

Probability distributions of molecular observables computed from Markov models. II. Uncertainties in observables and their time-evolution

John D. Chodera^{1,a)} and Frank Noé^{2,b)}

¹*Research Fellow, California Institute of Quantitative Biosciences (QB3), University of California, Berkeley, 260J Stanley Hall, Berkeley, California 94720, USA*

²*DFG Research Center Matheon, FU Berlin, Arnimallee 6, 14195 Berlin, Germany*

(Received 23 October 2009; accepted 14 June 2010; published online 8 September 2010)

Discrete-state Markov (or master equation) models provide a useful simplified representation for characterizing the long-time statistical evolution of biomolecules in a manner that allows direct comparison with experiments as well as the elucidation of mechanistic pathways for an inherently stochastic process. A vital part of meaningful comparison with experiment is the characterization of the statistical uncertainty in the predicted experimental measurement, which may take the form of an equilibrium measurement of some spectroscopic signal, the time-evolution of this signal following a perturbation, or the observation of some statistic (such as the correlation function) of the equilibrium dynamics of a single molecule. Without meaningful error bars (which arise from both approximation and statistical error), there is no way to determine whether the deviations between model and experiment are statistically meaningful. Previous work has demonstrated that a Bayesian method that enforces microscopic reversibility can be used to characterize the statistical component of correlated uncertainties in state-to-state transition probabilities (and functions thereof) for a model inferred from molecular simulation data. Here, we extend this approach to include the uncertainty in observables that are functions of molecular conformation (such as surrogate spectroscopic signals) characterizing each state, permitting the full statistical uncertainty in computed spectroscopic experiments to be assessed. We test the approach in a simple model system to demonstrate that the computed uncertainties provide a useful indicator of statistical variation, and then apply it to the computation of the fluorescence autocorrelation function measured for a dye-labeled peptide previously studied by both experiment and simulation. © 2010 American Institute of Physics. [doi:10.1063/1.3463406]

I. INTRODUCTION

A large variety of biophysical experimental techniques are currently in use, providing information about various aspects of biomolecular structure and dynamics. Unfortunately, these methods often provide only indirect information about the microscopic origins of processes of interest. Spectroscopic methods, for example, are inherently limited by the need to collect a sufficient quantity of photons in order to distinguish relevant features of the system from statistical fluctuations of the measurement process. Assaying macroscopic collections of molecules provides one solution to collecting sufficient photons while achieving high time-resolution, at the cost of sacrificing information about heterogeneity, as only average signals are observed. Conversely, single-molecule experiments can observe the statistical behavior of individual molecules, but at the loss of temporal resolution due to the need to collect sufficient photons to precisely resolve an observable.

Molecular simulation is a powerful complementary tool for probing molecular behavior in biology. With a forcefield sufficiently representative of physical reality, it is possible to probe the stochastic behavior of individual molecules and

their associated dynamical processes in atomic detail with high time-resolution, filling in the gaps between experiment and simplified theoretical models. In practice, however, this requires that practitioners must both validate the simulation against experimental measurements and collect sufficient data to provide mechanistic insight into the behaviors of interest. This process is frustrated by the statistical heterogeneity of dynamics (requiring many realizations be observed) and the long timescales involved in processes of interest (which may greatly exceed our capacity to simulate on all but the largest computers or simplest problems).

The presence of metastable conformational states,^{1,2} often the root cause of slow relaxation times and heterogeneous dynamics in many biomolecular systems, also provides a solution to this problem. A separation of timescales between fast relaxation within a metastable state and long waiting times between interstate transitions allows the statistical behavior of the system to be well-described by a discrete-state Markov model (or master equation model) over times longer than the fast mixing time within states.³ Even if no strong separation of timescales is present, an appropriately constructed Markov model can still provide an excellent approximation to the long-time statistical dynamics of the system.⁴ Because only transition rates between pairs of dynamically connected states need be estimated, these mod-

^{a)}Electronic mail: jchodera@berkeley.edu.

^{b)}Electronic mail: frank.noe@fu-berlin.de.

els can be efficiently constructed using molecular simulation data, where the dynamical trajectories employed need only be longer than the lag time required for a Markov model to well approximate the original continuous dynamics (here referred to as the *Markov time* τ_{eq}).^{5–8} In strongly metastable systems, this Markov time may be orders of magnitude shorter than the slowest relaxation times of the system, thus making Markov models an efficient tool to model slow processes with a large (but usually feasible) number of short trajectories, some or all of which may be generated in parallel. Methods for the construction^{7–9} of such models and determination of an appropriate Markov time^{4–7,10–12} are described in detail elsewhere.

Because these models are constructed from a finite quantity of simulation data, there will be a degree of statistical uncertainty in the model parameters and in the characterization of each discrete-state. This uncertainty will consequently induce uncertainties in quantities computed from the model, often in a complex way. Accurate characterization of these uncertainties is paramount to either comparison with existing experimental data (to validate the model) or the computation of new, unobserved properties (to predict unknown quantities). It is impossible to know if the model and experimental data are significantly discrepant, for example, without some estimate of both the uncertainty in the computed quantities arising from model uncertainty and the error in the experimental measurements.

Previously, a Bayesian framework was proposed for characterizing how finite quantities of simulation data produce uncertainties in the transition probabilities for a Markov model of physical systems satisfying detailed balance.¹³ Here, we extend this framework to include the uncertainty in observables (such as spectroscopically observable quantities) for each of these states, allowing the full uncertainty in equilibrium and time-resolved experimentally observable quantities computed from the model to be characterized. We stress, however, that the statistical error is only part of the contribution to the total *prediction error* of the model—the systematic discretization error from a choice of finite number of states, a particular state partitioning, and a fixed Markov time also contributes to this error,⁴ and the interaction between statistical and approximation error is a topic of further study. In this work, we restrict ourselves to the consideration of the statistical component of the prediction error, which, in a number of practical cases, may dominate the prediction error.

This article is organized as follows. Section II reviews the statistical mechanics of common measurements and their computation from Markov models. Section III briefly describes the Bayesian inference framework for sampling transition matrices that satisfy detailed balance. Section IV introduces the proposed approach for incorporating errors in estimated observables into this framework by adding a second sampling step to the Bayesian procedure in Sec. III. Section V validates the new procedure on an analytically tractable model system, and Sec. VI applies it to the calculation of spectroscopically observable quantities for a fluorescent peptide.

II. OBSERVABLES AND MARKOV MODELS

Biophysical spectroscopic measurements, although varied in terms of the quantity being measured, commonly fall into one of three categories.

Equilibrium measurements (e.g., equilibrium circular dichroism, infrared spectroscopy, and fluorescence measurements) represent ensemble-averages (or *expectations*) of a conformation-sensitive spectroscopic signal $A(x)$ over some equilibrium configurational probability density $\pi(x)$, where x denotes the instantaneous molecular configuration,

$$\langle A \rangle = \int dx \pi(x) A(x), \quad (1)$$

Nonequilibrium relaxation experiments (e.g., laser-induced temperature-jump experiments) monitor the relaxation of an ensemble-averaged spectroscopic signal following a perturbation and can be described in terms of the initial nonequilibrium density $\rho_0(x)$ and the time-evolution operator $T(x_t, t; x_0, 0)$ that gives the probability of finding a molecule in configuration x_t at time t given that it was initially in configuration x_0 at time 0, where we assume the process described by $T(x, t; y, t')$ is stationary,

$$\langle A(t) \rangle_{\rho_0} = \int dx_0 \rho_0(x_0) \int dx_t T(x_t, t; x_0, 0) A(x_t). \quad (2)$$

Equilibrium correlation measurements [e.g., fluorescence correlation spectroscopy (FCS), spectral density functions] probe the auto- or cross-correlation function of some experimental observables A and B at equilibrium,

$$\langle A(0)B(t) \rangle = \int dx_0 \pi(x_0) A(x_0) \int dx_t T(x_t, t; x_0, 0) B(x_t). \quad (3)$$

A discrete-state Markov model is defined by a partition of the molecular configuration space Ω into M rapidly equilibrating states \mathcal{S}_i , such that $\mathcal{S}_i \cap \mathcal{S}_j = 0$ if $i \neq j$ and $\bigcup_{i=1}^M \mathcal{S}_i = \Omega$. It is convenient to define a set of *characteristic* or *membership functions* $\phi_i(x)$ which assume a value of unity within the state and zero without

$$\phi_i(x) = \begin{cases} 1 & \text{if } x \in \mathcal{S}_i \\ 0 & \text{otherwise.} \end{cases} \quad (4)$$

Taken together, the $\phi_i(x)$ form a *partition of unity*, such that $\sum_{i=1}^M \phi_i(x) = 1$ for all $x \in \Omega$ (see, e.g., Ref. 14).

