Obtaining Rates from Biased Molecular Dynamics Simulations Using Poor Collective Variables

by

Nicodemo Mazzaferro

A dissertation submitted in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

Department of Chemistry

New York University

August, 2026

Dr. Glen M. Hocky

Dedication

To my family; Luigi Mazzaferro, Emma Mazzaferro, and Tommaso Mazzaferro

Acknowledgments

First and foremost, I am grateful to my thesis advisor Prof. Glen Hocky. He has provided the best guidance and encouragement for all of my research, and has taught me so much over the course of my graduate career. Working with him has always been a pleasure, and through him I have met other great researchers in the field. Second, I am also grateful to my collaborator and co-mentor Dr. Pilar Cossio from the Flatiron Institute, without whom the very premise of my thesis research would not have been possible. She has also provided excellent advice and encouragement throughout my time working with her.

I would like to extend my appreciation to the members of my dissertation committee, which includes Prof. Yingkai Zhang, Prof. Tania Lupoli, and Prof. Stefano Martiniani in addition to Prof. Glen Hocky and Dr. Pilar Cossio, for their feedback on this dissertation and their support throughout my graduate studies.

I must also thank the members of the Hocky group for providing such a friendly research environment. Specifically, I appreciate Willmor Peña Ccoa and Subarna Sasmal for their help in my research projects, for their contributions have greatly strengthened those projects. I also appreciate Yuvraj Singh and Gaurav Mitra for their conversations when I first joined the group, as well as the whole of the research group, past and present, for their friendliness and feedback.

For the research presented in Chapter 2, I also must thank Karen Palacio-Rodriguez, Fabio Pietrucci, Lukas S. Stelzl, Attila Szabo, Gerhard Hummer, and David Limmer for taking the time to discuss my ideas. Also, I must thank Barak Hirshberg and Ofir Blumer for providing their

chignolin miniprotein data for analysis, and Robert B. Best for providing the parameters for the Gō-like protein model that was used both in Chapter 2 and in Chapter 4.

I would also like to thank the Flatiron Institute for welcoming me throughout my collaboration with Dr. Pilar Cossio. I am specifically glad to have met Luke Evans from her research group there, who has had many great conversations with me, both regarding research and not. Also, I would like to thank Prof. Pratyush Tiwary and his graduate student Suemin Lee for collaborating with me on a joint research project. Suemin is such a driven researcher, and I wish her the best for the future.

I appreciate the National Institutes of Health for their grant funding as well as the Simons Center for Computational Physical Chemistry for their financial support in my research through fellowships. I also thank NYU for additional funding through the Margaret Strauss Kramer Fellowship, as well as the Center for Computational Mathematics at the Flatiron Institute for funding my research in the summer during my time working for them. I also would like to thank the NYU HPC team for access to the Greene and Torch supercomputer clusters, which were necessary to carry out my research.

Finally, I would like to thank my parents for their endless support and encouragement, not only for my graduate studies but for all I pursue, as well as my twin, Tommaso Mazzaferro, who has always supported me and is a great brother.

Abstract

Obtaining rate constants from molecular dynamics simulations is difficult because of the inaccessibility of the necessary time scales to observe rare events. To compute properties on these long time scales, we employ enhanced sampling techniques which lower the energy barriers for these events. A major area of interest is computing the unbiased kinetics for slow biomolecular processes. While this is possible using existing methods, many of the efficient approaches rely on the assumption that there is no energy added to the barrier, which is unlikely when a poor choice of collective variable is used to describe the transition process.

The work presented in this thesis produces and develops a new method to obtain rates from molecular dynamics with enhanced sampling called the exponential average time-dependent rate (EATR) method. The EATR method obtains rate constants by fitting observed transition times using a model for how the transition rate constant observed in the simulations is affected by the amount of bias potential present. This model contains an extra fitting parameter γ which controls the sensitivity of the rate constant to the bias potential, but also reproduces the rate constants predicted by the well-established infrequent metadynamics (iMetaD) estimator when $\gamma = 1$.

We demonstrate that the EATR method does reproduce the iMetaD rate constants when $\gamma = 1$ and that EATR outperforms iMetaD for cases where γ is significantly less than 1. We show that γ corresponds to the known quality of the collective variables used in biasing, and that the method works well for both simple one-dimensional systems and fully atomistic protein folding models. In Chapter 3, we then show that EATR can be used to track the improvement of a machine-learned

reaction coordinate in the case of protein-drug unbinding using γ as a quality metric. This also allowed us to develop new heuristics for checking whether these data-driven reaction coordinates are improving between iterations. Finally, in Chapter 4, we report a generalization of EATR which gives accurate rates in cases where the bias and observed rate are static or quasistatic over the course of a simulation. This corrects a problem in standard EATR where quasistatic biasing does not provide enough information to uniquely determine the rate constant and γ. In doing so, we introduce the idea to combine data from multiple sets of simulations to gain information which would not have been available otherwise. This concept allows us to consider the use of EATR for a much wider set of enhanced sampling approaches for computing accurate kinetics of slow processes.

Contents

List of Figures

List of Tables

List of Appendices

1 Introduction

1.1 Motivation

All matter is composed of atoms, and the properties, reactions, and phase transitions we observe in chemistry are ultimately determined by the motions and interactions of these atoms. However, we are incapable of directly observing these over the course of a chemical process. For this reason, we turn to molecular dynamics (MD) simulations[1], the so-called "computational microscope" which can shed light on the mechanisms that drive whichever process we may be interested in.

MD simulations are based in physics and, therefore, are predictive. MD is often used to make predictions for various properties of chemical systems, such as binding free energies[2, 3], diffusion coefficients[4, 5], thermal conductivity[6, 7], and rates[8, 9]. Predicting rates is particularly interesting in a pharmaceutical context, as it has been shown that the unbinding rate of a drug molecule from a target is a good predictor for the efficacy of the drug[10]. However, collecting enough simulation data to make these predictions is usually infeasible within a reasonable time period.

The goal of the research presented in this thesis is to provide a new method to obtain predictions for state-transition rates from MD simulations which were accelerated by enhanced sampling methods. This would allow for the use of rates not only in pharmaceutical lead optimization, but also in determining the accuracy of MD force fields. This work also refines a metric for the

accuracy of reaction coordinates, and therefore yields a better understanding of transition processes.

1.2 Molecular Dynamics

In MD, we model the atoms in a material as classical point-particles following Newton's laws of motion. Denoting the vector of atomic positions as the *configuration* $\vec{X}$, the forces acting between atoms are obtained by taking the negative gradient of a potential energy function called the *force field* $U(\vec{X})$.

Because all forces are derived from a potential energy function, solving Newton's or Hamilton's equations of motion conserves the total energy. However, in experiments, the systems we wish to study exchange heat with their surroundings. In MD simulations, the interactions of the system with its surroundings are modeled using a *thermostat*. Whether the thermostat in use is stochastic as in Langevin dynamics[11], or introduces external degrees of freedom as in the Nosé-Hoover thermostat[12, 13], ultimately the evolution of the system can no longer be predicted from its current internal state and must be modeled probabilistically. In constant temperature conditions, the atomic configuration becomes distributed according to the Boltzmann distribution:

$$p(\vec{X}) = \frac{1}{Z} e^{-\beta U(\vec{X})}, \tag{1.1}$$

where Z is the canonical partition function, which for our purposes may be regarded merely as a normalization constant, and $\beta = 1/k_B T$ where k_B is Boltzmann's constant and T is the temperature of the thermostat.

Experimental observations are associated with *ensemble averages* $\langle O \rangle$, which are averages taken over the distribution of configurations:

$$\langle O \rangle = \frac{1}{Z} \int_{\Omega} O(\vec{X}) e^{-\beta U(\vec{X})} \, d\vec{X}, \tag{1.2}$$

where $O(\vec{X})$ is an observable quantity and the integral is taken over Ω, which is the space of all atomic configurations. We can also take averages over subsets of Ω if needed.

Assuming that all atomic configurations in Ω can be adequately sampled with proper statistical frequency within an MD simulation, we can approximate $\langle O \rangle$ as an average of $O(\vec{X})$ over the steps of the simulation:

$$\langle O \rangle \approx \frac{1}{S} \sum_{s=1}^{S} O(\vec{X}(t_s)), \tag{1.3}$$

where S is the number of steps in the simulation and t_s is the simulation time at step s.

1.3 Metastable States and Transitions Between Them

1.3.1 Collective Variables and the Potential of Mean Force

Often, we are only interested in one or a few degrees of freedom in our systems. For example, for a folding protein we would only be interested in some folding measure, such as the compactness of the protein or the number of residues in contact. For a ligand bound to a substrate, we would be interested in the distance between the ligand and the binding site. For an atom diffusing across a metal surface, we would keep track of the position of the atom along the surface, and for a chemical reaction we would find and consider a reaction coordinate.

Generally, we are not interested in the atomic positions, but in a quantity which describes the collective behavior of the system. These *collective variables* (CVs) can be expressed as a function of the atomic positions, $\xi(\vec{X})$.

The *potential of mean force* (PMF) or *free energy* along ξ is the low-dimensional potential that is consistent with the total Boltzmann distribution for the system after integrating out other

degrees of freedom:

$$F(\xi) \equiv -\frac{1}{\beta} \ln p(\xi), \tag{1.4}$$

$$p(\xi) = \int \delta(\xi(X) - \xi) e^{-\beta U(\vec{X})} d\vec{X}. \tag{1.5}$$

This allows us not only to greatly simplify our analysis of the system, but also to determine the locations of the free energy minima. Regions of low free energy are, by definition, the regions where the system is most likely to be. These regions correspond to metastable states, which are collections of configurations in which the system stays for an extended period of time. Assuming that ξ is able to distinguish these states, then metastable states will correspond to convex regions of the PMF.

1.3.2 Barriers to Transitions, and Enhanced Sampling

Two separate macroscopic states correspond to two disconnected regions of low free energy, which implies a region of high free energy between them. In order to traverse from one state to another, the system must surmount these high energy barriers. However, the system is very unlikely to reach the top of these barriers at any given moment. This is so unlikely, in fact, that it is often statistically impossible to observe a transition between states in an MD simulation.

To observe transitions in MD simulations, we often employ a class of enhanced sampling methods that add an additional term to the force field:

$$U'(\vec{X}) = U_0(\vec{X}) + V(\vec{X}), \tag{1.6}$$

where $V(\vec{X})$ is the new term, called the *bias potential*, or simply the *bias*. For any variable A, A_0 will denote the value of A that would be observed in MD simulations with no bias, and A' will denote the value in MD simulations with bias. The bias is often a function of one or multiple CVs.

Given a time-independent bias, the ensemble average of an observable in an unbiased simulation can be obtained from the ensemble average in the biased simulation by performing a weighted average[14] over the frames of the simulation:

$$\langle O \rangle_0 = \frac{\langle O\, e^{\beta V} \rangle'}{\langle e^{\beta V} \rangle'} \approx \frac{\sum_s O(\vec{X}(t_s)) e^{\beta V(\vec{X}(t_s))}}{\sum_s e^{\beta V(\vec{X}(t_s))}}. \tag{1.7}$$

However, there are multiple enhanced sampling methods which add bias potentials which are time-dependent, such as metadynamics (MetaD)[15] and on-the-fly probability enhanced sampling (OPES)[16].

1.3.3 Metadynamics

MetaD[15] introduces a time-dependent bias potential which is a sum of Gaussians:

$$V(\xi, t) = \sum_{g=0}^{G(t)} h_g e^{-\frac{(\xi - \xi_g)^2}{2\sigma^2}}, \tag{1.8}$$

where $G(t)$ is the number of Gaussians that were added up to time t and h_g, ξ_g, and σ_g are the height, mean, and standard deviation of the g-th Gaussian, respectively. These Gaussians are added at a regular interval in the simulation and are centered at the immediate value of the CV when they are added.

In the original formulation of MetaD, the height and standard deviation of the Gaussians were all the same. One popular variant of MetaD is well-tempered MetaD (WT-MetaD)[17], where the height of the Gaussian g deposited at step s is decreased based on the amount of bias already present at $\xi(\vec{X}(t_s))$ according to

$$h_g = h e^{-V(\xi(\vec{X}(t_g)), t_g)/k_B \Delta T}, \tag{1.9}$$

where h is the base height for all Gaussians, t_g is the time that Gaussian g was added, and ΔT is a tempering factor that controls the degree to which the Gaussian height decays. This ensures that the bias converges to a finite value in the long-time limit [18]. This allows the reweighting in Eq. 1.7 to work for WT-MetaD after first waiting for the bias to converge.

1.3.4 On-the-fly Probability Enhanced Sampling

OPES[16] introduces a time-dependent bias that is chosen to sample a target probability distribution:

$$V(\xi, t) = \frac{1}{\beta} \ln \frac{\hat{p}(\xi, t)}{p_t(\xi)}, \tag{1.10}$$

where $\hat{p}(\xi, t)$ is the estimate of the unbiased probability distribution obtained from kernel density estimation and $p_t(\xi)$ is the target distribution. Upon choosing the biased probability distribution of WT-MetaD as the target, the bias becomes

$$V(\xi, t) = \left(1 - \frac{1}{B}\right) \frac{1}{\beta} \ln \left(\frac{\tilde{p}(\xi, t)}{Z(t)} + \epsilon \right), \tag{1.11}$$

where $B = (T+\Delta T)/T$ is called the "bias factor"[1], $\tilde{p}(\xi, t)$ is the unnormalized estimated probability distribution, $Z(t)$ is the normalization factor for $\tilde{p}(\xi, t)$, and ϵ is a regularization term that is introduced to ensure that the bias remains finite whenever $\tilde{p}(\xi, t) = 0$.

The regularization term is chosen to be $\epsilon = e^{-\beta \Delta E/(1-1/B)}$, where ΔE is a parameter that affects the maximum value of the bias potential. Wherever $\tilde{p}(\xi, t) = 0$ the bias $V(\xi, t) = -\Delta E$, and at the maximum value of $\tilde{p}(\xi, t)$, the bias will consistently be slightly over 0. This not only gives control over the amount of bias to be added, but in practice also causes the amount of bias in the simulation to reach its final value effectively instantly, assuming a sufficiently small update pace. This rapid convergence means that the reweighting in Eq. 1.7 starts applying quickly.

[1] Note that the bias factor is usually labeled γ, but here we require a different label to avoid confusion with the γ parameter in EATR.

In Ref. [16], the implementation of OPES includes the ability to select a region where bias cannot be added, which provides more control over the bias and is useful in obtaining rate constants.

1.4 Computational Methods to Obtain Rate Constants from Molecular Dynamics Simulations

Because of free energy barriers separating metastable states, transitions are often "rare events," where the system first spends a long time in its current state before a rare fluctuation causes a transition to another basin. Assuming that these events are independent from one another, we can model them as a Poisson process with transition times distributed as

$$P(n;t) = \frac{(kt)^n e^{-kt}}{n!}, \tag{1.12}$$

where $P(n;t)$ is the probability that n transitions occur in an interval of length t and k is the rate at which transitions occur. For first-order reactions with a single molecule, k is the rate constant. In MD simulations, a natural way to determine the rate constant is to consider the probability from Eq. 1.12 that zero transitions occur within time t:

$$S(t) \equiv P(0;t) = e^{-kt}, \tag{1.13}$$

which we call the *survival function*. In principle, we could run a large number of simulations and keep track of the fraction of them which have not experienced at least one transition, then estimate k by either fitting to Eq. 1.13 or by computing the reciprocal of the mean first transition time.

In practice, for real problems of interest we can simulate a single copy of our system for at most

hundreds of microseconds of MD simulation per month, using even the best available resources. This prevents us from properly sampling the transition time distribution in cases where the mean transition time is on the order of milliseconds, or even seconds, as is often the case for biological processes of interest. This requires alternate methods to estimate the rate for a given transition process.

1.4.1 Transition Path Sampling

Transition path sampling [19, 20] obtains rates by sampling an ensemble of trajectories with importance sampling. Given a trajectory $\vec{X}_i(t)$ from state A to state B, we can produce a new trajectory by selecting a step s^* in the trajectory where $\vec{X}_i(t_{s^*})$ is near the barrier, changing the velocities at that step, and simulating the system forward and backward in time from that step. This is called a *shooting move*. The resultant trajectory $\vec{X}_{i+1}(t)$ is accepted in the importance sampling with the probability

$$P_{accept} = \min\left(1, \frac{w(\vec{X}_i, s^*, \delta)h_A(\vec{X}_{i+1}(0))h_B(\vec{X}_{i+1}(T))}{w(\vec{X}_{i+1}, s^*, -\delta)h_A(\vec{X}_i(0))h_B(\vec{X}_i(T))}\right), \tag{1.14}$$

where $w(\vec{X}, s, \delta)$ is the probability of selecting the step s in $\vec{X}(t)$ and changing the velocity at that step by δ, $h_A(\vec{X})$ is the indicator function that is 1 only when $\vec{X}$ is in A, and T is the total length of the trajectory. If $\vec{X}_{i+1}(0) \notin A$ or $\vec{X}_{i+1}(T) \notin B$, the trajectory is automatically rejected. With enough shooting moves, we construct an ensemble of trajectories distributed according to

$$p[\vec{X}(t)] = p_{eq}[\vec{X}(t)]h_A(\vec{X}(0))h_B(\vec{X}(T)), \tag{1.15}$$

where $p_{eq}[\vec{X}(t)]$ is the probability density for obtaining the trajectory $\vec{X}(t)$ in a simulation.

The rate constant k from A to B can be calculated from this ensemble by the following equa-

tion:

$$k = \frac{1}{T}\frac{\langle h_A(0)h_B(T)\rangle}{\langle h_A\rangle}, \tag{1.16}$$

where the averages are taken over the path ensemble. This is valid as long as T is long enough to always capture the full transition path, yet short enough that there will only be one transition.

This method does not require an additional bias potential, and therefore the transition paths which are sampled are guaranteed to be physically reasonable. However, this method requires an initial transition path near the transition path ensemble, and it can be difficult and/or computationally expensive to obtain a representative set of transition paths[20]. To sample a large number of trajectories efficiently, we need to select steps for shooting moves that are near the barrier, because otherwise almost every trajectory will start and end in the same state. This requires a measure for how likely a trajectory is to reach one state versus another, which is captured in the *committor probability*[20–22]. However, while recent methods have made significant progress towards obtaining it[23, 24], the committor probability is generally very expensive to obtain for an arbitrary system.

1.4.2 Hyperdynamics

Hyperdynamics [25, 26] or conformational flooding [27] is a method derived from transition state theory to obtain unbiased rate constants from molecular dynamics simulations with a static bias. According to transition state theory, the rate constant is given by the following ensemble average:

$$k_0 = \langle |v|\delta_A(\vec{X})\rangle_{0,A}, \tag{1.17}$$

where the average is taken over the initial state A, $|v|$ is the component of the system velocity perpendicular to the boundary of A, and $\delta_A(\vec{X})$ is the Dirac delta function at the boundary of A.

In a biased simulation, we can perform reweighting with Eq. 1.7 to get the following:

$$\langle |v|\delta_A(\vec{X})\rangle_{0,A} = \frac{\langle |v|\delta_A(\vec{X})\rangle'_A}{\langle e^{\beta V}\rangle'_A} = \frac{k'}{\langle e^{\beta V}\rangle'_A}, \tag{1.18}$$

where we can eliminate the $e^{\beta V}$ from the numerator of Eq. 1.7 only if there is no bias at the boundary of A. In Ref. [25] it was shown that this is equivalent to scaling the duration Δt of each step s of a biased MD simulation according to

$$\Delta t_0(s) = \Delta t' e^{\beta V(\vec{X}(t_s))}. \tag{1.19}$$

By rescaling every step in a set of MD simulations, we can sample the unbiased transition time distribution of a transition process, and obtain the unbiased rate constant.

1.4.3 Infrequent Metadynamics and OPES-Flooding

Infrequent MetaD (iMetaD)[28, 29] is a rate estimation method where we apply the ideas of hyperdynamics to a set of simulations biased with MetaD. Typically, instead of rescaling each timestep and adding them together to get a total time, an equivalent calculation is done where the scaling factor is averaged across time, then is multiplied by the total biased time.

$$t_0 = t' \left(\frac{1}{t'} \sum_s e^{\beta V(\vec{X}(t_s))} \Delta t' \right) \equiv t' \times \alpha, \tag{1.20}$$

where α is called the *acceleration factor*.

There are two considerations that are important when applying hyperdynamics to MetaD simulations. Because the bias is not truly static, the reweighted ensemble average in Eq. 1.18 is no longer exact. However, as long as the Gaussians in MetaD are added infrequently, the time rescaling argument will mostly hold while the bias is static between the Gaussian additions. The second thing to consider is that we generally do not control where the bias is added in MetaD. As

such, it is possible that at some point the bias on the boundary of A will not be zero. However, again as long as the Gaussians are added infrequently, it will be very unlikely that a Gaussian will be added within the short time when the system is at the boundary.

The ideas from iMetaD can also be applied to OPES simulations, in which case the method is called OPES-flooding, or OPESf[30]. The above concerns are better addressed in OPESf, as the bias may be considered “pseudo-static" because it converges quickly, and we can select a region of space where bias is not allowed to be added.

To use hyperdynamics, we need to apply a static bias, which requires some prior knowledge of the system to construct. For iMetaD, we use MetaD to apply the bias, which only needs a CV to construct the bias on the fly. Because of this, iMetaD is easy to use and provides reasonable results in many cases. However, it turns out that iMetaD requires an ideal CV to give the correct unbiased rate constant. As an illustration, consider a two-dimensional double well system with coordinates x and y as shown in Fig. 1.1, and consider using x as the biasing CV $\xi(x, y) = x$. MetaD would apply a bias $V(x)$ that is a function of only x, so the same amount of bias is added to all points with the same value of x. If a Gaussian is added when the system is at point A, bias will also be added to the transition boundary at point B, which violates the requirement that the bias at the boundary is zero.

1.4.4 Kramers' Time-dependent Rate

The Kramers' time-dependent rate method (KTR)[31] is a rate estimation method based off of Kramers' rate theory. According to Kramers' rate theory, the rate of escape over a harmonic barrier I for a one-dimensional system in a harmonic well A is given by

$$k_0 = \frac{\omega_A \omega_I}{2\pi\zeta} e^{-\beta\Delta U} = k_{\text{pre}} e^{-\beta\Delta U} \tag{1.21}$$

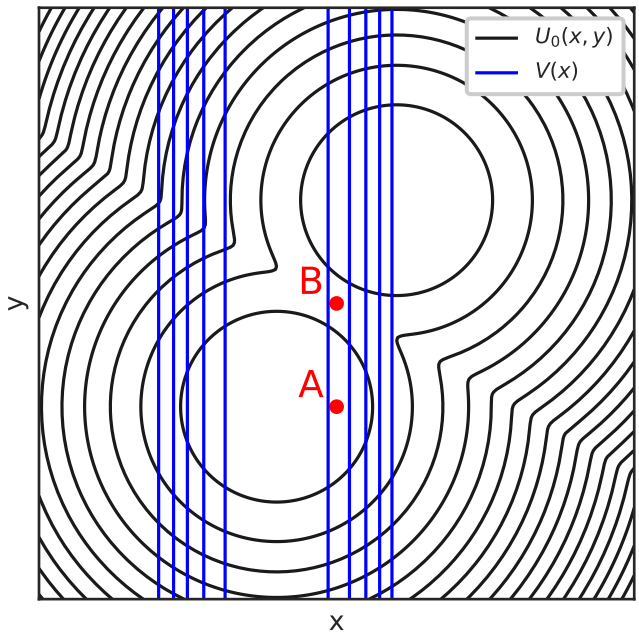

Figure 1.1: The contour plot of the potential energy for a two-dimensional two-state system, $U_0(x, y)$. The contour plot for the bias $V(x)$ is overlaid on top of it. Point A is a reasonable location for the system when it is in the lower state and point B is on the boundary of the lower state. Points A and B share a value of x, so the bias at A is equal to the bias at B.

where ω_i is the resonant frequency of the harmonic i, ζ is the friction coefficient, and ΔU is the difference in energy between the well minimum and the barrier maximum. For the purposes of KTR, the prefactor is bundled into a single constant k_{pre}. Upon adding a bias $V(\xi)$ that is zero on the barrier, the rate becomes

$$k' = k_{\text{pre}} e^{-\beta(\Delta U - V_A)} = k_{\text{pre}} e^{-\beta \Delta U} e^{\beta V_A} = k_0 e^{\beta V_A} \tag{1.22}$$

where V_A is the value of the bias at the well minimum. This is only valid if we assume that the bias does not significantly alter k_{pre}. A few changes are made to apply this to a simulation biased with MetaD without prior knowledge of the PMF, leading to:

$$k'(t) = k_0 e^{\beta \gamma \overline{V_{\max}}(t)}, \tag{1.23}$$

where $\overline{V_{max}}(t)$ is the maximum value of the bias potential at time t averaged over all simulations, and $\gamma \in [0, 1]$ is a scaling factor for the bias to account for the possibility that the bias is nonzero on the barrier.

Because the biased rate changes as bias is added over time, the Poisson distribution is no longer directly applicable. We instead need to model the transition as an *inhomogeneous* Poisson process, with a survival function of the form

$$S(t) = e^{-\int_0^t k(t')dt'}, \tag{1.24}$$

which, upon plugging in $k'(t)$ from KTR, yields

$$S(t) = e^{-k_0 \int_0^t e^{\beta \gamma \overline{V_{\max}}(t')} dt'}. \tag{1.25}$$

In Ref. [31], values for k_0 and γ were obtained using maximum likelihood estimation (MLE).

In Ref. [32], I demonstrated that these values could also be obtained from directly fitting the first transition times to the cumulative density function (CDF), $h(t) = 1 - S(t)$.

The γ value we obtain from this method is selected automatically while estimating k_0 to be the best estimate for the fraction of the bias that went into decreasing the barrier. As such, γ can be considered to be a measure for the quality of the CV in describing the transition process.

1.5 Thesis Overview

The research work presented in this thesis produced a variation of the KTR method called the exponential average time-dependent rate (EATR) method, which connects the methodology presented in the KTR method to the theory used in hyperdynamics, iMetaD, and OPESf.

In Chapter 2, we generalize the KTR method to apply to any model for how the rate constant changes when a bias potential is added, and derive the EATR method from that. The model for the rate's dependence on the bias is chosen to ensure that there is agreement with iMetaD when the CV quality measure $\gamma = 1$. After this, we show that the EATR method makes accurate predictions for the unbiased rate constant and CV quality in a one-dimensional potential, a coarse-grained protein folding model, and a fully atomistic protein folding model.

In Chapter 3, we apply EATR to a recently developed machine-learned CV[33] to measure the CV quality measure γ over the course of the training process. We are able to observe the CVs improve during training on multiple protein-ligand systems, demonstrating that γ does accurately measure CV quality.

In Chapter 4, we demonstrate that EATR requires a varying bias potential to properly estimate the effect that a bias has on the observed rate. We, then, extend EATR to work for time-independent bias potentials by combining data from multiple sets of simulations with different amounts of bias. We demonstrate that the extended method, EATR-flooding, works well on a coarse-grained protein folding system and a cavity-ligand toy model system.

2 | The Exponential Average Time-dependent Rate Method

This chapter is adapted from the published paper[32]:

Mazzaferro, Sasmal, Cossio, Hocky, *J. Chem. Theory Comput.* **2024**

DOI: 10.1021/acs.jctc.4c00425

2.1 Abstract

Our ability to calculate rates of biochemical processes using molecular dynamics simulations is severely limited by the fact that the time scales for reactions, or changes in conformational state, scale exponentially with the relevant free-energy barriers. In this Chapter, we improve upon a recently proposed rate estimator that allows us to predict transition times with molecular dynamics simulations biased to rapidly explore one or several collective variables (CVs). This approach relies on the idea that not all bias goes into promoting transitions, and along with the rate, it estimates a concomitant scale factor for the bias termed the "CV biasing efficiency" γ. First, we demonstrate mathematically that our new formulation allows us to derive the commonly used Infrequent Metadynamics (iMetaD) estimator when using a perfect CV, where $\gamma = 1$. After testing it on a model potential, we then study the unfolding behavior of a previously well characterized coarse-grained protein, which is sufficiently complex that we can choose many different collective

variables to bias, but which is sufficiently simple that we are able to compute the unbiased rate directly. For this system, we demonstrate that predictions from our new Exponential Average Time-Dependent Rate (EATR) estimator converge to the true rate constant more rapidly as a function of bias deposition time than does the previous iMetaD approach, even for bias deposition times that are short. We also show that the γ parameter can serve as a good metric for assessing the quality of the biasing coordinate. We demonstrate that these results hold when applying the methods to an atomistic protein folding example. Finally, we demonstrate that our approach works when combining multiple less-than-optimal bias coordinates, and adapt our method to the related 'OPES flooding' approach. Overall, our time-dependent rate approach offers a powerful framework for predicting rates from biased simulations.

2.2 Introduction

A major challenge in biomolecular simulation is to be able to accurately assess the transition rate constant (inverse of the mean residence time in a state) of complex processes, including conformational transitions and the binding/unbinding of macromolecules and their ligands. Processes of interest often involve rare events, where the system spends a large amount of time in a metastable state and rarely transitions to another relevant one, so the transition path time is typically orders of magnitude shorter than the time spent in either state [34]. Because of this, extracting rates of such processes directly from unbiased simulation is out of reach for all but the simplest of systems.

Numerous methodologies have been developed to accelerate rare conformational transitions, with the primary purpose being to compute ensemble-averaged observables [35, 36]. A major subclass of such methods operate by adding an additional biasing potential to the system's Hamiltonian, usually in terms of a small set of collective variables (CVs) which are believed or determined to be good descriptors of states of interest, or the path between them [35, 36]. Common exam-

ples of such methods include Umbrella Sampling, Adaptive Bias Force, Metadynamics (MetaD), Variationally Enhanced Sampling, and On-the-fly Probability Enhanced Sampling (OPES), among others [15–17, 37–42]. All of these methods pay the price of distorting the system's dynamics to obtain a much more rapid estimate of the underlying free-energy landscape as a function of the CVs.

Most methods that tackle the problem of computing rates of rare transitions seek to generate a set of unbiased trajectories, either through combining direct sampling of many short trajectories from different starting points [43, 44] as done in Markov State Modeling, through Monte Carlo in trajectory space as in Transition Path Sampling[19], or by generating trajectories that progress in a particular coordinate as in Forward Flux Sampling [45], Steered Transition Path Sampling [46], Weighted Ensemble [47], Transition Interface Sampling [48], and Milestoning [49]. However, these methods are computationally expensive, and some scale poorly with system size, making them challenging to apply for the complex biophysical problems we are interested in studying, such as finding the timescale for protein-drug unbinding[50], for the RBD opening of the SARS-CoV-2 Spike protein [51] or for the unbinding of cytoskeletal adhesion proteins under force [9, 52, 53].

As such, we are interested in approaches that build on CV biasing methods, which have been used to probe conformational transitions with sufficient computational efficiency even for relatively complex biological assemblies. The challenge already mentioned is that these methods alter the dynamics, which prevents any obvious solutions to inferring the unbiased time scales of events. However, starting with the Hyperdynamics method of Voter, it was shown that the first passage time of rare events could be approximately predicted using biased simulations if bias is not applied during the actual crossing through the transition state, by formulating an ansatz for how time is accelerated [25, 26]. This approach was originally developed using a time-independent potential defined on the whole system of interest, but later in the Infrequent Metadynamics (iMetaD) approach the same ideas were extended to CV biasing. MetaD [15] works

by updating an external bias with a Gaussian centered at the current position in CV space every Δ time steps (see Sec. 1.3.3 for details). iMetaD solves the problem of not biasing the transition over the barrier by only rarely updating the bias potential, such that it is unlikely to add bias on a high barrier during a fast crossing. iMetaD also introduces an additional approximation since the system is experiencing a time-dependent bias rather than a static one. The difficulty of avoiding adding bias during barrier crossings can also be mitigated by MetaD variants that only add bias within a region or up to a certain energy level [54, 55], which is now particularly easy to implement in the OPES variant of MetaD [16, 30]. iMetaD and similar approaches have now been used and benchmarked for many different problems, especially for protein-ligand unbinding problems[56], as reviewed in Ref. [57].

