

1 **Bayesian hidden Markov model analysis of single-molecule force**
2 **spectroscopy: Characterizing kinetics under measurement uncertainty**

3 John D. Chodera,^{1,*} Phillip Elms,^{1,2,3} Frank Noé,⁴ Bettina Keller,⁴
4 Christian M. Kaiser,^{1,5} Aaron Ewall-Wice,⁶ Susan Marqusee,^{1,7,3}
5 Carlos Bustamante,^{1,7,3,5,8,9} and Nina Singhal Hinrichs¹⁰

6 ¹*California Institute of Quantitative Biosciences (QB3),*
7 *University of California, Berkeley, CA 94720, USA*

8 ²*Biophysics Graduate Group, University of California, Berkeley, CA 94720, USA*

9 ³*Jason L. Choy Laboratory of Single Molecule Biophysics,*
10 *Institute for Quantitative Biosciences,*
11 *University of California, Berkeley, CA 94720, USA*

12 ⁴*DFG Research Center Matheon, FU Berlin,*
13 *Arnimallee 6, 14195 Berlin, Germany*

14 ⁵*Department of Physics, University of California, Berkeley, CA 94720, USA*

15 ⁶*University of Chicago, IL 60637, USA*

16 ⁷*Department of Molecular & Cell Biology,*
17 *University of California, Berkeley, CA 94720, USA*

18 ⁸*Department of Chemistry, University of California, Berkeley, CA 94720, USA*

19 ⁹*Howard Hughes Medical Institute, University of California, Berkeley, CA 94720, USA*

20 ¹⁰*Departments of Statistics and Computer Science,*
21 *University of Chicago, IL 60637, USA*

22 (Dated: August 5, 2011)

Abstract

Single-molecule force spectroscopy has proven to be a powerful tool for studying the kinetic behavior of biomolecules. Through application of an external force, conformational states with small or transient populations can be stabilized, allowing them to be characterized and the statistics of individual trajectories studied to provide insight into biomolecular folding and function. Because the observed quantity (force or extension) is not necessarily an ideal reaction coordinate, individual observations cannot be uniquely associated with kinetically distinct conformations. While maximum-likelihood schemes such as hidden Markov models have solved this problem for other classes of single-molecule experiments by using temporal information to aid in the inference of a sequence of distinct conformational states, these methods do not give a clear picture of how precisely the model parameters are determined by the data due to instrument noise and finite-sample statistics, both significant problems in force spectroscopy. We solve this problem through a *Bayesian* extension that allows the experimental uncertainties to be directly quantified, and build in detailed balance to further reduce uncertainty through physical constraints. We illustrate the utility of this approach in characterizing the three-state kinetic behavior of an RNA hairpin in a stationary optical trap.

23 INTRODUCTION

24 Recent advances in biophysical measurement have led to an unprecedented ability to monitor
25 the dynamics of single biological macromolecules, such as proteins and nucleic acids [1]. As a
26 new approach to probing the behavior of biological macromolecules, these experiments promise
27 to change the way we study folding, dynamics, catalysis, association, transcription, translation,
28 and motility, providing otherwise-inaccessible information about microscopic kinetics, energetics,
29 mechanism, and the stochastic heterogeneity inherent in these processes. Advances in instrumen-
30 tation for optical force spectroscopy in particular have produced instruments of extraordinary
31 stability, precision, and temporal resolution [2, 3] that have already demonstrated great utility in
32 the study of biomolecules in the presence of externally perturbative forces [4–6]. Under external
33 force, it becomes possible to stabilize and characterize short-lived conformational states, such as
34 protein folding and unfolding intermediates [7–9].

35 In a typical single-molecule optical trapping experiment, a protein or nucleic acid is tethered to
36 two polystyrene beads by dsDNA handles that prevent the molecule under study from interacting
37 with the beads (see Figure 1). The handle-biomolecule-handle assembly—referred to as a *fiber*—
38 is associated with the beads through tight noncovalent interactions, with one bead held in an
39 optical trap and the other either suctioned to a micropipette or held in a second optical trap.
40 During an experiment, the position of the bead within the laser trap is monitored, and either the
41 relative displacement from the trap center or the total force on the bead is recorded, resulting in
42 a timeseries such as the one depicted in Figure 2. The instrument can generally be operated in
43 several modes: a *force ramp* mode, in which the trap is translated rapidly enough to potentially
44 carry the system out of equilibrium; an equilibrium *passive* mode, in which the trap is held fixed;
45 and a *constant force-feedback* mode, in which the trap is continually repositioned to maintain a set
46 constant force on the fiber. Here, we concern ourselves with the latter two classes of experiment,
47 though nonequilibrium experiments remain an exciting topic of active research [10].

48 Often, the dynamics observed in these experiments appears to be dominated by stochastic
49 transitions between two or more strongly metastable conformational states [11, 12]—regions of
50 conformation space in which the system remains for long times before making a transition to
51 another conformational state. These transitions are generally well-described by first-order ki-
52 netics [13]. While visual inspection of the dynamics may suggest the clear presence of multiple

metastable states, quantitative characterization of these states is often difficult. First, the observed force or extension is unlikely to correspond to a true reaction coordinate easily able to separate all metastable states [14–17], and second, measurement noise may further broaden the force or extension signatures of individual states, increasing their overlap. Attempting to separate these states by simply dividing the observed force or extension range into regions, following current practice [18, 19], can often lead to a high degree of state mis-assignment that results in the estimated rate constants and state distributions containing a significant amount of error [20] (see *Supplementary Material: Comparison with threshold model*).

Hidden Markov models (HMMs) [21], which use temporal information in addition to the instantaneous value of the observable (force or extension) to determine which conformational states the system has visited during the experiment, have provided an effective solution to the hidden state problem in many other classes of single-molecule experiments, such as ion channel currents [22–25], single-molecule FRET [26–30], and the stepping of motor proteins [31–33]. In applying hidden Markov modeling to the analysis of single-molecule force spectroscopy data, the observed force or extension trace is assumed to come from a realization of an underlying Markov chain, where the system makes history-independent transitions among a set of discrete conformational states with probabilities governed by a transition or rate matrix. Data, in the form of force or bead-to-bead extension measurements, is sampled at an interval that ensures that sequential observations satisfy the Markov property of history-independence, though the appropriate interval depends on the properties of the experimental configuration. Under a given set of external force conditions, each state has a distribution of forces or extensions associated with it. Given observed timeseries data for forces or extensions, the maximum likelihood estimate (MLE) of the model parameters (transition rates and state force or extension distributions) and sequence of hidden states corresponding to the observed data can be determined by standard methods [34, 35], as demonstrated in recent work [36].

Unfortunately, this approach has a number of significant drawbacks. Due to technical limitations, experiments often suffer from limited statistics—the events of interest (transitions between states or visits to rare states) may occur only a few times during the course of the measurement, and data for additional fibers is time-consuming to collect. As a result, while the MLE yields the most *likely* set of model parameters, there may be enormous uncertainty in some of these parameters, and the uncertainty in multiple parameters may be correlated in complex nonlin-

ear ways. While methods exist for estimating the standard error or confidence intervals from MLHMMs [37], these schemes can be prohibitively costly for long traces, and may still significantly underestimate the statistical error for short traces due to the normally-distributed error approximation inherent in the approach. The high cost (both in terms of instrument and experimenter time) of collecting additional data also means that it is not a simple task to judge *how much* data need be collected to test a particular hypothesis in a statistically meaningful way. Worse yet, the standard algorithms employed to find the MLE may not even find the true maximum likelihood solution, instead converging to a local maximum in likelihood that is far from optimal [38].

Here, we resolve this issue through the use of a *Bayesian* extension of hidden Markov models [39–42] applicable to single molecule force experiments. By sampling over the posterior distribution of model parameters and hidden state assignments instead of simply finding the most likely values, the experimenter is able to accurately characterize the correlated uncertainties in both the model parameters (transition rates and state force or extension distributions) and hidden state sequences corresponding to observed data. Additionally, prior information (either from additional independent measurements or physical constraints) can be easily incorporated. We also include a reversibility constraint on the transition matrix—in which microscopic detailed balance is imposed on the kinetics, as dictated by the physics of equilibrium systems [48]—which has been shown to significantly reduce statistical uncertainties in data-poor conditions [43, 44]. The framework we present is based on Gibbs sampling [45, 46], allowing simple swap-in replacement of models for observable distributions, extension to multiple observables, and alternative models for state transitions. Additionally, the Bayesian method provides a straightforward way to model the statistical outcome and assess the utility of additional experiments given some preliminary data, allowing the experimenter a powerful tool for assessing whether the cost of collecting additional data is outweighed by their benefits. A Matlab implementation of this approach is available online [<http://simtk.org/home/bhmm>].

