

## Approximating nonequilibrium processes using a collection of surrogate diffusion models

Christopher P. Calderon<sup>1,a)</sup> and Riccardo Chelli<sup>2,3</sup>

<sup>1</sup>*Department of Statistics and Department of Computational and Applied Mathematics, Rice University, Houston, Texas 77005-1892, USA*

<sup>2</sup>*Dipartimento di Chimica, Università di Firenze, Via della Lastruccia 3, I-50019 Sesto Fiorentino, Italy*

<sup>3</sup>*European Laboratory for Non-Linear Spectroscopy (LENS), Via Nello Carrara 1, I-50019 Sesto Fiorentino, Italy*

(Received 6 August 2007; accepted 7 March 2008; published online 11 April 2008)

The surrogate process approximation (SPA) is applied to model the nonequilibrium dynamics of a reaction coordinate (RC) associated with the unfolding and refolding processes of a deca-alanine peptide at 300 K. The RC dynamics, which correspond to the evolution of the end-to-end distance of the polypeptide, are produced by steered molecular dynamics (SMD) simulations and approximated using overdamped diffusion models. We show that the collection of (estimated) SPA models contain structural information “orthogonal” to the RC monitored in this study. Functional data analysis ideas are used to correlate functions associated with the fitted SPA models with the work done on the system in SMD simulations. It is demonstrated that the shape of the nonequilibrium work distributions for the unfolding and refolding processes of deca-alanine can be predicted with functional data analysis ideas using a relatively small number of simulated SMD paths for calibrating the SPA diffusion models. © 2008 American Institute of Physics.

[DOI: [10.1063/1.2903439](https://doi.org/10.1063/1.2903439)]

### I. INTRODUCTION

Single-molecule experiments<sup>1</sup> offer the possibility of exploring dynamical responses in systems without having to resort to ensemble measurements.<sup>2–6</sup> In recent years, these types of investigations have received increasing attention due partially to developments in nonequilibrium statistical mechanics.<sup>7,8</sup> In the context of protein folding, vastly different dynamical trajectories can be observed when monitoring the nonequilibrium time evolution of a reaction coordinate (RC) obtained from independent single-molecule measurements carried out at identical experimental conditions. The differences in dynamical response cannot be usually attributed only to classical thermal noise, but also to the different atomistic “collective” coordinates<sup>9</sup> (e.g., radius of gyration) consistent with the conditions of the system at the initial time of the various nonequilibrium experiments. In many cases, these types of collective variables strongly influence the nature of observed dynamical response. This type of phenomenon has been quantified by modeling with “multiple states”<sup>10–12</sup> or by using the so-called “dynamic disorder” description.<sup>13</sup> Often, both descriptions can be reasonably thought of as artifacts resulting from an imperfect RC.<sup>12,14</sup>

One concrete example that exhibits the concepts expressed above is the extensively studied<sup>5,6</sup> process of unfolding and refolding of deca-alanine (a decapeptide consisting of alanine residues). In Fig. 1, we report some snapshots of deca-alanine taken from two forced refolding realizations along the end-to-end distance (the RC) produced by steered

molecular dynamics (SMD) simulations.<sup>6</sup> By comparing the time evolution of the RC in the two computational experiments (see panels in the figure), one would be inclined to classify the two nonequilibrium trajectories as (statistically) identical or, at least, very similar. However, when the snapshots of the entire molecule are observed, it is easy to distinguish between the dynamical responses (“proper folding” and “misfolding”). Another instance of this type of phenomenon is documented in Ref. 15. This situation is relevant to a variety of different systems because when all-atom descriptions are used, it is usually difficult to select RCs that accurately reflect all of the relevant details of the system.<sup>5,16,17</sup> Even if good RCs are known and easily manipulatable, *a priori* quantitative knowledge of how the system responds to external stimuli that vary these RCs is usually lacking.

This paper aims at providing this type of information by applying pathwise time series techniques<sup>18,19</sup> to the output of nonequilibrium single-molecule computer simulations. The numerical procedures presented here also have applicability to wet-laboratory times series data.<sup>2,3,20</sup> Nonequilibrium simulations are attractive because they sometimes allow a researcher to obtain detailed equilibrium information about portions of phase space that, if the system were unforced, would not be explored in the typical simulation time scale.<sup>5,16,17,21</sup> The system studied throughout is a SMD simulation of deca-alanine *in vacuo* maintained at constant temperature by a Nosé–Hoover thermostat.<sup>22</sup> We systematically manipulate the end-to-end distance RC of the deca-alanine peptide, referring to it as the “steered” RC.<sup>5</sup>

The purpose of this study is twofold. The first goal is to present a computational approach that attempts to learn about conformational dynamics not directly contained in a

<sup>a)</sup>Electronic mail: calderon@rice.edu.

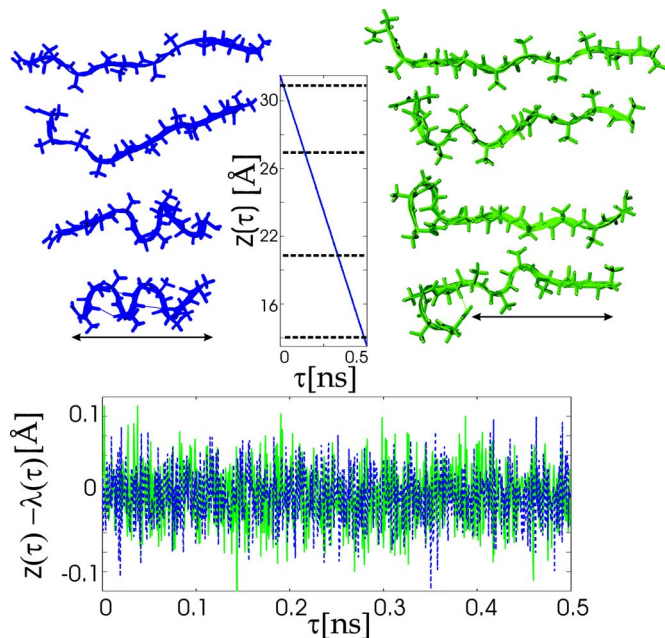


FIG. 1. (Color online) Snapshots of deca-alanine during two independent refolding processes obtained from SMD simulations published in Ref. 6. In the right side, we report snapshots from a process where misfolding occurs and the work done on the system is greater than that obtained from a reversible process. In the left side, we report snapshots from a process that instead microscopically violates the second law of thermodynamics (Ref. 51). The center panel contains the RC  $Z(\tau)$  for both simulations as a function of time. In the bottom panel, we plot  $Z(\tau) - \lambda(\tau)$ , where  $\lambda(\tau)$  is the deterministic protocol used for steering the RC. The snapshots were generated with the VMD program (Ref. 54).

single trajectory of the steered RC (but is indirectly contained in a collection of trajectories). Several configurational degrees of freedom, such as dihedral angles, may be indeed important in determining the dynamics of the studied system. These degrees of freedom can be thought of as “other” RCs and are not directly included in the dynamical models calibrated from the observed data. However, these RCs can relatively slowly evolve (compared to the time scale of the simulation) and may therefore effectively modulate the dynamics of the steered RC. Because of these other RCs, independent trajectories may be characterized by significantly different dynamics yielding a *collection* of diffusion models.<sup>23</sup> In this study, ideas from functional data analysis<sup>24</sup> (FDA) are used in our models to include this important source of “dynamical heterogeneity.” In particular, we develop a procedure that we call “FDA bootstrapping” (which is a specific instance of “type-II bootstrapping” discussed in Ref. 19). This bootstrapping is able to correlate the work done on the system during an unfolding or refolding realization to appropriate functions associated with a surrogate process approximation (SPA) model.<sup>19</sup> This type of treatment is attractive because it does not require the identification or manipulation of the slowly evolving collective coordinates (this is particularly relevant to wet-laboratory studies).