To extract the transition rate, these methods assume that a "good" CV is used, and validate the rate estimates using a Kolmogorov-Smirnov (KS) test between the empirical and theoretical survival distributions. Unfortunately, for large and complex transitions, the CV that is used may be poor because finding a good CV is challenging. Moreover, CV quality indicators, such as the committor [20], are expensive or intractable to compute. The Kramers time-dependent rate (KTR) method[31] was recently developed to extract transition rates from biased simulations, such as those used for iMetaD, but with much less sensitivity to CV choice. It introduced a new parameter γ, called the CV biasing efficiency, that scales the effect of the added potential. In that work, it was shown that γ had a lower value for a poor CV in a simple 2D double-well potential, and as such it was assumed to relate to the CV quality [31]. However, this has not been systematically demonstrated, and KTR has not been benchmarked on a problem where many CVs could be tested. Moreover, a direct connection between the KTR and iMetaD estimators has not been established.

In this work, we introduce a more general framework for computing rates from time-dependent biasing protocols, which allows us to treat the iMetaD and KTR estimators on the same footing. We then use this framework to propose a revision to KTR termed the Exponential Average Time-dependent Rate (EATR) method that bridges the two approaches. The EATR approach is shown to

give the correct Kramers' rate when $\gamma = 1$ for an idealized 1D potential. Then, we use a Gō-protein system as a model to show how the prediction of rates depends on the choice of bias coordinate, and compare EATR's results to the true intrinsic rate. Importantly, we find that γ correlates with the intuition of CV quality. We find that for the poor biasing coordinates, the original KTR and EATR results are comparable and they enable an accurate recovery of the unbiased rates. Surprisingly, this is often true even in the frequent-biasing regime. These same general conclusions hold when applied to the folding of the small peptide chignolin, using biased trajectories along three CVs from Ref. [58] provided by the authors.

The chapter is organized as follows. First, we present a general theory for rate calculations from time-dependent biased simulations. We relate it to iMetaD and KTR, and then formulate the EATR approach. Then, we show results for a 1D overdamped Langevin dynamics simulation, and for the unfolding process of two proteins for which we can easily compare biasing different CVs. We also adapt our approach to be used with OPES rather than a MetaD biasing protocol. We show that the method can be extended beyond one biasing coordinate, presenting accurate results on protein G unfolding when biasing two CVs simultaneously. We end with conclusions and future perspectives of the work.

2.3 Theory

2.3.1 Transition rate for rare events

The rare event problem constitutes the stochastic crossing of a single free-energy barrier, where typically the waiting time to cross the barrier is much longer than the transition time over it. For a high barrier, the survival function $S(t)$, which is the probability of a transition not occurring before time t, is given by an exponential distribution characterized by a single

transition rate constant k_0,

$$S(t) = e^{-k_0 t} \,. \tag{2.1}$$

Note that this survival probability is related to the probability of a transition occurring *at* time t via

$$p(t) = -\frac{dS(t)}{dt} \,, \tag{2.2}$$

and it is also related to the cumulative distribution function (CDF), the probability that a transition occurred *by* time t,

$$\mathrm{CDF}(t) = 1 - S(t) \,. \tag{2.3}$$

For Brownian dynamics, Kramers' rate theory[59–62] can be used with several approximations to calculate the barrier crossing rate,[3] k_0, from the bottom of a well on a potential surface $U(x)$ containing a single high barrier,

$$k_0 = D \left[\int_{\text{well}} e^{-\beta U(x)} dx \int_{\text{barrier}} e^{\beta U(x)} dx \right]^{-1} \,, \tag{2.4}$$

where D is the diffusion coefficient.

However, for most systems of interest, the diffusion coefficient and underlying potential (or free-energy landscape) are not known and, therefore, one cannot directly use Eq. 2.4 to estimate the rate. Instead, the transition rate constant can be calculated using the survival function and a set of simulations $i = 1, \ldots, N$ where $M \leq N$ have crossed the barrier and $N - M$ have not. Let t_i be the time the i-th simulation crossed the barrier, and $t_i = T_i$ the total simulation time for simulations $i = M + 1, \ldots, N$. For right censored transition times, the likelihood is given by

$$\mathscr{L} = \prod_{i=1}^{M} p(t_i) \prod_{i=M+1}^{N} S(t_i) \,, \tag{2.5}$$

[3] By convention, the term "rate" is often used in the Kramers' problem when referring to the rate constant/coefficient, and so we also use the term rate when not specifically referring to the rate constant, k.[62]

which is the product of the probabilities of the transitions occurring at times $\{t_i\}_{i=1}^{M}$ and the probabilities of not transitioning before the times $\{t_i = T_i\}_{i=M+1}^{N}$ for $N - M$ simulations [63].

An estimate for the transition rate constant can be obtained by substituting Eq. 2.1 and Eq. 2.2 into Eq. 2.5 and maximizing the logarithm of the likelihood with respect to k_0,

$$k_0^* = \frac{M}{\sum_{i=1}^{N} t_i} \,. \tag{2.6}$$

Note that the summation in the denominator takes into account the simulations that did not transition. When all simulations have crossed the barrier, Eq. 2.6 reduces to the inverse of the average barrier-crossing time, $k_0 = (\overline{t_i})^{-1}$ where the overline denotes the average over the simulations.

The transition rate can also be calculated by fitting the CDF (Eq. 2.3). To do so for the same set of simulations $i = 1, \ldots, N$, we construct an empirical CDF $h(t_i)$ which is the number of simulations that have transitioned before t_i over the total number of simulations. The theoretical CDF can then be fit to the empirical CDF with a least-squares method by optimizing k_0.

2.3.2 Expression for the general time-dependent rate

For time-dependent biased simulations, such as a MetaD simulation, the transition rate is no longer a constant, and hence we would like to formulate a general expression for the survival function in the case of a time varying potential (similar to a situation considered by Zwanzig in Ref. [64]). Without loss of generality, we can write

$$S(t) = e^{-\int_0^t k(t')dt'} \equiv e^{-k_0 \int_0^t f(t')dt'} \tag{2.7}$$

which can be used in Eq. 2.3 for fitting to an empirical CDF. Here, we introduce a time-dependent rate constant $k(t)$, and then re-express $k(t)$ as the unbiased k_0 scaled by a function of time,

$$k(t) = k_0 f(t) \,. \tag{2.8}$$

Because $S(t) = e^{\log(S(t))}$, this is equivalent to defining $k(t) = -\frac{d\log(S(t))}{dt}$, and we can do this because we expect $\log(S(t))$ to be differentiable at $t > 0$ for a physically realizable process.

Substituting Eq. 2.7 into Eq. 2.5 results in a general likelihood given by

$$\mathscr{L} = \prod_{i=1}^{M} k_0 f(t_i) e^{-k_0 \int_0^{t_i} f(t')dt'} \prod_{i=M+1}^{N} e^{-k_0 \int_0^{t_i} f(t')dt'} , \tag{2.9}$$

which can be simplified by taking its logarithm

$$\log(\mathscr{L}) = M \log(k_0) + \sum_{i=1}^{M} \log(f(t_i)) - k_0 \sum_{i=1}^{N} \int_0^{t_i} f(t')dt' . \tag{2.10}$$

Similarly to the unbiased case, we can maximize this expression with respect to k_0, to obtain k_0^* as an estimator for the true rate constant,

$$k_0^* = \frac{M}{\sum_{i=1}^{N} \int_0^{t_i} f(t')dt'} . \tag{2.11}$$

Maximization of the log likelihood function answers the question: what k_0 best describes the observed biased survival times assuming the particular form of the time-dependent rate constant given by Eq. 2.8.

2.3.3 Relation to iMetaD

In the hyperdynamics method [25, 26], the rate from transition state theory is scaled by the acceleration factor $\alpha = \left\langle e^{\beta V(x)} \right\rangle$ where $V(x)$ is a fixed bias function added to the system's Hamiltonian as a function of the system's full coordinates. This α arises by considering the average effect in many individual trajectories whose time is dilated by a factor $e^{\beta V_i(t)}$, where $V_i(t)$ is the bias experienced by the system during simulation i at time t. Hyperdynamics then corresponds

to a rate scaling function of the form

$$f(t) = \left\langle e^{\beta V(x)} \right\rangle = \alpha, \tag{2.12}$$

and replacing it in Eq. 2.7 results in the survival probability

$$S(t) = e^{-k_0 \alpha t} \, . \tag{2.13}$$

Using this expression and assuming all simulations transitioned ($M = N$), the likelihood maximization gives

$$k_0^* = \frac{1}{\alpha \, \overline{t_i}} \, . \tag{2.14}$$

In iMetaD, the form of the bias is also changing in time along with the configuration of the system in a history-dependent manner. Therefore, in iMetaD an acceleration factor for each simulation is approximated as a time average over that simulation instead of calculating it as a configurational average, $\alpha_i = \frac{1}{t_i} \int_0^{t_i} e^{\beta V_i(t')} dt'$. In the Supplementary Information Sec. 2.8.1, we show that if we use this expression in our time-dependent rate formulation, we recover the standard rescaling formula used in iMetaD,

$$k_0^* = \frac{1}{(\overline{\alpha_i t_i})} \, . \tag{2.15}$$

We note that this result is derived using the LM approach for the case where all simulations have transitioned.

In Ref. [29], it was shown that directly fitting the theoretical CDF (obtained from Eq. 2.13) is less sensitive to outliers in the tail of the distribution. The KS test can be used to assess whether the transition distribution is well-described by the theoretical CDF (see Sec. 2.6.4). Recently, it was also suggested that the short time information from the CDF can be fit to get a more robust

estimate of the rate [58]. We note that the results from likelihood maximization (LM) and CDF fitting need not coincide, as we will describe below.

2.3.4 Kramers time-dependent rate and the CV biasing efficiency

Most of the transition-rate methods for biased simulations, such as those described above, are formulated assuming that is possible to apply the bias along a perfect CV, where all added bias accelerates the barrier crossing event. However, for large biomolecular systems, choosing *a priori* a perfect CV for accelerating transitions to another targeted state is almost impossible. In practice, the bias is applied along non-ideal CVs, which insert bias along useless directions that are not aligned with the transition path.

To overcome this issue, the KTR theory [31] introduces a parameter, $\gamma \in [0, 1]$, to account for the efficiency of the biased CVs. In addition to the unbiased rate, γ will also be estimated from the simulation transition times, and it will inform about the quality of the CV with $\gamma \to 0$ reflecting poor CVs and $\gamma \sim 1$ good ones. In the KTR approach as previously implemented, the efficiency of CVs is accounted for by defining the scaling function as

$$f(t) = e^{\beta\gamma \overline{\max V_i(t)}}, \tag{2.16}$$

where $\overline{\max V_i(t)}$ is the maximum bias applied at any point up to time t in simulation i, averaged over all simulations (denoted $V_{MB}(t)$ in Ref. [31]). This form of treating the biasing potential was inspired by rate-calculation methods developed for force-spectroscopy[65, 66], where the barrier is reduced due to the external force, and therefore, by using $\overline{\max V_i(t)}$ it was assumed that the bias only affects the barrier height. Inserting Eq. 2.16 into Eq. 2.7 gives

$$S(t) = e^{-k_0 \int_0^t e^{\beta\gamma \overline{\max V_i(t)}} dt'}, \tag{2.17}$$

which can be used directly in a CDF fit. Substituting this expression into the log-likelihood from Eq. 2.10, and maximizing it with respect to k_0 results in a γ-dependent expression for the unbiased rate constant,

$$k_0^*(\gamma) = \frac{M}{\sum_{i=1}^{N} \int_0^{t_i} e^{\beta\gamma \overline{\max V_i(t)}} dt'} \,. \tag{2.18}$$

To obtain the maximum likelihood estimate for both γ and k_0, Eq. 2.18 is substituted back into the log-likelihood function and it is maximized numerically with respect to γ.

2.3.5 Exponential average time-dependent rate (EATR)

While rate constants computed by the KTR approach are accurate (as shown in Ref. [31] and in the following sections), we show below in Sec. 2.4.1 that it has the undesirable property that it does not agree with the iMetaD estimator when $\gamma = 1$, whereas we expect the iMetaD estimator to be correct for an ideal coordinate with a very high barrier and slow deposition time. The reason for this discrepancy is the way in which the average effect of the bias is defined—for iMetaD $e^{\beta V_i(t)}$ is averaged, whereas for KTR the maximum of the biasing potential is averaged.

To unify the two theories, we propose the following modification to the KTR method, which will have the desired property of producing the same rates as iMetaD in the case where $\gamma = 1$. To do so, we introduce the scaling function

$$f(t) = \overline{e^{\beta\gamma V_i(t)}} \,. \tag{2.19}$$

This gives the time-dependent rate

$$k(t) = k_0 \overline{e^{\beta\gamma V_i(t)}} \,, \tag{2.20}$$

and the survival probability

$$S(t) = e^{-k_0 \int_0^t \overline{e^{\beta\gamma V_i(t')}} dt'} \,. \tag{2.21}$$

Substituting this expression in Eq. 2.10 results in a log-likelihood of the form

$$\begin{aligned}\log \mathscr{L} =& M \log k_0 + \sum_{i=1}^{M} \log \overline{e^{\beta\gamma V_i(t_i)}} \\ & - k_0 \sum_{i=1}^{N} \int_0^{t_i} \overline{e^{\beta\gamma V_i(t')}} dt' \ .\end{aligned} \tag{2.22}$$

In the case where all simulations transition ($M = N$), the optimal unbiased k_0 as a function of γ is given by,

$$k_0^*(\gamma) = \frac{1}{\int_0^{t_i} e^{\beta\gamma V_i(t')} dt'} \ , \tag{2.23}$$

where we have benefited from the idempotence of averages to rewrite the average of an average as a single average. Observing that when $\gamma = 1$, the term within brackets in the denominator is equivalent to $t_i \alpha_i$, this estimator is then identical to the standard iMetaD estimator in Eq 2.15. Similarly as with the KTR, we can now substitute Eq. 2.23 into the expression into the log-likelihood and numerical maximize it with respect to γ to obtain estimates for both the unbiased rate constant and the efficiency of CVs. Incidentally, it might seem from the final step of this derivation that taking the average in Eq. 2.19 was redundant; however, in Sec. 2.8.2 we show that without doing this average, the log-likelihood cannot be maximized with respect to γ.

Importantly, we note that Eq. 2.21 also provides us the option to numerically fit the biased empirical CDF to find the best values of k_0 and γ. As initial guesses, we use the LM estimates for k_0 and γ, and optimize using the Levenberg-Marquardt algorithm implemented in the SciPy Python package[67, 68] to fit the empirical CDF to the theoretical CDF obtained from Eq. 2.21. The same can be done for the KTR method using the theoretical CDF from Eq. 2.17. We note that these optimization procedures are stable for time-dependent biases. In the Supporting Information, we explore their combination with OPES flooding [30] that effectively has a time independent bias.

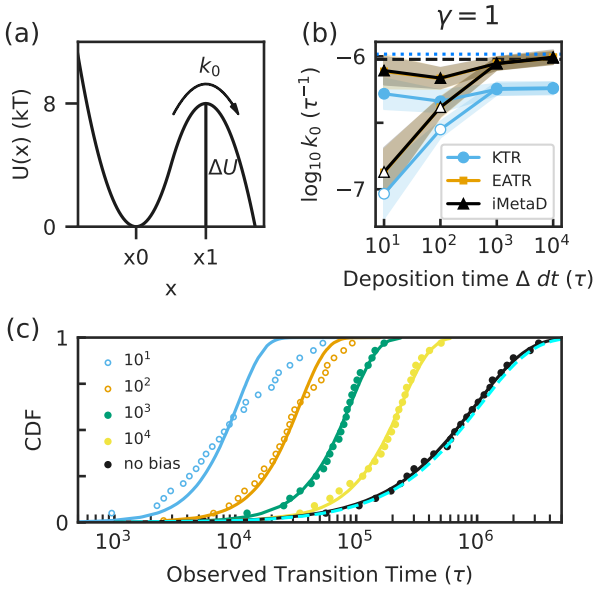

Figure 2.1: (a) The potential energy profile of the 1D matched-harmonic potential from Eq. 2.24. (b) The unbiased rate constant from Kramers' rate theory and from unbiased simulations are shown as the dashed black line and the dotted blue line, respectively. We compare these with predictions from iMetaD, and from the KTR and EATR methods by asserting $\gamma = 1$ as a function of the bias-deposition time (Δ dt). Maximum likelihood estimates are represented by open symbols while CDF-fit estimates are filled symbols. Error bars are from a bootstrap analysis as described in Sec. 2.6.3. (c) The empirical CDFs of observed transition times over the barrier are shown with their EATR-CDF fits; different bias deposition times are indicated in the caption with curves with fastest to slowest biasing appearing from left to right. Fits that fail the KS test are represented with open symbols. The unbiased empirical CDF (black points) is shown together with the Poisson-process distribution fit (black line) and the predicted distribution using the Kramers' analytical expression, Eq. 2.25 (cyan dashed line).

2.4 Results and Discussion

2.4.1 Benchmarking on a 1D Potential

The rate methods were tested first on the one-dimensional matched-harmonic potential illustrated in Fig. 2.1a given by

$$\begin{aligned} U(x) = &\left[\frac{1}{2}\omega_0^2(x - x_0)^2 - \frac{\Delta U}{2}\right](1 - \Theta(x)) \\ &- \left[\frac{1}{2}\omega_1^2(x - x_1)^2 - \frac{\Delta U}{2}\right]\Theta(x)\,, \end{aligned} \tag{2.24}$$

where the subscript 0 corresponds to the well and the subscript 1 corresponds to the barrier, $\Theta(x)$ is the Heaviside step function, and $\omega_i = \sqrt{\Delta U}/x_i$, which is needed to make the potential continuous. Full simulation details for this model are given in Sec. 2.6.1.

To estimate the unbiased rate constant, many Langevin dynamics simulations were performed on the potential from Eq. 2.24, starting from the bottom of the well (Fig. 2.1a). The first barrier-crossing time for each simulation was recorded, and the empirical CDF was calculated as described in Sec. 2.3. The unbiased rate constant was extracted by fitting the distribution to the expected Poisson process. The 2-sample KS test was performed to assess whether this transition is accurately described by a Poisson process. The p-value for the KS statistic is 0.97, demonstrating that the transition times are likely Poisson-distributed data. This yielded a $\log_{10}$ value of -5.98 ± 0.04 where rate constants are in units of τ^{-1}, with τ as the time unit, and the error is the standard deviation obtained from bootstrap analysis[69] (see Sec. 2.6.3). The $\log_{10}$ estimate calculated with Kramers' theory using Eq. 2.4 is -6.02, which agrees with the empirical rate constant within error.

We then performed well-tempered metadynamics (WT-MetaD) simulations (see Sec. 1.3.3) for this system to predict the rates using iMetaD, KTR, and EATR using different bias-deposition times $\Delta\, dt$ with dt the MD timestep (varying from 10 to 10000 τ, which corresponds to fractions of the mean-first passage time varying from 10^{-6} to 10^{-2}). In Fig. 2.1, we compare the methods for both the LM and CDF-fit for the situation where $\gamma = 1$ is enforced, because (*i*) this allows us to only assess the quality of the different time-dependent rate metrics, and (*ii*) for a 1D potential, all bias should go into promoting the transition as there are no orthogonal degrees of freedom; results obtained from fitting both k_0 and γ are shown in Fig. 2.7. We find that the original KTR method does not give a rate consistent with the empirical unbiased rate when $\gamma = 1$. On the other hand, we find that both the iMetaD and EATR methods are consistent with the expected values for the rate. Moreover, in agreement with Ref. [29], we find that fitting the CDF provides more accurate rate estimates than LM at small Δ for all three methods, with the discrepancy

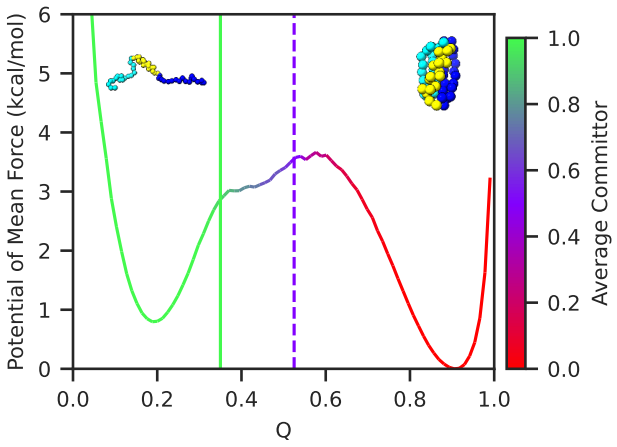

Figure 2.2: The potential of mean force along the fraction of native contacts Q for the Gō-like model of the B1 domain of protein G colored according to the average committor function. The value of Q where the average committor is 0.5 is marked with the dashed line, while the critical value for unfolding $Q = 0.35$ is marked with the solid line.

between the two fits negligible for large Δ. The rate estimates for each fitting procedure improve as Δ increases, which is consistent with the principles of iMetaD. This is also consistent with the results of the KS test. The 2-sample KS test was performed on iMetaD and the 1-sample test was performed for KTR and EATR as explained in Sec. 2.6.4. The KS tests failed for the CDF fits at $\Delta\, dt = 10^1\, \tau$ and $10^2\, \tau$ and passed for the CDF fits at $\Delta\, dt = 10^3\, \tau$ and $10^4\, \tau$. These KS test failures are shown for EATR in Fig. 2.1c as open circles. Given that the rate estimates are more accurate for fitting the CDF (even when the KS test fails), we report results from the CDF-fitting procedure below.

2.4.2 Protein G unfolding

2.4.2.1 Application of KTR and EATR under WT-MetaD bias

We now focus on a more complex system, the unfolding of the B1 domain of protein G using a Gō-like potential and MD simulations. This system has the advantage that it is possible to obtain an unbiased estimate of the unfolding rate, while having a rich unfolding landscape complexity,

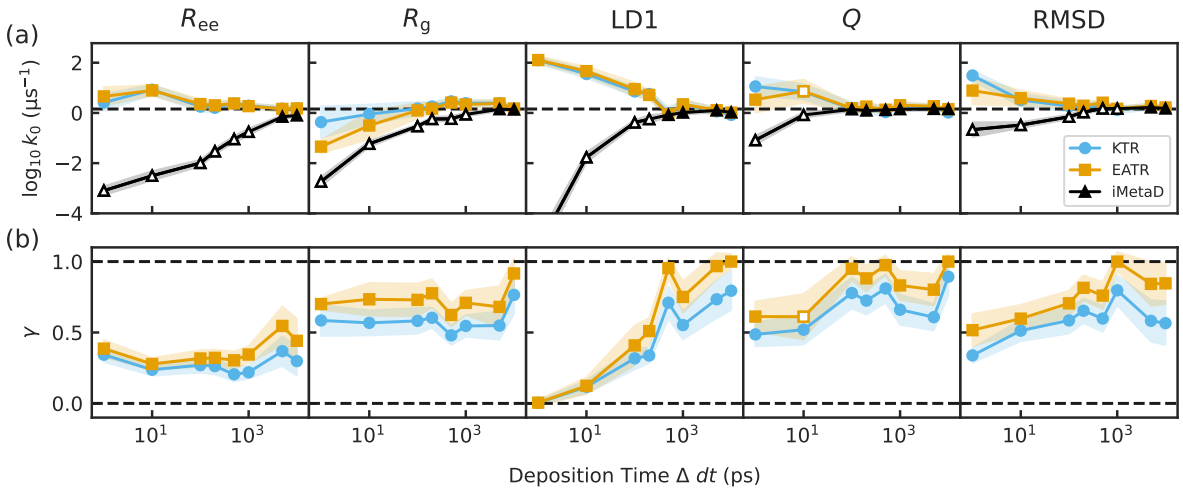

Figure 2.3: (a) The rate constants obtained from fitting the CDF for iMetaD (black), KTR (blue) and EATR (orange) for each CV at various deposition times ($\Delta\, dt$). The horizontal dashed line represents the empirical rate obtained from unbiased simulations. Open shapes indicate where the KS test failed. (b) γ values obtained for the KTR and EATR methods. Horizontal dashed lines represent the bounds placed on γ. In both panels, the error bars are computed from a bootstrap analysis as described in Sec. 2.6.3.

and many possible choices of CVs to characterize the transition [70]. Below, we will evaluate the quality of several good and bad CVs for predicting rates. For this study, we first considered the fraction of native contacts Q and distance between the ends of the protein (R_{ee}) which were shown to be a good and bad coordinate, respectively, for characterizing the folding of this protein in Ref. [70]. In addition to these two CVs, we will consider the radius of gyration (R_g), the root-mean-squared-deviation from the native state (RMSD), and a recently developed linear-discriminant analysis coordinate maximally separating states as defined through a clustering analysis [71] (LD1, see Sec. 2.6.2.2, Fig. 2.8).

We first performed a 120 μs-long unbiased simulation to study the system's behavior. For this unbiased trajectory, we computed the potential of mean force (PMF) along each CV by taking the negative log of the histogram of observed CV values (Eq. 2.26). The PMF for Q is shown in Fig. 2.2, and along all CVs in the SI Fig. 2.9, revealing a range of potential profiles and apparent barrier heights. Although the PMFs of each CV exhibit two wells, we know that R_{ee} is a poor CV

for characterizing unfolding because the unfolded ensemble contains configurations with small values of R_{ee} comparable to the folded state, resulting in an unusually shaped basin at small values of the CV.

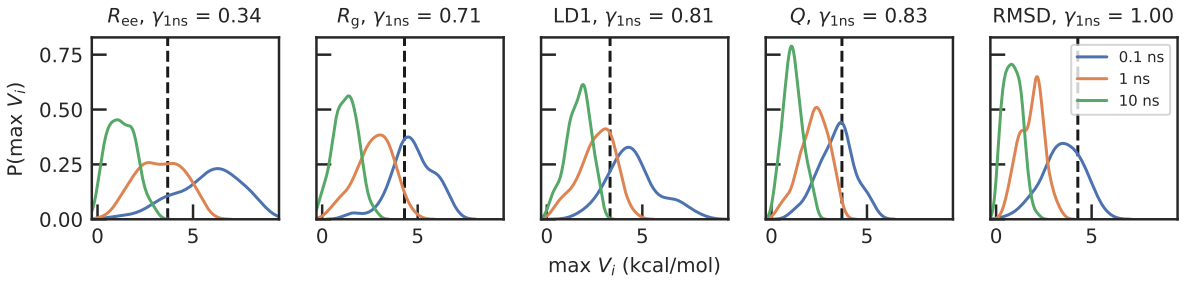

Figure 2.4: Histograms of the maximum bias deposited for the WT-MetaD simulations for the five CVs and different bias-deposition times. As the biasing efficiency γ increases, there is a decrease in the amount of bias needed for transition.

In contrast, we expect Q to be a good CV for unfolding [72], and so we used Q to define when our system has transitioned out of the folded state by using the average committor in the unbiased simulation [20–22]. To do so, we computed for every frame whether the simulation next reached the metastable free-energy minimum on the left (unfolded) before reaching the global minimum on the right (folded), and computed the average by binning these values as a function of the corresponding value of Q in that frame. Based on this result, shown in Fig. 2.2, we defined the unfolded region to be where $Q < 0.35$ because the average committor to the left of this point is effectively 1.0. To estimate the unbiased rate, we ran 200 unbiased simulations of the protein with randomized initial velocities, and stopped these simulations when Q dropped below $Q = 0.35$ using the `COMMITTOR` function of PLUMED [73]. The residence time for the folded state was recorded for each simulation. The unbiased rate constant was determined to be $1.4 \pm 0.1 \mu s^{-1}$ using the CDF-fitting procedure previously described. The KS test using a Poisson distribution passed with $p = 0.65$, demonstrating a good fit.

We then performed 100 biased simulations for each CV at various bias-deposition times to determine how sensitive each method was to biasing speed. The recovered rate constants from the

methods are shown in Fig. 2.3a for each of the biased CVs as a function of the bias-deposition time $\Delta\, dt$. Bias deposition times varied from 1 ps to 10 ns, corresponding to fractions of the mean first passage time ranging from $\sim 10^{-6}$ to $\sim 10^{-2}$ as in the case of the simple potential in the previous section. As expected, longer hill deposition times are observed to generally increase the accuracy of all rate calculations. However, for intermediate to fast-deposition times, KTR and EATR predict unbiased rate constants closer to the true rate than does iMetaD, especially for the three CVs shown on the left of Fig. 2.3a. We also performed a similar study using untempered MetaD and find that, similarly, all methods work well, with KTR and EATR slightly out-performing iMetaD in the fast biasing regime (SI Fig. 2.10).

The KTR and EATR methods also give a measure of the CV biasing efficiency γ, which is shown in Fig. 2.3b. We find that Q and RMSD typically give higher values of γ than the end-to-end distance (R_{ee}) and the radius of gyration (R_g). This coincides with the physical intuition of protein unfolding CVs, where the number of contacts and the similarity to the folded structure should be most relevant. This also agrees with Ref. [72], which found that Q is a good CV for this system. The CV obtained from linear discriminant analysis (LD1, Ref. [71]) appears to have a large value of γ for slow biasing and a small value of γ for fast biasing. A similar trend appears for all CVs tested, but this is most prominent in LD1, and we are still investigating the reason γ for LD1 is so much more sensitive than the other CVs here, while still serving as a very good CV for distinguishing folded and unfolded states (as proposed in our previous study [71]). We note that the discrepancy between iMetaD and the true rate constant is most pronounced when KTR or EATR predict lower values of γ.

Our intuitive expectation is that a bad CV would require significant amounts of extra bias to be deposited before the system can overcome the apparent barrier in the FES for that CV, necessitating a low value of γ to compensate in our rate calculation. To check this, we computed the histogram of the maximum bias across the different simulations at different deposition rates (Fig. 2.4). CVs with high γ and slow deposition have maximum bias that do not exceed the

apparent barrier, while fast biasing and poor CVs require substantial extra energy to be injected into the system to effect transitions. Accordingly, if we look at the average bias as a function of position along a poor and good CV (R_{ee} and Q, respectively, in Fig. 2.11), we find that the amount of bias applied near the transition state is much larger for R_{ee} than Q. Interestingly, even if some amount of bias is added within the transition region, EATR and KTR are still able to recover the true rate for most cases. We note that we use $\max V_i$ in Fig. 2.4 because it is a direct ingredient in the KTR method. Another interesting quantity to compute here would be the amount of non-equilibrium work performed by the MetaD bias, which has been recently exploited in another estimator of rates from time-dependent biased simulations [74].

For all these CVs, the KTR and EATR methods performed comparably well and consistently performed better than the iMetaD method. Interestingly, the KS test seems not to be as sensitive for KTR and EATR as it was in iMetaD. Fitting the CDF for iMetaD results in failed KS tests even where the error in the rate is small, but KTR never failed the KS test for these CVs and EATR only failed for one condition. This may be an effect of introducing γ as an additional fitting parameter, so it is possible to get fits closer to the empirical CDF with worse rate estimates. In Fig. 2.12, we show that the introduction of γ allows us to make good fits to the CDF; indeed, many pairs of (k_0, γ) can be used to fit these data; however, we note that the resulting predicted rate constants are still quite close to the most confident prediction, so this small amount of flexibility is not a problem here in practice.

