Title
Mountains and rivers: rare events in noisy systems and the forces that shape them

Author
Kuznets-Speck, Benjamin

Publication Date
2023

Peer reviewed|Thesis/dissertation

Mountains and rivers: rare events in noisy systems and the forces that shape them

by

Benjamin Kuznets-Speck

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Biophysics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Carlos Bustamante, Co-Chair
Professor David Limmer, Co-chair
Professor Michael DeWeese
Professor Hernan Garcia

Spring 2023

Mountains and rivers: rare events in noisy systems and the forces that shape them

Abstract

Mountains and rivers: rare events in noisy systems and the forces that shape them

by

Benjamin Kuznets-Speck

Doctor of Philosophy in Biophysics

University of California, Berkeley

Professor Carlos Bustamante, Co-Chair

Professor David Limmer, Co-chair

Rare events are ubiquitous in noisy complex systems throughout the physical sciences and to large extent determine their function and regulation. Dissipative outside forces often work hand in hand with equilibrium structure to shape the mechanism and frequency of such improbable fluctuations, but little is known about how to codify the influence of non-equilibrium on reaction rates and their mechanisms. In the last quarter century we have seen paradigm shifting breakthroughs in reaction rate theory that have allowed for the study of rare transitions in complex many particle systems. At the same time, development of the statistical mechanics of trajectories has revolutionized how we study the behavior, response and functional limits of systems away from equilibrium. Here, we develop a trajectory theory of how reaction rates respond to nonequilibrium forces, allowing us both to probe how non-equilibrium systems regulate their function, and leverage optimally designed forces to sample reaction rates from finite time driven trajectories for the first time.

To the family and friends who walk beside me always.

The world is very dusty, uncle. Let us work.
One day the sickness shall pass from the earth for good.
The orchard will bloom; someone will play the guitar.
Our work will be seen as strong and clean and good.
And all that we suffered through having existed
Shall be forgotten as though it had never existed.
–Donald Justice

Contents

Acknowledgments

I don't know where I would be without the help of my advisors throughout my jouney to a PhD. I would like to thank David Limmer for teaching me statistical mechanics and how to write and revise a paper start to finish. I thank Michael Hinczewski for wonderful collaborations and a constant stream of advice and support along the way. Finally, I thank Carlos Bustamante for welcoming me into his lab with open arms and fostering my love for single molecule biophysics.

I have been blessed with many amazing colleagues who have made even better friends. For their time, patience, laughter and love, I thank Adrianne Zhong, Trevor Grand Pre, Avishek Das, Aditya Singh, Robert Sosa, JeongHoon Kim and Francesca Burgos.

My friends and family have been my rock and for their support I thank my mothers Naomi Kuznets and Sherri Speck, my aunt Gina Pingitore, and Timothy Cater who for much of my PhD was a constant companion and the largest source of light in my life.

Chapter 1

Introduction

The biophysical world is random on the length scales over which proteins and molecular machines operate. This inherent randomness, due to thermal fluctuations, causes molecules to diffuse in their aqueous surroundings. On rare occasions, this noise can build up to affect large scale transitions between conformational states that would otherwise be impossible. The function and regulation of such a molecular system, i.e. the probabilities of being in each state and how they change in time, is determined by the rates at which the system switches between states. If one were able to tune transition rates, by injecting external energy through additional forces, they could control the function of such systems and design new ones to achieve specific tasks and preferentially go towards or steer away from target states. It is therefore natural to ask how transition rates change when time-dependent external forces are applied to a noisy physical system. Throughout this work, we will ask time and time again how rates are shaped by nonequilibrium external forces, how such forces regulate the mechanism of transitions and how they can be used to determine rates when no external forces are present at all.

1.1 Configurations and Trajectories

Configurations $\mathbf{x}$(t) are snapshots of physical systems at specific points in time, and trajectories $\mathbf{X} = \{\mathbf{x}(t_0), ...\mathbf{x}(t_f)\}$ are how these snapshots are strung together to make a movie. $\mathbf{x}$(t) could represent the full state of the system– the positions and velocities of all constituent particles– or it could be a collective variable, or order parameter, such as the end to end distance of a protein. A stochastic trajectory of a particle diffusing in a system with two stable states is shown in Fig. 1.1. In equilibrium, it is enough to study these snapshots by themselves, devoid of the forces that join them together through time. Statistical mechanics has traditionally been the study of configurational distributions, averaged over time but fundamentally outside of it. Only in recent years has it been realized that the statistical mechanics of trajectories holds a key to probing nonequilibriumn systems with a framework analogous to equilibrium thermodynamics.