Transition dynamics on the space of states is here characterized by a discrete-time Markov model, described by a row-stochastic transition matrix $\mathbf{T}(\tau)$ that evolves a row-vector of state occupation probabilities $\mathbf{p}(t)$ by a fixed observation interval (or *lag time*) τ ,

$$\mathbf{p}(\tau) = \mathbf{p}(0)\mathbf{T}(\tau). \quad (5)$$

The transition matrix $\mathbf{T}(\tau)$ has a well-defined meaning in terms of correlation functions,⁵

$$T_{ij}(\tau) \equiv \frac{\langle \phi_i(0) \phi_j(\tau) \rangle}{\langle \phi_i \rangle}. \quad (6)$$

If dynamics is well-approximated by a Markovian process \mathcal{S} , the transition probability for an integral multiplier n of the time τ is given by exponentiating the transition matrix,

$$\mathbf{T}(n\tau) \approx [\mathbf{T}(\tau)]^n \quad \text{for } \tau \geq \tau_{\text{eq}}, \quad (7)$$

where the minimal lag time τ for which this relationship holds approximately is termed the *Markov time* τ_{eq} .⁷

If the states are sufficiently metastable, the statistical dynamics can be equivalently described in continuous time in terms of a phenomenological rate matrix \mathbf{K} (see, e.g., Ref. 11), whose off-diagonal elements contain strictly positive rate constants,^{15–17} related to the transition matrix by $\mathbf{T}(\tau) = e^{\mathbf{K}\tau}$, where the exponential denotes the formal matrix exponential $e^{\mathbf{A}} \equiv \sum_{n=0}^{\infty} \mathbf{A}^n / n!$. It is critical to note that, while the statistical evolution can be described in continuous time, the model still provides no information about processes occurring on timescales shorter than the Markov time τ_{eq} because they have been coarse-grained out by aggregating configuration space into discrete-states. Here, we focus on the discrete-time transition matrix \mathbf{T} , but the method proposed in Sec. IV for modeling uncertainties in state observables can be combined with the inference of rate matrices \mathbf{K} (Ref. 18) in a straightforward manner.

Because each state mixes quickly compared to the transition time between states, the average value of some spectroscopic signal $A(x)$ over each state can be easily estimated, often from short trajectories that mix rapidly within the state,

$$A_i \equiv \langle A \rangle_i = \int dx \pi_i(x) A(x), \quad (8)$$

with

$$\pi_i(x) \equiv \pi(x) \phi_i(x) / \pi_i, \quad \pi_i \equiv \int dx \pi(x) \phi_i(x), \quad (9)$$

where $\pi_i(x)$ is the equilibrium probability restricted to the state defined by the membership function $\phi_i(x)$ and π_i is the equilibrium probability of occupying state i , which can also be determined from the equilibrium eigenvector of $\mathbf{T}(\tau)$.

The Markovian approximation model defined above allows us to write the three types of experimental observables [Eqs. (1)–(3)], computed from the Markov model defined by $\mathbf{T}(\tau)$ for a fixed $\tau \geq \tau_{\text{eq}}$, as

$$\langle A \rangle \approx \sum_{i=1}^M \pi_i A_i, \quad (10)$$

$$\langle A(n\tau) \rangle_{\rho_0} \approx \sum_{i=1}^M \sum_{j=1}^M \rho_{0i} ([\mathbf{T}(\tau)]^n)_{ij} A_j, \quad (11)$$

$$\langle A(0)B(n\tau) \rangle \approx \sum_{i=1}^M \sum_{j=1}^M \pi_i A_i ([\mathbf{T}(\tau)]^n)_{ij} B_j. \quad (12)$$

Because the transition matrix $\mathbf{T}(\tau)$ is *a priori* unknown, it must be estimated or *inferred* from some data, such as a sample of transition events between discrete-states observed

in one or more molecular dynamics simulations. Since the total number of observed events in a finite timeseries is necessarily finite, the true value of $\mathbf{T}(\tau)$, defined in terms of transition probabilities, will be uncertain.

For a given set of observed transitions from trajectory data, what is the uncertainty of \mathbf{T} and how does this affect the uncertainty of equilibrium or kinetic properties computed from the model? Adequate assessment of this uncertainty and how it influences (potentially complicated, nonlinear) observables of interest is critical in determining when a sufficient quantity of data has been collected for the hypotheses in question to be decided, and whether any discrepancy observed with experimental observation is to be deemed statistically significant. Additionally, one may further exploit knowledge of these uncertainties by planning new simulations such as to most reduce the uncertainties in the observables of interest, and thus to get converged observables with a minimal amount of simulation effort (see, e.g., Refs. 19–22).

Due to the expense of generating the simulation data required to construct a reliable model, it is generally not practical to partition the data into independent subsets to assess the variation in properties of interest among smaller data sets. Bootstrap methods have been applied to the estimation of error in Markov models,⁶ but it is challenging to correctly preserve the correlation structure in the data unless the model is constructed from numerous independent trajectories sampled from equilibrium.

Much effort has therefore been devoted to Bayesian approaches to uncertainty analysis for Markov models. Bayesian methods provide a powerful and simple framework for describing the posterior uncertainty in both model parameters and observables of interest computed from the model, in addition to the potential to quantify uncertainties with information-theoretic tools. Previously, Bayesian approaches have been proposed for inferring transition matrices that do not satisfy detailed balance,¹⁹ but recent work has shown that use of prior knowledge that the dynamics must obey detailed balance²³ can significantly improve the quality of the resulting inference.²⁴ To this end, Bayesian approaches for inferring either reversible transition matrices^{13,24} or rate matrices^{18,25} have been proposed.

Here, we consider the extension of the method for sampling reversible transition matrices proposed in Ref. 13 to also incorporate the uncertainty due to insufficient characterization of the averages of molecular observables within the states. Section III compactly describes the original sampling scheme for transition matrices, while Sec. IV extends this procedure by also addressing the uncertainty in the averages of observables restricted to individual states [Eq. (8)].

III. UNCERTAINTY IN TRANSITION PROBABILITIES

Consider a discrete-state Markovian trajectory of length $L\tau$ starting in a given state s_0 , with observations of the current state made with a time-resolution $\tau \geq \tau_{\text{eq}}$, denoted $\mathbf{s} \equiv \{s_0, s_1, \dots, s_L\}$. Given one or more such trajectories, the $M \times M$ transition count matrix $\mathbf{C} \equiv (c_{ij})$, where c_{ij} is the number of times a discrete trajectory \mathbf{s} appears in state i at

some time index t and state j at time index $t+1$, is a sufficient statistic for capturing all information about the stochastic behavior of this system.¹³ For one trajectory, \mathbf{C} can be written

$$c_{ij} = \sum_{t=1}^L \delta_{is_{t-1}} \delta_{js_t}. \quad (13)$$

In the limit of an infinitely long trajectory, the elements of the true transition matrix are given by the trivial estimator,

$$\hat{T}_{ij}(\tau) = \frac{c_{ij}}{\sum_{k=1}^M c_{ik}} = \frac{c_{ij}}{c_i}, \quad (14)$$

where $c_i \equiv \sum_{k=1}^M c_{ik}$ gives the total number of observed transitions leaving state i . For a trajectory of finite length, the underlying transition matrix $\mathbf{T}(\tau)$ cannot be computed exactly from the observed transition count matrix \mathbf{C} . Instead, we can compute the probability that the true (unknown) transition matrix \mathbf{T} generated the observed counts, assuming the process is Markovian (henceforth suppressing the argument τ),

$$p(\mathbf{C}|\mathbf{T}) = \prod_{t=1}^L T_{s_{t-1}s_t} = \prod_{i,j=1}^M T_{ij}^{c_{ij}}. \quad (15)$$

By Bayes' theorem, the probability that \mathbf{T} is the true transition matrix that generated the data \mathbf{C} is then

$$p(\mathbf{T}|\mathbf{C}) \propto p(\mathbf{C}|\mathbf{T})p(\mathbf{T}) = \left[\prod_{i,j=1}^M T_{ij}^{c_{ij}} \right] p(\mathbf{T}), \quad (16)$$

where $p(\mathbf{T})$ is the prior probability of transition matrices, reflecting our knowledge about \mathbf{T} before observing any data. For uniform prior $p(\mathbf{T})$, the maximum of the likelihood function Eq. (15) is given by the trivial estimator of Eq. (14). In the limit that the number of samples $N \rightarrow \infty$, $p(\mathbf{T}|\mathbf{C})$ becomes progressively more peaked around $\hat{\mathbf{T}}(\tau)$.