The second aim of this paper is to show how the proposed approach (coupled with ideas from Ref. 19) can be applied to recover the non-Gaussian work distributions associated with nonequilibrium SMD folding/unfolding of deca-alanine. This may possibly assist experiments or simulations

where path samples are so expensive or time consuming that obtaining an adequate statistical sample is difficult. The statistical validity of approximating deterministic (chaotic) dynamics using diffusion approximations has also been addressed here given its relevancy to a variety of atomistic systems.<sup>22,25,26</sup> However, in order to keep the presentation streamlined, we have relegated these results to the Supplementary Material.<sup>27</sup>

The outline of this article follows. In Sec. II, we review the basic concepts underlying the SPA model and the FDA methodology and provide the computational details regarding the SMD simulations of deca-alanine. In Sec. III, we report the results regarding the determination of the work distributions using diffusion models. In this section, we also report data and discuss issues regarding the connection between the dynamical information from SPA models and other possible RCs “orthogonal”<sup>28</sup> to the steered one. A summary and some potential future directions of research are finally discussed in Sec. IV.

## II. THEORY AND METHODS

### A. SMD simulations of deca-alanine

As stated in the Introduction, the dynamical system considered in the present study is the process of unfolding (and refolding) of one deca-alanine peptide at finite temperature. As in a previous work,<sup>6</sup> SMD simulations were used as a device for producing the forced nonequilibrium dynamics of the steered RC,  $Z$ , which is taken to be the end-to-end distance of deca-alanine. More specifically,  $Z$  corresponds to the distance between the N atom of the N-terminus amino acid (constrained to a fixed position) and the N atom of the C-terminus amino acid (constrained to move along a fixed direction). The values of  $Z$  in the completely folded and unfolded states of deca-alanine are assumed to be 15.5 and 31.5 Å, respectively. In the former state, deca-alanine is found in an  $\alpha$ -helix form, while an elongated “random coil”<sup>5</sup> configuration characterizes the latter state. We have arbitrarily assumed the unfolding and refolding processes as forward ( $F$ ) and reverse ( $R$ ), respectively.

It should be noted that, in general, the end-to-end distance does not uniquely determine the configurational state of polypeptides. However, in the specific case of deca-alanine, the equilibrium distribution at  $Z=15.5$  Å corresponds to an ensemble of microstates tightly peaked around the  $\alpha$ -helix form, as for this end-to-end distance alternative structures are highly unlikely under the associated stationary distribution.<sup>29</sup> The same holds true for the state corresponding to  $Z=31.5$  Å (i.e., there is a tightly peaked stationary probability density corresponding to an ensemble of elongated random coil configurations). Because of this, the thermodynamically relevant phase space points consistent with the  $\alpha$ -helix (and elongated random coil) to be used as initial microstates of the SMD simulations can readily be obtained by sampling from the stationary distribution associated with a constrained RC molecular dynamics simulation. However, as discussed in the Introduction, subtle configurational differences in these initial microstates may yield vastly different dynamical behaviors when the system is driven out

of equilibrium. The initial microstates of deca-alanine for the  $F$  and  $R$  realizations were randomly drawn from two “equilibrated” biased molecular dynamics simulations constraining  $Z$  to 15.5 and 31.5 Å, respectively, by means of a stiff harmonic potential (force constant  $k_{\text{pull}}$  equal to 800 kcal mol<sup>-1</sup> Å<sup>-2</sup>). In both equilibrium simulations and in the subsequent SMD simulations, constant temperature (300 K) has been enforced using a Nosé–Hoover thermostat.<sup>22</sup> Force field parameters have been taken from Ref. 30. For each type of process,  $F$  and  $R$ , we have generated 10<sup>3</sup> realizations guiding  $Z$  from 15.5 to 31.5 Å ( $F$  paths) or from 31.5 to 15.5 Å ( $R$  paths) using a time-dependent harmonic potential with the force constant reported above. The parameter driving the RC, which we denote with  $\lambda(t)$ , evolves in time with an externally determined constant rate of 0.032 Å ps<sup>-1</sup> (the same rate is used for both  $F$  and  $R$  processes). The harmonic guiding potential  $V(t)$  is the following:  $V(t) = 0.5k_{\text{pull}} [Z(t) - \lambda(t)]^2$ . We stress that  $\lambda(t)$  is a deterministic function common to all SMD simulations, and  $Z(t)$  represents a stochastic process. Moreover, we note that in order to generate a large amount of temporal data needed to reliably calibrate diffusion models and still have the time spacing between observations largely relative to the faster time scales,<sup>31–35</sup> we have performed longer SMD simulations than those reported in Ref. 6.

The ORAC program was used for all simulations.<sup>36</sup> It uses a reversible REference System Propagation Algorithm (r-RESPA), to integrate the equations of motion. In particular, forces between atoms whose distance is greater than 12 Å have been updated every 5 fs; forces between atoms whose distance falls in the range of 6–12 Å have been updated every ~1.667 fs; forces between atoms distant less than 6 Å, torsional forces, and forces between atoms in positions 1–4 have been updated every ~0.556 fs; and bending and stretching forces have been updated every ~0.278 fs. Simulation data needed for modeling, i.e., external work and  $Z(t)$ , have been stored every 50 fs.

## B. Surrogate process approximation and bootstrapping

If we assume that only a relatively small number of paths or realizations are available from an experiment/calculation, it is natural to try to use the rich amount of dynamical data contained in the paths in order to calibrate a surrogate model and then turn to this model to augment the limited data set for improving the quality of the sample. This is the basic idea behind the SPA method used in the present work and introduced in Ref. 19. We attempt to fit time series data coming from constant velocity SMD simulations of deca-alanine to time-dependent overdamped diffusion models using locally parametric maximum likelihood methods. The underlying complex system used to generate the time series is a system of deterministic chaotic ordinary differential equations, as opposed to a Langevin-type process, which was originally used in Ref. 19.

Specifically, a single RC trajectory obtained from a SMD simulation is used to calibrate a collection of local diffusion models, each model being characterized by a diffusion and a




FIG. 2. (Color online) Drift and diffusion functions (panels A and B, respectively) estimated using 15 SMD paths in the  $F$  and  $R$  directions (black and red curves, respectively). The sampling step of the SMD paths is 150 fs.

drift coefficient.<sup>19</sup> For the transition density expansions, we utilize the simulated maximum likelihood<sup>37</sup> estimator. Local overdamped diffusion models are estimated at the “center points” {15, 16, ..., 32} Å of the RC (for a total of 18 local models) using a neighborhood radius of 0.6 Å for each local model.<sup>19</sup> The diffusion and the drift coefficients of the local diffusion models are then stitched together to form two continuous functions, i.e., the diffusion and the drift functions representing the global diffusion model of the RC trajectory. These functions are obtained using a piecewise polynomial smoothing routine available from MATLAB’s spline toolbox (csaps, with a smoothing parameter equal to 0.4). In Fig. 2, we show the drift and diffusion functions calibrated using 15 forward and reverse SMD trajectories of deca-alanine. It should be noted that the diffusion model fitted to the data is identical to that used in Ref. 19. Section II of the aforementioned reference provides a more comprehensive discussion on the diffusion models.