HIDDEN MARKOV MODELS FOR FORCE SPECTROSCOPY

Suppose the temporal behavior of some observable $O(x)$ that is a function of molecular conformation x —here, generally force or molecular extension—is observed at temporal intervals Δt to produce a timeseries o_t , where $t = 0, 1, \dots, L$. An instantaneous observation o_t does not nec-

essarily contain enough information to unambiguously identify the current conformational state the molecule occupies; to infer the hidden state, we must also make use of the temporal information in the observed trace. We restrict ourselves to consideration of scalar functions $O(x)$, but the generalization to multidimensional probes (or multiple probes, such as combined force and fluorescence measurements [47]) and multiple observed temporal traces is straightforward.

We presume the system under study has M kinetically distinct states, in the sense that the system generally remains in a given state for several observation intervals Δt , but these states may not necessarily represent highly populated states of the system at equilibrium. We treat these conformational states as the *hidden states* of the model, because we cannot *directly* observe the identity of the metastable state in which the system resides. The hidden Markov model presumes the observed data $\mathbf{O} \equiv \{o_t\}$ was generated according to the following model dependent on parameters $\Theta \equiv \{\mathbf{T}, \mathbf{E}\}$, where \mathbf{T} is an $M \times M$ row-stochastic transition matrix and \mathbf{E} a set of emission parameters governing the observable (force or extension) distributions for each of the M hidden states, and prior information about the initial state distribution ρ ,

$$\begin{aligned}\mathbb{P}(s_0) &= \rho_{s_0} \\ \mathbb{P}(s_t | s_{t-1}, \mathbf{T}) &= T_{s_{t-1}s_t}, \quad t \geq 1 \\ \mathbb{P}(o_t | s_t, \mathbf{e}_{s_t}) &= \varphi(o_t | \mathbf{e}_{s_t}).\end{aligned}\tag{1}$$

In diagrammatic form, the observed state data $\{o_t\}$ and corresponding hidden state history $\{s_t\}$ can be represented

$$\begin{array}{ccccccc} \xrightarrow{\rho} & s_0 & \xrightarrow{\mathbf{T}} & s_1 & \xrightarrow{\mathbf{T}} & s_2 & \xrightarrow{\mathbf{T}} \cdots \xrightarrow{\mathbf{T}} s_L \\ & \downarrow \varphi & & \downarrow \varphi & & \downarrow \varphi & & \downarrow \varphi \\ & o_0 & & o_1 & & o_2 & & o_L \end{array}\tag{2}$$

The initial state distribution ρ reflects our knowledge of the initial conditions of the experiment that collected data \mathbf{o} . In the case that the experiment was prepared in equilibrium, ρ corresponds to the equilibrium distribution π of the model transition matrix \mathbf{T} ; if the experiment was prepared out of equilibrium, ρ may be chosen to reflect some other prior distribution (e.g. the uniform prior).

State transitions ($s_{t-1} \rightarrow s_t$) are governed by the discrete transition probability $T_{s_{t-1}s_t}$. The *Markov property* of HMMs prescribes that the probability that a system originally in state i at time

136 t is later found in state j at time $t + 1$ is dependent only on knowledge of the state i , and given by
 137 the corresponding matrix element T_{ij} of the (row-stochastic) transition matrix \mathbf{T} . Alternatively,
 138 one could instead use the rate matrix \mathbf{K} , related to the transition matrix \mathbf{T} through the equation
 139 $\mathbf{T} = e^{\mathbf{K}\Delta t}$. If the processes described by \mathbf{T} or \mathbf{K} are slow compared to the observation interval
 140 Δt , then we can easily estimate the rate matrix from the associated transition matrix in a way that
 141 avoids the matrix logarithm, through the expansion $\mathbf{K} \approx (\mathbf{T} - \mathbf{I})/\Delta t$, where \mathbf{I} denotes the $M \times M$
 142 identity matrix.

143 The probabilistic “emission” of observables from each state ($s_t \rightarrow o_t$) is governed by the con-
 144 tinuous emission probability $\varphi(o_t \mid \mathbf{e}_{s_t})$, parametrized by observable *emission* parameters \mathbf{e} . For
 145 example, in the force spectroscopy applications described here, $\varphi(o \mid \mathbf{e}_s)$ is taken to be a univariate
 146 normal (Gaussian) distribution parameterized by a mean μ and variance σ^2 that characterize each
 147 state, such that $\mathbf{e}_i \equiv \{\mu_i, \sigma_i^2\}$. Other choices of observable distribution can easily be substituted in
 148 a modular way without affecting the structure of the algorithms presented here.

149 Given the HMM process specified in Eq. 1, the probability of observing data \mathbf{O} given the model
 150 parameters Θ is then,

$$P(\mathbf{O} \mid \Theta) = \sum_{\mathbf{S}} \rho_{s_0} \varphi(o_0 \mid \mathbf{e}_{s_0}) \prod_{t=1}^L T_{s_{t-1}s_t} \varphi(o_t \mid \mathbf{e}_{s_t}), \quad (3)$$

151 where the sum over hidden state histories \mathbf{S} is shorthand for

$$\sum_{\mathbf{S}} \equiv \sum_{s_0=1}^M \sum_{s_1=1}^M \cdots \sum_{s_L=1}^M. \quad (4)$$

152 If multiple independent traces $\{o_t\}$ are available, the probability $P(\mathbf{O} \mid \Theta)$ is simply the product
 153 of Eq. 3 for the independent traces.

154 Maximum likelihood hidden Markov model (MLHMM)

155 The standard approach to construct an HMM from observed data is to compute the *maximum*
 156 *likelihood estimator* (MLE) for the model parameters $\Theta \equiv \{\mathbf{T}, \mathbf{E}\}$, which maximize the probability
 157 of the observed data \mathbf{O} given the model,

$$\hat{\Theta} = \arg \max_{\Theta} P(\mathbf{O} \mid \Theta), \quad (5)$$

158 yielding MLE estimates of transition matrix $\hat{\mathbf{T}}$ and state emission parameters $\hat{\mathbf{E}}$. Typically, deter-
 159 mination of the model parameters Θ is carried out using the Baum-Welch algorithm [34].

160 Once the MLE parameters $\hat{\Theta}$ are determined, the most likely hidden state history that pro-
 161 duced the observations \mathbf{O} can be determined using these parameters:

$$\hat{\mathbf{S}} = \arg \max_{\mathbf{S}} P(\mathbf{S} | \mathbf{O}, \hat{\Theta}). \quad (6)$$

162 This is typically carried out using the Viterbi algorithm [35], a classic example of dynamic pro-
 163 gramming.

164 **Bayesian hidden Markov model (BHMM)**

165 Instead of simply determining the model that maximizes the likelihood of observing the data
 166 \mathbf{O} given the model parameters Θ , we can make use of Bayes' theorem to compute the *posterior*
 167 distribution of model parameters given the observed data:

$$P(\Theta | \mathbf{O}) \propto P(\mathbf{O} | \Theta)P(\Theta). \quad (7)$$

168 Here, $P(\Theta)$ denotes a *prior* distribution that encodes any *a priori* information we may have about
 169 the model parameters Θ . This prior information might include, for example, physical constraints
 170 (such as ensuring the transition matrix satisfies detailed balance) or prior rounds of inference from
 171 other independent experiments.

172 Making use of the likelihood (Eq. 3), the model posterior is then given by,

$$P(\Theta | \mathbf{O}) \propto P(\Theta) \sum_{\mathbf{S}} \rho_{s_0} \varphi(o_0 | \mathbf{e}_{s_0}) \prod_{t=1}^L T_{s_{t-1}s_t} \varphi(o_t | \mathbf{e}_{s_t}). \quad (8)$$

173 Drawing samples of Θ from this distribution will, in principle, allow the *confidence* with which
 174 individual parameters and combinations thereof are known, given the data (and subject to the
 175 validity of the model of Eq. 1 in correctly representing the process by which the observed data
 176 is generated). While the posterior $P(\Theta|\mathbf{O})$ is complex, we could in principle use a Markov chain
 177 Monte Carlo (MCMC) approach [46] to sample it. In its current form, however, this would be
 178 extremely expensive due to the sum over all hidden state histories \mathbf{S} appearing in ratios involving
 179 Eq. 8. Instead, we introduce the hidden state histories \mathbf{S} as an auxiliary variable, sampling from
 180 the augmented posterior,

$$P(\Theta, \mathbf{S} | \mathbf{O}) \propto \left[\rho_{s_0} \varphi(o_0 | \mathbf{e}_{s_0}) \prod_{t=1}^L T_{s_{t-1}s_t} \varphi(o_t | \mathbf{e}_{s_t}) \right] P(\Theta). \quad (9)$$

181 which makes it much less costly to compute the ratios required for MCMC on the augmented
 182 (Θ, \mathbf{S}) parameter space.