2.4.2.2 Adapting the approach to OPES flooding

As mentioned in the introduction, the MetaD-like method OPES offers a promising alternative approach to computing rate constants via the estimator given in Eq. 2.15. That is because OPES has a parameter specifying a maximum amount of bias to add, allowing the method to be kept below any apparent barriers to the reaction of interest, and it can be adapted to not include bias outside a pre-specified region of CV space [16, 30]. A brief technical description of the OPES-

MetaD biasing procedure is given in Sec. 1.3.4. Computing rates in this fashion (referred to as OPES flooding by Ref. [30]) satisfies the assumptions required to derive the iMetaD rate estimator without using the infrequent biasing in iMetaD. The OPES flooding bias rapidly converges to its final value, leading in principle to faster observations of the rare events [30, 57]. We expect that this approach should work better than traditional iMetaD for high dimensional systems such as conformational changes in large biomolecular systems.

In principle, our time-dependent rate framework should apply to OPES flooding without modification. We performed six sets of 100 OPES simulations each for the R_{ee} and Q CVs of the protein G model to compare the different methods. Each set of simulations used a different barrier parameter ΔE, ranging between 1 and 4 kcal/mol, with resulting rates shown in Fig. 2.5a.

OPES-flooding performs comparably to standard iMetaD for both CVs (Fig. 2.5a, Fig. 2.13). Surprisingly, we found that the rapid convergence of the bias function results in a log-likelihood for both KTR and EATR which is insensitive to γ, making it very difficult to maximize; a discussion on why this is the case is given in Sec. 2.8.9. Attempting to fit both in practice often causes instabilities with γ tending toward the extreme values 0 or 1 (Fig. 2.5b), and leading to the inaccurate rate constant estimates shown in Fig. 2.5a. Here, we found that EATR performed better than the standard estimator only in some cases and KTR consistently performed worse. Although we could not find stable solutions for both γ and k_0 by optimizing the likelihood, we discovered that we could fit γ and k_0 by determining how the observed unfolding rate constant scaled with the average of a bias measure as described in Sec. 2.8.9 and Fig. 2.13. This 'OPES-EATR-slope' approach results in rate constant predictions accurate to within 5% for both CVs (Fig. 2.5a), but required combining the results from simulations biased using different OPES flooding barriers.

2.4.2.3 Simultaneous biasing of two collective variables

MD studies involving large biomolecules or molecular assemblies will typically have many slow degrees of freedom characterizing transitions between important states, and hence we ex-

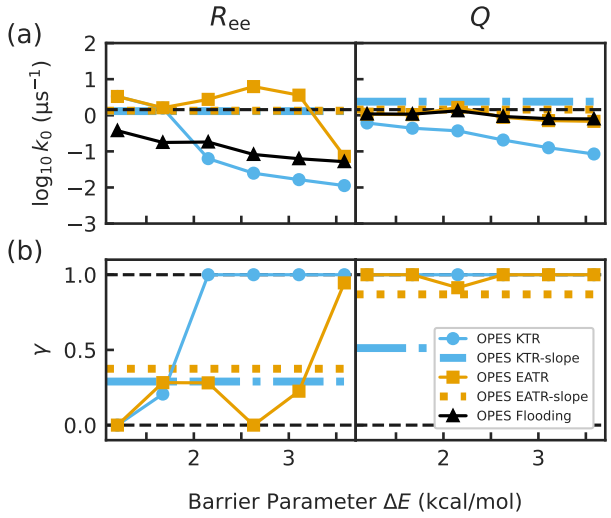

Figure 2.5: (a) The rate constant estimates obtained from applying the iMetaD, KTR, and EATR methods to OPES simulations when biasing R_{ee} (left) or Q (right), varying the OPES barrier parameter. (b) The values of γ obtained from these KTR and EATR fits. Colored horizontal dashed lines in both panels represent the more accurate estimates obtained from fitting the observed OPES unfolding rate constant and γ as described in Sec. 2.8.9, denoted OPES-EATR-slope or OPES-KTR-slope, using different OPES barrier simulations; see Fig. 2.13.

pect the need to use multiple CVs to bias the system in order to promote transitions in a reasonably short amount of simulation time. We wanted to assess whether KTR and EATR still work for this case, despite the fact that the role of γ in characterizing CV quality is less direct. To do so, we performed WT-MetaD simulations while simultaneously biasing the end-to-end distance R_{ee} and the radius of gyration R_g, and summarize the results in Sec. 2.8.10 and Fig. 2.14. Overall, both KTR and EATR recovered the rate equally well apart from EATR for the fastest bias-deposition time, where the KS test failed. For intermediate values of bias deposition time, the value of γ extracted is higher than that of either CV alone, which could be connected to the fact that biasing multiple CVs simultaneously increases the efficiency of the biasing to produce transitions, and this is something we will investigate more rigorously in the future.

2.4.3 Chignolin miniprotein unfolding

To ensure that our method is robust for more complicated atomistic systems, we applied the KTR and EATR methods to the chignolin miniprotein data presented in Ref. [58]. The CVs which were biased in that paper were the C-alpha RMSD to the folded state, the radius of gyration, and a CV obtained from harmonic linear discriminant analysis (HLDA), and for each CV and bias deposition time, 1000 iMetaD calculations were performed.[75] We emphasize that this HLDA CV is different from the LDA CV used in Sec. 2.4.2, and it was the CV that worked the best in Ref. [58]. The rate constants and γ estimates are given in Fig. 2.6. We found that fitting the CDFs for KTR and EATR both dramatically improve rate estimates as compared to the iMetaD estimator (performing at least comparably to the short time fitting in Ref. [58]). Moreover, in this approach we are able to extract a γ parameter showing that, for chignolin, the RMSD to the native is a worse CV for predicting rates of unfolding than radius of gyration, in contrast to our Go-model example. As expected based on the more accurate iMetaD results, γ for HLDA is much closer to 1, validating that it is a good CV for this problem. Although this CV already gives fairly good rate constant predictions for slow deposition, accounting for the fitted γ when using HLDA in both EATR and KTR results in nearly perfect rate predictions at all deposition times.

2.5 Conclusions

In this work, we have developed a general rate theory for time-dependent biased simulations that encompasses several of the existing methods by using the scaling factor $f(t)$ (from Eq. 2.11) with different analytical formulations. In practice, LM and CDF-fitting are two different manners for estimating the unbiased rate constant from an ensemble of simulations launched from a single state. Although we demonstrate in this work that the previously proposed KTR approach works robustly for realistic problems, it does not predict as accurate results as iMetaD for the case of an

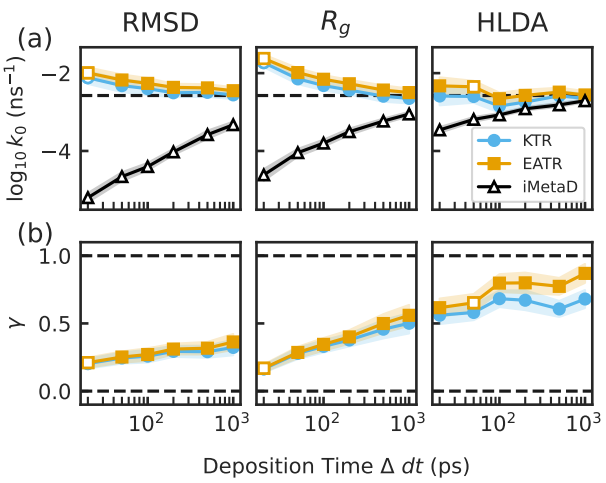

Figure 2.6: (a) The unfolding rate constants obtained from fitting the CDF for iMetaD (black), KTR (blue) and EATR (orange) at various deposition times for each of the biased CVs of all-atom simulations for the chignolin miniprotein from Ref. [58]. (b) The values of γ for each CV for each deposition time. Open shapes indicate where the KS test failed. The error bars were obtained by bootstrapping with subsamples of 200 simulations out of the 1000 in each set.

ideal 1D CV. Therefore, we proposed the EATR formulation, which exactly coincides with iMetaD in the ideal CV case but which gives good rate estimates for bad coordinates more robustly than does iMetaD, with the additional benefit of reporting on the efficiency of that parameter through the parameter γ. Having a high value of γ seems to suggest that KTR or EATR will provide high quality rate constant predictions. We also foresee using γ as a metric that can help us in iteratively optimize our choice of CV(s) for rate calculations. We do note though that having a good CV for a rate calculation may not be the same as having a good CV for predicting a full free energy surface, as it may only be a good way of describing the starting state and configurations on the way to a transition state.

We have validated the methods over the more complex landscape of G-protein and chignolin unfolding where we could still have ground truth. We have found that both the KTR and EATR methods offer accurate rate measurements from biased simulations. The accuracy of the rates determined from these methods are surprisingly insensitive to biasing rate and CV quality, even for frequent-biasing regimes where the average maximum bias likely exceeds the true free energy

barrier. We find that the γ computed from both KTR and EATR report CV efficiencies γ that correlate with our qualitative intuition of what is a good or bad biasing coordinates. Overall, we find that the CDF fit is better than LM to obtain the rate constant, however, many solution pairs of $\{\gamma, k_0\}$ could pass the KS test, so getting a good fit is not sufficient to guarantee that the optimal unbiased rate constant was obtained; however, in practice the predicted values were all very close to the highest confidence prediction (as shown in SI Fig. 2.12). Indeed, for cases where the biasing potential is well approximated as time-independent (such as OPES flooding), we find that optimizing both k_0 and γ simultaneously brings unstable solutions. Nonetheless, by analyzing the unfolding rate constant as a function of a rate scaling function for different OPES barrier parameter values, we are able to overcome this, obtaining stable and accurate estimates.

We find that for 2D biasing landscapes all methods perform reasonably well, which is an indication that biasing many directions might be helpful for barrier-crossing enhancement. In the future, we hope to test the methods on more complex systems using multiple biasing dimensions and, if necessary, extend the theory to multiple dimensions as was done for force-spectroscopy in Refs. [76] and [77].

Despite this success, there are several aspects of the EATR method that can still be improved through future work. For example, we could take into account the effect of the bias on the pre-exponential factor. Although the method works well in the case of our coarse-grained model of protein G for very fast biasing, we have not solved the general problem of how to compute rates in the over-biasing regime where γ times the bias could still be larger than the true barrier, as for example for LDA at fast deposition times (Fig. 2.4), which could lead to the overestimated rate constants and small γ (Fig. 2.3). Addressing this issue will be crucial for systematically using the EATR method for large systems with many slow degrees of freedom. It might be that using a slightly more time-consuming procedure like the OPES flooding EATR-slope fit with multiple barrier heights is the best solution. Going forward, we would like to determine whether it is possible to find a theoretical interpretation for γ in multiple dimensions, e.g. whether it can

be derived considering projection operator approaches, and investigate whether it is rigorously connected to non-equilibrium estimators of the effect of a time-varying bias on the rates [74].

2.6 Computational Methods

2.6.1 Overdamped Langevin dynamics on a 1D potential

We ran overdamped Langevin dynamics over the potential given in Eq. 2.24 using $x_0 = -3$, $x_1 = 3$, and $\Delta U = 8$ $k_B T$. We used an integration time step of $dt = 0.01$ τ where τ is the time unit. The friction parameter used was 0.02 τ, which corresponds to the friction coefficient $\zeta = 50$ τ^{-1} and the diffusion coefficient $D = 0.02$ $\tau k_B T$. All simulations were started from $x = -3$. Because the diffusion coefficient and potential are known, we can derive the standard Kramer's expression in the Smoluchowski limit [61] from Eq. 2.25

$$k_0 = \frac{\omega_0 \omega_1}{2\pi\zeta} e^{-\beta \Delta U} , \tag{2.25}$$

and use that to determine the unbiased rate constant. For this specific system, the theoretical rate constant is 9.49×10^{-7} τ^{-1}. These Langevin dynamics simulations were performed using the PESMD tool in PLUMED [73]. Some of these simulations were biased using WT-MetaD with a starting hill height of 1 $k_B T$, a σ of 0.5, and a biasfactor of 2.0. MetaD simulations were performed for four bias-deposition times (Δ dt): 10^1 τ, 10^2 τ, 10^3 τ, and 10^4 τ.

2.6.2 Gō-like model of protein G

A Gō-like coarse grained model of the B1 domain of protein G was prepared to assess the accuracy of the rate extraction methods, starting from PDB ID 1PGB. This system was selected because it was previously used as a paradigmatic example of a two state folder with known good and bad reaction coordinates [70, 78, 79]. In a Gō-like model, each residue is modeled as a bead at

the position of the α-carbon. The force field for this model treats pseudo-bonds and angles harmonically, and pseudo-dihedrals using a Fourier series. Non-covalent interactions, as in Ref. [80], depend on whether the residues are in contact in the native structure, which is determined by whether the side chains of two residues contain heavy atoms within 4.5 Å of each other. The force field parameters for the model in Refs. [70, 79] were provided by the authors. Our implementation of the potential in LAMMPS [81] and input files for all simulations are provided in the GitHub for the article adapted in this chapter (see Data Availability below).

2.6.2.1 Molecular dynamics simulations

The MD simulations of the Gō-like model were performed using the LAMMPS software package [81]. The software was updated partway through the project and the version used for each set of simulations is shown in Table 2.1. All simulations used a time step of dt =10 fs, and the temperature was held constant at 312 K using the Nosé-Hoover chain thermostat [13] with a damping factor of 1 ps and a chain length of 3. All simulations started from the folded structure.

For the unbiased simulations, we ran 200 replicates and ended the simulations when the protein model unfolded, which was defined to be when the CV Q decreased past 0.35 as described above. The empirical CDF for the transition times to the unfolded state were fit to Eq. 2.1 as explained in the Theory section to obtain the observed unbiased rate for this system.

2.6.2.2 Collective Variables

A variety of CVs were analyzed for the Gō-like model, which were used for the biased simulations and during the rate analysis. The first of these is the fraction of native contacts (Q), which captures the degree to which the protein is folded. This CV is the fraction of the contacts present in the native structure which are still present, and was defined as in Ref. [79]. The end-to-end distance (R_{ee}) was also used, as it was previously determined to be a poor coordinate[70]. The RMSD of the protein with respect to the native structure and the radius of gyration (R_g) were

included to compare with the previously used CVs for this system.

To define the LD1 coordinate, first we performed a cluster scan on the unbiased trajectory using the shapeGMM clustering algorithm [82] with 50,000 frames for training, 3 training sets and 15 attempts each, for cluster sizes (K) = 2,. . . ,6. The training curve with cross validation from the scan is shown in Supplementary Fig. 2.8. We used the positions of the beads as input features for shapeGMM. We did a 5 state shapeGMM fit on the entire trajectory ($\sim$ 1.2M frames) with 15 attempts to identify the distinct clusters. We then performed an iterative global alignment of the trajectory to the global mean and covariance. Multi-state Linear Discriminant Analysis (LDA) was performed on the globally aligned trajectory with frames from all 5 clusters. Only the first coordinate (LD1) out of four resulting LD coordinates has been used in this study.[71] In Supplementary Fig. 2.8b,c we show that this coordinate completely separates the folded and unfolded states, with the other states appearing as intermediates.

2.6.2.3 Biased simulations

The collective variables and biasing for protein G were handled using PLUMED [73]. The version of PLUMED used for each set of simulations is shown in Table 2.1. As is the case for the 1D potential, WT-MetaD was used to bias the simulations. A set of untempered MetaD simulations were also performed, the results of which are provided in the Supplementary Information. The parameters used for the WT- and untempered MetaD simulations are given in Supplementary Tab. 2.2. The values of σ were chosen for WT-MetaD according to the standard deviation of the biased CV in the folded state, and for untempered MetaD σ was chosen to be less than that used in WT-MetaD. We performed simulations at eight different bias deposition times ($\Delta\ dt$): 1 ps, 10 ps, 100 ps, 200 ps, 500 ps, 1 ns, 5 ns, and 10 ns. 100 simulations were performed for each $\Delta\ dt$. The simulations were halted when the protein was determined to have unfolded, or when either wall-clock time reached 48 hours or a total simulation time of 10 μs was reached.

We also performed 6 sets of 100 OPES simulations each for Q and R_{ee}. These sets were run

with ΔE values of 5, 7, 9, 11, 13, and 15 kJ/mol. We used the values of σ from the WT-MetaD simulations as the width of the kernels and used a kernel update time of 1 ns. We excluded bias in the region $Q < 0.65$ when biasing Q and the region $R_{ee} > 2.9$ when biasing R_{ee}.

2.6.2.4 Potential of mean force and committor analysis

A long simulation of protein G was performed and the potentials of mean force (PMFs) along various CVs were determined from the unbiased simulation data using

$$A(\xi) = -k_B T \log P(\xi) , \tag{2.26}$$

where ξ is the CV along which the potential of mean force is computed and $P(\xi)$ is the probability density of ξ obtained by computing a normalized histogram.

Committor analysis[20–22] along Q was done on this long simulation by assigning either 0 or 1 to each frame of the trajectory depending on whether the system visits the folded or unfolded state next, then taking the average for all frames associated with each value of Q. In order to prevent incorrect assignments to either state, for the committor analysis the system was considered to be in the unfolded state when $Q < 0.25$ and to be in the unfolded state when $Q > 0.85$. From this, $Q = 0.35$ was decided to be the critical value for unfolding, as illustrated in Fig. 2.2. This was chosen to be less than the transition state to prevent counting cases where the system enters the transition region, but fails to unfold.

2.6.3 Bootstrap analysis

Errors were obtained from bootstrap analysis[69]. For this analysis, a new set of transition times was constructed by choosing random simulations from the original set with replacement. Once the new set had the same size as the original set, the rate calculation was performed on the new set. This was repeated 100 times and the standard deviation of the log of the rate and γ

across these new sets is reported.

2.6.4 Kolmogorov-Smirnov test

The KS test was performed to assess whether the transition distribution is accurately described by the expected theoretical distribution. For the case of unbiased or iMetaD, it is a Poissonian distribution. The 2-sample KS test was used for the unbiased and iMetaD analyses. This version of the test determines the maximum deviation of the observed CDF from two samples and gives a p-value which, when sufficiently low, allows us to conclude that the samples most likely did not come from the same underlying distribution. We consider the empirical and theoretical distributions to coincide if $p > 0.05$. The 1-sample KS test was used for the KTR and EATR analyses, as generating large random samples from their distributions took a significant amount of time. This version of the test determines the maximum deviation of the observed CDF for one sample from a theoretical CDF, and gave the same results as the 2-sample test in all the cases that were checked.

2.7 Data availability

Inputs for simulations, LAMMPS code for the Gō-like model, code for analysis, data for each system consisting of the value of the CVs and bias versus time, a script for clustering and generating the LDA coordinate, and code for generating figures are available at https://github.com/hocky-research-group/EATR-paper-2024.

2.8 Supplementary Information

2.8.1 iMetaD from the general time-dependent rate theory

Let simulation i have a history-dependent bias $V_i(t)$ which scales the rate for that simulation as

$$f_i(t) = e^{\beta V_i(t)}\,, \tag{2.27}$$

where the "rate for a simulation" is the inverse of the mean observed transition time we would expect for a set of simulations experiencing the same bias potential. Substituting this into Eqs. 2.7 and 2.11 yields

$$S_i(t_i) = e^{-k_0 \int_0^{t_i} e^{\beta V_i(t')}\, dt'}\,, \tag{2.28}$$

with

$$k_0^* = \frac{M}{\sum_{i=1}^{N} \int_0^{t_i} e^{\beta V_i(t')}\, dt'}\,. \tag{2.29}$$

If all simulations transition and we define the acceleration factor α_i to be

$$\alpha_i = \frac{1}{t_i} \int_0^{t_i} e^{\beta V_i(t')}\, dt'\,, \tag{2.30}$$

we can reinterpret the scaling of the rate constant as a scaling of time by α_i

$$S(t_i) = e^{-k_0 \alpha_i t_i} = e^{-k_0 t_i^{\text{rescaled}}}\,, \tag{2.31}$$

with

$$k_0^* = \frac{N}{\sum_{i=1}^{N} \alpha_i t_i} = \frac{1}{\overline{t_i^{\text{rescaled}}}}\,, \tag{2.32}$$

which are the survival function and rate estimator for iMetaD, respectively.

2.8.2 Adding γ directly to the iMetaD rate scaling function

Consider if we had simply added γ to Eq. 2.27 to form a new rate theory. This would yield the rate scaling function

$$f_i(t) = e^{\beta\gamma V_i(t)}\,, \tag{2.33}$$

which, when substituted into Eq. 2.11, gives the LM estimate for the rate constant

$$k_0^*(\gamma) = \frac{M}{\sum_{i=1}^{N}\int_0^{t_i} e^{\beta\gamma V_i(t')}\,dt'}\,, \tag{2.34}$$

which can be substituted along with Eq. 2.33 into Eq. 2.10 to get a likelihood function which could in principle be maximized. However, if we do this and take the derivative with respect to γ, we find

$$\frac{d(\log \mathscr{L})}{d\gamma} = \sum_{i=1}^{M}\beta V_i(t_i) - k_0\sum_{i=1}^{N}\int_0^{t_i}\beta V_i(t')e^{\beta\gamma V_i(t')}dt'. \tag{2.35}$$

The first term is guaranteed to be zero if no bias is deposited at the barrier, since the system is at or past the barrier when it transitions at time t_i, and the second term is strictly less than zero. Therefore, there is no maximum in the likelihood function, so no LM estimate for the rate exists.

2.8.3 Results for a matched harmonic potential where γ is fit rather than being restricted to $\gamma = 1$

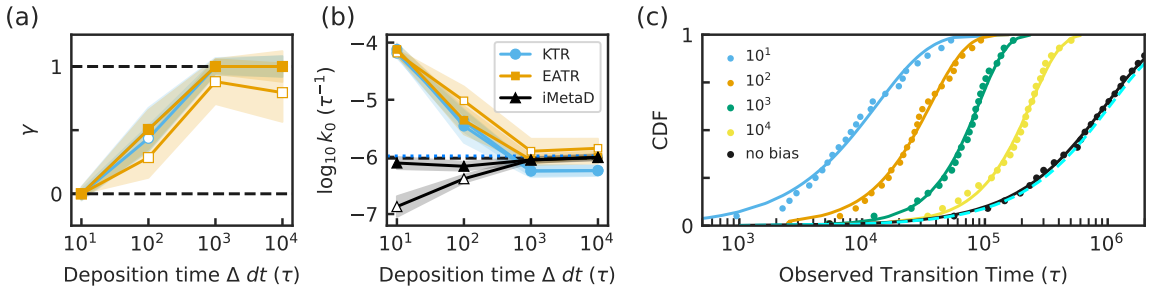

Figure 2.7: (a) γ obtained from the KTR and EATR methods, where the maximum likelihood estimates are represented by open symbols and the least-squares CDF fits are represented by filled symbols. (b) The rate constants obtained from the iMetaD, KTR, and EATR methods. iMetaD provides better rate estimates in this case because this is a perfect CV, which iMetaD assumes while KTR and EATR do not. (c) The best-fit EATR CDFs for each value of deposition time, Δ dt.

2.8.4 ShapeGMM and Position-LDA clustering analysis

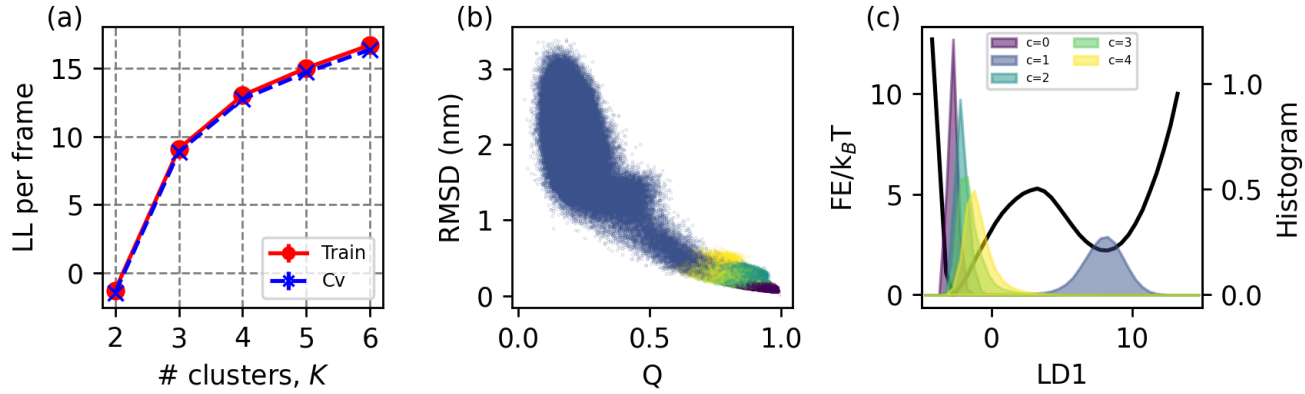

Figure 2.8: (a) Log-likelihoods of data coming from a shapeGMM model trained on K clusters as described in Sec. 2.6.2.2. The agreement between training and cross-validation prediction suggests that the data is not over-fit for any of these numbers of clusters. (b) 2D scatter plot of sampled conformations colored according to their cluster assignments from ShapeGMM using a $K = 5$ model. Based on the position of each cluster in (Q,RMSD) space, we assign cluster 0 to be the folded state and 1 the unfolded state (colors defined in c). (c) 1D PMF profile (solid black) along the LD1 coordinate. Histograms of LD1 corresponding to separate clusters, normalized individually and colored according to cluster id are shown in transparent colors.

2.8.5 PMFs calculated for Gō-like model of Protein G from an unbiased trajectory

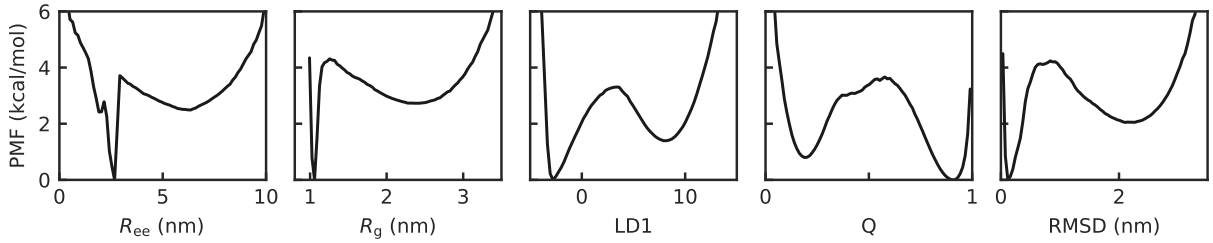

Figure 2.9: PMFs of the Gō-like model of Protein G computed along each of the five CVs presented in the Computational Methods subsection.

2.8.6 Rate estimates using untempered MetaD

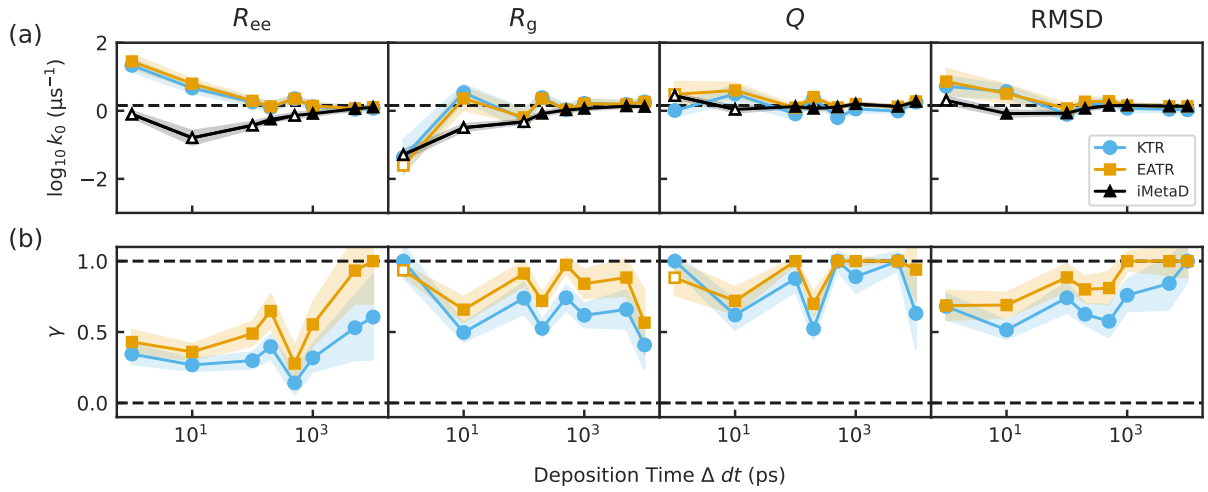

Figure 2.10: (a) The rate constants obtained from the four CVs biased using untempered MetaD using the iMetaD, KTR, and EATR methods. These methods provide accurate results in this biasing scheme. (b) The values of γ obtained from the KTR and EATR methods for each of the CVs.

2.8.7 Distribution of bias deposited during MetaD

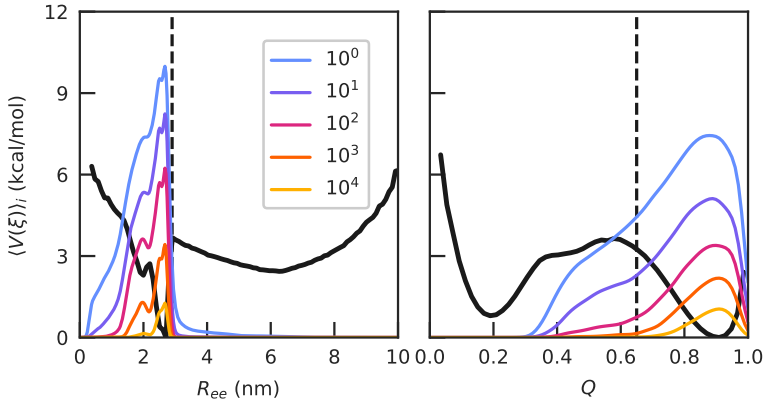

Figure 2.11: The averaged bias as a function of R_{ee} and Q for well-tempered MetaD simulations of protein G unfolding, for different values of $\Delta\ dt$ (in ps) shown in the legend. The vertical dashed lines represent the values of these CVs which were used for the excluded region in OPES flooding. The simulations with smaller $\Delta\ dt$ show overbiasing.

2.8.8 Multiple k_0 and γ pairs pass the KS test

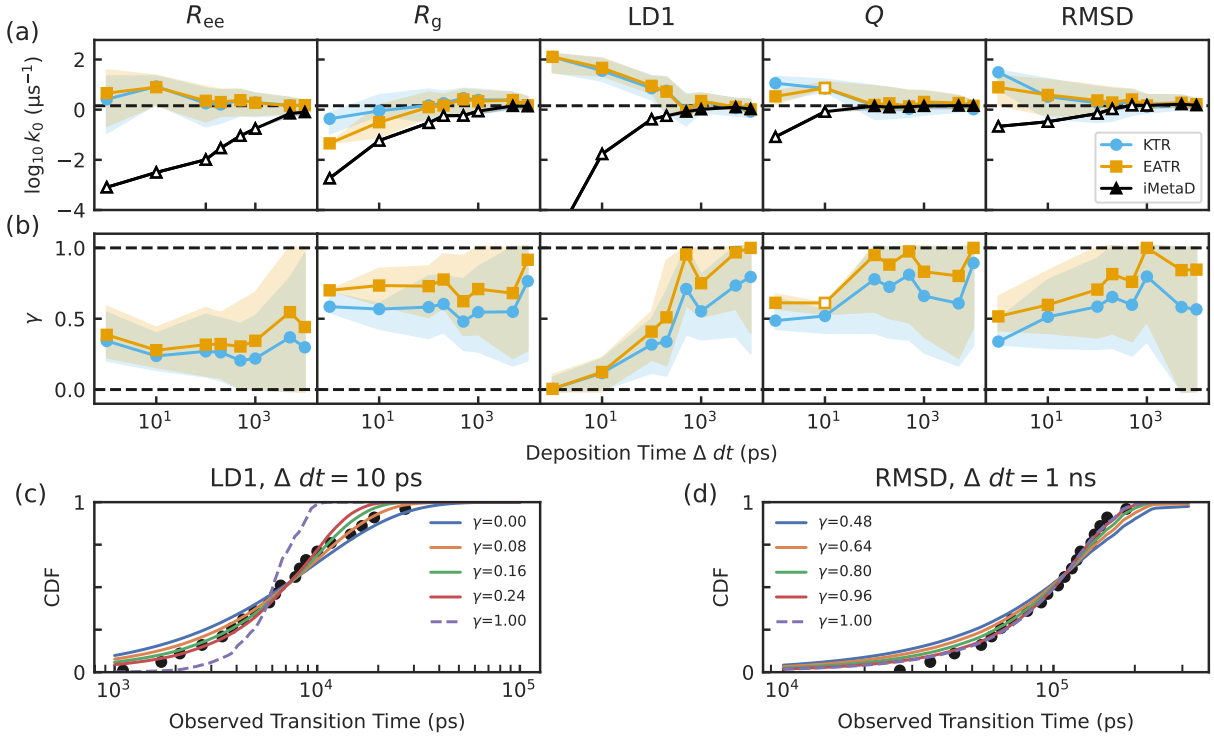

Figure 2.12: Same data as Fig. 2.3. (a) Shaded regions represent the ranges of rate constants for each of the five CVs from the iMetaD, KTR, and EATR methods which pass the KS test for some value of γ. (b) Shaded regions represent the ranges of γ values from the KTR and EATR methods which pass the KS test for some value of the rate constant. (c) Several EATR CDFs which pass the KS test for the LD1 CV at Δ dt = 10 ps. (d) Several EATR CDFs which pass the KS test for the RMSD CV at Δ dt = 1 ns.