For a physical system, the condition of detailed balance [$\pi_i T_{ij}(\tau) = \pi_j T_{ji}(\tau)$] must be satisfied for any τ .²³ We presume also that the Markov chain described by \mathbf{T} is *indecomposable*, so that the unit eigenvalue is nondegenerate. Such a stochastic matrix has a single, dominant eigenvector of 1, and the corresponding left eigenvector $\boldsymbol{\pi}$,

$$\boldsymbol{\pi} \mathbf{T} = \boldsymbol{\pi}, \quad (17)$$

which yields the stationary distribution of the transition matrix when normalized to a 1-norm of 1, such that $\sum_j \pi_j = 1$. All other eigenvalues are real and lie on the interval $(-1, +1)$.

We wish to impose the following constraints on the transition matrix \mathbf{T} through the choice of prior $p(\mathbf{T})$:

C1: Elementwise non-negativity $0 \leq T_{ij} \quad \forall i, j$,

C2: Row-stochasticity $\sum_{j=1}^M T_{ij} = 1 \quad \forall i$,

C3: Detailed Balance $\pi_i T_{ij} = \pi_j T_{ji} \quad \forall i, j$.

Formally, the addition of these constraints to the row-wise Dirichlet prior on \mathbf{T} results in a new prior,

$$p(\mathbf{T}) \propto \prod_{i,j=1}^M T_{ij}^{(b_{ij}-1)} \delta(\pi_i T_{ij} - \pi_j T_{ji}) h(T_{ij}), \quad (18)$$

where $\delta(x)$ is the Dirac delta function, $h(x)$ is the Heaviside function (assuming a value of zero for $x < 0$ and unity for $x \geq 0$) and $\mathbf{B} \equiv (b_{ij})$ is an $M \times M$ matrix of prior *pseudocounts*. Typical choices of these pseudocounts \mathbf{B} include what corresponds to a uniform prior in the absence of constraints C1–C3, $b_{ij} = 1 \quad \forall i, j$,¹³ and the choice corresponding to the null prior, $b_{ij} = 0 \quad \forall i, j$.²⁶ The uniform prior assigns a uniform *a priori* distribution to all matrix elements by adding a full pseudocount to each of them, which takes much observation data to be overridden if the system consists of many states. The null prior, on the other hand, gives most impact to the observed data, as it forces the maximum likelihood estimator and mean of the transition matrix to coincide, and all transition matrix elements in which no count has been observed to zero.

To sample the distribution given by Eq. (16), a sampling procedure based on the Metropolis–Hasting algorithm is used. Given a current matrix \mathbf{T} and a proposed new matrix \mathbf{T}' , the Metropolis–Hastings acceptance probability is

$$P_{\text{accept}}(\mathbf{T} \rightarrow \mathbf{T}') = \min \left\{ 1, \frac{p(\mathbf{T}' \rightarrow \mathbf{T}) p(\mathbf{T}'|\mathbf{C})}{p(\mathbf{T} \rightarrow \mathbf{T}') p(\mathbf{T}|\mathbf{C})} \right\},$$

where $p(\mathbf{T} \rightarrow \mathbf{T}')$ and $p(\mathbf{T}' \rightarrow \mathbf{T})$ denote the probabilities to propose a stochastic move to \mathbf{T}' given \mathbf{T} , and vice versa. Two types of proposal steps (which together generate an ergodic chain) are used and are described briefly below. To ensure the ratio of priors remains calculable, the moves are restricted to only generate proposals \mathbf{T}' that conform to constraints C1–C3. This is achieved by the two move types described in the subsequent subsections, each of which is chosen with a 50% probability in each iteration. For numerical stability, the logarithms of the Metropolis–Hastings acceptance probability are usually computed.

The sampling procedure is initialized with a matrix that fulfills detailed balance,

$$T_{ij}^{\text{init}} = \frac{c_{ij} + c_{ji}}{\sum_{k=1}^M [c_{ik} + c_{ki}]}. \quad (19)$$

When the transition counts c_{ij} come from trajectories that are initiated far from equilibrium, the transition matrix \mathbf{T}^{init} can be very different from the probability maximum or even outside the region of main probability mass. Therefore, it is necessary to run the sampler for a “burn-in” phase long enough until the estimated properties of interest (e.g., uncertainties of the expectation or the correlation values) become stable. This burn-in phase is discarded and the actual estimation of the uncertainties is based on the subsequent production phase of the sampler. Monitoring of the log posterior probability is usually sufficient to detect end of the initial transient decay to the relevant model space, although more sophisticated automatic methods can also be used (e.g., Ref. 27).

Below, we describe the two types of moves used in sampling from the posterior for the transition matrix \mathbf{T} . Further details and formal proofs of correctness of this sampling procedure are given in Ref. 13.

A. Reversible element shift

Consider a pair of states (i, j) , $i \neq j$. The changed elements in the proposed transition matrix \mathbf{T}' after a move parameterized by Δ are given by

$$T'_{ij} = T_{ij} - \Delta, \quad T'_{ji} = T_{ji} - \frac{\pi_i}{\pi_j} \Delta,$$

$$T'_{ii} = T_{ii} + \Delta, \quad T'_{jj} = T_{jj} + \frac{\pi_i}{\pi_j} \Delta.$$

If \mathbf{T} fulfills the detailed balance condition (C3), \mathbf{T}' will also fulfill C3, and the stationary distribution $\boldsymbol{\pi}$ of \mathbf{T}' remains unchanged,

$$\boldsymbol{\pi}' = \boldsymbol{\pi}.$$

In order to maintain the stochasticity of \mathbf{T} (constraint C2), the parameter Δ is restricted to

$$\Delta \in \left[\max \left(-T_{ii}, -\frac{\pi_i}{\pi_j} T_{jj} \right), T_{ij} \right].$$

If Δ is chosen uniformly from this range, the proposal probabilities for the reversible element shift are given by

$$\frac{p(\mathbf{T}' \rightarrow \mathbf{T})}{p(\mathbf{T} \rightarrow \mathbf{T}')} = \sqrt{\frac{(T'_{ij})^2 + (T'_{ji})^2}{(T_{ij})^2 + (T_{ji})^2}}. \quad (20)$$

The ratio of posterior probabilities used in the Metropolis-Hastings criterion is given by

$$\frac{p(\mathbf{T}'|\mathbf{C})}{p(\mathbf{T}|\mathbf{C})} = \left(\frac{T'_{ii}}{T_{ii}} \right)^{C_{ii}} \left(\frac{T'_{ij}}{T_{ij}} \right)^{C_{ij}} \left(\frac{T'_{ji}}{T_{ji}} \right)^{C_{ji}} \left(\frac{T'_{jj}}{T_{jj}} \right)^{C_{jj}}. \quad (21)$$

B. Row shift

In a row shift move, a row i of \mathbf{T} is selected with uniform probability $1/M$, and probability mass is moved from the diagonal element T_{ii} to all outgoing probabilities T_{ij} for $j=1, \dots, M$,

$$T'_{ij} = \alpha T_{ij},$$

$$T'_{ii} = \alpha(T_{ii} - 1) + 1.$$

To maintain row-stochasticity (C2), the parameter α is drawn uniformly from the range

$$\alpha \in \left[0, \frac{1}{1 - T_{ii}} \right].$$

The ratio of proposal probabilities is given by

$$\frac{p(\mathbf{T}' \rightarrow \mathbf{T})}{p(\mathbf{T} \rightarrow \mathbf{T}')} = \alpha^{(m-2)}, \quad (22)$$

where m is the number of nonzero transition probabilities in the modified row.

The ratio of posterior probabilities is given by

$$\frac{p(\mathbf{T}' \rightarrow \mathbf{T})}{p(\mathbf{T} \rightarrow \mathbf{T}')} = \alpha^{(c_i - c_{ii})} \left(\frac{1 - \alpha(1 - T_{ii})}{T_{ii}} \right)^{c_{ii}}. \quad (23)$$

The row shift operation will change the stationary distribution of $\boldsymbol{\pi}'$, but it may be efficiently updated,

$$\pi'_i = \frac{\pi_i}{\pi_i + \alpha(1 - \pi_i)}, \quad \pi'_j = \frac{\alpha \pi_j}{\pi_i + \alpha(1 - \pi_i)}.$$

Since this update scheme is incremental, it will accumulate numerical errors over time that cause the updated $\boldsymbol{\pi}$ to drift away from the stationary distribution of the current transition matrix. To avoid this, $\boldsymbol{\pi}$ is recomputed from the current sample of the transition matrix in regular intervals (here, every 100 sampling steps).

IV. UNCERTAINTY IN MOLECULAR OBSERVABLES

While the above procedure allows us to treat the uncertainty in the transition matrix \mathbf{T} and properties computed directly from it (such as lifetimes, relaxation times, eigenvalues, and eigenvectors), it does not prescribe how we might treat the statistical error of a large class of equilibrium and temporal quantities that involve molecular observables. We therefore describe a simple extension of this Bayesian sampling scheme that allows us to include a treatment of statistical error in quantities involving these observables.

