

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

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

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

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

³Jason L. Choy Laboratory of Single Molecule Biophysics,

Institute for Quantitative Biosciences, University of California, Berkeley, CA 94720, USA

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

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

⁶University of Chicago, IL 60637, USA

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

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

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

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

(Dated: July 30, 2011)

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.

INTRODUCTION

Recent advances in biophysical measurement have led to an unprecedented ability to monitor the dynamics of single biological macromolecules, such as proteins and nucleic acids [3]. As a new approach to probing the behavior of biological macromolecules, these experiments promise to change the way we study folding, dynamics, catalysis, association, transcription, translation, and motility, providing otherwise-inaccessible information about microscopic kinetics, energetics, mechanism, and the stochastic heterogeneity inherent in these processes. Advances in instrumentation for optical force spectroscopy in particular have produced instruments of extraordinary stability, precision, and temporal resolution [4, 5] that have already demonstrated great utility in the study of biomolecules in the presence of externally perturbative forces [6–8]. Under external force, it becomes possible to stabilize and characterize short-lived conformational states, such as protein folding and unfolding intermediates [9–11].

In a typical single-molecule optical trapping experiment, a protein or nucleic acid is tethered to two polystyrene beads by dsDNA handles that prevent the molecule under study from interacting with the beads (see Figure 1). The handle-biomolecule-handle assembly—referred to as a *fiber*—is associated with the beads through tight noncovalent interactions, with one bead held in an optical trap and the other either suctioned to a micropipette (as in Figure 1) or held in a second optical trap. During an experiment, the position of the bead

within the laser trap is monitored, and either the relative displacement from the trap center or the total force on the bead is recorded, providing a series of displacement or force measurements equally spaced in time, resulting in a timeseries such as the one depicted in Figure 2. The instrument can generally be operated in several modes: a *force ramp* mode, in which the trap is translated rapidly enough to potentially carry the system out of equilibrium; an equilibrium *passive* mode, in which the trap is held fixed; and a constant *force-feedback* mode, in which the trap is continually repositioned to maintain a set constant force on the fiber. Here, we concern ourselves with the latter two classes of experiment, though nonequilibrium experiments remain an exciting topic of current research [12].

Often, the dynamics observed in these experiments appears to be dominated by stochastic interconversions between two or more strongly metastable conformational states [13, 14]—regions of conformation space in which the system remains for long times before making a transition to another conformational state. These transitions are generally well-described by first-order kinetics [15]. 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 [16–19], 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 [20, 21], can of-




FIG. 1. **Single-molecule optical trapping configuration for p5ab RNA hairpin.** 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.

ten 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 [22] (see *Supporting Information*).

Hidden Markov models (HMMs) [23], 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 [24–27], single-molecule FRET [28–32], and the stepping of motor proteins [33–35]. 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 [36, 37], as demonstrated in recent work [38].

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 nonlinear ways. While methods exist for estimating the standard error or confidence intervals from MLHMMs [39], 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 [40].

Here, we resolve this issue through the use of a *Bayesian* extension of hidden Markov models [41–44] 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—which has been shown to significantly reduce statistical uncertainties in data-poor conditions [45, 46]. The framework we present is based on Gibbs sampling [47, 48], 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 configuration 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 necessarily 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 [49]) 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

169 states may not necessarily represent highly populated states
 170 of the system at equilibrium. We treat these conformational
 171 states as the *hidden states* of the model, because we cannot *di-*
 172 *rectly* observe the identity of the metastable state in which the
 173 system resides. The hidden Markov model presumes the ob-
 174 served data $\mathbf{O} \equiv \{o_t\}$ was generated according to the follow-
 175 ing model dependent on parameters $\Theta \equiv \{\mathbf{T}, \mathbf{E}\}$, where \mathbf{T}
 176 is an $M \times M$ row-stochastic transition matrix and \mathbf{E} a set of
 177 emission parameters governing the observable (force or exten-
 178 sion) distributions for each of the M hidden states, and prior
 179 information about the initial state distribution ρ ,

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

180 In diagrammatic form, the observed state data $\{o_t\}$ and corre-
 181 sponding hidden state history $\{s_t\}$ can be represented

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

182 The initial state distribution ρ reflects our knowledge of the
 183 initial conditions of the experiment that collected data \mathbf{o} . In the
 184 case that the experiment was prepared in equilibrium, ρ cor-
 185 responds to the equilibrium distribution π of the model tran-
 186 sition matrix \mathbf{T} ; if the experiment was prepared out of equilib-
 187 rium, ρ may be chosen to reflect some other prior distribution
 188 (e.g. the uniform prior).

189 State transitions ($s_{t-1} \rightarrow s_t$) are governed by the discrete
 190 transition probability $T_{s_{t-1}s_t}$. The *Markov property* of HMMs
 191 prescribes that the probability that a system originally in state
 192 i at time t is later found in state j at time $t+1$ is dependent only
 193 on knowledge of the state i , and given by the corresponding
 194 matrix element T_{ij} of the (row-stochastic) transition matrix \mathbf{T} .
 195 Alternatively, one could instead use the rate matrix \mathbf{K} , related
 196 to the transition matrix \mathbf{T} through the equation $\mathbf{T} = e^{\mathbf{K}\Delta t}$.
 197 If the processes described by \mathbf{T} or \mathbf{K} are slow compared to
 198 the observation interval Δt , then we can easily estimate the
 199 rate matrix from the associated transition matrix in a way that
 200 avoids the matrix logarithm, through the expansion $\mathbf{K} \approx (\mathbf{T} -$
 201 $\mathbf{I})/\Delta t$, where \mathbf{I} denotes the $M \times M$ identity matrix.

202 The probabilistic “emission” of observables from each state
 203 ($s_t \rightarrow o_t$) is governed by the continuous emission probability
 204 $\varphi(o_t | \mathbf{e}_{s_t})$, parametrized by observable *emission* parameters \mathbf{e} .
 205 For example, in the force spectroscopy applications described
 206 here, $\varphi(o | \mathbf{e}_s)$ is taken to be a univariate normal (Gaussian)
 207 distribution parameterized by a mean μ and variance σ^2 that
 208 characterize each state, such that $\mathbf{e}_i \equiv \{\mu_i, \sigma_i^2\}$. Other choices
 209 of observable distribution can easily be substituted in a mod-
 210 ular way without affecting the structure of the algorithms pre-
 211 sented here.

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

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

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

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

215 If multiple independent traces $\{o_t\}$ are available, the probabil-
 216 ity $P(\mathbf{O} | \Theta)$ is simply the product of Eq. 3 for the independent
 217 traces.

