

Marginal process framework: A model reduction tool for Markov jump processesLeo Bronstein and Heinz Koepl^{*}*Department of Electrical Engineering and Information Technology, Technische Universität Darmstadt, 64283 Darmstadt, Germany*

(Received 10 February 2018; published 26 June 2018)

Markov jump process models have many applications across science. Often these models are defined on a state space of product form and only one of the components of the process is of direct interest. In this paper we extend the marginal process framework, which provides a marginal description of the component of interest, to the case of fully coupled processes. We use entropic matching to obtain a finite-dimensional approximation of the filtering equation, which governs the transition rates of the marginal process. The resulting equations can be seen as a combination of two projection operations applied to the full master equation so that we obtain a principled model reduction framework. We demonstrate the resulting reduced description on the totally asymmetric exclusion process. An important class of Markov jump processes are stochastic reaction networks, which have applications in chemical and biomolecular kinetics, ecological models, and models of social networks. We obtain a particularly simple instantiation of the marginal process framework for mass-action systems by using product Poisson distributions for the approximate solution of the filtering equation. We investigate the resulting approximate marginal process analytically and numerically.

DOI: [10.1103/PhysRevE.97.062147](https://doi.org/10.1103/PhysRevE.97.062147)**I. INTRODUCTION**

Markov jump processes (MJPs) have many applications across science and engineering. The master equation (ME) [1], which governs the time evolution of the probability distribution of the process, is generally too complicated to solve analytically and frequently infeasible to solve numerically. A popular alternative is the stochastic simulation (Gillespie) algorithm [2], which produces samples of trajectories of the process. For larger systems, however, this approach can be computationally very expensive, especially if the system exhibits multiscale behavior.

In many cases, the MJP is defined on a product-form state space $\mathbb{X} \times \hat{\mathbb{X}}$ and only one component is of direct interest while the other component can be considered a nuisance variable. The question arises whether it is possible to derive a reduced description for the stochastic process corresponding to the component of interest only. Mathematically, the remaining component then has to be marginalized out of the full stochastic process. This approach has been used for reaction networks in the case when the variables of interest do not influence the nuisance variables [3] and also in other contexts [4], mostly to speed up stochastic simulations.

A large number of other model reduction methods have been proposed in the literature. Many of these are based on timescale separation or abundance separation, a nonexhaustive list being [5–11]. Other approaches based on marginalization have recently been published [12–14] and focus on stochastic differential equation models.

In this article we extend the marginalization approach of [3,4] to a general MJP with full coupling between the variable of interest and the nuisance variable. Marginalization

requires the solution of the (in general, infinite-dimensional) filtering equation, which describes the evolution of the conditional probability of the nuisance variable given the trajectory of the marginal process. We use entropic matching [15,16] to obtain a finite-dimensional approximation. The filtering equation and entropic matching can be interpreted as the result of projection operations consecutively applied to the full ME of the joint process. In this way, we obtain a principled model reduction method. Our focus is on the marginal process framework as a theoretical tool for model reduction rather than as a method for more efficient stochastic simulation.

A particularly important class of MJPs with applications in chemical kinetics, biological population models, and models of social interactions are reaction networks. For reaction networks with mass-action kinetics, a particularly simple reduced description is obtained when a product Poisson ansatz distribution is used for the approximate solution of the filtering equation. We call the resulting reduced model the Poisson marginal process and investigate it in detail. Analogously, for exclusion processes, a product Bernoulli ansatz distribution leads to what is the simplest possible reduced model within our framework. We investigate this reduced process for the example of the totally asymmetric exclusion process (TASEP) on the line with open boundaries.

This paper is organized as follows. After describing the problem setting in Sec. II, we provide an outline of the proposed method in Sec. III, using a simple model of constitutive gene expression as a running example. The general form of the marginal process framework is derived in Sec. IV. The finite-dimensional approximations necessary for a tractable description of the marginal process are discussed in Sec. V, where we also apply our method to the TASEP as a first example. The Poisson marginal process for mass-action reaction networks is discussed in Sec. VI.

^{*}heinz.koepl@bcs.tu-darmstadt.de

II. SETTING

We consider an MJP $(X, \hat{X}) = (X(t), \hat{X}(t))_{t \geq 0}$ on a product-form state space $\mathbb{X} \times \hat{\mathbb{X}}$, where \mathbb{X} and $\hat{\mathbb{X}}$ are countable sets. The marginal probability distribution $p_t(x, \hat{x}) = \Pr(X(t) = x, \hat{X}(t) = \hat{x})$ of such a process [with initial distribution $p_0(x, \hat{x})$ at time 0] is governed by the ME

$$\begin{aligned} \frac{d}{dt} p_t(x, \hat{x}) &= [\mathcal{L} p_t](x, \hat{x}) \\ &= \sum_{y, \hat{y}} \{L(x, \hat{x} | y, \hat{y}) p_t(y, \hat{y}) - L(y, \hat{y} | x, \hat{x}) p_t(x, \hat{x})\}, \end{aligned}$$


where $L(x, \hat{x} | y, \hat{y})$ is the rate of transitioning from state (y, \hat{y}) to state (x, \hat{x}) . We set $L(x, \hat{x} | x, \hat{x}) = 0$ for all $(x, \hat{x}) \in \mathbb{X} \times \hat{\mathbb{X}}$ for convenience. We will also require the backward evolution operator \mathcal{L}^\dagger , which acts on functions $\psi : \mathbb{X} \times \hat{\mathbb{X}} \rightarrow \mathbb{R}$ and is given by

$$[\mathcal{L}^\dagger \psi](x, \hat{x}) = \sum_{y, \hat{y}} L(y, \hat{y} | x, \hat{x}) \{\psi(y, \hat{y}) - \psi(x, \hat{x})\}. \quad (1)$$

It is the adjoint of \mathcal{L} with respect to the pairing $(p, \psi) := \sum_{x, \hat{x}} p(x, \hat{x}) \psi(x, \hat{x})$. Recall that \mathcal{L}^\dagger governs the moment equations of the stochastic process. Thus, for a function $\psi(x, \hat{x})$, the expectation $\langle \psi \rangle_t$ with respect to the distribution $p_t(x, \hat{x})$ evolves according to

$$\begin{aligned} \frac{d}{dt} \langle \psi \rangle_t &= \sum_{x, \hat{x}} \psi(x, \hat{x}) \frac{d}{dt} p_t(x, \hat{x}) = \sum_{x, \hat{x}} \psi(x, \hat{x}) [\mathcal{L} p_t](x, \hat{x}) \\ &= \sum_{x, \hat{x}} p_t(x, \hat{x}) [\mathcal{L}^\dagger \psi](x, \hat{x}) = \langle \mathcal{L}^\dagger \psi \rangle_t. \end{aligned} \quad (2)$$

We are particularly interested in reaction networks, consisting of $N + \hat{N}$ species and R reactions that are specified as



for $j = 1, \dots, R$. Here we have divided the set of all species into the species X_1, \dots, X_N of interest, the subnet, and the remaining species $\hat{X}_1, \dots, \hat{X}_{\hat{N}}$, the environment. We are interested in the case when the copy numbers of some of the species might be small, so a fully stochastic description in terms of an MJP is necessary. The process $X = (X_1, \dots, X_N)$ then describes the state of the subnet species, while $\hat{X} = (\hat{X}_1, \dots, \hat{X}_{\hat{N}})$ describes the state of the environment species. The state space is given by $\mathbb{X} = \mathbb{N}_0^N$ and $\hat{\mathbb{X}} = \mathbb{N}_0^{\hat{N}}$. For a state (x, \hat{x}) , for each $j = 1, \dots, R$ there exists a transition to the state $(x + v_j, \hat{x} + \hat{v}_j)$ with rate $h_j(x, \hat{x})$ and change vector (v_j, \hat{v}_j) with $v_j = (r_{1j} - s_{1j}, \dots, r_{Nj} - s_{Nj})$ and $\hat{v}_j = (\hat{r}_{1j} - \hat{s}_{1j}, \dots, \hat{r}_{\hat{N}j} - \hat{s}_{\hat{N}j})$. The operators \mathcal{L} and \mathcal{L}^\dagger take the form

$$\begin{aligned} [\mathcal{L} p](x, \hat{x}) &= \sum_{j=1}^R \{h_j(x - v_j, \hat{x} - \hat{v}_j) p(x - v_j, \hat{x} - \hat{v}_j) \\ &\quad - h_j(x, \hat{x}) p(x, \hat{x})\}, \\ [\mathcal{L}^\dagger \psi](x, \hat{x}) &= \sum_{j=1}^R h_j(x, \hat{x}) \{\psi(x + v_j, \hat{x} + \hat{v}_j) - \psi(x, \hat{x})\}. \end{aligned}$$

In all examples treated in this article, we will employ mass-action kinetics, which are given by $h_j(x, \hat{x}) = \Omega c_j f_j(x) \hat{f}_j(\hat{x})$ with

$$f_j(x) = \prod_{n=1}^N \frac{(x_n)_{s_{nj}}}{\Omega^{s_{nj}}}, \quad \hat{f}_j(\hat{x}) = \prod_{n=1}^{\hat{N}} \frac{(\hat{x}_n)_{\hat{s}_{nj}}}{\Omega^{\hat{s}_{nj}}}, \quad (4)$$

where $(x)_s = x(x-1) \cdots (x-s+1)$ denotes the falling factorial, c_j is a reaction rate constant, and we have introduced the system size Ω in terms of which we will analyze the behavior of the approximate marginal process.


We will also require a description of the process (X, \hat{X}) in terms of the R reactions. We associate with each reaction channel j a counting process $Y_j(t)$ which counts the number of firings of reaction j over the time interval $[0, t]$. The process (Y_1, \dots, Y_R) can again be seen as a reaction network of the form (3) with values in \mathbb{N}_0^R , which can only change by increments of size 1 in any one of its components at a single time. The state of the original process (X, \hat{X}) is recovered from the state of these counting processes via

$$X(t) = \sum_{j=1}^R Y_j(t) v_j, \quad \hat{X}(t) = \sum_{j=1}^R Y_j(t) \hat{v}_j.$$

III. OUTLINE OF THE METHOD

As explained in the Introduction, our goal is to derive a marginal process description for the process of interest X . While the joint process (X, \hat{X}) is Markovian, this is no longer the case for the marginal process X . The effect of the nuisance variables \hat{X} is implicitly contained in the memory of the process X . We now illustrate our proposed method on a simple reaction network from biology. We focus on the underlying ideas and postpone derivations to later sections.

A very simple model of constitutive gene expression is given by the reaction network



Assuming that we are interested primarily in the protein dynamics, we will consider the mRNA to be a nuisance species. Our goal is to obtain a marginal description of the protein dynamics. Thus, the mRNA plays the role of the nuisance variable \hat{X} and the protein the role of the variable of interest X . Assuming mass-action kinetics, the transition rates of the four reactions in the state (x, \hat{x}) are given in Table I.

The steps to obtain a tractable approximate description of the marginal process are as follows. (i) Determine how the transition rates of the marginal process at time t depend on the process history $x_{[0,t]}$. (ii) Find a description for these marginal transition rates in terms of an evolution equation driven by the process X . The resulting equations are generally

TABLE I. Transition rates for the reaction network (5) when the process is in state (x, \hat{x}) .