183 If we presume the prior is separable, such that $P(\Theta) \equiv P(\mathbf{T})P(\mathbf{E})$, we can sample from the
 184 augmented posterior (Eq. 9) using the framework of *Gibbs sampling* [46], in which the augmented
 185 model parameters are updated by sampling from the conditional distributions,

$$\begin{aligned}
 P(\mathbf{S} \mid \mathbf{T}, \mathbf{E}, \mathbf{O}) &\propto \rho_{s_0} \varphi(o_0 \mid \mathbf{e}_{s_0}) \prod_{t=1}^L T_{s_{t-1}s_t} \varphi(o_t \mid \mathbf{e}_{s_t}) \\
 P(\mathbf{T} \mid \mathbf{E}, \mathbf{S}, \mathbf{O}) &= P(\mathbf{T} \mid \mathbf{S}) \propto P(\mathbf{T}) \prod_{t=1}^L T_{s_{t-1}s_t} \\
 P(\mathbf{E} \mid \mathbf{S}, \mathbf{T}, \mathbf{O}) &= P(\mathbf{E} \mid \mathbf{S}, \mathbf{O}) \propto P(\mathbf{E}) \prod_{t=0}^L \varphi(o_t \mid \mathbf{e}_{s_t}).
 \end{aligned} \tag{10}$$

186 The equalities on the second and third lines reflect the conditional independence of the hidden
 187 Markov model defined by Eq. 1. When only the model parameters $\Theta \equiv \{\mathbf{T}, \mathbf{E}\}$ or the hidden state
 188 histories \mathbf{S} are of interest, we can simply marginalize out the uninteresting variables by sampling
 189 from the augmented joint posterior for $\{\mathbf{T}, \mathbf{E}, \mathbf{S}\}$ and examine only the variables of interest. In
 190 addition, the structure of the Gibbs sampling scheme above allows individual components (such
 191 as the observable distribution model $\varphi(o \mid \mathbf{e})$ or transition probability matrix \mathbf{T}) to be modified
 192 without affecting the structure of the remainder of the calculation.

193 In the illustrations presented here, we employ a Gaussian observable distribution model for
 194 $\varphi(o \mid \mathbf{e})$,

$$\varphi(o \mid \mathbf{e}) = \varphi(o \mid \mu, \sigma^2) = \frac{1}{\sqrt{2\pi}\sigma} \exp \left[-\frac{1}{2} \frac{(o - \mu)^2}{\sigma^2} \right], \tag{11}$$

195 where μ is the mean force or extension characterizing a particular state, and σ is the standard
 196 deviation or width of forces or extensions corresponding to that state. We note that marginal
 197 posterior distributions of each mean $P(\mu_i \mid \mathbf{O})$ reflect the statistical uncertainty in how well the
 198 mean force or position is determined, and need not correspond to the standard deviation σ_i , which
 199 may be much broader (or narrower, depending on the situation).

200 ALGORITHMS

201 Generating an initial model

202 To initialize either computation of the MLHMM or sampling from the posterior for the BHMM,
 203 an initial model that respects any constraints imposed in the model prior $P(\Theta)$ must be selected.
 204 Here, we employ a Gaussian observable distribution model for $\varphi(o | \mathbf{e})$ (Eq. 11) and enforce that
 205 the transition matrix \mathbf{T} satisfy detailed balance.

206 Observable parameter estimation

207 We first initialize the observed distributions of each state by fitting a Gaussian mixture model
 208 with M states to the pooled observed data \mathbf{O} , ignoring temporal information:

$$P(\mathbf{O} | \boldsymbol{\pi}, \mathbf{E}) = \prod_{t=0}^L \sum_{m=1}^M \pi_m \varphi(o_t | \mu_m, \sigma_m^2), \quad (12)$$

209 where the state observable emission probability vector $\mathbf{E} \equiv \{\mathbf{e}_1, \dots, \mathbf{e}_M\}$ and $\mathbf{e}_m \equiv \{\mu_m, \sigma_m^2\}$
 210 with μ_m denoting the observable mean and σ_m^2 the variance for state m for the Gaussian mixture
 211 model. The vector $\boldsymbol{\pi}$ is composed of equilibrium state populations $\{\pi_1, \dots, \pi_M\}$ with $\pi_m \geq 0$ and
 212 $\sum_{m=1}^M \pi_m = 1$.

213 A first approximation to $\boldsymbol{\pi}$ and \mathbf{E} is computed by pooling and sorting the observed o_t , and
 214 defining M indicator functions $h_m(o)$ that separate the data into M contiguous regions of the ob-
 215 served range of o of roughly equal population. Let $N_m \equiv \sum_{t=0}^L h_m(o_t)$ denote the total number of
 216 observations falling in region m , and $N_{\text{tot}} = \sum_{m=1}^M N_m$. The initial parameters are then computed
 217 as,

$$\pi_m = N_m / N_{\text{tot}}$$

$$\mu_m = N_m^{-1} \sum_{t=0}^L o_t h_m(o_t) \quad (13)$$

$$\sigma_m^2 = N_m^{-1} \sum_{t=0}^L (o_t - \mu_m)^2 h_m(o_t). \quad (14)$$

218 This approximation is then improved upon by iterating the expectation-maximization proce-

219 dure described by Bilmes [49],

$$\begin{aligned}
\pi'_m &= N_{\text{tot}}^{-1} \sum_{t=0}^L \chi_m(o_t, \mathbf{E}, \boldsymbol{\pi}) \\
\mu'_m &= (\pi'_m N_{\text{tot}})^{-1} \sum_{t=0}^L o_t \chi_m(o_t, \mathbf{E}, \boldsymbol{\pi}) \\
\sigma'^2_m &= (\pi'_m N_{\text{tot}})^{-1} \sum_{t=0}^L (o_t - \mu'_m)^2 \chi_m(o_t, \mathbf{E}, \boldsymbol{\pi})
\end{aligned} \tag{15}$$

220 where the function $\chi_m(o, \mathbf{E}, \boldsymbol{\pi})$ is given by the fuzzy membership function,

$$\chi_m(o, \mathbf{E}, \boldsymbol{\pi}) = \frac{\pi_m \varphi(o \mid \mathbf{e}_m)}{\sum_{l=1}^M \pi_l \varphi(o \mid \mathbf{e}_l)}. \tag{16}$$

221 The iterative procedure is terminated at iteration j when the change in the parameters $\{\boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\sigma}^2\}$
222 falls below a certain relative threshold, such as $\|\boldsymbol{\pi}^{[j]} - \boldsymbol{\pi}^{[j-1]}\|_2 / \|\boldsymbol{\pi}^{[j]}\|_2 < 10^{-4}$.

223 *Transition matrix estimation*

224 Once initial state observable emission parameters \mathbf{E} are determined, an initial transition ma-
225 trix is estimated using an iterative likelihood maximization approach that enforces detailed bal-
226 ance [50]. First, a matrix of fractional transition counts $\mathbf{C} \equiv (c_{ij})$ is estimated using the member-
227 ship function:

$$c_{ij} = \sum_{t=1}^L \chi_i(o_{t-1}, \mathbf{E}, \boldsymbol{\pi}) \chi_j(o_t, \mathbf{E}, \boldsymbol{\pi}) \tag{17}$$

228 A symmetric $M \times M$ matrix $\mathbf{X} \equiv (x_{ij})$ is initialized by

$$x_{ij} = x_{ji} = c_{ij} + c_{ji}. \tag{18}$$

229 The iterative procedure described in Algorithm 1 of [50] is then applied. For each update iteration,
230 we first update the diagonal elements of \mathbf{X} :

$$x'_{ii} = \frac{c_{ii}(x_{i*} - x_{ii})}{c_{i*} - c_{ii}} ; \quad c_{i*} = \sum_{j=1}^M c_{ij} ; \quad x_{i*} = \sum_{j=1}^M x_{ij}, \tag{19}$$

231 followed by the off-diagonal elements:

$$x'_{ij} = x'_{ji} = \frac{-b + \sqrt{b^2 - 4ac}}{2a} \tag{20}$$

where the quantities a , b , and c are computed from \mathbf{X} and \mathbf{C} ,

$$\begin{aligned} a &\equiv c_{i*} - c_{ij} + c_{j*} - c_{ji} \\ b &\equiv c_{i*}(x_{j*} - x_{ji}) + c_{j*}(x_{i*} - x_{ij}) \\ &\quad - (c_{ij} + c_{ji})(x_{i*} - x_{ij} + x_{j*} - x_{ji}) \\ c &\equiv -(c_{ij} + c_{ji})(x_{i*} - x_{ij})(x_{j*} - x_{ji}). \end{aligned} \quad (21)$$

Once a sufficient number of iterations j have been completed to compute a stable estimate of \mathbf{X} (such as the relative convergence criteria $\|\mathbf{X}^{[j]} - \mathbf{X}^{[j-1]}\|_2 / \|\mathbf{X}^{[j]}\|_2 < 10^{-4}$, the maximum likelihood transition matrix estimate \mathbf{T} is computed as

$$T_{ij} = x_{ij} / x_{i*}. \quad (22)$$

Note that the equilibrium probability vector $\boldsymbol{\pi}$ computed during the Gaussian mixture model fitting is not respected during this step.