218 Maximum likelihood hidden Markov model (MLHMM)

219 The standard approach to construct an HMM from ob-
 220 served data is to compute the *maximum likelihood estimator*
 221 (MLE) for the model parameters $\Theta \equiv \{\mathbf{T}, \mathbf{E}\}$, which maxi-
 222 mize the probability of the observed data \mathbf{O} given the model,

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

223 yielding MLE estimates of transition matrix $\hat{\mathbf{T}}$ and state emis-
 224 sion parameters $\hat{\mathbf{E}}$. Typically, determination of the model
 225 parameters Θ is carried out using the Baum-Welch algo-
 226 rithm [36].

227 Once the MLE parameters $\hat{\Theta}$ are determined, the most
 228 likely hidden state history that produced the observations \mathbf{O}
 229 can be determined using these parameters:

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

230 This is typically carried out using the Viterbi algorithm [37], a
 231 classic example of dynamic programming.

232 Both the Baum-Welch and Viterbi schemes are described in
 233 detail in *Algorithms*.

234 Bayesian hidden Markov model (BHMM)

235 Instead of simply determining the model that maximizes
 236 the likelihood of observing the data \mathbf{O} given the model param-
 237 eters Θ , we can make use of Bayes’ theorem to compute the
 238 *posterior* distribution of model parameters given the observed
 239 data:

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

240 Here, $P(\Theta)$ denotes a *prior* distribution that encodes any *a pri-*
 241 *ori* information we may have about the model parameters Θ .
 242 This prior information might include, for example, physical
 243 constraints (such as ensuring the transition matrix satisfies de-
 244 tailed balance) or prior rounds of inference from other inde-
 245 pendent experiments.

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

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

248 Drawing samples of Θ from this distribution will, in princi-
 249 ple, allow the *confidence* with which individual parameters and
 250 combinations thereof are known, given the data (and subject
 251 to the validity of the model of Eq. 1 in correctly representing
 252 the process by which the observed data is generated). While
 253 the posterior $P(\Theta | \mathbf{O})$ is complex, we could in principle use
 254 a Markov chain Monte Carlo (MCMC) approach [48] to sam-
 255 ple it. In its current form, however, this would be extremely
 256 expensive due to the sum over all hidden state histories \mathbf{S}
 257 appearing in ratios involving Eq. 8. Instead, we introduce

258 the hidden state histories \mathbf{S} as an auxiliary variable, sampling
259 from the augmented posterior,

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

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

262 If we presume the prior is separable, such that $P(\Theta) \equiv$
263 $P(\mathbf{T})P(\mathbf{E})$, we can sample from the augmented posterior
264 (Eq. 9) using the framework of *Gibbs sampling* [48], in which the
265 augmented model parameters are updated by sampling from
266 the conditional distributions,

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

267 The equalities on the second and third lines reflect the condi-
268 tional independence of the hidden Markov model defined
269 by Eq. 1. When only the model parameters $\Theta \equiv \{\mathbf{T}, \mathbf{E}\}$
270 or the hidden state histories \mathbf{S} are of interest, we can sim-
271 ply marginalize out the uninteresting variables by sampling
272 from the augmented joint posterior for $\{\mathbf{T}, \mathbf{E}, \mathbf{S}\}$ and examine
273 only the variables of interest. In addition, the structure of the
274 Gibbs sampling scheme above allows individual components
275 (such as the observable distribution model $\varphi(o \mid \mathbf{e})$ or transi-
276 tion probability matrix \mathbf{T}) to be modified without affecting the
277 structure of the remainder of the calculation.

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

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

280 where μ is the mean force or extension characterizing a particu-
281 lar state, and σ is the standard deviation or width of forces or
282 extensions corresponding to that state. We note that marginal
283 posterior distributions of each mean $P(\mu_i \mid \mathbf{O})$ reflect the statis-
284 tical uncertainty in how well the mean force or position is
285 determined, and need not correspond to the standard devia-
286 tion σ_i , which may be much broader (or narrower, depending
287 on the situation).

288 ALGORITHMS

289 Generating an initial model

290 To initialize either computation of the MLHMM or sam-
291 pling from the posterior for the BHMM, an initial model that
292 respects any constraints imposed in the model prior $P(\Theta)$
293 must be selected. Here, we employ a Gaussian observable dis-
294 tribution model for $\varphi(o \mid \mathbf{e})$,

$$\varphi(o \mid \mathbf{e}) = \varphi(o \mid \mu, \sigma^2) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left[-\frac{1}{2} \frac{(o - \mu)^2}{\sigma^2}\right], \quad (12)$$

295 and enforce that the transition matrix \mathbf{T} satisfy detailed bal-
296 ance¹.

Observable parameter estimation

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

$$P(\mathbf{O} \mid \boldsymbol{\pi}, \mathbf{E}) = \prod_{t=0}^L \sum_{m=1}^M \pi_m \varphi(o_t \mid \mu_m, \sigma_m^2), \quad (13)$$

301 where the state observable emission probability vector $\mathbf{E} \equiv$
302 $\{\mathbf{e}_1, \dots, \mathbf{e}_M\}$ and $\mathbf{e}_m \equiv \{\mu_m, \sigma_m^2\}$ with μ_m denoting the ob-
303 servable mean and σ_m^2 the variance for state m for the Gaus-
304 sian mixture model. The vector $\boldsymbol{\pi}$ is composed of equilibrium
305 state populations $\{\pi_1, \dots, \pi_M\}$ with $\pi_m \geq 0$ and $\sum_{m=1}^M \pi_m =$
306 1.

307 A first approximation to $\boldsymbol{\pi}$ and \mathbf{E} is computed by pooling
308 and sorting the observed o_t , and defining M indicator func-
309 tions $h_m(o)$ that separate the data into M contiguous regions
310 of the observed range of o of roughly equal population. Let
311 $N_m \equiv \sum_{t=0}^L h_m(o_t)$ denote the total number of observations
312 falling in region m , and $N_{\text{tot}} = \sum_{m=1}^M N_m$. The initial param-
313 eters are then computed as,

$$\begin{aligned} \pi_m &= N_m / N_{\text{tot}} \\ \mu_m &= N_m^{-1} \sum_{t=0}^L o_t h_m(o_t) \end{aligned} \quad (14)$$

$$\sigma_m^2 = N_m^{-1} \sum_{t=0}^L (o_t - \mu_m)^2 h_m(o_t). \quad (15)$$

314 This approximation is then improved upon by utilizing the
315 expectation-maximization procedure described by Bilmes [51],

$$\begin{aligned} \pi'_m &= N_{\text{tot}}^{-1} \sum_{t=0}^L \chi_m(o_t, \mathbf{E}, \boldsymbol{\pi}) \\ \mu'_m &= (\pi'_m N_{\text{tot}})^{-1} \sum_{t=0}^L o_t \chi_m(o_t, \mathbf{E}, \boldsymbol{\pi}) \\ \sigma'^2_m &= (\pi'_m N_{\text{tot}})^{-1} \sum_{t=0}^L (o_t - \mu'_m)^2 \chi_m(o_t, \mathbf{E}, \boldsymbol{\pi}) \end{aligned} \quad (16)$$