The resulting global diffusion model can then be used to generate a collection of new paths taking the same initial condition as the SMD trajectories but using a batch of random number sequences. The different random number sequences used to create Brownian paths are intended to incorporate variability that can be attributed to physically unimportant fast degrees of freedom, such as vibrational degrees of freedom associated with covalent bonds, which are present in the full atomistic system. The term “type-I bootstrapping” is used here to describe the generation of new paths from the procedure described above.

There is interest in understanding the consequences of approximating highly structured chaotic “noise” with a fairly crude diffusion model<sup>33–35</sup> when the data are frequently sampled in time (see the Supplementary Material for an extensive discussion along these lines). In principle, any stochastic process could be fitted to the time series data and could serve as the surrogate, e.g., underdamped diffusions<sup>32</sup> and non-Markovian processes.<sup>31,38</sup> However, the motivation for focusing on the estimation of overdamped models stems from two facts. (1) In macromolecules, large collective mo-

tions typically<sup>9</sup> cause an effectively overdamped system even when the traditional hydrodynamic “damping” (or viscosity) is very low.<sup>39</sup> (2) Typically, in wet-laboratory experiments, only simple positionlike coordinates are available, which force the researcher into a framework where overly simplified dynamical models are used because the information to calibrate more sophisticated models is simply not accessible (or perhaps deemed not worth pursuing).<sup>40</sup>

Characterizing the errors introduced by ignoring degrees of freedom in a low-dimensional RC calculated from observations of a detailed atomistic system is a difficult and open problem.<sup>31,33,38</sup> The multiscale nature of the signals that come from atomistic systems makes the legitimacy of using a single low-dimensional diffusion model questionable. The approach described in the next section aims at remedying this problem by enhancing the sampling of *models* that a batch of fully atomistic paths sometimes yields. A physics-based description of the presence of a collection of models would assume that the effective free energy<sup>34</sup> is a function of two or more RCs, but we only observe and model the dynamics of one. Unlike the RC explicitly considered in the model, the other RCs are not controlled either at the beginning of the nonequilibrium experiment or during it. This implies that each RC path may give rise to a dynamical model that depends heavily on unobserved RCs. Therefore, many RC paths may result in a collection of models, whose “heterogeneous dynamical response” can be attributed to the unconstrained other RCs in the nonequilibrium experiment. In this paper, we present a technique for using the information contained in this distribution of models to make refined predictions of the work distribution from a small data set. We also demonstrate that this distribution of models can provide indirect information about other important RCs. This is made possible by FDA ideas.

### C. Using FDA ideas to bootstrap work paths

The basic goal of FDA (Refs. 24 and 41) is to characterize the distribution of a family of curves. Instead of statistically analyzing finite-dimensional random vectors (e.g., Euclidean vectors), one focuses instead on the distribution of objects living in an infinite-dimensional space. In the present case, we deal with the space of continuous functions that describe the dynamics of the various SMD trajectories. One may associate three functions to each SMD trajectory: The work function corresponding to the time sequence of the external work performed on the system during the simulation and the two SPA coefficient functions (see Fig. 2). Here, we assume that the observed work functions are independent and identically distributed (iid).

Unfortunately, one usually does not have the luxury of continuous time observation in many applications.<sup>18,24</sup> However, when the underlying objects are believed to be relatively smooth functions (as it seems to be the case of the deca-alanine system), it makes sense to discretely sample the external work at a moderate sampling frequency. Then, one can attempt dimension reduction by using principal component analysis (PCA) tools.<sup>24</sup> To determine the correlation be-

tween the SPA model functions and the observed work functions, one could create a raw observation matrix of the form

$$\begin{bmatrix} \mathcal{W}^{(1)}(\tau_1) & \mathcal{W}^{(2)}(\tau_1) & \cdots & \mathcal{W}^{(N)}(\tau_1) \\ \mathcal{W}^{(1)}(\tau_2) & \mathcal{W}^{(2)}(\tau_2) & \cdots & \mathcal{W}^{(N)}(\tau_2) \\ & \vdots & & \\ \mathcal{W}^{(1)}(\tau_D) & \mathcal{W}^{(2)}(\tau_D) & \cdots & \mathcal{W}^{(N)}(\tau_D) \\ \tilde{A}^{(1)}[Z(\tau_1)] & \tilde{A}^{(2)}[Z(\tau_1)] & \cdots & \tilde{A}^{(N)}[Z(\tau_1)] \\ \tilde{A}^{(1)}[Z(\tau_2)] & \tilde{A}^{(2)}[Z(\tau_2)] & \cdots & \tilde{A}^{(N)}[Z(\tau_2)] \\ & \vdots & & \\ \tilde{A}^{(1)}[Z(\tau_D)] & \tilde{A}^{(2)}[Z(\tau_D)] & \cdots & \tilde{A}^{(N)}[Z(\tau_D)] \\ \tilde{C}^{(1)}[Z(\tau_1)] & \tilde{C}^{(2)}[Z(\tau_1)] & \cdots & \tilde{C}^{(N)}[Z(\tau_1)] \\ \tilde{C}^{(1)}[Z(\tau_2)] & \tilde{C}^{(2)}[Z(\tau_2)] & \cdots & \tilde{C}^{(N)}[Z(\tau_2)] \\ & \vdots & & \\ \tilde{C}^{(1)}[Z(\tau_D)] & \tilde{C}^{(2)}[Z(\tau_D)] & \cdots & \tilde{C}^{(N)}[Z(\tau_D)] \end{bmatrix} = \begin{bmatrix} \mathcal{W} \\ \tilde{\mathbf{A}} \\ \tilde{\mathbf{C}} \end{bmatrix},$$

where  $\mathcal{W}^{(i)}(\tau_j)$ ,  $\tilde{A}^{(i)}[Z(\tau_j)]$ , and  $\tilde{C}^{(i)}[Z(\tau_j)]$  are the work, drift, and diffusion functions at time  $\tau_j$  corresponding to the  $i$ th SMD trajectory. The  $N$  matrix columns correspond to the function “samples” observed discretely at the times  $\tau_1, \tau_2, \dots, \tau_D$ .<sup>42</sup> After projecting these functions onto a basis, the weights (i.e., magnitude of the projection onto a finite member basis set) are then “stacked” upon each other (this simply means that the weights of the different projections are concatenated) to form a new matrix with blocks  $\mathcal{W}'$ ,  $\tilde{\mathbf{A}}'$ , and  $\tilde{\mathbf{C}}'$ .<sup>43</sup> The specific weight that should be assigned to the different blocks is not an easy question to answer because introducing an appropriate metric (or semimetric) is difficult<sup>41</sup> without *a priori* information on the smoothness of the curves. In this work, we use a fairly simple weighting procedure made of the following steps. (1) Within each block, the population average (across the  $N$  samples) of the coefficients (a vector) was subtracted from the individual coefficient vectors. (2) Within each block, the coefficients resulting from step 1 were “normalized” by the corresponding block standard deviation.<sup>44</sup> The intention is to give each block the same relative importance in a PCA-type decomposition. After the PCA was carried out on the data matrix, each stacked observation (of basis weights) was then projected onto the constructed principal component basis (three PCA modes were retained in all cases reported) and the PCA weights of each projection were then recorded. In order to get the distribution of the weights, two methods were used.