Consider the expectation of some molecular observable $A(x)$ computed for each state,

$$\langle A \rangle_i \equiv \int dx \pi_i(x) A(x). \quad (24)$$

From a single trajectory of length $L\tau$, the straightforward estimator of this quantity is simply the sample mean over those samples within state i ,

$$\hat{A}_i = \frac{\sum_{t=1}^L \phi_i(x_t) A(x_t)}{\sum_{t=1}^L \phi_i(x_t)}. \quad (25)$$

Temporally sequential samples $A_t \equiv A(x_t)$ collected with a temporal resolution of the Markov time τ_{eq} are presumed to be uncorrelated because, by definition, this is the maximum time required for the system to decorrelate within any state. We presume that the set of samples $A(x_t)$ for those configurations x_t appearing in state i are collected in the set $\{A_n\}_{n=1}^N$ in the remainder of this section, generally abbreviated as $\{A_n\}$.

Because only a finite number of samples are collected for each state, there will be a degree of uncertainty in this estimate. Unlike the problem of inferring the transition matrix elements, however, we cannot write an exact expression for the probability of observing the A_n in terms of a simple parametric form, since its probability distribution may be arbitrarily complex,

$$p_i(A_n) = \int dx \delta(A_n - A(x)) \pi_i(x). \quad (26)$$

Despite this, the central limit theorem states that the behavior of \hat{A}_i approaches a normal distribution (generally very rap-

idly) as the number of samples N increases. We will therefore make the assumption that $p_i(A_n)$ is *normal*—that is, that

$$p_i(A_n; \mu_i, \sigma_i^2) = (2\pi)^{-1/2} \sigma_i^{-1} \exp\left[-\frac{1}{2\sigma_i^2}(A_n - \mu_i)^2\right], \quad (27)$$

and demonstrate that this allows us to do a very good job of inferring the distribution of the error in $\delta\hat{A}_i \equiv \hat{A}_i - \langle A \rangle_i$ for a very reasonable number of samples and generally gives an overestimate of the error (which is arguably less dangerous than an underestimate) for smaller sample sizes.

Consider the sample mean estimator for $\langle A \rangle_i$,

$$\hat{\mu} = \frac{1}{N} \sum_{n=1}^N A_n. \quad (28)$$

The asymptotic variance of $\hat{\mu}$, which provides a good estimate of the statistical uncertainty in $\hat{\mu}$ in the large-sample limit, is given as a simple consequence of the central limit theorem,

$$\delta^2 \hat{\mu} \equiv E[(\hat{\mu} - E[\hat{\mu}])^2] = \frac{\text{var } A_n}{N} \approx \frac{\hat{\sigma}^2}{N}, \quad (29)$$

where the unbiased estimator for the variance $\sigma^2 \equiv \text{var } A_n$ is given by

$$\hat{\sigma}^2 \equiv \frac{1}{N-1} \sum_{n=1}^N (A_n - \hat{\mu})^2. \quad (30)$$

Suppose we now *assume* that the distribution of A from state i is normal [described by Eq. (27)],

$$A|\mu, \sigma^2 \sim N(\mu, \sigma^2). \quad (31)$$

Were this to be a reasonable model, we could model the timeseries of the observable $A_t \equiv A(x_t)$ by the hierarchical process,

$$s_t | s_{t-1}, \mathbf{T} \sim \text{BS}(T_{s_{t-1}}, \dots, T_{s_{t-1}N}), \quad (32)$$

$$A_t | \mu_{s_t}, \sigma_{s_t}^2 \sim N(\mu_{s_t}, \sigma_{s_t}^2).$$

Here, the notation $\text{BS}(\pi_1, \dots, \pi_N)$ denotes a Bernoulli scheme where discrete outcome n has associated probability π_n and $N(\mu, \sigma^2)$ denotes the normal distribution with mean μ and variance σ^2 . We will demonstrate below how this model does, in fact, recapitulate the expected behavior in the limit where there are sufficient samples from each state.

We choose the (improper) Jeffreys prior,²⁸

$$p(\mu, \sigma^2) \propto \sigma^{-2}, \quad (33)$$

because it satisfies intuitively reasonable reparametrization²⁸ and information-theoretic²⁹ invariance principles. Note that this prior is uniform in $(\mu, \log \sigma)$.

The posterior is then given by




FIG. 1. Approach to normality for marginal distribution of the mean $p(\mu|\{A_n\})$. For fixed $\hat{\mu}$ and $\hat{\sigma}^2$, the marginal posterior distribution of μ (red), a scaled and shifted Student t-distribution, rapidly approaches the normal distribution (black) expected from asymptotic statistics. The PDF is shown for sample sizes of $N=5$ (the broadest), 10, 20, and 30.

$$\begin{aligned} p(\mu, \sigma^2 | \{A_n\}) &\propto \left[\prod_{n=1}^N p(A_n | \mu, \sigma^2) \right] p(\mu, \sigma^2) \\ &\propto \sigma^{-(N+2)} \exp\left[-\frac{1}{2\sigma^2} \sum_{n=1}^N (A_n - \mu)^2\right]. \end{aligned} \quad (34)$$

Rewriting in terms of the sample statistics $\hat{\mu}$ and $\hat{\sigma}^2$, we obtain

$$\begin{aligned} p(\mu, \sigma^2 | \{A_n\}) &\propto \sigma^{-(N+2)} \exp\left\{-\frac{1}{2\sigma^2} \left[\sum_{n=1}^N (A_n - \hat{\mu})^2 + N(\hat{\mu} - \mu)^2 \right]\right\} \\ &\propto \sigma^{-(N+2)} \exp\left\{-\frac{1}{2\sigma^2} [(N-1)\hat{\sigma}^2 + N(\hat{\mu} - \mu)^2]\right\}. \end{aligned} \quad (35)$$

The posterior has marginal distributions,

$$\sigma^2 | \{A_n\} \sim \text{Inv} - \chi^2(N-1, \hat{\sigma}^2), \quad (36)$$

$$\mu | \{A_n\} \sim t_{N-1}(\hat{\mu}, \hat{\sigma}^2/N),$$

where σ^2 is distributed according to scaled inverse chi-square distribution with $N-1$ degrees of freedom and μ according to Student's t-distribution with $N-1$ degrees of freedom that has been shifted to be centered about $\hat{\mu}$ and whose width has been scaled by $\hat{\sigma}^2/N$.

As can be seen in Fig. 1, as the number of degrees of freedom increases, the marginal posterior for μ approaches the normal distribution with the asymptotic behavior expected from standard frequentist analysis for the standard error of the mean, namely,

$$\mu \rightarrow N(\hat{\mu}, \hat{\sigma}^2/N). \quad (37)$$

At low sample counts, the t-distribution is lower and wider than the normal distribution, meaning that confidence intervals computed from this distribution will be somewhat larger than those of the corresponding normal estimate for small samples. In some sense, this partly compensates for $\hat{\sigma}^2$ being

a poor estimate of the true variance for small sample sizes, which would naturally lead to underestimates of the statistical uncertainty. In any case, this is also far from the asymptotic limit where the normal distribution with variance $\hat{\sigma}^2/N$ is expected to model the uncertainty well.

The posterior can also be decomposed as

$$p(\mu, \sigma^2 | \{A_n\}) = p(\mu | \sigma^2, \{A_n\}) p(\sigma^2 | \{A_n\}). \quad (38)$$

This readily suggests a two-step sampling scheme for generating uncorrelated samples of (μ, σ^2) , in which we first sample σ^2 from its marginal distribution, and then μ from its distribution conditional on σ^2 ,

$$\sigma^2 | \{A_n\} \sim \text{Inv} - \chi^2(N-1, \hat{\sigma}^2), \quad (39)$$

$$\mu | \sigma^2, \{A_n\} \sim N(\hat{\mu}, \sigma^2/N).$$

Alternatively, if the scaled inverse-chi-square distribution is not available, the χ^2 -distribution (among others) can be used to sample σ^2 ,

$$(N-1)(\hat{\sigma}^2/\sigma^2) | \{A_n\} \sim \chi^2(N-1), \quad (40)$$

where the first argument is the shape parameter and the second argument is the scale parameter.

V. VALIDATION IN A MODEL SYSTEM

Using the procedures described in Secs. III and IV, we are now equipped with a scheme to sample from the joint posterior describing our confidence in that a Markov model characterized by a transition matrix \mathbf{T} and state expectations μ_i , $i=1, \dots, M$, produced the observed trajectory data. Using a set of models sampled from this posterior, we can characterize the statistical component of the uncertainty as it propagates into equilibrium averages, nonequilibrium relaxations, and (non)equilibrium correlation measurements computed from the Markov model. To ensure the correctness of this procedure, however, we first test its ability to correctly characterize the posterior distribution for a finite-size sample from a true Markovian model system.

