Eigenvector method for umbrella sampling enables error analysis

Erik H. Thiede,^{1,a)} Brian Van Koten,^{2,a)} Jonathan Weare,^{2,b)} and Aaron R. Dinner^{1,c)}

¹*Department of Chemistry and James Franck Institute, The University of Chicago, Chicago, Illinois 60637, USA*

²*Department of Statistics and James Franck Institute, The University of Chicago, Chicago, Illinois 60637, USA*

(Received 14 March 2016; accepted 28 July 2016; published online 30 August 2016)

Umbrella sampling efficiently yields equilibrium averages that depend on exploring rare states of a model by biasing simulations to windows of coordinate values and then combining the resulting data with physical weighting. Here, we introduce a mathematical framework that casts the step of combining the data as an eigenproblem. The advantage to this approach is that it facilitates error analysis. We discuss how the error scales with the number of windows. Then, we derive a central limit theorem for averages that are obtained from umbrella sampling. The central limit theorem suggests an estimator of the error contributions from individual windows, and we develop a simple and computationally inexpensive procedure for implementing it. We demonstrate this estimator for simulations of the alanine dipeptide and show that it emphasizes low free energy pathways between stable states in comparison to existing approaches for assessing error contributions. Our work suggests the possibility of using the estimator and, more generally, the eigenvector method for umbrella sampling to guide adaptation of the simulation parameters to accelerate convergence. *Published by AIP Publishing.* [<http://dx.doi.org/10.1063/1.4960649>]

I. INTRODUCTION

One of the main uses of molecular simulations is the calculation of equilibrium averages. For understanding reaction processes, the free energy projected onto selected coordinates (collective variables) is of special interest. It relates directly to the probabilities of the coordinates taking particular values, and it can provide valuable information about the stable states, the barriers between them, and the origin of their stabilization. Furthermore, it is the starting point for most rate theories. Although in principle the free energy can be estimated from a long unbiased simulation, in practice doing so is challenging because bottlenecks slow the exploration of the configuration space. In other words, transitions between regions of the space are very infrequent in comparison to local fluctuations.

Various methods have been introduced to overcome this problem. Here, we consider one of the oldest and still most widely used such methods, umbrella sampling (US).^{1,2} In this approach, the collective-variable interval of interest is covered by a series of simulations, in each of which the system is biased such that sampling is restricted to a relatively narrow window of values of the collective variables. This can be accomplished by addition of a biasing potential that is small in the window and large outside it. The information from the different simulations must be combined, and the effect of the bias removed, to obtain the overall free energy profile. This requires consistently normalizing the probabilities in different windows, a task that is complicated by the fact that the simulations are run independently.

Considerable effort has been devoted to determining how best to combine the results from different simulations. Initially, researchers manually adjusted the zero of free energy in each window to make the full free energy profile continuous and, often, smooth; conflicting results arising from limited sampling at the window peripheries were removed. The desire to use all the simulation data motivated the introduction of estimators that allow for systematically combining the data from different simulations. By far, the most widely used of these in chemical physics applications is the weighted histogram analysis method (WHAM). The multistate Bennett acceptance ratio (MBAR) method, as it is referred to in the molecular-simulation literature and will be referred to here, is closely related but does not rely on binning the data.^{3–6} Both WHAM and MBAR can be derived from maximum-likelihood or minimum asymptotic variance principles, assuming independent, identically distributed sampling in each window, and have corresponding statistical optimality properties under those conditions. Recent extensions seek to improve performance when the sampling is limited and to extend the algorithm to more general ensembles.^{7,8}

In the present paper, we introduce an alternative scheme for estimating the free energy from US simulation data. In this approach, the normalization constants needed to combine information from separate simulation windows are the components of the eigenvector of a stochastic matrix that can be constructed from running averages in the windows. We thus term our method Eigenvector Method for Umbrella Sampling (EMUS). The advantage of our method is that it lends itself to error analysis. Following previous work,^{9–11} we measure error with the asymptotic variance.

Our paper is organized as follows. After giving some background on US in Section II, we formulate EMUS in

^{a)}E. H. Thiede and B. Van Koten contributed equally to this work.

^{b)}Electronic mail: weare@uchicago.edu

^{c)}Electronic mail: dinner@uchicago.edu

Section III. In Sections IV and V, we show that EMUS performs comparably to WHAM and MBAR, and discuss its connection with the latter. In Section VI, we use scaling arguments with simplifying assumptions to show that accounting for the error associated with combining the data is important and limits the speedups that can be achieved by increasing the number of simulation windows. In Section VII, we provide the full numerical analysis, which applies generally, without simplifying assumptions. Specifically, we derive a central limit theorem for averages from EMUS and use it to develop a means for estimating the error contributions from individual windows. We demonstrate the method for the free energy projected onto the ϕ and ψ dihedral angles of the alanine dipeptide and compare the error contributions with those from an estimator introduced by Zhu and Hummer.¹² We conclude in Section VIII.

II. BACKGROUND ON UMBRELLA SAMPLING

Here, we review umbrella sampling and establish basic terms and notation. The goal is the calculation of an average of an observable g over a time-independent probability distribution π ,

$$\langle g \rangle = \int g(x) \pi(x) dx. \quad (1)$$

At thermal equilibrium, π is the Boltzmann distribution,

$$\pi(x) = \frac{\exp(-H_0(x)/k_B T)}{\int \exp(-H_0(x)/k_B T) dx}, \quad (2)$$

where H_0 is the system Hamiltonian, k_B is Boltzmann's constant, and T is the temperature. In particular, we can express the free energy difference between two states S_1 and S_2 as

$$\Delta G = -k_B T \ln \left(\frac{\langle \mathbb{1}_{S1} \rangle}{\langle \mathbb{1}_{S2} \rangle} \right), \quad (3)$$

where $\mathbb{1}$ is the indicator function

$$\mathbb{1}_S(x) = \begin{cases} 1 & \text{if } x \in S, \text{ and} \\ 0 & \text{otherwise.} \end{cases} \quad (4)$$

Similarly, the reversible work to constrain a collective variable $q(x)$ to a particular value q' , also known as the potential of mean force (PMF), may be written as

$$W(q') = -k_B T \ln \langle \delta(q - q') \rangle. \quad (5)$$

For complex systems, averages of the form in (1) must be evaluated numerically. Typically, this is done by generating a chain of related configurations, X_t , using Metropolis Monte Carlo methods or molecular dynamics, and by assuming ergodicity. Namely, as the number of configurations N goes to infinity, $\langle g \rangle$ is the limit of the sample mean,

$$\bar{g} = \frac{1}{N} \sum_{t=0}^{N-1} g(X_t). \quad (6)$$

In all practical sampling methods, successive configurations are strongly correlated. While ergodicity guarantees that

sample means converge to averages over π , convergence can be extremely slow if the correlation between subsequent points is strong. This is the case when sampling π relies on visiting low-probability states, such as transition states of chemical reactions.

US methods address this issue by enforcing sampling of different regions of configuration space (windows), introducing L nonnegative bias functions ψ_i and then using L independent simulations to sample from the biased probability distributions,

$$\pi_i(x) = \frac{\psi_i(x) \pi(x)}{\int \psi_i(x) \pi(x) dx}. \quad (7)$$

The essential idea is that sampling each π_i is fast because ψ_i is chosen so that relatively likely states under ψ_i are not separated by relatively unlikely states. This is accomplished by restricting the set of states on which ψ_i is non-negligible so that π is closer to constant on that set. In Section VI C we make this point more carefully by examining a regime in which umbrella sampling can be shown to be exponentially more efficient than direct simulation. A popular choice is to use bias functions that take a Gaussian form,

$$\psi_i(q) = \exp \left(-\frac{1}{2} k_i (q - q_i^0)^2 / k_B T \right), \quad (8)$$

such that

$$\pi_i(x) \propto \exp \left[- \left(H_0(x) + \frac{1}{2} k_i (q - q_i^0)^2 \right) / k_B T \right]. \quad (9)$$

This corresponds to adding a harmonic potential centered at q_i^0 with spring constant k_i to the system Hamiltonian. We call the relative normalization constant of the i th biased distribution z_i ,

$$z_i = \frac{\int \psi_i(x) \pi(x) dx}{\sum_{k=1}^L \int \psi_k(x) \pi(x) dx}. \quad (10)$$

We also define the free energy in window i as

$$G_i = -k_B T \ln z_i. \quad (11)$$

We denote averages over the biased distributions by

$$\langle g \rangle_i = \int g(x) \pi_i(x) dx. \quad (12)$$

Overall averages of interest, $\langle g \rangle$, can be estimated as z_i -weighted sums of averages computed in each of the windows. We detail our prescription in Section III.

III. THE EIGENVECTOR METHOD FOR UMBRELLA SAMPLING

In this section, we present the Eigenvector Method for Umbrella Sampling (EMUS). We begin by defining

$$g^* \equiv \frac{g}{\sum_{k=1}^L \psi_k}. \quad (13)$$

for any function g . Then, we observe that

$$\begin{aligned}
\langle g \rangle &= \int g(x) \pi(x) dx \\
&= \int g(x) \left\{ \frac{\sum_{i=1}^L \psi_i(x) \left[\frac{\int \psi_i(x) \pi(x) dx}{\int \psi_i(x) \pi(x) dx} \right]}{\sum_{k=1}^L \psi_k(x)} \right\} \pi(x) dx \\
&= \sum_{i=1}^L \int \psi_i(x) \pi(x) dx \frac{\int g^*(x) \psi_i(x) \pi(x) dx}{\int \psi_i(x) \pi(x) dx} \\
&= \sum_{i=1}^L z_i \left(\sum_{k=1}^L \int \psi_k(x) \pi(x) dx \right) \langle g^* \rangle_i. \quad (14)
\end{aligned}$$

The factor in parentheses can be taken out of the sum over i . To express this factor in terms of computable averages, we repeat the same steps with $g = 1$,

$$\sum_{k=1}^L \int \psi_k(x) \pi(x) dx = \frac{1}{\sum_{i=1}^L z_i \langle 1^* \rangle_i}. \quad (15)$$

Substituting (15) into (14),

$$\langle g \rangle = \frac{\sum_{i=1}^L z_i \langle g^* \rangle_i}{\sum_{i=1}^L z_i \langle 1^* \rangle_i}. \quad (16)$$

Consequently, if we can evaluate the z_i , the $\langle 1^* \rangle_i$, and the $\langle g^* \rangle_i$, then we can assemble the original average $\langle g \rangle$ of interest. The averages $\langle g^* \rangle_i$ can be computed from sequences X_t^i (typically independent for each i) that sample the π_i . Umbrella sampling methods differ primarily in how the z_i are computed.