$\emptyset \rightarrow \text{mRNA}$	$\text{mRNA} \rightarrow \emptyset$	$\text{mRNA} \rightarrow \text{mRNA} + \text{P}$	$\text{P} \rightarrow \emptyset$
Ωc_1	$c_2 \hat{x}$	$c_3 \hat{x}$	$c_4 x$

infinite dimensional but provide an exact description of the marginal process. (iii) Choose an approximation to obtain finite-dimensional equations. We will now carry out these steps for our simple example network.

(i) *Description of the marginal process.* Since the first two reactions in Table I do not change the state of X , the marginal process consists of two reactions, corresponding to the last two reactions in Table I. Generally, since the marginal process is no longer Markovian, the transition rates for these reactions will depend on the entire history $x_{[0,t]}$ instead of just on the current state $x(t)$. However, the transition rate of the fourth reaction in Table I does not depend on the mRNA abundance. Consequently, its marginal transition rate remains unchanged and is given by $c_4 x(t)$. In particular, it depends only on the current state $x(t)$ of the marginal process. In contrast to this, the rate for the third reaction does depend on the mRNA abundance. As will be derived in Sec. IV A, the corresponding marginal transition rate is given by $c_3 \mathbf{E}[\hat{X}(t) | x_{[0,t]}]$. This is an intuitive result, expressing the fact that in absence of information about the mRNA abundance, the marginal transition rate is given by the expectation of the transition rate conditional on all available information, i.e., conditional on the entire process history $x_{[0,t]}$.

(ii) *Filtering equation.* We are now tasked with computing the expectation $\mathbf{E}[\hat{X}(t) | x_{[0,t]}]$. A convenient way to do this is to derive an evolution equation, driven by the marginal process $X(t)$, for the so-called filtering distribution $\pi_t(\hat{x}) := \Pr(\hat{X}(t) = \hat{x} | x_{[0,t]})$ with respect to which this expectation is computed. The resulting equation is called the filtering equation. As the process X is a jump process, the trajectory $x_{[0,t]}$ is piecewise constant. The filtering equation for $\pi_t(\hat{x})$ will thus consist of two parts: continuous evolution (described by a differential

equation) as long as X remains constant and discontinuous jumps whenever X jumps. This is schematically illustrated in Fig. 1. As will be derived in Sec. IV C, the continuous evolution is given by

$$\begin{aligned} \frac{d}{dt} \pi_t(\hat{x}) = & \Omega c_1 [\pi_t(\hat{x} - 1) - \pi_t(\hat{x})] \\ & + c_2 [(\hat{x} + 1) \pi_t(\hat{x} + 1) - \hat{x} \pi_t(\hat{x})] \\ & - c_3 [\hat{x} - \langle \hat{x} \rangle_t] \pi_t(\hat{x}). \end{aligned} \quad (6)$$

Here the expectation $\langle \hat{x} \rangle_t$ is computed with respect to the distribution $\pi_t(\hat{x})$ itself. The first two terms on the right-hand side of (6) simply correspond to the ME for the mRNA alone, the dynamics of which do not depend on the protein abundance. The last term describes how the information that is contained in the trajectory $x_{[0,t]}$ impacts our state of knowledge about the mRNA abundance. Note that the right-hand side of this equation does not depend on the state of the marginal process X . This is because the reaction network has a feedforward structure. In general, the filtering equation will depend on the state of X .

As explained above, the filtering distribution will also jump instantaneously whenever the driving process X jumps (see Fig. 1). At a jump of X at time t , the corresponding jump $\pi_{t+} - \pi_{t-}$ of the filtering distribution will depend on which reaction caused the change in X . Since the protein decay reaction does not depend on the mRNA abundance, no information about mRNA abundance is obtained when this reaction fires. Therefore, $\pi_{t+} - \pi_{t-} = 0$ in this case, in analogy to the continuous part (6) of the filtering equation in which the protein abundance likewise plays no role. When the protein abundance increases via the third reaction in Table I, however, we instantaneously receive a finite amount of information about the mRNA state. To understand this, note, for example, that the third reaction can fire only if there is at least one mRNA molecule present. Thus, the filtering distribution immediately after the jump, $\pi_{t+}(\hat{x})$, certainly has to satisfy $\pi_{t+}(0) = 0$. As will be shown in Sec. IV C, the jump in the filtering distribution when the third reaction fires is given by

$$\pi_{t+}(\hat{x}) = \frac{\hat{x}}{\langle \hat{x} \rangle_{t-}} \pi_{t-}(\hat{x}). \quad (7)$$

In principle, (6) and (7) provide a full, exact description of the marginal process, allowing us to compute the marginal transition rates at any time t from the history $x_{[0,t]}$ of the marginal process X . For some simple processes, the corresponding equations can be solved in closed form, as will be demonstrated in Sec. IV D. In general, however, these equations constitute an infinite-dimensional system that does not provide a sufficiently simple description of the marginal process dynamics. We thus have to look for finite-dimensional approximations.

(iii) *Finite-dimensional approximation.* Since we are interested only in the expectation $\langle \hat{x} \rangle_t = \mathbf{E}[\hat{X}(t) | x_{[0,t]}]$ of the filtering distribution $\pi_t(\hat{x})$, it seems reasonable to consider the first-order moment equations for (6) and (7). However, the equations for the mean $\langle \hat{x} \rangle_t$ are not closed, because the second-order moment $\langle \hat{x}^2 \rangle_t$ enters: We obtain

$$\frac{d}{dt} \langle \hat{x} \rangle_t = \Omega c_1 - c_2 \langle \hat{x} \rangle_t - c_3 (\langle \hat{x}^2 \rangle_t - \langle \hat{x} \rangle_t^2) \quad (8)$$




FIG. 1. Schematic illustration of the concepts involved in the construction of the marginal process, based on the example network (5). Bottom: Bold curve and shaded area show mean and plus or minus one standard deviation of the filtering distribution. The marginal process trajectory (top) drives the evolution of the filtering distribution and the filtering distribution mean determines the transition rates of the marginal process X . Note that jumps of the filtering distribution occur only if the marginal process increases by a jump.

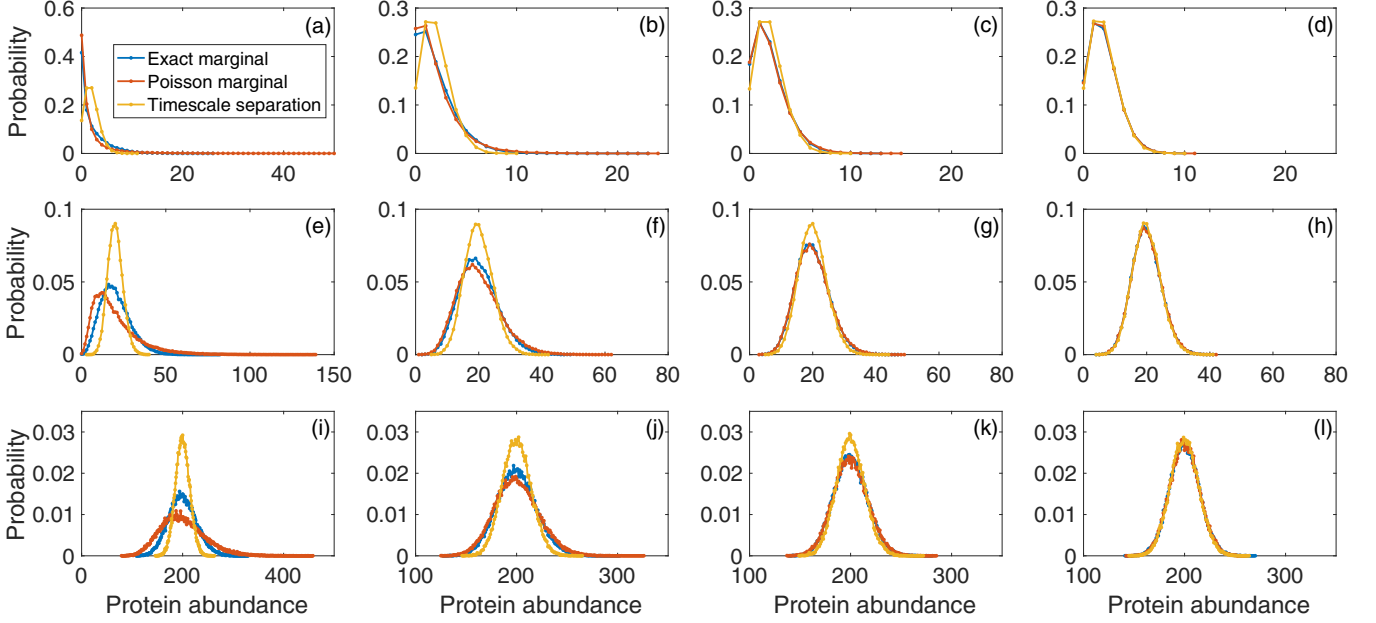


FIG. 2. Monte Carlo evaluation of the approximation quality of the Poisson marginal process and of timescale separation. Distributions of protein abundance at stationarity from 50 000 Monte Carlo runs for each case. The parameters are $c_1 = \gamma$, $c_2 = \gamma/2$, $c_3 = 1$, and $c_4 = 0.1$. Rows correspond to system sizes $\Omega = 0.1$ for (a–d), $\Omega = 1$ for (e–h) and $\Omega = 10$ for (i–l). Columns correspond to mRNA process speeds of (a), (e), and (i) $\gamma = 0.5$, (b), (f), and (j) $\gamma = 2$, (c), (g), and (k) $\gamma = 5$, and (d), (h), and (l) $\gamma = 25$. Note that the Poisson marginal process has a somewhat heavier right tail than the exact marginal process, especially at low system size and low value of γ . The algorithm used for stochastic simulation of the marginal process is explained in Appendix A.

from (6) and

$$\langle \hat{x} \rangle_{t+} = \frac{\langle \hat{x}^2 \rangle_{t-}}{\langle \hat{x} \rangle_{t-}} \quad (9)$$

from (7). To find a tractable description of the marginal process, we employ moment closure to obtain a finite-dimensional system of equations. As will be explained in more detail in Sec. VI, in this article we want to obtain the simplest possible description of the (approximate) marginal process and so choose a first-order closure, incorporating the mean of the filtering distribution only. A natural choice for such a closure ansatz is the Poisson distribution. Writing $\theta(t)$ for the mean of the Poisson ansatz distribution, we obtain

$$\frac{d}{dt}\theta(t) = \Omega c_1 - c_2\theta(t) - c_3\theta(t) \quad (10)$$

from (8) and

$$\theta(t+) = \theta(t-) + 1 \quad (11)$$

from (9). These equations complete our description of the approximate marginal process, which we denote by X' (and refer to as the Poisson marginal process) to distinguish it from the exact marginal process X . Using (10) and (11), we can compute the marginal transition rates at time t based on the full history $x_{[0,t]}$ of the approximate marginal process X' . We use the fact that knowing the history $x_{[0,t]}$ is equivalent to knowing the histories $(y_3)_{[0,t]}$ and $(y_4)_{[0,t]}$ of the two processes Y_3 and Y_4 (as defined in Sec. II) that count firings of the third and fourth reactions. Solving (10) and (11) in terms of the process histories, we obtain for the marginal rate of the third reaction

the expression

$$c_3 \left[e^{-(c_2+c_3)t}\theta(0) + \int_0^t e^{-(c_2+c_3)(t-\tau)} \{ \Omega c_1 d\tau + dy_3(\tau) \} \right].$$

The Stieltjes integral here reduces to a sum, because $(y_3)_{[0,t]}$ is piecewise constant.