Fitting a maximum likelihood HMM

The HMM model parameters $\boldsymbol{\Theta} \equiv \{\mathbf{T}, \mathbf{E}\}$ are fit to the observed data \mathbf{O} through use of the expectation-maximization (EM) algorithm [51]. This is an iterative procedure, where the model parameters are subsequently refined through successive iterations. The initial HMM is usually quick to compute, and can give the experimenter a rough idea of the model parameters, as well as providing a useful starting point for sampling models from the Bayesian posterior.

During each iteration, the Baum-Welch algorithm [34] is used to compute $\Xi \equiv (\xi_{tij})$, which represents the probability that the system transitions from hidden state i at time $t - 1$ to hidden state j at time t , and γ_{ti} , the probability that the system occupied state i at time t . This is accomplished by first executing the *forward algorithm*,

$$\alpha_{tj} = \begin{cases} \rho_j \varphi(o_0 | \mathbf{e}_j) & t = 0 \\ \varphi(o_t | \mathbf{e}_j) \sum_{i=1}^M \alpha_{(t-1)i} T_{ij} & t = 1, \dots, L \end{cases} \quad (23)$$

followed by the *backward algorithm*,

$$\beta_{ti} = \begin{cases} 1 & t = L \\ \sum_{j=1}^M T_{ij} \varphi(o_{t+1} | \mathbf{e}_j) \beta_{(t+1)j} & t = (L-1), \dots, 0 \end{cases} \quad (24)$$

249 The $L \times M \times M$ matrix Ξ is then computed for $t = 0, \dots, (L - 1)$ as,

$$\xi_{tij} = \alpha_{ti} \varphi(o_{t+1} | \mathbf{e}_i) T_{ij} \beta_{(t+1)j} / \sum_{i=1}^M \alpha_{Ti} \quad (25)$$

$$\gamma_{ti} = \sum_{j=1}^M \xi_{tij} \quad (26)$$

250 In practice, the logarithms of these quantities are computed instead to avoid numerical underflow.

251 The aggregate matrix of expected transition counts $\mathbf{C} \equiv (c_{ij})$ is then computed from Ξ as,

$$c_{ij} = \sum_{t=0}^{L-1} \xi_{tij}. \quad (27)$$

252 This count matrix is used to update the maximum-likelihood transition matrix \mathbf{T} using the
253 method of Prinz et al. [50] described in the previous section.

254 The state observable distribution parameters \mathbf{E} are then updated from the γ_{ti} . For the uni-
255 variate normal distribution applied to force spectroscopy data here, we update the mean μ_i and
256 variance σ_i^2 for state i using the scheme,

$$\mu'_i = \frac{\sum_{t=0}^L o_t \gamma_{ti}}{\sum_{t=0}^L \gamma_{ti}} ; \sigma'^2_i = \frac{\sum_{t=0}^L (o_t - \mu'_i)^2 \gamma_{ti}}{\sum_{t=0}^L \gamma_{ti}}. \quad (28)$$

257 Once the model parameters have been fitted by iteration of the above update procedure to
258 convergence (which may only converge to a local maximum of the likelihood), the most likely
259 hidden state sequence can be determined given the observations \mathbf{O} and the MLE model $\hat{\Theta}$ using
260 the Viterbi algorithm [35]. Like the forward-backward algorithm employed in the Baum-Welch
261 procedure, the Viterbi algorithm also has a forward recursion component,

$$\epsilon_{jt} = \begin{cases} \rho_j \varphi(o_t | \mathbf{e}_j) & t = 0 \\ \varphi(o_t | \mathbf{e}_j) \max_i \epsilon_{i(t-1)} T_{ij} & t = 1, \dots, L \end{cases} \quad (29)$$

$$\Phi_{jt} = \begin{cases} 1 & t = 0 \\ \arg \max_i \epsilon_{i(t-1)} T_{ij} & t = 1, \dots, L \end{cases}$$

262 as well as a reverse reconstruction component to compute the most likely state sequence $\hat{\mathbf{S}}$,

$$\hat{s}_t = \begin{cases} \arg \max_i \epsilon_{it} & t = L \\ \Phi_{\hat{s}_{t+1}(t+1)} & t = (L - 1), \dots, 0 \end{cases} \quad (30)$$

263 Sampling from the posterior of the BHMM

264 Sampling from the posterior of the BHMM (Eq. 8) proceeds by rounds of Gibbs sampling,
 265 where each round consists of an update of the augmented model parameters $\{\mathbf{T}, \mathbf{E}, \mathbf{S}\}$ by sam-
 266 pling

$$\begin{aligned}\mathbf{S}' \mid \mathbf{T}, \mathbf{E}, \mathbf{O} &\sim P(\mathbf{S}' \mid \mathbf{T}, \mathbf{E}, \mathbf{O}) \\ \mathbf{T}' \mid \mathbf{S}' &\sim P(\mathbf{T}' \mid \mathbf{S}') \\ \mathbf{E}' \mid \mathbf{S}', \mathbf{O} &\sim P(\mathbf{E}' \mid \mathbf{S}', \mathbf{O})\end{aligned}$$

267 where the conditional probabilities are given by Eq. 10.

268 *Updating the hidden state sequences*

269 We use a modified form of the Viterbi process to generate an independent sample of the hidden
 270 state history \mathbf{S} given the transition probabilities \mathbf{T} , state observable distribution parameters \mathbf{E} ,
 271 and observed data \mathbf{O} . Like the Viterbi scheme, a forward recursion is applied to each observation
 272 trace \mathbf{o} , but instead of computing the most *likely* state history on the reverse pass, a new hidden
 273 state history \mathbf{S} is drawn from the distribution $P(\mathbf{S} \mid \mathbf{O}, \mathbf{T}, \mathbf{E})$. The forward recursion uses the
 274 same forward algorithm as used in Baum-Welch [34],

$$\alpha_{tj} = \begin{cases} \rho_j \varphi(o_0 \mid \mathbf{e}_j) & t = 0 \\ \varphi(o_t \mid \mathbf{e}_j) \sum_{i=1}^M \alpha_{(t-1)i} T_{ij} & t = 1, \dots, L \end{cases} \quad (31)$$

275 In the reverse recursion, we now *sample* a state sequence by sampling each hidden state from the
 276 conditional distribution $s_t \sim P(s_t \mid s_{t+1}, \dots, s_L)$ starting from $t = L$ and proceeding down to
 277 $t = 0$, where the conditional distribution is given by,

$$\begin{aligned} &P(s_t = i \mid s_{t+1}, \dots, s_L) \\ &\propto \begin{cases} \alpha_{ti} / \sum_{j=1}^M \alpha_{tj} & t = L \\ \alpha_{ti} T_{is_{t+1}} / \sum_{j=1}^M \alpha_{tj} T_{js_{t+1}} & t = (L-1), \dots, 0 \end{cases} \end{aligned} \quad (32)$$

278 It is straightforward to show the result of these sampling steps reconstitutes the probability dis-
 279 tribution $P(\mathbf{S} \mid \mathbf{T}, \mathbf{E}, \mathbf{O})$ (see *Supplementary Material: Proof of state history sampling scheme*).

281 If no detailed balance constraint is used and the prior $P(\mathbf{T})$ is Dirichlet in each row of the
 282 transition matrix \mathbf{T} , it is possible to generate an independent sample of the transition matrix from
 283 the conditional distribution $P(\mathbf{T}' | \mathbf{S}')$ by sampling each row of the transition matrix from the
 284 conjugate Dirichlet posterior using the transition counts from the sampled state sequence \mathbf{S}' [43].
 285 However, because physical systems in the absence of energy input through an external driving
 286 force should satisfy detailed balance, we make use of this constraint in updating our transition
 287 probabilities, since this has been demonstrated to substantially reduce parameter uncertainty in
 288 the data-limited regime [43].