To express the constants z_i in terms of averages over the biased distributions, we take $g(x) = \psi_j(x)$ in (14). Then, z_i solves

$$z_j = \sum_{i=1}^L z_i F_{ij},$$

where

$$F_{ij} = \langle \psi_j^* \rangle_i. \quad (17)$$

That is, the vector of normalization constants z is a left eigenvector of the matrix F with eigenvalue one. Under conditions to be elaborated in Section III B, the solution to (17) is uniquely specified when we notice that

$$\sum_{i=1}^L z_i = 1. \quad (18)$$

A. Computational procedure

In the EMUS algorithm, we estimate the entries of F and the averages $\langle g^* \rangle_i$ and $\langle 1^* \rangle_i$ by sample means, then assemble the estimate of $\langle g \rangle$ using (16). To be precise, we denote the sample means by

$$\bar{g}_i^* = \frac{1}{N_i} \sum_{t=0}^{N_i-1} \frac{g(X_t^i)}{\sum_k \psi_k(X_t^i)}, \quad (19)$$

$$\bar{1}_i^* = \frac{1}{N_i} \sum_{t=0}^{N_i-1} \frac{1}{\sum_k \psi_k(X_t^i)}, \text{ and} \quad (20)$$

$$\bar{F}_{ij} = \frac{1}{N_i} \sum_{t=0}^{N_i-1} \frac{\psi_j(X_t^i)}{\sum_k \psi_k(X_t^i)}. \quad (21)$$

EMUS proceeds as follows:

1. Choose the biasing functions ψ_i .
2. Compute trajectories that sample states X_t^i from the biased distributions π_i .
3. Calculate the matrix $\bar{F}$ and the averages $\bar{g}_i^*$ and $\bar{1}_i^*$.
4. Calculate the vector of estimated normalization constants z^{EMUS} as the solution to

$$z_j^{\text{EMUS}} = \sum_{i=1}^L z_i^{\text{EMUS}} \bar{F}_{ij} \quad \text{with} \quad \sum_{i=1}^L z_i^{\text{EMUS}} = 1. \quad (22)$$

We use QR factorization as in Golub and Meyer.¹³ See also Section VII B 1.

5. Compute the estimate of $\langle g \rangle$,

$$\langle g \rangle^{\text{EMUS}} = \frac{\sum_{i=1}^L z_i^{\text{EMUS}} \bar{g}_i^*}{\sum_{i=1}^L z_i^{\text{EMUS}} \bar{1}_i^*} \quad (23)$$

by substituting z^{EMUS} and the sample means in (16).

We have provided an implementation of this algorithm online, along with implementations of the iteration described in Section IV and the asymptotic error estimate found in Section VII B 1.¹⁴ We remark that when one wishes to compute a free energy difference or a ratio of two observables, as in Equation (3), it is not necessary to compute the 1_i^* . Instead, one may use the formula

$$\frac{\langle g_1 \rangle}{\langle g_2 \rangle} = \frac{\sum_{i=1}^L z_i \langle g_1 \rangle_i}{\sum_{i=1}^L z_i \langle g_2 \rangle_i} \quad (24)$$

where g_1 and g_2 are arbitrary functions.

B. The eigenvector problem

In this section, we give conditions under which the eigenvector problem has a unique solution. First, we show that F is a stochastic matrix; that is, each element F_{ij} is nonnegative and every row of F sums to one:

$$\sum_{j=1}^L F_{ij} = \sum_{j=1}^L \left\langle \frac{\psi_j}{\sum_{k=1}^L \psi_k} \right\rangle_i = \left\langle \frac{\sum_{j=1}^L \psi_j}{\sum_{k=1}^L \psi_k} \right\rangle_i = 1. \quad (25)$$

The entries of F are nonnegative since we require that the bias functions be nonnegative. One can show that the matrix $\bar{F}$ is also stochastic by similar arguments.

A stochastic matrix J has a unique eigenvector with eigenvalue one if it is irreducible: for every possible grouping of the indices into two distinct sets, A and B , $J_{ij} \neq 0$ for some $i \in A$ and $j \in B$.¹⁵ In fact, this statement remains true when J is nonnegative with largest eigenvalue equal to one. For any such matrix we let $z(J)$ denote the continuous function returning the unique left eigenvector of J corresponding to eigenvalue one.

In the case of the particular stochastic matrix F defined in (17) these statements imply that if, for any division of the indices into sets A and B , there is a sufficient overlap between the sets $\cup_{i \in A} \{x : \psi_i(x) > 0\}$ and $\cup_{j \in B} \{x : \psi_j(x) > 0\}$ then there will be a unique solution $z(F)$ to (17) which necessarily equals the relative normalization constants z defined in (10). Because $z(J)$ is a continuous

function of its arguments, $z^{\text{EMUS}} = z(\bar{F})$ converges to z as $\bar{F}$ converges to F . Consequently, EMUS produces a consistent estimator in the sense that if the sample averages used to estimate the entries F_{ij} and $\langle g^* \rangle_i$ converge (in probability or with probability one) to the true values, then the estimate of $\langle g \rangle$ also converges (in the same sense).

IV. THE CONNECTION BETWEEN EMUS AND MBAR

Building upon earlier work in the statistics literature,^{3,4,16} Shirts and Chodera⁶ suggested a class of algorithms for estimating free energy differences between states, which they termed MBAR. This method is similar to WHAM but does not require binning the simulation data to form histograms (see Tan *et al.*⁹). In this section, we explain the relation between EMUS and MBAR.⁶ We also derive a new iterative method for solving the MBAR equations, and we show that our iteration leads naturally to a new family of related consistent estimators.

The starting point of Shirts and Chodera⁶ is the identity (see their (5))

$$z_j \sum_{i=1}^L \langle \alpha_{ij}(x) \psi_i(x) \pi(x) \rangle_j = \sum_{i=1}^L z_i \langle \alpha_{ij}(x) \psi_j(x) \pi(x) \rangle_i, \quad (26)$$

where $\alpha_{ij}(x)$ is an arbitrary function. They proposed the choice

$$\alpha_{ij}^{\text{MBAR}}(x) = \frac{n_i/z_i}{\sum_k \psi_k(x) \pi(x) n_k/z_k}, \quad (27)$$

where n_i is the number of uncorrelated samples in window i . Substituting (27) into (26) gives

$$z_j = \sum_{i=1}^L z_i \left\langle \frac{\psi_j(x) n_i/z_i}{\sum_k \psi_k(x) n_k/z_k} \right\rangle_i. \quad (28)$$

We can cast (28) in a form reminiscent of EMUS by writing

$$z_j = \sum_{i=1}^L z_i F_{ij}(z), \quad (29)$$

where

$$F_{ij}(w) = \left\langle \frac{\psi_j(x) n_i/w_i}{\sum_k \psi_k(x) n_k/w_k} \right\rangle_i \quad (30)$$

for any vector w with positive entries. EMUS corresponds to setting $w = n$ so that

$$\alpha_{ij}^{\text{EMUS}}(x) = \frac{1}{\sum_k \pi(x) \psi_k(x)}, \quad (31)$$

and (26) reduces to the eigenproblem (17).

In practice, one must replace the matrix $F_{ij}(w)$ in (30) by the sample mean approximation

$$\bar{F}_{ij}(w) = \frac{1}{N_i} \sum_{t=0}^{N_i-1} \left[\frac{\psi_j(X_t^i) n_i/w_i}{\sum_k \psi_k(X_t^i) n_k/w_k} \right]. \quad (32)$$

Substituting $\bar{F}_{ij}(z)$ for $F_{ij}(z)$ in (29) yields the equation

$$z_j^{\text{MBAR}} = \sum_{i=1}^L z_i^{\text{MBAR}} \bar{F}_{ij}(z^{\text{MBAR}}) \quad (33)$$

for z^{MBAR} , which we refer to here as the MBAR estimator. If the samples X_t^i are independent, MBAR is the nonparametric maximum-likelihood estimator of z .³

In practice, the samples X_t^i are not independent for a given i , and the n_i must be estimated from data. Several algorithms for estimating the n_i have been proposed.^{17–19} Shirts and Chodera⁶ base their estimates on the integrated autocorrelation times of physically motivated coordinates, and we follow this common practice here. In fact, once the n_i have been estimated, Shirts and Chodera⁶ suggest replacing sample averages over all N_i points by sample averages over the n_i points obtained by including only every N_i/n_i -th sample along the trajectory. We note that both the subsampling approach and the one in (32) correspond to approximations of expression (26) with (27), and we regard both as variations on the MBAR estimator. When the samples are independent, the two approaches are the same. In tests of the iterative EMUS algorithm introduced below, we find estimates to be insensitive to the choice of n_i and they can be set equal to 1, though in that case the estimator no longer corresponds directly to MBAR.