In order to evaluate the quality of our chosen approximation, we can compute the mean and the variance of the approximate marginal process X' and of the exact marginal process X . Using results from Sec. VIA, we find that the means of the exact and approximate marginal processes coincide at all times, assuming the initial conditions are chosen appropriately. At stationarity, the means are given by $\langle x \rangle_\infty = \langle x' \rangle_\infty = \Omega c_1 c_3 / c_2 c_4$. The variances of the processes, however, differ. We compute the relative error of the variance approximation at stationarity and find

$$\frac{\langle x'^2 \rangle_\infty - \langle x^2 \rangle_\infty}{\langle x^2 \rangle_\infty - \langle x \rangle_\infty^2} = \frac{c_3^2}{2c_2(c_2 + c_3 + c_4)}.$$

One particular regime where the error vanishes is timescale separation, when $c_1, c_2 \rightarrow \infty$ with c_1/c_2 constant. It is thus natural to compare our approach with an approximation that directly invokes timescale separation. As mentioned in the Introduction, there exist a large number of approaches. For our simple network, however, there is one particularly natural option (e.g., [5]): We consider the process $\emptyset \xrightarrow{k_3} \text{protein} \xrightarrow{c_4} \emptyset$ with the rate constant $k_3 = \Omega c_3 c_1 / c_2$. One easily checks that at stationarity, the timescale separation ansatz reproduces the correct mean. We compare the approximation of the full distributions at stationarity numerically in Fig. 2. We see

that the Poisson marginal process systematically improves on timescale separation.

IV. MARGINAL PROCESS FRAMEWORK

In this section we introduce the marginal process framework in full generality and derive the necessary equations for the case of a general MJP defined on a product-form state space. We then specialize to the case of reaction networks, where it is useful to additionally introduce a slightly modified version of the marginal process.

A. Marginal process

As explained in Sec. III, the marginal process X is in general no longer Markovian, so the transition rates at time t will depend on the entire history $x_{[0,t]}$ of the process over the time interval $[0,t]$, instead of just on the current state $x(t)$. We now proceed to compute these marginal transition rates in a way analogous to [17].

For the marginal process, the probability for a transition into the state $y \in \mathbb{X}$ to happen in the time interval $[t, t + \Delta t]$, conditional on the process history $x_{[0,t]}$ [and assuming $y \neq x(t)$], is given by

$$\begin{aligned} & \Pr(X(t + \Delta t) = y \mid x_{[0,t]}) \\ &= \sum_{\hat{x}, \hat{y}} \Pr(X(t + \Delta t) = y, \hat{X}(t + \Delta t) = \hat{y} \mid X(t) = x(t), \\ & \quad \hat{X}(t) = \hat{x}) \Pr(X(t) = x(t), \hat{X}(t) = \hat{x} \mid x_{[0,t]}) \\ &= \sum_{\hat{x}, \hat{y}} L(y, \hat{y} \mid x(t), \hat{x}) \Pr(X(t) = x(t), \\ & \quad \hat{X}(t) = \hat{x} \mid x_{[0,t]}) \Delta t + o(\Delta t) \\ &= \mathbf{E}[K(y \mid x(t), \hat{X}(t)) \mid x_{[0,t]}] \Delta t + o(\Delta t), \end{aligned}$$

where $K(y \mid x, \hat{x}) = \sum_{\hat{y}} L(y, \hat{y} \mid x, \hat{x})$ is the total rate for jumps from the state (x, \hat{x}) leading to any state in $\{y\} \times \hat{\mathbb{X}}$. Thus, the marginal transition rate is given by

$$\mathbf{E}[K(y \mid x(t), \hat{X}(t)) \mid x_{[0,t]}], \quad (12)$$

i.e., by the expectation of the total transition rate conditional on the entire history of the marginal process up to time t .

The distribution $\Pr(\hat{X}(t) = \hat{x} \mid x_{[0,t]})$ with respect to which the expectation is computed is the filtering distribution for the stochastic process \hat{X} given the ‘‘observed’’ trajectory $x_{[0,t]}$ of the stochastic process X . The filtering distribution is the solution to the problem of estimating the state of the unobserved variable $\hat{X}(t)$ given the available information $x_{[0,t]}$ about the observed variable.

We see that, in order to obtain a useful description of the marginal process, we require a sufficiently simple description

of the filtering distribution, or at least of the marginal transition rates $\mathbf{E}[K(y \mid x(t), \hat{X}(t)) \mid x_{[0,t]}]$ that are computed with respect to the filtering distribution. One way to obtain such a description is to formulate an evolution equation for the filtering distribution driven by the marginal process X . For the case of two fully coupled Markov jump processes, we are not aware of the required results existing in the literature, so we provide an elementary derivation. For an overview of stochastic filtering in general, see [18].

B. Filtering equation

The filtering distribution $\pi_t(\hat{x}) := \Pr(\hat{X}(t) = \hat{x} \mid x_{[0,t]})$ is, in principle, defined over the state space $\hat{\mathbb{X}}$ of the nuisance variable. It is however convenient and natural to consider it as a distribution over the joint state space $\mathbb{X} \times \hat{\mathbb{X}}$ via

$$\begin{aligned} \pi_t(x, \hat{x}) &:= \Pr(X(t) = x, \hat{X}(t) = \hat{x} \mid x_{[0,t]}) \\ &= \delta_{x(t), x} \Pr(\hat{X}(t) = \hat{x} \mid x_{[0,t]}), \end{aligned}$$

where $\delta_{x,y}$ is the Kronecker delta. This simply expresses the fact that conditional on $x_{[0,t]}$, the state of $X(t)$ is known to be $x(t)$ with probability one. Depending on the situation, either of these two views will be more convenient, so in the following we will repeatedly switch between considering the filtering distribution to be defined either on $\hat{\mathbb{X}}$ or on $\mathbb{X} \times \hat{\mathbb{X}}$.

For the derivations below, the following two operators will be useful: a summation operator \mathcal{S} and an evaluation operator \mathcal{P}_y (which depends on a state $y \in \mathbb{X}$), both of which act on functions $\psi : \mathbb{X} \times \hat{\mathbb{X}} \rightarrow \mathbb{R}$. They are defined by

$$\begin{aligned} [\mathcal{S}\psi] &= \sum_{x, \hat{x}} \psi(x, \hat{x}), \\ [\mathcal{P}_y\psi](x, \hat{x}) &= \delta_{y,x} \psi(y, \hat{x}). \end{aligned} \quad (13)$$

We can now derive the filtering equation. The filtering distribution π_t will evolve according to a differential equation in between jumps of the process X and will jump whenever X jumps. The intuition here is that, over an infinitesimal time interval dt , if the observed process X does not jump, we receive only an infinitesimal amount of information so that the change in the filtering distribution should also be infinitesimal. When, however, X does jump, we receive a finite amount of information and, correspondingly, the filtering distribution has to jump too (see also Fig. 1).

Assuming that we have observed the process X over a time interval $[0, t + \Delta t]$, these observations can be partitioned into the observations $x_{[0,t]}$ up to time t and the observation $x(t + \Delta t)$. We assume Δt sufficiently small such that at most one jump occurred during the time interval $[t, t + \Delta t]$. Using Bayes’ theorem, we have

$$\begin{aligned} & \Pr(\hat{X}(t + \Delta t) = \hat{x} \mid x(t + \Delta t), x_{[0,t]}) \\ &= \sum_{y, \hat{y}} \Pr(X(t + \Delta t) = x(t + \Delta t), \hat{X}(t + \Delta t) = \hat{x} \mid X(t) = y, \hat{X}(t) = \hat{y}) \frac{\Pr(X(t) = y, \hat{X}(t) = \hat{y} \mid x_{[0,t]})}{\Pr(X(t + \Delta t) = x(t + \Delta t) \mid x_{[0,t]})} \\ &= \frac{[e^{\Delta t \mathcal{L}} \pi_t](x(t + \Delta t), \hat{x})}{\sum_{\hat{y}} [e^{\Delta t \mathcal{L}} \pi_t](x(t + \Delta t), \hat{y})} = \frac{\pi_t(x(t + \Delta t), \hat{x}) + \Delta t [\mathcal{L} \pi_t](x(t + \Delta t), \hat{x}) + o(\Delta t)}{\sum_{\hat{y}} \{\pi_t(x(t + \Delta t), \hat{y}) + \Delta t [\mathcal{L} \pi_t](x(t + \Delta t), \hat{y})\} + o(\Delta t)}. \end{aligned}$$

Multiplying this equation by $\delta_{x(t+\Delta t),x}$, using that

$$\pi_{t+\Delta t}(x, \hat{x}) = \delta_{x(t+\Delta t),x} \Pr(\hat{X}(t+\Delta t) = \hat{x} \mid x(t+\Delta t), x_{[0,t]}),$$

and noting the definition of \mathcal{P} and \mathcal{S} in (13), we find

$$\pi_{t+\Delta t} = \frac{\mathcal{P}_{x(t+\Delta t)}\pi_t + \Delta t[\mathcal{P}_{x(t+\Delta t)}\mathcal{L}\pi_t] + o(\Delta t)}{\delta_{x(t+\Delta t),x(t)} + \Delta t[\mathcal{S}\mathcal{P}_{x(t+\Delta t)}\mathcal{L}\pi_t] + o(\Delta t)}. \quad (14)$$

In the denominator, we also used that

$$\sum_{\hat{y}} \pi_t(x(t+\Delta t), \hat{y}) = \delta_{x(t+\Delta t),x(t)}.$$

We now have to distinguish the cases $x(t+\Delta t) = x(t)$ and $x(t+\Delta t) \neq x(t)$. When $x(t+\Delta t) = x(t)$, i.e., X remained constant over the time interval $[t, t+\Delta t]$, we have $\mathcal{P}_{x(t+\Delta t)}\pi_t = \pi_t$. Subtracting π_t from (14), dividing by Δt , and taking the limit $\Delta t \rightarrow 0$, we obtain

$$\frac{d}{dt}\pi_t(x, \hat{x}) = [\mathcal{P}_{x(t)}\mathcal{L}\pi_t](x, \hat{x}) - \pi_t(x, \hat{x})[\mathcal{S}\mathcal{P}_{x(t)}\mathcal{L}\pi_t]. \quad (15)$$

This is the differential equation that the filtering distribution satisfies in between jumps of the process X . It turns out that (15) can also be obtained as an orthogonal projection of the full (joint) ME computed with respect to the Fisher-Rao information metric. This point of view is described in Appendix B and will allow us to better understand the finite-dimensional approximation of the filtering equation introduced in Sec. VB below.

When $x(t+\Delta t) \neq x(t)$, i.e., when X jumps during the time interval $[t, t+\Delta t]$, we have $\mathcal{P}_{x(t+\Delta t)}\pi_t = 0$. Taking the limit $\Delta t \rightarrow 0$ in (14), we obtain an expression for the filtering distribution immediately after the jump, π_{t+} , in terms of the filtering distribution immediately before the jump, π_{t-} , given by

$$\pi_{t+}(x, \hat{x}) = \frac{[\mathcal{P}_{x(t+)}\mathcal{L}\pi_{t-}](x, \hat{x})}{[\mathcal{S}\mathcal{P}_{x(t+)}\mathcal{L}\pi_{t-}]}, \quad (16)$$

where $x(t+)$ is the value of X after the jump.