316 where the function $\chi_m(o, \mathbf{E}, \boldsymbol{\pi})$ is given by the fuzzy member-
317 ship function,

$$\chi_m(o, \mathbf{E}, \boldsymbol{\pi}) = \frac{\pi_m \varphi(o \mid \mathbf{e}_m)}{\sum_{l=1}^M \pi_l \varphi(o \mid \mathbf{e}_l)}. \quad (17)$$

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

¹ Physical systems that are not driven in time by an external force or fed by an energy reservoir should satisfy detailed balance [50], and its use has been shown to provide a large reduction in transition matrix uncertainty in data-poor conditions [45, 46]. Detailed balance specifies that $\pi_i T_{ij} = \pi_j T_{ji}$ for all i, j , where $\boldsymbol{\pi}$ is the equilibrium distribution of the row-stochastic transition matrix \mathbf{T} .

321

Transition matrix estimation

322 Once initial state observable emission parameters \mathbf{E} are de-
 323 termined, an initial transition matrix is estimated using an iter-
 324 ative likelihood maximization approach that enforces detailed
 325 balance [52]. First, a matrix of fractional transition counts
 326 $\mathbf{C} \equiv (c_{ij})$ is estimated using the membership function:

$$c_{ij} = \sum_{t=1}^L \chi_i(o_{t-1}, \mathbf{E}, \boldsymbol{\pi}) \chi_j(o_t, \mathbf{E}, \boldsymbol{\pi}) \quad (18)$$

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

$$x_{ij} = x_{ji} = c_{ij} + c_{ji} \quad (19)$$

328 and a vector of row sums

$$x_{i*} = \sum_{j=1}^M x_{ij}. \quad (20)$$

329 Then, the iterative procedure described in Algorithm 1 of [52]
 330 is applied. For each update iteration, we first update the diag-
 331 onal elements of \mathbf{X} :

$$x'_{ii} = \frac{c_{ii}(x_{i*} - x_{ii})}{c_{i*} - c_{ii}} \quad (21)$$

332 where

$$c_{i*} = \sum_{j=1}^M c_{ij} \quad (22)$$

333 followed by the off-diagonal elements:

$$x'_{ij} = x'_{ji} = \frac{-b + \sqrt{b^2 - 4ac}}{2a} \quad (23)$$

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

$$\begin{aligned} a &\equiv c_{i*} - c_{ij} + c_{j*} - c_{ji} \\ b &\equiv c_{i*}(x_{j*} - x_{ji}) + c_{j*}(x_{i*} - x_{ij}) \\ &\quad - (c_{ij} + c_{ji})(x_{i*} - x_{ij} + x_{j*} - x_{ji}) \\ c &\equiv -(c_{ij} + c_{ji})(x_{i*} - x_{ij})(x_{j*} - x_{ji}) \end{aligned} \quad (24)$$

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

$$T_{ij} = \frac{x_{ij}}{x_{i*}}. \quad (25)$$

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

343

Fitting a maximum likelihood HMM

344 The HMM model parameters $\boldsymbol{\Theta} \equiv \{\mathbf{T}, \mathbf{E}\}$ are fit to the ob-
 345 served data \mathbf{O} through use of the expectation-maximization
 346 (EM) algorithm [53]. This is an iterative procedure, where the
 347 model parameters are subsequently refined through succes-
 348 sive iterations. The initial HMM is usually quick to compute,

349 and can give the experimenter a rough idea of the model pa-
 350 rameters, as well as providing a useful starting point for sam-
 351 pling models from the Bayesian posterior.

352 During each iteration, the Baum-Welch algorithm [36] is
 353 used to compute $\boldsymbol{\Xi} \equiv (\xi_{tij})$, which represents the probabil-
 354 ity that the system transitions from hidden state i at time $t - 1$
 355 to hidden state j at time t , and γ_{ti} , the probability that the sys-
 356 tem occupied state i at time t . This is accomplished by first
 357 executing the *forward algorithm*,

$$\alpha_{tj} = \begin{cases} \rho_j \varphi(o_0 | \mathbf{e}_j) & t = 0 \\ \varphi(o_t | \mathbf{e}_j) \sum_{i=1}^M \alpha_{(t-1)i} T_{ij} & t = 1, \dots, L \end{cases} \quad (26)$$

358 followed by the *backward algorithm*,

$$\beta_{ti} = \begin{cases} 1 & t = L \\ \sum_{j=1}^M T_{ij} \varphi(o_{t+1} | \mathbf{e}_j) \beta_{(t+1)j} & t = (L-1), \dots, 0 \end{cases} \quad (27)$$

359 The $L \times M \times M$ matrix $\boldsymbol{\Xi}$ is then computed for $t = 0, \dots, (L-1)$
 360 as,

$$\xi_{tij} = \alpha_{ti} \varphi(o_{t+1} | \mathbf{e}_i) T_{ij} \beta_{(t+1)j} / \sum_{i=1}^M \alpha_{Ti} \quad (28)$$

$$\gamma_{ti} = \sum_{j=1}^M \xi_{tij} \quad (29)$$

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

363 The aggregate matrix of expected transition counts
 364 $\mathbf{C} \equiv (c_{ij})$ is then computed from $\boldsymbol{\Xi}$ as,

$$c_{ij} = \sum_{t=0}^{L-1} \xi_{tij}. \quad (30)$$

365 This count matrix is used to update the maximum-likelihood
 366 transition matrix \mathbf{T} using the method of Prinz et al. [52] de-
 367 scribed in the previous section.