As written above, the MBAR estimator (33) resembles an eigenvector problem. However, the dependence of $\bar{F}(z)$ on z implies that the solution must be obtained self-consistently. The approach advocated by Shirts and Chodera for computing the MBAR estimator corresponds in the framework described here to solving (33) by a Newton-type iteration. However, the eigenvector form of (33) suggests an alternative approach. Rather than Newton's method, we employ the following algorithm:

1. As an initial guess for z^{MBAR} , choose a vector z^0 with positive entries. Estimate the n_i . Set $m = 0$.
2. (a) Calculate $\bar{F}_{ij}(z^m)$ according to (32).
(b) Calculate a new estimate z^{m+1} of z^{MBAR} by solving the eigenproblem

$$z_j^{m+1} = \sum_{i=1}^L z_i^{m+1} \bar{F}_{ij}(z^m). \quad (34)$$

3. If $\max_i |z_i^{m+1} - z_i^m|/z_i^m > \text{Tolerance}$,
(a) Increment m ;
(b) Go to Step 2.

A similar algorithm was proposed by Meng and Wong.²⁰

To show that this iteration makes sense, we must prove that the eigenproblem (34) always has a unique solution and that z^m converges to z^{MBAR} as m goes to infinity. To see that the eigenproblem has a solution, first observe that if $\bar{F}_{ij}(w)$ is irreducible for one vector with positive entries, w , then it is irreducible for all vectors with positive entries. When applying the EMUS method, we thus assume that $\bar{F}(w)$ is irreducible. Moreover, observe that for any positive vector w , the vector with entries n_i/w_i is a right eigenvector of $\bar{F}(w)$ with eigenvalue one and positive entries. It follows from the Perron-Frobenius theorem that the matrix $\bar{F}(w)$ has a unique left eigenvector $z(\bar{F}(w))$ with eigenvalue one and

that $z(\bar{F}(w))$ has positive entries. Thus, the eigenproblem always has a unique solution. We do not have a proof that the iterates converge. However, since z^{MBAR} is a fixed point of the iteration, if the iterates do converge, their limit must be z^{MBAR} . In practice, we find that the iteration converges quickly, usually to a relative error of 10^{-6} within 10 iterates.

In addition to its apparently rapid convergence, another argument in favor of the algorithm that we introduce above for solving (33) is that each iteration of the scheme results in a new consistent estimator. We will use the term iterative EMUS to refer to this family of estimators. With the initial guess $z^0 = n$, the result, z^1 , of the first iteration is the EMUS estimator defined in Section III. In the Appendix, we show that for any fixed finite number of iterations m , z^m is also a consistent estimator of the vector z of normalization constants. By contrast, other schemes, such as Newton's method, for solving (33) may require that the number of iterations goes to infinity to obtain a consistent estimate. We also remark that the consistency result in the Appendix holds as long as the n_i converge to non-random, positive values with increasing numbers of samples N_i . They can be chosen as described above, or simply set to a fixed value.

Differences between the iterative EMUS scheme above and the application of Newton's method proposed by Shirts and Chodera⁶ are mostly matters of implementation. As we will see in Section V, the results are not very sensitive to these computational details; most of the accuracy in the iterative EMUS approach is achieved in the first step. In any case, we remind the reader that the primary goal of this paper is to characterize those properties of the broader umbrella sampling approach that are essential to its success, not to analyze details of implementation.

While we focus here on potentials of mean force, the MBAR estimator has been applied to a broader category of free energy problems, including the analysis of single-molecule pulling experiments and alchemical free energy calculations.^{6,21,22} The close relation between EMUS and MBAR indicates that error analysis of EMUS may provide insight into the sources of error in MBAR for these problems, but we do not pursue this idea further in the present work.

V. NUMERICAL COMPARISON

To test the algorithm numerically, we performed 100 independent umbrella sampling calculations for the PMF of the ϕ coordinate of the alanine dipeptide (i.e., *N*-acetyl-alanyl-*N'*-methylamide) in vacuum. Simulations were run using GROMACS version 5.1.1 with harmonic bias potentials applied using the PLUMED 2.2.0 software package.^{23,24} The molecule was represented by the CHARMM 27 force field without CMAP corrections,²⁵ with covalent bonds to hydrogen atoms constrained by the LINCS algorithm.²⁶ Twenty windows were evenly spaced along the ϕ dihedral angle. The force constant $k_i = 0.007\,605\,35 \times 10^{-2}$ kcal mol⁻¹ deg⁻² such that the standard deviation of the Gaussian bias functions was 9°. In each window, we integrated the equations of motion with the GROMACS leap-frog Langevin integrator with a 1 fs time step and a time constant of 0.1 ps. The temperature

was 300 K. The system was equilibrated for 40 ps and then sampled for 100 ps, saving structures every 10 fs.

The data were then analyzed with EMUS, Grossfield's implementation of WHAM,²⁷ and the algorithm proposed by Zhu and Hummer (ZH) (see Equation (A1) and the discussion following it in Appendix A of that paper¹²). The data were

FIG. 1. Comparison of umbrella sampling methods applied to simulation data for the alanine dipeptide. (a) Average window free energies, G_i , for the indicated methods. Error bars are estimated standard deviations of the means. (b) Standard deviation of each method relative to that of the WHAM algorithm. Colors are the same as in (a). (c) EMUS as the first step in a self-consistent iteration to solve the MBAR equations (see text). The number of uncorrelated samples in each window (n_i) was estimated by calculating the integrated autocorrelation of the ϕ dihedral angle from each trajectory. Results shown are for identical molecular dynamics data (see text for simulation details); the methods differ only with respect to combination of the data to estimate the free energies.

also analyzed with pyMBAR;⁶ as pyMBAR gave results virtually identical to WHAM, the results are not shown. In Figure 1, we show the resulting average potentials of mean force, as well as the standard deviation of the estimates over the 100 runs. WHAM and EMUS converge to the same result. This is to be expected, as both algorithms are consistent (i.e., they converge to the exact result as the amount of samples in each window tends to infinity; see Section VII), although WHAM exhibits a small bias from the binning of data for the histograms.⁶ The standard deviations of the free energies are slightly higher for EMUS than for the other two algorithms, but these differences are negligible compared with those expected for physically weighted simulations. Moreover, the relative performances are likely to be problem dependent. We note that ZH is based on thermodynamic integration, and the finite number of integration points causes quadrature error,²⁸ and, in turn, a systematic error in the barrier height.

In Figure 1(c), we apply the self-consistent iteration described in Section IV. For this calculation, we estimate the number of independent samples in each window (n_i) from the integrated autocorrelation time of the ϕ dihedral angle time series. We plot the standard deviation of the values of z calculated after the first iteration (EMUS), the second iteration, and after convergence to a relative residual smaller than 10^{-6} . In general, convergence is achieved after an average of 9 iterations; none of the 100 data sets required more than 15 iterations. However, we note that after two iterations, the estimates of z already have a standard deviation equivalent to that of the WHAM algorithm. In this article we focus on the scheme corresponding to the first iteration only and do not attempt to analyze the improvement due to multiple iterations. In our tests the performance gain from multiple iterations is negligible compared to the improvement over direct approximation of free energy differences using long unbiased trajectories.

VI. JUSTIFICATION FOR UMBRELLA SAMPLING BY SCALING ARGUMENTS

The quality of a statistical estimate from umbrella sampling depends strongly on the choices made for the simulation windows. In this section, we discuss how the error scales as properties of the simulation change. We begin in Section VI A with a description of a prevalent justification for the use of US. We show in Section VI B that this argument is incomplete and, in turn, misleading. In Section VI C, we provide an alternative justification; namely, we show that in the low temperature limit, the cost to achieve a fixed accuracy by US grows slowly compared to direct simulation. In this section we make several simplifying assumptions that allow us to draw precise conclusions about the scaling properties of EMUS. In Section VII we provide error bounds for EMUS under much more general assumptions.

A. Scaling in the limit of many windows

To justify umbrella sampling, it is often suggested that the total computational time required to accurately sample

statistics is inversely proportional to the number of windows, L .^{17,29–32} The argument for this scaling proceeds as follows.

- Divide a one-dimensional collective variable space into L windows of equal length, inversely proportional to L (i.e., L^{-1}).
- Assume the windows are small enough that no free energy barriers exist in each window. The time to explore a window should be diffusion limited and proportional to the length of the window squared. Therefore, the simulation time required to accurately sample statistics in one window is also proportional to L^{-2} .
- Because there are L windows, the total simulation time required to compute averages to fixed accuracy should scale as $L \times L^{-2} = L^{-1}$.

While this argument is now standard,^{17,29–32} Virnau and Müller³³ observed that the error for computing the free energy difference between phases of Lennard-Jones particles with an approximately fixed amount of sampling was insensitive to the number of windows in practice, and they noted that the argument above neglects the error associated with combining the data from different simulation windows. Nguyen and Minh recently made a similar suggestion for a related class of methods.³⁴ This intuition is supported by our analysis in Section VI B, which shows that the total computational cost to achieve a fixed accuracy should be at best insensitive to the choice of L , so long as it is sufficiently large.

B. A simple model problem

To perform a more precise analysis, we make a number of simplifying assumptions. We emphasize that these assumptions are in force only for the purposes of the scaling arguments in this section. We provide more general error bounds for EMUS in Section VII.

Assumption 6.1. The total computation time, N , is divided equally among the windows such that $N_i = N/L$.

Assumption 6.2. The ψ_i are functions on the one-dimensional interval $[0, 1]$, and the set of points where ψ_i is non-zero, $\{q : \psi_i > 0\}$, is an interval of length $|\{q : \psi_i > 0\}| \leq \gamma/L$. We also assume that $\psi_i \psi_j = 0$ unless $|j - i| \leq 1$. Consequently, both the exact matrix F and the sample mean $\bar{F}$ are tri-diagonal. This assumption clearly does not hold when the bias functions ψ_i are Gaussian. Nonetheless, the rapid decay of Gaussian bias functions away from their peaks guarantees that entries of F and $\bar{F}$ far from the diagonal are very small, such that we expect our conclusions to still hold (though their justification would be more complicated).

Assumption 6.3. The overlap of ψ_i and $\psi_{i\pm 1}$ (i.e., the integral of their product) is large enough that

$$\min\{F_{i,i+1}, F_{i,i-1}\} > \delta > 0 \quad (35)$$

for all L and for all $i \leq L$. If our last assumption holds, but this one does not, then we can find more than one vector z satisfying Equations (17) and (18). This assumption is a

slightly stronger version of the notion of irreducibility that we defined earlier (see Section III B). Note that we require the irreducibility to hold uniformly in the large L limit, and we thus introduce the δ , which is independent of L .