How can we test a Bayesian posterior distribution? One of the more powerful features of a Bayesian model is its ability to provide confidence intervals that correctly reflect the level of certainty that the true value will lie within it. For example, if the experiment were to be repeated many times, the true value of the parameter being estimated should fall within the confidence interval for a 95% confidence level 95% of the time. As an illustrative example, consider a biased coin where the probability of turning heads is θ . From an observed sample of N coin flips, we can estimate θ using a Binomial model for the number of coin flips that turn up heads and a conjugate beta Jeffreys prior.^{28,29} Each time we run experiment and generate a new independent collection of N samples, we get a different posterior estimate for θ and a different confidence interval (Fig. 2, top). If we run many trials and record what fraction of the time the true (unknown) value of θ falls within the confidence interval estimated from that trial, we can see if our model is correct. If correct, the observed confidence level should match the desired confidence level (Fig. 2, bottom right). Deviation from parity




FIG. 2. Testing the posterior for inference of a biased coin flip experiment. Top: Posterior distribution for inferring the probability of heads, θ , for a biased coin from a sequence of $N=1000$ coin flips (dark line) with 95% symmetric confidence interval about the mean (shaded area). The true probability of heads is 0.3 (vertical thick line). Posteriors from five different experiments are shown as dotted lines. Bottom left: Desired and actual confidence levels for an idealized normal posterior distribution that either overestimates (upper left curves) or underestimates (bottom right curves) the true posterior variance by different degrees. Bottom right: Desired and actual confidence levels for the binomial-beta posterior for the coin flip problem depicted in upper panel. Error bars show 95% confidence intervals estimates from 1000 independent experimental trials. For inference, we use a likelihood function such that the observed number of heads is $N_H | \theta \sim \text{Binomial}(N_H, N, \theta)$ and conjugate Jeffreys prior (Refs. 28 and 29) $\theta \sim \text{Beta}(1/2, 1/2)$ which produces posterior $\theta | N_H \sim \text{beta}(N_H + 1/2, N_T + 1/2)$ along with constraint $N_H + N_T = N$.

means that the posterior is either too broad or too narrow, and that the statistical uncertainty is being either over- or underestimated (Fig. 2, bottom left).

We performed a similar test on a three-state model system, using a model (reversible, row-stochastic) transition matrix for one Markov time is given by

$$\mathbf{T}(1) = \begin{bmatrix} 0.86207 & 0.12931 & 0.00862 \\ 0.15625 & 0.83333 & 0.01041 \\ 0.00199 & 0.00199 & 0.99602 \end{bmatrix}. \quad (41)$$

Each state is characterized by a mean value of the observable $A(x)$, fixed to 3, 2, and 1 for the first, second, and third states, respectively. The equilibrium populations are $\pi \approx [0.1625 \ 0.1345 \ 0.7031]$. Simulation from this model involves a stochastic transition according to the transition element T_{ij} followed by observation of the value of $A(x)$ sampled independent and identically distributed from the current state's probability distribution $p_i(A)$. Multiple independent realizations of this process were carried out and subjected to the Bayesian inference procedure for transition matrices and observables described in Secs. III and IV. The nonequilibrium relaxation $\langle A \rangle_{\rho_0}$ from the initial condition




FIG. 3. Observables for three-state model system. Top: Relaxation of $\langle A(t) \rangle_{\rho_0}$ (solid line) from initial distribution $\rho_0 = [1 \ 0 \ 0]$ to equilibrium expectation $\langle A \rangle$ (dash-dotted line). Bottom: Equilibrium autocorrelation function $\langle A(0)A(t) \rangle$ (solid line) to $\langle A \rangle^2$ (dash-dotted line). The estimates of both $\langle A(t) \rangle_{\rho_0}$ and $\langle A(0)A(t) \rangle$ at 50 timesteps (red vertical line) were assessed in the validation tests described here.

$\rho_0 = [100]$ in which all density is concentrated in state 1, as well as the autocorrelation function $\langle A(0)A(t) \rangle$, is shown in Fig. 3.

With the means of $p_i(A)$ within each state fixed as above, we considered models for $p_i(A)$ that were either normal or exponential, using the probability density functions,

$$p_i(A) = (2\pi)^{-1/2} \sigma_i^{-1} \exp\left[-\frac{1}{2\sigma_i^2}(A - \mu_i)^2\right] \quad (\text{normal}),$$

$$p_i(A) = \mu_i^{-1} \exp[-A/\mu_i], \quad A \geq 0 \quad (\text{exponential}).$$

While the normal output distribution for $p_i(A)$ corresponds to the hierarchical Bayesian model that forms the basis for our

approach, the exponential distribution is significantly different and represents a challenging test case.

Figure 4 depicts the resulting uncertainty estimates for both normal (top) and exponential (bottom) densities for the observable A . In both cases, the confidence intervals are *underestimated* for short trajectory lengths (1000 steps) where, in many realizations, few samples are observed in one or more states, so that the variance is underestimated or the effective asymptotic limit has not yet been reached. As the simulation length is increased to 10 000 or 100 000 steps so that it is much more likely there are a sufficient number of samples in each state to reach the asymptotic limit, however, the confidence intervals predicted by the Bayesian posterior become quite good. For the exponential model for observing values of A (which might be the case in, say, fluorescence lifetimes), we observe similar behavior. Except for what appears to be a slight, consistent underestimation of $\langle A(t) \rangle_{\rho_0}$ (much less than half a standard deviation), there appears to be excellent agreement between the expected and observed confidence intervals, confirming that this method is expected to be a useful approach to modeling statistical uncertainties in equilibrium and kinetic observables.

VI. APPLICATION TO FLUORESCENCE CORRELATION IN A PEPTIDE

Time-resolved single-molecule fluorescence experiments provide a way to monitor the microscopic features of folding by probing the fluctuations of a conformation-sensitive spectroscopic probe for individual molecules. In contrast to ensemble measurements, which yield information only on average properties, single-molecule experiments provide information on distributions and time trajectories of properties that would otherwise be hidden.³⁰ Due to recent ad-




FIG. 4. Confidence interval tests for model system. Top: Expected and observed confidence intervals for three-state system with normal distribution for observable A with unit variance for simulations of length 1000 (left), 10 000 (middle), and 100 000 (right) steps. Confidence intervals were estimated from generating 10 000 samples from the Bayesian posterior. Estimates of the fraction of observed times the true value was within the confidence interval estimated from the Bayesian posterior were computed from generating 1000 independent experimental realizations. The resulting curves are shown for the equilibrium estimate $\langle A \rangle$ (red), nonequilibrium relaxation $\langle A \rangle_{\rho_0}$ (green), and the equilibrium correlation function $\langle A(0)A(t) \rangle$ (blue). Bottom: Same as top, except an exponential distribution with the same mean was used for the probability of observing a particular value of A within each state.




FIG. 5. Fluorescent peptide MR121-GSGS-W. The large fluorescent dye molecule, MR121, is the fused multiring structure visible at the N-terminus on the left of the figure.

vances in instrumentation, single-molecule fluorescence spectroscopy has received a particular surge of interest. In contrast to traditional fluorescence experiments, nonequilibrium perturbations to synchronize an ensemble of molecules are avoided, allowing equilibrium conformational dynamics to be studied.^{30–32} FCS takes advantage of Brownian diffusion to bring one or a few fluorescent molecules into the observation volume at any given time, recording bursts of fluorescent emission. By calculating the autocorrelation function of the fluorescence intensity fluctuations, information about the temporal statistics of dynamics can be obtained for a wide range of time scales spanning nanoseconds to seconds.³² Because the fluorescence signal is only an indirect probe of molecular conformation, molecular dynamics simulations (typically limited to nanosecond to microsecond timescales) have proven to be of great utility in the interpretation of fluorescence autocorrelation experiments.³³ By employing Markov models to describe the long-time dynamics, it is possible to extend the utility of molecular dynamics simulation for predicting and interpreting these experiments into the microsecond regime and beyond.²⁶

Here, we illustrate the use of the Bayesian uncertainty analysis procedure described above in interpreting an experimental study of the loop-closure dynamics of a small fluorescently labeled peptide using a Markov model constructed from molecular dynamics simulation data. To probe the fastest processes in protein folding, Krieger *et al.*³⁴ and Neuwiler *et al.*³⁵ collected single-molecule fluorescence data on a series of small peptides incorporating fluorophores and quenchers. In the experiment considered here, end-to-end contact-formation is studied by monitoring selective fluorescence quenching of an N-terminal fluorescent oxazine derivative MR121 by a C-terminal tryptophan residue, an efficient natural amino acid quencher, with an intervening Gly-Ser-Gly-Ser repeat (hereafter called MR121-GSGS-W, see Fig. 5). To first order, when the tryptophan is sufficiently close to the MR121 dye, fluorescence is quenched (resulting in an “off” state); conversely, when the tryptophan is far from the dye and the peptide is in a more extended conformation, the dye is fluorescent (resulting in an “on” state). The quenching process has been shown to be diffusion limited,^{35,36} enabling the underlying contact-formation kinetics to be probed by FCS with nanosecond time-resolution.^{34,35}