289 The transition matrix is updated using the reversible transition matrix sampling scheme of
 290 Noé [43, 52]. Here, an adjusted count matrix $\mathbf{C} \equiv (c_{ij})$ is computed using the updated hidden
 291 state sequence \mathbf{S}' ,

$$c_{ij} = b_{ij} + \sum_{t=1}^L \delta_{is_{t-1}} \delta_{js_t}, \quad (33)$$

292 where the Kronecker $\delta_{ij} = 1$ if $i = j$ and zero otherwise, and $\mathbf{B} \equiv (b_{ij})$ is a matrix of prior
 293 pseudocounts, which we take to be zero following the work of Noé et al. [13]. Using the adjusted
 294 count matrix \mathbf{C} , a Metropolis-Hastings Monte Carlo procedure [53] is used to update the matrix
 295 and produce a new sample from $P(\mathbf{T}' | \mathbf{S}')$. Two move types are attempted, selected with equal
 296 probability, and 1000 moves are attempted to generate a new sample \mathbf{T}' that is approximately
 297 uncorrelated from the previous \mathbf{T} . Prior to starting the Monte Carlo procedure, the vector of
 298 equilibrium probabilities for all states $\boldsymbol{\pi}$ is computed according to

$$\mathbf{T}^T \boldsymbol{\pi} = \boldsymbol{\pi}. \quad (34)$$

The first move type is a *reversible element shift*. A pair of states (i, j) , $i \neq j$, are selected with uniform probability, and a random number Δ is selected uniformly over the interval,

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

The changed elements in the proposed transition matrix \mathbf{T}' are then given by:

$$\begin{aligned} T'_{ij} &= T_{ij} - \Delta ; \quad T'_{ji} = T_{ji} - \frac{\pi_i}{\pi_j} \Delta \\ T'_{ii} &= T_{ii} + \Delta ; \quad T'_{jj} = T_{jj} + \frac{\pi_i}{\pi_j} \Delta. \end{aligned}$$

299 This move is accepted with probability

$$P_{\text{accept}}(\mathbf{T}'|\mathbf{T}) = \min \left\{ 1, \sqrt{\frac{(T'_{ij})^2 + (T'_{ji})^2}{(T_{ij})^2 + (T_{ji})^2}} \right. \\ \left. \times \left(\frac{T'_{ii}}{T_{ii}} \right)^{c_{ii}} \left(\frac{T'_{ij}}{T_{ij}} \right)^{c_{ij}} \left(\frac{T'_{jj}}{T_{jj}} \right)^{c_{jj}} \left(\frac{T'_{ji}}{T_{ji}} \right)^{c_{ji}} \right\}. \quad (35)$$

300 This move will leave the vector of stationary probabilities π unchanged.

The second move type is a *row shift*. A row i of \mathbf{T} is selected with uniform probability, and a random number η chosen uniformly over the interval

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

301 and used to update row i of \mathbf{T} according to

$$T'_{ij} = \begin{cases} \eta T_{ij} & j = 1, \dots, M, j \neq i \\ \eta(T_{ii} - 1) + 1 & j = i \end{cases} \quad (36)$$

302 This move is accepted with probability

$$P_{\text{accept}}(\mathbf{T}'|\mathbf{T}) = \min \left\{ 1, \eta^{(M-2)} \eta^{(c_{i*} - c_{ii})} \left(\frac{1 - \eta(1 - T_{ii})}{T_{ii}} \right)^{c_{ii}} \right\}. \quad (37)$$

The row shift operation will change the stationary distribution of π' , but it may be efficiently updated with

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

303 Since this update scheme is incremental, it will accumulate numerical errors over time that cause
304 the updated π to drift away from the stationary distribution of the current transition matrix. To
305 avoid this, π is recomputed from the current sample of the transition matrix in regular intervals
306 (here, every 100 sampling steps).

307 *Updating the observable distribution parameters*

308 Following the update of the transition matrix \mathbf{T} , the observable distribution parameters \mathbf{E} are
309 updated by sampling \mathbf{E} from the conditional probability $P(\mathbf{E}' | \mathbf{S}', \mathbf{O})$. The conditional proba-
310 bility for the observable distribution parameters for state m , denoted \mathbf{e}_m , is given in terms of the

311 output model $\varphi(o | \mathbf{e})$ by Bayes' theorem,

$$P(\mathbf{E} | \mathbf{O}, \mathbf{S}) = \left[\prod_{t=0}^L \varphi(o_t | \mathbf{e}_{s_t}) \right] P(\mathbf{E}). \quad (38)$$

312 An important choice must be made with regards to the prior, $P(\mathbf{E})$. If the prior is chosen to be
313 composed of independent priors for each state, as in

$$P(\mathbf{E}) = \prod_{m=1}^M P(\mathbf{e}_m), \quad (39)$$

314 then the full BHMM posterior (Eq. 8) will be invariant under any permutation of the states. This
315 behavior might be undesirable, as the states may switch labels during the posterior sampling
316 procedure; this will require any analysis of the models sampled from the posterior to account for
317 the possible permutation symmetry in the states. On the other hand, breaking this symmetry (e.g.,
318 by enforcing an ordering on the state mean observables) can artificially restrict the confidence
319 intervals of the states, which might additionally complicate data analysis.

320 Here, we make the choice that the prior be separable (Eq. 39), which has the benefit of allowing
321 the conditional probability for \mathbf{E} (Eq. 38) to be decomposed into a separate posterior for each state.
322 For each state m , collect all the observations o_t whose updated hidden state labels $s_t' = m$ into
323 a single dataset $\mathbf{o} \equiv \{o_n\}_{n=1}^{N_m}$, where N_m is the total number of times state m is visited, for the
324 purposes of this update procedure. Then, the observable parameters \mathbf{e} for this state are given by

$$P(\mathbf{e} | \mathbf{o}) = P(\mathbf{o} | \mathbf{e})P(\mathbf{e}) = \left[\prod_{n=1}^{N_m} \varphi(o_n | \mathbf{e}) \right] P(\mathbf{e}). \quad (40)$$

325 In the application presented here, we use a Gaussian output model (Eq. 11) for the state ob-
326 servable distributions $P(o | \mathbf{e})$, where $\mathbf{e} \equiv \{\mu, \sigma^2\}$, with μ the state mean observable and σ^2
327 the variance (which will include both the distribution of the observable characterizing the state
328 and any broadening from measurement noise). Other models (including multidimensional or
329 multimodal observation models) are possible, and require replacing only the observation model
330 $\varphi(o | \mathbf{e})$ and corresponding prior $P(\mathbf{e})$.

331 We use the (improper) Jeffreys prior [54] which has the information-theoretic interpretation
332 as the prior that maximizes the information content of the data [55], (suppressing the state index
333 subscript m),

$$P(\mathbf{e}) \propto \sigma^{-1}, \quad (41)$$

334 which produces the posterior

$$P(\mathbf{e} \mid \mathbf{o}) \propto \sigma^{-(N+1)} \exp \left[-\frac{1}{2\sigma^2} \sum_{n=1}^N (o_n - \mu)^2 \right], \quad (42)$$

335 where we remind the reader that here and in the remainder of this section, the symbols \mathbf{e} , \mathbf{o} , σ , μ ,

336 and N refer to \mathbf{e}_m , \mathbf{o}_m , σ_m , μ_m , and N_m , respectively.

337 Updating $\{\mu, \sigma^2\}$ also proceeds by a Gibbs sampling scheme, alternately updating μ and σ , as
338 earlier described in Ref. [52],

$$\begin{aligned} \mu &\sim P(\mu \mid \sigma^2, \mathbf{o}) \\ \sigma^2 &\sim P(\sigma^2 \mid \mu, \mathbf{o}) \end{aligned} \quad (43)$$

339 The conditional distribution of the mean μ is then given by

$$P(\mu \mid \sigma^2, \mathbf{o}) \propto \exp \left[-\frac{1}{2(\sigma^2/N)} (\mu - \hat{\mu})^2 \right] \quad (44)$$

340 where $\hat{\mu}$ is the sample mean for \mathbf{o} , the samples in state m ,

$$\hat{\mu} \equiv \frac{1}{N} \sum_{n=1}^N o_n \quad (45)$$

341 This allows us to update μ according to

$$\mu' \sim \mathcal{N}(\hat{\mu}, \sigma^2/N) \quad (46)$$

342 The conditional distribution of the variance σ^2 is given by

$$P(\sigma^2 \mid \mu, \mathbf{o}) \propto \sigma^{-(N+1)} \exp \left[-\frac{N\hat{\sigma}^2}{2\sigma^2} \right] \quad (47)$$

343 where the quantity $\hat{\sigma}^2$, which is *not* in general identical to the sample variance, is given by

$$\hat{\sigma}^2 \equiv \frac{1}{N} \sum_{n=1}^N (o_n - \mu)^2. \quad (48)$$

344 A convenient way to update $\sigma^2 \mid \mu, \mathbf{o}$ is to sample a random variate y from the chi-square distri-
345 bution with $N - 1$ degrees of freedom,

$$y \sim \chi^2(N - 1) \quad (49)$$

346 and then update σ^2 as

$$\sigma'^2 = N\hat{\sigma}^2/y. \quad (50)$$

347 Note that μ and σ^2 can be updated in either order, but the *updated* values of μ or σ^2 must be used
348 in sampling the not-yet-updated σ^2 or μ , and vice-versa.