Assumption 6.4. Sample averages computed in different windows are independent, i.e., $\bar{F}_{i,i\pm 1}$ and $\bar{F}_{j,j\pm 1}$ for $j \neq i$ are independent. We do not assume (here or anywhere else in this paper) that samples generated within a single window are independent. Indeed, even if the samples from π_i are independent, $\bar{F}_{i,i+1}$ and $\bar{F}_{i,i-1}$ are dependent random variables.

As an example average, let us consider the error in the free energy difference between the first and last windows,

$$\Delta\bar{G}_{L,1} = -k_B T \ln \left(\frac{z_L}{z_1} \right). \quad (36)$$

Assumption 6.2 is sufficient for $z(\bar{F})$ to be in detailed balance with $\bar{F}$ (Kelly Ref. 35, Lemma 1.5 and Section 1.3),

$$z_{i+1}(\bar{F}) = z_i(\bar{F}) \frac{\bar{F}_{i,i+1}}{\bar{F}_{i+1,i}}. \quad (37)$$

Using (37) recursively,

$$\begin{aligned} \Delta\bar{G}_{L,1} &= -k_B T \ln \left(\frac{\prod_{i=1}^{L-1} \bar{F}_{i,i+1}}{\prod_{i=1}^{L-1} \bar{F}_{i+1,i}} \right) \\ &= -k_B T \ln \bar{F}_{1,2} + k_B T \ln \bar{F}_{L,L-1} \\ &\quad + k_B T \sum_{i=2}^{L-1} \ln \left(\frac{\bar{F}_{i,i+1}}{\bar{F}_{i,i-1}} \right). \end{aligned} \quad (38)$$

To understand the error (variance) of the terms in (38), we must further specify $F_{i,i+1}$ and $F_{i,i-1}$.

Assumption 6.5. For N_{min} and L_{min} sufficiently large, when $N_i \geq N_{min}$ and $L \geq L_{min}$,

$$\frac{K_{min}}{N_i L^2} \leq \text{var} \left(\ln \left(\frac{\bar{F}_{i,i+1}}{\bar{F}_{i,i-1}} \right) \right) \leq \frac{K_{max}}{N_i L^2} \quad (39)$$

for $i = 2, 3, \dots, L-1$, and the same upper and lower bounds hold for $\text{var}(\ln F_{1,2})$ and $\text{var}(\ln F_{L,L-1})$. This is just a precise interpretation of the diffusion limited sampling assumption made in the standard justification of US reproduced in Section VI A. Under such an assumption we expect both $\bar{F}_{i,i+1}$ and $\bar{F}_{i,i-1}$ to have variance on the order of $1/(N_i L^2)$ and, in light of (35), the function $\ln(x/y)$ is smooth near $(x, y) = (F_{i,i+1}, F_{i,i-1})$. These considerations are closely related to Lemma 7.2 in Section VII A.

With all the assumptions in hand, we now complete the argument by taking the variance of both sides of (38). Since samples from different windows are independent, the variance of $\bar{G}_{L,1}$ is a sum of contributions from each window,

$$\begin{aligned} \text{var}(\Delta\bar{G}_{L,1}) &= \text{var}(k_B T \ln \bar{F}_{1,2}) + \text{var}(k_B T \ln \bar{F}_{L,L-1}) \\ &\quad + \sum_{i=2}^{L-1} \text{var} \left(k_B T \ln \left(\frac{\bar{F}_{i,i+1}}{\bar{F}_{i,i-1}} \right) \right). \end{aligned} \quad (40)$$

Using (39) and substituting $N_i = N/L$, we find that, as long as $N/L \geq N_{min}$ and $L \geq L_{min}$,

FIG. 2. The scaling of umbrella sampling error with number of windows on a flat potential. A Brownian particle on a flat, one-dimensional potential was simulated for 480 identical runs, and the free energy difference between the first and last windows was calculated, as described in the text. Here, the mean square error from the exact result is plotted against the number of windows. The lines show the scaling in error predicted by the L^{-1} and L^0 scalings. Fitting the data on a log-log scale give a scaling exponent of -0.026 ± 0.028 .

$$(k_B T)^2 \frac{K_{min}}{N} \leq \text{var}(\Delta\bar{G}_{L,1}) \leq (k_B T)^2 \frac{K_{max}}{N}. \quad (41)$$

To verify that this conclusion carries over to harmonic bias potentials, we performed multiple umbrella sampling calculations for a Brownian particle on a flat potential on the interval $[0, 1]$ with a stepsize of 1.0×10^{-6} and $k_B T = 1.0$ using Gaussian bias functions with a standard deviation of $1/L$. The number of windows was varied from $L = 10$ to 46 in steps of 2. For each value, a total of 10^7 steps were distributed equally in the windows, and the US calculation was repeated 480 times. We then calculated the mean square error of the free energy difference between the first and last window over the 480 replicates and determined how the mean square error scaled with L . Rather than the mean square error varying inversely with the number of windows, the data plotted in Figure 2 support a scaling of L^0 , consistent with (41).

It is worth noting that the inverse scaling with total cost N in (41) is exactly the scaling one would expect for the variance of an estimate of the free energy difference $\Delta\bar{G}_{L,1}$ constructed from a molecular dynamics trajectory of length N . Because US and direct simulations of comparable total numbers of steps require comparable computational effort (ignoring the overhead associated with combining the simulation data, which is typically small in comparison with the computational cost of the sampling), the benefits of US must be encoded in the constants K_{min} and K_{max} . A dramatic demonstration of this observation is the purpose of Section VI C.

C. The low temperature limit

To understand the benefits of umbrella sampling, we must study its performance in the presence of free energy barriers. In particular, we compare the performance of umbrella sampling to physically weighted sampling as the temperature goes

to zero. In this limit, the cost of direct sampling increases exponentially with $1/T$, while, as we show, the cost of umbrella sampling increases only algebraically. A formal discussion is given in a separate publication;³⁶ here, we present a simple plausibility argument.

Owing to the free energy barriers, the assumption of diffusive dynamics in each window no longer holds. Instead, we expect a form typical of reaction rate theories in each window. We define ΔW_i as the maximum difference in the PMF in window i ,

$$\Delta W_i = \max_{\{q: \psi_i(q) > 0\}} \{W(q)\} - \min_{\{q: \psi_i(q) > 0\}} \{W(q)\}. \quad (42)$$

Assumption 6.5': We now replace the upper and lower bounds in (39) by the upper bound

$$\text{var} \left(\ln \left(\frac{\bar{F}_{i,i+1}}{\bar{F}_{i,i-1}} \right) \right) \leq \frac{K}{k_B T N_i L^2} \exp \left(\frac{\Delta W_i}{k_B T} \right) \quad (43)$$

for $i = 2, 3, \dots, L-1$ with analogous replacements for $i = 1$ and $i = L$, as long as $N_i \geq N_{\min}$ and $L \geq L_{\min}$. The constant K here is assumed to be independent of temperature. This bound captures the diffusion limited sampling assumption when L is very large, but is more detailed than (39) in that it captures (crudely) the increasing difficulty of the sampling problem as the temperature decreases with all other parameters held fixed. Under reasonable additional assumptions on the underlying potential, the bias functions ψ_i , and the sampling scheme, one can rigorously establish an asymptotic (large N_i) bound of the form in (43).³⁶

Substituting this new bound into (40), we find that, if $L \geq L_{\min}$ and $N/L \geq N_{\min}$, then

$$\text{var}(\Delta \bar{G}_{L,1}) \leq \frac{K k_B T}{N L} \sum_{i=1}^L \exp \left(\frac{\Delta W_i}{k_B T} \right). \quad (44)$$

As the temperature decreases, we choose to increase L such that $\Delta W_i/k_B T$ is bounded above. This can be achieved by scaling L linearly with $1/T$: if the derivative of the PMF is bounded (in absolute value) by $W'_{\max}$, choosing L so that

$$L \geq \frac{W'_{\max}}{\Omega k_B T} \quad (45)$$

ensures that $\Delta W_i/k_B T$ is bounded by Ω (since we have assumed that the argument of W is in $[0, 1]$). On the other hand, our assumption that the length of $\{q : \psi_i > 0\}$ (Assumption 6.2) does not exceed γ/L implies that

$$\frac{\Delta W_i}{k_B T} \leq \frac{W'_{\max} \gamma}{k_B T L}. \quad (46)$$

Consequently, as long as (45) holds,

$$\frac{\Delta W_i}{k_B T} \leq \Omega \gamma. \quad (47)$$

Finally, substituting this result into (44) we find that if $L \geq L_{\min}$ and $N/L \geq N_{\min}$, then

$$\text{var}(\Delta \bar{G}_{L,1}) \leq \frac{K k_B T \exp(\Omega \gamma)}{N} \leq \frac{K L k_B T \exp(\Omega \gamma)}{N_{\min}}. \quad (48)$$

With the best possible (smallest) choice of L allowed by (45), this bound becomes

$$\text{var}(\Delta \bar{G}_{L,1}) \leq \frac{K W'_{\max} \exp(\Omega \gamma)}{\Omega N_{\min}}. \quad (49)$$

The remarkable feature of the bound in (49) is that it is independent of T . This does not mean that the cost to achieve a fixed accuracy is independent of T . However, it does imply that as the temperature is decreased, we do not have to increase $N_{\min}$ to maintain a fixed accuracy. Expression (45) and the fact that $N_i \geq N_{\min}$ together imply that, under the assumptions of this section, the computational cost of obtaining an accurate estimate of $\Delta \bar{G}_{L,1}$ by US increases algebraically with $(k_B T)^{-1}$. That scaling is to be compared to exponential in $(k_B T)^{-1}$ to achieve the same accuracy by direct simulation.