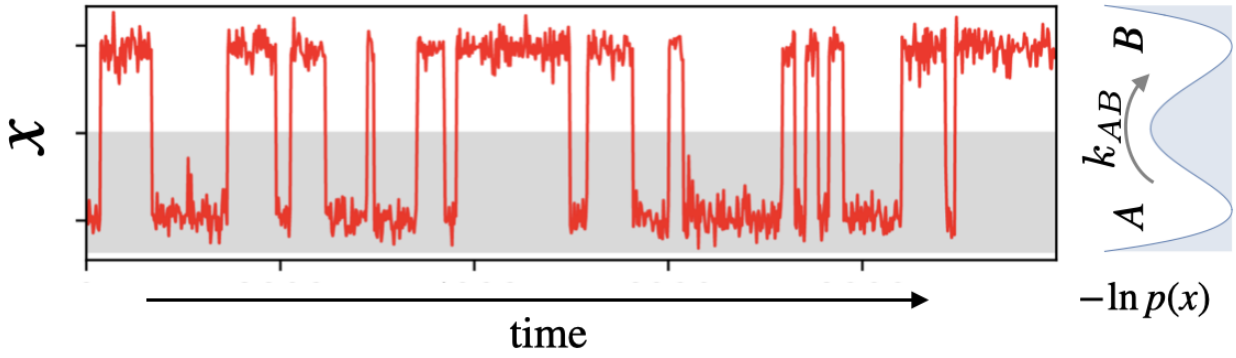

Figure 1.1: An order parameter x which describes the system state transitions between two stable regions of phase space, A and B, shaded and white respectively. We can histogram the trajectory over time to yield the probability distribution $p(x)$, and $-\ln p(x)$ can be viewed as a free energy surface in which the system diffuses. Transitions from A to B occur at rate k_{AB} and require crossing a free energy barrier.

1.2 Transition rate theory

The rates at which stable states transition between one another determines how the states are distributed. Transition rates can be defined as the inverse first passage time for a given transition to occur

$$k = 1/\langle t \rangle. \tag{1.1}$$

This description is consistent with defining the rate as a cumulative transition probability per unit time $k = P[B|A]/\tau$ whenever the transition is rare.[1] Here the transition path time τ is the typical time to cross from A to B when the transition has been initiated, and it can usually be described by the local correlation timescale of state A. Since each transition attempt is uncorrelated once time τ has elapsed, and the system typically transitions only after many attempts $m \gg 1$ at time $t = m\tau \gg \tau$, the distribution of waiting times $\rho[B|A](t)$ is Poissonian

$$\rho[B|A](t) = (1 - k\tau)^m k = (1 - kt/m)^m k \approx k\tau e^{-kt} \tag{1.2}$$

so the cumulative transition probability

$$P[B|A](t) = \int_0^t dt' k e^{-kt'} = 1 - e^{-kt} \approx kt, \quad \tau \leq t \ll k^{-1} \tag{1.3}$$

is linear in time on the appropriate timescale. The transition probability $P[B|A](t)$ can also be represented by a correlation function between regions of phase space A and B [1]

$$kt = \frac{\langle h_A(0) h_B(t) \rangle}{\langle h_A \rangle} \equiv \langle h_{B|A} \rangle, \quad \tau \leq t \ll k^{-1} \tag{1.4}$$

where the average is taken over stochastic paths of length t,

$$\langle h_A(0)h_B(t)\rangle = \int D[\mathbf{X}(t)]h_A(0)h_B(t)P[\mathbf{X}(t)] \tag{1.5}$$

is the joint probability of being in state A at time 0 and state B at time t, $\langle h_A\rangle$ is the probability to find the system in state A, and h_i denotes an indicator function

$$h_i(t) = \begin{cases} 1 & \text{if } \mathbf{x}(t) \in i \\ 0 & \text{else} \end{cases}. \tag{1.6}$$

In Fig. 1.2, we illustrate that the time derivative of the side-side correlation function encodes the rate by showing how it plateaus to the time independent rate constant.