- (1) Under the *ad hoc* assumption that the PCA weights are statistically independent<sup>45</sup> (by the construction of the PCA decomposition, the weights are linearly uncorrelated), the empirical distributions (across the functional curve population) were recorded for each PCA projection and new samples were created by resampling the PCA weight distributions. This new set of weights augments the original small set of weights (the basic idea being to “interpolate” between function curves by resampling from the empirical cumulative distribution functions).

- (2) In an effort to computationally remove any dependence between the weights, and still under the assumption given in Ref. 43, we employed tools, such as the independent component analysis<sup>46,47</sup> (ICA) or projection-pursuit methods, to transform the weights into roughly “decoupled” coordinates (we used the FASTICA package<sup>47</sup>). We then empirically measured the cumulative distribution function of the weights and used the resulting distributions to generate a new set of weights. The resulting weights were finally transformed back to the original coordinates, and the resulting work paths were recorded.<sup>48</sup>

The type of methods described above are referred to as FDA bootstrapping. The basic steps of the overall method using one of the above methods are summarized below.

- Let  $\mathbf{X} \equiv [\mathcal{W}', \tilde{\mathbf{A}}', \tilde{\mathbf{C}}']^T$  and denote the corresponding linear basis by  $\{\mathbf{u}_1, \mathbf{u}_2, \dots\}$ . Each column of  $\mathbf{X}$  is then projected onto each column of the linear basis to obtain an ensemble of weights  $\equiv \{\mathbf{w}_1, \mathbf{w}_2, \dots\}$ .
- The empirical cumulative distribution function associated with each (raw or processed) row of the  $\mathbf{w}_i$  vectors is determined (i.e., across the  $N$  samples, a cumulative distribution function of each component of the weight vector is measured).
- Using the above cumulative distribution function, a large number of “bootstrapped” weights are then resampled.
- Synthetic work paths and diffusion models are finally produced using “FDA bootstrapped models.” The cost of this operation is usually marginal in relation to the overall data production procedure.

### III. RESULTS

#### A. Failure of direct work process approximation

SPA models were calibrated by observing 15 randomly selected SMD paths. The time-dependent distribution of  $Z$  was accurately modeled with our stochastic dynamical models (see Figs. 1 and 2 in the Supplementary Material). The accuracy in the  $Z$  distribution approximation gives us partial hope in using the estimated SPA models to generate a collection of surrogate processes from which work can *directly* be measured by type-I bootstrapping (as was done in Ref. 19).  $10^4$  such paths were generated and the work done on the system (pulling in the  $R$  direction) was recorded as a function of time. The “pulling” parameters used in these type-I bootstrapping SPA simulations were identical to those used in the corresponding SMD paths. For comparison, the work done on the system as a function of time from the 1000 SMD paths has also been computed. Both SPA and SMD average work profiles, along with their standard deviations, are shown in Fig. 3. The standard deviation of the SPA work paths is far too large compared to the SMD process it aims at approximating. This can be very problematic if one wants to use the approximate work distribution to compute free en-




FIG. 3. (Color online) Mean  $\mu$  and standard deviation  $\sigma$  for  $R$  work paths using the SMD and SPA processes. The curves labeled  $W_i$  correspond to a few randomly chosen work paths (see the legend).

ergy differences using a nonequilibrium method, such as the Jarzynski equality.<sup>7</sup>

A few representative work paths are also plotted for both the SPA and SMD cases. When individual paths are compared, one clearly observes that the SPA work paths are too “rough” in relation to the relatively smooth SMD work paths. This smoothness is likely caused by structured fast degrees of freedom not explicitly included in our surrogate models. Accurately approximating this structure would be a formidable time series modeling problem and an interesting future research direction. However, here, we turn instead to the tools of FDA to model the work process. The estimated SPA models are still useful (despite the fact that they cannot directly simulate useful work paths) because we demonstrate that the information in the collection of SPA models indirectly summarizes some features of the complex SMD process.

#### B. Applying FDA bootstrapping to approximate work distributions

We recall that one of the motivations of this study is to use the information contained in the estimated SPA models for approximating the shape of the work distributions associated with SMD simulations. In several early computational studies,<sup>5,21,49</sup> Gaussian work distributions were observed in pulling nonequilibrium processes. Most of the arguments in the literature that attempt to justify this fact either are heuristic<sup>49</sup> or make the assumption that a *single* overdamped diffusion can be used to model the time series coming from nonequilibrium SMD simulation data.<sup>5</sup> However, it has also been shown that even seemingly Gaussian work distributions may not be Gaussian. This is just the case of the work distribution for the unfolding process of deca-alanine whose shape, though very similar to a Gaussian profile, is inherently asymmetric.<sup>6</sup> This brings up a subtle point we want to stress. In this system, a single overdamped model (calibrated at this time scale) is not adequate to describe the dynamics of each nonequilibrium response. This is likely caused by slowly evolving degrees of freedom outside of the selected RC that have not been sufficiently “averaged out.”<sup>32,50</sup> In our system, this causes one to estimate *multiple* overdamped diffusion

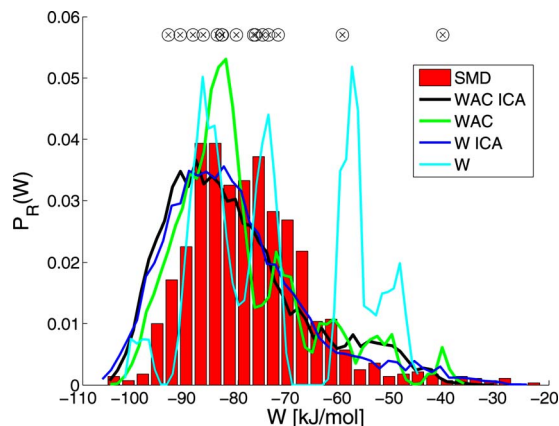


FIG. 4. (Color online) Work distributions in the  $R$  direction obtained from 1000 SMD simulations and from the FDA bootstrapping methods (see the legend). The circles indicate the work values corresponding to the 15 SMD paths used for calibrating the SPA models used in FDA bootstrapping.

models. The approximation we introduced takes this model variability into account and should be viewed as a new type of stochastic process that is different from a single “simple” overdamped diffusion. This new stochastic process violates the conditions that guarantee a Gaussian work distribution stated in Ref. 5. Because guaranteeing (or enforcing) a Gaussian work distribution can be difficult with certain RCs, it would be useful to have a tool that can be used to check the validity of the Gaussian work assumption without requiring a large number of nonequilibrium realizations. In the deca-alanine system, the FDA motivated techniques presented here are shown to be useful in this context (see below).