2.8.9 Rate estimation for OPES flooding

2.8.9.1 k_0 and γ simultaneous optimization fails for constant bias

Because OPES produces a rough estimation of the free energy very quickly, the biasing potential averaged over many trajectories is effectively time-independent. While this is excellent for enhanced sampling applications, the KTR and EATR methods' k_0 and γ simultaneous optimization fails in cases where the rate is not time-dependent. To see why, consider what would happen if the bias $V_i(t) = V(\xi(t))$ were a constant function of the CV. Then, according to KTR,

the biased observed rate constant (denoted k_{obs}) can be expressed as

$$k_{\text{obs}} = k_0 e^{\beta\gamma V_{\max}} , \tag{2.36}$$

where the survival function is

$$S(t) = e^{-k_0 e^{\beta\gamma V_{\max}} t} . \tag{2.37}$$

Clearly, we cannot fit both k_0 and γ, as they ultimately have the same effect on the survival function: changing the decay constant. Attempting to fit both, even for EATR (where $\left\langle e^{\beta\gamma V(\xi)} \right\rangle$ is constant instead), often causes unstable solutions in the optimization, e.g. shown in Fig. 2.5 where γ tends to 0 or 1. This leads to poor rate estimates.

2.8.9.2 OPES-slope fit

We can overcome the unstable optimization solutions by using the barrier parameter in OPES to determine the relationship between the amount of bias and the biased observed rate constant in a procedure we call OPES-slope. To do this, we ran sets of OPES simulations with different values for the barrier parameter. We measured the observed rate for each set of simulations as the inverse of the biased mean residence time for the simulations. We assume that the observed rate constant is given by

$$k_{\text{obs}} = k_0 e^{\gamma g(\{V_i(t)\})} , \tag{2.38}$$

where $\{\cdot\}$ denotes all values in a set of simulations, and the bias measure $g(\{V_i(t)\})$ is given by the rate scaling functions for KTR and EATR (Eqs. 2.16 and 2.19, respectively). For KTR $g(\{V_i(t)\}) = \beta \langle V_{\max} \rangle$, and for EATR we approximate $g(\{V_i(t)\}) \approx \log \left\langle e^{\beta V_i(t)} \right\rangle$, where $\langle \cdot \rangle$ is an average over the simulations in a set . Although plugging the EATR bias measure into Eq. 2.36 does not exactly lead to Eq. 2.19, we found that this approximation gave good results. In OPES, these bias measures are effectively constant, but to obtain a single value we took the time average for each simulation

set. Finally, taking the natural logarithm of Eq. 2.38, we fit the observed rates as a function of the bias measures, for the sets of simulations, to

$$\log k_{\text{obs}} = \log k_0 + \gamma g(\{V_i(t)\}) \,, \tag{2.39}$$

to obtain estimates for k_0 and γ. In Fig. 2.13, we show $\log(k_{\text{obs}})$ as a function of $g(\{V_i(t)\})$ for OPES simulations with different barrier parameter values for protein G unfolding along R_{ee} and Q. The fit of these points with Eq. 2.39 results in an accurate estimate of the unbiased rate and γ. The values from fitting are: KTR-slope: $k_0 = 2.38\ \mu\text{s}^{-1}$ and $\gamma = 0.51$ for Q, $k_0 = 1.30\ \mu\text{s}^{-1}$ and $\gamma = 0.29$ for R_{ee}; EATR-slope: $k_0 = 1.44\ \mu\text{s}^{-1}$ and $\gamma = 0.87$ for Q, $k_0 = 1.34\ \mu\text{s}^{-1}$ and $\gamma = 0.37$ for R_{ee}. The estimates of the unbiased rate constants are within 5% of the true rate.

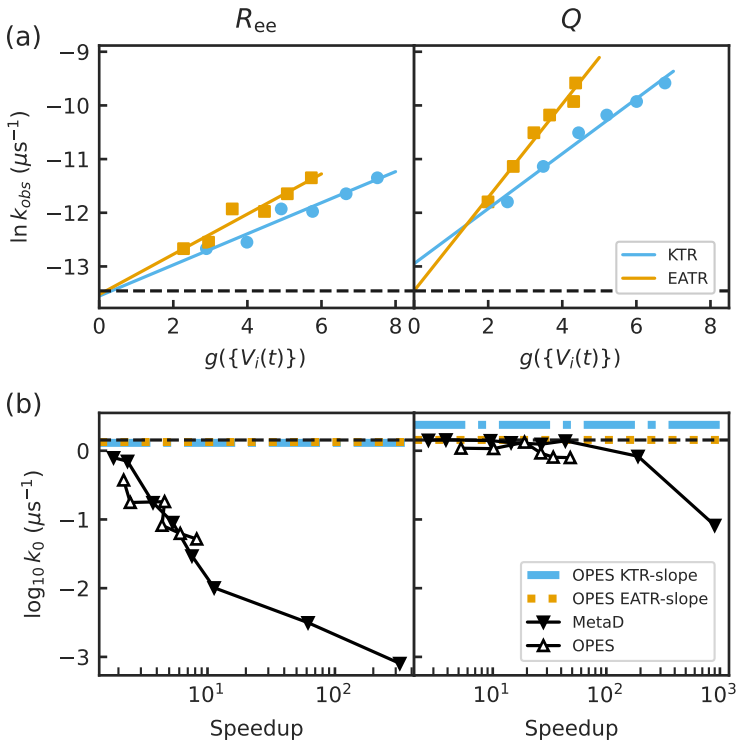

Figure 2.13: (a) The observed rate (inverse mean residence time in the biased simulation) for each set of simulations as a function of the bias measure $g(\{V_i(t)\})$ calculated for KTR and EATR. The black horizontal dashed line marks the true value of $\ln k_0$. (b) A comparison of the rate estimates for MetaD and OPES using the iMetaD estimator. The rates are plotted against "speedup", $\tau_{\text{unbiased}}/\tau_{\text{obs}}$ (see also Ref. [58]). Both biasing schemes give similar results at any given amount of speedup. The blue and orange-dashed lines are the KTR and EATR rate estimates from the fit in panel (a).

2.8.10 2D MetaD on R_{ee} and R_g

We performed MetaD simulations while simultaneously biasing the end-to-end distance R_{ee} and the radius of gyration R_g. For these two CVs, a combined PMF from the long unbiased trajectory is shown in Fig. 2.14a. The rate constants obtained from each method are shown in Fig. 2.14b, and all methods recovered the rate constant equally well, except for EATR at the fastest bias-deposition time, where the KS test failed. The corresponding CDFs and EATR fits are shown in Fig. 2.14c, confirming that the method is able to predict the distribution of transition times. For intermediate $\Delta\, dt \geqslant 10^2$ ps, the value for γ improved in this case over only biasing R_{ee} and improved slightly over only biasing R_g. This suggests that biasing multiple CVs simultaneously increases the efficiency of biasing without affecting the rate estimation. This might also be the reason why iMetaD preforms well in this case, while it failed for the same bias-deposition times when biasing individual CVs (see Fig. 2.3).

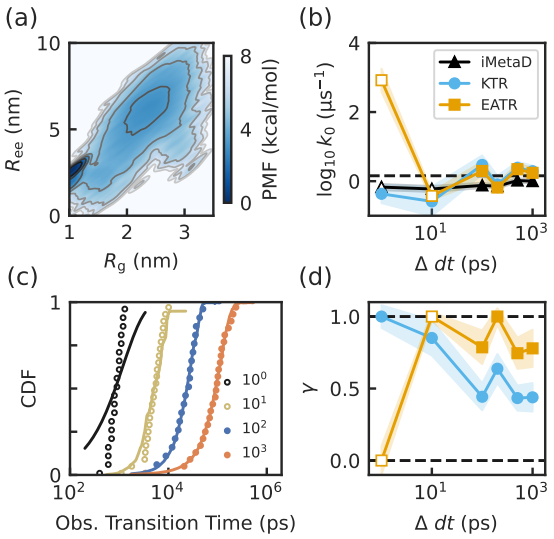

Figure 2.14: (a) The potential of mean force projected onto the 2D R_g/R_{ee} space. (b) The rate constants obtained from each method at various deposition times. The horizontal dashed line represents the true unbiased rate for this system. (c) The empirical and EATR-fit CDFs for each deposition time. The deposition times in the legend are given in ps. (d) The γ values obtained for KTR and EATR at various deposition times. In panels b, c, and d, open shapes indicate where the KS test failed. We performed some simulations with even slower biasing with the result that not every simulation transitioned; while the resulting rate constant estimate was unaffected, the estimated value of γ dropped significantly. The effects on the rate and γ are shown in SI Fig. 2.15.

2.8.11 Effect of un-transitioned simulations on k_0 and γ

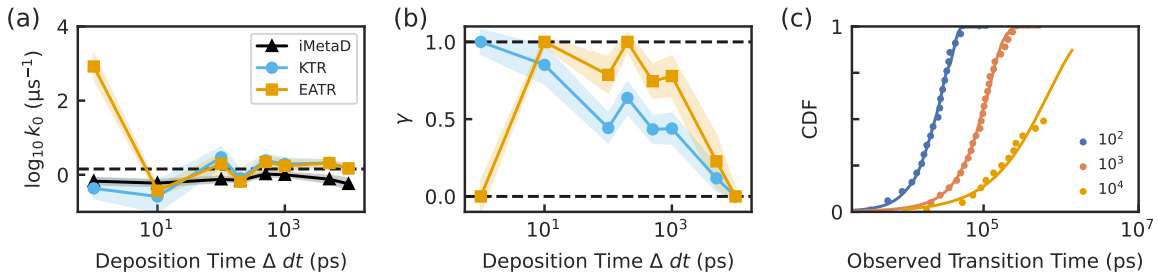

Figure 2.15: Same as Fig. 2.14, with additional slower deposition times where not all simulations transitioned. (a) The rate constants obtained from fitting the CDF for iMetaD, KTR, and EATR for the 2D R_{ee}/R_g CV. The rate constant estimates for $\Delta\, dt = 5$ ns and 10 ns are as accurate as the previous points, although only 66% and 50% of the simulations transitioned, respectively. (b) The values of γ obtained from KTR and EATR. The γ estimates for $\Delta\, dt = 5$ ns and 10 ns are significantly lower than the previous points. (c) The empirical CDF and the best fit EATR CDF for $\Delta\, dt = 100$ ps, 1 ns, and 10 ns. The KS test is not reported for incomplete simulation sets due to how the KS test is implemented in SciPy.

2.8.12 Simulation software and parameters

Simulations were initially performed with the LAMMPS version from 5 May 2020 and PLUMED 2.6.6. We switched to a newer version of LAMMPS and PLUMED to perform simulations with the LDA coordinate, and due to an error in how the older version of PLUMED was computing Q, which we corrected in PLUMED 2.8[1]. The exact versions used for all of our simulations are given in Tab. 2.1.

[1] https://github.com/plumed/plumed2/pull/951

Simulations	LAMMPS	PLUMED
Unbiased 1D	—	v2.8.1
Biased 1D	—	v2.8.1
Unbiased GB1	5 May 2020	—
R_{ee}	5 May 2020	v2.6.6
R_g	5 May 2020	v2.6.6
LD1	23 Jun 2022	v2.8.3
RMSD	5 May 2020	v2.6.6
Q	23 Jun 2022	v2.8.3
R_{ee}/R_g (2D)	23 Jun 2022	v2.8.3
Q and R_{ee} OPES	23 Jun 2022	v2.8.3

Table 2.1: The version of each software used in each simulation set.

Simulations	WT HEIGHT	WT SIGMA	WT BIASFACTOR	UT HEIGHT	UT SIGMA
1D	1.0 k_BT	0.5	2.0	—	—
R_{ee}	0.4 kJ/mol	0.06 nm	10.0	0.2 kJ/mol	0.01 nm
R_g	0.4 kJ/mol	0.02 nm	10.0	0.2 kJ/mol	0.01 nm
LD1	1.0 kJ/mol	0.5	6.0	—	—
RMSD	0.4 kJ/mol	0.02 nm	10.0	0.2 kJ/mol	0.01 nm
Q	0.6 kJ/mol	0.02	10.0	0.2 kJ/mol	0.01
R_{ee}/R_g (2D)	0.6 kJ/mol	0.06 nm/0.02 nm	10.0	—	—

Table 2.2: The MetaD parameters used for the well-tempered and untempered MetaD simulations for each CV.

3 | Measuring the Quality of a Machine Learned Collective Variable

This chapter is adapted from the published paper[83]:

Mazzaferro†, Lee†, Cossio, Tiwary, Hocky, *J. Phys. Chem. B* **2025**

DOI: 10.1021/acs.jpcb.5c04626

†These authors contributed equally to the reproduced part in this thesis.

3.1 Abstract

Having an accurate reaction coordinate (RC) is essential for reliable kinetic characterization of molecular processes, but there are few quantitative metrics to evaluate RC quality. In this chapter, we consider the dimensionless γ metric from the Exponential Average Time-dependent Rate (EATR) method, which represents the fraction of a biasing potential along the RC that contributes to increasing the rate constant. We demonstrate that γ can be used to test whether the utility of a reaction coordinate for predicting kinetics with a Metadynamics bias improves as the coordinate is iteratively updated to include new data. We evaluate reaction coordinates approximated via the iterative State Predictive Information Bottleneck (SPIB) approach, which was previously shown to be accurate across six protein–ligand dissociation systems. For these same systems, we compute γ values and mean accelerated times $\overline{\tau}_{\text{accel}}$. After systematically scanning over fit-

ting parameters, the results show that γ increases closer to 1, while $\overline{\tau}_{\text{accel}}$ decreases, revealing a consistent inverse correlation. These results demonstrate that γ serves as a practical criterion for RC evaluation and offers guidance for selecting SPIB-derived coordinates yielding quantitative kinetic predictions.

3.2 Introduction

Determining biomolecular kinetics is crucial in many areas of biochemistry and biophysics, such as protein folding, ligand unbinding, membrane permeation, and other cellular events. Molecular dynamics (MD) simulations offer the promise of being able to predict kinetics of such processes from simple physical principles.[1, 84, 85] In practice, however, the dynamical processes of interest are "rare events," the typical timescales of which are dominated by crossing over high free energy barriers, and also dictated by the need to diffuse throughout a metastable state before escaping. Consequently, direct MD simulations that access the timescales of microseconds are rarely sufficient to sample enough configurational space to estimate the transition times of key biological processes, which can range from microseconds to seconds. [86, 87]

To address the sampling problem, a wide range of enhanced-sampling approaches based on adding additional bias to a set of collective variables (CVs), most notably umbrella sampling, metadynamics (MetaD), and on-the-fly probability enhanced sampling (OPES) have been developed to accelerate exploration.[16, 17, 37, 88] After sampling with an applied bias along one or several CVs, equilibrium averages in the unbiased ensemble can be obtained by reweighting.

A separate challenge arises when one is interested in computing dynamical quantities dominated by rare events. Because enhanced-sampling approaches accelerate exploration, they also distort the underlying kinetics of the simulated system. Surprisingly, it is still possible to extract meaningful kinetics from CV-based biasing methods with several approximations. Approaches for doing so include hyperdynamics, infrequent metadynamics (iMetaD), OPES flooding, and ap-

plying the Kawasaki relation.[26, 28, 30, 89] While many of these approaches have been carefully validated against model and real world systems,[57] most of these methods require a good CV for biasing that approximates the true reaction coordinate (RC). This is because applying a bias potential along a poor CV has a less understood effect on the slow transition of interest, which must be accounted for. Thus identifying good CVs remains a critical bottleneck, especially for complex systems with many degrees of freedom.

A good biasing variable that approximates the RC should both distinguish the metastable basins of interest and characterize the dynamics of the slow collective modes of a molecular process of interest. The most ideal RC is considered to be the 'committor', a function that reports the probability of ending within a product or reactant state at some later time for a current configuration. It is thus the best variable to describe a transition process, as it separates the metastable states of interest by definition and includes the kinetic information necessary to estimate rates. [22, 90–92] In practice, despite recent progress to compute the committor using machine learning approaches [93–96], the committor is still difficult to obtain, and it is not easily generalizable across different conditions. [90] Instead of using the committor as a direct RC, many other approaches have been proposed to learn an approximate RC, including transition-path sampling, diffusion maps, Linear Discriminant Analysis, the variational approach to conformational dynamics, time-lagged independent component analysis, VAMPnets, SGOOP, and others.[19, 20, 71, 97–103] Building upon these methods, recent studies have shown that incorporating deep learning approaches can improve on RC discovery.

In particular, dimensionality-reduction approaches integrated with deep neural networks have shown much success. [102–105] However, compressing a high-dimensional process into a single coordinate can lead to the loss of information to accurately describe the reaction mechanism. This loss of information is a major obstacle in extracting a good RC and building a reliable kinetic model. Tiwary and coworkers addressed this challenge by performing an iterative process through the State Predictive Information Bottleneck (SPIB) framework,[106] to learn the

RC associated with kinetic rates.[33] SPIB employs a deep neural network that identifies a low-dimensional representation of the system's high-dimensional dynamics that is maximally predictive of the future state of the system after a specified lag time. This approach ensures that the learned RC captures the most predictive features of the system's slow dynamics, effectively approximating aspects of the committor function. By retraining the network after each round of accelerated sampling, it is believed that SPIB progressively refines the RC. In other words, this proposed algorithm, using SPIB, learns from newly sampled trajectories and then guides subsequent simulation rounds toward more relevant features. Here, we ask the following question: Does this iterative cycle genuinely improve the RC?

For rare events such as ligand dissociation, it is difficult to determine the quality of a collective variable since we do not typically have the true RC to compare with. These considerations raise a broader issue: How can we systematically evaluate the quality of RCs for obtaining rates from biased simulations? For unbiased simulations, one way to determine whether an RC captures the rare event of interest is by performing a Kolmogorov-Smirnov test and using the p-value to determine if the transition time distribution is Poissonian.[29]. Another way is to calculate the committor probability as a function of the RC to ensure that the reactant and product states are well-separated.[20] We can also rank RCs using the Bayesian transition-path probability $p(\mathrm{TP}|x)$, as has been demonstrated by Best and Hummer.[92] However, applying these metrics requires unbiased simulation data that includes hundreds of transition events, which is often a great luxury to have for a rare event such as ligand dissociation. To evaluate the RC quality in our biased simulations, we instead use the exponential average time-dependent rate (EATR) method,[32] which provides the dimensionless metric γ that was originally introduced in the Kramers time-dependent rate (KTR) approach [31]. This γ quantifies how the RC influences transition times and thus captures how well the bias potential is applied along the RC in the rate-determining directions. When $\gamma = 1$, the bias acts along the true transition path and has the greatest effect on the transition times, indicating a perfect RC; when $\gamma = 0$, the bias acts orthogonally to the

transition path, indicating a poor RC. In this manner, γ serves as a quantitative benchmark of RC performance (see Sec. 2.3.5 for details). This SPIB + EATR pair, therefore, offers a self-contained, inexpensive, and quantitative solution to observe improvement during RC training. In this context, our study is driven by two central questions: *i*) Does the iterative SPIB procedure enhance the RC for kinetic studies performed here using metadynamics, and *ii*) Can the EATR metric γ reliably quantify the quality of any candidate RC? Our results answer both questions affirmatively: γ offers a sensitive benchmark for RC comparison, and each round of SPIB refinement leads to improvement in RC quality for more accurate kinetic predictions.

The remainder of the chapter is organized as follows. In Sec. 3.3, we review the background on SPIB, infrequent metadynamics (iMetaD), and EATR. We then describe the workflow for assessing RC quality using the EATR γ values for the SPIB RC iterations. In Sec.3.4, we report the EATR γ values for the iterations of the SPIB RC and present further analysis on the γ metric. Finally, in Sec. 3.5, we summarize our key findings, discuss the limitations, and propose future directions.

3.3 Background and method

In this section, we briefly review the concepts of SPIB,[106] iMetaD,[28] and the EATR[32] γ parameter. We then present our integrated workflow, detailing how SPIB-derived reaction coordinates guide iMetaD simulations and how EATR is subsequently applied to quantify RC quality (Fig. 3.1).

3.3.1 State Predictive Information Bottleneck (SPIB)

The SPIB method approximates a low-dimensional RC that preserves the metastable dynamics present in the system's original high-dimensional space. [106] SPIB is inspired by another information bottleneck-based framework, Reweighted Autoencoded Variational Bayes for Enhanced Sampling (RAVE)[107, 108], but improves upon this formulation. The SPIB approach has

demonstrated notable success in learning effective RCs across diverse applications, ranging from crystal nucleation and ligand permeation through lipid membranes to protein conformational changes.[33, 109–113] SPIB CVs can be used with the PLUMED open-source enhanced sampling library [73], and a tutorial for training and using SPIB CVs is available from the PLUMED-tutorials resource [114].

SPIB determines a low-dimensional RC that effectively predicts the future metastable states of the system while relying on minimal information from the current configuration. SPIB uses raw MD trajectory data through an encoder-decoder network framework, where the encoder compresses the input into a bottleneck variable that minimizes mutual information with the input, and the decoder maximizes mutual information between the bottleneck and future states after a time lag Δt. By trading off the minimization of input information and the maximization of future predictive information, SPIB simultaneously learns the RC and discovers the location and number of metastable states on-the-fly. A mixture-of-Gaussians prior is used for the bottleneck variable, where the number of Gaussians automatically adjusts to equal the number of metastable states. We emphasize that the number of metastable states in a system is not an intrinsic property of the system, but depends on the time resolution being used to filter out the dynamics. In this spirit, the number of metastable states learned in SPIB decreases as a function of the time-lag parameter.

A notable success of the SPIB framework is its use in protein–ligand dissociation kinetics. In our previous work of Ref. [33], we implemented an iterative SPIB pipeline where each newly learned RC guides a subsequent round of iMetaD (see Sec. 1.4.3), and the resulting trajectories are used to retrain SPIB. This iterative process not only accelerates the observation of the unbinding process but also demonstrates success in SPIB refinements to a more descriptive RC, yielding more accurate dissociation-rate constant estimates.

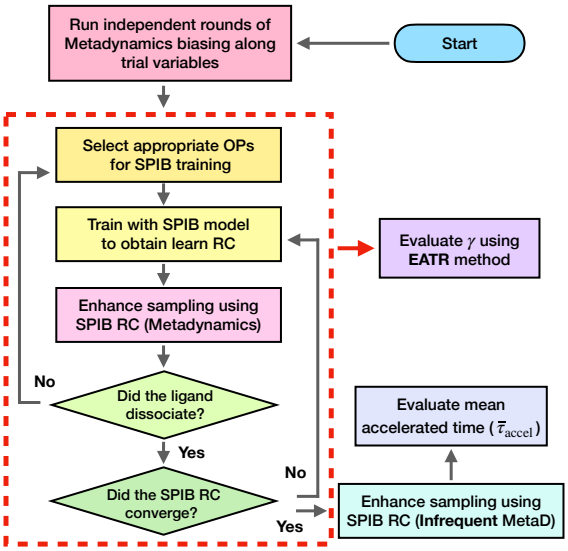

Figure 3.1: Workflow of SPIB combined with EATR. The gray arrows in the protocol indicate the original SPIB–iMetaD workflow used to compute the residence time of protein–ligand complexes. In Ref. [33], we first performed initial trial runs of WT-MetaD biasing along trial order parameters (OPs). We then iteratively applied SPIB and metadynamics until the accelerated time was quantitatively reduced for the SPIB RC. Finally, we ran iMetaD biasing along the converged RC. For each RC at every SPIB iteration, we conducted an additional step of γ evaluations, which is indicated as a red arrow. The red dotted box highlights the iterative process of the SPIB.

3.3.2 Overall workflow

Building on the background described above, we now present an integrated SPIB–EATR framework for quantifying RC quality across multiple protein–ligand dissociation systems. We examine six distinct protein-ligand dissociation systems where residence times span approximately 12 orders of magnitude from ~ 10 ns to $\sim 10^3$ s. The systems studied are FKBP–DMSO, FKBP–DSS, T4 Lysozyme–benzene, ABL kinase wild-type (WT) bound to Imatinib, and two ABL kinase mutants (N368S and L364I) bound to Imatinib. Biased trajectories of SPIB iterations for each system were taken directly from Ref. [33], with additional data added for some cases indicated below. All data were generated using GROMACS [115] and PLUMED [73] using the protocol described in Ref. [33].

The overall workflow, summarized in Fig. 3.1, begins with well-tempered metadynamics along trial order parameters (OPs) chosen by chemical intuition. The gray arrow in the workflow diagram denotes the steps originally described in Ref. [33]. Trajectories from this initial metadynamics run seed the first SPIB training round; the RC learned in that round is then used to bias along new MetaD simulations. This iterative cycle repeats until the average accelerated time no longer increases. A final set of iMetaD simulations are carried out along the converged RC. Detailed choices of OPs, SPIB hyperparameters, and MD setup were used as reported in Ref. [33].

Next, we apply the EATR analysis to the MetaD simulations performed during the SPIB iterations. For each system, we extract the biased trajectory data from each SPIB iteration, typically eight trajectories per iteration, and compute the mean accelerated time $\overline{\tau}_{\mathrm{accel}}$. We then use the corresponding bias potentials to evaluate the γ metric at each iteration, thereby quantifying RC performance without performing any extra simulations after those used in the SPIB iterations. In the following section, we present these numerical results, describe the EATR parameterization, and discuss how the γ values inform RC quality assessment.

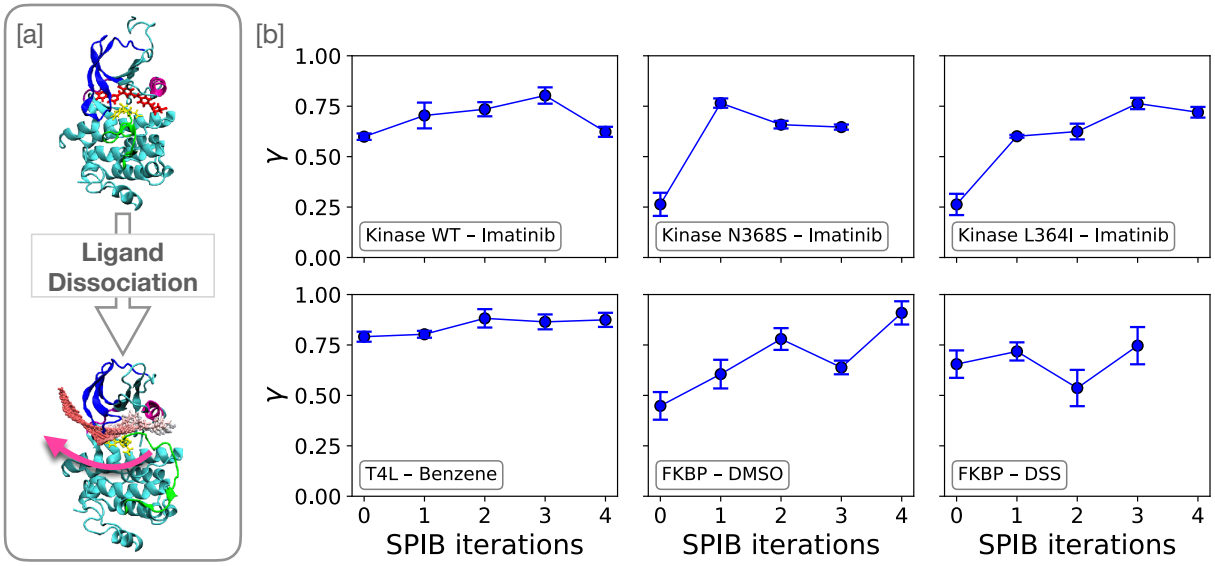

Figure 3.2: [a] Schematic of Imatinib dissociation from the Abl-kinase. [b] Values of the EATR metric γ over SPIB iterations for six protein–ligand systems while fixing k_0 to the experimental value. ABL kinase Wild Type (WT) - Imatinib, ABL kinase mutant N368S - Imatinib, ABL kinase mutant L364I - Imatinib, T4 Lysozyme - Benzene, FKBP - DMSO, FKBP - DSS. Error bars denote standard deviations of fitting uncertainty in γ.

3.4 Results

In this section, we present two main results. First, we analyze how the RC produced by SPIB at successive iterations relates to the EATR metric γ. Second, we provide brief mathematical insights on correlations between the γ and the average accelerated time. Finally, we demonstrate that these results are not strongly dependent on the rate at which bias is applied to a particular RC.

3.4.1 Evaluation of SPIB RC via EATR γ

To quantify how successive SPIB iterations refine the RC, we first computed the EATR-derived metric γ at every iteration for the six protein–ligand systems examined.[33] In the original EATR approach, both the experimental rate constant k_0 and γ are treated as free parameters. Here we compared two fitting protocols. In the first approach, the unbiased rate constant k_0 was held fixed

at the experimental value reported in SI Table 3.1 for each system. The true unbiased simulation rate constant may be different from the experimental rate constant for these systems due to force field inaccuracies, which will introduce an error in γ. However, we feel the difference in the 'true' rate constant for the simulation system and experiment will not affect the relative values of γ between SPIB iterations within a system.

Fixing the unbiased rate allows us to fit the survival function only to a single parameter γ, which makes the fitting more reliable by removing the possible correlations between γ and the rate constant k_0. Although this approach would not apply to all cases, we use this single parameter approach to assess whether SPIB RC quality improves across iterations. We also applied our general EATR approach to fit both γ and k_0, with qualitatively similar results for whether the CV is improving, but much lower γ values showing there is likely still more that can be done to get more efficient rates as discussed in the conclusions (see SI Fig. 3.5).

As shown in Fig. 3.2, the results indicate that γ increases with SPIB iteration for all systems converging toward 1, indicating the improvement during the SPIB cycles. Despite the overall upward trend in the plots, we note the final SPIB iteration in Fig. 3.2 does not always converge to the highest γ. In subplot (a) Abl kinase WT–Imatinib, for instance, we extended several iterations beyond the point where optimal RC was chosen from previous work of Ref. [33] to see if γ would continue to rise toward 1. However, as the results demonstrate, simply prolonging SPIB iterations does not guarantee the highest γ. This indicates that indefinite retraining does not necessarily yield an optimal RC in several systems. This is, however, not a problem in practical terms as one can work with the RC with the highest γ value from all iterations.