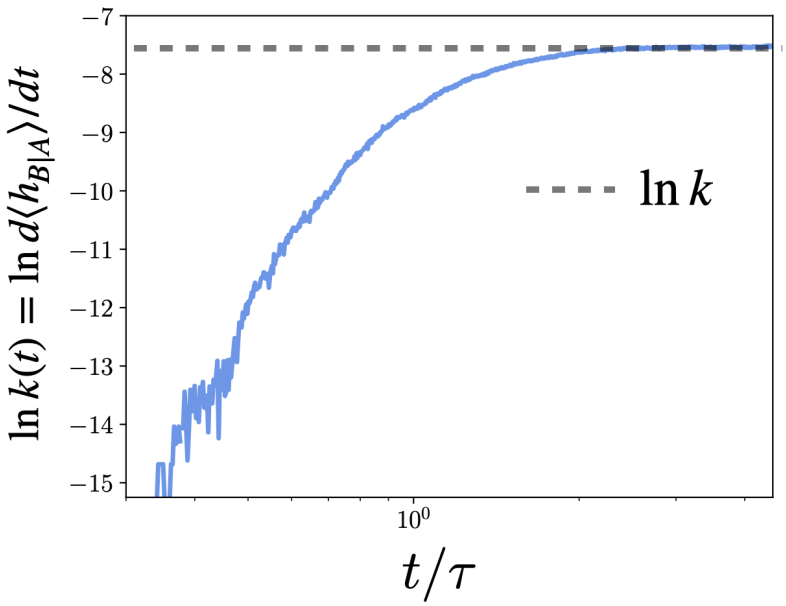

Figure 1.2: Transition rates are encoded in the temporal correlations between different areas, A and B, of configuration space. The side-side correlation function $\langle h_{B|A}\rangle = kt$ grows linearly in time with a proportionality constant equal to the transition rate k. We plot the log of the rate, $k(t) = d\langle h_{B|A}\rangle/dt$, which plateaus to a constant value for times $\tau \leq t \ll k^{-1}$ on the order of the transition path time τ.

Rare rates can also be described by the probability of traversing a so called transition state atop a barrier in free energy (negative log-probability, $-\ln p(x)$) between states A and B, as shown in Fig. 1.1. Integrating by parts, $\langle h_A(0)\dot{h}_B(t)\rangle = -\langle \dot{h}_A(0)h_B(t)\rangle$ since $\langle h_A(0)\dot{h}_B(0)\rangle = \langle \dot{h}_A(t)h_B(t)\rangle = 0$. Now using the chain rule $\dot{h}_A(0) = \dot{\mathbf{x}}(0)\nabla_{\mathbf{x}}h_A(\mathbf{x}) = -\dot{\mathbf{x}}(0)\delta(\mathbf{x}(0) - \mathbf{x}^{\ddagger})$, where the transition state $\mathbf{x}^{\ddagger}$ defines the border between A and B and

$\delta(\mathbf{x}(0) - \mathbf{x}^\ddagger)$ is the Dirac delta function. Multiplying and dividing by the transition state probability $\langle\delta(\mathbf{x} - \mathbf{x}^\ddagger)\rangle = p(\mathbf{x}^\ddagger)$, we arrive at transition state theory [2]

$$k(t) = \frac{\langle\dot{\mathbf{x}}(0)\delta(\mathbf{x}(0) - \mathbf{x}^\ddagger)h_B(t)\rangle}{2\langle\delta(\mathbf{x}(0) - \mathbf{x}^\ddagger)\rangle}\frac{\langle\delta(\mathbf{x}(0) - \mathbf{x}^\ddagger)\rangle}{\langle h_A\rangle} \equiv \frac{v(t)}{2}\frac{p(\mathbf{x}^\ddagger)}{\langle h_A\rangle}. \tag{1.7}$$

Here $v(t) = \langle\dot{\mathbf{x}}(0)h_B(t)\rangle_{\mathbf{x}(0)=\mathbf{x}^\ddagger}$ is the reactive flux through the transition state $\mathbf{x}^\ddagger$ at time $\tau \leq t \ll k^{-1}$, when it has become quasi-stationary, and $p(\mathbf{x}^\ddagger) \propto e^{-\Delta F(\mathbf{x}^\ddagger)}$ gives the expected Arrhenius rate behavior in equilibrium.

1.3 Stochastic path ensembles

As we just saw, we can access transition rates through the probability of finding a reactive path. Path ensembles have proven useful for their ability to make broad reaching general statements about fluctuations far from equilibrium. Under the Markovian assumption, [3] the probability of observing a stochastic path $\mathbf{X}(t)$ can be broken down into the probabilities of observing changes in individual conformations over independent time steps of length Δt

$$P[\mathbf{X}(t)] = P[\mathbf{x}(\Delta t), \mathbf{x}(2\Delta t), ..., \mathbf{x}(m\Delta t)] = p_0(\mathbf{x})\prod_m^{t/\Delta t} P[\mathbf{x}_{m+1}|\mathbf{x}_m] \tag{1.8}$$

which we label by m. Here $p_0(\mathbf{x})$ is the probability to observe an initial condition of configuration $\mathbf{x}$. To study systems with N particles in continuous phase space with dimensions labeled by α, we typically employ the Langevin equation

$$\dot{\mathbf{x}}^i = \mu^i\mathbf{F}^i_\lambda(\mathbf{x}, t) + \boldsymbol{\eta}^i, \quad 1 \leq i \leq N, \tag{1.9}$$

where μ^i is a vector of mobility coefficients and $\boldsymbol{\eta}^i$ is white noise with $\langle\boldsymbol{\eta}^i(t)\rangle = 0$ and $\langle\boldsymbol{\eta}^i(t)\boldsymbol{\eta}^{j\beta}(t')\rangle = 2D^i\delta_{ij}\delta_{\alpha\beta}\delta(t - t')$. The strength of the noise fluctuations is proportional to the particle diffusion $\langle\boldsymbol{\eta}^i(t)\boldsymbol{\eta}^i(t')\rangle \propto D^i = k_\mathrm{B}T\mu^i$, given by the Einstein relation.