The work distributions recovered from the 1000 SMD simulations and from the various FDA bootstrapping approaches (see Sec. II C) for the  $R$  realizations are shown in

Fig. 4. Fifteen SMD trajectories, taken randomly from the overall ensemble of SMD trajectories, were used to calibrate 15 SPA models used for the various FDA bootstrapping approaches. The work values corresponding to these SMD trajectories are also reported in Fig. 4. The FDA bootstrapping types are classified as follows. (1a) The FDA bootstrapping used diffusion and drift coefficient functions from the SPA models along with the corresponding work function from SMD trajectories ( $\mathcal{WAC}$  bootstrapping). (1b) Same as (1a), but ICA decoupling has been used on weights ( $\mathcal{WAC}$  ICA bootstrapping). (2a) Only the work function from the SMD trajectories has been used in the FDA bootstrapping ( $\mathcal{W}$  bootstrapping). (2b) Same as (2a), but ICA decoupling has been used on weights ( $\mathcal{W}$  ICA bootstrapping). Three PCA modes were used in all cases reported. The number of surrogate paths extracted from each FDA bootstrapping type is  $10^4$ . In addition, the same random number stream was used for each case in order to minimize sampling errors and to attribute differences in the results solely to the systematic differences in the computational methods.

We note that a naive bootstrapping ( $\mathcal{W}$ ) falsely predicts multiple modes, whereas the other bootstrapping schemes fairly accurately capture the shape of the SMD work distribution. This is indeed quite surprising given that only 15 SMD paths underly the SPA models. The fact that a structured tube of functions is observed in the SPA models (see Fig. 2) and that using this information in the FDA bootstrapping improves the reproduction of the work distribution suggests that the work and the diffusion dynamics of the RC (accounted for by the SPA models) are indeed significantly correlated.

Increasing the number of SMD paths underlying the SPA models improves the estimation of the low order moments (see Table I), but the overall shape of the predicted work

TABLE I. Work statistics. The SMD label corresponds to cases in which SMD trajectories have been used to evaluate the work performed on the system during the whole unfolding path. In all of the FDA bootstrapping cases,  $10^4$  bootstrapped processes have been used. The number of SMD paths used to calibrate each collection of SPA models is reported in parentheses (in the SMD cases, this corresponds to the number of paths from which work was directly measured). In this table, the larger samples are built incrementally from the smaller ones. For example,  $\mathcal{W}$  (30) contains the 15 paths from  $\mathcal{W}$  (15).

Case	$\hat{\mu}$	$\hat{\mu}$ -median	Mode	$\hat{\sigma}$	Skew	Kurtosis
SMD (15)	-77.785	2.478	-75.270	13.390	1.491	4.915
SMD (30)	-79.362	1.915	-82.014	11.827	1.210	5.254
SMD (60)	-77.481	1.424	-82.795	12.166	1.235	5.382
SMD (100)	-78.123	2.642	-87.845	12.154	0.993	4.583
SMD (1000)	-77.661	1.301	-84.889	11.844	1.127	5.385
$\mathcal{WAC}$ ICA (15)	-80.591	2.677	-88.129	13.408	0.856	3.228
$\mathcal{WAC}$ ICA (30)	-80.100	0.789	-82.127	11.634	0.358	2.994
$\mathcal{WAC}$ ICA (60)	-80.151	1.366	-86.195	12.165	0.510	2.985
$\mathcal{W}$ ICA (15)	-80.736	2.105	-85.370	13.428	1.032	4.233
$\mathcal{W}$ ICA (30)	-81.688	1.300	-84.595	11.538	0.948	4.900
$\mathcal{W}$ ICA (60)	-80.042	1.119	-82.647	12.161	0.736	4.279
$\mathcal{WAC}$ (15)	-79.560	3.027	-83.593	12.693	1.024	3.726
$\mathcal{WAC}$ (30)	-78.427	0.563	-80.560	11.351	0.110	3.417
$\mathcal{WAC}$ (60)	-75.564	0.980	-82.097	11.561	0.306	2.859
$\mathcal{W}$ (15)	-73.196	2.376	-86.172	14.201	0.240	1.832
$\mathcal{W}$ (30)	-76.847	0.426	-74.422	10.621	0.338	3.488
$\mathcal{W}$ (60)	-75.353	1.416	-83.797	12.163	0.623	3.305




FIG. 5. (Color online)  $\mathcal{WAC}$  ICA work distributions in the  $R$  and  $F$  directions (three solid curves and two dashed curves, respectively). In the inset, the  $R$  and  $F$  work distributions obtained from the SMD simulations are reported (black and green histograms, respectively). The work value at the intersection point of these last work distributions corresponds to the free energy difference ( $91.7 \text{ kJ mol}^{-1}$ ). We stress that each FDA bootstrapping (approximate) work density used different iid SMD paths as the underlying data. For example, in the  $R$  direction, 45 total SPA models were estimated, and this collection was split into three sets of 15. These batches were used to create the approximate  $R$  work densities shown here.

distributions does not substantially change. We point out that  $\mathcal{W}$  ICA and  $\mathcal{WAC}$  ICA methods seem to work better than  $\mathcal{W}$  and  $\mathcal{WAC}$  methods because there are multiple clusters (discussed in detail in Sec. III C) of observed dynamical responses (i.e., SPA models) that a linear method such as PCA cannot directly account for.<sup>46,47</sup> The SPA model information aids in detecting these clusters, thus allowing for an appropriate assignment of the weights to the work paths. Though the weighting induced by the SPA functions resulting in an improved work distribution estimation is just a fortuitous coincidence, in the future, we hope to develop schemes that systematically correlate these SPA “function clusters” with slowly evolving collective degrees of freedom and assign appropriate weights to the paths using empirically determined correlations.

In the  $\mathcal{WAC}$  and  $\mathcal{W}$  bootstrapping methods, small batches of 15 SMD paths were drawn (without replacement) and used to calibrate 15 SPA models, and then these models are used to approximate the nonequilibrium work distribution. When this procedure is repeated (with new batches of iid SMD sample paths), the mode variability in the resulting work distribution is significant (see Table I). The mean and mode in the  $\mathcal{WAC}$  ICA approach are relatively stable, and when independent batches of 15 SMD paths are used in FDA bootstrapping, the resulting work distributions have a non-Gaussian shape with roughly the same mode. This is demonstrated in Fig. 5, where  $\mathcal{WAC}$  ICA bootstrapping is applied for calculating the  $R$  and  $F$  work distributions using independent batches of SMD paths (three batches of 15 paths for the  $R$  direction and two batches of 15 paths for the  $F$  direction). For comparison, in the inset in Fig. 5, we report the  $F$  and  $R$  work distributions recovered from the 1000 SMD paths. The

free energy difference between folded and unfolded states of deca-alanine can be predicted using a random pair of  $F$  and  $R$  work distributions contained in Fig. 5 (free energy difference actually corresponds to the intersection point of the  $P_R(-W)$  and  $P_F(W)$  distributions). Such free energy differences are roughly consistent with the free energy difference obtained from the SMD work distributions (inset in Fig. 5) and those reported in Ref. 6. A more detailed analysis about free energy determination using the various FDA models is provided in the Supplementary Material.