A. Simulation details and Markov model construction

Here, a 3 μ s molecular dynamics simulation of MR121-GSGS-W is analyzed. The simulations were performed in explicit water at 293 K using the GROMOS96 force field 43a1 (Refs. 37 and 38) and the GROMACS program version 3.2.1.^{39,40} Partial atomic charges for the dye MR121 were taken from Vaiana *et al.*³⁶ One peptide molecule in an extended conformation was solvated with SPC water⁴¹ and placed in a periodic rhombic dodecahedron box large enough to contain the peptide molecule and ≈ 1.0 nm of solvent on all sides at a liquid density of 55.32 mol/l (≈ 1 g/cm³), resulting in 1155 water molecules in the simulation box. The simulation volume was held fixed, and thermal control was enforced (note that the Berendsen weak-coupling algorithm does not generate a true NVT ensemble and only produces statistics equivalent to a NVT ensemble in the thermodynamic limit^{42–44}) using the Berendsen weak-coupling algorithm⁴⁵ with a coupling time of 0.1 ps. All bond lengths were fixed using the LINCS algorithm of order of 4 with a tolerance of 10^{-4} nm (Ref. 46) and a time step of 2 fs for numerical integration was used. Periodic boundary conditions were applied to the simulation box and the long-range electrostatic interactions were treated with the particle mesh Ewald method⁴⁷ using a grid spacing of 0.12 nm combined with a fourth-order B-spline interpolation to compute the potential and forces in between grid points and an Ewald tolerance of 10^{-5} . The real space cutoff distance was set to 0.9 nm. The C-terminal end of the peptide was modeled as COO⁻ to reproduce a pH of about 7 as in the experimental conditions.³⁵ No counterions were added since the simulation box was already neutral (one positive charge on MR121 and one negative charge on the terminal COO⁻). The coordinates were saved every $\Delta t=0.2$ ps.

As a crude model of fluorescence quenching, all configurations in which the heavy atoms of the rings systems of MR121 and the tryptophan had a nearest-neighbor distance of greater than 0.45 nm were defined to be fluorescent, and the remainder as dark (quenched). A cutoff of 0.45 nm was selected since quenching occurs upon van der Waals contact.³⁵ Thus, a fluorescence observable can be defined by the fluorescence observable $f(x)$,

$$f(x) = \begin{cases} 0 & x \text{ dark} \\ 1 & x \text{ fluorescent,} \end{cases} \quad (42)$$

while the fluorescence of state i is given by

$$f_i = \int dx \pi_i(x) f(x), \quad (43)$$

which, of course, must be estimated from the simulation data collected within state i . [Note that, for this particular case, the observable $f(x)$ could be modeled *exactly* by a Bernoulli trial where the state fluorescence f_i are the (unknown) parameters. However, due to the broad applicability of the normal model to many types of observables, we do not examine the Bernoulli model here. In any case, any significant difference between the two models should vanish as the number of samples reaches the asymptotic limit.]

From the simulation trajectory, 1000 configurations

equally spaced in time were used to define the *generators* of an initial partitioning of configuration space into Voronoi polytopes (as in the microstates of Ref. 7), where each structure in the trajectory is assigned to the generator it is closest to using the least-squares-aligned root mean square deviation (RMSD)^{48,49} as a metric⁵⁰ using the efficient RMSD calculation procedure of Theobald.⁵¹ A lag time of $\tau=1$ ns was then used, as examination of the implied timescales⁷ showed they become approximately independent of lag time for larger values of τ . In order to obtain reasonable estimates of the per-state fluorescence f_i and transition probabilities T_{ij} , the state space was further coarse-grained by computing the right eigenvectors of the 1000×1000 row-stochastic transition matrix and lumping states that had a Euclidean distance of less than 0.05 in the space of the 20 dominant eigenvectors which were normalized to have unit 2-norm. This generated a reduced state space of 163 states, and it was verified that the implied timescales did not significantly deteriorate.

The unnormalized fluorescence autocorrelation function is given in terms of f_i and T_{ij} as

$$\langle f(0)f(n\tau) \rangle = \sum_{i=1}^M \sum_{j=1}^M \pi_i f_i ([T(\tau)^n]_{ij}) f_j, \quad (44)$$

which corresponds to the part of the fluorescence autocorrelation curve that is measured by FCS experiments. Note that the measured FCS curves will also contain additional contributions due to effects such as triplet states and diffusion which are not considered in our analysis.

The normalized form of the fluorescence autocorrelation function, which decays from an initial value of unity to zero at large t , is given by

$$C(t) \equiv \frac{\langle f(0)f(n\tau) \rangle - \langle f \rangle^2}{\langle f^2 \rangle - \langle f \rangle^2} = \frac{\sum_{i=1}^M \sum_{j=1}^M \pi_i f_i ([T(\tau)^n]_{ij}) f_j - \sum_{i=1}^M (\pi_i f_i)^2}{\sum_{i=1}^M \pi_i f_i^2 (1 - \pi_i)}. \quad (45)$$

B. Uncertainties in model-derived observables

To determine how the uncertainties in the computed unnormalized and normalized fluorescence correlation functions depend on the quantity of simulation data collected, the initial portions of the clustered molecular dynamics trajectory consisting of 0.1, 0.25, 0.5, 1, 2, and 3 μ s were used to construct the count matrix \mathbf{C} and number of fluorescent samples per state f_i . The sampling procedure described above was used to sample from the model posterior, discarding 10^6 initial samples for burn-in. Due to the local nature of the Monte Carlo transition matrix moves detailed in Sec. III, subsequent transition matrix samples are generally strongly correlated. While it is asymptotically correct to use all such samples to estimate expectations and uncertainties, computing correlation functions for strongly correlated transition matrix samples is in practice inefficient. Here, 10^5 sampling steps were conducted in between each transition matrix used to estimate the fluorescence correlation function, and 10^3 fluorescence correlation functions were sampled in this way

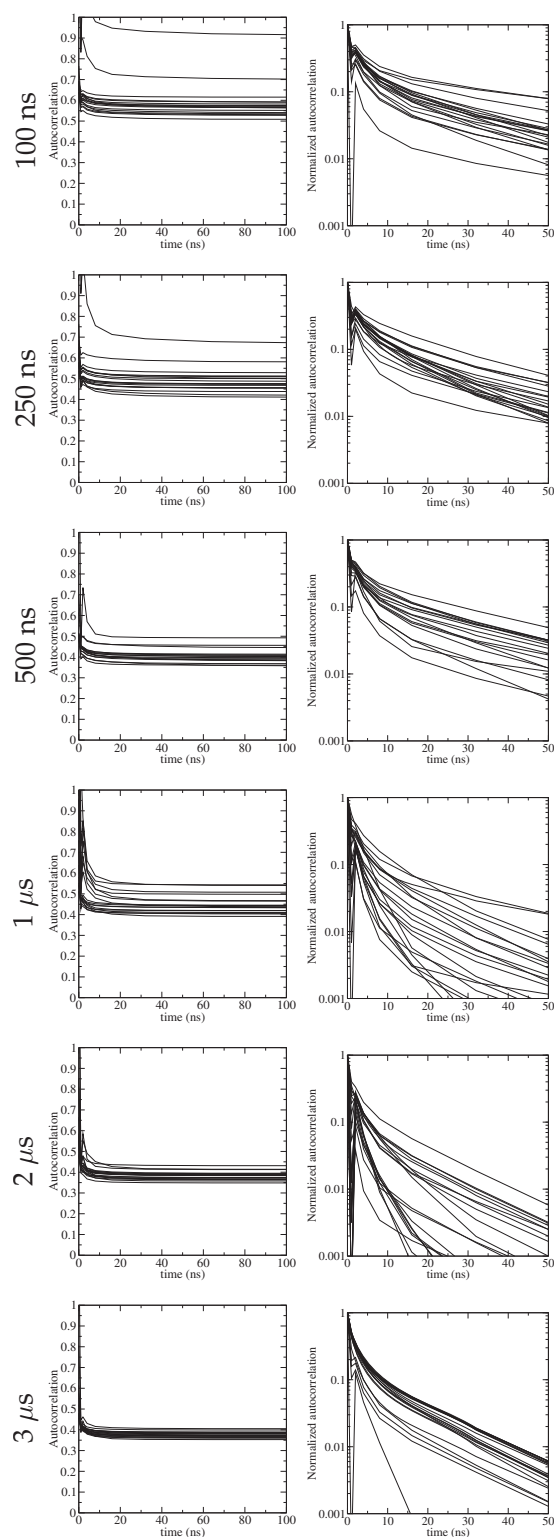


FIG. 6. Sampled fluorescence autocorrelation functions for fluorescent peptide for different lag times. Twenty posterior samples of the computed fluorescence autocorrelation function are shown to indicate the uncertainty in the autocorrelation. Left: Unnormalized autocorrelation function on a linear scale. Right: Normalized autocorrelation function on a logarithmic scale, to better illustrate decay components.