We now write down expressions (15) and (16) explicitly in terms of the transition rates. The explicit expressions are simpler if we regard the filtering distribution as being defined only over \mathbb{X} , i.e., $\pi_t = \pi_t(\hat{x})$. We define

$$R(x, \hat{x}) = \sum_{y \neq x} \sum_{\hat{y}} L(y, \hat{y} \mid x, \hat{x}),$$

the total rate of those transitions out of state (x, \hat{x}) that change the \mathbb{X} component. In between jumps of X , we then have

$$\begin{aligned} \frac{d}{dt}\pi_t(\hat{x}) &= \sum_{\hat{y}} \{L(x(t), \hat{x} \mid x(t), \hat{y})\pi_t(\hat{y}) \\ &\quad - L(x(t), \hat{y} \mid x(t), \hat{x})\pi_t(\hat{x})\} \\ &\quad - \{R(x(t), \hat{x}) - \langle R(x(t), \hat{x}) \rangle_t\}\pi_t(\hat{x}), \end{aligned} \quad (17)$$

where $\langle R(x(t), \hat{x}) \rangle_t = \sum_{\hat{x}} R(x(t), \hat{x})\pi_t(\hat{x})$ denotes the expectation computed using the filtering distribution π_t . The first term on the right-hand side of (17) is an ME for the nuisance component \hat{X} involving only those transitions which do not change the \mathbb{X} component of the state. Note that the corresponding transition rates can still depend on the current state of X . The last term in (17) accounts for the observations. Here the

observations contain information by virtue of the fact that X does not jump as long as (17) is in effect. From this equation we also see that the effect of feedback from the variable of interest to the nuisance variable is very simple: Because X is constant between its jumps, X is simply fixed to its current value in the transition rates in (17). When X does jump so that $x(t+) \neq x(t-)$, the corresponding jump in the filtering distribution is given by

$$\pi_{t+}(\hat{x}) = \frac{\sum_{\hat{y}} L(x(t+), \hat{x} \mid x(t-), \hat{y})\pi_t(\hat{y})}{\sum_{\hat{x}', \hat{y}} L(x(t+), \hat{x}' \mid x(t-), \hat{y})\pi_t(\hat{y})}. \quad (18)$$


The combination of (17) and (18) with the marginal transition rates (12) provides a full description of the marginal process X . For simple processes, these expressions can be evaluated and solved in closed form, as will be demonstrated for a simple reaction network in Sec. IV D. Before discussing the example, we specialize the discussion to reaction networks.

C. Reaction networks

As we can see from (17), the transitions of the MJP are naturally partitioned into two groups: those that change the state of the \mathbb{X} component and those that do not. We will denote by $J_{\mathbb{X}} \subseteq \{1, \dots, R\}$ the subset of indices of those reactions which can modify X and by $\bar{J}_{\mathbb{X}} = \{1, \dots, R\} \setminus J_{\mathbb{X}}$ the indices of all remaining reactions. This partitioning also results in a partitioning of the counting processes Y_1, \dots, Y_R (defined in Sec. II) into two processes $Y = (Y_j)_{j \in J_{\mathbb{X}}}$ and $\hat{Y} = (Y_j)_{j \in \bar{J}_{\mathbb{X}}}$ with the former containing the reactions in $J_{\mathbb{X}}$ and the latter the remaining reactions in $\bar{J}_{\mathbb{X}}$. Note that the state of the subnet can then be recovered from Y alone, while the state of the environment generally requires knowledge of both Y and \hat{Y} :

$$X(t) = \sum_{j \in J_{\mathbb{X}}} Y_j(t)v_j, \quad \hat{X}(t) = \sum_{j \in J_{\mathbb{X}}} Y_j(t)\hat{v}_j + \sum_{j \in \bar{J}_{\mathbb{X}}} Y_j(t)\hat{v}_j.$$

We now specialize the results obtained for the marginal process for general MJPs to the case of reaction networks. At this point, there arises an issue regarding the precise definition of the history of the marginal process on which we condition in the marginal transition rates (12). Generally, it can happen that a reaction network contains two different reactions, with different change vectors (v_i, \hat{v}_i) and (v_j, \hat{v}_j) , for which however the components corresponding to the subnet are identical, $v_i = v_j$. For example, this is the case for the simple gene expression model with negative feedback



Here G_0 and G_1 are the two possible states of a gene and X is the gene product which is produced when the gene is in state G_1 . The gene product can also reversibly bind to the gene and switch it to state G_0 , in which production of X is no longer possible. When the gene product X is considered to constitute the subnet, the reactions $G_1 \rightarrow G_1 + X$ and $G_0 \rightarrow G_1 + X$ both lead to an increase of X of size 1. Similarly, $X \rightarrow \emptyset$ and $G_1 + X \rightarrow G_0$ both lead to a decrease of X of size 1.

For such a reaction network, we obtain two different marginal processes depending on whether the history of the process is defined to be just the trajectory $x_{[0,t]}$ (as was done in Sec. IV A) or the trajectory $y_{[0,t]}$ of the counting processes




FIG. 3. Example of a sampled trajectory of the process (X, Θ) . The solid red curve shows the abundance of gene product X over time. The dashed blue curve shows the state of the filtering distribution mean Θ . Note that, as can be seen from (25), jumps in Θ occurring when $\Theta = 1$ always lead to the same value of $\Theta = c_2/(c_2 + c_3)$, indicated by the dotted line. Whether Θ increases or decreases after the jump depends on the value of X . The parameters are $c_1 = 1$, $c_2 = 0.25$, $c_3 = 0.5$, and $c_4 = 2$. The initial state was $X(0) = 0$ and $\Theta(0) = 1$.

Y of all reactions which change the subnet. In the former case we will speak of the marginal process X and in the latter case of the marginal process Y . Both marginal processes are meaningful and only minor changes in the derivations presented in Secs. IV A and IV B are necessary. We will present expressions for both cases because each version of the marginal process has advantages and disadvantages.

The marginal transition rate for the process Y for reaction $j \in J_{\mathbb{X}}$ is given by

$$\mathbf{E}[h_j(x(t), \hat{X}(t)) \mid y_{[0,t]}], \quad (20)$$

where $x(t) = \sum_{j \in J_{\mathbb{X}}} y_j(t) v_j$. This is different from (12), which for reaction networks reads

$$\sum_j \mathbf{E}[h_j(x(t), \hat{X}(t)) \mid x_{[0,t]}]$$

for a transition with change vector v and where the summation runs over all $j \in J_{\mathbb{X}}$ such that $v_j = v$.

The filtering equation, similarly, exists in two variants, depending on which form of the marginal process we consider. It turns out however that the continuous part (15) of the filtering equation is the same for both variants and explicitly reads

$$\begin{aligned} \frac{d}{dt} \pi_t(\hat{x}) &= \sum_{j \in \overline{J_{\mathbb{X}}}} \{h_j(x(t), \hat{x} - \hat{v}_j) \pi_t(\hat{x} - \hat{v}_j) - h_j(x(t), \hat{x}) \pi_t(\hat{x})\} \\ &\quad - \sum_{j \in J_{\mathbb{X}}} \{h_j(x(t), \hat{x}) - \langle h_j(x(t), \hat{x}) \rangle_t\} \pi_t(\hat{x}). \end{aligned} \quad (21)$$

For the marginal process X as defined in Sec. IV A, the jump in π_t when X jumps is given by

$$\pi_{t+}(\hat{x}) = \frac{\sum_j h_j(x(t-), \hat{x} - \hat{v}_j) \pi_{t-}(\hat{x} - \hat{v}_j)}{\sum_j \langle h_j(x(t-), \hat{x}) \rangle_{t-}}, \quad (22)$$

where the sums in numerator and denominator each run over all reaction indices $j \in J_{\mathbb{X}}$ such that $v_j = x(t+) - x(t-)$. If instead we consider the marginal process Y , a transition $j \in J_{\mathbb{X}}$ leads to a jump in the filtering distribution given by

$$\pi_{t+}(\hat{x}) = \frac{h_j(x(t-), \hat{x} - \hat{v}_j) \pi_{t-}(\hat{x} - \hat{v}_j)}{\langle h_j(x(t-), \hat{x}) \rangle_{t-}}. \quad (23)$$

The absence of summations in (23) will be useful in Sec. VI A. Here we proceed to discuss a simple example for which only the marginal process X is useful.

D. Example: A case with finite-dimensional filtering equations

We consider the simple gene expression model (19), with the gene product X chosen to constitute the subnet. For this model, every reaction changes the state of X , so the marginal process Y would be equal to the full process and thus of no interest. Consequently, we instead consider the (one-dimensional) marginal process X . This process has two reactions $\emptyset \rightarrow X$ and $X \rightarrow \emptyset$, with rates at time t given by

$$c_1 \theta(t) + c_4 [1 - \theta(t)], \quad [c_2 + c_3 \theta(t)] x(t),$$

respectively, where $\theta(t) = \langle g_1 \rangle_t$ is the filtering distribution mean of the gene state G_1 and we assumed that only a single copy of the gene is present. The filtering distribution $\pi_t(g_0, g_1)$, initially defined on $\{0, 1\} \times \{0, 1\}$, is fully determined by a single number due to the conservation relation $G_0 + G_1 = 1$. Similarly, for the expectation values with respect to $\pi_t(g_0, g_1)$ we have $\theta(t) = \langle g_1 \rangle_t = \pi_t(0, 1) = 1 - \langle g_0 \rangle_t$. We can now write down the (one-dimensional) filtering equation using (21) and (22). In between jumps of X , the result reads

$$\frac{d}{dt} \theta(t) = c(x(t)) [1 - \theta(t)] \theta(t), \quad (24)$$

where $c(x(t)) = c_4 - [c_3 x(t) + c_1]$. This is solved, for an initial value of $\theta(t_0)$ at time t_0 , by

$$\theta(t) = \frac{\theta(t_0) e^{c(x(t_0))(t-t_0)}}{1 + \theta(t_0) (e^{c(x(t_0))(t-t_0)} - 1)},$$

where we used that $x(t)$ is constant and equal to $x(t_0)$ in between jumps. When the reaction $\emptyset \rightarrow X$ fires, the filtering distribution mean θ jumps to 1. This is clear because both reactions of (19) that cause a change in X of size +1 lead to the gene being in state G_1 . More interesting is the case when the reaction $X \rightarrow \emptyset$ fires. Then the jump in the filtering distribution mean is given by

$$\theta(t+) = \frac{c_2 \theta(t-)}{c_2 + c_3 \theta(t-)}. \quad (25)$$

This completes the description of the marginal process. If we consider $\theta(t)$ as an auxiliary variable and use it to augment the process state, the resulting process (X, Θ) is a piecewise-deterministic Markov process with two reactions and deterministic evolution in between jumps given by (24). We show a sample from this augmented process in Fig. 3.

While simple systems such as the one discussed in this section can be treated without approximation, more complicated

systems will require an approximate solution of the filtering equations, as was already mentioned in Sec. III. We address this issue next.

V. FINITE-DIMENSIONAL APPROXIMATIONS OF THE FILTERING EQUATION

The filtering equation is in general infinite dimensional and, just as the ME, far too complicated for either analytical or numerical solution. Thus, regardless of whether one is interested in the marginal process for analytical investigation or for stochastic simulation, an approximate treatment of the filtering equation is necessary.

A. Moment equations and moment closure

A standard approach to obtain a finite-dimensional approximation is moment closure [19], which in the context of stochastic filtering is also known as assumed-density filtering [20]. While moment closures have often been considered *ad hoc* approximations, a simple variational derivation has been obtained recently [16].