3.4.2 γ CORRELATES WITH $\overline{\tau}_{\text{accel}}$

Figure 3.2 shows an increase in the EATR metric γ over successive SPIB iterations. In contrast, our earlier study demonstrated that the mean accelerated time, which is the unbiased time predicted by iMetaD, denoted here as $\overline{\tau}_{\text{accel}}$ (average of the accelerated time, computed directly,

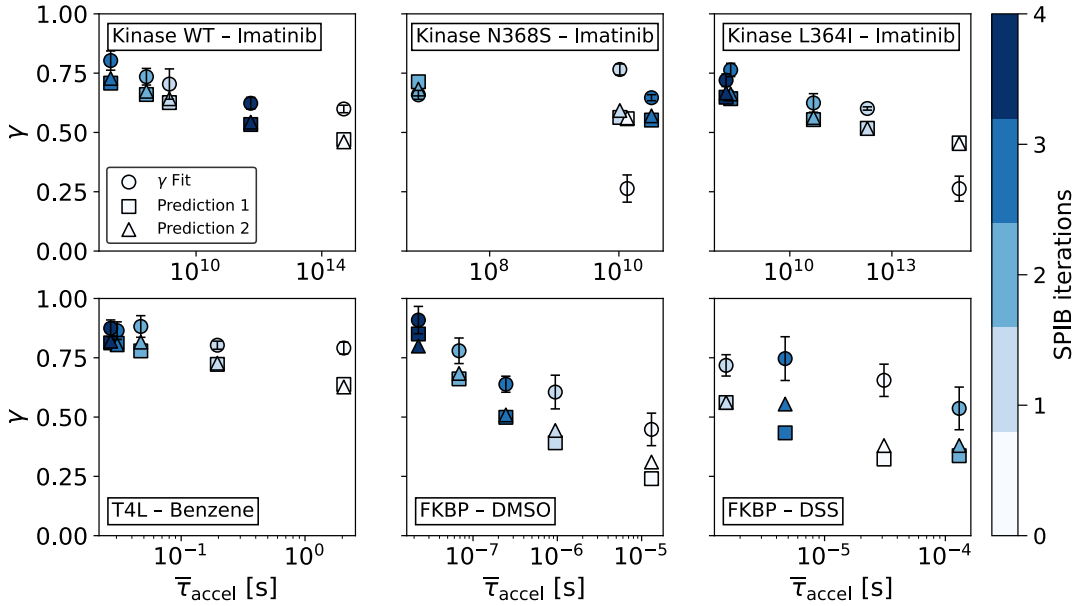

Figure 3.3: Inverse correlation between γ and average accelerated time $\overline{\tau}_{\text{accel}}$ across SPIB iterations. Scatter plots are shown for six benchmark systems (arranged from top left to bottom right): Kinase WT–Imatinib, Kinase N368S–Imatinib, Kinase L364I–Imatinib, T4 lysozyme–Benzene, FKBP–DMSO, and FKBP–DSS. Circles denote γ values obtained from the original EATR fit, squares ("prediction 1") show γ evaluated with Eq. 3.1, and triangles ("prediction 2") show γ evaluated with Eq. 3.2. Different colors indicate the SPIB iteration, ranging from light blue (early iterations) to dark blue (later iterations). Error bars denote standard deviations of fitting uncertainty in γ. An alternative presentation of these data as a function of iteration is given in Fig. 3.6.

without fitting to a Poisson distribution, to facilitate straightforward comparison) decreases as the RC is refined. Taken together, these opposing trends point to an inverse relationship between γ and $\overline{\tau}_{\text{accel}}$, as illustrated for all six protein–ligand systems in Fig. 3.3 (color scale denotes SPIB iteration index). To replace this empirical observation with a more rigorous argument, here we derive an analytical expression that links γ to the mean accelerated time $\overline{\tau}_{\text{accel}}$ and then compared the predictive quantities against the fitted γ values. Starting from the EATR formalism and isolating the terms involving γ, we obtain

$$\gamma \approx 1 - \frac{\log(k_0(\overline{\tau}_{\text{accel}} - \text{cov}(\tau', \alpha)))}{\log(\overline{\alpha})}, \tag{3.1}$$

where τ' is the transition time observed in the biased simulation and $\tau_{\text{accel}} = \alpha\tau'$ is the unbiased time predicted with iMetaD. This can be further reduced to

$$\gamma \approx 1 - \frac{\log(k_0\overline{\tau}_{\text{accel}})}{\log(\overline{\alpha})}, \tag{3.2}$$

when the covariance term $\text{cov}(\tau', \alpha)$ is negligibly small compared to $\overline{\tau}_{\text{accel}}$. The detailed derivation can be found in the SI.

To assess how well the analytical expressions reproduce the fitted results, Fig. 3.3 compares the γ values predicted by Eq. 3.1 (squares, "prediction 1") and Eq. 3.2 (triangles, "prediction 2") with the values obtained from direct EATR fitting approach (circles). The two analytical approximations practically overlap, confirming that $\text{cov}(\tau', \alpha)$ is indeed small in most cases. Although both predictions show some, system-specific deviations from the fitted γ, they capture the same overall trend, demonstrating that either expression can serve as a fast diagnostic tool when a full survival-curve fit is impractical.

Collectively, the results demonstrate the reciprocal relationship between γ and $\overline{\tau}_{\text{accel}}$. In early iterations ($\gamma \approx 0.25 - 0.5$), the RC is non-optimal, whereas in later iterations ($\gamma \approx 0.7 - 0.9$) the RC captures much better the true slow mode, and the required acceleration is minimal. The

consistency of this reciprocal relationship across all six systems confirms that iterative SPIB systematically improves RC quality and kinetic accuracy, and demonstrates that γ is, in fact, a robust metric for evaluating RC performance.

3.4.3 Effect of bias deposition pace on the RC

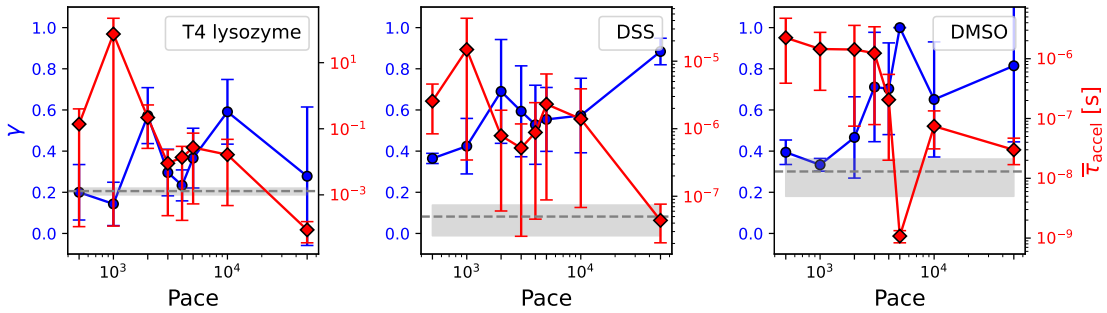

Figure 3.4: EATR γ metric (blue circle, left axis) and mean accelerated time $\langle\tau\rangle$ (red diamond, right axis, log scale) as functions of bias-deposition pace for three protein–ligand systems. (a) T4 Lysozyme–Benzene (b) FKBP–DSS (c) FKBP–DMSO. The gray dashed line represents the experimental residence time, and the shaded gray band indicates the error margin. The blue error bar represents the standard deviation of γ. The red error bars represent the error of $\overline{\tau}_{\text{accel}}$, computed using the bootstrapping method with a 95% confidence interval. In earlier figures, standard MetaD simulations correspond to a pace of 1000 (2 ps) and iMetaD simulations correspond to a pace of 5000 (10 ps).

As an extension to the previous analysis, we examined whether the EATR metric γ depends on the frequency with which Gaussian hills are deposited when the RC is held fixed. This is of immense practical utility as being able to add bias more frequently while still being able to recover accurate kinetics could have a non-linear effect on computational efficiency. We selected the optimal RC for each system, i.e., the one associated with the minimum average accelerated time $\overline{\tau}_{\text{accel}}$ over all SPIB iterations, and carried out short infrequent-metadynamics simulations across a range of bias deposition paces.

Fig. 3.4 presents γ as a function of deposition pace for three representative systems: T4 Lysozyme–benzene, FKBP–DSS, and FKBP–DMSO. The prefactor k_0 was kept fixed so that the effect of bias frequency on computed γ could be isolated. As $\overline{\tau}_{\text{accel}}$ approaches the experimental

residence time (gray dashed line), γ converges toward 1. Increasing the deposition pace up to roughly 5000 steps (i.e., decreasing the frequency of bias deposition) tends to produce a monotonic rise in γ; beyond this threshold, γ no longer improves across all three systems.

From these findings, two conclusions follow: (i) very frequent bias deposition (pace $\leqslant$1000=2 ps) drives γ toward 0, leading to a suboptimal RC for kinetic studies (ii) Within a practical pace range ($\approx$ 5000–50000), γ remains largely invariant to the deposition frequency, underscoring its robustness as a metric for assessing RC quality. These results are consistent with the trends observed in Ref. [32].

3.5 Conclusion

In this study, we performed a systematic evaluation of multiple iterations of an approximate RC derived from the SPIB framework for obtaining unbiased kinetics using iMetaD using the γ metric of EATR. Two questions motivated the work: (i) does iterative SPIB training measurably improve RCs, and (ii) can the γ metric from the EATR approach provide a good score for the evaluation? Across six different protein–ligand benchmarks, we observed an increase in γ along with a corresponding decrease in the mean accelerated time $\overline{\tau}_{\text{accel}}$ in each SPIB iteration. Importantly, we found that the quality of an RC as measured by γ was not strongly dependent on the pace of bias deposition, demonstrating robustness. Our reported approximate formula for CV quality (Eq. 3.2) is a good heuristic for whether the chosen CVs are working well to predict kinetics, but requires some knowledge of k_0. Moreover, if successive iterations of an RC do not show an inverse relationship between γ and $\overline{\tau}_{\text{accel}}$, then the RC should be treated with caution. Examining these two complementary metrics side by side thus provides a straightforward, orthogonal strategy for evaluating RC quality.

As outlined in Sec. 3.3.2, the data evaluated in this work were generated using a standard bias deposition frequency for earlier iterations, and a low bias deposition frequency for only the final

RC obtained through SPIB training. Ideally, if computational resources are available, γ should be evaluated under after executing additional iMetaD simulations in each round, but this will not be practical in many realistic studies. We feel that skipping this step will not significantly alter the conclusions obtained when testing whether the CV is improving. As shown in previous work [32] and also assessed in Fig. 3.4, the EATR framework is robust for both iMetaD and regular MetaD, with the biasing pace playing only a minor role. The decisive factor for accuracy remains the choice of RC.

In general, it should be possible to directly apply EATR to iMetaD data to obtain both a good estimate of γ and k_0. While this works well for some cases, in some systems a low value of γ is obtained by this approach even though the simpler $\overline{\tau}_{\text{accel}}$ is relatively close to the experimental value. This may simply reflect a limitation of EATR in the case where limited data is available for prediction (i.e. tens of simulations to fit a rate). We are now developing an alternative approach based on the EATR formalism for obtaining even more reliable estimates of rate constants and γ through a combination of data from OPES-MetaD applied with multiple biasing strengths [30], which may alleviate this problem, but that requires further investigation.

Overall, our framework is broadly applicable. First, beyond its use as a scoring metric, γ can be monitored during SPIB training to adaptively optimize the RC and to serve as a stopping criterion, to avoid unnecessary computation. Second, when a full fit of the γ metric is impractical, the predictive equations we derived offer a rapid diagnostic for assessing relative RC quality. Finally, the approach extends beyond protein–ligand kinetics to systems undergoing folding–unfolding transitions and even intrinsically disordered proteins, where suitable reference coordinates are elusive.

3.6 Data availability

A Python implementation for different γ evaluation methods is available on GitHub repository: https://github.com/hocky-research-group/SPIB_EATR_paper_data.

3.7 Supplementary Information

3.7.1 Unbiased rates used in calculating γ

System	Transition Time (s)
FKBP-DMSO	13×10^{-9}
FKBP-DSS	50×10^{-9}
T4 Lysozyme L99A-Benzene	1.3×10^{-3}
Abl Kinase WT-Imatinib	1200
Abl Kinase N368S-Imatinib	400
Abl Kinase L364I-Imatinib	700

Table 3.1: The fixed values of the experimental residence times used in EATR to determine γ during the SPIB iterations. Values taken from Ref. [33].

3.7.2 Fitting k_0 in EATR

We can perform a global fit for the unbiased rate k_0 in addition to γ by calculating the sum of square errors for each (k_0, γ) pair, then taking the minimum. This gives the results in Fig. 3.5.

3.7.3 Derivation of a relation between γ and $\overline{\tau}_{\text{accel}}$

Considering the observed relationship between the average accelerated time and γ, we provide an argument for why this relationship should exist. We first connect the average accelerated time $\overline{\tau}_{\text{accel}}$ to the average biased simulation time $\overline{\tau}'$:

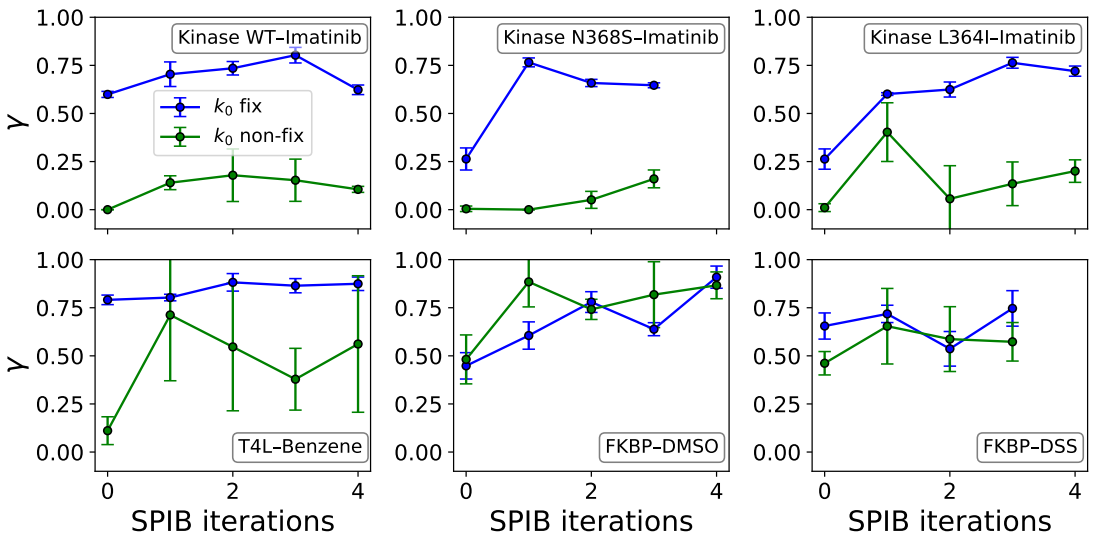

Figure 3.5: Comparison of fitted γ values with fixed k_0 versus free k_0. Each subplot shows γ plotted against SPIB iteration for six different protein–ligand systems. Error bars denote the standard deviation of the fitting uncertainty in γ. The blue plot represents γ value in which k_0 was held fixed at its experimental value (as in Fig. 3.2), while the green plot represents γ value in which both k_0 and γ were treated as free parameters.

$$\overline{\tau}_{\text{accel}} = \overline{\tau'\alpha} = \overline{\tau}'\overline{\alpha} + \text{cov}(\tau', \alpha) \,,$$

$$\overline{\alpha}\,\overline{\tau}' = \overline{\tau}_{\text{accel}} - \text{cov}(\tau', \alpha) \,, \tag{3.3}$$

where $\text{cov}(X, Y)$ is the covariance between X and Y.

We now wish to demonstrate a connection between $\overline{\tau}_{\text{accel}}$ and γ. Using the relationship between the biased rate and the biased simulation time, $k = 1/\overline{\tau}'$, and applying Eq. 2.19, we can rewrite the left hand side of Eq. 3.3:

$$\overline{\alpha}\,\overline{\tau}' = \overline{\alpha}/(k_0\overline{e^{\beta\gamma V}}) \tag{3.4}$$

To solve for γ, we note that we can approximately take γ outside of the ensemble average for large values of $e^{\beta V}$ (when $0 < \gamma < 1$):

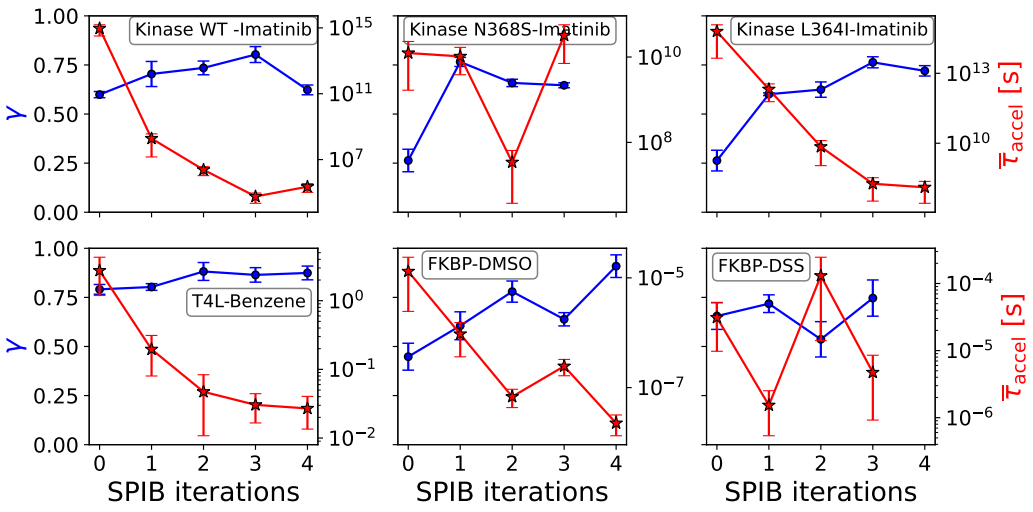

Figure 3.6: Same data as Fig. 3.3 represented in a different way. Here, results of the EATR metric γ (blue circle, left axis) and the mean accelerated time $\overline{\tau}_{\text{accel}}$ (red star, right axis, log scale) over SPIB iterations for six protein–ligand systems with k_0 fixed to the experimental value. Top left to bottom right: ABL kinase Wild Type (WT) - Imatinib, ABL kinase mutant N368S - Imatinib, ABL kinase mutant L364I - Imatinib, T4 Lysozyme - Benzene, FKBP - DMSO, FKBP - DSS. Error bars denote standard deviations of fitting uncertainty in γ and standard deviations in $\overline{\tau}_{\text{accel}}$.

$$\frac{k_0 \overline{e^{\beta\gamma V}}}{\overline{\alpha}} \approx \frac{k_0 \overline{\alpha}^{\gamma}}{\overline{\alpha}} = k_0\, \overline{\alpha}^{\,\gamma-1}, \tag{3.5}$$

because $\frac{\overline{e^{\beta\gamma V}} - \overline{e^{\beta V}}^{\gamma}}{\overline{e^{\beta V}}} \approx 0$.

Combining Eqs. 3.4 and 3.5 we get,

$$\Rightarrow \overline{\alpha}\,\overline{\tau}' \approx \overline{\alpha}^{\,1-\gamma} k_0^{-1}. \tag{3.6}$$

Equating Eqs. 3.3 and 3.6 gives us

$$k_0\,(\overline{\tau}_{\text{accel}} - \text{cov}(\tau', \alpha)) \approx \overline{\alpha}^{\,1-\gamma}$$

and taking the logarithm on both sides and solving for γ, we obtain the following two equations presented in Sec. 3.4.2:

$$\gamma \approx 1 - \frac{\log\left(k_0(\overline{\tau}_{\text{accel}} - \text{cov}(\tau', \alpha))\right)}{\log(\overline{\alpha})} \tag{3.7a}$$

$$\gamma \approx 1 - \frac{\log\left(k_0\,\overline{\tau}_{\text{accel}}\right)}{\log(\overline{\alpha})}. \tag{3.7b}$$

4 | Applying the Method to Time-Independent Biases

This chapter is adapted from the preprint version of "Stepping up enhanced rate calculations with EATR-flooding"[116]

Mazzaferro, Peña Ccoa, Cossio, Hocky, submitted to *J. Chem. Theory Comput.*

arXiv:2604.27136

4.1 Abstract

Several recent methods have shown that it is possible to compute rate constants of very slow biomolecular processes using simulations where a time-dependent bias is added along one or several collective variables (CVs). We previously reported the exponential average time-dependent rate (EATR) method, which can improve upon these approaches by accounting for how efficiently the external biasing potential modifies the observed rate using a learned CV-quality factor γ. This results in more accurate rate estimates using the same data when biasing a suboptimal coordinate. However, as formulated EATR depended on the biasing potential varying over time to properly determine the biasing efficiency, which limits the method's applicability to quasi-static biasing schemes such as "flooding" or on-the-fly probability enhanced sampling (OPES). Here, we present the EATR-flooding approach, which generalizes our method by replacing the need

for a time-dependent bias by instead varying (stepping up) the strength of the biasing potential across multiple sets of simulations. We implement this approach as an open-source Python library, and demonstrate that this approach is accurate without substantial loss of efficiency compared to standard EATR for a coarse-grained protein system, and also show good performance on a fully atomistic cavity-ligand model. Two additional appealing features of EATR-flooding are an internal check for over-biasing and the fact that only a single γ parameter is predicted for a given choice of CVs, as compared to our earlier results where γ empirically depended on biasing rate. Finally, we believe EATR-flooding applies not only to OPES simulations but more generally to CV biasing enhanced sampling approaches, making it broadly useful.

4.2 Introduction

The development of computational approaches to predict protein-ligand binding kinetics is an active area of research, due to the growing recognition of the importance of drug lifetime in controlling the efficacy and durability of pharmacological treatments [117–122]. Classical molecular dynamics (MD) simulations provide a powerful approach to predict the behavior of biomolecules in solution as well as their interaction with ligands or other biomolecules [1, 123]. Despite the name, MD simulations are more often used to estimate short timescale properties of a system by integrating the equations of motion to produce statistical samples of the system at fixed temperature or fixed temperature and pressure (distorting the microscopic kinetics), rather than long-timescale dynamical properties [36]. For many biophysical problems of interest, however, the dynamics of the most relevant configurational changes are dominated by transitions over free energy barriers between metastable states, meaning the kinetics can still be computed by statistical approaches.

Therefore, atomistic MD simulations can still provide a powerful technique for predicting the time scale or average rate constant for transitions between these states, as long as sufficient

transitions can be observed within accessible computational time. Unfortunately, there is often a substantial difference in timescales between what can be accessed in MD simulations as compared to the time scales of the properties of interest. Modern simulations often generate microseconds of trajectory data for relatively simple systems, whereas the processes of biomolecular interest occur on timescales ranging from milliseconds to seconds. The transitions of interest are so-called rare events, in the sense that the system spends long periods in a metastable state before crossing into a new state in a relatively fast manner. Because of this undesired tendency to remain in a single region of state space for long periods of time, a number of enhanced sampling methods have been developed to quickly drive systems to escape free energy minima by adding an extra bias along certain collective coordinates or variables (CVs) describing the system [15–17, 35, 36, 38, 42]. This extra bias effectively lowers the barrier to transitions, allowing one to reconstruct the underlying free energy landscape through inverting the effect of the bias on the frequency of observed configurations. The two such methods that we will focus on in this work are Metadynamics (MetaD) and On-the-fly probability enhanced sampling (OPES), which both work by constructing a bias on the fly which is tailored to push the system away from already explored regions of CV space [15–17, 41].

In both approaches, the consequence of adding this external bias is to significantly and intentionally reduce the long dwell times before the transitions. However, it is not always clear that rate constants of interest can be recovered under these biases in the same way as thermodynamic properties. Several methods have been developed to recover kinetics from CV biasing techniques, relying on using a thermodynamic model for the kinetic process of interest, such as Kramers theory [26, 28–30, 57, 58]. In essence, doing so requires assuming that the bias is applied only to destabilize the starting state and does not directly affect the transition over the barrier, in which case transition times are sped up approximately with the average of the exponential of the applied bias. For example, in Infrequent Metadynamics (iMetaD) [29], this is achieved by only updating the biasing potential on a time scale longer than the barrier crossing time. In OPES-flooding, this

is achieved by capping the highest level of bias applied to the system to fall below a previously estimated barrier level, ΔE [30]. All of these methods also have an implicit assumption that the bias is applied to a CV that is a "good" reaction coordinate for the process of interest, since bias applied perpendicular to the ideal reaction coordinate should not be included in estimating such a speedup factor. Despite these stringent requirements, iMetaD and OPES-flooding shine in the context of giving a computationally tractable way to estimate kinetics of very slow biomolecular processes (milliseconds to days) using only nanoseconds or microseconds of total simulation [57].

In our work, we have concentrated on relaxing the requirement to choose a good CV, since in practice this is challenging to do *a priori* - or even *a posteriori*. We have pursued a strategy whereby we extract the rate constant from the data along with an additional parameter γ that takes into account the efficiency of the chosen CV for producing desired transitions [31, 32]. In our Exponential Average Time-dependent Rate (EATR) approach—described in detail in the next section—we estimate both the rate and γ by averaging our estimate of the speed up factor across many MetaD simulations, and then fitting transition times either via maximum likelihood or by fitting the cumulative distribution of observed transition times.

While EATR works well for simulations biased with MetaD, it failed when applied to simulations biased with OPES [32]. This failure was due to the rapid convergence of OPES to a final biasing function in the starting basin, which prevents our fitting approach from disambiguating the unbiased rate from the estimate of CV quality, as explained below. To solve this, we propose a new method called EATR-flooding which estimates γ using the biased rates from multiple sets of simulations, each with a different average amount of bias applied before transitions are observed; in OPES this is achieved by varying the value of ΔE. By "stepping up" the value of this parameter, which limits the average amount of bias that is added, we are able to probe the dependence of the observed rate on the amount of bias added, which can allow us to extract information about the quality of the biased CV(s) with minimal extra cost. We first speculated about this approach in the supplementary information of our earlier study [32] and showed how by varying the amount

of bias, we could, with an uncontrolled approximation, extract an approximate rate and γ. In this chapter, we put this method on more rigorous grounds. We then test our new method on a coarse-grained model for protein folding as well as a fully atomistic model of a ligand binding problem, and show that it provides efficient and accurate estimates for the unbiased rate, using a single γ parameter to account for the quality of the biased CV. We also show that the method generalizes beyond OPES, working for the case of a fixed external bias as originally envisioned in the conformational flooding or hyperdynamics methods [26, 27], and also for a time-dependent bias applied with iMetaD.

The rest of this chapter is organized as follows. In Sec. 4.3, we explain the necessity for and the theory behind the EATR-flooding method. In Sec. 4.4, we present the results of applying the EATR-flooding method to a coarse-grained protein and model cavity-ligand system. In Sec. 4.5, we present our conclusions from the results. Finally, in Sec. 4.6, we provide details of our computational methods.

4.3 EATR-flooding (EATRf)

If the biasing potential is constant in time, the simulation average in Eq. 2.20 is approximately the ensemble average:

$$\overline{e^{\beta\gamma V(\xi(t))}} \approx \left\langle e^{\beta\gamma V(\xi)} \right\rangle \equiv \alpha_\gamma \,, \tag{4.1}$$

which is constant. In this case, the transition becomes Poissonian:

$$S(t) = e^{-k_0 \int_0^t \left\langle e^{\beta\gamma V(\xi)} \right\rangle dt'} = e^{-k_0 \alpha_\gamma t} \,, \tag{4.2}$$

and there is not a unique combination of k_0 and γ consistent with the data, since any decrease to α_γ can be compensated for by an increase in k_0. As such, EATR as previously formulated struggles to properly estimate k_0 and γ when employing a static (or quasi-static) biasing potential, such as

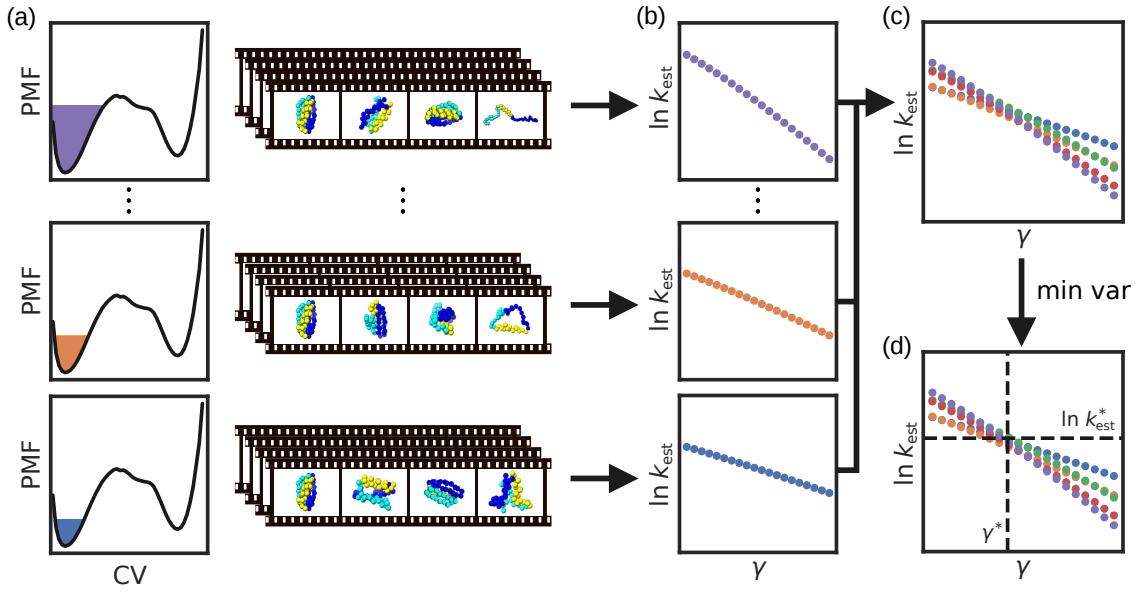

Figure 4.1: A schematic diagram for the EATR-flooding method. (a) Multiple sets of simulations are run until a transition occurs, where each set is run with a different average amount of applied bias, controlled by ΔE in OPES. Sets are color-coded for clarity. (b) The estimated log-rate for the transition is computed as a function of γ, $\ln k_{\text{est}} = \ln k_{\text{obs}} - \ln \left\langle e^{\beta\gamma V(\xi)} \right\rangle$. (c) The rate estimates are combined for all sets, and (d) the log-rate and γ which minimize the variance over the sets are taken to be the EATR-flooding estimates.

that used in OPES.

In order to obtain estimates for k_0 and γ using OPES, we need a formalism which does not assume that the biasing potential is time-dependent. We can do this by first applying Eq. 4.1 to Eq. 2.20, taking the logarithm on both sides, and solving for the logarithm of the estimated rate constant, yielding

$$\ln k_{\text{est}}(\gamma) \equiv \ln k_{\text{obs}} - \ln \left\langle e^{\beta\gamma V(\xi)} \right\rangle , \tag{4.3}$$

where k_{est} is the estimate for k_0, and k_{obs} is the constant value of the biased rate $k(t)$. Note that this rate estimate depends on γ.

Given that there is only one true value of $\ln k_0$, we hypothesize that there exists a corresponding unique value for γ^* which produces only one value of $\ln k_0^*$. Therefore, we propose to

determine γ^* and k_0^* by comparing the results of multiple *sets* of simulations, each with a different amount of bias, as shown in Fig. 4.1a, and computing the effect of this different amount of bias on the observed rate constant. The unbiased rate for each set is estimated as a function of γ using Eq. 4.3, then the unbiased rate and γ estimates for the full collection of simulation sets are obtained by finding the value of γ where all sets of simulations have approximately the same unbiased rate estimate, or equivalently,

$$\gamma^* = \text{argmin}_\gamma \text{Var}(\ln k_{\text{est}}(\gamma)) \,, \tag{4.4}$$

when the variance of the estimates across the sets is minimized. The best estimated value of the true rate is then $k_0^* = \overline{k_{\text{est}}(\gamma^*)}$, where the average is done over the simulation sets. This is illustrated in Fig. 4.1b-d.