(corresponding to 10^8 sampling steps in total). This procedure required 2–3 min on a standard 2.4 GHz Intel CPU core in total.

Figure 6 shows a sample of 20 unnormalized (left) and

normalized (right) fluorescence correlation functions computed from the Bayesian model posterior. The most obvious source of uncertainty in the unnormalized correlation functions is the uncertainty in the stationary distribution π , which results in a visually apparent dispersion in the plateau values. As expected, the variance in the plateau values becomes smaller for longer simulation times, especially when the simulation times exceed $2 \mu\text{s}$. The uncertainty in the rate of decay, however, appears to be rather similar despite this. This is more easily seen in the normalized fluorescence correlation functions (Fig. 6, right, shown in log-scale) which has the effect of removing the effects of differences in asymptotic value. The normalized autocorrelation is plotted on a log-scale, such that a single-exponential relaxation would appear as a straight line. All curves show an asymptotic single-exponential relaxation with a timescale and rate determined by the slowest process implied by the corresponding transition matrix. At shorter times of about 10–20 ns, however, all curves appear to be multiexponential. While multiexponential behavior was not seen in Ref. 35, but this timescale is close to the resolution limit of this experiment. For all simulation times, the rate of decay of the slow components appears to be similar, although the amplitude (i.e., the relative amplitude of slow and fast components) varies. Note that the visually apparent increase in variance at larger at long simulation times ($\geq 1 \mu\text{s}$) is simply a consequence of the logarithmic scale used for plotting the normalized autocorrelation function. However, it seems that the short simulation times, the slowest processes are estimated to be slower than at long simulation times.

The autocorrelation function $\langle A(0)A(n\tau) \rangle$ [Eq. (3)] can be written in terms of the left eigenvectors \mathbf{l}_i , right eigenvectors \mathbf{r}_i , and eigenvalues λ_i of $\mathbf{T}(\tau)$,

$$\begin{aligned} \langle A(0)A(n\tau) \rangle &= \sum_{i,j=1}^M A_i([\mathbf{T}(\tau)]^n)_{ij} A_j = \sum_{i,j} A_i \left(\sum_k r_{ki} \lambda_k^n l_{kj} \right) A_j \\ &= \sum_k \left[\left(\sum_i A_i r_{ki} \right) \left(\sum_j A_j l_{kj} \right) \right] \lambda_k^n \equiv \sum_k a_k \lambda_k^n = 1 \\ &= \sum_{k=2}^M a_k e^{-n\tau/t_k}. \end{aligned} \quad (46)$$

The timescales t_i associated with the two slowest processes were computed by

$$t_i = -\frac{\tau}{\log \lambda_i}, \quad i = 2, \dots, K, \quad (47)$$

where $1 = \lambda_1 > \lambda_2 > \dots > \lambda_K$ are the ordered real eigenvalues of $\mathbf{T}(\tau)$. The relative amplitudes of the two processes in the fluorescence correlation function,

$$a'_i = \frac{a_i}{a_2 + a_3} \quad (48)$$

with $i=2,3$, were computed as well. The amplitudes a'_2, a'_3 are shown in Fig. 7. At simulation times of $1 \mu\text{s}$ or less, the relative amplitudes of the slowest process is somewhat smaller than that of the second-slowest process. This changes at 2 or $3 \mu\text{s}$, where the slowest process is revealed as the




FIG. 7. FCS relative intensity and timescale of the slowest (black) and the second-slowest (red) process.

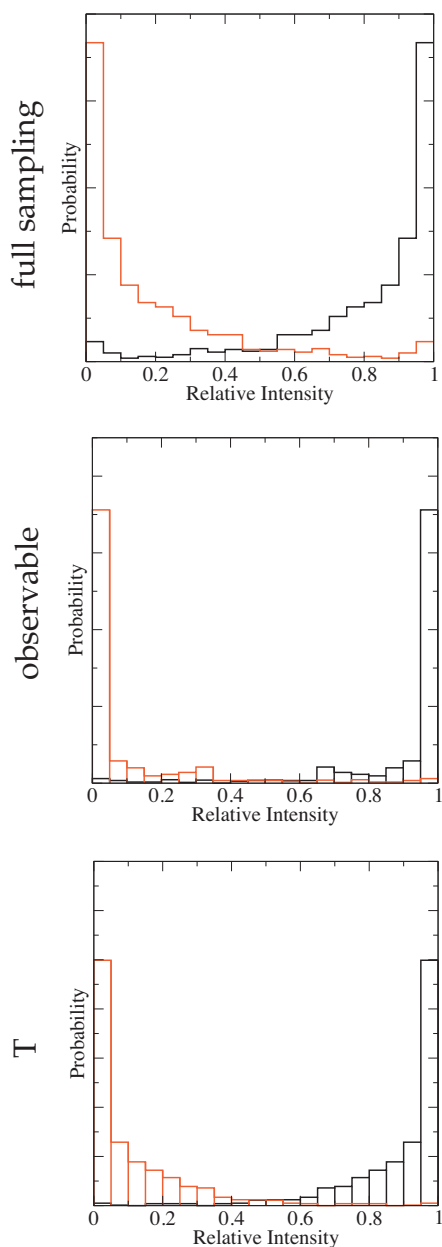


FIG. 8. Effect of sampling only in observable or transition matrix space on the FCS relative intensity and timescale of the slowest (black) and the second-slowest (red) process using 500 ns of trajectory data.

one with the largest amplitude. As seen in the previous plot, the timescales, in particular, the slowest timescale, are overestimated at short simulation times, while their uncertainty, being a factor of 2–3 at 100 ns, rapidly decreases. At 2 and 3 μ s, the timescales appear well-determined, and are around 15 and 20 ns. Note that although these timescales are above the resolution limit of the experiment in Ref. 35, in the presence of noise, a signal consisting of two exponentials with similar timescales can generally not be distinguished from a single-exponential.

Figure 8 illustrates the relative contributions to the overall uncertainty estimate from the uncertainty in the observable expectations within states A_i and the uncertainty in the transition matrix $\mathbf{T}(\tau)$. To assess this, the sampler is either (1) fixed at the maximum likelihood reversible transition matrix and the state-dependent expectation values of the fluo-

rescence are sampled or (2) the expectation values of the fluorescence observable are fixed at their sample means while the transition matrix is sampled. Only the differences in the amplitudes of the fluorescence correlation components are depicted, as the timescales are only affected by the transition matrix uncertainties. While there is some nontrivial contribution to the uncertainty in the relative intensities of the slowest and second-slowest processes due to the uncertainty in the state observable expectations, it is clear that, in this case, the dominant contribution is due to the uncertainty in the transition matrix.

VII. DISCUSSION

We have demonstrated how a normal model for the probability of generating observations from each state in a Markov model, even though approximate, can give useful estimates for the uncertainties in expectations computed for this state. Furthermore, this procedure provides an algorithmically simple and convenient means for propagating these uncertainties into the uncertainties in more complicated quantities (such as correlation functions and nonequilibrium relaxation spectra) predicted from the model. As the number of samples observed within each state approaches the asymptotic limit (which need not be more than a few tens of uncorrelated samples), the uncertainties computed from the normal model approach those expected from asymptotic theory. While in particular situations, a different (non-normal) model may be applicable to modeling the distribution of observables from each state given unknown parameters, the generality and simplicity of the normal model presented here are expected to be of broadest utility.

In conjunction with transition matrix sampling to incorporate the statistical uncertainties of the transition probabilities between Markov model states, the posterior distribution of virtually any observable accessible with experiment can be computed, provided two conditions are met. (1) The Markov time τ for which the model accurately reproduces dynamics is shorter than the timescales of the process of interest, and (2) it is possible to compute a *surrogate* for the spectroscopic signal as a function of conformation to an accuracy greater than the uncertainty of the experimental measurement. Even if the accuracy of the surrogate for the spectroscopic signal is not high, it is possible that the *timescales* present in the experimental signal may still be compared.

In the present simulations it was found that the uncertainty of the transition matrix is the main contribution to uncertainties of derived molecular observables. This is first due to the fact that the transition matrix uncertainty is exclusively responsible for uncertainties in the stationary distribution that affects stationary properties and also in the relaxation timescales that affect the decay of correlation or relaxation functions. Moreover, for the present example of fluorescence correlation spectroscopy, the transition matrix uncertainty was also the main contributor to uncertainties in the amplitudes of the fluorescence correlation components. Nevertheless, it is conceivable that this relative contribution is reversed in other applications that have only few samples of

the observable per state or in the case of observables that depend on the functions of state in a highly nonlinear fashion.