368 The state observable distribution parameters \mathbf{E} are then up-
 369 dated from the γ_{ti} . For the univariate normal distribution ap-
 370 plied to force spectroscopy data here, we update the mean μ_i
 371 and variance σ_i^2 for state i using the scheme,

$$\mu'_i = \frac{\sum_{t=0}^L o_t \gamma_{ti}}{\sum_{t=0}^L \gamma_{ti}}; \quad \sigma'^2_i = \frac{\sum_{t=0}^L (o_t - \mu'_i)^2 \gamma_{ti}}{\sum_{t=0}^L \gamma_{ti}}. \quad (31)$$

372 Once the model parameters have been fitted by iteration of
 373 the above update procedure to convergence (which may only
 374 converge to a local maximum of the likelihood), the most likely
 375 hidden state sequence can be determined given the observa-
 376 tions \mathbf{O} and the MLE model $\hat{\boldsymbol{\Theta}}$ using the Viterbi algorithm [37].
 377 Like the forward-backward algorithm employed in the Baum-
 378 Welch procedure, the Viterbi algorithm also has a forward re-
 379 cursion component,

$$\epsilon_{jt} = \begin{cases} \rho_j \varphi(o_t | \mathbf{e}_j) & t = 0 \\ \varphi(o_t | \mathbf{e}_j) \max_i \epsilon_{i(t-1)} T_{ij} & t = 1, \dots, L \end{cases} \quad (32)$$

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

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

$$\hat{s}_t = \begin{cases} \arg \max_i \epsilon_{it} & t = L \\ \Phi_{\hat{s}_{t+1}(t+1)} & t = (L-1), \dots, 0 \end{cases} \quad (33)$$

382 Sampling from the posterior of the BHMM

383 Sampling from the posterior of the BHMM (Eq. 8) proceeds
384 by rounds of Gibbs sampling, where each round consists of
385 an update of the augmented model parameters $\{\mathbf{T}, \mathbf{E}, \mathbf{S}\}$ by
386 sampling

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

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

388 Updating the hidden state sequences

389 We use a modified form of the Viterbi process to generate
390 an independent sample of the hidden state history \mathbf{S} given the
391 transition probabilities \mathbf{T} , state observable distribution param-
392 eters \mathbf{E} , and observed data \mathbf{O} . Like the Viterbi scheme, a for-
393 ward recursion is applied to each observation trace \mathbf{o} , but in-
394 stead of computing the most *likely* state history on the reverse
395 pass, a new hidden state history \mathbf{S} is drawn from the distri-
396 bution $P(\mathbf{S} | \mathbf{O}, \mathbf{T}, \mathbf{E})$. The forward recursion uses the same
397 forward algorithm as used in Baum-Welch [36],

$$\alpha_{tj} = \begin{cases} \rho_j \varphi(o_0 | \mathbf{e}_j) & t = 0 \\ \varphi(o_t | \mathbf{e}_j) \sum_{i=1}^M \alpha_{(t-1)i} T_{ij} & t = 1, \dots, L \end{cases} \quad (34)$$

398 In the reverse recursion, we now *sample* a state sequence by
399 sampling each hidden state from the conditional distribution
400 $s_t \sim P(s_t | s_{t+1}, \dots, s_L)$ starting from $t = L$ and proceeding
401 down to $t = 0$, where the conditional distribution is given by,

$$P(s_t = i | s_{t+1}, \dots, s_L) \propto \begin{cases} \alpha_{ti} / \sum_{j=1}^M \alpha_{tj} & t = L \\ \alpha_{ti} T_{is_{t+1}} / \sum_{j=1}^M \alpha_{tj} T_{js_{t+1}} & t = (L-1), \dots, 0 \end{cases} \quad (35)$$

402 It is straightforward to show the result of these sampling
403 steps reconstitutes the probability distribution $P(\mathbf{S} | \mathbf{T}, \mathbf{E}, \mathbf{O})$
404 (see *Supplementary Material*).

405 Updating the transition probabilities

406 If no detailed balance constraint is used and the prior $P(\mathbf{T})$
407 is Dirichlet in each row of the transition matrix \mathbf{T} , it is possi-
408 ble to generate an independent sample of the transition matrix
409 from the conditional distribution $P(\mathbf{T}' | \mathbf{S}')$ by sampling each
410 row of the transition matrix from the conjugate Dirichlet pos-
411 terior using the transition counts from the sampled state se-
412 quence \mathbf{S}' [45]. However, because physical systems in the ab-
413 sence of energy input through an external driving force should

414 satisfy detailed balance, we make use of this constraint in up-
415 dating our transition probabilities, since this has been demon-
416 strated to substantially reduce parameter uncertainty in the
417 data-limited regime [45].

418 The transition matrix is updated using the reversible transi-
419 tion matrix sampling scheme of Noé [45, 54]. Here, an adjusted
420 count matrix $\mathbf{C} \equiv (c_{ij})$ is computed using the updated hidden
421 state sequence \mathbf{S}' ,

$$c_{ij} = b_{ij} + \sum_{t=1}^L \delta_{is_{t-1}} \delta_{js_t}, \quad (36)$$

422 where the Kronecker $\delta_{ij} = 1$ if $i = j$ and zero otherwise, and
423 $\mathbf{B} \equiv (b_{ij})$ is a matrix of prior pseudocounts, which we take
424 to be zero following the work of Noé et al. [15]. Using the
425 adjusted count matrix \mathbf{C} , a Metropolis-Hastings Monte Carlo
426 procedure [55] is used to update the matrix and produce a new
427 sample from $P(\mathbf{T}' | \mathbf{S}')$. Two move types are attempted, se-
428 lected with equal probability, and 1000 moves are attempted to
429 generate a new sample \mathbf{T}' that is approximately uncorrelated
430 from the previous \mathbf{T} . Prior to starting the Monte Carlo proce-
431 dure, the vector of equilibrium probabilities for all states $\boldsymbol{\pi}$ is
432 computed according to

$$\mathbf{T}^T \boldsymbol{\pi} = \boldsymbol{\pi}. \quad (37)$$

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:

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

433 This move is accepted with probability

$$P_{\text{accept}}(\mathbf{T} \rightarrow \mathbf{T}') = \min \left\{ 1, \sqrt{\frac{(T'_{ij})^2 + (T'_{ji})^2}{(T_{ij})^2 + (T_{ji})^2}} \right. \\ \left. \times \left(\frac{T'_{ii}}{T_{ii}} \right)^{c_{ii}} \left(\frac{T'_{ij}}{T_{ij}} \right)^{c_{ij}} \left(\frac{T'_{jj}}{T_{jj}} \right)^{c_{jj}} \left(\frac{T'_{ji}}{T_{ji}} \right)^{c_{ji}} \right\} \quad (38)$$

434 This move will leave the vector of stationary probabilities $\boldsymbol{\pi}$
435 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 uni-
formly over the interval

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

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

$$T'_{ij} = \begin{cases} \eta T_{ij} & j = 1, \dots, M, j \neq i \\ \eta(T_{ii} - 1) + 1 & j = i \end{cases} \quad (39)$$