Finally, we remark that while our analysis provides a convincing explanation for the performance benefits offered by umbrella sampling, it neglects a number of practical realities. First, we have ignored the cost of equilibrating the simulation in each window, which can be challenging. Second, we have ignored the practical difficulties that arise when the number of windows grows large. The bias functions (restraints) can introduce additional free energy barriers that slow mixing in the degrees of freedom orthogonal to the collective variable, and, if sufficiently restrictive, could in principle necessitate modifying the elementary simulation step sizes. Our analysis reveals, however, that these difficulties are *not* responsible for the slow (L^0) error scaling observed for large L .

VII. ANALYSIS OF THE ERROR OF EMUS

In this section, we study the error of EMUS in full generality, without imposing the simplifying assumptions of Section VI. Our main results are a central limit theorem for EMUS (Theorem 7.4) and an easily computed, practical error estimator which reveals the contributions of the different windows to the total error. These results may be used to compare the efficiency of EMUS and other methods and to study how the efficiency of EMUS depends on parameters such as the number of samples allocated to each window.

A. A central limit theorem for EMUS

Before developing the error analysis, we define a single notation for EMUS which incorporates both the case of a free-energy difference and the case of an ensemble average. In either case, one must compute $\bar{F}$ and also $\bar{g}_{1,i}^*$ and $\bar{g}_{2,i}^*$ for two real valued functions g_1 and g_2 . To compute a free energy difference, we choose based on (3),

$$g_1 = \mathbb{1}_{S_1} \text{ and } g_2 = \mathbb{1}_{S_2}. \quad (50)$$

To compute an ensemble average $\langle g \rangle$, we choose based on (16),

$$g_1 = g \text{ and } g_2 = 1. \quad (51)$$

We furthermore define the function

$$v_i(x) = (\psi_1^*(x), \dots, \psi_L^*(x), g_{1,i}^*(x), g_{2,i}^*(x)) \quad (52)$$

so that

$$\bar{v}_i = \frac{1}{N_i} \sum_{t=0}^{N_i-1} v_i(X_t^i) = (\bar{F}_{i1}, \dots, \bar{F}_{iL}, \bar{g}_{1,i}^*, \bar{g}_{2,i}^*), \quad (53)$$

where we remind the reader that, for each i , the process X_t^i samples the biased distribution π_i . Define

$$\bar{v} = (\bar{v}_1, \dots, \bar{v}_L), \quad (54)$$

and let

$$\langle v \rangle = (\langle v_1 \rangle_1, \dots, \langle v_L \rangle_L) \quad (55)$$

denote the corresponding vector of exact averages. Using the notation defined in Section III B, the EMUS estimator takes the form $B(\bar{v})$, where for a free-energy difference,

$$B(\bar{v}) = -k_B T \ln \left(\frac{\sum_{i=1}^L z_i(\bar{F}) \bar{g}_{1,i}^*}{\sum_{i=1}^L z_i(\bar{F}) \bar{g}_{2,i}^*} \right), \quad (56)$$

and for an ensemble average,

$$B(\bar{v}) = \frac{\sum_{i=1}^L z_i(\bar{F}) \bar{g}_{1,i}^*}{\sum_{i=1}^L z_i(\bar{F}) \bar{g}_{2,i}^*}. \quad (57)$$

We now proceed with the error analysis. First, we characterize the error of the sample means over the biased distributions. As discussed by Frenkel and Smit (Ref. 17, Appendix D), the variance of a sample mean may be expanded in terms of the integrated autocovariance of the process. We define the autocovariance function of $v_i(X_t^i)$ to be

$$C_i(t) = \langle (v_i(X_0^i) - \langle v_i \rangle_i)(v_i(X_t^i) - \langle v_i \rangle_i)^T \rangle_i, \quad (58)$$

where T denotes a vector transpose, and here the outer $\langle \dots \rangle_i$ denotes the exact average not only over X_0^i sampled from π_i but also subsequent points of the sequence X_t^i . Note that $C_i(t)$ is a $(L+2) \times (L+2)$ matrix. We define the integrated autocovariance to be

$$\Sigma_i = \sum_{t=-\infty}^{\infty} C_i(t). \quad (59)$$

The integrated autocovariance is the leading order coefficient in an expansion of the covariance $\bar{v}_i$ (see Ref. 17, D.1.3),

$$\text{cov}(\bar{v}_i) = \frac{\Sigma_i}{N_i} + o\left(\frac{1}{N_i}\right), \quad (60)$$

where $o(1/N_i)$ denotes terms that go to zero faster than $1/N_i$ (i.e., $N_i o(1/N_i) \rightarrow 0$).

Under certain conditions on the process X_t^i , one can strengthen the expansion of the covariance (60) to a central limit theorem (CLT) for $\bar{v}_i$. We expect such a CLT to hold for most problems and most sampling methods in computational statistical physics. However, to avoid a lengthy and technical digression, we simply take the CLT as an assumption; we justify this assumption in more detail in another work,³⁶ and we refer to the work of Lelièvre *et al.* (Ref. 37, Section 2.3.1.2) for a general discussion of the CLT in the context of computational statistical physics.

Assumption 7.1 (Central limit theorem for $\bar{v}_i$). We assume that

$$\sqrt{N_i}(\bar{v}_i - \langle v_i \rangle) \xrightarrow{d} \mathbf{N}(0, \Sigma_i), \quad (61)$$

where $\Sigma_i \in \mathbb{R}^{(L+2) \times (L+2)}$ is the integrated autocovariance matrix defined in (59). The symbol $\xrightarrow{d}$ denotes convergence in distribution as $N_i \rightarrow \infty$. Notice that when the elements of the sequence X_t^i are independent and drawn from π_i then $\Sigma_i = \langle (v_i - \langle v_i \rangle_i)(v_i - \langle v_i \rangle_i)^T \rangle_i / N_i$. More generally, samples are correlated, so Σ_i includes a factor that accounts for the time to decorrelate.

Having characterized the errors in the sample means, we now study how these errors propagate through the EMUS algorithm. Our goal is to prove a CLT for EMUS. We accomplish this using the delta method.

Lemma 7.2 (The delta method; Proposition 6.2 of Bilodeau and Brenner³⁸). Let θ_N be a sequence of random variables taking values in $\mathbb{R}^d$. Assume that a central limit theorem holds for θ_N with mean $\mu \in \mathbb{R}^d$ and asymptotic covariance matrix $\Sigma \in \mathbb{R}^{d \times d}$; that is, assume

$$\sqrt{N}(\theta_N - \mu) \xrightarrow{d} \mathbf{N}(0, \Sigma). \quad (62)$$

Let $\Phi : \mathbb{R}^d \rightarrow \mathbb{R}$ be a function differentiable at μ . Then we have the central limit theorem

$$\sqrt{N}(\Phi(\theta_N) - \Phi(\mu)) \xrightarrow{d} \mathbf{N}(0, \nabla \Phi^T(\mu) \Sigma \nabla \Phi(\mu)) \quad (63)$$

for the sequence of random variables $\Phi(\theta_N)$.

To motivate the delta method, we observe that if X has distribution $\mathbf{N}(\mu, \Sigma)$, then $\nabla \Phi(\mu)^T X$ has distribution $\mathbf{N}(\Phi(\mu), \nabla \Phi^T(\mu) \Sigma \nabla \Phi(\mu))$. That is, according to the delta method, the asymptotic distribution of $\Phi(X)$ is the linearization of Φ at μ applied to the asymptotic distribution of X . Thus, one may regard the delta method as a rigorous version of the standard error propagation formula based on linearization.

We prove the CLT for EMUS by applying the delta method with $\bar{v}$ taking the place of θ_N and with the function B taking the place of Φ . We require the following assumptions in addition to Assumption 7.1.

Assumption 7.3. We assume the following:

1. *The proportion of the total number of samples drawn from each window is constant in the limit as $N \rightarrow \infty$; that is,*

$$\lim_{N \rightarrow \infty} N_i/N = \kappa_i. \quad (64)$$

2. *Sampling in different windows is independent; that is, $\bar{v}_i$ is independent of $\bar{v}_j$ when $j \neq i$.*
3. *The biasing functions ψ_i are chosen so that F is irreducible; see Section III B.*

We now give the CLT for EMUS.

Theorem 7.4 (Central Limit Theorem for EMUS). *Let Assumptions 7.1 and 7.3 hold. Let*

$$\frac{\partial B}{\partial \bar{v}_i} = \left(\frac{\partial B}{\partial \bar{F}_{i1}}, \dots, \frac{\partial B}{\partial \bar{F}_{iL}}, \frac{\partial B}{\partial \bar{g}_{1,i}^*}, \frac{\partial B}{\partial \bar{g}_{2,i}^*} \right) \in \mathbb{R}^{L+2} \quad (65)$$

denote the partial derivative of B with respect to $\bar{v}_i$. Under the assumptions stated above,

$$\sqrt{N}(B(\bar{v}) - B(\langle v \rangle)) \xrightarrow{d} \mathbf{N}(0, \sigma^2), \quad (66)$$

where

$$\sigma^2 = \sum_{i=1}^L \frac{1}{\kappa_i} \left(\frac{\partial B}{\partial \bar{v}_i} \Sigma_i \frac{\partial B}{\partial \bar{v}_i} \right). \quad (67)$$

We refer to σ^2 as the asymptotic variance of EMUS.

Proof. First, we write down a central limit theorem for $(\bar{v}_1, \dots, \bar{v}_L)$. We have that

$$\sqrt{N}(\bar{v}_i - \langle v_i \rangle) \xrightarrow{d} \mathbf{N}(0, \kappa_i^{-1} \Sigma_i) \quad (68)$$

by Assumption 7.1 and (64). Since the sampling in different windows is assumed to be independent, (68) implies

$$\sqrt{N}(\bar{v} - \langle v \rangle) \xrightarrow{d} \mathbf{N}(0, \Sigma), \quad (69)$$

where $\Sigma \in \mathbb{R}^{L(L+2) \times L(L+2)}$ is the block diagonal matrix

$$\Sigma = \begin{bmatrix} \Sigma_1/\kappa_1 & 0 & 0 & \dots \\ 0 & \Sigma_2/\kappa_2 & 0 & \dots \\ 0 & 0 & \Sigma_3/\kappa_3 & \dots \\ \vdots & \vdots & \vdots & \ddots \end{bmatrix}. \quad (70)$$

Second, we verify that B is differentiable at $\bar{v}$. Since F is assumed to be an irreducible stochastic matrix, $z(\bar{F})$ is differentiable at $\bar{F}$. We refer to Thiede *et al.*³⁹ Lemma 3.1 for a complete explanation. It follows from the chain rule that B is differentiable at $\bar{v}$.