The Langevin equation has discretized increments

$$\Delta\mathbf{x}^i(t) \equiv \mathbf{x}^i(t + \Delta t) - \mathbf{x}^i(t) = \mu^i\mathbf{F}^i_\lambda(t)\Delta t + \sqrt{2D^i\Delta t}\boldsymbol{\eta}^i, \tag{1.10}$$

where $\boldsymbol{\eta}^i$ is a vector of Gaussian random numbers with zero-mean and unit variance. Since we are dealing with white noise, we can write down an explicit expression for the probability of observing a single noise transition. The white noise is the only place randomness occurs in the equations of motion, so

$$-\ln P[\mathbf{x}^i_{m+1}|\mathbf{x}^i_m] + \text{const.} = \frac{(\boldsymbol{\eta}^i)^2}{2\langle\delta^2\boldsymbol{\eta}^i\rangle} = \Delta t\frac{(\Delta\mathbf{x}^i/\Delta t - \mu^i\mathbf{F}^i_\lambda)^2}{4D^i}. \tag{1.11}$$

In the limit of small time steps, the path probability

$$P[\mathbf{X}(t)] = p_0(\mathbf{x})e^{U_\lambda(t)}/Z_\lambda \tag{1.12}$$

can be written in terms of a path-action [4]

$$U_\lambda = \sum_i -\frac{1}{4D^i}\int_0^t dt'(\dot{\mathbf{x}}^i - \mu^i \mathbf{F}_\lambda^i)^2, \tag{1.13}$$

and a path partition function

$$Z_\lambda = \int D[\mathbf{X}] d\mathbf{x} p_0(\mathbf{x}) e^{U_\lambda(t)}. \tag{1.14}$$

To evaluate averages of path observables, we sum over paths with weight $P[\mathbf{X}(t)]$

$$\langle O \rangle(t) = \int D[\mathbf{X}(t)] O P[\mathbf{X}(t)]. \tag{1.15}$$

1.4 Trajectory reweighting

We often care about how an average, such as a side-side correlation function for example, changes when a force is added to or removed from a stochastic system. We can relate averages in the presence and absence of a driving force λ by reweighting stochastic trajectories by their path distributions. Inserting a factor of unity, and assuming the initial distributions are the same, we arrive at the reweighted path average [5]

$$\langle O \rangle_\lambda = \int D[\mathbf{X}(t)] O P_\lambda[\mathbf{X}(t)] = \int D[\mathbf{X}(t)] O P_0[\mathbf{X}(t)] \frac{P_\lambda[\mathbf{X}(t)]}{P_0[\mathbf{X}(t)]} \equiv \langle O e^{\Delta U_\lambda} \rangle_0. \tag{1.16}$$

Here

$$\Delta U_\lambda(t) = U_\lambda - U_0 = \sum_i -\frac{1}{4D^i}\int_0^t dt'(\mu^i \boldsymbol{\lambda}^i)^2 - 2\mu^i \boldsymbol{\lambda}^i(\dot{\mathbf{x}}^i - \mu^i \mathbf{F}_0^i) \tag{1.17}$$

is what we will refer to as the relative action.

1.5 Fluctuation theorems and second laws

Basic universal conclusions about the nonequilibrium nature of a system can be uncovered by studying how ΔU_λ fluctuates. When the observable in Eq. 1.16 is unity, $O = 1$,

$$1 = \langle e^{\Delta U_\lambda} \rangle_0 = \langle e^{-\Delta U_\lambda} \rangle_\lambda \tag{1.18}$$

by normalization of the path distribution. This fluctuation theorem [6, 7] in the second equality is illustrated in Fig. 1.3 for the bi-stable process shown in Fig. 1.1. Jensen's inequality (a convex function of a mean is less than the mean of the function) further tells us that

$$\langle \Delta U_\lambda \rangle_0 \leq 0 \leq \langle \Delta U_\lambda \rangle_\lambda. \tag{1.19}$$