Given that the underlying decoupling ICA methods aim at separating signals into non-Gaussian sources,<sup>47</sup> it is reassuring to note that the  $\mathcal{WAC}$  ICA approach is also able to recover Gaussian-like work distributions, such as those in the  $F$  direction (Fig. 5). This is quantified by the skew and by the difference between the mean and the median of the two  $F$  work distributions, which are very small (skews: 0.0862 and  $-0.0142$ ; difference between the mean and the median:  $-0.0979$  and  $0.0249$ ). It should also be noted that the ICA methods can give rise to convergence problems for small sample sizes (this computational decoupling method usually requires larger “training sets”<sup>46,47</sup>). We notice, however, that in the present case good results are obtained for both  $\mathcal{W}$  ICA and  $\mathcal{WAC}$  ICA bootstrapping even using small sample sizes. The positive performance of the small sample  $\mathcal{WAC}$  case is reassuring because this method can be used even when the sample size in hand is not large enough for reliable use of ICA methods.

### C. Physical interpretation and collective variables orthogonal to the RC

Previous works have already made some connections between classical statistical mechanics and nonequilibrium work relations.<sup>51</sup> The time series studied here suggest that paths exhibiting extreme work values, i.e., paths whose related work falls in one of the tails of the measured work distribution, are associated with stochastic dynamics, which more strongly deviate from the “average SPA model.” These types of relationships were first discovered by the data driven methods presented earlier (using only the end-to-end distance RC). Physics-based intuitive explanations were then tested by analyzing the sequence of SMD all-atom snapshots. The FDA methods revealed that there are clusters (subpopulations) in our data sets. The fact that the clusters appear to have physical relevance (see discussion below) suggests that SPA output may be useful for a functional classification tool<sup>41</sup> in systems where the underlying molecular details are not well understood. We expand on this speculative idea below, but first some established physical facts are presented.

When a process takes place far from equilibrium, classical statistical mechanics predicts that paths on average dissipate more work than they would if the process were reversibly carried out. We label these paths as the “uninteresting paths.” We label those that microscopically violate the second law of thermodynamics, i.e., with negative dissipated work, as the “interesting paths.”<sup>51</sup> The labels are only used here to distinguish between the two cases in terms of the contribution each of them gives to the exponential average

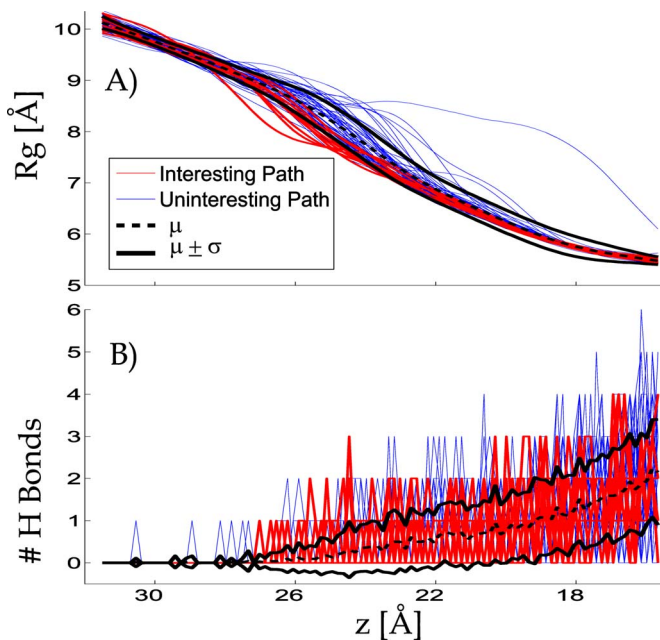


FIG. 6. (Color online) Alternative RCs, i.e., radius of gyration (panel A) and number of hydrogen bonds in deca-alanine (panel B), as a function of the steered RC (from  $R$  SMD simulations). The black dashed line denotes the mean population curve  $\mu$  calculated using 100 SMD paths. The solid black lines denote the standard deviation curves  $\mu + \sigma$  and  $\mu - \sigma$ . Ten interesting and 50 uninteresting paths are plotted in red and blue, respectively. Both quantities were calculated using the VMD program (Ref. 54). The radius of gyration was computed using all atoms of deca-alanine. The hydrogen bonds were computed using the VMD subroutine H bonds using a 3 Å  $\text{HN} \cdots \text{OC}$  bond distance and 20°  $\text{O} \cdots \text{N-H}$  angle cutoffs.

used for recovering free energy differences.<sup>7</sup> We notice that in our system, it is difficult to distinguish between interesting and uninteresting paths by visually inspecting the dynamics of the steered RC alone (see the bottom of Fig. 1). However, there are other configurational coordinates that we could use to make the differences more apparent. Figure 6 shows some of these possible configurational coordinates, namely, the radius of gyration and the number of hydrogen bonds in deca-alanine, as a function of the steered RC. Specifically, in Fig. 6, we report the path population mean  $\mu$  and the curves that indicate the standard deviation,  $\mu + \sigma$  and  $\mu - \sigma$ , of these two quantities. Moreover, the blue lines denote a set of uninteresting paths, while the red lines denote a set of interesting paths. Using the number of hydrogen bonds, it is difficult to distinguish between interesting and uninteresting paths, but by using the radius of gyration, one observes that the interesting paths form a tube that is systematically lower than the mean of the entire population. It is noteworthy that the radius of gyration, associated with both interesting and uninteresting paths, has significantly different behavior in the range  $22 < Z < 26$  Å. When the total amount of nonequilibrium work is used to index the radius of gyration paths, a “clusterlike” trend emerges. This indicates that in this system, the work produced during nonequilibrium processes can help in classifying trajectories. Here, we aim to show that indexing SPA models on the basis of the nonequilibrium work also induces clusterlike behavior in the empirically measured SPA functions, and that such a clustering contains partial information about a reaction coordinate not directly




FIG. 7. (Color online) Drift (panel A) and diffusion (panel B) functions (with population mean subtracted) of SPA models calibrated on interesting and uninteresting  $R$  SMD paths (solid and dashed lines, respectively).

modeled, namely the radius of gyration. This observation is likely unique to deca-alanine or to a small class of similar macromolecules. The general concept, which we believe is applicable to a wider class of systems, is that clusters of responses observed in the SPA model correlate with some signature of a more complex (initially unobserved) RC. In the case of simulations, one is free to investigate the detailed trajectories to see if this is so, and in wet-laboratory experiments, one may want to consider testing a variety of hypotheses to see if the cluster of SPA models correlates to a physically interesting quantity.