437 This move is accepted with probability

$$P_{\text{accept}}(\mathbf{T} \rightarrow \mathbf{T}') = \min \left\{ 1, \eta^{(M-2)} \eta^{(c_{i*} - c_{ii})} \left(\frac{1 - \eta(1 - T_{ii})}{T_{ii}} \right)^{c_{ii}} \right\} \quad (40)$$

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

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

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

Updating the observable distribution parameters

Following the update of the transition matrix \mathbf{T} , the observable distribution parameters \mathbf{E} are updated by sampling \mathbf{E} from the conditional probability $P(\mathbf{E}' | \mathbf{S}', \mathbf{O})$. The conditional probability for the observable distribution parameters for state m , denoted \mathbf{e}_m , is given in terms of the output model $\varphi(o | \mathbf{e})$ by Bayes' theorem,

$$P(\mathbf{E} | \mathbf{O}, \mathbf{S}) = \left[\prod_{t=0}^L \varphi(o_t | \mathbf{e}_{s_t}) \right] P(\mathbf{E}). \quad (41)$$

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

$$P(\mathbf{E}) = \prod_{m=1}^M P(\mathbf{e}_m), \quad (42)$$

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

Here, we make the choice that the prior be separable (Eq. 42), which has the benefit of allowing the conditional probability for \mathbf{E} (Eq. 41) to be decomposed into a separate posterior for each state. For each state m , collect all the observations o_t whose updated hidden state labels $s_t' = m$ into 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 purposes of this update procedure. Then, the observable parameters \mathbf{e} for this state are given by

$$P(\mathbf{e} | \mathbf{o}) = P(\mathbf{o} | \mathbf{e})P(\mathbf{e}) = \left[\prod_{n=1}^{N_m} \varphi(o_n | \mathbf{e}) \right] P(\mathbf{e}). \quad (43)$$

In the application presented here, we use a Gaussian output model (Eq. 12) for the state observable distributions $P(o | \mathbf{e})$, where $\mathbf{e} \equiv \{\mu, \sigma^2\}$, with μ the state mean observable and σ^2 the variance (which will include both the distribution of the observable characterizing the state and any broadening from measurement noise). Other models (including multidimensional or multimodal observation models) are possible, and require replacing only the observation model $\varphi(o | \mathbf{e})$ and corresponding prior $P(\mathbf{e})$.

We use the (improper) Jeffreys prior [56] which has the information-theoretic interpretation as the prior that maximizes the information content of the data [57], (suppressing the state index subscript m),

$$P(\mathbf{e}) \propto \sigma^{-1}, \quad (44)$$

which produces the posterior

$$P(\mathbf{e} | \mathbf{o}) \propto \sigma^{-(N+1)} \exp \left[-\frac{1}{2\sigma^2} \sum_{n=1}^N (o_n - \mu)^2 \right], \quad (45)$$

where we remind the reader that here and in the remainder of this section, the symbols \mathbf{e} , \mathbf{o} , σ , μ , and N refer to \mathbf{e}_m , \mathbf{o}_m , σ_m , μ_m , and N_m , respectively.

Updating $\{\mu, \sigma^2\}$ also proceeds by a Gibbs sampling scheme, alternately updating μ and σ , as earlier described in Ref. [54],

$$\begin{aligned} \mu &\sim P(\mu | \sigma^2, \mathbf{o}) \\ \sigma^2 &\sim P(\sigma^2 | \mu, \mathbf{o}) \end{aligned} \quad (46)$$

The conditional distribution of the mean μ is then given by

$$\begin{aligned} P(\mu | \sigma^2, \mathbf{o}) &\equiv \frac{P(\mu, \sigma^2 | \mathbf{o})}{\int d\mu P(\mu, \sigma^2 | \mathbf{o})} \\ &\propto \exp \left[-\frac{1}{2(\sigma^2/N)} (\mu - \hat{\mu})^2 \right] \end{aligned} \quad (47)$$

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

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

This allows us to update μ according to

$$\mu' \sim \mathcal{N}(\hat{\mu}, \sigma^2/N) \quad (49)$$

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

$$\begin{aligned} P(\sigma^2 | \mu, \mathbf{o}) &= \frac{p(\mu, \sigma^2 | \mathbf{o})}{\int d\sigma^2 p(\mu, \sigma^2 | \mathbf{o})} \\ &\propto \sigma^{-(N+1)} \exp \left[-\frac{1}{2\sigma^2} \sum_{n=1}^N (o_n - \mu)^2 \right] \\ &\propto \sigma^{-(N+1)} \exp \left[-\frac{N\hat{\sigma}^2}{2\sigma^2} \right] \end{aligned} \quad (50)$$

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

$$\hat{\sigma}^2 \equiv \frac{1}{N} \sum_{n=1}^N (o_n - \mu)^2. \quad (51)$$

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

$$y \sim \chi^2(N - 1) \quad (52)$$

and then update σ^2 as

$$\sigma'^2 = \frac{N\hat{\sigma}^2}{y}. \quad (53)$$

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

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




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.

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

Property		True Value	Estimated Model Parameters		
			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}

VALIDATION USING SYNTHETIC DATA

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

As an example synthetic model, consider the three-state system intended to mimic a protein with (1) a highly-compliance, low-force unfolded state, (2) a moderately compliant low-population intermediate at intermediate force, and (3) a low-compliance, high-force folded state. Here, the term “compli-

ance” refers to the width of the force or extension distribution characterizing the state. Parameters of the model are given in Table I, and the observation interval was taken to be $\tau = 1$ ms. An example realization of a model trajectory, along with the MLHMM state assignment, is shown in Figure 2. We generated a trajectory of 100 000 observations, and characterized the BHMM mean parameter estimate and 95% confidence intervals for a subset of this trajectory of varying lengths. The results, shown in Table I, show that the confidence intervals contract as trajectory length increases, as expected, and the BHMM-computed 95% confidence intervals contain the true model parameters with the expected statistics. In contrast, a model created from simply segmenting the observed forces into disjoint regions and assigning state membership based on the force value alone estimates model parameters with significant bias even for 1 000 000 observations (see *Supporting Information*).

As a more rigorous test, we sampled 50 random models from the prior $P(\Theta)$ with two to six states, generated a 10 000 observation synthetic trajectory for each, and accumulated statistics on the observed fraction of time the true model pa-




FIG. 3. **Validation of confidence intervals using randomly-generated synthetic test data.** Observed confidence intervals (points) are plotted as a function of the desired confidence intervals for equilibrium probabilities (π_i), transition probabilities (T_{ij}), state means (μ_i), and state standard deviations (σ_i). The black diagonal line indicates perfect agreement between expected and observed confidence intervals, while observed confidence intervals above the diagonal indicate overestimates of the uncertainty, and below the diagonal indicate underestimates. Because only 50 random models were evaluated, error bars denote a 95% confidence interval in the estimated observed confidence intervals.