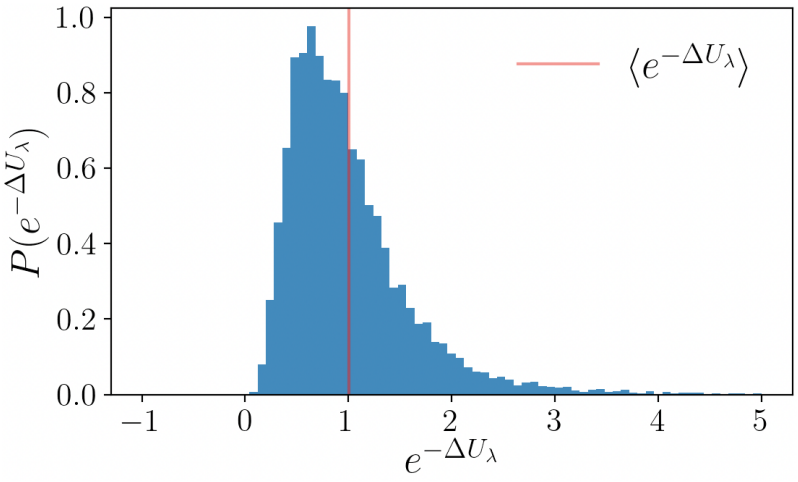

Figure 1.3: An illustration of the integral fluctuation theorem: in a double-well potential driven by an athermal active force, the average value of $e^{-\Delta U_\lambda}$ is unity for short-time trajectories $t \sim \tau$ started in state A (see Fig. 1.1).

When $\boldsymbol{\lambda}$ time reverses the unperturbed system, $\dot{\mathbf{x}}(t) \mapsto -\dot{\mathbf{x}}(t_f - t)$ and the relative action evaluates to the entropy production $-\beta(W - \Delta F)$, and we get back a statement about equilibrium free energy differences called Jarzynski equality [8]

$$\beta \Delta F = -\ln \langle e^{-\beta W} \rangle_F = \ln \langle e^{\beta W} \rangle_R$$

and the second law of thermodynamics

$$-\langle W \rangle_R \leq \Delta F \leq \langle W \rangle_F, \tag{1.20}$$

where

$$\Delta F = F_B - F_A, \quad F_{A,B} = k_B T \ln \int d\mathbf{x} e^{-\beta V_{A,B}(\mathbf{x})} \tag{1.21}$$

for equilibrium potentials $V_{A,B}$, and

$$W(t) = \int_0^t dt' \frac{\partial}{\partial t'} V(\mathbf{x}(t'), t'). \tag{1.22}$$

is the work done on the system due to driving. As we will see later on, O is only 1 sometimes, we can arrive at similar (second law type) expressions constraining transition rates.

1.6 Optimal control and gradients of noisy trajectory observables

Second law type statements like those in Eq. 1.20 are variational in the sense that to get the closest value to for instance ΔF by measuring nonequilibrium work, one should minimize

$\langle W \rangle_F$ To minimize the average value of a path observable $\langle O(t) \rangle_\theta$ over a set of control parameters θ, we represent the time-dependent external force as

$$\lambda(\mathbf{x}, t) = \sum_\mu \theta_\mu \lambda^\mu(\mathbf{x}, t) \tag{1.23}$$

with parameters θ and basis functions $\lambda^\mu(\mathbf{x}, t)$ [9]. Once the control force is parameterized, we can perform gradient descent

$$\theta^{i+1} = \theta^i + l \nabla_\theta \langle O(t) \rangle_\theta \tag{1.24}$$

where here l is a learning rate, to iteratively update the parameters. Altering the parameters changes the distribution of paths in phase space, so parameter gradients of average observables

$$\partial_\theta \langle O \rangle_\theta = \langle \partial_\theta O \rangle_\theta + \langle O \partial_\theta \ln P[\mathbf{X}(t)] \rangle_\theta \equiv \langle \partial_\theta O \rangle_\theta + \langle O q_\theta \rangle_\theta \tag{1.25}$$

can be taken by keeping track of so called Malliavin weights q_μ [10, 11, 9].

For the systems we consider here

$$\dot{\mathbf{x}}^i = \mu^i \mathbf{F}^i_\theta(\mathbf{x}, t) + \boldsymbol{\eta}^i \tag{1.26}$$

the path distribution is

$$\ln P[\mathbf{X}(t)] = \sum_i -\frac{1}{4D^i} \int_0^t dt' (\dot{\mathbf{x}}^i - \mu^i \mathbf{F}^i_\theta)^2 - \ln Z_\theta \tag{1.27}$$

where Z_θ is a partition function. Differentiating the above equation with respect to θ yields [11]

$$q_\theta(t) = y_\theta(t) - \langle y_\theta(t) \rangle, \quad y_\theta(t) = \int_0^t dt' \frac{\beta}{2} \sum_i \frac{\boldsymbol{\eta}^i(t')}{\Delta t} \partial_\theta \mathbf{F}^i_\theta. \tag{1.28}$$