We note that this approach does not require that the bias potential is static or quasi-static, but only that it is possible to converge an acceleration factor α_γ for each set of biasing parameters. Our EATRf method is therefore based on the same ansatz as the original EATR approach but is more general.

4.4 Results and Discussion

4.4.1 Coarse-grained model for protein folding

To evaluate our hypothesis in the preceding section, we first test our approach on a system sufficiently complicated that we can bias a wide range of CVs to obtain kinetic information, but sufficiently simple that we can obtain an unbiased rate constant estimate. We therefore adopted the same system used in our prior development of EATR, a Gō-like model of the B1 domain of protein G, which had previously been established as a useful system for studying the quality of CVs for describing protein folding.[32, 70, 79]. As described in the Methods, we ran MD simula-

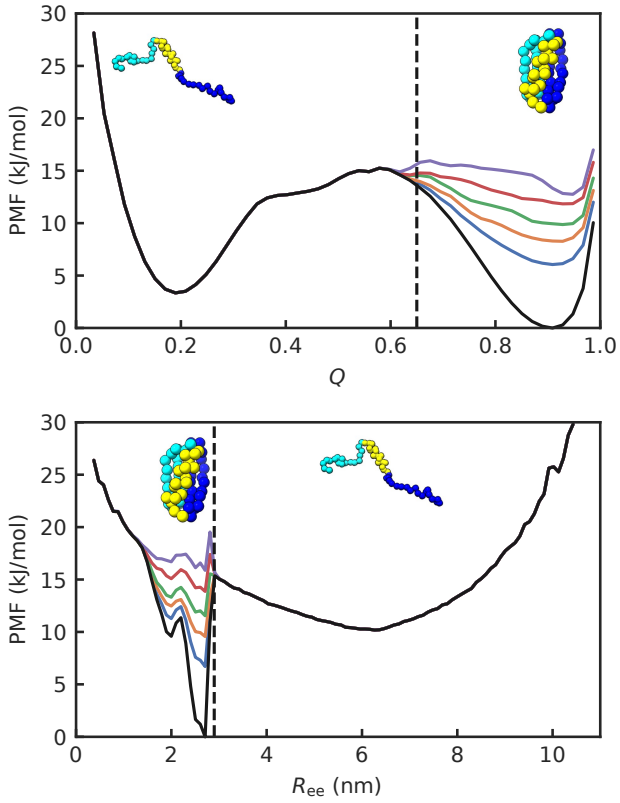

Figure 4.2: The potential of mean force (PMF) along the fraction of native contacts CV Q (top) and the end-to-end distance CV R_{ee} (bottom). The colored lines indicate the PMF with the average final biasing potential from OPES with various values of the ΔE parameter. From bottom to top, ΔE = 5, 7, 9, 11, and 13 kJ/mol. The vertical dashed lines represent the boundary of the excluded region.

tions where an OPES bias was applied to one of the following CVs: the fraction of native contacts Q, the end-to-end distance R_{ee}, and the radius of gyration R_g. We also performed simulations in which we biased R_{ee} and R_g simultaneously.

For each choice of CV, we ran five sets of simulations where we varied the maximum level of OPES bias, ΔE. This resulted in a different biasing potential in each set. As an example of how this looks when projected along two of the CVs, in Fig. 4.2 we plot the average final bias obtained from OPES on top of the underlying PMF computed in our previous work [32]. As can be seen, the amount of bias/forcing increases in concert with the ΔE parameter, and depends on the choice of CV. In Fig. 4.8, we show that the effective average bias quickly plateaus to a static

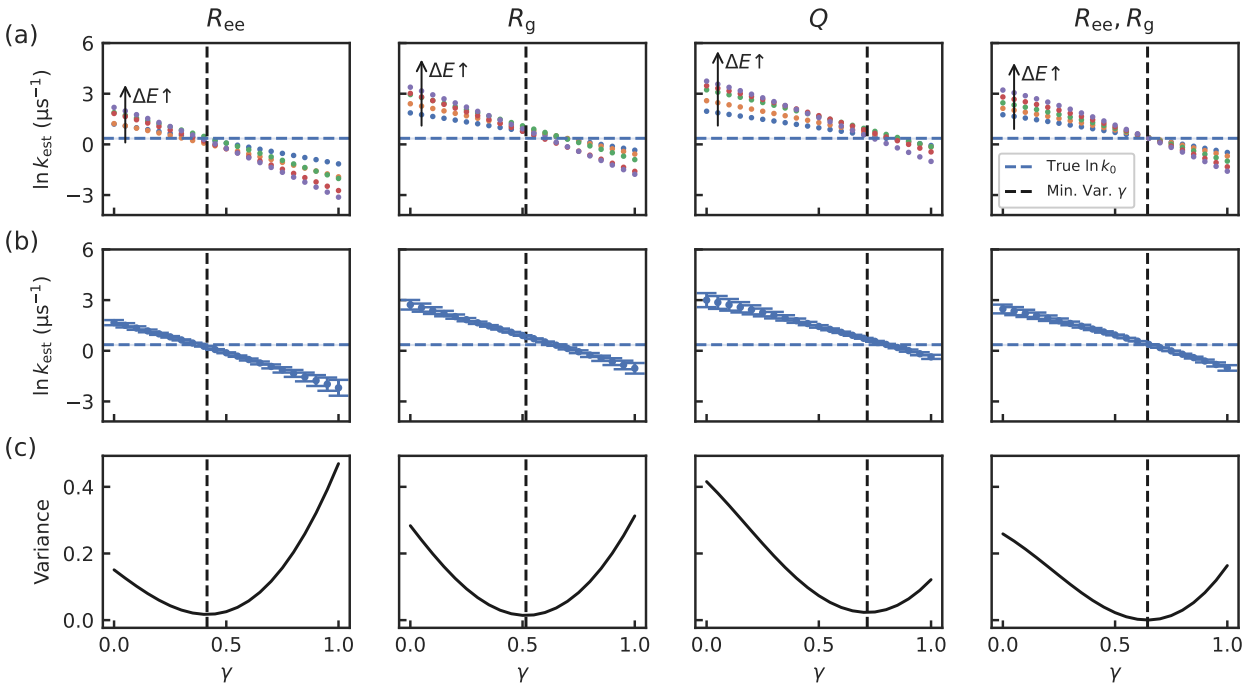

Figure 4.3: The results of EATR-flooding on the protein G model. (a) The values of $\ln k_0$ predicted by Eq. 4.3 for each set of OPES simulations with each ΔE parameter for various values of γ. (b) The value of $\ln k_0$ averaged over the sets of simulations for each potential value of γ. The variance of $\ln k_0$ across the sets is represented by error bars. (c) The variance in the predicted value of $\ln k_0$ across the sets of simulations. The horizontal dashed lines represent the true value of $\ln k_0$, the uncertainty in which is too small to represent in this plot. The vertical dashed lines represent the minimum variance value of γ, which is the value which causes all sets to give the same value of $\ln k_0$.

value, showing why we are not able to apply our original EATR method.

We therefore proceed to apply Eq. 4.3 to these data while varying γ over the range 0–1, which gives us the results in Fig. 4.3a. In this Figure, the horizontal dashed line shows the log of the unbiased rate computed previously [32]. With these data, we can directly see that the curves for each value of ΔE cross near the true rate. The crossover behavior can be explained as follows: for $\gamma = 0$, simulations with high ΔE are on top, which is simply a reflection of having faster observed transition times when the system is biased more strongly; in contrast, $\gamma = 1$ is equivalent to iMetaD which over-predicts transition times for non-ideal CVs [32, 83] and consequently larger over-predictions (lower rate constants) arise for higher ΔE.

In Fig. 4.3b,c, we show the average and variance across data sets obtained from the values computed via Eq. 4.3. With a vertical dashed line, we indicate the γ which has the lowest variance, as computed using an optimization algorithm as described in the methods. Visually, we can observe that the minimum variance occurs close to when the estimated log rate value intersects that of the unbiased rate estimate.

Inspecting the minimum-variance γ values, we find that they are similar to those computed using our original EATR approach using data from iMetaD [32]. As expected, Q has the highest value of γ while R_{ee} has the lowest. Simultaneously biasing the two lower quality CVs, which individually have γ values of 0.41 and 0.52, results in a combined γ value of 0.64, showing the synergistic value of multidimensional biasing. Of note, the γ value for Q is somewhat less for these OPES simulations than we observed with iMetaD, where γ could approach 1 for slow biasing along Q, which may reflect an interesting difference between how OPES and iMetaD build up the bias histogram, which is worthy of further investigation in the future.

In addition to this γ analysis, we can also see the quality of the CVs visually by plotting the log of the observed rate versus the log acceleration factor, $\ln\langle e^{\beta V}\rangle$, as shown in Fig. 4.9. Since the OPESf acceleration factor is accurate for good CVs such as Q, we expect a linear relationship with near unit slope between these two terms. As discussed in Ref. [32], lower slopes indicate a bad or inefficient CV. This is discussed more in the next section.

Having demonstrated that EATR-flooding is an accurate approach, we then wanted to investigate whether the additional cost of running with multiple values of ΔE made the method prohibitively expensive compared to OPESf. To do so, we recomputed the predicted rates by EATRf using subsampling with smaller numbers of simulations per batch. As shown in Fig. 4.4, the typical predicted rate constant is unchanged even when using only 10 simulations per batch, corresponding to 50x less simulation time than used for the full analysis. However, the typical error in doing so is likely to be larger, and so there is some benefit to obtaining more data. From this analysis, we directly observe that the OPESf estimator is systematically incorrect for poor

Table 4.1: The rates and γ values obtained for the protein G model by applying EATR-flooding to different CVs, compared to the rates obtained from conventional OPES-flooding, combining data from all sets with the ΔE values 5, 7, 9, 11, and 13 kJ/mol. The 'unbiased' entry is obtained from running 100 equilibrium MD simulations, from Ref. [32].

OPES-flooding		
CV	$\ln k_0^{\text{est}}$ (μs^{-1})	
R_{ee}	-1.92 ± 0.07	
R_{g}	-0.82 ± 0.05	
R_{g}, R_{ee}	-0.63 ± 0.05	
Q	0.02 ± 0.05	
EATR-flooding (OPES)		
CV	$\ln k_0^{\text{est}}$ (μs^{-1})	γ
R_{ee}	0.23 ± 0.23	0.41 ± 0.06
R_{g}	0.85 ± 0.16	0.52 ± 0.04
R_{g}, R_{ee}	0.42 ± 0.17	0.64 ± 0.04
Q	0.62 ± 0.19	0.74 ± 0.06
EATR-flooding (fixed bias)		
CV	$\ln k_0^{\text{est}}$ (μs^{-1})	γ
R_{ee}	0.40 ± 0.18	0.31 ± 0.07
Q	0.53 ± 0.16	0.88 ± 0.07
EATR-flooding (iMetaD)		
CV	$\ln k_0^{\text{est}}$ (μs^{-1})	γ
R_{ee}	0.43 ± 0.06	0.28 ± 0.01
Q	0.47 ± 0.17	0.67 ± 0.06
unbiased:	0.36 ± 0.07	

CVs, and additional sampling does not alleviate the issue. In contrast, the mean predicted value of k_0 from EATRf is generally close to the unbiased value for all CVs, always within a factor of 2, and seems to have near perfect accuracy for 2d biasing of R_{g} and R_{ee} (see also Tab. 4.1).

As just discussed, Fig. 4.4 shows that our approach of combining data from multiple sets of simulations biased by different amounts is successful and efficient. In that analysis, we always combined results from five different values of ΔE, and showed that accurate predictions of the rate could be obtained even with only a small number of simulations per set. We also wanted to investigate what the minimum number of different ΔE values we could use is, and also whether

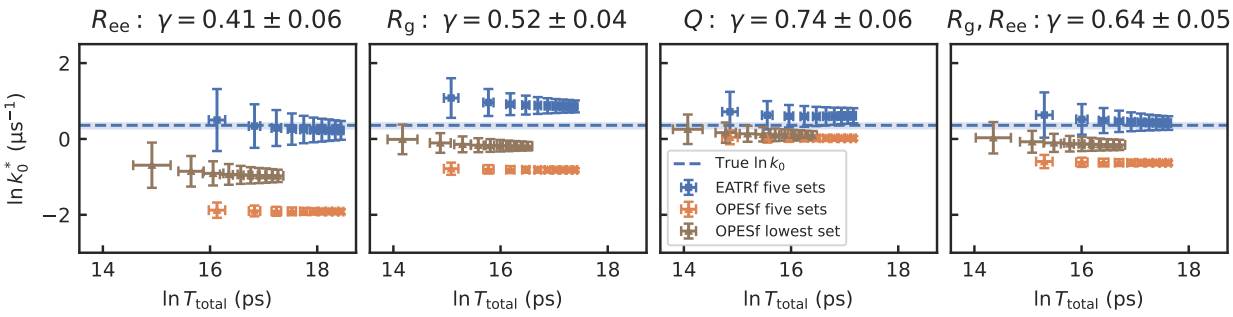

Figure 4.4: The predicted value of $\ln k_0$ from EATR-flooding (in blue) and OPES-flooding (in orange) as a function of the logarithm of the total simulation time with the lowest five values for ΔE. Various numbers of simulations per set, from 10 (left side of each plot) to 100 (right side of each plot). The blue dashed line represents the rate obtained from unbiased simulations. Error bars indicate the standard error in bootstrap analysis.

it was required to have the lowest ΔE values (which require the most simulation time to observe transitions). Fig. 4.5 shows the result of this analysis, where for the worst CV, R_{ee}, we systematically eliminate either 0, 1, 2, or 3 of our data sets, while maintaining a constant number of simulations within each data set. As was the case in Fig. 4.4, we see that the estimate from the OPES-flooding estimator is systematically wrong in all cases. Interestingly, as seen in the top panel, adding more data sets with higher ΔE makes the OPESf estimate worse; in the lower panel, we see the first point using only the highest two ΔE is particularly bad, but adding more data at lower barriers slightly improves the estimate. In stark contrast, using the same data the EATRf estimator is relatively accurate when using just the lowest 2 barriers, and becomes systematically better as more data is added. It is also 2 log units better than the OPESf estimate with even just the highest two ΔE sets included, albeit with a large variance.

Given that EATRf was conceived of in the context of quasi-static biasing, we also performed simulations where we added a truly static Gaussian bias rather than using OPES. We performed otherwise identical simulations and analysis, adding a constant external potential to either R_{ee} or Q, described in Sec. 4.8.1.1 in the SI. Fig. 4.10 shows the EATRf analysis for these data, which shows equally or more accurate results as compared to the OPES data, despite the naive way in which the bias was constructed (Tab. 4.2).

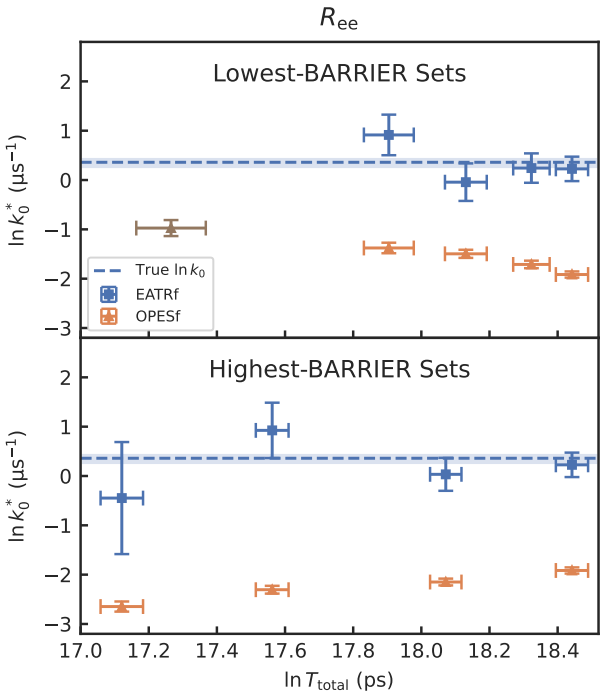

Figure 4.5: The predicted value of $\ln k_0$ from EATR-flooding and OPES-flooding on the R_{ee} CV with the lowest (top) and highest (bottom) values for ΔE with 100 simulations per set. The blue dashed line represents the rate obtained from unbiased simulations. Error bars indicate the standard error in bootstrap analysis.

Finally, to further demonstrate the generality of EATRf, we applied our workflow to data generated by iMetaD for Ref. [32], where the different forcing amounts are controlled by the frequency of hill deposition alone. As described in SI Sec. 4.8.1.2 and seen in Fig. 4.11 and Tab. 4.1, this approach is as accurate as our EATRf results using OPES.

4.4.2 Cavity-Ligand Model

Although EATRf seems promising on a toy model, we want to ensure that it works for a larger and fully atomistic system. A persistent challenge is finding a challenging benchmark system that still admits a known "ground truth" and which also has multiple possible choices of CVs to bias. For this study, we adopted a simple model of receptor-ligand unbinding that was originally developed to study the effect of collective dewetting on ligand dissociation[124].

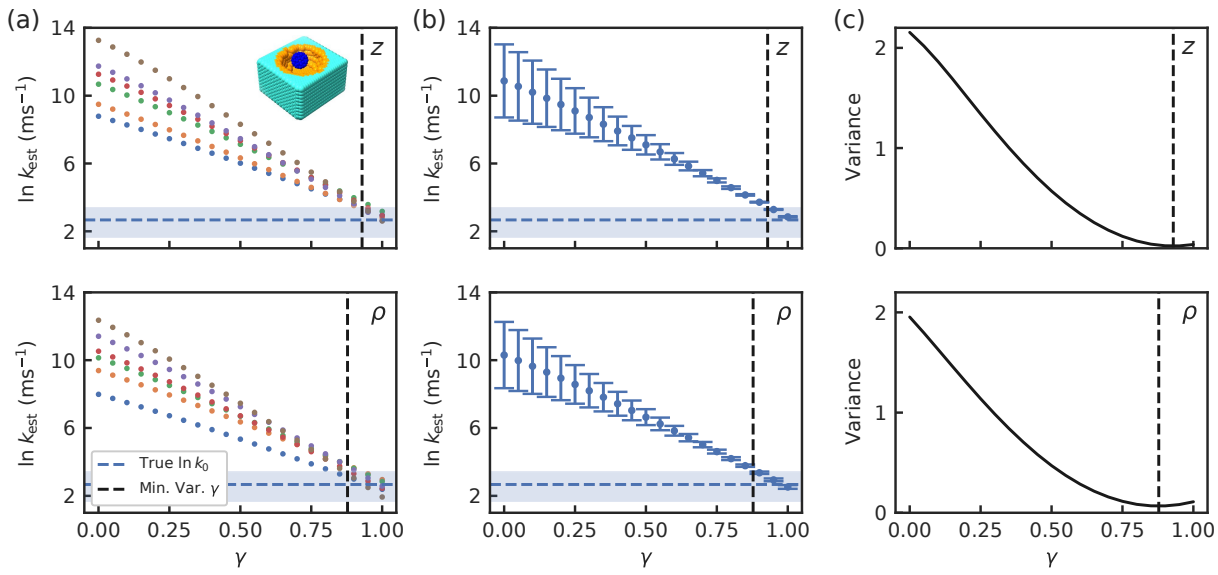

Figure 4.6: The results of EATR-flooding on the weakened cavity-ligand model. (a) The values of $\ln k_0$ predicted by Eq. 4.3 for each set of OPES simulations with each ΔE for various values of γ. (b) The value of $\ln k_0$ averaged over the sets for each value of γ. The variance of $\ln k_0$ across the sets is represented by error bars. (c) The variance in the predicted value of $\ln k_0$ across the sets of simulations. The horizontal dashed lines represent the true value of $\ln k_0$ with the 95% credible interval. The vertical dashed lines represent the minimum variance value of γ.

This system consists of a buckyball in a hydrophobic cavity, and it is known that the typical unbinding reaction occurs with the buckyball sliding sideways in the pocket before escaping, which allows the pocket to become hydrated, rather than coming directly out which creates a temporary vacuum [125]. We have previously used this model to study the effect of mechanical pulling forces on the unbinding rates using iMetaD [9]. However, we do not know the "true" unbinding rate, and estimates from an alternative approach, Markov state modeling, gives an estimate orders of magnitude faster than that from iMetaD [126]. For this study, we therefore decided to modify the parameters of this model to produce a series of cavity-ligand models with lower binding affinities, from which we can select a system that displays the same qualitative behavior but for which unbinding occurs fast enough to allow us to estimate an unbiased rate.

As detailed in the SI, we reduced the ϵ parameters of the non-bonded interactions between the receptor and ligand atoms. Three different models with decreasing attractive interactions were evaluated. We computed the unbinding PMF using the MetaD procedure from our previous study[9], with results shown in Fig. 4.12. As ϵ is lowered slightly (SI Tables 4.2 and 4.3) we see a systematic decrease in the barrier to unbinding. The third model (ϵ_3) tested has a barrier approximately 10 kcal/mol lower and we proceeded to assess whether we could approximate its unbinding rate constant with unbiased MD[9]. At the same time, as shown in Fig. 4.12a,b, the underlying free energy landscape has the same shape as in the original model, meaning we still expect unbinding to follow the same type of nontrivial unbinding pathway.

To estimate an unbiased rate constant for our model, we performed 360 total microseconds of unbiased MD launched from 48 different sampled starting points. This produced only five transitions. Using these data we made a maximum likelihood estimate for the unbiased log-rate of 2.67 (with k_0 in ms^{-1}), with a highest-density 95% credible interval of 1.62 to 3.44 determined according to Sec. 4.6.3.

We then performed sets of 48 OPES simulations each for a series of ΔE values from 3 to 6 kcal/mol. We biased two CVs: the perpendicular (z) and orthogonal ($\rho = \sqrt{x^2 + y^2}$) distance from

the ligand to the center of the cavity. For larger and more complex systems like this, we first examined the change in the logarithm of the rate constant with the logarithm of the average acceleration factor, which is presented in Fig. 4.13. In practice, we have observed that when the system is not over-biased, this curve will be linear with a slope approximately equal to γ (equivalent to taking γ outside the average in Eq. 4.3). For very large ΔE, we observe that the points start to curve down, and we take this as a heuristic of which points to exclude for our EATRf analysis, as these likely correspond to over-biasing regimes.

Taking 6 values of ΔE (3, 3.5, 4, 4.5, 5, and 6 kcal/mol) we performed the EATRf analysis with results shown in Fig. 4.6. When biasing the z-coordinate, we obtained 3.47 ± 0.34 as the estimate for $\ln k_0$ with k_0 in ms^{-1} and $\gamma = 0.93 \pm 0.04$. When biasing the ρ-coordinate, we obtained $\ln k_0 = 3.50 \pm 0.36$ and $\gamma = 0.88 \pm 0.04$. We find that z, the natural unbinding coordinate, has a high γ and a predicted rate constant within the credible intervals of our unbiased estimate. For ρ, we get a slightly lower γ as expected since it doesn't directly drive the unbinding. Using the traditional OPESf analysis from just our lowest barrier data produced $\ln k_0 = 2.89 \pm 0.06$ for z and $\ln k_0 = 2.50 \pm 0.07$ for ρ.

Although we do not have an estimate of the unbiased rate for the original "strong" model, we did perform EATR-flooding analysis as shown in Fig. 4.14 in Sec. 4.8.3 of the SI. In this case, we find that z alone is predicted to be a less ideal CV than in our weaker case ($\gamma = 0.64$) and this is not substantially improved by biasing z and ρ simultaneously (also $\gamma = 0.64$). The predicted unbinding times fall between those for MSM and previous iMetaD studies and are probably underestimates given our predicted difference in the barrier heights of the original and new models.

Returning to our weaker cavity model, because z and ρ were both found to be good CVs for biasing in this system, we decided to define a set of CVs to demonstrate that EATRf still works even if the quality of the CV is worse than that of ρ. We first created a nuisance CV θ that represents the zenith angle of a chosen carbon atom (defined in PLUMED using a dihedral between that atom, the center of the buckyball, the center of the cavity, and an atom in the

cavity that is aligned to the vertical axis), then we created four CVs, each of which is of the form $\xi = a\,z + b\,\theta$, where $a^2 + b^2 = 1$. We used $a = 1.00, 0.95, 0.90$, and 0.87. The results of performing EATRf and OPESf on these simulations are given in Fig. 4.7. Adding weight to θ and removing weight from z in the CV is expected to make the CV worse for biasing, and in Fig. 4.7, we can clearly see the effect of systematically making our biased CV worse, as the slope decreases and deviates from 1, which is the OPESf approximation.

Performing our full EATRf analysis, we find that γ decreases systematically along with the fraction of θ introduced in Fig. 4.7b. At the same time, this leads to a decrease in the OPESf rate estimate, i.e. a severe overestimate of the binding lifetime. In contrast, we were pleased to see that the rate prediction from EATRf is roughly the same even as the CV gets substantially worse.

4.5 Conclusions

In this study, we developed a variant of the EATR method that allows for a more general way of obtaining accurate rate constants when applying biasing to CVs of non-ideal quality, including methods that apply a quasi-static bias where our previous formulation failed. The primary innovation is to run many sets of biased simulations, each with a different amount of bias. These multiple sets of biased data allow us to visualize whether the bias coordinate is efficient. Moreover, we can use the relationship between the amount of bias, the observed rate, and the unbiased rate to estimate the unbiased rate for each set of simulations as a function of the γ parameter. The value of γ that minimizes the variance in these estimates is taken as the γ estimate, and the average of the corresponding rate estimates at that γ is taken as the unbiased rate estimate for the method. We found that we can use multiple sets of simulations to measure the sensitivity of the observed rate on the bias, which determines the γ parameter for EATR.

For the cases we have evaluated, EATR-flooding produces consistent and accurate rate constant estimates when varying the choice of biased CV or CVs. While this requires performing

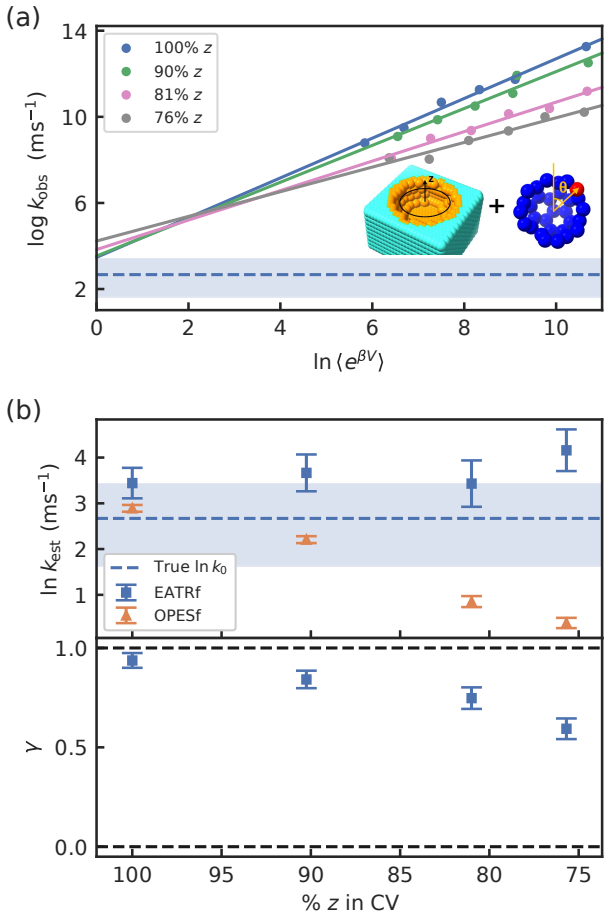

Figure 4.7: (a) The observed value of the log rate in the simulation sets while biasing a CV defined as linear combinations of z and θ, plotted against $\ln \langle e^{\beta V} \rangle$, which acts as a measure of the amount of bias added. For worse CVs that have a larger contribution from θ, adding the same amount of bias has less of an effect on the observed rate. Solid lines are linear fits to the data, with slopes from top to bottom of 0.94, 0.84, 0.75, and 0.59. (b) The predicted value of $\ln k_0$ (top) and γ (bottom) from EATR-flooding and OPESf while biasing each of the CVs. The x-axis is the value of a^2 where a is the coefficient of z in the linear combination. Mean γ values are 0.92, 0.86, 0.68, and 0.57, quite similar to the values of slope in (a), in accordance with the hypothesis in Ref. [32].

multiple sets of simulations, reasonable rates can be obtained with relatively few simulations per set and a small number of sets such that there is only a relatively small increase in cost for a potentially substantial gain in accuracy. It was also shown that when enough bias has been added to reach the overbiasing regime, the rate becomes less sensitive to the further addition of bias, allowing us to diagnose when the system is being truly over-biased in a manner that may prevent

extracting a true rate.

By extending the EATR method to work with OPES simulations, we have not only improved the capability for rapid biasing, but we have also opened up the possibility to gain control over the biasing with OPES, using its inbuilt limitation in maximum bias, ΔE, and optionally the use of an excluded region for applying bias. This will allow us to tackle kinetics computations for challenging biomolecular problems, such as in our previous work evaluating the force-dependence of the unbinding of the protein vinculin from actin, where OPESf seemed to provide more stable estimates of kinetics than iMetaD [127]. We would now like to revisit those calculations to determine whether additional OPES simulations varying ΔE will substantially change our estimates of vinculin binding lifetimes, and whether this approach can allow us to find a better set of CVs for efficient lifetime computations.

As described in the theory section, we are not limited to combining sets of OPES simulations, in principle this should apply to other biasing methods as well, including time-dependent approaches like MetaD. While our preliminary results in Fig. 4.11 show that EATRf can be applied directly to iMetaD data using, as an example, different hill addition rates, further work is required to determine whether optimal choices for iMetaD parameters are required. Moreover, we would like to investigate whether we can combine our two approaches to take advantage of both the increase of bias between sets of simulations and the increase of bias over the course of each simulation to determine the sensitivity of the rate to bias, and thus obtain a more accurate estimate of γ.

4.6 Computational Methods

MD simulations were performed in LAMMPS [81] and GROMACS [128] using standard protocols described in the methods. Enhanced sampling simulations using both MetaD and OPES were performed with the PLUMED open source sampling library [129]. We have previously con-

tributed an example of computing rates of such systems under applied mechanical force in the PLUMED-tutorials resource [114].

4.6.1 MD Simulation details

The molecular dynamics simulations of the Gō-like model of protein G were performed using the 23 Jun 2022 version of the LAMMPS[81] software package patched with PLUMED 2.8.3.[129] The OPES simulation data where the Q and R_{ee} CVs were used for biasing were the same data as used in the Supporting Information of Ref. [32]. The data where the R_g and R_g/R_{ee} CVs were used for biasing were collected using similar scripts. The simulations were run with a 10 fs time step, using the Nosé-Hoover chain thermostat at 312 K with a damping parameter of 1 ps and a chain length of 3. Each simulation was stopped once the protein unfolded, which we defined to occur once the Q CV dropped below 0.35.

The cavity-ligand model was adapted from the model used in Ref. [9], which was the same as in Ref. [126]. The model is composed of a cube of hexagonally close-packed carbon-like atoms with an ellipsoidal cavity and a buckyball (C60) ligand, solvated with TIP4P water molecules. The cavity and ligand atoms have attractive non-bonded interactions. The model used in this project was adjusted to weaken these interactions, as described in the SI.

The MD simulations for the cavity-ligand model were performed using GROMACS 2023.2 [128] patched with PLUMED 2.9.[129]. The simulations were run in a periodic box with a 2 fs time step at 300 K, using the V-Rescale temperature coupling. Each simulation ended once the ligand unbound from the cavity or 500 ns had been simulated. Ligand unbinding was defined to occur once the z-distance increased past 15 Å.

