, t_2, \dots are the jump times.

Algorithm 3 The Stochastic Simulation Algorithm to simulate samples from a Markov chain p .

Input: simulation length T , Markov chain p , initial distribution p_0

Output: $\mathbf{n}_{[0,T]}$ - sampled trajectory

sample $\mathbf{n}_0 \sim p_0$

$t \leftarrow 0, i \leftarrow 0$

while $t < T$ **do**

sample $\tau_{i+1} \sim \text{Exp}(1/p_{\leftarrow \mathbf{n}_i})$

sample $\mathbf{n}_{i+1} = \mathbf{m}$ **with probability** $p_{\mathbf{m} \leftarrow \mathbf{n}_i} / p_{\leftarrow \mathbf{n}_i}$

$t += \tau_{i+1}$

return states $(\mathbf{n}_0, \mathbf{n}_1, \dots)$, jump times (t_1, t_2, \dots)

B Kullback-Leibler Divergence between Markov Chains

In [29] the authors derive an expression for the Kullback-Leibler divergence between two continuous-time Markov chains q and p defined on a common state space \mathcal{X} . For any integer $N > 1$ let $q^{(N)}$ resp. $p^{(N)}$ be the time discretisations of q and p with time step $\delta t = T/N$. These define probability distributions on \mathcal{X}^{N+1} with KL divergence

$$\text{KL}(q^{(N)} \parallel p^{(N)}) = \sum_{\mathbf{n}} q_0(\mathbf{n}) \log \frac{q_0(\mathbf{n})}{p_0(\mathbf{n})} + \sum_{i=1}^N \sum_{\mathbf{n}, \mathbf{m}} q_{i-1}^{(N)}(\mathbf{n}) q_i^{(N)}(\mathbf{m} | \mathbf{n}) \log \frac{q_i^{(N)}(\mathbf{m} | \mathbf{n})}{p_i^{(N)}(\mathbf{m} | \mathbf{n})} \quad (46)$$

The transition rates are related to the discrete-time transition probabilities as follows:

$$q_i^{(N)}(\mathbf{n} | \mathbf{n}) = 1 - (\delta t) q_{\leftarrow \mathbf{n}}(iT/N) + o(\delta t) \quad (47)$$

$$q_i^{(N)}(\mathbf{m} | \mathbf{n}) = (\delta t) q_{\mathbf{m} \leftarrow \mathbf{n}}(iT/N) + o(\delta t) \quad (\mathbf{m} \neq \mathbf{n}) \quad (48)$$

Using these identities it can be verified that the KL divergence (46) converges to the following as $N \rightarrow \infty$:

$$\text{KL}(q \| p) = \text{KL}(q_0 \| p_0) - \int_0^T dt \sum_{\mathbf{n}} q_t(\mathbf{n}) (q_{\leftarrow \mathbf{n}}(t) - p_{\leftarrow \mathbf{n}}(t)) \quad (49)$$

$$+ \int_0^T dt \sum_{m \neq \mathbf{n}} q_t(\mathbf{n}) q_{m \leftarrow \mathbf{n}}(t) (\log q_{m \leftarrow \mathbf{n}}(t) - \log p_{m \leftarrow \mathbf{n}}(t)) \quad (50)$$

In general the KL divergence can be written as the difference of the the cross-entropy $H(q; p)$ and the entropy $H(q; p)$, and we define the cross-entropy of two continuous-time Markov chains as

$$H(q; p) = H(q_0; p_0) + \int_0^T dt \sum_{\mathbf{n}} q_t(\mathbf{n}) p_{\leftarrow \mathbf{n}}(t) - \int_0^T dt \sum_{m \neq \mathbf{n}} q_t(\mathbf{n}) q_{m \leftarrow \mathbf{n}}(t) \log p_{m \leftarrow \mathbf{n}}(t). \quad (51)$$

The above derivation is valid only if $q_{m \leftarrow \mathbf{n}} \neq 0$ implies $p_{m \leftarrow \mathbf{n}} \neq 0$; otherwise the cross-entropy and KL divergence are both infinite.

C Likelihoods of Reduced Trajectories

Let q be a continuous-time Markov chain defined on the discrete state space \mathcal{X} with initial distribution q_0 , and let \tilde{q} be its projection onto $\tilde{\mathcal{X}}$. Computing the log probability of a trajectory $\tilde{\mathbf{n}}_{[0,T]}$ under \tilde{q} requires integrating over all possible full trajectories $\mathbf{n}_{[0,T]} = (\tilde{\mathbf{n}}_{[0,T]}, \mathbf{z}_{[0,T]})$ that are compatible with $\tilde{\mathbf{n}}_{[0,T]}$, that is,

$$\tilde{q}(\tilde{\mathbf{n}}_{[0,T]}) = \int_{\mathbf{z}_{[0,T]}} q(\tilde{\mathbf{n}}_{[0,T]}, \mathbf{z}_{[0,T]}) d\mathbf{z}_{[0,T]} \quad (52)$$

In order to compute the integral on the right-hand side we consider the time-discretisations of q and \tilde{q} . The marginal likelihood in this case can be obtained using the well-known forward algorithm for HMMs, which sequentially computes the joint probabilities $q^{(N)}(\tilde{\mathbf{n}}_{0:i}, \mathbf{z}_i)$:

$$q^{(N)}(\tilde{\mathbf{n}}_0, \mathbf{z}_0) = q_0(\tilde{\mathbf{n}}_0, \mathbf{z}_0) \quad (53)$$

$$q^{(N)}(\tilde{\mathbf{n}}_{0:i+1}, \mathbf{z}_{i+1}) = \sum_{\mathbf{z}} q^{(N)}(\tilde{\mathbf{n}}_{0:i}, \mathbf{z}_i = \mathbf{z}) q^{(N)}(\tilde{\mathbf{n}}_{i+1}, \mathbf{z}_{i+1} | \tilde{\mathbf{n}}_i, \mathbf{z}_i = \mathbf{z}) \quad (54)$$

The marginal probability of the entire trajectory $\tilde{\mathbf{n}}_{0:N}$ can then be computed by marginalising after the last step:

$$q^{(N)}(\tilde{\mathbf{n}}_{0:N}) = \sum_{\mathbf{z}_N} q^{(N)}(\tilde{\mathbf{n}}_{0:N}, \mathbf{z}_N). \quad (55)$$

Passing to the limit $N \rightarrow \infty$ and rewriting the visible transition rates in terms of reaction propensities we arrive at (23a)–(23c).