By keeping track of q_θ for each control parameter, we can minimize the expectation of any observable. We note that since the noise is uncorrelated to the force at the same timestep, $\langle y_\theta(t) \rangle \to 0$ in the infinite trajectory limit.[12]

Chapter 2

Bringing the heat: dissipation sets a speed limit on driven transitions

This work was done in a collaboration with David Limmer that resulted in the publication "Dissipation bounds the amplification of transition rates far from equilibrium." [13]

Dissipation is the energetic cost of maintaining nonequilibrium. As otherwise static systems are injected with energy from external forces, they lose a portion of that energy to the thermal bath around them through dissipation. Using trajectory reweighting, we link dissipation to how transition rates are enhanced or suppressed under time dependent external force, and find that this fundamental cost sets a limit on transition rate response.

Complex systems can convert energy imparted by nonequilibrium forces to regulate how quickly they transition between long lived states. While such behavior is ubiquitous in natural and synthetic systems, currently there is no general framework to relate the enhancement of a transition rate to the energy dissipated, or to bound the enhancement achievable for a given energy expenditure. We employ recent advances in statistical thermodynamics to build such a framework, which can be used to gain mechanistic insight on transitions far from equilibrium. We show that under general conditions, there is a basic speed-limit relating the typical excess heat dissipated throughout a transition and the rate amplification achievable. We illustrate this trade-off in canonical examples of diffusive barrier crossings in systems driven with autonomous and deterministic external forcing protocols. In both cases, we find that our speed limit tightly constrains the rate enhancement.

The natural world is full of systems in which the rate of a rare dynamical event is enhanced through coupling to a dissipative process. [14, 15] *In vivo*, molecular chaperones accelerate protein folding and assembly so that otherwise slow transitions occur on biologically relevant timescales, at the energetic cost of maintaining chemical potential gradients [16]. Shear forces drive colloidal assemblies and polymer films to order rapidly enough for viable synthesis, at

the expense of applying external forces [17]. Such behavior is leveraged across physical and biological systems, but there are few known principles available to act as guides or constrain possibilities. Here we use nonequilibrium stochastic thermodynamics to demonstrate that dissipation bounds the enhancement of the rate of a transition away from equilibrium. The bound is sharp near equilibrium and for large barriers, holds arbitrarily far from equilibrium, and can be tightened with additional knowledge of kinetic factors. Our work thus elucidates a fundamental trade-off between speed and energy consumption.

In equilibrium, the rate of a transition between two long lived states is determined by the likelihood that a thermal fluctuation provides sufficient energy to the system to overcome a free energy barrier. Away from equilibrium, external forces and nonthermal fluctuations can mitigate this constraint, modulating the rate relative to its equilibrium value. Departures from thermal equilibrium make it difficult to predict the extent to which a dissipative process can influence a transition, as traditional rate theories are grounded in equilibrium statistical mechanics. For instance, both classical transition state theory [2] and Kramer's theory [18] require information on the probability to reach a rare dividing surface, or transition state. In equilibrium the Boltzmann distribution supplies that probability, but within a nonequilibrium steady-state that information is generally unavailable. Freidlin-Wentzell theory [19], and transition path theory [1] supply formal means of estimating rates away from equilibrium through the consideration of path ensembles. However, rate calculations within these formalisms require complex optimizations or partition function evaluations, and do not encode simple relationships between rates and other measurable quantities.

Using principles from stochastic thermodynamics, we develop a general theory of nonequilibrium rate enhancement, deriving exact relations and fundamental bounds[3]. Stochastic thermodynamics has supplied a number of relationships that constrain fluctuations away from equilibrium in terms of measurable energetic observables [8]. The fluctuation theorems illustrate fundamental time-reversal symmetries [6], and thermodynamic uncertainty relations bound response [20]. In this work, we show that the rate enhancement achievable away from equilibrium is bounded by the heat dissipated over the course of the transition,

$$\frac{k_{\rm neq}}{k_{\rm eq}} \leq e^{\beta \bar{\mathcal{Q}}/2}, \tag{2.1}$$

where $k_{\rm neq}/k_{\rm eq}$ is the ratio of the nonequilibrium to equilibrium transition rates, and deviation from equilibrium due to broken detailed balance is codified by $\bar{\mathcal{Q}}$, the average excess heat released over the transition due to the nonequilibrium process, in units of $k_{\rm B}T$, where $k_{\rm B}$ is Boltzmann's constant and T the temperature of the bath. Our theory demonstrates that the rate enhancement achievable by coupling a system to a dissipative process, an essential dynamical quantity, is limited by general thermodynamic constraints. To test the theory, we study paradigmatic two-state continuous force systems, driving them from equilibrium with both deterministic and autonomous forces.