349 Other output probabilities, such as *mixtures* of normal distributions or other distributions, can
350 be substituted by simply changing $P(\mathbf{E} \mid \mathbf{O}, \mathbf{S})$ and the scheme by which \mathbf{E} is updated.

351 VALIDATION USING SYNTHETIC DATA

352 To verify that our BHMM posterior sampling scheme reflects the *true* uncertainty in the model
353 parameters, we tested the scheme on synthetic data generated from a model with known pa-
354 rameters Θ^* . Given observed data \mathbf{O} generated from $P(\mathbf{O} \mid \Theta^*)$, sampling from the posterior
355 $P(\Theta \mid \mathbf{O})$ using the scheme described in *Sampling from the posterior of the BHMM* will provide us
356 with confidence intervals $[\theta_{\text{low}}, \theta_{\text{high}}]$ for a specified confidence interval size $\alpha \in [0, 1]$. If these
357 computed confidence intervals are accurate, we should find that the true model parameter θ^* lies
358 in the computed confidence interval $[\theta_{\text{low}}^{(\alpha)}, \theta_{\text{high}}^{(\alpha)}]$ with probability α . This can be tested by gener-
359 ating synthetic observed data \mathbf{O} from $P(\mathbf{O} \mid \Theta^*)$ and verifying that we find $\theta^* \in [\theta_{\text{low}}^{(\alpha)}, \theta_{\text{high}}^{(\alpha)}]$ in a
360 fraction α of these synthetic experiments.

361 As an example synthetic model, consider the three-state system intended to mimic a pro-
362 tein with (1) a highly-compliance, low-force unfolded state, (2) a moderately compliant low-
363 population intermediate at intermediate force, and (3) a low-compliance, high-force folded state.
364 Here, the term “compliance” refers to the width of the force or extension distribution character-
365 izing the state. Parameters of the model are given in Table I, and the observation interval was
366 taken to be $\tau = 1$ ms. An example realization of a model trajectory, along with the MLHMM
367 state assignment, is shown in Figure 2. We generated a trajectory of 100 000 observations, and
368 characterized the BHMM mean parameter estimate and 95% confidence intervals for a subset of
369 this trajectory of varying lengths. The results, shown in Table I, show that the confidence intervals
370 contract as trajectory length increases, as expected, and the BHMM-computed 95% confidence in-
371 tervals contain the true model parameters with the expected statistics. In contrast, a model created
372 from simply segmenting the observed forces into disjoint region and assigning state membership
373 based on the force value alone estimates model parameters with significant bias even for 1 000
374 000 observations (see *Supporting Information*).

375 As a more rigorous test, we sampled 50 random models from the prior $P(\Theta)$ with two to six
376 states, generated a 10 000 observation synthetic trajectory for each, and accumulated statistics on
377 the observed fraction of time the true model parameters were within the BHMM confidence in-
378 tervals for various values of the confidence interval width α . The results of this test are depicted
379 in *Supplementary Figure 1*. We expect that the plot traces the diagonal if the observed and expected
380 confidence intervals are identical; an overestimate of the confidence interval will be above the

381 diagonal, and an underestimate will fall below it. Because only a finite number of independent
382 replicates of the experiment are conducted, there is some associated uncertainty with the observed
383 confidence intervals. The results show that the observed confidence intervals line up with the ex-
384 pected confidence intervals to within statistical error, suggesting the BHMM confidence intervals
385 neither underestimate nor overestimate the actual uncertainty in model parameters.

386 RNA HAIRPIN KINETICS IN A PASSIVE OPTICAL TRAP

387 We illustrate the BHMM approach applied to real force spectroscopy data by characterizing
388 the average forces and transition rates among kinetically distinct states of the p5ab RNA hairpin
389 in an optical trap under passive (equilibrium) conditions.

390 The p5ab RNA hairpin from *Tetrahymena thermophila* was provided by Jin-Der Wen, and pre-
391 pared as previously described [56]. Within the population of RNA hairpin molecules in the ex-
392 amined sample, there were two chemically distinct species present in the sample (i.e. as a result
393 of post-transcriptional or other covalent modification during sample storage), exhibiting either
394 apparent two-state (as reported previously [56]) or three-state behavior (studied here). For the
395 purposes of testing this method, we examined a fiber that appeared to consistently exhibit three-
396 state behavior upon visual inspection of the force timeseries data.

397 The instrument used in this experiment was a dual-beam counter-propagating optical trap
398 with a spring constant of 0.1 pN/nm. A piezoactuator controlled the position of the trap and
399 allowed position resolution to within 0.5 nm [57]. Drift in the instrument was less than 1
400 nm/minute resulting in a constant average force within 0.1 pN over the course of a typical 60
401 s experiment. For these constant trap position experiments, higher frequency data was recorded
402 at 50 kHz recording the voltage corresponding to the force on the tether directly from the position-
403 sensitive detectors. To ensure sequential samples obeyed Markovian statistics, these data were
404 subsampled down to 1 kHz for analysis by the BHMM framework after examination of autocorre-
405 lation functions for trap positions where the hairpin appeared to remain in a single conformational
406 state (see *Supplementary Material: Choice of observation interval*).

407 A single observed force trajectory at a fixed trap position adequate to cause hopping among
408 multiple states is shown in Figure 3. The most likely state trajectory from the MLHMM fit with
409 three states is shown by coloring the observations most likely to be associated with each state, with

bands of color indicating the mean and standard deviation about the mean force characterizing each state.

Table II lists the BHMM posterior means and confidence intervals characterizing the three-state model extracted from this single 60 s observed force trace. Several things are notable about the estimated model parameters. Surprisingly, while there is a clearly-resolved intermediate-force state (*state 2*) through which most of the flux from the high- and low-force states passes (as seen from large K_{12} and K_{23}), there are nontrivial rate constants connecting the high and low force states directly (K_{13}), indicating that while a sequential mechanism involving passing through the intermediate state is preferred, it may not be an obligatory step in hairpin formation under these conditions. While the state mean forces are clearly distinct, the state standard deviations—which reflect the width of the observed force distribution characterizing each state, rather than the uncertainty in state means—possess overlapping confidence intervals. These standard deviations reflect not only contributions from both the distribution of extensions sampled by the hairpin in each conformational state, but also from fluctuations in the handles and beads, and other sources of mechanical and electrical noise in the measurement. As we would expect the unfolded hairpin to be more compliant (i.e. possess a wider distribution of forces) than the folded hairpin, the inability to distinguish the standard deviations among states is suggestive that, for this experimental configuration and observation time, the predominant contribution to the observed force distributions for each state may be in the form of handle or bead fluctuations or other sources of measurement noise.

Finally, the lifetime of the intermediate-force state is significantly shorter than for the low- and high-force states by nearly an order of magnitude, and only a few times longer than the observation interval of 1 ms—despite this, the lifetime appears to be well-determined, as indicated by the narrow confidence intervals.

DISCUSSION

We have described an approach to determining the first-order kinetic parameters and observable (force or extension) distributions characterizing conformational states in single-molecule force spectroscopy. By use of a Bayesian extension of hidden Markov models, we are able to characterize the experimental uncertainty in these parameters due to instrument noise and finite-

size datasets. The use of a detailed balance constraint additionally helps reduce the experimental uncertainty over standard hidden Markov models, as both transitions into and out of conformational states provide valuable information about state kinetics and populations in data-poor conditions [43, 44]. Additionally, the Gibbs sampling framework used to sample from the Bayesian posterior can be easily extended to incorporate additional nuisance parameters, such as stochastic models of instrument drift or laser power fluctuations.

We have opted to make use of a reversible transition matrix to describe the statistical kinetic behavior between the observation intervals Δt , but it is possible to use a reversible rate matrix instead by substituting a rate matrix sampling scheme [58] in the appropriate stage of the Gibbs sampling updates.