4.6.2 EATR-flooding

The rate estimations were performed on collections of OPES simulation sets, where each simulation set was run with a different value for ΔE, the `BARRIER` parameter in PLUMED. The protein G simulations were run with ΔE ranging from 5 to 13 kJ/mol with 2 kJ/mol increments. The excluded regions for the bias in the protein G simulations for each CV were $Q < 0.65$, $R_g > 1.15$ nm, and $R_{ee} > 2.9$ nm. For the 2D R_g, R_{ee} CV, bias was excluded from the union of the excluded regions for the separate CVs. The cavity-ligand model simulations shown in the main text were run with ΔE ranging from 3 to 6 kcal/mol in 1 kcal/mol increments, plus 3.5 and 4.5 kcal/mol. Those in the SI for the original model from Ref. [125] used ΔE from 13 to 16 kcal/mol in 1 kcal/mol increments, plus 13.5 and 14.5 kcal/mol. No excluded region was used for the cavity models.

While analyzing the protein G simulation sets, the biased rate constants were determined by fitting the final transition times to a Poisson distribution. The same was done for the original cavity-ligand model. However, for the new cavity-ligand model, we instead used maximum likelihood estimation to determine the biased rate constants because some simulations were stopped prematurely, which would influence the observed rate constant.

The ensemble average $\langle e^{\beta\gamma V} \rangle$ for a given value for ΔE was estimated by first averaging the quantity $e^{\beta\gamma V(t)}$ across all running simulations in the set at time t, then averaging the result over time. This ordering was done to be consistent with our earlier EATR approach, but we also tested doing the calculation in the other order with minor quantitative differences. We applied Eq. 4.3 to each simulation set to get a set of $\ln k_0$ estimates as a function of γ, then minimized the variance in the $\ln k_0$ estimates using the bounded version of Brent's algorithm implemented in the SciPy Python package.[67] The minimizing value of γ and the mean of the corresponding $\ln k_0$ values are reported as the γ and $\ln k_0$ estimates in this method.

4.6.3 Error Analysis

All error values, unless otherwise stated, were obtained by bootstrapping [69]. In this method, we calculate the error in some statistic from a number of datasets constructed by randomly sampling the original dataset with replacement. We resampled 1000 times and calculated the relevant quantity for each resampling (usually the log-rate or γ), then took the standard deviation across the resamplings as the error in that quantity. The error in the unbiased unfolding rate of the protein G model is a 95% confidence interval obtained using the bootstrapping method implemented in SciPy.[67]

For the estimation of the unbiased rate for the weakly bound cavity model, we used Bayesian analysis to estimate a 95% credibility interval from our set of simulations where we observed 5 transitions within 360 microseconds. The posterior distribution, $\pi(k_0|N, T_{\text{total}})$, given a number of transitions N and a total simulation time T_{total} was determined following Ref. [130] using a noninformative gamma prior distribution $\pi(k_0) \propto 1/k_0$, leading to

$$\pi(k_0|N, T_{\text{total}}) \propto k_0^{N-1} \exp\left(-k_0 T_{\text{total}}\right) . \tag{4.5}$$

This was then converted to the posterior distribution for $\ln k_0$. The highest-density credible interval is obtained from this posterior distribution by finding two $\ln k_0$-values a and b where $P(a < \ln k_0 < b|N, T_{\text{total}}) = 95\%$ and $\pi(\ln k_0 = a|N, T_{\text{total}}) = \pi(\ln k_0 = b|N, T_{\text{total}})$.

4.7 Data availability

Data and analysis scripts are available from the GitHub repository for the paper adapted in this chapter, https://github.com/hocky-research-group/EATRf-paper-2026. The general EATR analysis framework software is available from https://github.com/hocky-research-group/EATR-rate-analysis.

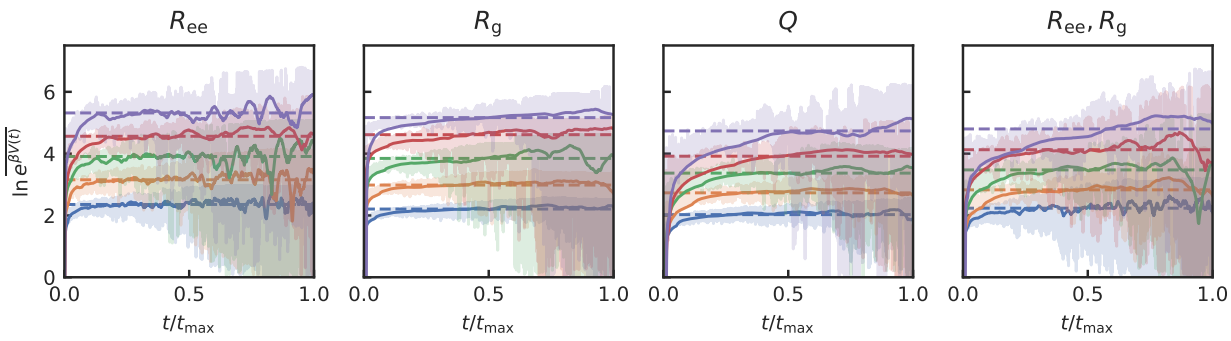

Figure 4.8: The log of the value of $e^{\beta V(t)}$ averaged over all simulations in each set (low opacity). The solid lines represent the log of the time average of $\overline{e^{\beta V(t)}}$ over a window of 25000 steps, and the horizontal dashed lines represent the time average from 0 up to the maximum transition time, which approximates the log of the ensemble average $\ln \langle e^{\beta V} \rangle$ in the folded state.

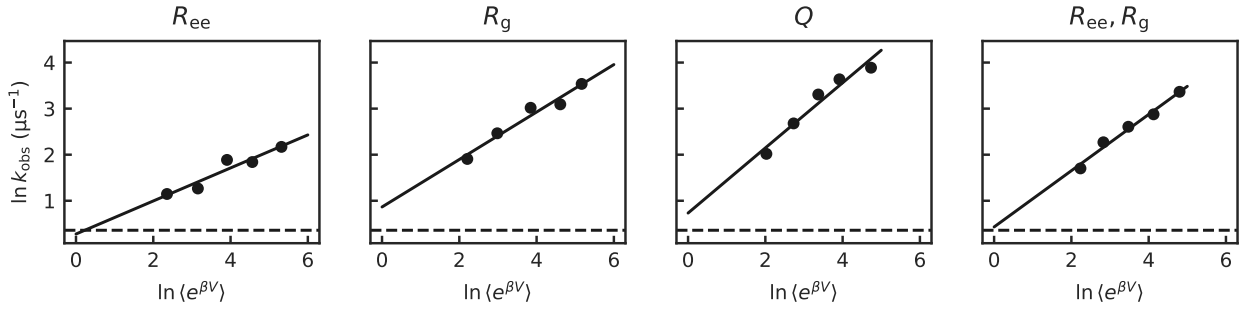

Figure 4.9: The logarithm of the observed rate plotted against $\ln \langle e^{\beta V} \rangle$ in the folded state of the protein G model, estimated as the log of the time average of $\overline{e^{\beta V(t)}}$. The horizontal dashed line represents the value of $\ln k_0$. Increasing $\ln \langle e^{\beta V} \rangle$ for poor CVs has less of an effect on $\ln k_{obs}$ than for good CVs. The EATR-slope approach takes the y-intercept of the best fit line to be the estimate for $\ln k_0$ and the slope to be the estimate for γ.

4.8 Supplementary Information

4.8.1 Further results for protein G

4.8.1.1 Static biasing

Because the EATR-flooding theory requires the bias to be effectively static, we decided to also run the analysis on a collection of simulation sets with exactly static biasing potentials. We ran six sets of simulations each for the Q and R_{ee} CVs of the protein G system, each with nominal

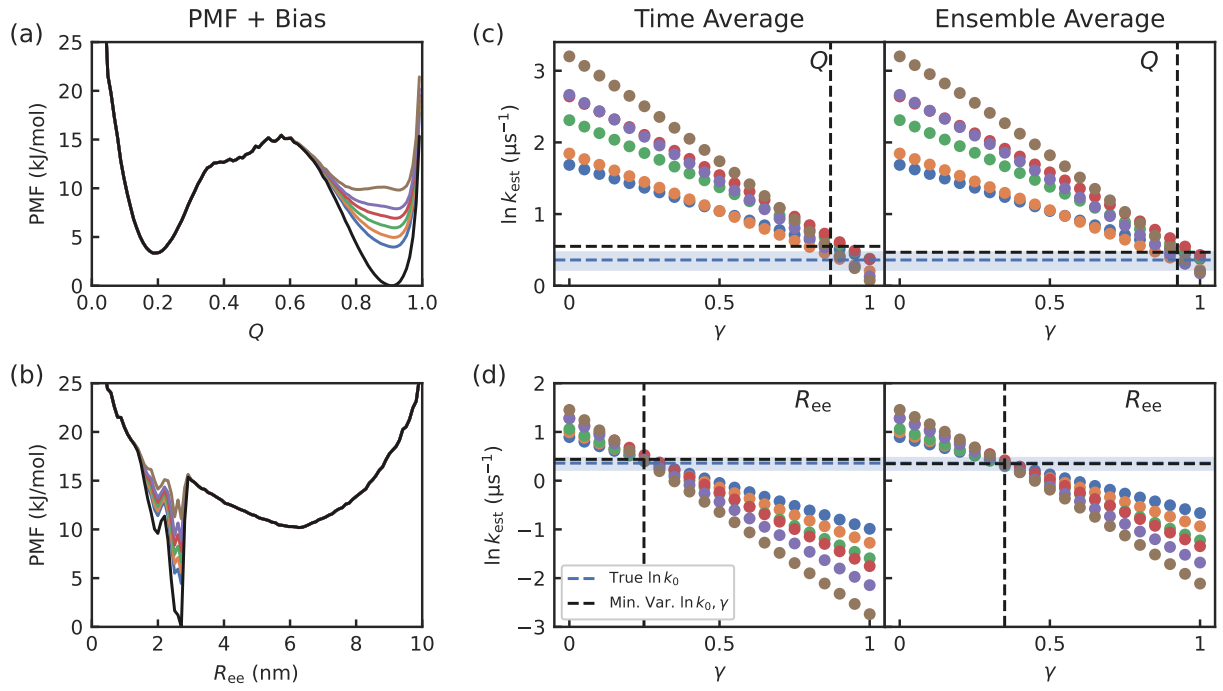

Figure 4.10: (a), (b) The PMFs of the protein G model for Q and R_{ee}. The unbiased PMF is shown in black while the biased PMF for each bias level is colored. (c), (d) The values of $\ln k_0$ predicted by Eq. 4.3 for each set of simulations for various values of γ, by approximating the expected value $\langle e^{\beta\gamma V} \rangle$ using the time average of the simulation average (left) and the directly estimated ensemble average (right).

bias levels of 4, 5, 6, 7, 8, and 10 kJ/mol.

For the Q CV, we applied a bias potential of the form

$$V(Q) = H \exp\left(-(Q-\mu)^2/a^2\right) ,$$

where H is the nominal bias level. We used $\mu = 0.9$ and $a = 0.13$.

Using the same EATR-flooding procedure as in the main text, we obtain $\ln k_0 = 0.53 \pm 0.16$ and $\gamma = 0.88 \pm 0.07$. The biased PMFs and the values of $\ln k_0$ obtained from Eq. 4.3 are given in Fig. 4.10a,c.

Because we have the exact expression for the bias, as well as a good estimate for the potential of mean force (PMF) in the folded state, we decided to also directly calculate the $\ln\langle e^{\beta\gamma V}\rangle$ term in Eq. 4.3 using

$$\langle A\rangle \approx \frac{\sum_i A(\xi_i) e^{-\beta(F(\xi_i)+V(\xi_i))}}{\sum_i e^{-\beta(F(\xi_i)+V(\xi_i))}} ,$$

where ξ_i are evenly spaced values for the CV ξ in the reactant state and $F(\xi_i)$ is the PMF of ξ at ξ_i. Using this instead yields $\ln k_0 = 0.48 \pm 0.14$ and $\gamma = 0.92 \pm 0.06$.

For the R_{ee} CV, we applied a bias potential of the form

$$V(R_{ee}) = H\left(0.9 \exp\left(-(R_{ee}-\mu_1)^2/a_1^2\right) + \exp\left(-(R_{ee}-\mu_2)^2/a_2^2\right) + 0.5 \exp\left(-(R_{ee}-\mu_3)^2/a_3^2\right)\right) ,$$

where we used $\mu_1 = 2.45$ nm, $\mu_2 = 2.68$ nm, $\mu_3 = 1.9$ nm, $a_1 = 0.2$ nm, $a_2 = 0.12$ nm, and $a_3 = 0.3$ nm.

Estimating $\ln\langle e^{\beta\gamma V}\rangle$ as in the main text results in $\ln k_0 = 0.40 \pm 0.18$ and $\gamma = 0.31 \pm 0.07$. Estimating it as above instead results in $\ln k_0 = 0.34 \pm 0.18$ and $\gamma = 0.35 \pm 0.07$. The biased PMFs and the values of $\ln k_0$ obtained from Eq. 4.3 are given in Fig. 4.10b,d.

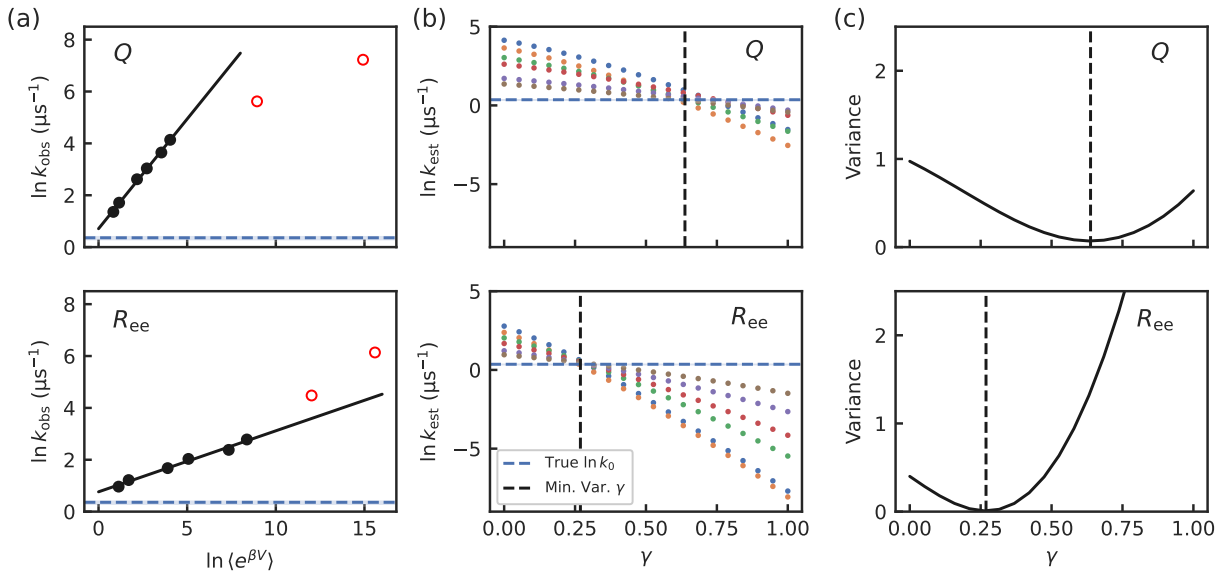

Figure 4.11: (a) The observed rate constant $\ln k_{\text{obs}}$ plotted against the $\ln \left\langle e^{\beta V} \right\rangle$ in the folded state of the protein G model, estimated as the average acceleration factor $\ln \overline{\alpha_i}$. The red open points are the sets of simulations which deviate from linearity, and are discarded. (b) The values of $\ln k_0$ predicted by Eq. 4.3 for each set of simulations for various values of γ (c) The variance in the predicted value of $\ln k_0$ across the sets of simulations. The horizontal dashed lines represent the true value of $\ln k_0$, with the 95% confidence interval represented by the blue shaded region. The vertical dashed lines represent the minimum variance value of γ.

4.8.1.2 Metadynamics biasing

We also decided to run the analysis on a set of MetaD simulations to see if this procedure also works for non-static biasing potentials. We applied the EATR-flooding approach to the MetaD simulations of the protein G model which were run with different bias deposition paces and were analyzed in Ref.[32]. The simulations were run using well-tempered MetaD using 10.0 for the BIASFACTOR parameter, with a hill height of 0.4 kJ/mol and σ of 0.06 nm for R_{ee}, and a hill height of 0.6 kJ/mol and σ of 0.02 for Q. Each set of simulations was run using a different PACE parameter, from 1 ps to 10 ns per hill deposition.

First, the linear fit of $\ln k_{obs}$ as a function of $\ln \left\langle e^{\beta V} \right\rangle$ was assessed as shown in Fig. 4.11a. The value of $\ln k_{obs}$ was obtained by taking the reciprocal of the mean transition time, as the transition time distribution was expected to be non-Poissonian and therefore fitting the CDF would not be appropriate. The value of $\ln \left\langle e^{\beta V} \right\rangle$ was estimated by taking the final value of the acceleration factor for each simulation

$$\alpha_i = \frac{1}{t_i} \int_0^{t_i} e^{\beta V(\xi(t'),t')} \, dt'$$

and averaging across all simulations in the set. The fastest paces corresponding to deposition times of 1 and 10 ps were observed to deviate from the expected linear behavior, and were thus discarded.

The EATR-flooding protocol, where $\ln \left\langle e^{\beta \gamma V} \right\rangle$ was estimated by taking the simulation average before the time average, was then applied to the remaining data, as shown in Fig. 4.11b,c. For Q, the least-variance value for γ was 0.67 ± 0.06 and the corresponding value for $\ln k_0^*$ was 0.47 ± 0.17 with k_0^* in (μs^{-1}). For R_{ee}, γ^* was 0.28 ± 0.01 and the corresponding value for $\ln k_0^*$ was 0.43 ± 0.06.

Table 4.2: Lennard-Jones parameters for the receptor and ligand atoms. CW, CP, and CF refer to the wall, cavity, and fullerene atoms respectively

atom-type	σ(nm)	$\epsilon(\frac{\text{kJ}}{\text{mol}})$
CF	0.3500	0.276144
CW	0.4152	0.002400
CP	0.4152	0.008000

Table 4.3: Lennard-Jones parameters for non-bonded interactions between receptor and ligand atoms. Parameter ϵ_1 was used in the original model, while ϵ_3 was used for the fast unbinding model. The non-bonded interactions of the CP-CP, CW-CP, and CW-CW pairs were excluded from the total energy by setting their LJ parameters to 0 and the entire lattice was position-restrained with a 1000 $\frac{kJ}{mol\,nm^2}$ force constant.

i	j	σ(nm)	$\epsilon_1(\frac{\text{kJ}}{\text{mol}})$	$\epsilon_2(\frac{\text{kJ}}{\text{mol}})$	$\epsilon_3(\frac{\text{kJ}}{\text{mol}})$	$\epsilon_4(\frac{\text{kJ}}{\text{mol}})$
CF	CP	0.3812	0.04700	0.03600	0.02500	0.01850
CF	CW	0.3812	0.02574	0.01972	0.01369	0.01015
CW	CW	0.4152	0.00000	0.00000	0.00000	0.00000
CP	CP	0.4152	0.00000	0.00000	0.00000	0.00000
CP	CW	0.4152	0.00000	0.00000	0.00000	0.00000

4.8.2 Further results for the cavity-ligand model

The receptor is a semi-hollow cube that has a cavity which is half an ellipsoid of radius 11 Å. Moreover, the cube consists of two types of hydrophobic atoms; the cavity atoms (CP) and the wall atoms (CW) (the atoms in the cavity have a higher attraction to the ligand atoms (CF) than do the wall atoms), and the whole complex model is solvated with TIP4P water. Parameters for non-bonded interactions between receptor and ligand atoms were determined using GROMACS' combination rule 3, where $\sigma_{ij} = \sqrt{\sigma_i \sigma_j}$ and $\epsilon_{ij} = \sqrt{\epsilon_i \epsilon_j}$. The ϵ values were reduced while keeping the ratio between $\epsilon_{\text{CP-CF}}$ and $\epsilon_{\text{CW-CF}}$ fixed.

In Fig. 4.12C,D, Topology 1 (top. 1) describes the original model and has the highest barrier to unbinding on the z CV. Topology 2 (top. 2) has a lower barrier but remains high and kinetics were still too slow. Topology 4 has the lowest barrier, but the timescale of the unbinding kinetics was reduced to less than a microsecond. Thus, the second-weakest model (top. 3) was chosen for

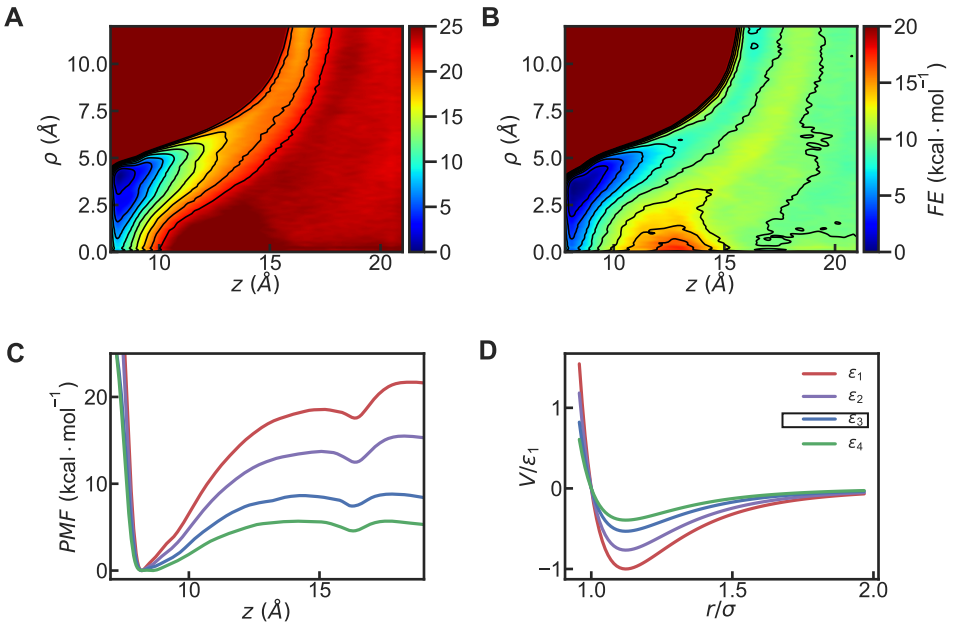

Figure 4.12: FES of original (A) and fast unbinding model, ϵ_3 (B), as function of ρ and z. (C) Projection of FES onto transverse distance, z, for all models. (D) LJ potential for interactions between cavity and ligand atoms (CP-CF) in original and modified models, potential was scaled by ϵ_1 (see Table 4.3) and distance was scaled by σ.

benchmarking.

Using the WTMetaD approach, FES calculations were performed for all models, including the original, biasing both the transverse distance, z, and radial distance, ρ. The biasing parameters were: HEIGHT=0.478, SIGMA=0.3,0.1, BIASFACTOR=10, and PACE=300. for the modified models. For the original model, only one parameter was different, BIASFACTOR=12, which was higher to obtain comparable sampling within 100 ns. Restraining potentials (PLUMED UPPER_WALLS) were applied for the transverse and radial distance CVs at 21 Å and 12 Å, respectively to limit sampling within the periodic box.

4.8.3 Original cavity-ligand model results

For the original cavity-ligand model, we performed simulations with a set of ΔE values where we biased z alone, both z and ρ, and ρ alone. For all collections of simulation sets, we ran 6 sets of

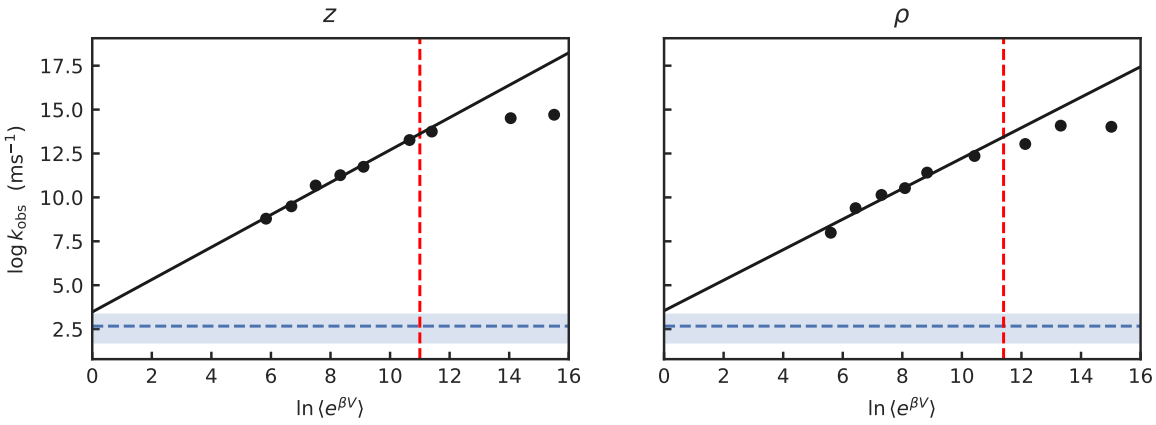

Figure 4.13: The logarithm of the observed rate plotted against $\ln \left\langle e^{\beta V} \right\rangle$ in the bound state of the cavity-ligand model, estimated as the log of the time average of $\overline{e^{\beta V(t)}}$. The horizontal dashed line represents the value of $\ln k_0$. Only points to the left of the red vertical dashed line were used in the EATRf analysis in the main text, as the rate becomes less sensitive to bias past that point, demonstrated by the deviation from linear behavior.

simulations with ΔE varying from 13 to 16 kcal/mol, except we did not run the 13 kcal/mol set for ρ alone. Because of a few outliers in the data set, the observed unbinding rates in the simulations could not be estimated using the maximum likelihood estimator. $\ln k_{\text{obs}}$ was estimated by fitting the CDF instead, where in the case where a simulation ended early without transitioning, the CDF is fit only up to the time where that simulation ended. The results from EATRf are given in Fig. 4.14. The minimum-variance residence time we obtained from EATRf is $\tau_0 = 0.023$ s ($\gamma = 0.64$) when biasing z alone, $\tau_0 = 0.021$ s ($\gamma = 0.64$) when biasing z and ρ together, and $\tau_0 = 0.097$ s ($\gamma = 0.67$) when biasing ρ alone. This is between the residence time estimates from Markov state modeling ($\tau_0 = 0.00055$ s)[126] and previous infrequent metadynamics studies ($\tau_0 = 3863$ s, $\tau_0 = 2000$ s).[124, 125]

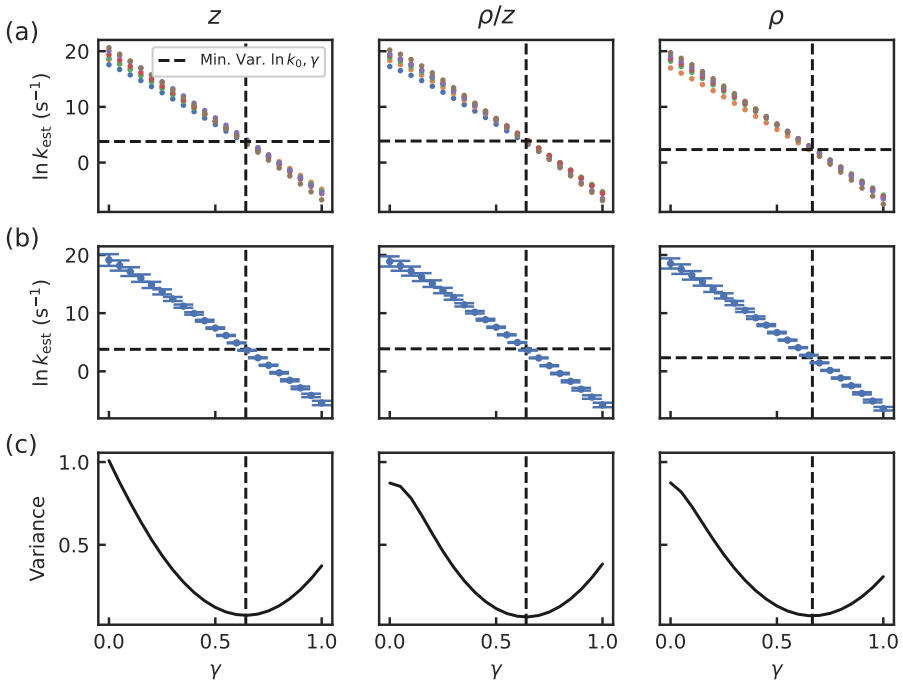

Figure 4.14: The results of EATRf on the original cavity-ligand model. (a) The values of $\ln k_0$ predicted by Eq. 4.3 for each set of OPES simulations with each BARRIER parameter for various values of γ. (b) The value of $\ln k_0$ averaged over the sets of simulations for each potential value of γ. The variance of $\ln k_0$ across the sets is represented by error bars. (c) The variance in the predicted value of $\ln k_0$ across the sets of simulations. The dashed lines represent the minimum variance value of γ, which is the value that causes all sets to give the same value of $\ln k_{est}$, and the corresponding $\ln k_0^*$.

5 | Conclusion

5.1 Summary

The goal of this thesis was to develop and apply a new method for obtaining transition rates from molecular dynamics simulations biased with enhanced sampling. The method we developed provides robust rates even when the CV chosen is not suited to biasing.

First, we derived the EATR method by generalizing the KTR method by replacing $e^{\beta\gamma\overline{V_{max}}(t)}$ with a general rate scaling function $f(t)$ then selecting $f(t) = \overline{e^{\beta\gamma V_i(t)}}$, which was chosen to be consistent with iMetaD when $\gamma = 1$. We demonstrated that EATR is consistent with iMetaD for a one-dimenstional system with a perfect CV, then that EATR either outperformed or performed as well as iMetaD when predicting unfolding rates for a coarse-grained protein model, depending on the CV chosen for biasing as well as the frequency that bias is added in MetaD. Specifically, EATR, as well as KTR, consistently gave the same rate as was obtained from unbiased simulations while iMetaD consistently underestimated it. We showed that EATR still gives the correct results for the unfolding of the chignolin miniprotein.

Second, we applied the EATR method to a set of six atomistic protein-ligand systems to track the improvement of the SPIB reaction coordinate described in Ref. [33] during training. For this, because we were not interested in the rates for these processes, we discovered we could set the unbiased rate and only fit the γ parameter to ensure stable predictions. We saw significant improvement toward the beginning of the training process in all six cases. We also found that γ was

inversely correlated with the rescaled times predicted by iMetaD, and provided the connection between them.

Third, we found that the KTR and EATR methods required a time-dependent bias to estimate γ, and therefore these methods are not suitable for obtaining unbiased rates from MD simulations with static biasing or OPES. Using the rate scaling function of EATR, we proposed a new technique to obtain rates from these types of simulations. By combining the data from multiple sets of simulations, each set with a systematically different amount of bias, we were able to recover the information needed to determine the effect of the bias on the observed rate, and therefore to estimate γ. By asserting that all sets of simulations have the same underlying unbiased rate, we obtained rate estimates by minimizing the variance in the calculated rates. This new method, EATR-flooding, was demonstrated to work on the coarse grained protein unfolding model studied in the original EATR paper, as well as on a cavity-ligand model.

5.2 Future Directions

For all of the theory in this thesis, we had assumed that the biased rate is a simple function or functional of the bias, and that the biased rate at some moment of a MetaD simulation is the same as the biased rate for a statically-biased simulation with the same bias potential. However, we expect that increasing the potential energy during the simulation will have effects on the dynamics which are not currently being taken into account. While the increase in energy will be dissipated by the thermostat over long time scales, for the frequent-biasing regime this would need to be considered.