Finally, applying Lemma 7.2 with B playing the role of Φ and $\bar{v}$ the role of θ_N concludes the proof. $\square$

The asymptotic variance σ^2 appearing in Theorem 7.4 measures the rate at which the error of EMUS decreases with the number of samples. To make this precise, we observe that Theorem 7.4 is equivalent to the following asymptotic result concerning confidence intervals. For every $\alpha > 0$,

$$\lim_{N \rightarrow \infty} \mathbf{P} \left[|B(\bar{v}) - B(\langle v \rangle)| \leq \frac{\alpha \sigma}{\sqrt{N}} \right] = \text{erf} \left(\frac{\alpha}{\sqrt{2}} \right), \quad (71)$$

where $\mathbf{P}$ denotes a probability and erf denotes the error function.

The asymptotic variance is commonly used to measure the efficiency of an estimator. We refer to the work of van der Vaart⁴⁰ for an explanation and for a discussion of other possibilities. In Section VII B, we explain how the proportion κ_i of samples allocated to each window may be adjusted to minimize the asymptotic variance of EMUS, thereby maximizing efficiency.

We note that a central limit theorem similar to Theorem 7.4 has been proved for the MBAR estimator by Gill *et al.* (Ref. 4, Proposition 2.2). However, the authors of this work do not study the dependence of the asymptotic variance on the parameters, as we do. In fact, the MBAR estimator is significantly more complicated than EMUS, and

its dependence on the number of windows and the allocation of samples is harder to understand.

We³⁶ use a result similar to Theorem 7.4 to generalize the conclusions of Section VI to periodic and multi-dimensional reaction coordinates and to a wider class of observables than free energy differences. We show both that the asymptotic variance is constant in the limit of large L and that the work required to compute an average to fixed precision increases only algebraically in the low temperature limit. In addition, we use recently developed perturbation estimates for Markov chains³⁹ to quantify the dependence of the asymptotic variance of EMUS on the degree to which the bias functions overlap.

B. Estimating the asymptotic variance of EMUS

Our goal in this section is to derive a computable estimate $\bar{\sigma}^2$ of the asymptotic variance σ^2 , which can be decomposed to assess the contributions from individual windows to errors in averages. We recall that formula (67) for σ^2 involves partial derivatives of B . Our estimate $\bar{\sigma}^2$ of σ^2 requires explicit formulas for these partial derivatives. We provide the appropriate expressions, both for ensemble averages and for free-energy differences, in Lemma 7.5. Following the partial derivatives, we present an algorithm for evaluating $\bar{\sigma}^2$ and demonstrate it for the alanine dipeptide. Finally, we compare with the output of a procedure from Zhu and Hummer (ZH)¹² in Section VII B 3.

Lemma 7.5. We have the following formulas for $\partial B / \partial \bar{v}_i$:

1. When EMUS is used to compute an ensemble average $\langle g \rangle$, B is defined by (57), and we have

$$\frac{\partial B}{\partial \bar{F}_{ij}}(\bar{v}) = \frac{\sum_k z_i(\bar{F})(I - \bar{F})_{jk}^\# (\bar{g}_k^* - B(\bar{v}) \bar{\mathbf{I}}_k^*)}{\sum_k z_k(\bar{F}) \bar{\mathbf{I}}_k^*}, \quad (72)$$

$$\frac{\partial B}{\partial \bar{g}_{1,i}^*}(\bar{v}) = \frac{z_i(\bar{F})}{\sum_k z_k(\bar{F}) \bar{\mathbf{I}}_k^*}, \text{ and} \quad (73)$$

$$\frac{\partial B}{\partial \bar{g}_{2,i}^*}(\bar{v}) = -\frac{B(\bar{v}) z_i(\bar{F})}{\sum_k z_k(\bar{F}) \bar{\mathbf{I}}_k^*}, \quad (74)$$

where $\#$ denotes the group inverse (see Ref. 13 and discussion below).

2. When EMUS is used to compute a free-energy difference, B is defined by (56), and we have

$$\frac{\partial B}{\partial \bar{F}_{ij}}(\bar{v}) = k_B T z_i(\bar{F}) \left(\frac{\sum_k (I - \bar{F})_{jk}^\# \bar{\mathbf{I}}_{S2,k}^*}{\sum_k z_k(\bar{F}) \bar{\mathbf{I}}_{S2,k}^*} - \frac{\sum_k (I - \bar{F})_{jk}^\# \bar{\mathbf{I}}_{S1,k}^*}{\sum_k z_k(\bar{F}) \bar{\mathbf{I}}_{S1,k}^*} \right), \quad (75)$$

$$\frac{\partial B}{\partial \bar{g}_{1,i}^*}(\bar{v}) = k_B T \frac{z_i(\bar{F})}{\sum_k z_k(\bar{F}) \bar{\mathbf{I}}_{S1,k}^*}, \text{ and} \quad (76)$$

$$\frac{\partial B}{\partial \bar{g}_{2,i}^*}(\bar{v}) = -k_B T \frac{z_i(\bar{F})}{\sum_k z_k(\bar{F}) \bar{\mathbf{I}}_{S2,k}^*}. \quad (77)$$

3. When EMUS is used to compute the free energy of the k th window, B is defined by (11), and we have

$$-k_B T \frac{\partial B}{\partial \bar{F}_{ij}}(\bar{v}) = \frac{z_i(\bar{F})}{z_k(\bar{F})} (I - \bar{F})_{jk}^\# \text{ and} \quad (78)$$

$$\frac{\partial B}{\partial \bar{g}_{i,1}^*}(\bar{v}) = \frac{\partial B}{\partial \bar{g}_{i,2}^*}(\bar{v}) = 0. \quad (79)$$

Note that for a free energy difference between windows, we can simply subtract derivatives for the corresponding windows.

Proof. We begin by reminding the reader that the output of EMUS is the vector of window normalization constants, z , which depends on the sample mean $\bar{F}$. Because all other averages and, in turn, their derivatives rely on z , we need to determine the sensitivity of each element of z to each element of $\bar{F}$ (i.e., $\partial z_k / \partial \bar{F}_{ij}$). Since $\bar{F}$ is a stochastic matrix, some care must be taken in defining this derivative. We resolve the technical difficulties in detail elsewhere; see Ref. 36 and Ref. 39 (Lemma 3.1). Here, to obtain the derivative $\partial z_k / \partial \bar{F}_{ij}$, evaluated at $\bar{F}$, we perturb around $\bar{F}$,

$$\left. \frac{d}{d\varepsilon} \right|_{\varepsilon=0} z_k(\bar{F} + \varepsilon E) = \sum_{i,j=1}^L \frac{\partial z_k}{\partial \bar{F}_{ij}}(\bar{F}) E_{ij}, \quad (80)$$

where E is an arbitrary matrix, ε is a scalar, and we assume that the sum $\bar{F} + \varepsilon E$ is also a stochastic matrix. The right hand side follows from the chain rule, effectively treating each element of the matrix as a separate argument to each element z . Then, we employ a relation from Golub and Meyer (Ref. 13, Theorem 3.1),

$$\left. \frac{d}{d\varepsilon} \right|_{\varepsilon=0} z_k(\bar{F} + \varepsilon E) = z(\bar{F})^T E (I - \bar{F})^\# e_k, \quad (81)$$

where $\#$ denotes the group inverse, a generalized matrix inverse similar to the Moore-Penrose inverse. It is defined as satisfying $AA^\#A = A$, $A^\#AA^\# = A^\#$, $AA^\# = A^\#A$. We refer to Golub and Meyer¹³ for further discussion of the group inverse and an algorithm for computing it. Finally, we equate (80) and (81) and solve for the derivative of interest,

$$\frac{\partial z_k}{\partial \bar{F}_{ij}}(\bar{F}) = z_i(\bar{F}) (I - \bar{F})_{jk}^\#. \quad (82)$$

Thus the sensitivity of each element of z to each element of $\bar{F}$ can be computed from linear algebra operations.

With (82), we can now compute the derivatives of B . We derive the formulas for the free-energy difference explicitly; the other cases are similar. In this case,

$$B(\bar{v}) = k_B T \ln \left(\sum_{k=1}^L z_k(\bar{F}) \bar{g}_{2,k}^* \right) - k_B T \ln \left(\sum_{k=1}^L z_k(\bar{F}) \bar{g}_{1,k}^* \right). \quad (83)$$

By the chain rule,

$$\frac{\partial}{\partial \bar{g}_{1,i}^*} \ln \left(\sum_{k=1}^L z_k(\bar{F}) \bar{g}_{1,k}^* \right) = \frac{z_i(\bar{F})}{\sum_{k=1}^L z_k(\bar{F}) \bar{g}_{1,k}^*} \quad (84)$$

and

$$\frac{\partial}{\partial \bar{F}_{ij}} \ln \left(\sum_{k=1}^L z_k(\bar{F}) \bar{g}_{1,k}^* \right) = \frac{\sum_{k=1}^L \frac{\partial z_k}{\partial \bar{F}_{ij}}(\bar{F}) \bar{g}_{1,k}^*}{\sum_{k=1}^L z_k(\bar{F}) \bar{g}_{1,k}^*}. \quad (85)$$

The stated result follows by substituting $g_1 = \mathbb{1}_{S1}$, $g_2 = \mathbb{1}_{S2}$, and the expression in (82) for $\partial z_k / \partial \bar{F}_{ij}$. $\square$

1. Computational procedure

We now provide a practical procedure that uses the derivatives above to estimate σ^2 from trajectories that sample the distributions π_i . For clarity, we assume that the system is equilibrated (i.e., X_0^i has distribution π_i , so that the process X_t^i is stationary) throughout this section.