We first derive the filtering moment equations in their general form, starting from (15) and (16). The moments of the filtering distribution will evolve according to a differential equation in between jumps of the marginal process and will jump whenever the marginal process jumps. Again considering π_t to be defined on $\mathbb{X} \times \hat{\mathbb{X}}$, we consider the moment equation for a function $\psi : \mathbb{X} \times \hat{\mathbb{X}} \rightarrow \mathbb{R}$, which can be obtained as demonstrated in (2) and is given by

$$\begin{aligned} \frac{d}{dt} \langle \psi \rangle_t &= \sum_{x, \hat{x}} \psi(x, \hat{x}) \{ [\mathcal{P}_{x(t)} \mathcal{L} \pi_t](x, \hat{x}) - \pi_t(x, \hat{x}) [\mathcal{S} \mathcal{P}_{x(t)} \mathcal{L} \pi_t] \} \\ &= \langle \mathcal{L}^\dagger \mathcal{P}_{x(t)} [\psi - \langle \psi \rangle_t] \rangle_t \end{aligned} \quad (26)$$

in between jumps of X , where we used that $\mathcal{P}_{x(t)}^\dagger = \mathcal{P}_{x(t)}$. When X jumps we have

$$\langle \psi \rangle_{t+} = \frac{\langle \mathcal{L}^\dagger \mathcal{P}_{x(t+)} \psi \rangle_{t-}}{\langle \mathcal{L}^\dagger \mathcal{P}_{x(t+)} 1 \rangle_{t-}}. \quad (27)$$

We skip explicit expressions in terms of rates for general MJPs and instead write down the simpler explicit expressions for reaction networks, which read (now for a function $\phi : \hat{\mathbb{X}} \rightarrow \mathbb{R}$)

$$\begin{aligned} \frac{d}{dt} \langle \phi \rangle_t &= \sum_{j \in J_{\hat{\mathbb{X}}}} \{ \langle h_j(x(t), \hat{x}) \phi(\hat{x} + \hat{v}_j) \rangle_t - \langle h_j(x(t), \hat{x}) \phi(\hat{x}) \rangle_t \} \\ &\quad - \sum_{j \in J_{\mathbb{X}}} \{ \langle h_j(x(t), \hat{x}) \phi(\hat{x}) \rangle_t - \langle h_j(x(t), \hat{x}) \phi(\hat{x}) \rangle_t \} \end{aligned} \quad (28)$$

in between jumps. Focusing on the marginal process Y , at a jump of Y via reaction $j \in J_{\hat{\mathbb{X}}}$ we have

$$\langle \phi \rangle_{t+} = \frac{\langle h_j(x(t-), \hat{x}) \phi(\hat{x} + \hat{v}_j) \rangle_{t-}}{\langle h_j(x(t-), \hat{x}) \rangle_{t-}}. \quad (29)$$

The moment equations are, as is generally the case, not closed: Choosing, for instance, $\phi(\hat{x}) = \hat{x}_n$ for a reaction network to obtain first-order moments, the resulting equations will depend on moments of order higher than one. In this way, an

infinite hierarchy of moment equations is obtained. We require a way to close the system of equations.

In this paper we consider a special case of moment closure, which has a dual interpretation based on minimization of relative entropy [15,16] and on projection using the Fisher-Rao information metric [20]. We next present a derivation analogous to [16], which allows for a unified treatment of the continuous and discrete parts of the filtering equation.

B. Entropic matching

A finite-dimensional approximation of a distribution $p(\hat{x})$ can be obtained by choosing a distribution from within a finite-dimensional parametric family $p_\theta(\hat{x})$ with parameters θ . There are strong arguments [21] for choosing this approximation so that it minimizes the relative entropy

$$D[p \parallel p_\theta] = \sum_{\hat{x}} p(\hat{x}) \ln \frac{p(\hat{x})}{p_\theta(\hat{x})}.$$

In the context of the filtering equation, we proceed as follows [16]. Choose a parametric family of probability distributions $p_\theta(\hat{x})$ depending on parameters θ ranging in some open subset of \mathbb{R}^K . Assume that, at time t , we have an approximation $p_{\theta(t)}(\hat{x})$ of the filtering distribution $\pi_t(\hat{x})$ available. As for the filtering distribution itself, we identify the approximation $p_\theta(\hat{x})$ on $\hat{\mathbb{X}}$ with $p_\theta(x, \hat{x}) = \delta_{x(t), x} p_\theta(\hat{x})$ on $\mathbb{X} \times \hat{\mathbb{X}}$.

We first consider the continuous part of the filtering equation. Then a short time step Δt later, $p_{\theta(t)}$ will have evolved to

$$\begin{aligned} p(x, \hat{x}) &= p_{\theta(t)}(x, \hat{x}) + \Delta t [\mathcal{P}_{x(t)} \mathcal{L} p_{\theta(t)}](x, \hat{x}) \\ &\quad - \Delta t p_{\theta(t)}(x, \hat{x}) [\mathcal{S} \mathcal{P}_{x(t)} \mathcal{L} p_{\theta(t)}]. \end{aligned}$$

We will obtain an approximation to $p(x, \hat{x})$ that lies in the parametric family p_θ by choosing parameters $\theta(t + \Delta t)$ to minimize the relative entropy $D[p \parallel p_{\theta(t+\Delta t)}]$. We then take the limit $\Delta t \rightarrow 0$ to obtain an ordinary differential equation (ODE) for the parameters θ . Write, for brevity, $\theta = \theta(t)$ and $\tilde{\theta} = \theta(t + \Delta t)$. Then we have, to first order in Δt ,

$$\begin{aligned} D[p \parallel p_{\tilde{\theta}}] &= \left\langle \ln \frac{p_\theta + \Delta t \{ \mathcal{P}_{x(t)} \mathcal{L} p_\theta - p_\theta [\mathcal{S} \mathcal{P}_{x(t)} \mathcal{L} p_\theta] \}}{p_{\tilde{\theta}}} \right\rangle_p \\ &= \left\langle \ln \frac{p_\theta}{p_{\tilde{\theta}}} \right\rangle_\theta \\ &\quad + \Delta t \left[\left\langle \left\{ \frac{\mathcal{P}_{x(t)} \mathcal{L} p_\theta}{p_\theta} - [\mathcal{S} \mathcal{P}_{x(t)} \mathcal{L} p_\theta] \right\} \ln \frac{p_\theta}{p_{\tilde{\theta}}} \right\rangle_\theta + \text{const} \right], \end{aligned}$$

where const denotes terms independent of $\tilde{\theta}$ and $\langle \cdot \rangle_\theta$ denotes an expectation taken with respect to the distribution p_θ . The first term is simply equal to $D[p_\theta \parallel p_{\tilde{\theta}}]$, which to second order in $\tilde{\theta} - \theta$ is given by

$$D[p_\theta \parallel p_{\tilde{\theta}}] = \frac{1}{2} (\tilde{\theta} - \theta)^\dagger G(\theta) (\tilde{\theta} - \theta),$$

where $G(\theta)$ is the Fisher information matrix of the parametric family p_θ at parameter value θ , the components of which are given by

$$G_{kl}(\theta) = \left\langle \frac{\partial \ln p_\theta}{\partial \theta_k} \frac{\partial \ln p_\theta}{\partial \theta_l} \right\rangle_\theta, \quad k, l = 1, \dots, K. \quad (30)$$

To minimize $D[p \parallel p_{\tilde{\theta}}]$, we take the derivative with respect to $\tilde{\theta}$ and obtain

$$\begin{aligned} 0 &= G(\theta)(\tilde{\theta} - \theta) \\ &\quad - \Delta t \left\langle \left\{ \frac{\mathcal{P}_{x(t)} \mathcal{L} p_{\theta}}{p_{\theta}} - [\mathcal{S} \mathcal{P}_{x(t)} \mathcal{L} p_{\theta}] \right\} \nabla_{\tilde{\theta}} \ln p_{\tilde{\theta}} \right\rangle_{\theta} \\ &= G(\theta)(\tilde{\theta} - \theta) - \Delta t \langle \mathcal{L}^{\dagger} \mathcal{P}_{x(t)}^{\dagger} \nabla_{\tilde{\theta}} \ln p_{\tilde{\theta}} \rangle_{\theta} \\ &\quad + \Delta t [\mathcal{S} \mathcal{P}_{x(t)} \mathcal{L} p_{\theta}] \langle \nabla_{\tilde{\theta}} \ln p_{\tilde{\theta}} \rangle_{\theta}. \end{aligned}$$

Dividing by Δt , taking the limit $\Delta t \rightarrow 0$, and using that $\langle \nabla_{\theta} \ln p_{\theta} \rangle_{\theta} = 0$, we get

$$\frac{d}{dt} \theta = G(\theta)^{-1} \langle \mathcal{L}^{\dagger} \mathcal{P}_{x(t)}^{\dagger} \nabla_{\theta} \ln p_{\theta} \rangle_{\theta}. \quad (31)$$

This is a closed equation for the parameters θ . Using the resulting approximate solution p_{θ} of the filtering equation, all necessary expectations, in particular the marginal transition rates, can be computed.

When the process X jumps, the filtering distribution jumps according to (16), so the approximation $p_{\theta(t-)}$ immediately before the jump is updated to

$$p = \frac{\mathcal{P}_{x(t+)} \mathcal{L} p_{\theta(t-)}}{[\mathcal{S} \mathcal{P}_{x(t+)} \mathcal{L} p_{\theta(t-)}]}.$$

Here too we can obtain an updated approximation within the parametric family by minimizing the relative entropy, i.e., choosing $\theta(t+)$ to minimize $D[p \parallel p_{\theta(t+)}]$. In general, this will be impractical. However, usually one will choose p_{θ} to be an exponential family

$$p_{\theta}(\hat{x}) = \frac{1}{Z(\theta)} \exp \left\{ \sum_{k=1}^K \theta_k \phi_k(\hat{x}) \right\} q(\hat{x}). \quad (32)$$

In this case, minimizing the relative entropy amounts to matching moments, i.e., choosing $\theta(t+)$ so that $\langle \phi_k \rangle_{\theta(t+)} = \langle \phi_k \rangle_p$ for $k = 1, \dots, K$, which is practical.

The entropic matching equations (applied in the context of filtering for stochastic differential equations) were proposed in [20] and derived using a projection argument employing the Fisher-Rao information metric. This geometrical approach to (31), which we describe in Appendix B, is completely analogous to the projection leading to the filtering equation. In this way, entropic matching is seen to be a very natural way to produce a finite-dimensional approximation to the filtering equation, in addition to the justification provided above.

C. Example: Totally asymmetric exclusion process

In this section we apply the marginal process framework to the TASEP on the line with open boundaries. The TASEP [22] describes particles hopping on N sites X_1, \dots, X_N where each site can be occupied by at most one particle. We take $X_n = 1$ when site X_n is occupied and $X_n = 0$ otherwise. If the first site X_1 is empty, a particle can enter at a rate α . If site X_{n+1} is empty and site X_n occupied, a particle can move from X_n to X_{n+1} with rate c . Finally, a particle at the last site X_N can leave the system with rate β .

We consider the situation where only the dynamics of the last site X_N is of interest to us, which might serve as a proxy for, say, the flux through the entire system.