Plugging Eq. 2.19 into Eq. 2.8, we can "solve" for γ using the same approximation as was used for Eq. 3.5, yielding $\gamma \approx (\ln k'/k_0)/(\ln \overline{e^{\beta V}})$. This implies that it is possible to obtain an analytical expression for γ from any rate theory that connects the potential energy to the rate. This could lead to a new understanding of γ and how the choice of CV to bias influences it. If such an

analytical expression exists that could be applied to any system, we could eliminate γ as a fitting parameter and instead know the correct value to use *a priori*. However, γ appears to depend not only on the choice of CV, but also on the bias potential.

Bibliography

[1] Martin Karplus and J Andrew McCammon. "Molecular dynamics simulations of biomolecules". In: *Nat. Struct. Mol. Biol.* 9.9 (2002), pp. 646–652.

[2] Tjerk P Straatsma and J Andrew McCammon. "Computational Alchemy". In: *Annual Review of Physical Chemistry* 43.Volume 43, (1992), pp. 407–435. ISSN: 1545-1593.

[3] Johan Åqvist, Victor B. Luzhkov, and Bjørn O. Brandsdal. "Ligand Binding Affinities from MD Simulations". In: *Accounts of Chemical Research* 35.6 (2002), pp. 358–365.

[4] C Bruin. "Transport coefficients from molecular dynamics for a very small system". In: *Physics Letters A* 28.11 (1969), pp. 777–778. ISSN: 0375-9601.

[5] William L Jorgensen, Jayaraman Chandrasekhar, Jeffry D Madura, Roger W Impey, and Michael L Klein. "Comparison of simple potential functions for simulating liquid water". In: *J. Chem. Phys.* 79.2 (1983), pp. 926–935.

[6] Christian Carbogno, Rampi Ramprasad, and Matthias Scheffler. "Ab Initio Green-Kubo Approach for the Thermal Conductivity of Solids". In: *Phys. Rev. Lett.* 118 (17 Apr. 2017), p. 175901.

[7] Paolo Pegolo, Enrico Drigo, Federico Grasselli, and Stefano Baroni. "Transport coefficients from equilibrium molecular dynamics". In: *J. Chem. Phys.* 162.6 (Feb. 2025), p. 064111. ISSN: 0021-9606.

[8] Pratyush Tiwary, Vittorio Limongelli, Matteo Salvalaglio, and Michele Parrinello. “Kinetics of protein–ligand unbinding: Predicting pathways, rates, and rate-limiting steps”. In: *Proc. Natl. Acad. Sci.* 112.5 (2015), E386–E391.

[9] Willmor J Peña Ccoa and Glen M Hocky. “Assessing models of force-dependent unbinding rates via infrequent metadynamics”. In: *J. Chem. Phys.* 156.12 (2022), p. 125102.

[10] Robert A. Copeland, David L. Pompliano, and Thomas D. Meek. “Drug–target residence time and its implications for lead optimization”. In: *Nat. Rev. Drug Desc.* 5.9 (Sept. 2006), pp. 730–739. ISSN: 1474-1784.

[11] Giovanni Bussi and Michele Parrinello. “Accurate sampling using Langevin dynamics”. In: *Phys. Rev. E* 75 (5 May 2007), p. 056707.

[12] Shūichi Nosé. “A molecular dynamics method for simulations in the canonical ensemble”. In: *Mol. Phys.* 52.2 (1984), pp. 255–268.

[13] Glenn J Martyna, Michael L Klein, and Mark Tuckerman. “Nosé–Hoover chains: The canonical ensemble via continuous dynamics”. In: *J. Chem. Phys.* 97.4 (1992), pp. 2635–2643.

[14] J. P. Valleau and S. G. Whittington. “A Guide to Monte Carlo for Statistical Mechanics: 1. Highways”. In: *Statistical Mechanics: Part A: Equilibrium Techniques*. Ed. by Bruce J. Berne. Boston, MA: Springer US, 1977, pp. 137–168. ISBN: 978-1-4684-2553-6.

[15] Alessandro Laio and Michele Parrinello. “Escaping free-energy minima”. In: *Proc. Natl. Acad. Sci. U.S.A.* 99.20 (2002), pp. 12562–12566.

[16] Michele Invernizzi and Michele Parrinello. “Rethinking metadynamics: From bias potentials to probability distributions”. In: *J. Phys. Chem. Lett.* 11.7 (2020), pp. 2731–2736.

[17] Alessandro Barducci, Giovanni Bussi, and Michele Parrinello. “Well-tempered metadynamics: a smoothly converging and tunable free-energy method”. In: *Phys. Rev. Lett.* 100.2 (2008), p. 020603.

[18] James F Dama, Michele Parrinello, and Gregory A Voth. “Well-tempered metadynamics converges asymptotically”. In: *Phys. Rev. Lett.* 112.24 (2014), p. 240602.

[19] Christoph Dellago, Peter G Bolhuis, Félix S Csajka, and David Chandler. “Transition path sampling and the calculation of rate constants”. In: *J. Chem. Phys.* 108.5 (1998), pp. 1964–1977.

[20] Peter G Bolhuis, David Chandler, Christoph Dellago, and Phillip L Geissler. “Transition path sampling: Throwing ropes over rough mountain passes, in the dark”. In: *Annu. Rev. Phys. Chem.* 53.1 (2002), pp. 291–318.

[21] Rose Du, Vijay S Pande, Alexander Yu Grosberg, Toyoichi Tanaka, and Eugene S Shakhnovich. “On the transition coordinate for protein folding”. In: *J. Chem. Phys.* 108.1 (1998), pp. 334–350.

[22] Ao Ma and Aaron R Dinner. “Automatic method for identifying reaction coordinates in complex systems”. In: *J. Phys. Chem. B* 109.14 (2005), pp. 6769–6779.

[23] Andrew R. Mitchell and Grant M. Rotskoff. “Committor Guided Estimates of Molecular Transition Rates”. In: *Journal of Chemical Theory and Computation* 20.21 (2024). PMID: 39420582, pp. 9378–9393.

[24] Chatipat Lorpaiboon, Jonathan Weare, and Aaron R. Dinner. “An Exact Multiple-Time-Step Variational Formulation for the Committor and the Transition Rate”. In: *The Journal of Physical Chemistry B* 130.1 (2026). PMID: 41404981, pp. 155–170.

[25] Arthur F Voter. “A method for accelerating the molecular dynamics simulation of infrequent events”. In: *J. Chem. Phys.* 106 (1997), pp. 4665–4677.

[26] Arthur F Voter. “Hyperdynamics: Accelerated molecular dynamics of infrequent events”. In: *Phys. Rev. Lett.* 78.20 (1997), p. 3908.

[27] Helmut Grubmüller. “Predicting slow structural transitions in macromolecular systems: Conformational flooding”. In: *Phys. Rev. E* 52 (3 Sept. 1995), pp. 2893–2906.

[28] Pratyush Tiwary and Michele Parrinello. “From Metadynamics to Dynamics”. In: *Phys. Rev. Lett.* 111 (23 Dec. 2013), p. 230602.

[29] Matteo Salvalaglio, Pratyush Tiwary, and Michele Parrinello. “Assessing the Reliability of the Dynamics Reconstructed from Metadynamics”. In: *J. Chem. Theory Comput.* 10.4 (2014), pp. 1420–1425.

[30] Dhiman Ray, Narjes Ansari, Valerio Rizzi, Michele Invernizzi, and Michele Parrinello. “Rare event kinetics from adaptive bias enhanced sampling”. In: *J. Chem. Theor. Comput.* 18.11 (2022), pp. 6500–6509.

[31] Karen Palacio-Rodriguez, Hadrien Vroylandt, Lukas S. Stelzl, Fabio Pietrucci, Gerhard Hummer, and Pilar Cossio. “Transition Rates and Efficiency of Collective Variables from Time-Dependent Biased Simulations”. In: *J. Phys. Chem. Lett.* 13.32 (2022), pp. 7490–7496.

[32] Nicodemo Mazzaferro, Subarna Sasmal, Pilar Cossio, and Glen M. Hocky. “Good Rates From Bad Coordinates: The Exponential Average Time-dependent Rate Approach”. In: *J. Chem. Theory Comput.* 20.14 (2024). PMID: 38954555, pp. 5901–5912.

[33] Suemin Lee, Dedi Wang, Markus A. Seeliger, and Pratyush Tiwary. “Calculating Protein–Ligand Residence Times through State Predictive Information Bottleneck Based Enhanced Sampling”. In: *J. Chem. Theor. Comput.* 20.14 (2024), pp. 6341–6349.

[34] David Chandler. “Barrier crossings: classical theory of rare but important events”. In: *Classical and quantum dynamics in condensed phase simulations* 523 (1998).

[35] Jérôme Hénin, Tony Lelièvre, Michael Shirts, Omar Valsson, and Lucie Delemotte. "Enhanced Sampling Methods for Molecular Dynamics Simulations [Article v1. 0]". In: *LiveCoMS* 4.1 (2022), pp. 1583–1583.

[36] Mark E Tuckerman. *Statistical mechanics: theory and molecular simulation*. Oxford university press, 2023.

[37] Glenn M Torrie and John P Valleau. "Nonphysical sampling distributions in Monte Carlo free-energy estimation: Umbrella sampling". In: *J. Comp. Phys.* 23.2 (1977), pp. 187–199.

[38] Johannes Kästner. "Umbrella sampling". In: *Wiley Interdisciplinary Reviews: Computational Molecular Science* 1.6 (2011), pp. 932–942.

[39] Eric Darve and Andrew Pohorille. "Calculating free energies using average force". In: *J. Chem. Phys.* 115.20 (2001), pp. 9169–9183.

[40] Jeffrey Comer, James C Gumbart, Jérôme Hénin, Tony Lelièvre, Andrew Pohorille, and Christophe Chipot. "The adaptive biasing force method: Everything you always wanted to know but were afraid to ask". In: *J. Phys. Chem. B* 119.3 (2015), pp. 1129–1151.

[41] Giovanni Bussi and Alessandro Laio. "Using metadynamics to explore complex free-energy landscapes". In: *Nat. Rev. Phys.* 2.4 (2020), pp. 200–212.

[42] Omar Valsson and Michele Parrinello. "Variational approach to enhanced sampling and free energy calculations". In: *Phys. Rev. Lett.* 113.9 (2014), p. 090601.

[43] Brooke E Husic and Vijay S Pande. "Markov State Models: From an Art to a Science". In: *J. Am. Chem. Soc.* 140.7 (2018), pp. 2386–2396.

[44] Erik H Thiede, Dimitrios Giannakis, Aaron R Dinner, and Jonathan Weare. "Galerkin approximation of dynamical quantities using trajectory data". In: *J. Chem. Phys.* 150.24 (2019).

[45] Rosalind J Allen, Chantal Valeriani, and Pieter Rein Ten Wolde. "Forward flux sampling for rare event simulations". In: *J. Phys. Cond. Mat.* 21.46 (2009), p. 463102.

[46] Nicholas Guttenberg, Aaron R Dinner, and Jonathan Weare. "Steered transition path sampling". In: *J. Chem. Phys.* 136.23 (2012).

[47] Daniel M. Zuckerman and Lillian T. Chong. "Weighted Ensemble Simulation: Review of Methodology, Applications, and Software". In: *Annu. Rev. Biophys.* 46.1 (2017), pp. 43–57.

[48] Titus S. van Erp, Daniele Moroni, and Peter G. Bolhuis. "A novel path sampling method for the calculation of rate constants". In: *J. Chem. Phys.* 118.17 (Apr. 2003), pp. 7762–7774. ISSN: 0021-9606.

[49] Ron Elber. "Milestoning: An efficient approach for atomically detailed simulations of kinetics in biophysics". In: *Annu. Rev. Biophys.* 49 (2020), pp. 69–85.

[50] Rodrigo Casasnovas, Vittorio Limongelli, Pratyush Tiwary, Paolo Carloni, and Michele Parrinello. "Unbinding kinetics of a p38 MAP kinase type II inhibitor from metadynamics simulations". In: *J. Am. Chem. Soc.* 139.13 (2017), pp. 4780–4788.

[51] Miłosz Wieczór, Phu K Tang, Modesto Orozco, and Pilar Cossio. "Omicron mutations increase interdomain interactions and reduce epitope exposure in the SARS-CoV-2 spike". In: *Iscience* 26.2 (2023).

[52] David Gomez, Willmor J Peña Ccoa, Yuvraj Singh, Enrique Rojas, and Glen M Hocky. "Molecular paradigms for biological mechanosensing". In: *J. Phys. Chem. B* 125.44 (2021), pp. 12115–12124.

[53] Fatemah Mukadum, Willmor J Peña Ccoa, and Glen M Hocky. "Molecular simulation approaches to probing the effects of mechanical forces in the actin cytoskeleton". In: *Cytoskeleton* (2024), pp. 1–10.

[54] Michael McGovern and Juan De Pablo. “A boundary correction algorithm for metadynamics in multiple dimensions”. In: *J. Chem. Phys.* 139.8 (2013), p. 084102.

[55] James F Dama, Glen M Hocky, Rui Sun, and Gregory A Voth. “Exploring valleys without climbing every peak: more efficient and forgiving metabasin metadynamics via robust on-the-fly bias domain restriction”. In: *J. Chem. Theor. Comput.* 11.12 (2015), pp. 5638–5650.

[56] Pratyush Tiwary, Jagannath Mondal, and Bruce J Berne. “How and when does an anticancer drug leave its binding site?” In: *Sci. Adv.* 3.5 (2017), e1700014.

[57] Dhiman Ray and Michele Parrinello. “Kinetics from Metadynamics: Principles, Applications, and Outlook”. In: *J. Chem. Theor. Comput.* 19.17 (2023), pp. 5649–5670.

[58] Ofir Blumer, Shlomi Reuveni, and Barak Hirshberg. “Short-Time Infrequent Metadynamics for Improved Kinetics Inference”. In: *J. Chem. Theor. Comput.* 20.9 (2024). PMID: 38668722, pp. 3484–3491.

[59] H.A. Kramers. “Brownian motion in a field of force and the diffusion model of chemical reactions”. In: *Physica* 7.4 (1940), pp. 284–304. ISSN: 0031-8914.

[60] Alexander M. Berezhkovskii and Attila Szabo. “Population Fluctuations at Equilibrium and Kramers’ Rate of Diffusive Barrier Crossing”. In: *J. Phys. Chem. B* 127.22 (2023), pp. 5084–5088.

[61] Peter Hänggi, Peter Talkner, and Michal Borkovec. “Reaction-rate theory: fifty years after Kramers”. In: *Rev. Mod. Phys.* 62.2 (1990), p. 251.

[62] Robert Zwanzig. *Nonequilibrium statistical mechanics*. Oxford university press, 2001.

[63] Wayne Nelson. “Maximum Likelihood Analysis of Multiply Censored Data”. In: *Applied Life Data Analysis*. John Wiley & Sons, Ltd, 1982. Chap. 8, pp. 313–404. ISBN: 9780471725237.

[64] Robert Zwanzig. “Rate processes with dynamical disorder”. In: *Acc. Chem. Res.* 23.5 (1990), pp. 148–152.

[65] Olga K Dudko, Gerhard Hummer, and Attila Szabo. “Intrinsic rates and activation free energies from single-molecule pulling experiments”. In: *Phys. Rev. Lett.* 96.10 (2006), p. 108101.

[66] Pilar Cossio, Gerhard Hummer, and Attila Szabo. “Kinetic ductility and force-spike resistance of proteins from single-molecule force spectroscopy”. In: *Biophys. J.* 111.4 (2016), pp. 832–840.

[67] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. “SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python”. In: *Nat. Methods* 17 (2020), pp. 261–272.

[68] Jorge J Moré. “The Levenberg-Marquardt algorithm: implementation and theory”. In: *Numerical analysis: proceedings of the biennial Conference held at Dundee, June 28–July 1, 1977*. Springer. 2006, pp. 105–116.

[69] B. Efron. “Bootstrap Methods: Another Look at the Jackknife”. In: *Ann. Stat.* 7.1 (1979), pp. 1–26.

[70] Olga K. Dudko, Thomas G. W. Graham, and Robert B. Best. “Locating the Barrier for Folding of Single Molecules under an External Force”. In: *Phys. Rev. Lett.* 107 (20 Nov. 2011), p. 208301.

[71] Subarna Sasmal, Martin McCullagh, and Glen M. Hocky. “Reaction Coordinates for Conformational Transitions Using Linear Discriminant Analysis on Positions”. In: *J. Chem. Theor. Comput.* 19.14 (2023), pp. 4427–4435.

[72] Robert B. Best, Gerhard Hummer, and William A. Eaton. “Native contacts determine protein folding mechanisms in atomistic simulations”. In: *Proc. Natl. Acad. Sci. U.S.A.* 110.44 (2013), pp. 17874–17879.

[73] Massimiliano Bonomi, Giovanni Bussi, Carlo Camilloni, Gareth A Tribello, Pavel Banáš, Alessandro Barducci, Mattia Bernetti, Peter G Bolhuis, Sandro Bottaro, Davide Branduardi, et al. “Promoting transparency and reproducibility in enhanced molecular simulations”. In: *Nat. Methods* 16.8 (2019), pp. 670–673.

[74] Benjamin Kuznets-Speck and David T. Limmer. “Inferring equilibrium transition rates from nonequilibrium protocols”. In: *Biophys. J.* 122 (2023), pp. 1659–1664. ISSN: 0006-3495.

[75] Dan Mendels, Giovannimaria Piccini, Z Faidon Brotzakis, Yi I Yang, and Michele Parrinello. “Folding a small protein using harmonic linear discriminant analysis”. In: *J. Chem. Phys.* 149.19 (2018), p. 194113.

[76] Pilar Cossio, Gerhard Hummer, and Attila Szabo. “On artifacts in single-molecule force spectroscopy”. In: *Proc. Natl. Acad. Sci. USA* 112.46 (2015), pp. 14248–14253.

[77] Pilar Cossio, Gerhard Hummer, and Attila Szabo. “Transition paths in single-molecule force spectroscopy”. In: *J. Chem. Phys.* 148.12 (2018), p. 123309.

[78] John Karanicolas and Charles L Brooks. “The origins of asymmetry in the folding transition states of protein L and protein G”. In: *Prot. Sci.* 11.10 (2002), pp. 2281–2531.

[79] Thomas G. W. Graham and Robert B. Best. “Force-Induced Change in Protein Unfolding Mechanism: Discrete or Continuous Switch?” In: *J. Phys. Chem. B* 115.6 (2011), pp. 1546–1561.

[80] John Karanicolas and Charles L Brooks. “Improved Go-like models demonstrate the robustness of protein folding mechanisms towards non-native interactions”. In: *J. Mol. Biol.* 334.2 (2003), pp. 309–325.

[81] Aidan P Thompson, H Metin Aktulga, Richard Berger, Dan S Bolintineanu, W Michael Brown, Paul S Crozier, Pieter J In’t Veld, Axel Kohlmeyer, Stan G Moore, Trung Dac Nguyen, et al. “LAMMPS-a flexible simulation tool for particle-based materials modeling at the atomic, meso, and continuum scales”. In: *Comp. Phys. Comm.* 271 (2022), p. 108171.

[82] Heidi Klem, Glen M. Hocky, and Martin McCullagh. “Size-and-Shape Space Gaussian Mixture Models for Structural Clustering of Molecular Dynamics Trajectories”. In: *J. Chem. Theory Comput.* 18.5 (2022), pp. 3218–3230.

[83] Nicodemo Mazzaferro, Suemin Lee, Pilar Cossio, Pratyush Tiwary, and Glen M Hocky. “Using Time Dependent Rate Analysis to Evaluate the Quality of Machine Learned Reaction Coordinates for Biasing and Computing Kinetics”. In: *J. Phys. Chem. B* 129.42 (2025), pp. 10967–10974.

[84] Jinan Wang, Hung N. Do, Kushal Koirala, and Yinglong Miao. “Predicting Biomolecular Binding Kinetics: A Review”. In: *J. Chem. Theor. Comput.* 19.8 (2023), pp. 2135–2148.

[85] Anupam Anand Ojha, Lane W. Votapka, and Rommie E. Amaro. “Advances and Challenges in Milestoning Simulations for Drug–Target Kinetics”. In: *J. Chem. Theor. Comput.* 20.22 (2024), pp. 9759–9769.

[86] E.A. Carter, Giovanni Ciccotti, James T. Hynes, and Raymond Kapral. “Constrained reaction coordinate dynamics for the simulation of rare events”. In: *Chem. Phys. Lett.* 156.5 (1989), pp. 472–477. ISSN: 0009-2614.

[87] Aditya N. Singh, Avishek Das, and David T. Limmer. “Variational Path Sampling of Rare Dynamical Events”. In: *Annu. Rev. Phys. Chem.* 76.Volume 76, 2025 (2025), pp. 639–662. ISSN: 1545-1593.

[88] Omar Valsson, Pratyush Tiwary, and Michele Parrinello. “Enhancing Important Fluctuations: Rare Events and Metadynamics from a Conceptual Viewpoint”. In: *Annu. Rev. Phys. Chem.* 67.1 (2016), pp. 159–184.

[89] Benjamin Kuznets-Speck and David T Limmer. “Inferring equilibrium transition rates from nonequilibrium protocols”. In: *Biophys. J.* 122.9 (2023), pp. 1659–1664.

[90] Baron Peters. “Reaction Coordinates and Mechanistic Hypothesis Tests”. In: *Annu. Rev. Phys. Chem.* 67.Volume 67, 2016 (2016), pp. 669–690. ISSN: 1545-1593.

[91] Weinan E. and Eric Vanden-Eijnden. “Towards a Theory of Transition Paths”. In: *J. Stat. Phys.* 123.3 (May 2006), pp. 503–523.

[92] Robert B Best and Gerhard Hummer. “Reaction coordinates and rates from transition paths”. In: *Proc. Natl. Acad. Sci.* 102.19 (2005), pp. 6732–6737.

[93] Qianxiao Li, Bo Lin, and Weiqing Ren. “Computing committor functions for the study of rare events using deep learning”. In: *J. Chem. Phys.* 151.5 (Aug. 2019), p. 054112. ISSN: 0021-9606.

[94] Yuehaw Khoo, Jianfeng Lu, and Lexing Ying. “Solving for high-dimensional committor functions using artificial neural networks”. In: *Res. Math. Sci.* 6 (2019), pp. 1–13.

[95] Haochuan Chen, Benoıt Roux, and Christophe Chipot. “Discovering reaction pathways, slow variables, and committor probabilities with machine learning”. In: *J. Chem. Theory Comput.* 19.14 (2023), pp. 4414–4426.

[96] Peilin Kang, Enrico Trizio, and Michele Parrinello. “Computing the committor with the committor to study the transition state ensemble”. In: *Nat. Comp. Sci.* 4.6 (2024), pp. 451–460.

[97] Mary A Rohrdanz, Wenwei Zheng, Mauro Maggioni, and Cecilia Clementi. "Determination of reaction coordinates via locally scaled diffusion map". In: *J. Chem. Phys.* 134.12 (2011), 03B624.

[98] R. R. Coifman, S. Lafon, I. G. Kevrekidis, M. Maggioni, and B. Nadler. "Diffusion maps, reduction coordinates, and low dimensional representation of stochastic systems". In: *Multiscale Modeling and Simulation* 7.2 (2008), pp. 842–864. ISSN: 1540-3459.

[99] Dan Mendels, GiovanniMaria Piccini, and Michele Parrinello. "Collective variables from local fluctuations". In: *J. Phys. Chem. Lett.* 9.11 (2018), pp. 2776–2781.

[100] Subarna Sasmal, Martin McCullagh, and Glen M. Hocky. "Improved Data-Driven Collective Variables for Biased Sampling through Iteration on Biased Data". In: *The Journal of Physical Chemistry B* 129.25 (2025). PMID: 40506852, pp. 6163–6171.

[101] L. Molgedey and H. G. Schuster. "Separation of a mixture of independent signals using time delayed correlations". In: *Phys. Rev. Lett.* 72 (23 June 1994), pp. 3634–3637.

[102] Andreas Mardt, Luca Pasquali, Hao Wu, and Frank Noé. "VAMPnets for deep learning of molecular kinetics". In: *Nature Commun.* 9.1 (2018), pp. 1–11.

[103] Pratyush Tiwary and B. J. Berne. "Spectral gap optimization of order parameters for sampling complex molecular systems". In: *Proc. Natl. Acad. Sci.* 113.11 (2016), pp. 2839–2844.

[104] Christoph Wehmeyer and Frank Noé. "Time-lagged autoencoders: Deep learning of slow collective variables for molecular kinetics". In: *J. Chem. Phys.* 148.24 (Mar. 2018), p. 241703. ISSN: 0021-9606.

[105] Shams Mehdi, Zachary Smith, Lukas Herron, Ziyue Zou, and Pratyush Tiwary. "Enhanced Sampling with Machine Learning". In: *Annual Review of Physical Chemistry* 75.Volume 75, 2024 (2024), pp. 347–370. ISSN: 1545-1593.

[106] Dedi Wang and Pratyush Tiwary. “State predictive information bottleneck”. In: *J. Chem. Phys.* 154.13 (2021), p. 134111.

[107] Joao Marcelo Lamim Ribeiro, Pablo Bravo, Yihang Wang, and Pratyush Tiwary. “Reweighted autoencoded variational Bayes for enhanced sampling (RAVE)”. In: *J. Chem. Phys.* 149.7 (2018), p. 072301.

[108] Yihang Wang, João Marcelo Lamim Ribeiro, and Pratyush Tiwary. “Past–future information bottleneck for sampling molecular reaction coordinate simultaneously with thermodynamics and kinetics”. In: *Nature Communications* 10.1 (2019), p. 3573.

[109] Ziyue Zou, Eric R. Beyerle, Sun-Ting Tsai, and Pratyush Tiwary. “Driving and characterizing nucleation of urea and glycine polymorphs in water”. In: *Proc. Natl. Acad. Sci.* 120.7 (2023), e2216099120.

[110] Shams Mehdi, Dedi Wang, Shashank Pant, and Pratyush Tiwary. “Accelerating All-Atom Simulations and Gaining Mechanistic Understanding of Biophysical Systems through State Predictive Information Bottleneck”. In: *J. Chem. Theor. Comput.* 18.5 (2022), pp. 3231–3238.

[111] Bodhi P. Vani, Akashnathan Aranganathan, Dedi Wang, and Pratyush Tiwary. “AlphaFold2-RAVE: From Sequence to Boltzmann Ranking”. In: *J. Chem. Theor. Comput.* 19.14 (2023), pp. 4351–4354.

[112] Bodhi P. Vani, Akashnathan Aranganathan, and Pratyush Tiwary. “Exploring Kinase Asp-Phe-Gly (DFG) Loop Conformational Stability with AlphaFold2-RAVE”. In: *J. Chem. Inf. Model.* 64.7 (2023), pp. 2789–2797.

[113] Nancy D. Pomarici, Shams Mehdi, Patrick K. Quoika, Suemin Lee, Johannes R. Loeffler, Klaus R. Liedl, Pratyush Tiwary, and Monica L. Fernández-Quintero. “Learning high-dimensional reaction coordinates of fast-folding proteins using State Predictive information bottleneck and Bias Exchange Metadynamics”. In: *bioRxiv* (2023).

[114] Gareth A Tribello, Massimiliano Bonomi, Giovanni Bussi, Carlo Camilloni, Blake I Armstrong, Andrea Arsiccio, Simone Aureli, Federico Ballabio, Mattia Bernetti, Luigi Bonati, et al. "PLUMED Tutorials: A collaborative, community-driven learning ecosystem". In: *J. Chem. Phys.* 162.9 (2025).

[115] David Van Der Spoel, Erik Lindahl, Berk Hess, Gerrit Groenhof, Alan E Mark, and Herman JC Berendsen. "GROMACS: fast, flexible, and free". In: *J. Comp. Chem.* 26.16 (2005), pp. 1701–1718.

[116] Nicodemo Mazzaferro, Willmor J Pena Ccoa, Pilar Cossio, and Glen M. Hocky. *Stepping up enhanced rate calculations with EATR-flooding*. 2026.

[117] Noelia Ferruz and Gianni De Fabritiis. "Binding kinetics in drug discovery". In: *Mol. Inform.* 35.6-7 (2016), pp. 216–226.

[118] Samuel D Lotz and Alex Dickson. "Unbiased molecular dynamics of 11 min timescale drug unbinding reveals transition state stabilizing interactions". In: *J. Am. Chem. Soc.* 140.2 (2018), pp. 618–628.

[119] Mattia Bernetti, Matteo Masetti, Walter Rocchia, and Andrea Cavalli. "Kinetics of drug binding and residence time". In: *Annu. Rev. Phys. Chem.* 70.1 (2019), pp. 143–171.

[120] Ariane Nunes-Alves, Daria B Kokh, and Rebecca C Wade. "Recent progress in molecular simulation methods for drug binding kinetics". In: *Curr. Opin. Struct. Biol.* 64 (2020), pp. 126–133.

[121] Surl-Hee Ahn, Anupam A Ojha, Rommie E Amaro, and J Andrew McCammon. "Gaussian-accelerated molecular dynamics with the weighted ensemble method: A hybrid method improves thermodynamic and kinetic sampling". In: *J. Chem. Theory Comput.* 17.12 (2021), pp. 7938–7951.

[122] Zachary Smith, Davide Branduardi, Dmitry Lupyan, Giulia D'Arrigo, Pratyush Tiwary, Lingle Wang, and Goran Krilov. "Toward automated physics-based absolute drug residence time predictions". In: *J. Chem. Inf. Model.* 65.24 (2025), pp. 13360–13373.

[123] Tamar Schlick, Stephanie Portillo-Ledesma, Christopher G. Myers, Lauren Beljak, Justin Chen, Sami Dakhel, Daniel Darling, Sayak Ghosh, Joseph Hall, Mikaeel Jan, Emily Liang, Sera Saju, Mackenzie Vohr, Chris Wu, Yifan Xu, and Eva Xue. "Biomolecular Modeling and Simulation: A Prospering Multidisciplinary Field". In: *Ann. Rev. Biophys.* 50.1 (2021). PMID: 33606945, pp. 267–301.

[124] Jagannath Mondal, Joseph A. Morrone, and B. J. Berne. "How hydrophobic drying forces impact the kinetics of molecular recognition". In: *Proc. Natl. Acad. Sci. USA* 110.33 (2013), pp. 13277–13282.

[125] Pratyush Tiwary, Jagannath Mondal, Joseph A. Morrone, and B. J. Berne. "Role of water and steric constraints in the kinetics of cavity–ligand unbinding". In: *Proc. Natl. Acad. Sci. USA* 112.39 (2015), pp. 12015–12019.

[126] Navjeet Ahalawat, Satyabrata Bandyopadhyay, and Jagannath Mondal. "On the role of solvent in hydrophobic cavity–ligand recognition kinetics". In: *J. Chem. Phys.* 152.7 (Feb. 2020), p. 074104. ISSN: 0021-9606.

[127] Willmor J Peña Ccoa, Fatemah Mukadum, Aubin Ramon, Guillaume Stirnemann, and Glen M Hocky. "A direct computational assessment of vinculin–actin unbinding kinetics reveals catch-bonding behavior". In: *Proc. Natl. Acad. Sci. USA* 122.21 (2025), e2425982122.

[128] Mark James Abraham, Teemu Murtola, Roland Schulz, Szilárd Páll, Jeremy C Smith, Berk Hess, and Erik Lindahl. "GROMACS: High performance molecular simulations through multi-level parallelism from laptops to supercomputers". In: *SoftwareX* 1 (2015), pp. 19–25.

[129] Gareth A Tribello, Massimiliano Bonomi, Davide Branduardi, Carlo Camilloni, and Giovanni Bussi. “PLUMED 2: New feathers for an old bird”. In: *Comput. Phys. Comm.* 185.2 (2014), pp. 604–613.

[130] Joseph G Ibrahim, Ming-Hui Chen, and Debajyoti Sinha. *Bayesian survival analysis*. Springer Science & Business Media, 2001.