Another advantage conferred by having a fully *generative* model—in which complete realizations of the experiment can be produced artificially—is the possibility for predicting how additional data collection schemes will most rapidly reduce the uncertainty in the quantities desired from the model. For example, if a particular nonequilibrium relaxation spectrum is of interest, Bayesian experimental design techniques⁵² could be employed in order to determine how new simulations can be initiated in a way expected to most rapidly reduce the uncertainty in this quantity. Similar approaches have already been described by Singhal and Pande, for example, in designing experiments to optimally reduce variances in properties computed from the transition matrix alone.^{19,20}

ACKNOWLEDGMENTS

The authors thank Vijay S. Pande and Sergio Bacallado (Stanford), as well as the two anonymous referees, for helpful feedback on the manuscript. J.D.C. acknowledges support from a distinguished postdoctoral fellowship from the California Institute for Quantitative Biosciences (QB3) at the University of California, Berkeley. F.N. acknowledges support from DFG Research Center Matheon.

- ¹C. Schütte, A. Fischer, W. Huisinga, and P. Deuffhard, *J. Comput. Phys.* **151**, 146 (1999).
- ²C. Schütte and W. Huisinga, in *Handbook of Numerical Analysis*, edited by P. G. Ciarlet and J.-L. Lions (Elsevier, New York, 2002), Vol. X.
- ³F. Noé and S. Fischer, *Curr. Opin. Struct. Biol.* **18**, 154 (2008).
- ⁴M. Sarich, F. Noé, and C. Schütte, *Multiscale Model. Simul.* **8**, 1154 (2010), <http://proteomics-berlin.de/771/>.
- ⁵W. C. Swope, J. W. Pitera, and F. Suits, *J. Phys. Chem. B* **108**, 6571 (2004).
- ⁶J. D. Chodera, W. C. Swope, J. W. Pitera, and K. A. Dill, *Multiscale Model. Simul.* **5**, 1214 (2006).
- ⁷J. D. Chodera, N. Singhal, V. S. Pande, K. A. Dill, and W. C. Swope, *J. Chem. Phys.* **126**, 155101 (2007).
- ⁸F. Noé, I. Horenko, C. Schütte, and J. C. Smith, *J. Chem. Phys.* **126**, 155102 (2007).
- ⁹V. Schultheis, T. Hirschberger, H. Carstens, and P. Tavan, *J. Chem. Theory Comput.* **1**, 515 (2005).
- ¹⁰S. Park and V. S. Pande, *J. Chem. Phys.* **124**, 054118 (2006).
- ¹¹N.-V. Buchete and G. Hummer, *J. Phys. Chem. B* **112**, 6057 (2008).
- ¹²S. Bacallado, J. D. Chodera, and V. Pande, *J. Chem. Phys.* **131**, 045106 (2009).
- ¹³F. Noé, *J. Chem. Phys.* **128**, 244103 (2008).
- ¹⁴M. Weber, Ph.D. thesis, Freie Universität Berlin—Fachbereich Mathematik und Informatik, 2006, <http://www.zib.de/weber/Promotion.pdf>.
- ¹⁵D. Chandler, *J. Chem. Phys.* **68**, 2959 (1978).
- ¹⁶J. E. Adams and J. D. Doll, *Surf. Sci.* **111**, 492 (1981).
- ¹⁷A. F. Voter and J. D. Doll, *J. Chem. Phys.* **82**, 80 (1985).
- ¹⁸G. Hummer, *New J. Phys.* **7**, 34 (2005).
- ¹⁹N. Singhal and V. S. Pande, *J. Chem. Phys.* **123**, 204909 (2005).
- ²⁰N. S. Hinrichs and V. S. Pande, *J. Chem. Phys.* **126**, 244101 (2007).
- ²¹F. Noé, M. Oswald, G. Reinelt, S. Fischer, and J. C. Smith *Multiscale Model. Simul.* **5**, 393 (2006).
- ²²F. Noé, M. Oswald, and G. Reinelt, in *Operations Research Proceedings*, edited by J. Kalcsics and S. Nickel (Springer, New York, 2007), pp. 435–440.
- ²³N. G. van Kampen, *Stochastic Processes in Physics and Chemistry*, 2nd ed. (Elsevier, New York, 1997).
- ²⁴P. Metzner, F. Noé, and C. Schütte, *Phys. Rev. E* **80**, 021106 (2009).
- ²⁵S. Sriraman, I. G. Kevrekidis, and G. Hummer, *J. Phys. Chem. B* **109**, 6479 (2005).
- ²⁶F. Noé, C. Schütte, E. Vanden-Eijnden, L. Reich, and T. R. Weikl, *Proc. Natl. Acad. Sci. U.S.A.* **106**, 19011 (2009).
- ²⁷A. Gelman and D. B. Rubin, *Stat. Sci.* **7**, 457 (1992).
- ²⁸H. Jeffreys, *Proc. R. Soc. London, Ser. A* **186**, 453 (1946).
- ²⁹P. Goyal, in *Bayesian Inference and Maximum Entropy Methods in Science and Engineering*, edited by K. H. Knuth, A. E. Abbas, R. D. Morris, and J. P. Castle (American Institute of Physics, New York, 2005), pp. 366–373.
- ³⁰X. Michalet, S. Weiss, and M. Jäger, *Chem. Rev. (Washington, D.C.)* **106**, 1785 (2006).
- ³¹B. Schuler, *ChemPhysChem* **6**, 1206 (2005).
- ³²P. Tinnefeld and M. Sauer, *Angew. Chem., Int. Ed.* **44**, 2642 (2005).
- ³³I.-C. Yeh and G. Hummer, *J. Am. Chem. Soc.* **124**, 6563 (2002).
- ³⁴F. Krieger, B. Fierz, O. Bieri, M. Drewello, and T. Kiefhaber, *J. Mol. Biol.* **332**, 265 (2003).
- ³⁵H. Neuweiler, M. Löllmann, S. Doose, and M. Sauer, *J. Mol. Biol.* **365**, 856 (2007).
- ³⁶A. C. Vaiana, H. Neuweiler, A. Schulz, J. Wolfrum, M. Sauer, and J. C. Smith, *J. Am. Chem. Soc.* **125**, 14564 (2003).
- ³⁷W. F. van Gunsteren, S. R. Billeter, A. A. Eising, P. H. Hünenberger, P. Krueger, A. E. Mark, W. R. P. Scott, and I. G. Tironi, *Biomolecular Simulation: The GROMOS96 Manual and User Guide* (Vdf Hochschulverlag, AG Zurich, 1996).
- ³⁸W. R. P. Scott, P. H. Hünenberger, I. G. Tironi, A. E. Mark, S. R. Billeter, J. Fennen, A. E. Torda, T. Huber, P. Kruger, and W. F. van Gunsteren, *J. Phys. Chem. A* **103**, 3596 (1999).
- ³⁹H. J. C. Berendsen, D. van der Spoel, and R. van Druren, *Comput. Phys. Commun.* **91**, 43 (1995).
- ⁴⁰E. Lindahl, B. Hess, and D. van der Spoel, *J. Mol. Model.* **7**, 306 (2001).
- ⁴¹H. J. C. Berendsen, J. R. Grigera, and T. P. Straatsma, *J. Phys. Chem.* **91**, 6269 (1987).
- ⁴²T. Morishita, *J. Chem. Phys.* **113**, 2976 (2000).
- ⁴³M. D'Alessandro, A. Tenenbaum, and A. Amadei, *J. Phys. Chem. B* **106**, 5050 (2002).
- ⁴⁴A. Mudi and C. Chakravarty, *Mol. Phys.* **102**, 681 (2004).
- ⁴⁵H. J. C. Berendsen, J. P. M. Postma, W. F. van Gunsteren, A. DiNola, and J. R. Haak, *J. Chem. Phys.* **81**, 3684 (1984).
- ⁴⁶B. Hess, H. Bekker, H. J. C. Berendsen, and J. G. E. M. Fraaije, *J. Comput. Chem.* **18**, 1463 (1997).
- ⁴⁷T. Darden, D. York, and L. Pedersen, *J. Chem. Phys.* **98**, 10089 (1993).
- ⁴⁸W. Kabsch, *Acta Crystallogr.* **A32**, 922 (1976).
- ⁴⁹W. Kabsch, *Acta Crystallogr., Sect. A: Cryst. Phys., Diffr., Theor. Gen. Crystallogr.* **A34**, 827 (1978).
- ⁵⁰B. Steipe, *Acta Crystallogr., Sect. A: Found. Crystallogr.* **A58**, 506 (2002).
- ⁵¹D. L. Theobald, *Acta Crystallogr., Sect. A: Found. Crystallogr.* **A61**, 478 (2005).
- ⁵²K. J. Ryan, *Journal of Computational and Graphical Studies* **12**, 585 (2003).