549 rameters were within the BHMM confidence intervals for various values of the confidence interval width α . The results of this test are depicted in Figure 3. We expect that the plot traces the diagonal if the observed and expected confidence intervals are identical; an overestimate of the confidence interval will be above the diagonal, and an underestimate will fall below it. Because only a finite number of independent replicates of the experiment are conducted, there is some associated uncertainty with the observed confidence intervals. The results show that the observed confidence intervals line up with the expected confidence intervals to within statistical error, suggesting the BHMM confidence intervals neither underestimate nor overestimate the actual uncertainty in model parameters.

562 RNA HAIRPIN KINETICS IN A PASSIVE OPTICAL TRAP

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

564 The p5ab RNA hairpin from *Tetrahymena thermophila* was

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}

569 provided by Jin-Der Wen, and prepared as previously described [58]. Within the population of RNA hairpin molecules in the examined sample, there were two chemically distinct species present in the sample (i.e. as a result of post-transcriptional or other covalent modification during sample storage), exhibiting either apparent two-state (as reported previously [58]) or three-state behavior (studied here). For the purposes of testing this method, we examined a fiber that appeared to consistently exhibit three-state behavior upon visual inspection of the force timeseries data.

579 The instrument used in this experiment was a dual-beam counter-propagating optical trap with a spring constant of 0.1 pN/nm. A piezoactuator controlled the position of the trap and allowed position resolution to within 0.5 nm [59]. Drift in the instrument was less than 1 nm/minute resulting in a constant average force within 0.1 pN over the course of a typical 60 s experiment. For these constant trap position experiments, higher frequency data was recorded at 50 kHz recording the voltage corresponding to the force on the tether directly from the position-sensitive detectors. To ensure sequential samples obeyed Markovian statistics, these data were subsampled down to 1 kHz for analysis by the BHMM framework after examination of autocorrelation functions for trap positions where the hairpin appeared to remain in a single conformational state (see *Supplementary Information*).

594 A single observed force trajectory at a fixed trap position adequate to cause hopping among multiple states is shown in Figure 4. The most likely state trajectory from the MLHMM fit with three states is shown by coloring the observations most likely to be associated with each state, with bands of color indi-




FIG. 4. **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.

599 cating the mean and standard deviation about the mean force
600 characterizing each state.

601 Table II lists the BHMM posterior means and confidence
602 intervals characterizing the three-state model extracted from
603 this single 60 s observed force trace. Several things are no-
604 table about the estimated model parameters. Surprisingly,
605 while there is a clearly-resolved intermediate-force state (*state*
606 2) through which most of the flux from the high- and low-force
607 states passes (as seen from large K_{12} and K_{23}), there are non-
608 trivial rate constants connecting the high and low force states
609 directly (K_{13}), indicating that while a sequential mechanism
610 involving passing through the intermediate state is preferred,
611 it may not be an obligatory step in hairpin formation under
612 these conditions. While the state mean forces are clearly distinct,
613 the state standard deviations—which reflect the width
614 of the observed force distribution characterizing each state,
615 rather than the uncertainty in state means—possess overlap-
616 ping confidence intervals. These standard deviations reflect
617 not only contributions from both the distribution of extensions
618 sampled by the hairpin in each conformational state, but also
619 from fluctuations in the handles and beads, and other sources
620 of mechanical and electrical noise in the measurement. As
621 we would expect the unfolded hairpin to be more compliant
622 (i.e. possess a wider distribution of forces) than the folded hair-
623 pin, the inability to distinguish the standard deviations among
624 states is suggestive that, for this experimental configuration
625 and observation time, the predominant contribution to the ob-
626 served force distributions for each state may be in the form of
627 handle or bead fluctuations or other sources of measurement
628 noise.

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

DISCUSSION

636 We have described an approach to determining the first-
637 order kinetic parameters and observable (force or extension)
638 distributions characterizing conformational states in single-
639 molecule force spectroscopy. By use of a Bayesian extension
640 of hidden Markov models, we are able to characterize the ex-

641 perimental uncertainty in these parameters due to instrument
642 noise and finite-size datasets. The use of a detailed balance
643 constraint additionally helps reduce the experimental uncer-
644 tainty over standard hidden Markov models, as both tran-
645 sitions into and out of conformational states provide valu-
646 able information about state kinetics and populations in data-
647 poor conditions. Additionally, the Gibbs sampling framework
648 used to sample from the Bayesian posterior can be easily ex-
649 tended to incorporate additional nuisance parameters, such as
650 stochastic models of instrument drift or laser power fluctua-
651 tions.

652 We have opted to make use of a reversible transition ma-
653 trix to describe the statistical kinetic behavior between the ob-
654 servation intervals Δt , but it is possible to use a reversible
655 rate matrix instead by substituting a rate matrix sampling
656 scheme [60] in the appropriate stage of the Gibbs sampling up-
657 dates.

658 While the experimenter must currently choose the number
659 of conformational states by hand, a number of extensions of
660 Bayesian hidden Markov models can be used to automati-
661 cally determine the number of states best supported by the
662 data, including reversible-jump schemes [61, 62] and varia-
663 tional Bayes methods [63, 64].

664 We note that the experimenter in principle has access to
665 the full posterior distribution of models given the observed
666 data, so that instead of looking at the confidence of single pa-
667 rameters, confidence intervals in more complex functions of
668 parameters—such as the rates or lifetimes in Table II—can be
669 computed, or joint posterior distributions of multiple param-
670 eters examined. It is also possible to generate *synthetic* data
671 from the current model, or family of models, to examine how
672 the collection of additional data will further reduce uncertain-
673 ties or allow discrimination among particular hypotheses. The
674 field of *Bayesian experimental design* [65] holds numerous pos-
675 sibilities for selecting how future experiments can maximize
676 information gain, and whether the information gain from the
677 collection of additional data will be of sufficient utility to jus-
678 tify the expense.

ACKNOWLEDGMENTS

680 The authors thank Sergio Bacallado (Stanford University)
681 for helpful feedback on this manuscript, and Steve Presse
682 (UCSF) for engaging discussions on this topic. JDC acknowl-

683 edges support from a QB3-Berkeley Distinguished Postdoc-
684 toral Fellowship. FN acknowledges DFG Grant 825/2. This
685 work was supported in part by a grant from the NSF (SM).