Thus, the only transitions which will be retained are the two transitions corresponding to a particle entering or leaving site X_N . For simplicity, in the notation in this section, we do not distinguish between the variable of interest X_N and the remaining variables. The filtering moment equations for the mean occupancies read

$$\begin{aligned} \frac{d}{dt} \langle x_1 \rangle_t &= \alpha \langle 1 - x_1 \rangle_t - c \langle x_1 (1 - x_2) \rangle_t, \\ \frac{d}{dt} \langle x_n \rangle_t &= c \langle x_{n-1} (1 - x_n) \rangle_t - c \langle x_n (1 - x_{n+1}) \rangle_t, \end{aligned}$$

$$n = 2, \dots, N - 2,$$

$$\begin{aligned} \frac{d}{dt} \langle x_{N-1} \rangle_t &= c \langle x_{N-2} (1 - x_{N-1}) \rangle_t \\ &\quad - c [1 - x_N(t)] \langle x_{N-1} \rangle_t \langle 1 - x_{N-1} \rangle_t. \end{aligned}$$

As expected and is well known, these contain second-order moments. Here we are interested in obtaining the simplest possible approximate marginal process. Thus, we will obtain closed equations in terms of the first-order moments $\langle x_1 \rangle_t, \dots, \langle x_{N-1} \rangle_t$ only. A very natural approach to obtain such a closure is to use entropic matching with a product Bernoulli distribution ansatz

$$p_{\theta}(x) = \prod_{n=1}^{N-1} \theta_n^{x_n} (1 - \theta_n)^{1-x_n}.$$

We refer to the resulting approximate marginal process as the Bernoulli marginal TASEP. After application of product Bernoulli entropic matching, the closed filtering moment equations, in between observations, are given by

$$\begin{aligned} \frac{d}{dt} \theta_1(t) &= \alpha [1 - \theta_1(t)] - c \theta_1(t) [1 - \theta_2(t)], \\ \frac{d}{dt} \theta_n(t) &= c \theta_{n-1}(t) [1 - \theta_n(t)] - c \theta_n(t) [1 - \theta_{n+1}(t)], \\ n &= 2, \dots, N - 2, \\ \frac{d}{dt} \theta_{N-1}(t) &= c \theta_{N-2}(t) [1 - \theta_{N-1}(t)] \\ &\quad - c [1 - x_N(t)] \theta_{N-1}(t) [1 - \theta_{N-1}(t)]. \end{aligned}$$

Unsurprisingly, the resulting equations are identical to a naive mean-field approximation. Note that when $x_N(t) = 1$, i.e., the last site is occupied, the observation term (the last term of the last line) vanishes because a particle cannot enter the last site.

When a particle leaves site X_N , no update to the filtering distribution moments is required. When a particle enters site X_N at time t , the update is simply given by

$$\begin{aligned} \theta_n(t+) &= \theta_n(t-), \quad n = 1, \dots, N - 2, \\ \theta_{N-1}(t+) &= 0. \end{aligned}$$

This is intuitively clear: Site X_{N-1} is necessarily empty immediately after a particle enters site X_N . The means of the remaining sites are left unchanged because of the product-form closure employed.

We performed Monte Carlo simulations of both the Bernoulli marginal TASEP and the full TASEP to compare their behavior. In Fig. 4 we plot the distribution of waiting times between a particle leaving site X_N and the next particle entering

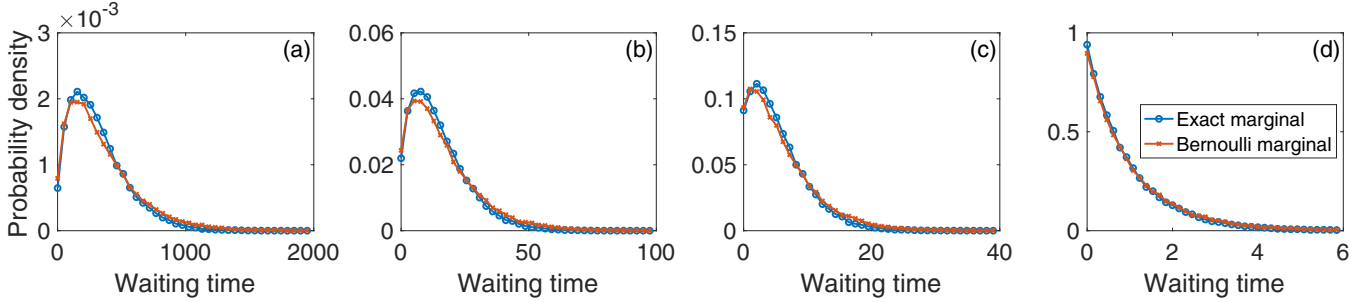


FIG. 4. Numerical evaluation of the accuracy of the Bernoulli marginal TASEP approximation on $N = 10$ sites at stationarity. Waiting-time distributions are shown for a particle to enter site X_N after the previous particle left X_N . The parameters are $\alpha = \beta = 1$ and (a) $c = 0.01$, (b) $c = 0.1$, (c) $c = 1$, and (d) $c = 2$. Distributions are from 100 000 samples.

X_N when the process is at stationarity. The waiting-time distribution varies depending on the parameters of the process. The Bernoulli marginal process reproduces the exact results with high accuracy, despite the fact that we have used a very simple closure for the filtering equation.

VI. PRODUCT POISSON MARGINAL PROCESS

In this section we will apply our results to general reaction networks. While the marginal process framework described in Secs. IV and V is very flexible, it does not provide any indication of how to close the filtering moment equations. Each choice of closure leads to a different (approximate) marginal process. In Sec. VC we chose what is presumably the simplest-possible nontrivial closure (depending on a single parameter for each variable of the filtering equation) for the TASEP.

In this section, in order to initiate the systematic study of the marginal process framework for reaction networks, we similarly investigate what is arguably the simplest nontrivial closure for reaction networks with mass-action kinetics, one of the most important classes of reaction networks. Throughout, we focus on the marginal process Y as introduced in Sec. IVC, which is more convenient here.

A. Product Poisson closure

We consider a general reaction network (3) with mass-action rates. In Sec. VC we employed product Bernoulli entropic matching, which lead to naive mean-field equations. In the context of general reaction networks, we note that a naive first-order mean-field closure, in which the variance is set to zero, leads to the vanishing of the second term (corresponding to the observations) in (28). Instead, we will obtain a principled closure by employing entropic matching using a product Poisson distribution

$$p_\theta(\hat{x}) = \prod_{n=1}^{\hat{N}} e^{-\theta_n} \frac{\theta_n^{\hat{x}_n}}{\hat{x}_n!}.$$

Applying product Poisson entropic matching to a mass-action reaction network leads to equations for the Poisson means that coincide with the macroscopic reaction rate equations. This result will then also hold for the first term in (28), which corresponds to the prior evolution of the environment species. In this sense, product Poisson entropic matching behaves similar (though not identical) to a naive mean-field

closure. However, unlike for a naive first-order closure, the term corresponding to the observations in (28) does not vanish. We obtain

$$\frac{d}{dt}\theta = \sum_{j \in \overline{J_X}} \Omega c_j f_j(x(t)) \hat{g}_j(\theta) \hat{v}_j - \sum_{j \in J_X} \Omega c_j f_j(x(t)) \hat{g}_j(\theta) \hat{s}_j \quad (33)$$

for the continuous part of the filtering equation and

$$\theta(t+) = \theta(t-) + \hat{r}_j \quad (34)$$

when the marginal process jumps via reaction j , where $\hat{r}_j = (\hat{r}_{1j}, \dots, \hat{r}_{\hat{N}j})$ and $\hat{s}_j = (\hat{s}_{1j}, \dots, \hat{s}_{\hat{N}j})$. Here \hat{g}_j are the mass-action rates for the environment species in their macroscopic form, given by

$$\hat{g}_j(\theta) = \prod_{n=1}^{\hat{N}} \left(\frac{\theta_n}{\Omega} \right)^{\hat{s}_{nj}}.$$

The simplicity of (34) is the reason for considering the marginal process Y . If we instead consider the marginal process X , the corresponding equation is more complicated and the results obtained in the following would not hold.

In order to better understand the Poisson marginal process, we can investigate its moment equations. For this purpose, we consider the augmented process (X', Θ) , where X' is the approximate marginal process and Θ is the stochastic process corresponding to the Poisson means θ . As in Sec. IVD, (X', Θ) is a piecewise deterministic Markov process. For such a process, from the known form of the backward evolution operator [23], we obtain the moment equation for a function $\psi(x, \theta)$ in the form

$$\begin{aligned} \frac{d}{dt}\langle \psi \rangle &= \sum_{j \in \overline{J_X}} \langle \Omega c_j f_j(x) \hat{g}_j(\theta) \hat{v}_j^\dagger \nabla_\theta \psi \rangle \\ &\quad - \sum_{j \in J_X} \langle \Omega c_j f_j(x) \hat{g}_j(\theta) \hat{s}_j^\dagger \nabla_\theta \psi \rangle \\ &\quad + \sum_{j \in J_X} \langle \Omega c_j f_j(x) \hat{g}_j(\theta) [\psi(x + \nu_j, \theta + \hat{r}_j) - \psi(x, \theta)] \rangle. \end{aligned} \quad (35)$$

Here and in the following we consider all vectors as column vectors. In particular [writing for brevity $\chi_j = \chi_j(x, \theta) = \Omega c_j f_j(x) \hat{g}_j(\theta)$], the first-order moment equations are

given by

$$\begin{aligned}\frac{d}{dt}\langle x \rangle &= \sum_{j \in J_{\mathbb{X}}} \langle \chi_j \rangle v_j = \sum_{j=1}^R \langle \chi_j \rangle v_j, \\ \frac{d}{dt}\langle \theta \rangle &= \sum_{j \in \overline{J_{\mathbb{X}}}} \langle \chi_j \rangle \hat{v}_j - \sum_{j \in J_{\mathbb{X}}} \langle \chi_j \rangle \hat{s}_j + \sum_{j \in J_{\mathbb{X}}} \langle \chi_j \rangle \hat{r}_j \\ &= \sum_{j=1}^R \langle \chi_j \rangle \hat{v}_j,\end{aligned}$$

where we used that $\hat{v}_j = \hat{r}_j - \hat{s}_j$ and, by definition of $\overline{J_{\mathbb{X}}}$, $v_j = 0$ for each $j \in \overline{J_{\mathbb{X}}}$. For a linear reaction network, these equations are identical to the first-order moment equations obtained for the full process (X, \hat{X}) . Since these equations are closed, we see that the Poisson marginal process, for a linear reaction network, reproduces the mean of the exact marginal process.