2.1 Stochastic thermodynamics of rate enhancement

To derive Eq. 2.1, we consider systems driven by a time-dependent force, $\lambda(t)$, either externally controlled or coupled to an additional nonthermal noise-source that evolves independently of the system-state, precluding feedback. Extensions to systems evolving in boundary driven nonequilibrium steady-states, though likely possible, are not explored in this work.

In the presence of the time-dependent force, the rate, k_λ, of transition between two long-lived states, is the probability that a transition occurs per unit time. For a system described by a configuration, $\mathbf{x}(t)$, at time t, we will consider initial and final states, A and B, that are collections of configurations defined by the indicator functions,

$$h_i(t) = \begin{cases} 1 & \text{if } \mathbf{x}(t) \in i \\ 0 & \text{else} \end{cases} \tag{2.2}$$

where $i \in \{A, B\}$, and we assume A and B are not intersecting. For times longer than the characteristic local relaxation time and much shorter than the inverse rate, k_λ derives from a ratio of path partition functions,

$$k_\lambda(A \to B) = \frac{d}{dt}\frac{Z_{AB}(\lambda)}{Z_A(\lambda)}. \tag{2.3}$$

Here,

$$Z_{AB}(\lambda) = \int D[\mathbf{X}(t)] h_A(0) h_{\mathrm{B}}(t) P_\lambda[\mathbf{X}(t)] \tag{2.4}$$

is the number of transition paths, $\mathbf{X}(t) = \{\mathbf{x}(0), ..., \mathbf{x}(t)\}$, starting in A and ending in B at time t, weighted with probability $P_\lambda[\mathbf{X}(t)]$, and

$$Z_A(\lambda) = \int D[\mathbf{X}(t)] h_A(0) P_\lambda[\mathbf{X}(t)] \tag{2.5}$$

is the corresponding number of paths starting in A [1].

The ratio in Eq. 2.3 is simply the conditional probability of the system being in state B given it started in A. Provided the transition is rare, consistent with A and B representing metastable states, there is a range of time over which $Z_{AB}(\lambda)$ increases linearly, and k_λ is constant. Specifically, the rate constant is defined for observation times $\tau_A \lesssim t \ll k_\lambda^{-1}$, where the transition path time is typically on the order of τ_{A}, the characteristic relaxation time within state A, and shorter than the timescale required for global relaxation. The probability of a path is the product of a distribution of initial conditions, $\rho_\lambda[\mathbf{x}(0)]$, and the conditional transition probability $P_\lambda[\mathbf{X}(t)|\mathbf{x}(0)]$, such that $P_\lambda[\mathbf{X}(t)] = P_\lambda[\mathbf{X}(t)|\mathbf{x}(0)]\rho_\lambda[\mathbf{x}(0)]$. While in general away from thermal equilibrium, $\rho_\lambda[\mathbf{x}(0)]$ is unknown, $P_\lambda[\mathbf{X}(t)|\mathbf{x}(0)]$ can be inferred, provided an equation of motion. For the specific model calculations discussed below, $P_\lambda[\mathbf{X}(t)|\mathbf{x}(0)]$ will take an Onsager-Machlup form[4].

Stochastic thermodynamics gives structure to path ensembles and relations to thermodynamic quantities. In an equilibrium system, the principle of microscopic reversibility implies

that the probability of a trajectory is equal to its time-reverse. Specifically, let $\tilde{P}_\lambda[\tilde{\mathbf{X}}(t)]$ denote the probability of observing a time-reversed trajectory $\tilde{\mathbf{X}}(t) = \{\tilde{\mathbf{x}}(t), ..., \tilde{\mathbf{x}}(0)\}$, where $\tilde{\mathbf{x}}(t)$ is a time-reversed configuration of the system at t, labeled in the forward time direction. In the absence of the dissipative protocol, $\lambda = 0$, the system is in equilibrium and $P_0[\mathbf{X}(t)] = \tilde{P}_0[\tilde{\mathbf{X}}(t)]$. The Crooks fluctuation theorem extends this notion to systems driven away from equilibrium by an arbitrary time dependent force $\lambda(t)$[6]. For a nonequilibrium system, microscopic reversibility is manifested by $P_\lambda[\mathbf{X}(t)|\mathbf{x}(0)] = \tilde{P}_\lambda[\tilde{\mathbf{X}}(t)|\mathbf{x}(t)] \exp[\beta(Q[\mathbf{X}(t)|\mathbf{x}(0)] + Q_{\rm rev}[(\mathbf{x}(t), \mathbf{x}(0)])]$. The first term in the exponential, βQ, is what we refer to as *dissipation*, as it is the excess heat transferred from the system to the bath along a trajectory driven from equilibrium over the corresponding heat transferred for a reversible process. For an equilibrium system, the reversible contribution to the heat is equal to the Shanon entropy, $\beta Q_{\rm rev} = \ln \rho_0[\mathbf{x}(t)]/\rho_0[\mathbf{x}(0)]$. It is generally derivable as the change in energy due to the conservative forces, and thus depends on only the trajectory's boundaries.