We begin by rewriting (67) as

$$\sigma^2 = \sum_{i=1}^L \frac{\chi_i^2}{\kappa_i}, \quad (86)$$

where

$$\begin{aligned} \chi_i^2 &= \frac{\partial B}{\partial \bar{v}_i}^T \Sigma_i \frac{\partial B}{\partial \bar{v}_i} \\ &= \frac{\partial B}{\partial \bar{v}_i}^T \left(\sum_{t=-\infty}^{\infty} C_i(t) \right) \frac{\partial B}{\partial \bar{v}_i}. \end{aligned} \quad (87)$$

Defining the sequence

$$\zeta_t^i = \left. \frac{\partial B}{\partial \bar{v}_i} \right|_{\bar{v}} \cdot (v_i(X_t^i) - \langle v_i \rangle_i) \quad (88)$$

we find that

$$\chi_i^2 = \sum_{t=-\infty}^{\infty} \langle \zeta_t^i \zeta_0^i \rangle \quad (89)$$

which is the integrated autocovariance of ζ_t^i .

We thus propose the following algorithm, given simulation data:

1. Compute $\bar{v}$.
2. Compute $z(\bar{F})$ and $(I - \bar{F})^\#$ using the algorithm of Golub and Meyer.¹³
3. Evaluate $\partial B / \partial \bar{v}_i$ at $\bar{v}$ using the formulas in Lemma 7.5.
4. Compute

$$\bar{\zeta}_t^i = \left. \frac{\partial B}{\partial \bar{v}_i} \right|_{\bar{v}_i} \cdot (v_i(X_t^i) - \bar{v}_i). \quad (90)$$

5. Compute an estimate $\bar{\chi}_i^2$ of the integrated autocovariance of $\bar{\zeta}_t^i$ using an algorithm such as ACOR.⁴¹
6. Compute the estimate of σ^2 ,

$$\bar{\sigma}^2 = \sum_{i=1}^L \frac{\bar{\chi}_i^2}{\kappa_i}. \quad (91)$$

Since $\bar{F}$, $\bar{v}$, and z are all computed in the process of obtaining the EMUS averages, estimating $\bar{\sigma}^2$ only requires one additional pass over the simulation data. This additional cost is insignificant compared with that of computing the trajectories.

Both (67) and its approximation (91) decompose the asymptotic variance of EMUS into a sum of contributions from each window. By comparing the sizes of terms in the sum, we can determine the degrees to which different windows contribute to the error. In principle, this information can be used to guide modification of the parameters of the simulation to improve efficiency. For instance, one might adjust the

proportion of samples allocated to each window, κ_i , to minimize the asymptotic variance. From (86), the asymptotic variance σ^2 is minimized when $\kappa_i \propto \chi_i$ (see Equation (42) of Ref. 12). Consequently, we can define the *relative importance* of window i as

$$\mu_i = L \frac{\chi_i}{\sum_{k=1}^L \chi_k}, \quad (92)$$

where the normalization is chosen so that $\mu_i = 1$, regardless of L , if all windows have the same importance. The relative importance represents how many samples would be allocated to a window to optimally estimate a specific observable, compared to a uniform distribution over all windows.

2. Numerical results

To study the behavior of these estimates, we performed a two-dimensional umbrella sampling calculation with restraints on the ϕ and ψ dihedral angles of the alanine dipeptide. Parameters were the same as in the one-dimensional calculation above, with the addition of 20 bias functions in the ψ dihedral with the same force constant, creating a grid of 400 windows. Each window was equilibrated for 40 ps and sampled for a further 150 ps, with the collective variable values output every 10 fs.

In Figures 3 and 4(a), we plot the two-dimensional PMF from EMUS and the importances for the free energy difference between two windows located at the C_7 equatorial and C_7 axial configurations. Comparison shows that the importances are high for windows on low free energy pathways between the two windows of interest. Two such pathways exist. In the representation in Figure 3, one proceeds up and to the left of the C_7 equatorial basin and then (via the periodic boundaries) enters the C_7 axial basin through transition state 1 (TS1 in Figure 3). The other pathway proceeds down

FIG. 3. Potential of mean force obtained from US with biases on the ϕ and ψ dihedral angles. Major basins and barriers on pathways connecting them are indicated. Angles are measured in degrees. The scale bar indicates PMF values in kcal/mol, and the contour spacing is $2 k_B T$. The surface is constructed from simulation data accumulated in histograms with 100 bins in each collective variable. See text for simulation details.

then right through transition state 2 (TS2 in Figure 3). Of these two pathways, the first has a lower free energy barrier. We observe that the EMUS importances are larger for windows located on this pathway. In contrast, windows off these pathways in regions with high free energies have very low importances presumably because, though sampling error may be large in those regions, they do not contribute significantly to the desired averages. We refer the reader to Thiede *et al.*³⁹ for a mathematical discussion of the sensitivities of the averages.

We expect the importances to depend on the computed average. To illustrate that this is the case numerically, we show the log importances for the free energy difference between a window in the C_7 axial basin and one located on TS1 in Figure 4(b). Compared to Figure 4(a), the importances are higher in the C_7 axial basin and lower in the C_7 equatorial basin, which highlights that the importances depend on the average computed and do not simply mirror the free energy. In Figures 4(c) and 4(d), we plot the importances for estimating the window free energy (not the free energy difference) of the window on TS1 and the window in the C_7 axial basin, respectively. We note that the importances in the C_7 equatorial basin are higher in Figures 4(c) and 4(d) than in 4(b). This suggests that when the free energy difference between the two windows is considered, there is some cancellation of the errors arising in the C_7 equatorial basin.

3. Comparison with other algorithms for determining error contributions

Zhu and Hummer¹² proposed an algorithm for determining window free energies by calculating the mean restraining forces for each window and using thermodynamic integration to estimate free energy differences between adjacent windows. These are combined using least squares to calculate window free energies. Like EMUS, this algorithm allows one to construct error estimates that can be decomposed into contributions from individual windows. The authors give an expression for the error in the free energy of one window. This expression can be easily extended to the free energy difference between two windows, giving

$$\text{var}(\Delta G_{ji}) = \sum_k \text{var} \left(\sum_{\alpha}^D (c_{jk\alpha} - c_{ik\alpha}) \bar{f}_{k\alpha} \right), \quad (93)$$

where $\bar{f}_{k\alpha}$ is the average force exerted by the bias function for window k in the α th dimension. The constants $c_{ik\alpha}$ and $c_{jk\alpha}$ are defined in Appendix A of Zhu and Hummer.¹² The authors propose that these error estimates are applicable to WHAM and other umbrella sampling algorithms.

Both (67) and (93) are sums of contributions from individual windows. Using the formalism introduced in Section VII B 1, we define the process

$$\zeta_t^{k,ZH} = \sum_{\alpha}^D (c_{jk\alpha} - c_{ik\alpha}) (f_{k\alpha}(X_t^k) - \langle f_{k\alpha} \rangle), \quad (94)$$

and $(\chi_i^{ZH})^2$ as the integrated autocovariance of $\zeta_t^{k,ZH}$. This allows us to define importance for the Zhu and Hummer algorithm analogously to those for EMUS (see (92)).

FIG. 4. EMUS relative importances. Angles are measured in degrees. (a) Relative importances for the free energy difference between windows in the C_7 axial and C_7 equatorial basins. The window in the C_7 equatorial basin is centered at $(-81^\circ, 81^\circ)$, and the window in the C_7 axial basin at $(63^\circ, -63^\circ)$. (b) Window importances for the free energy difference between windows in the C_7 axial basin and at TS1. Windows are centered at $(63^\circ, -63^\circ)$ and $(135^\circ, -117^\circ)$, respectively. (c) Importances for the free energy of the window at TS1. (d) Importances for the window in the C_7 axial basin.

We applied the ZH error analysis to the two-dimensional umbrella sampling data used in Section VII B 2 and calculated the importances for the same free energy difference as in Figure 4(a) (Figure 5(a)). Rather than falling along the low free energy pathways, as for EMUS, the ZH importances mirror the autocorrelation times (Figure 5(b)).

This indicates that windows have large ZH importances if they have large fluctuations in free energy. We thus see that different algorithms emphasize different windows in US. We can understand the behaviors of these two algorithms by considering (88) and (94). The factor $\partial B / \partial \bar{v}_i$ in (88) depends explicitly on the normalization constant for each window (see

FIG. 5. Comparison to the work of Zhu and Hummer. (a) ZH estimates for the relative importances for the free energy difference between windows in the C_7 axial and C_7 equatorial basins. Compare with Figure 4(a). (b) Autocorrelation times of the trajectory $\zeta_t^{k,ZH}$ in each window. The largest value observed is 3 ps, but the scale is limited to 1 ps for visual clarity.

Lemma 7.5). By contrast, the factor $(c_{jka} - c_{ika})$ in (94) depends only on the relative positions of the windows and not on their free energies.

VIII. CONCLUSIONS

The success of an umbrella sampling simulation depends on the choice of windows (i.e., how the system is biased) and the estimator used to determine the normalization constants of the windows from trajectory data. Here, we show that the normalization constants can be obtained from an eigenvector of a stochastic matrix. This eigenvector method for umbrella sampling (EMUS) can be viewed as the first step in an implementation of the MBAR estimator. In our experience, this first step is nearly converged, and machine precision is reached in only a few iterations. Moreover, each iteration yields a consistent estimate. Most importantly, error analysis is considerably easier for EMUS than MBAR because the elements of the stochastic matrix do not depend on the normalization constants.