Similarly, we can investigate the relation between the covariance matrices of the Poisson marginal and the full process by considering the second-order moment equations, which for the augmented Poisson marginal process (X', Θ) are given by

$$\begin{aligned}\frac{d}{dt}\langle xx^\dagger \rangle &= \sum_{j \in J_{\mathbb{X}}} \langle \chi_j [xv_j^\dagger + v_jx^\dagger + v_jv_j^\dagger] \rangle, \\ \frac{d}{dt}\langle x\theta^\dagger \rangle &= \sum_{j=1}^R \langle \chi_j x \hat{v}_j^\dagger \rangle + \sum_{j \in J_{\mathbb{X}}} \langle \chi_j [v_j\theta^\dagger + v_j\hat{r}_j^\dagger] \rangle, \\ \frac{d}{dt}\langle \theta\theta^\dagger \rangle &= \sum_{j=1}^R \langle \chi_j [\theta \hat{v}_j^\dagger + \hat{v}_j\theta^\dagger] \rangle + \sum_{j \in J_{\mathbb{X}}} \langle \chi_j \hat{r}_j \hat{r}_j^\dagger \rangle.\end{aligned}$$

Denote by

$$S = \begin{bmatrix} \langle xx^\dagger \rangle & \langle x\hat{x}^\dagger \rangle \\ \langle \hat{x}x^\dagger \rangle & \langle \hat{x}\hat{x}^\dagger \rangle \end{bmatrix}, \quad S' = \begin{bmatrix} \langle xx^\dagger \rangle & \langle x\theta^\dagger \rangle \\ \langle \theta x^\dagger \rangle & \langle \theta\theta^\dagger \rangle \end{bmatrix}$$

the matrices of second-order moments for the full process (X, \hat{X}) and for the augmented Poisson marginal process (X', Θ) , respectively. For a linear reaction network, these then evolve according to

$$\begin{aligned}\frac{d}{dt}S &= AS + SA^\dagger + B(t), \\ \frac{d}{dt}S' &= AS' + S'A^\dagger + B'(t),\end{aligned}\quad (36)$$

with matrices A , $B(t)$, and $B'(t)$. The difference between the matrices $B(t)$ and $B'(t)$ is given by

$$\begin{aligned}B(t) - B'(t) &= \sum_{j=1}^R \langle \chi_j \rangle_t \begin{bmatrix} 0 & -v_j \hat{s}_j^\dagger \\ -\hat{s}_j v_j^\dagger & \hat{v}_j \hat{v}_j^\dagger \end{bmatrix} \\ &\quad - \sum_{j \in J_{\mathbb{X}}} \langle \chi_j \rangle_t \begin{bmatrix} 0 & 0 \\ 0 & \hat{r}_j \hat{r}_j^\dagger \end{bmatrix}.\end{aligned}$$


Using variation of constants to solve (36), we find that the difference between second-order moments of exact and

approximate processes is given by

$$S(t) - S'(t) = \int_0^t e^{(t-\tau)A} [B(\tau) - B'(\tau)] e^{(t-\tau)A^\dagger} d\tau,$$

where we assumed $S(0) = S'(0)$. In particular, if $B(t) - B'(t)$ is, say, positive semidefinite for all $t \geq 0$, this will also hold for $S(t) - S'(t)$.

Even when the reaction network is not linear, the macroscopic rates \hat{g}_j coincide with the transition rates \hat{f}_j to leading order in the system size Ω when expressed in terms of concentrations. One might then expect that in the large system size limit, the mean of the Poisson marginal process will coincide with the mean of the exact marginal process. We now investigate these findings numerically on the Lotka-Volterra system



a simple model of predator-prey interaction with oscillatory dynamics. Here we take the prey species \hat{X} to be part of the environment, while the predator species X constitutes the subnet. Numerical results for various system sizes are shown in Fig. 5. We see the expected behavior: With increasing system size, the mean of the Poisson marginal process approaches the exact mean. We also see that the Poisson marginal process underestimates the variance of the exact marginal process.

B. Explicit representation of marginal rates

The representation of the marginal process that we have considered in the previous sections involves auxiliary variables, either the filtering distribution itself or the filtering distribution moments. It is interesting to represent the (approximate) marginal process in a way that explicitly shows its memory. A subclass of systems for which this is readily done for the Poisson marginal process are processes with transition rates linear in the environment variables. Note that this does not imply that the joint reaction network (3) is linear. For example, the Lotka-Volterra system (37) satisfies this condition. This will put the marginal process framework in a form more similar to other approaches for obtaining reduced models that have recently been investigated [12–14].

Assume that (33) has the form

$$\frac{d}{dt}\theta = F(x(t))\theta \quad (38)$$

for the appropriate matrix F , which is explicitly time dependent through $x(t)$. The extension of the following results to the case where (38) contains an inhomogeneity is obvious. We write

$$V(t, \tau) = \overleftarrow{\mathcal{T}} \exp \left\{ \int_\tau^t F(x(t')) dt' \right\}$$

for the time-ordered exponential [which reduces to a product of finitely many ordinary exponentials because $x(t)$ is piecewise constant]. Noting from (34) that the increments of θ at jumps of the marginal process are independent of θ , we can represent the solution of the filtering equation as

$$\theta(t) = V(t, 0)\theta(0) + \sum_{i \in J_{\mathbb{X}}} \hat{r}_i \int_0^t V(t, \tau) dy_i(\tau).$$

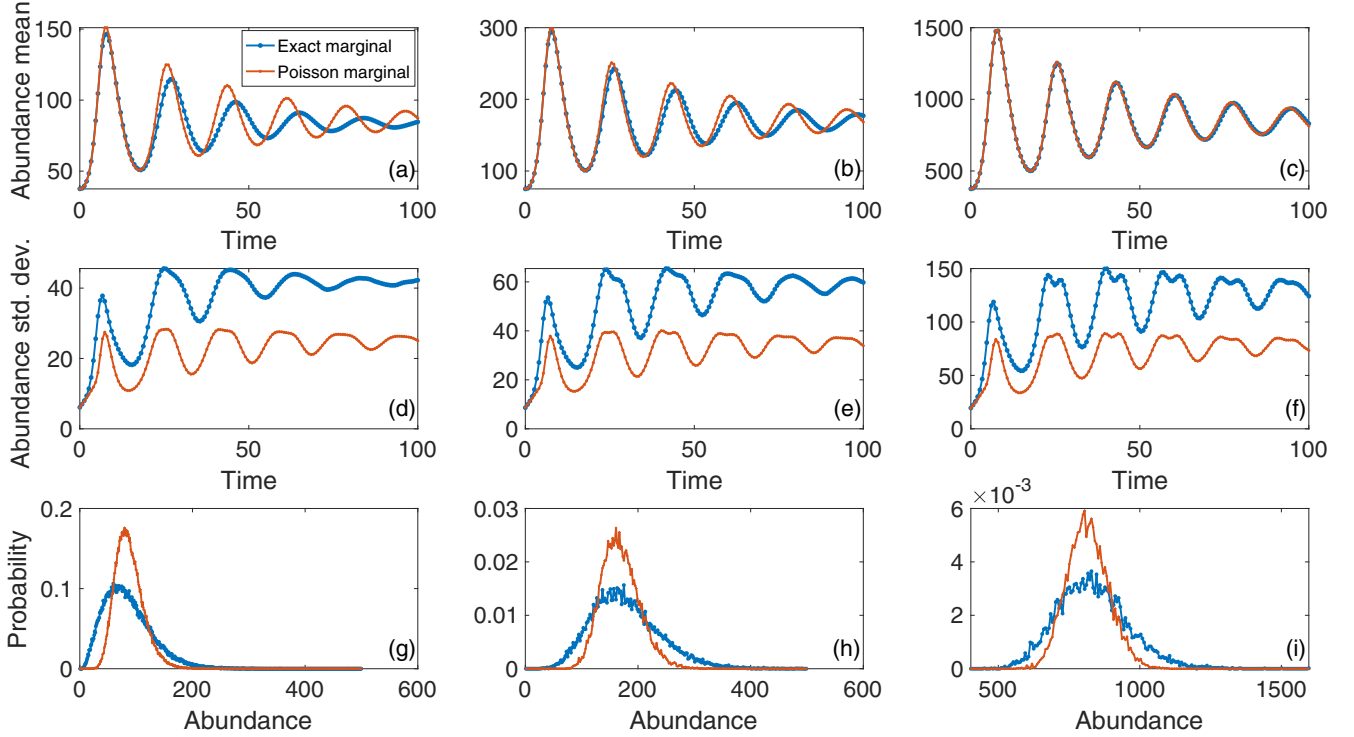


FIG. 5. Monte Carlo evaluation of the accuracy of the Poisson marginal process for the Lotka-Volterra system at system size (a), (d), and (g) $\Omega = 0.5$, (b), (e), and (h) $\Omega = 1$, and (c), (f), and (i) $\Omega = 5$. The mean of the Poisson marginal process approaches the true mean as the system size increases. In all three cases, the Poisson marginal process has a smaller standard deviation than the exact process. (g)–(i) show the full distributions at time $t = 100$. The parameters are $c_1 = 1$, $c_2 = 5$, $c_3 = 0.5$, $c_4 = 0.003$, and $c_5 = 0.3$. Initial conditions were of product Poisson form with means $\langle X \rangle = \langle \hat{X} \rangle = 75\Omega$. The number of simulated trajectories was 100 000 for $\Omega = 0.5$, 20 000 for $\Omega = 1$, and 10 000 for $\Omega = 5$.

Note that Y is a piecewise-constant process with jumps of size 1, so the Stieltjes integral reduces to a sum. Thus, a fully explicit representation for the marginal reaction rate of reaction channel $j \in J_{\mathbb{X}}$ at time t of the marginal process is

$$\langle h_j \rangle_\theta = h_j \left(x(t), V(t, 0)\theta(0) + \sum_{i \in J_{\mathbb{X}}} \hat{r}_i \int_0^t V(t, \tau) dy_i(\tau) \right).$$

We can apply this result to the Lotka-Volterra system (37). The equation for the Poisson mean (33) reads (setting $\Omega = 1$ for simplicity)

$$\frac{d}{dt}\theta = c_1 + c_3\theta - c_4x(t)\theta.$$

The only reaction with a rate depending on θ is $\hat{X} + X \xrightarrow{c_4} 2X$. Setting $v(t, \tau) = \int_\tau^t [c_3 - c_4x(t')]dt'$ and noting that $\hat{r}_2 = \hat{r}_4 = \hat{r}_5 = 0$, we obtain for the marginal reaction rate of this reaction the explicit representation

$$\langle h_4 \rangle_\theta = c_4x(t) \left[e^{v(t, 0)}\theta(0) + \int_0^t e^{v(t, \tau)} c_1 d\tau \right].$$

C. Limitations of the product Poisson closure

Using a product Poisson ansatz to close the filtering equation will not always be appropriate. The most obvious situation where this approach might fail is in the presence of conservation relations among the environment species. This

will be particularly problematic when there is no intrinsic noise in the environment. A simple example for this behavior would be the gene expression network (5) with input rate c_1 and decay rate c_2 for the mRNA set to zero. Irrespective of the initial distribution of mRNA at time zero, the filtering distribution will converge to a unit mass at the true mRNA abundance as the time interval over which the subnet process is observed tends to infinity. Since for a Poisson distribution the variance is equal to the mean, the vanishing of the variance of the true filtering distribution over time cannot be captured by the Poisson closure.

VII. CONCLUSION

In this article we have illustrated how the marginal process framework, in combination with entropic matching, results in a principled model reduction method for Markov jump processes. We derived the filtering equation for two fully coupled Markov jump processes in a transparent form that shows its relation to the full master equation, establishing the filtering equation as the result of the application of a projection operator. The application of a further projection, given by entropic matching, results in a finite-dimensional approximation to the filtering equation and thus in a tractable approximation for the marginal process. Apart from being a useful tool for the efficient stochastic simulation of the marginal process, it also provides a theoretical understanding of the marginal description. A particularly simple instantiation

of the marginal process framework for mass-action reaction networks, the Poisson marginal process, was investigated in detail. We derived analytical results for the approximation error for linear reaction networks. A similar approximation, based on product Bernoulli entropic matching, was employed for the TASEP.

An interesting question for future investigation is the theoretical analysis of approximations of the filtering equation more accurate than the product Poisson ansatz or the product Bernoulli ansatz. In particular, approximations based on first- and second-order moments might allow one to compare the marginal process framework to other marginalization approaches published previously [12–14]. For this purpose, a variant of the marginal process framework for subnet and environment both modeled by stochastic differential equations would be of interest.

ACKNOWLEDGMENTS

H.K. acknowledges support from the LOEWE research priority program CompuGene.