In Fig. 7, we show drift and diffusion coefficients of two subsets of SPA models taken from a population of 60 estimated models. The subsets correspond to SPA models whose work paths, or more precisely the work paths corresponding to the underlying SMD paths used to estimate the SPA models, fell into the two tails of the  $R$  work distribution (note that in the figure, the mean function calculated over the 60 SPA model coefficient functions has been subtracted from the individual functions). Sampling the paths in the tails of the work distribution allows us to select interesting (smallest work values) and uninteresting paths (largest work values). In the interesting paths, the drift function shows a faster contraction at the *beginning portion* of the path. A reasonable explanation for this would be that in the interesting paths, the elongated random coil is aligned in such a way that it experiences significantly more attractive forces from adjacent peptides (compared to a typical configuration consistent with the initial  $Z$ ). The converse applies to uninteresting paths. Perhaps a more intriguing trend is observed in the diffusion functions associated with interesting paths. Initially, they are low but then increase rapidly in a sort of “transition” that occurs around  $Z \approx 26$  Å. It is interesting to note that the function describing the noise magnitude undergoes this tran-

sition roughly at the same  $Z$  value at which the radius of gyration paths “diverge” (compare Figs. 6 and 7). Some physicists might attribute this fact to a metastable state that exists for  $Z$  near 26 Å.<sup>15</sup>

The possible connections existing between diffusion models and structural features in nonequilibrium processes, as we have shown here, suggest that incorporating dynamic response information (such as that contained in this collection of SPA models) along with other structural information may provide a useful tool in classifying configurational trajectories where there are important unresolved (due to lack of direct observation) configurational details that are believed to be important.<sup>10,38</sup> Identifying clusters of dynamical responses could potentially assist in correlating SPA model functions to both work and slowly evolving configurational details using more sophisticated FDA techniques (e.g., functional partial least squares or mixed linear models<sup>41</sup>). This is attractive in situations where small and large samples of nonequilibrium paths are available. Physically, the clusters of curves in Fig. 7 suggest that in this system, the microstates associated with the slowly evolving degrees of freedom can be effectively modeled as a continuum. For example, the radius of gyration appears to take a continuous range of values for a given temperature and fixed value of the RC. As a result of this, a continuous tube of SPA models is likely observed because the (roughly) continuously distributed radius of gyration RC modulates the dynamical response.<sup>52</sup> In more complex macromolecules, one should be warned that the estimated SPA models will likely depend on the configurational degrees of freedom in a more complicated fashion, but if they do measurably modulate the dynamics (e.g., see Ref. 4), then our technique may be helpful (but the nonequilibrium work will probably not be approximated as accurately as it was done here).

#### IV. CONCLUSIONS AND PERSPECTIVES

We have presented a SPA method<sup>19</sup> based on a diffusion process to model nonequilibrium paths of a given RC. We have applied the methodology to the nonequilibrium dynamics of the end-to-end distance of a deca-alanine peptide at finite temperature produced by SMD simulations. Special attention has been devoted to the refolding process of deca-alanine, which was previously demonstrated to yield markedly non-Gaussian work distributions.<sup>6</sup> In modeling the dynamical behavior of the RC of deca-alanine, two strategies have been proposed. One is based on the modeling of one SMD path through a series of local diffusion models, each describing a limited interval of the path. These local models are then stitched to give the global diffusion model (SPA model) for the given SMD path. A collection of global SPA models, modeled on various statistically independent SMD paths, has then been used to increase the statistical sampling, that is, to produce a large amount of surrogate paths. Unluckily, these surrogate paths have been demonstrated to furnish excessively broad work distributions. In order to correlate the dynamics of the RC resulting from the SPA models with the work paths observed in the underlying SMD paths, we have turned to FDA.<sup>23</sup> This procedure has

allowed us to recover good quality work distributions considering the relatively small number of SMD paths used for the modeling. This is indeed one of the goals of our study: When the number of direct observations (from experiments or calculations) is small, our approach allows one to obtain a reliable estimation of the work distribution shape. Given the fact that several recent methods for estimating free energy differences<sup>7,8</sup> depend on approximations of the work distributions obtained from a finite collection of nonequilibrium realizations along a given RC, any method for improving the quality of the work distributions may provide considerable help. Moreover, the determination of the work distribution shape can also furnish significant insights on the underlying dynamics of nonequilibrium processes. Taking advantage of this possibility, we currently investigate strategies for supplying work distribution reweighting algorithms [see, e.g., Eq. (6) of Ref. 53] with SPA protocols.

A further outcome of this study is to reveal the possibility of using SPA models to detect stochastic dynamical features of a system, e.g., trends in the drift or diffusion coefficient of an approximating SDE, from noisy reaction coordinate time series (obtained from nonequilibrium simulations). In fact, the dynamical behavior of slowly evolving degrees of freedom or of collective variables orthogonal to the chosen RC has been shown to be emphasized by SPA models. This is evident in the varied modulation of the drift and diffusion functions characterizing the various SPA models (see Fig. 2). However, the drift and diffusion functions of a SPA model alone (along a single path) can hardly reveal the occurrence of peculiar heterogeneities in the dynamics of the system. On the other hand, such features can be detected by analyzing clusters in a collection of SPA drift and diffusion coefficient functions. At least, this is the case observed in deca-alanine, where the dynamics of the radius of gyration appears to be significantly correlated to both drift and diffusion functions. The radius of gyration was shown to readily identify work paths that transiently violated the second law of thermodynamics. It should be noted that in a real experiment, a careful analysis of additional degrees of freedom, as we have done for deca-alanine, may be difficult or simply not feasible. The methods presented here provide an indirect way of quantifying the effect of difficult to measure degrees of freedom utilizing dynamical information contained in a collection of time series of a simple RC.

In summary, we have demonstrated that the measured diffusion functions offer an alternative method for using noise (estimated on a pathwise basis) in nonequilibrium simulations/experiments to detect subtle small scale transitions. One possible future application of this observation would be to develop a sampling method that constructs a basis (e.g., functional PCA) from a few “entire” pullings (i.e., span the entire RC range of interest). If after this step is completed, transition regions become apparent (or the location of transitions is known from established physical theories), one could exploit this finding and run a larger batch of “short” SMD simulations/experiments. By this, we mean that the SMD simulations/experiments would only need to be carried out until just after the transition region of the RC. In this way, one could project onto the previously determined(em-

pirical) basis and use the larger set of new trajectories to get a refined estimate of the frequency of certain important small scale molecular events by analyzing the distribution of the projections onto this basis.

## ACKNOWLEDGMENTS

The work of C.P.C. was supported by NSF Grant No. DMS-0240058 and No. ACI-0325081 and that of R.C. by the European Union Grant No. RII3-CT-2003-506350.