Within this framework, we revisited a common scaling argument for justifying umbrella sampling and showed that once the number of windows becomes sufficiently large, the scheme does not benefit from the addition of more windows (i.e., the variance is not further reduced for a fixed computational effort). We show that an alternative scaling regime in which temperature decreases (or, equivalently, free-energy barrier heights increase) as the number of windows increases best demonstrates the potential benefits of the umbrella sampling strategy; in that regime the efficiency improvement over direct simulation is exponential in the (inverse) temperature.

Our main theoretical result is a central limit theorem for the statistical averages obtained from EMUS. This result relies on the delta method, which we use to characterize the propagation of the asymptotic error through the solution of a stochastic matrix eigenproblem. The central limit theorem provides an expression for the asymptotic variance of the averages of interest. It is a sum of contributions from individual windows, and we use it to develop a prescription for estimating the relative importances of windows for averages from the trajectory data. For free energy differences of states of the alanine dipeptide, we find numerically that the importances are largest for low-free energy pathways that connect the specific states of interest. These results suggest that the importances could serve as the basis for adaptive schemes that focus computational effort on the windows of most importance. Even more interesting would be to adjust the bias functions as the simulation progresses. How best to do this remains an open area of investigation.

ACKNOWLEDGMENTS

This research was supported by National Institutes of Health (NIH) Grant No. 5 R01 GM109455-02. We wish to thank Jonathan Mattingly, Jeremy Tempkin, and Charlie Matthews for helpful discussions.

APPENDIX: CONSISTENCY OF ITERATIVE EMUS

Here, we prove that for fixed finite m , z^m is a consistent estimator of the vector of normalization constants z . With the initial guess $z^0 = n$, the result, z^1 , of the first iteration is the EMUS estimator. We now show that z^2 is also consistent in the sense that if the trajectory averages defining $\bar{F}(w)$ converge then z^2 converges to z . Because the various sequences in question are sequences of random variables, one must specify what is meant by convergence. The argument below applies when convergence refers either to convergence in probability or convergence with probability one (almost sure convergence) as long as the notion of convergence is consistent throughout. Consistency of z^m follows by induction on m using a similar argument.

For any positive vector w , we define

$$u_k = \prod_{j \neq k} \frac{w_j}{n_j}, \quad (\text{A1})$$

and we write

$$\bar{F}_{ij}(w) = \frac{1}{N_i} \sum_{t=0}^{N_i-1} h_{ij}(u, x), \quad (\text{A2})$$

where

$$h_{ij}(u, x) = \frac{u_i \psi_j(x)}{\sum_k u_k \psi_k(x)}. \quad (\text{A3})$$

We then observe that

$$\partial_{u_k} h_{ij}(u, x) = \begin{cases} \frac{1}{u_i} h_{ij}(u, x) [1 - h_{ii}(u, x)], & \text{if } k = i \\ -\frac{1}{u_i} h_{ij}(u, x) h_{ik}(u, x), & \text{if } k \neq i \end{cases}. \quad (\text{A4})$$

Because $h_{ij}(u, x) \leq u_i/u_j$,

$$|\partial_{u_k} h_{ij}(u, x)| \leq \frac{1}{u_j} \max \left\{ 1, \frac{u_i}{u_k} \right\}. \quad (\text{A5})$$

Therefore,

$$|h_{ij}(\tilde{u}, x) - h_{ij}(u, x)| \leq \gamma(u, \tilde{u}), \quad (\text{A6})$$

for a continuous function γ defined for positive vectors u and $\tilde{u}$ and such that $\gamma(u, u) = 0$ for any u . The function $\gamma(u, \tilde{u})$ must explode when the entries of u or $\tilde{u}$ approach 0. Now define

$$u_k = \prod_{j \neq k} \frac{z_j^1}{n_j} \text{ and } \tilde{u}_k = \prod_{j \neq k} \frac{z_j}{n_j}, \quad (\text{A7})$$

where z is the exact vector of normalization constants. By (A6), we have

$$|\bar{F}_{ij}(z^1) - \bar{F}_{ij}(z)| \leq \gamma(u, \tilde{u}). \quad (\text{A8})$$

As the number of samples N increases, $\bar{F}(z)$ converges to $F(z)$. Moreover, since z^1 is the EMUS estimate of z , z^1 converges to z . Therefore, u converges to $\tilde{u}$, and (A8) implies that $\bar{F}_{ij}(z^1)$ converges to $F_{ij}(z)$. Finally, since the function mapping an irreducible, stochastic matrix to its invariant vector is continuous, it follows that z^2 converges to the invariant vector of $F(z)$, which is z . This verifies the consistency of z^2 .

- ¹G. M. Torrie and J. P. Valleau, *J. Comput. Phys.* **23**, 187 (1977).
- ²C. Pangali, M. Rao, and B. Berne, *J. Chem. Phys.* **71**, 2975 (1979).
- ³Y. Vardi, *Ann. Stat.* **13**, 178 (1985).
- ⁴R. D. Gill, Y. Vardi, and J. A. Wellner, *Ann. Stat.* **16**, 1069 (1988).
- ⁵S. Kumar, D. Bouzida, R. H. Swendsen, P. A. Kollman, and J. M. Rosenberg, *J. Comput. Chem.* **13**, 1011 (1992).
- ⁶M. R. Shirts and J. D. Chodera, *J. Chem. Phys.* **129**, 124105 (2008).
- ⁷E. Rosta and G. Hummer, *J. Chem. Theory Comput.* **11**, 276 (2014).
- ⁸A. S. Mey, H. Wu, and F. Noé, *Phys. Rev. X* **4**, 041018 (2014).
- ⁹Z. Tan, E. Gallicchio, M. Lapelosa, and R. M. Levy, *J. Chem. Phys.* **136**, 144102 (2012).
- ¹⁰T. Lelièvre, G. Stoltz, and M. Rousset, *Free Energy Computations: A Mathematical Perspective* (World Scientific, 2010).
- ¹¹D. D. Minh and J. D. Chodera, *J. Chem. Phys.* **131**, 134110 (2009).
- ¹²F. Zhu and G. Hummer, *J. Comput. Chem.* **33**, 453 (2012).
- ¹³G. H. Golub and C. D. Meyer, Jr., *SIAM J. Algebraic Discrete Methods* **7**, 273 (1986).
- ¹⁴E. H. Thiede, EMUS, 2016, <https://github.com/ehthiede/EMUS>.
- ¹⁵H. Schneider, *Linear Algebra Appl.* **18**, 139 (1977).
- ¹⁶Z. Tan, *J. Am. Stat. Assoc.* **99**, 1027 (2004).
- ¹⁷D. Frenkel and B. Smit, *Understanding Molecular Simulation: From Algorithms to Applications* (Academic Press, 2001).
- ¹⁸C. J. Geyer, *Stat. Sci.* **7**, 473 (1992).
- ¹⁹H. Doss and A. Tan, *J. R. Stat. Soc.: Ser. B* **76**, 683 (2014).
- ²⁰X.-L. Meng and W. H. Wong, *Statistica Sinica* **6**, 831 (1996).
- ²¹H. Paliwal and M. R. Shirts, *J. Chem. Theory Comput.* **9**, 4700 (2013).
- ²²M. R. Shirts, D. L. Mobley, and J. D. Chodera, *Annu. Rep. Comput. Chem.* **3**, 41 (2007).
- ²³M. J. Abraham, T. Murtola, R. Schulz, S. Páll, J. C. Smith, B. Hess, and E. Lindahl, *SoftwareX* **1**, 19 (2015).
- ²⁴G. A. Tribello, M. Bonomi, D. Branduardi, C. Camilloni, and G. Bussi, *Comput. Phys. Commun.* **185**, 604 (2014).
- ²⁵A. D. MacKerell, N. Banavali, and N. Foloppe, *Biopolymers* **56**, 257 (2000).
- ²⁶B. Hess, H. Bekker, H. J. C. Berendsen, and J. G. E. M. Fraaije, *J. Comput. Chem.* **18**, 1463 (1997).
- ²⁷A. Grossfield, “WHAM: The weighted histogram analysis method (version 2.0.9),” 2013, <http://membrane.urmc.rochester.edu/content/wham>.
- ²⁸E. Süli and D. F. Mayers, *An Introduction to Numerical Analysis* (Cambridge University press, 2003).
- ²⁹D. Chandler, *Introduction to Modern Statistical Mechanics* (Oxford University Press, 1987).
- ³⁰C. Chipot and A. Pohorille, *Free Energy Calculations* (Springer, 2007), p. 86.
- ³¹J. van Duijneveldt and D. Frenkel, *J. Chem. Phys.* **96**, 4655 (1992).
- ³²W. Wojtas-Niziurski, Y. Meng, B. Roux, and S. Bernèche, *J. Chem. Theory Comput.* **9**, 1885 (2013).
- ³³P. Virnau and M. Müller, *J. Chem. Phys.* **120**, 10925 (2004).
- ³⁴T. H. Nguyen and D. D. Minh, *J. Chem. Theory Comput.* **12**, 2154 (2016).
- ³⁵F. P. Kelly, *Reversibility and Stochastic Networks* (Cambridge University Press, 2011).
- ³⁶A. R. Dinner, E. H. Thiede, B. Van Koten, and J. Weare, “Stratification of Markov chain Monte Carlo sampling” (unpublished).
- ³⁷T. Lelièvre, G. Stoltz, and M. Rousset, *Free Energy Computations: A Mathematical Perspective* (Imperial College Press, Hackensack, NJ, 2010), p. 458.
- ³⁸M. Bilodeau and D. Brenner, *Theory of Multivariate Statistics* (Springer Science & Business Media, 2008).
- ³⁹E. Thiede, B. Van Koten, and J. Weare, *SIAM J. Matrix Anal. Appl.* **36**, 917 (2015).
- ⁴⁰A. W. van der Vaart, *Asymptotic Statistics*, Cambridge Series on Statistical and Probabilistic Mathematics (Cambridge University Press, Cambridge, New York, 1998), p. 443.
- ⁴¹D. Foreman-Mackey and J. Goodman, ACOR 1.1.1, 2014, <https://pypi.python.org/pypi/acor/1.1.1>.