While the experimenter must currently choose the number of conformational states by hand, a number of extensions of Bayesian hidden Markov models can be used to automatically determine the number of states best supported by the data, including reversible-jump schemes [59, 60] and variational Bayes methods [61, 62].

We note that the experimenter in principle has access to the full posterior distribution of models given the observed data, so that instead of looking at the confidence of single parameters, confidence intervals in more complex functions of parameters—such as the rates or lifetimes in Table II—can be computed, or joint posterior distributions of multiple parameters examined. It is also possible to generate *synthetic* data from the current model, or family of models, to examine how the collection of additional data will further reduce uncertainties or allow discrimination among particular hypotheses. The field of *Bayesian experimental design* [63] holds numerous possibilities for selecting how future experiments can maximize information gain, and whether the information gain from the collection of additional data will be of sufficient utility to justify the expense.

ACKNOWLEDGMENTS

The authors thank Sergio Bacallado (Stanford University) for helpful feedback on this manuscript, and Steve Presse (UCSF) for engaging discussions on this topic. JDC acknowledges support from a QB3-Berkeley Distinguished Postdoctoral Fellowship. FN acknowledges DFG Grant 825/2. This work was supported in part by a grant from the NSF (SM).

* Corresponding author

- [1] F. Ritort, J. Phys.: Condens. Matter, **18**, R531 (2006).
- [2] K. C. Neuman and S. M. Block, Rev. Sci. Instrum., **75**, 2787 (2004).
- [3] J. R. Moffitt, Y. R. Chemla, S. B. Smith, and C. Bustamante, Annu. Rev. Biochem., **77**, 205 (2008).
- [4] C. Bustamante, Z. Bryant, and S. B. Smith, Nature, **421**, 423 (2003).
- [5] M. T. Woodside, C. García-Garcá, and S. M. Block, Curr. Opin. Chem. Biol., **12**, 640 (2008).
- [6] C. Bustamante, W. Cheng, and Y. X. Mejia, Cell, **144**, 480 (2011).
- [7] S. B. Fowler, R. B. Best, J. L. T. Herrera, T. J. Rutherford, A. Steward, E. Paci, M. Karplus, and J. Clarke, J. Mol. Biol., **322**, 841 (2002).
- [8] C. Cecconi, E. A. Shank, C. Bustamante, and S. Marqusee, Science, **309**, 2057 (2005).
- [9] J. C. M. Gebhardt, T. Bornschlöggl, and M. Rief, Proc. Natl. Acad. Sci. USA, **107**, 2013 (2010).
- [10] F. Ritort, Adv. Chem. Phys., **137**, 31 (2008).
- [11] C. Schütte and W. Huisinga, in *Handbook of Numerical Analysis - special volume on computational chemistry*, Vol. X, edited by P. G. Ciarlet and J.-L. Lions (Elsevier, 2002).
- [12] J. D. Chodera, N. Singhal, W. C. Swope, V. S. Pande, and K. A. Dill, J. Chem. Phys., **126**, 155101 (2007).
- [13] F. Noé, C. Schütte, E. Vanden-Eijnden, L. Reich, and T. R. Weikl, Proc. Natl. Acad. Sci. USA, **106**, 19011 (2009).
- [14] R. B. Best, E. Paci, G. Hummer, and O. K. Dudko, J. Phys. Chem. B, **112**, 5968 (2008).
- [15] Z. T. Yew, M. Schlierf, M. Rief, and E. Paci, Phys. Rev. E, **81**, 031923 (2010).
- [16] G. Morrison, C. Hyeon, M. Hinczewski, and D. Thirumalai, Phys. Rev. Lett., **106**, 138102 (2011).
- [17] J. D. Chodera and V. S. Pande, arXiv preprint, arXiv:1105.0710 (2011).
- [18] M. T. Woodside, P. C. Anthony, W. M. Behnke-Parks, K. Larizadeh, D. Herschlag, and S. M. Block, Science, **314**, 1001 (2006).
- [19] N. Forns, S. de Lorenzo, M. Manosas, J. M. Huguet, and F. Ritort, Biophys. J., **100**, 1765 (2011).
- [20] J. D. Chodera, P. Elms, W. C. Swope, F. Noé, and V. S. Pande, in preparation (2011).

- [21] L. R. Rabiner, Proceedings of the IEEE, **77**, 257 (1989).
- [22] J. D. Becker, J. Honerkamp, J. Hirsch, U. Fröbe, E. Schlatter, and R. Greger, Pflügers Arch., **426**, 328 (1994).
- [23] F. Qin, A. Auerbach, and F. Sachs, Biophys. J., **79**, 1915 (2000).
- [24] M. C. M. De Gunst, H. R. Künsch, and J. G. Schouten, J. Am. Stat. Assoc., **96**, 805 (2001).
- [25] F. Qin, Biophys. J., **86**, 1488 (2004).
- [26] M. Andrec, R. M. Levy, and D. S. Talaga, J. Phys. Chem. A, **107**, 7454 (2003).
- [27] S. McKinney, C. Joo, and T. Ha, Biophys. J., **91**, 1941 (2006).
- [28] T.-H. Lee, J. Phys. Chem. B, **113**, 11535 (2009).
- [29] Y. Li, X. Qu, A. Ma, G. J. Smith, N. F. Scherer, and A. R. Dinner, J. Phys. Chem. B, **113**, 7579 (2009).
- [30] I. V. Gopich and A. Szabo, J. Phys. Chem. B, **113**, 10965 (2009).
- [31] D. A. Smith, W. Steffen, R. M. Simmons, and J. Sleep, Biophys. J., **81**, 2795 (2001).
- [32] L. S. Miclescu, A. Yildiz, P. R. Selvin, and F. Sachs, Biophys. J., **91**, 3135 (2006).
- [33] F. E. Müllner, S. Syed, P. R. Sevin, and F. J. Sigworth, Biophys. J., **99**, 3684 (2010).
- [34] L. E. Baum, T. Petrie, G. Soules, and N. Weiss, Ann. Math. Statist., **41**, 164 (1970).
- [35] A. J. Viterbi, IEEE Trans. Info. Theory, **13**, 260 (1967).
- [36] M. Kruithof and J. van Noort, Biophys. J., **96**, 3708 (2009).
- [37] T. Aittokallio and E. Uusipaikka, *Computation of standard errors for maximum-likelihood estimates in hidden Markov models*, Tech. Rep. (University of Turku, 2008) technical Report No. 379.
- [38] B. Merialdo, IEEE Trans. Speech and Audio Proc., **1**, 255 (1993).
- [39] C. P. Robert, G. Celeux, and J. Diebolt, Stat. Prob. Lett., **16**, 77 (1993).
- [40] S. Chib, J. Econometrics, **75**, 79 (1996).
- [41] S. L. Scott, J. Am. Stat. Assoc., **97**, 337 (2002).
- [42] T. Rydén, Bayesian Analysis, **3**, 659 (2008).
- [43] F. Noé, J. Chem. Phys., **128**, 244103 (2008).
- [44] P. Metzner, F. Noé, and C. Schütte, Phys. Rev. E, **80**, 021106 (2009).
- [45] S. Geman and D. Geman, IEEE Trans. Pattern Anal., **6**, 721 (1984).
- [46] J. S. Liu, *Monte Carlo strategies in scientific computing*, 2nd ed. (Springer-Verlag, New York, 2002).
- [47] M. J. Comstock, T. Ha, and Y. R. Chemla, Nature Methods, **8**, 335 (2011).