- <sup>1</sup>C. Bustamante, J. Liphardt, and F. Ritort, *Phys. Today* **58**(7), 43 (2005).
- <sup>2</sup>D. Collin, F. Ritort, C. Jarzynski, S. B. Smith, I. Tinoco, Jr., and C. Bustamante, *Nature (London)* **437**, 231 (2005).
- <sup>3</sup>N. C. Harris, Y. Song, and C.-H. Kiang, *Phys. Rev. Lett.* **99**, 068101 (2007).
- <sup>4</sup>S. Paramore, G. S. Ayton, and G. A. Voth, *J. Chem. Phys.* **127**, 105105 (2007).
- <sup>5</sup>S. Park and K. Schulten, *J. Chem. Phys.* **120**, 5946 (2004).
- <sup>6</sup>P. Procacci, S. Marsili, A. Barducci, G. Signorini, and R. Chelli, *J. Chem. Phys.* **125**, 164101 (2006).
- <sup>7</sup>C. Jarzynski, *Phys. Rev. Lett.* **78**, 2690 (1997).
- <sup>8</sup>G. E. Crooks, *J. Stat. Phys.* **90**, 1481 (1998).
- <sup>9</sup>W. Wriggers, Z. Zhang, M. Shah, and D. C. Sorensen, *Mol. Simul.* **32**, 803 (2006).
- <sup>10</sup>X. Zhuang, H. Kim, M. J. B. Pereira, H. P. Babcock, N. G. Walter, and S. Chu, *Science* **296**, 1473 (2002).
- <sup>11</sup>K. A. Walther, F. Gräter, L. Dougan, C. L. Badilla, B. J. Berne, and J. M. Fernandez, *Proc. Natl. Acad. Sci. U.S.A.* **104**, 7916 (2007).
- <sup>12</sup>J. Wang, *Chem. Phys. Lett.* **418**, 544 (2006).
- <sup>13</sup>S. Kim, P. C. Blainey, C. M. Schroeder, and X. S. Xie, *Nat. Methods* **4**, 397 (2007).
- <sup>14</sup>M. Vendruscolo and C. M. Dobson, *Science* **313**, 1586 (2006).
- <sup>15</sup>L. Janosi, Doctoral Dissertation (Advisor: Ioan Kosztin), U. Missouri (Columbia) "Multiscale Modeling of Biomolecular Systems," available at: <http://edt.missouri.edu/Fall2007/Dissertation/JanosiL-083107-D9109/>
- <sup>16</sup>G. Hummer and A. Szabo, *Proc. Natl. Acad. Sci. U.S.A.* **98**, 3658 (2001).
- <sup>17</sup>F. Gräter and H. Grubmüller, *J. Struct. Biol.* **157**, 557 (2007).
- <sup>18</sup>B. L. S. Prakasa Rao, *Statistical Inference for Diffusion Type Processes* (Arnold, London, 1999).
- <sup>19</sup>C. P. Calderon, *J. Chem. Phys.* **126**, 084106 (2007).
- <sup>20</sup>M. Carrion-Vazquez, A. F. Oberhauser, T. E. Fisher, P. E. Marszalek, H. Li, and J. M. Fernandez, *Prog. Biophys. Mol. Biol.* **74**, 63 (2000).
- <sup>21</sup>I. Kosztin, B. Barz, and L. Janosi, *J. Chem. Phys.* **124**, 064106 (2006).
- <sup>22</sup>W. G. Hoover, *Phys. Rev. A* **31**, 1695 (1985).
- <sup>23</sup>Or, more generally, a collection of dynamical responses (one is free to use dynamical models more complicated than diffusion processes).
- <sup>24</sup>J. Ramsay and B. W. Silverman, *Functional Data Analysis* (Springer, New York, 2005).
- <sup>25</sup>D. J. Searles, L. Rondoni, and D. J. Evans, *J. Stat. Phys.* **128**, 1337 (2007).
- <sup>26</sup>H. Kushner, *Weak Convergence Methods and Singularly Perturbed Stochastic Control and Filtering Problems* (Birkhauser, Boston, 1990).
- <sup>27</sup>See EPAPS Document No. E-JCPSA6-128-001816 for all references to the Supplementary Material. For more information on EPAPS, see <http://www.aip.org/pubservs/epaps.html>.
- <sup>28</sup>We simply define two RCs as orthogonal when the information set is disjointed (i.e., the value of each coordinate contains relevant system information that cannot be determined from the value of the other coordinate).
- <sup>29</sup>We refrain from using the term "equilibrium distribution" because a biasing potential is used to maintain the configurations at the desired  $Z$  value.
- <sup>30</sup>A. D. Mackerell, D. Bashford, M. Bellot, R. L. Dunbrack, J. D. Evanseck, M. J. Field, J. Gao, H. Guo, S. Ha, D. Joseph-McCarthy, L. Kuchnir, K. Kuczera, F. T. K. Lau, C. Mattos, S. Michnick, T. Nog, D. T. Nguyen, B. Prodhom, W. E. Reiher, B. Roux, M. Schlenkrich, J. C. Smith, R. Stote, J. Straub, M. Watanabe, J. Wiorcikiewicz-Kuczera, and M. Karplus, *J. Phys. Chem. B* **102**, 3586 (1998).
- <sup>31</sup>A. Mamonov, M. Kurnikova, and R. Coalson, *Biophys. Chem.* **124**, 268 (2006).
- <sup>32</sup>R. Zwanzig, *Nonequilibrium Statistical Mechanics* (Oxford University Press, New York, 2001).
- <sup>33</sup>G. A. Pavliotis, A. M. Stuart, and K. C. Zygalakis, *Commun. Math. Sci.* **5**, 507 (2007).
- <sup>34</sup>G. A. Pavliotis and A. M. Stuart, *J. Stat. Phys.* **127**, 741 (2007).
- <sup>35</sup>C. P. Calderon, *Multiscale Model. Simul.* **6**, 656 (2007).
- <sup>36</sup>P. Procacci, T. A. Darden, E. Paci, and M. Marchi, *J. Comput. Chem.* **18**, 1848 (1997).
- <sup>37</sup>M. Bradnt and P. Santa-Clara, *J. Financ. Econ.* **63**, 161 (2002).
- <sup>38</sup>S. C. Kou and X. S. Xie, *Phys. Rev. Lett.* **93**, 180603 (2004).
- <sup>39</sup>A. Ansari, *J. Chem. Phys.* **110**, 1774 (1999).
- <sup>40</sup>Besides the statistical validity of the overdamped diffusion, approximation can be readily tested using information contained in our models (see the Supplementary Material).
- <sup>41</sup>F. Ferraty and P. Vieu, *Nonparametric Functional Data Analysis*, Springer Series in Statistics (Springer, New York, 2006).
- <sup>42</sup>It should be noted that the actual RC value  $[Z(\tau)]$  is very close to the deterministic target value  $[\lambda(\tau)]$  because of the high spring constant used in the SMD simulations. In the FDA analysis employed, there was a negligible difference between results obtained using  $Z$  and those obtained using  $\lambda$ .
- <sup>43</sup>The matrix we work with in this study is actually a block of the weights numerically obtained when our discretely observed functions are projected onto a Haar wavelet basis set.
- <sup>44</sup>Within each block, the variance at a grid point was first computed across the function population. The variance at all grid points within the block was then summed, and the square root of this quantity is what we call the "standard deviation."
- <sup>45</sup>Since the initial configurations are iid samples from a stationary distribution, there is no statistical correlation between the work, diffusion, and drift functions for two different SMD simulations.
- <sup>46</sup>J. V. Stone, J. Porrill, N. R. Porter, and I. W. Wilkinson, *Neuroimage* **15**, 407 (2002).
- <sup>47</sup>A. Hyvarinen, *IEEE Trans. Neural Netw.* **10**, 626 (1999).
- <sup>48</sup>Unfortunately, ICA typically requires very large sample sizes to decouple and separate signals into (roughly) independent components. However, if one has a large collection of paths and is interested in using ICA to classify trajectory clusters, then ICA might be more useful (and reliable).
- <sup>49</sup>F. Ritort, in "Work fluctuations and transient violations of the second law: Perspectives in theory and experiments," *Seminaire Poincaré* **2**, (Birkhäuser, Verlag, Basel, 2003), p. 193.
- <sup>50</sup>We do stress that the SPA models do accurately capture many statistical features of the true data (see the Supplementary Material).
- <sup>51</sup>C. Jarzynski, *Phys. Rev. E* **73**, 046105 (2006).
- <sup>52</sup>Of course, many other factors are relevant in determining the dynamical response, but the basic idea remains the same.
- <sup>53</sup>R. Chelli, S. Marsili, and P. Procacci, arXiv:0711.2726.
- <sup>54</sup>W. Humphrey, A. Dalke, and K. Schulten, *J. Mol. Graphics* **14**, 33 (1996).