686 * Corresponding author

- 687 [1]
688 [2] 00, 0. 0.
689 [3] Ritort, F., 2006. Single-molecule experiments in biological
690 physics: methods and applications. *J. Phys.: Condens.*
691 *Matter* 18:R531–R583.
692 [4] Neuman, K. C., and S. M. Block, 2004. Optical trapping.
693 *Rev. Sci. Instrum.* 75:2787.
694 [5] Moffitt, J. R., Y. R. Chemla, S. B. Smith, and C. Bustamante,
695 2008. Recent advances in optical tweezers. *Annu.*
696 *Rev. Biochem.* 77:205–228.
697 [6] Bustamante, C., Z. Bryant, and S. B. Smith, 2003. Ten
698 years of tension: single-molecule DNA mechanics. *Nature*
699 421:423–427.
700 [7] Woodside, M. T., C. García-Garcá, and S. M. Block, 2008.
701 Folding and unfolding single RNA molecules under tension.
702 *Curr. Opin. Chem. Biol.* 12:640–646.
703 [8] Bustamante, C., W. Cheng, and Y. X. Mejia, 2011. Re-
704 visiting the central dogma one molecule at a time. *Cell*
705 144:480–497.
706 [9] Fowler, S. B., R. B. Best, J. L. T. Herrera, T. J. Rutherford,
707 A. Steward, E. Paci, M. Karplus, and J. Clarke, 2002. Mechanical
708 unfolding of a titin Ig domain: Structure of unfolding intermediates
709 revealed by combining AFM, molecular dynamics simulations, NMR and
710 protein engineering. *J. Mol. Biol.* 322:841–849.
711 [10] Cecconi, C., E. A. Shank, C. Bustamante, and S. Marqusee,
712 2005. Direct observation of the three-state folding of a
713 single protein molecule. *Science* 309:2057–2060.
714 [11] Gebhardt, J. C. M., T. Bornschlöggl, and M. Rief, 2010. Full
715 distance-resolved folding energy landscape of one single
716 protein molecule. *Proc. Natl. Acad. Sci. USA* 107:2013–
717 2018.
718 [12] Ritort, F., 2008. Nonequilibrium fluctuations in small systems:
719 From physics to biology. *Adv. Chem. Phys.* 137:31–
720 123.
721 [13] Schütte, C., and W. Huisinga, 2002. Biomolecular conformations
722 can be identified as metastable states of molecular dynamics. In
723 P. G. Ciaret, and J.-L. Lions, editors, *Handbook of Numerical Analysis -*
724 *special volume on computational chemistry*, Elsevier, volume X.
725 [14] Chodera, J. D., N. Singhal, W. C. Swope, V. S. Pande, and
726 K. A. Dill, 2007. Automatic discovery of metastable states for the
727 construction of markov models of macromolecular conformational
728 dynamics. *J. Chem. Phys.* 126:155101.
729 [15] Noé, F., C. Schütte, E. Vanden-Eijnden, L. Reich, and T. R.
730 Weikl, 2009. Constructing the full ensemble of folding pathways
731 from short off-equilibrium simulations. *Proc. Natl. Acad. Sci. USA*
732 106:19011–19016.
733 [16] Best, R. B., E. Paci, G. Hummer, and O. K. Dudko, 2008. Pulling
734 direction as a reaction coordinate for the mechanical unfolding of
735 single molecules. *J. Phys. Chem. B* 112:5968–5976.
736 [17] Yew, Z. T., M. Schlierf, M. Rief, and E. Paci, 2010. Direct
737 evidence of the multidimensionality of the free-energy landscapes of
738 proteins revealed by mechanical probes. *Phys. Rev. E* 81:031923.
739 [18] Morrison, G., C. Hyeon, M. Hinczewski, and D. Thirumalai, 2011.
740 Compaction and tensile forces determine the accuracy of folding
741 landscape parameters from single molecule pulling experiments. *Phys. Rev. Lett.*
742 106:138102.
743 [19] Chodera, J. D., and V. S. Pande, 2011. Splitting probabilities as a
744 test of reaction coordinate choice in single-molecule experiments. *arXiv preprint*
745 arXiv:1105.0710.
746 [20] Woodside, M. T., P. C. Anthony, W. M. Behnke-Parks, K. Larizadeh,
747 D. Herschlag, and S. M. Block, 2006. Direct measurement of the full,
748 sequence-dependent folding landscape of a nucleic acid. *Science*
749 314:1001–1004.
750 [21] Forns, N., S. de Lorenzo, M. Manosas, J. M. Huguet, and F. Ritort,
751 2011. Improving signal/noise resolution in single-molecule experiments
752 using molecular constructs with short handles. *Biophys. J.* 100:1765–
753 1774.
754 [22] Chodera, J. D., P. Elms, W. C. Swope, F. Noé, and V. S. Pande,
755 2011. A robust approach to estimating rates from time-correlation
756 functions. *in preparation*.
757 [23] Rabiner, L. R., 1989. A tutorial on Hidden Markov models and selected
758 applications in speech recognition. *Proceedings of the IEEE* 77:257–
759 286.
760 [24] Becker, J. D., J. Honerkamp, J. Hirsch, U. Fröbe, E. Schlatter,
761 and R. Greger, 1994. Compaction and tensile forces determine the
762 accuracy of folding landscape parameters from single molecule pulling
763 experiments. *Pflügers Arch.* 426:328–332.
764 [25] Qin, F., A. Auerbach, and F. Sachs, 2000. A direct optimization
765 approach to hidden Markov modeling for single channel kinetics. *Biophys. J.*
766 79:1915–1927.
767 [26] De Gunst, M. C. M., H. R. Künsch, and J. G. Schouten, 2001. Statistical
768 analysis of ion channel data using hidden Markov models with correlated
769 state-dependent noise and filtering. *J. Am. Stat. Assoc.* 96:805–815.
770 [27] Qin, F., 2004. Restoration of single-channel currents using the
771 segmental k -means method based on hidden Markov modeling. *Biophys. J.*
772 86:1488–1501.
773 [28] Andrec, M., R. M. Levy, and D. S. Talaga, 2003. Direct determination
774 of kinetic rates from single-molecule photon arrival trajectories using
775 hidden Markov models. *J. Phys. Chem. A* 107:7454–7464.
776 [29] McKinney, S., C. Joo, and T. Ha, 2006. Analysis of single-molecule
777 FRET trajectories using hidden Markov models. *Biophys. J.* 91:1941–
778 1951.
779 [30] Lee, T.-H., 2009. Extracting kinetics information from single-molecule
780 fluorescence resonance energy transfer data using hidden Markov models. *J. Phys. Chem. B*
781 113:11535–11542.
782 [31] Li, Y., X. Qu, A. Ma, G. J. Smith, N. F. Scherer, and A. R. Dinner,
783 2009. Models of single-molecule experiments with periodic perturbations
784 reveal hidden dynamics in RNA folding. *J. Phys. Chem. B* 113:7579–
785 7590.
786 [32] Gopich, I. V., and A. Szabo, 2009. Decoding the pattern of photon
787 colors in single-molecule FRET. *J. Phys. Chem. B* 113:10965–
788 10973.
789 [33] Smith, D. A., W. Steffen, R. M. Simmons, and J. Sleep, 2001. Hidden-
790 Markov methods for the analysis of single-molecule actomyosin displacement
791 data: The variance-hidden-Markov method. *Biophys. J.* 81:2795–
792 2816.
793 [34] Milesescu, L. S., A. Yildiz, P. R. Selvin, and F. Sachs, 2006. Extracting
794 dwell time sequences from processive molecular motor data. *Biophys. J.*
795 91:3135.
796 [35] Müllner, F. E., S. Syed, P. R. Sevin, and F. J. Sigworth, 2010. Improved
797 hidden Markov models for molecular motors,