- 528 [48] N. G. van Kampen, *Stochastic processes in physics and chemistry*, 2nd ed. (Elsevier, 1997).
- 529 [49] J. A. Bilmes, *A gentle tutorial of the EM algorithm and its application to parameter estimation for*
530 *Gaussian mixture and hidden Markov models*, Tech. Rep. (University of California, Berkeley,
531 1998).
- 532 [50] J.-H. Prinz, H. Wu, M. Sarich, B. Keller, M. Fischbach, M. Held, J. D. Chodera, C. Schütte,
533 and F. Noé, *J. Chem. Phys.*, **134**, 174105 (2011).
- 534 [51] A. P. Dempster, N. M. Laird, and D. B. Rubin, *J. Royal Statist. Soc. B*, **39**, 1 (1977).
- 535 [52] J. D. Chodera and F. Noé, *J. Chem. Phys.*, **133**, 105012 (2010).
- 536 [53] W. K. Hastings, *Biometrika*, **57**, 97 (1970).
- 537 [54] H. Jeffreys, *Proc. Royal Soc. A*, **186**, 453 (1946).
- 538 [55] P. Goyal, in *Bayesian Inference and Maximum Entropy Methods in Science and Engineering*, edited
539 by K. H. Knuth, A. E. Abbas, R. D. Morriss, and J. P. Castle (American Institute of Physics,
540 2005) pp. 366–373.
- 541 [56] J.-D. Wen, M. Manosas, P. T. X. Li, S. B. Smith, C. Bustamante, F. Ritort, and J. Ignacio Tinoco,
542 *Biophys. J.*, **92**, 2996 (2007).
- 543 [57] C. J. Bustamante and S. B. Smith, “Light-force sensor and method for measuring axial optical-
544 trap forces from changes in light momentum along an optic axis,” (2006).
- 545 [58] G. Hummer, *New Journal of Physics*, **7**, 34 (2005).
- 546 [59] C. P. Robert, T. Rydén, and D. M. Titterton, *J. R. Statist. Soc. B*, **62**, 57 (2000).
- 547 [60] M. C. M. De Gunst and B. Schouten, *Bernoulli*, **9**, 373 (2003).
- 548 [61] M. J. Beal, *Variational algorithms for approximate Bayesian inference*, Master’s thesis, University
549 of Cambridge, UK (2003).
- 550 [62] J. E. Bronson, J. Fei, J. M. Hofman, R. L. G. Jr., and C. H. Wiggins, *Biophys. J.*, **97**, 3196 (2009).
- 551 [63] K. Chaloner and I. Verdinelli, *Statist. Sci.*, **10**, 273 (1995).

552 FIGURES

553

FIG. 1. Single-molecule optical trapping configuration. The biomolecule of interest—here, the p5ab RNA hairpin—is tethered to two polystyrene beads by dsDNA handles. The fluctuating force on one bead held in an optical trap is monitored, while the other bead is held suctioned to a micropipette tip. Conformational transitions of the hairpin—such as transitions among the three kinetically metastable states illustrated here—are observed indirectly through motion of the bead in the trap.

FIG. 2. Synthetic force trajectory and inferred state assignments in MLHMM. Observed samples are colored by their hidden state assignments. Dark horizontal lines terminating in triangles to the right denote state means, while lightly colored bands indicate one standard deviation on either side of the state mean. The gray histogram on the right side shows the total observed probability of samples, while the colored peaks show the weighted Gaussian output contribution from each state, and the black outline the weighted sum of the Gaussian output contributions from the HMM states.

FIG. 3. Experimental force trajectory of the p5ab hairpin and MLHMM state assignments. Observed samples are colored by their hidden state assignments. Dark horizontal lines terminating in triangles to the right denote state means, while lightly colored bands indicate one standard deviation on either side of the state mean. The gray histogram on the right side shows the total observed probability of samples, while the colored peaks show the weighted Gaussian output contribution from each state, and the black outline the weighted sum of the Gaussian output contributions from the HMM states.

TABLE I. Estimated mean model parameters and confidence intervals for synthetic timeseries data


| Property | | Estimated Model Parameters | | | |
|--------------------------|------------|----------------------------|---|---|---|
| | | True Value | 1 000 observations | 10 000 observations | 100 000 observations |
| stationary probability | π_1 | 0.308 | 0.228 ^{0.480} _{0.074} | 0.318 ^{0.407} _{0.244} | 0.324 ^{0.355} _{0.292} |
| | π_2 | 0.113 | 0.093 ^{0.172} _{0.042} | 0.124 ^{0.155} _{0.098} | 0.112 ^{0.121} _{0.104} |
| | π_3 | 0.579 | 0.679 ^{0.870} _{0.415} | 0.558 ^{0.648} _{0.455} | 0.564 ^{0.599} _{0.531} |
| transition probability | T_{11} | 0.980 | 0.970 ^{0.987} _{0.945} | 0.972 ^{0.978} _{0.966} | 0.979 ^{0.981} _{0.978} |
| | T_{12} | 0.019 | 0.023 ^{0.045} _{0.009} | 0.026 ^{0.032} _{0.021} | 0.020 ^{0.021} _{0.018} |
| | T_{13} | 0.001 | 0.007 ^{0.018} _{0.001} | 0.002 ^{0.003} _{0.001} | 0.001 ^{0.001} _{0.001} |
| | T_{21} | 0.053 | 0.054 ^{0.106} _{0.018} | 0.067 ^{0.082} _{0.053} | 0.057 ^{0.061} _{0.052} |
| | T_{22} | 0.900 | 0.868 ^{0.931} _{0.790} | 0.890 ^{0.907} _{0.870} | 0.897 ^{0.903} _{0.892} |
| | T_{23} | 0.050 | 0.078 ^{0.136} _{0.035} | 0.043 ^{0.056} _{0.033} | 0.046 ^{0.050} _{0.042} |
| | T_{31} | 0.001 | 0.002 ^{0.006} _{0.000} | 0.001 ^{0.002} _{0.000} | 0.001 ^{0.001} _{0.000} |
| | T_{32} | 0.009 | 0.010 ^{0.019} _{0.004} | 0.010 ^{0.012} _{0.007} | 0.009 ^{0.010} _{0.008} |
| | T_{33} | 0.990 | 0.988 ^{0.995} _{0.978} | 0.990 ^{0.992} _{0.987} | 0.990 ^{0.991} _{0.989} |
| state mean force (pN) | μ_1 | 3.000 | 2.947 ^{3.082} _{2.812} | 2.998 ^{3.033} _{2.963} | 3.001 ^{3.013} _{2.990} |
| | μ_2 | 4.700 | 4.666 ^{4.721} _{4.612} | 4.699 ^{4.716} _{4.683} | 4.702 ^{4.707} _{4.696} |
| | μ_3 | 5.600 | 5.597 ^{5.614} _{5.583} | 5.602 ^{5.607} _{5.596} | 5.602 ^{5.603} _{5.600} |
| state std dev force (pN) | σ_1 | 1.000 | 1.037 ^{1.134} _{0.951} | 0.992 ^{1.018} _{0.967} | 0.999 ^{1.007} _{0.991} |
| | σ_2 | 0.300 | 0.254 ^{0.300} _{0.217} | 0.287 ^{0.300} _{0.275} | 0.301 ^{0.305} _{0.296} |
| | σ_3 | 0.200 | 0.200 ^{0.211} _{0.190} | 0.203 ^{0.207} _{0.199} | 0.201 ^{0.203} _{0.200} |

TABLE II. BHMM model estimates for p5ab hairpin data.





| Property | Value | |
|---|----------|---|
| Equilibrium probability | π_1 | 0.215 ^{0.236} _{0.193} |
| | π_2 | 0.046 ^{0.050} _{0.041} |
| | π_3 | 0.740 ^{0.762} _{0.717} |
| Transition probability ($\Delta t = 1$ ms) | T_{11} | 0.954 ^{0.959} _{0.950} |
| | T_{12} | 0.033 ^{0.037} _{0.029} |
| | T_{13} | 0.013 ^{0.015} _{0.011} |

| | | |
|------------------------------|------------|----------------------------|
| | T_{21} | $0.154^{0.169}_{0.139}$ |
| | T_{22} | $0.650^{0.673}_{0.627}$ |
| | T_{23} | $0.196^{0.216}_{0.180}$ |
| | T_{31} | $0.004^{0.004}_{0.003}$ |
| | T_{32} | $0.012^{0.013}_{0.011}$ |
| | T_{33} | $0.984^{0.985}_{0.983}$ |
| State force mean (pN) | μ_1 | $12.549^{12.552}_{12.544}$ |
| | μ_2 | $13.016^{13.027}_{13.006}$ |
| | μ_3 | $13.849^{13.852}_{13.848}$ |
| State force std dev (pN) | σ_1 | $0.210^{0.213}_{0.207}$ |
| | σ_2 | $0.201^{0.208}_{0.193}$ |
| | σ_3 | $0.213^{0.214}_{0.211}$ |
| Transition rate (s^{-1}) | k_{12} | $41.4^{46.6}_{36.3}$ |
| | k_{13} | $9.1^{11.3}_{7.2}$ |
| | k_{21} | $194.7^{216.7}_{173.1}$ |
| | k_{23} | $243.7^{271.5}_{219.0}$ |
| | k_{31} | $2.6^{3.2}_{2.1}$ |
| | k_{32} | $15.0^{16.6}_{13.4}$ |
| State mean lifetime (ms) | τ_1 | $21.9^{24.1}_{20.0}$ |
| | τ_2 | $2.9^{3.1}_{2.7}$ |
| | τ_3 | $63.1^{68.5}_{58.4}$ |

Optical Trap



DNA




RNA






Folded

Partially
Folded

Unfolded

Pipette
Tip



-  = Streptavidin
-  = Anti-Digoxigenin Antibody
-  = Biotin
-  = Digoxigenin
-  = Polystyrene Bead