APPENDIX A: MARGINAL SIMULATION ALGORITHM

Here we describe one possible way to simulate the (approximate) marginal process for reaction networks. Let

$$\frac{d}{dt}\theta = v(\theta, x) \quad (\text{A1})$$

be the differential equation governing the parameters θ of the (approximate or exact) solution of the filtering equation in between jumps, which in general will depend on the marginal process state x . For example, for the Poisson marginal process, v is given by the right-hand side of (33). Similarly, let

$$\theta_+ = v_j(\theta_-, x_-) \quad (\text{A2})$$

be the equation specifying the update to the parameters θ when the subnet jumps via reaction j . For the Poisson marginal process, v_j is given by the right-hand side of (34).

An algorithm [24] based on the modified next reaction method [25] can be formulated as follows. The expected

Algorithm 1 Marginal stochastic simulation algorithm (modified next reaction method).

```

Set  $t \leftarrow 0, x \leftarrow x_0, \theta \leftarrow \theta_0.$  ▷ Initialization
for  $j \in J_{\mathbb{X}}$  do
    Sample  $u \sim \text{Uniform}(0,1).$ 
    Set  $\tau_j \leftarrow -\ln u.$ 
end for

while  $t < T$  do ▷ Main loop
    Solve (A1) and (A3) until the first variable  $\tau_{j_*}$  reaches 0 for
    some index  $j_* \in J_{\mathbb{X}}.$ 
    Update  $\theta \leftarrow v_{j_*}(\theta, x).$ 
    Update  $x \leftarrow x + v_{j_*}.$ 
    Sample  $u \sim \text{Uniform}(0,1).$ 
    Set  $\tau_{j_*} \leftarrow -\ln u.$ 
end while
    
```

reaction rates $\langle h_j \rangle_\theta, j \in J_{\mathbb{X}}$ of those reactions that modify the state of Y are functions of θ . We augment the ODE system (A1) to include auxiliary variables

$$\frac{d}{dt}\tau_j = -\langle h_j \rangle_\theta, \quad j \in J_{\mathbb{X}}. \quad (\text{A3})$$

The system can then be simulated using Algorithm 1, which samples a trajectory of the (approximate) marginal process over the time interval $[0, T]$ starting from an initial subnet state x_0 and initial parameters θ_0 for the filtering distribution at time 0. The algorithm has to find the time point at which a function of the ODE system state crosses a specified threshold (one of the variables τ_j reaches 0). This is a functionality provided by many ODE solvers, so the algorithm is straightforward to implement.

APPENDIX B: FILTERING EQUATION AND ENTROPIC MATCHING AS PROJECTION OPERATIONS

Here we discuss how the continuous part of the filtering equation (15) and the entropic matching equation (31) arise as an application of an orthogonal projection (using the Fisher-Rao information metric) applied to the vector field defined by the ME. Since the results in this paper do not actually depend on any of the results in this appendix, we restrict the discussion to a form which stresses the geometrical significance and neglects any technical difficulties. See [26] for a general treatment of information geometry.

For simplicity, assume that $\mathbb{X} \times \hat{\mathbb{X}}$ is finite and define the set of probability distributions on $\mathbb{X} \times \hat{\mathbb{X}}$,

$$\mathbb{P} = \{p : \mathbb{X} \times \hat{\mathbb{X}} \rightarrow [0,1] \mid \mathcal{S}p = 1\},$$

which inherits a manifold structure as a subset of finite-dimensional Euclidean space. The tangent space at a point $p \in \mathbb{P}$ is given by

$$T_p\mathbb{P} = \{p\} \times \{v : \mathbb{X} \times \hat{\mathbb{X}} \rightarrow \mathbb{R} \mid \mathcal{S}v = 0\}.$$

For an MJP (X, \hat{X}) on $\mathbb{X} \times \hat{\mathbb{X}}$, the master equation $dp_t(x, \hat{x})/dt = [\mathcal{L}p_t](x, \hat{x})$ defines a vector field on \mathbb{P} , the vector attached at a point p being $\mathcal{L}p$. We define a base-point-dependent metric by

$$g_p(v, w) = \sum_{x, \hat{x}} \frac{v(x, \hat{x})w(x, \hat{x})}{p(x, \hat{x})} \quad p(x, \hat{x}) \neq 0$$

for $v, w \in T_p\mathbb{P}$. When restricted to $p \in \mathbb{P}$ with $p > 0$ everywhere, this is the information metric. Our extension to other p is somewhat *ad hoc* but sufficient for our purposes.

It turns out that the continuous part of the filtering equation when X is in state $x(t)$ is obtained simply as an orthogonal projection of the vector field of the full master equation on the tangent space to the submanifold

$$\mathbb{P}_{x(t)} = \{p \in \mathbb{P} \mid p(x, \hat{x}) = \delta_{x(t), x} \hat{p}(\hat{x}) \text{ for some } \hat{p}(\hat{x})\}.$$

From now on, let $p \in \mathbb{P}_{x(t)}$ with $p(x, \hat{x}) = \delta_{x(t), x} \hat{p}(\hat{x})$ and assume $\hat{p} > 0$ everywhere. One easily checks that the linear operator

$$\mathcal{F}_p : T_p\mathbb{P} \rightarrow T_p\mathbb{P}, \quad \mathcal{F}_p v = (\mathcal{P}_{x(t)} - p\mathcal{S}\mathcal{P}_{x(t)})v$$

satisfies $\mathcal{F}_p^2 = \mathcal{F}_p$ and $g_p(\mathcal{F}_p v, w) = g_p(v, \mathcal{F}_p w)$ for all $v, w \in T_p \mathbb{P}$, so \mathcal{F}_p is an orthogonal projection. The projected vector field $p \mapsto (p, \mathcal{F}_p \mathcal{L} p)$ corresponds to the filtering equation.

The entropic matching equations can similarly be derived as an application of a further projection. They were derived in [20] using such a geometric approach. Considering a K -dimensional parametric family p_θ of probability distributions on $\hat{\mathbb{X}}$, the map $\theta \mapsto \delta_{x(t), x} p_\theta(\hat{x})$ defines a submanifold of $\mathbb{P}_{x(t)}$ which we denote by $\mathbb{P}'_{x(t)}$. From now on, let $p \in \mathbb{P}'_{x(t)}$ with $p(x, \hat{x}) = \delta_{x(t), x} p_\theta(\hat{x})$. The tangent space $T_p \mathbb{P}'_{x(t)}$ to this submanifold is spanned by the vectors

$$v_k = v_k(x, \hat{x}) = \delta_{x(t), x} \frac{\partial p_\theta(\hat{x})}{\partial \theta_k}, \quad k = 1, \dots, K.$$

We then find for the Gram matrix

$$g_p(v_k, v_l) = G_{kl}(\theta), \quad k, l = 1, \dots, K,$$

i.e., the information metric as defined by (30). The orthogonal projection $\mathcal{Q}_p : T_p \mathbb{P} \rightarrow T_p \mathbb{P}'_{x(t)}$ is

given by

$$\mathcal{Q}_p w = \sum_{k,l=1}^K g_p(w, v_k) [G(\theta)^{-1}]_{kl} v_l.$$

Because $T_p \mathbb{P}'_{x(t)}$ is a subspace of $T_p \mathbb{P}_{x(t)}$, we have $\mathcal{Q}_p \mathcal{F}_p = \mathcal{Q}_p$. The resulting projected ME, defined by the vector field

$$p \mapsto (p, \mathcal{Q}_p \mathcal{F}_p \mathcal{L} p) = (p, \mathcal{Q}_p \mathcal{L} p),$$

evolves on the manifold $\mathbb{P}'_{x(t)}$ when it is started there. When this is written in terms of the variables θ , we obtain the equations of entropic matching (31). Here we also see that entropic matching can be used to directly obtain a finite-dimensional approximation to the filtering equation from the ME, without deriving the filtering equation in an intermediate step. The derivation presented in Sec. VB could also be adapted in this way and would then be an application of variational inference [27].

-
- [1] N. G. van Kampen, *Stochastic Processes in Physics and Chemistry* (Elsevier, Amsterdam, 2007).
- [2] D. T. Gillespie, *J. Phys. Chem.* **81**, 2340 (1977).
- [3] C. Zechner and H. Koepl, *PLoS Comput. Biol.* **10**, e1003942 (2014).
- [4] K. Giesecke, H. Kakavand, and M. Mousavi, *Oper. Res.* **59**, 1233 (2011).
- [5] Y. Cao, D. T. Gillespie, and L. R. Petzold, *J. Chem. Phys.* **122**, 014116 (2005).
- [6] A. Hellander and P. Lötstedt, *J. Comput. Phys.* **227**, 100 (2007).
- [7] T. Jahnke, *Multiscale Model. Simul.* **9**, 1646 (2011).
- [8] P. Thomas, R. Grima, and A. V. Straube, *Phys. Rev. E* **86**, 041110 (2012).
- [9] G. W. Constable, A. J. McKane, and T. Rogers, *J. Phys. A* **46**, 295002 (2013).
- [10] H.-W. Kang and T. G. Kurtz, *Ann. Appl. Probab.* **23**, 529 (2013).
- [11] J. K. Kim, G. A. Rempala, and H.-W. Kang, *Multiscale Model. Simul.* **15**, 1376 (2017).
- [12] K. J. Rubin, K. Lawler, P. Sollich, and T. Ng, *J. Theor. Biol.* **357**, 245 (2014).
- [13] B. Bravi, P. Sollich, and M. Opper, *J. Phys. A* **49**, 194003 (2016).
- [14] B. Bravi and P. Sollich, *Phys. Biol.* **14**, 045010 (2017).
- [15] T. Ramalho, M. Selig, U. Gerland, and T. A. Enßlin, *Phys. Rev. E* **87**, 022719 (2013).
- [16] L. Bronstein and H. Koepl, *J. Chem. Phys.* **148**, 014105 (2018).
- [17] C. Zechner, M. Unger, S. Pelet, M. Peter, and H. Koepl, *Nat. Methods* **11**, 197 (2014).
- [18] A. Bain and D. Crisan, *Fundamentals of Stochastic Filtering* (Springer, Berlin, 2009).
- [19] C. Kuehn, in *Control of Self-Organizing Nonlinear Systems*, edited by E. Schöll, S. H. L. Klapp, and P. Hövel (Springer, Cham, 2016), pp. 253–271.
- [20] D. Brigo, B. Hanzon, and F. Le Gland, *Bernoulli* **5**, 495 (1999).
- [21] R. H. Leike and T. A. Enßlin, *Entropy* **19**, 402 (2017).
- [22] B. Derrida, *Phys. Rep.* **301**, 65 (1998).
- [23] A. Crudu, A. Debussche, A. Muller, and O. Radulescu, *Ann. Appl. Probab.* **22**, 1822 (2012).
- [24] A. Alfonsi, E. Cancès, G. Turinici, B. Di Ventura, and W. Huisinga, in *CEMRACS 2004—Mathematics and Applications to Biology and Medicine*, ESAIM: Proceedings, edited by E. Cancès and J.-F. Gerbeau (EDP Sciences, Les Ulis, 2005), Vol. 14, pp. 1–13.
- [25] D. F. Anderson, *J. Chem. Phys.* **127**, 214107 (2007).
- [26] N. Ay, J. Jost, H. Vân Lê, and L. Schwachhöfer, *Information Geometry* (Springer, Berlin, 2017).
- [27] M. J. Wainwright and M. I. Jordan, *Found. Trends Mach. Learn.* **1**, 1 (2008).