- part 1: Basic theory. *Biophys. J.* 99:3684–3695.
- [36] Baum, L. E., T. Petrie, G. Soules, and N. Weiss, 1970. A maximization technique occurring in the statistical analysis of probabilistic functions of Markov chains. *Ann. Math. Statist.* 41:164–171.
- [37] Viterbi, A. J., 1967. Error bounds for convolutional codes and an asymptotically optimum decoding algorithm. *IEEE Trans. Info. Theory* 13:260–269.
- [38] Kruithof, M., and J. van Noort, 2009. Hidden Markov analysis of nucleosome unwrapping under force. *Biophys. J.* 96:3708–3715.
- [39] Aittokallio, T., and E. Uusipaikka, 2008. Computation of standard errors for maximum-likelihood estimates in hidden Markov models. Technical report, University of Turku. Technical Report No. 379.
- [40] Merialdo, B., 1993. On the locality of the forward-backward algorithm. *IEEE Trans. Speech and Audio Proc.* 1:255–257.
- [41] Robert, C. P., G. Celeux, and J. Diebolt, 1993. Bayesian estimation of hidden Markov chains: A stochastic implementation. *Stat. Prob. Lett.* 16:77–83.
- [42] Chib, S., 1996. Calculating posterior distributions and modal estimates in Markov mixture models. *J. Econometrics* 75:79–97.
- [43] Scott, S. L., 2002. Bayesian methods for hidden Markov models: Recursive computing in the 21st century. *J. Am. Stat. Assoc.* 97:337–351.
- [44] Rydén, T., 2008. EM versus Markov chain Monte Carlo for estimation of hidden Markov models: A computational perspective. *Bayesian Analysis* 3:659–688.
- [45] Noé, F., 2008. Probability distributions of molecular observables computed from Markov models. *J. Chem. Phys.* 128:244103.
- [46] Metzner, P., F. Noé, and C. Schütte, 2009. Estimating the sampling error: Distribution of transition matrices and functions of transition matrices for given trajectory data. *Phys. Rev. E* 80:021106.
- [47] Geman, S., and D. Geman, 1984. Stochastic relaxation, Gibbs distributions and the Bayesian restoration of images. *IEEE Trans. Pattern Anal.* 6:721–741.
- [48] Liu, J. S., 2002. Monte Carlo strategies in scientific computing. Springer-Verlag, New York, 2nd ed. edition.
- [49] Comstock, M. J., T. Ha, and Y. R. Chemla, 2011. Ultrahigh-resolution optical trap with single-fluorophore sensitivity. *Nature Methods* 8:335–340.
- [50] van Kampen, N. G., 1997. Stochastic processes in physics and chemistry. Elsevier, second edition.
- [51] Bilmes, J. A., 1998. A gentle tutorial of the EM algorithm and its application to parameter estimation for Gaussian mixture and hidden Markov models. Technical report, University of California, Berkeley.
- [52] Prinz, J.-H., H. Wu, M. Sarich, B. Keller, M. Fischbach, M. Held, J. D. Chodera, C. Schütte, and F. Noé, 2011. Markov models of molecular kinetics: Generation and validation. *J. Chem. Phys.* 134:174105.
- [53] Dempster, A. P., N. M. Laird, and D. B. Rubin, 1977. Maximum-likelihood from incomplete data via the EM algorithm. *J. Royal Statist. Soc. B* 39:1–38.
- [54] Chodera, J. D., and F. Noé, 2010. Probability distributions of molecular observables computed from Markov models: II. uncertainties in observables and their time-evolution. *J. Chem. Phys.* 133:105012.
- [55] Hastings, W. K., 1970. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika* 57:97–109.
- [56] Jeffreys, H., 1946. An invariant form for the prior probability in estimation problems. *Proc. Royal Soc. A* 186:453–461.
- [57] Goyal, P., 2005. Prior probabilities: An information-theoretic approach. In K. H. Knuth, A. E. Abbas, R. D. Morriss, and J. P. Castle, editors, Bayesian Inference and Maximum Entropy Methods in Science and Engineering, American Institute of Physics, 366–373.
- [58] Wen, J.-D., M. Manosas, P. T. X. Li, S. B. Smith, C. Bustamante, F. Ritort, and J. Ignacio Tinoco, 2007. Force unfolding kinetics of RNA using optical tweezers. I. Effects of experimental variables on measured results. *Biophys. J.* 92:2996–3009.
- [59] Bustamante, C. J., and S. B. Smith, 2006. Light-force sensor and method for measuring axial optical-trap forces from changes in light momentum along an optic axis.
- [60] Hummer, G., 2005. Position-dependent diffusion coefficients and free energies from Bayesian analysis of equilibrium and replica molecular dynamics simulations. *New Journal of Physics* 7:34.
- [61] Robert, C. P., T. Rydén, and D. M. Titterton, 2000. Bayesian inference in hidden Markov models through the reversible jump Markov chain Monte Carlo method. *J. R. Statist. Soc. B* 62:57–75.
- [62] De Gunst, M. C. M., and B. Schouten, 2003. Model selection for hidden Markov models of ion channel data by reversible Markov chain Monte Carlo. *Bernoulli* 9:373–393.
- [63] Beal, M. J., 2003. Variational algorithms for approximate Bayesian inference. Master’s thesis, University of Cambridge, UK.
- [64] Bronson, J. E., J. Fei, J. M. Hofman, R. L. G. Jr., and C. H. Wiggins, 2009. Learning rates and states from biophysical time series: a Bayesian approach to model selection and single-molecule FRET data. *Biophys. J.* 97:3196–3205.
- [65] Chaloner, K., and I. Verdinelli, 1995. Bayesian experimental design: A review. *Statist. Sci.* 10:273–204.