Coupling the system to a dissipative process will generally change its dynamics. Using trajectory reweighting, we relate the transition rate in the presence and absence of the nonequilibrium force $\lambda(t)$. We consider two path probability distributions with support on the same $\mathbf{X}(t)$, so that the relative action

$$\beta \Delta U_\lambda[\mathbf{X}(t)|\mathbf{x}(0)] = \ln \frac{P_\lambda[\mathbf{X}(t)|\mathbf{x}(0)]}{P_0[\mathbf{X}(t)|\mathbf{x}(0)]}\,, \tag{2.6}$$

relating one to the other, is well-defined. Performing a change of measure, for a constant distribution of initial conditions, we express ratios of path partition functions in either ensemble as

$$\frac{Z_{AB}(\lambda)}{Z_{AB}(0)} = \left\langle e^{\beta \Delta U_\lambda} \right\rangle_0 = \left\langle e^{-\beta \Delta U_\lambda} \right\rangle_\lambda^{-1}\,, \tag{2.7}$$

where, unlike in chapter 1, the *brackets denote a conditional average in a transition path ensemble connecting states A and B in time t*, with path probability $P_0[\mathbf{X}(t)]$ in the first equality, or $P_\lambda[\mathbf{X}(t)]$ in the second equality.

When transitions in both path ensembles are rare, $k_\lambda t$ and $k_0 t \ll 1$, the overwhelming majority of paths originating from A will remain there on the timescales where the rate is time independent, so that $Z_A(\lambda) = Z_A(0)$. In Sec. S1., we consider generalizations away from this limit, showing that the ratio $Z_A(0)/Z_A(\lambda)$ cancels contributions in Eq. 2.7 due to different distributions of initial conditions. Thus, under mild assumptions, combining Eq. 2.3 with Eq. 2.7, we find

$$\frac{k_\lambda}{k_0} = \left\langle e^{\beta \Delta U_\lambda} \right\rangle_0 = \left\langle e^{-\beta \Delta U_\lambda} \right\rangle_\lambda^{-1} \tag{2.8}$$

which is an exact relation between transition rates in the presence or absence of the dissipative process. Lower and upper bounds can be read off by applying Jensen's inequality to each of these expressions,

$$\beta \left\langle \Delta U_\lambda \right\rangle_0 \leq \ln \frac{k_\lambda}{k_0} \leq \beta \left\langle \Delta U_\lambda \right\rangle_\lambda \tag{2.9}$$

constituting a fundamental envelope for the rate enhancement. This result is general provided any two trajectory ensembles have common support. In a suitably defined linear response regime, the ensembles are approximately equal, $\langle \Delta U_\lambda \rangle_\lambda \approx \langle \Delta U_\lambda \rangle_0$, so the bounds are saturated. This corresponds to a near-equilibrium regime where driving is small. While Eq. 2.9 is derived for constant initial conditions for simplicity of notation, the impact of different initial conditions on an upper bound is to add a positive constant equal to a symmetrized Kullback–Leibler divergence between initial distributions in the driven and equilibrium ensembles (Sec. S1.). In the linear response regime, and when transition paths are long enough to lose memory, a common occurrence for rare transitions which quickly relax in both A and B, changes to initial conditions can be neglected.

Generally, ΔU_λ contains thermodynamic and kinetic factors. To separate them, we decompose ΔU_λ into time-reversal-symmetric and asymmetric trajectory observables,

$$\Delta U_\lambda = \frac{Q + \Gamma}{2}, \tag{2.10}$$

where the excess heat $Q = \Delta U_\lambda[\mathbf{X}(t)|\mathbf{x}(0)] - \Delta U_\lambda[\tilde{\mathbf{X}}(t)|\tilde{\mathbf{x}}(0)]$ is odd under time-reversal, and the excess dynamical activity $\Gamma = \Delta U_\lambda[\mathbf{X}(t)|\mathbf{x}(0)] + \Delta U_\lambda[\tilde{\mathbf{X}}(t)|\tilde{\mathbf{x}}(0)]$ is even. On the whole, both the heat and the activity play important roles in response and stability of nonequilibrium systems [7, 5]. While the heat has a simple mechanical definition and is largely independent of the system's dynamics, the activity depends on details of the equation of motion, making it hard to generalize.

We find that for rare transitions across a host of physically relevant conditions, the activity can be neglected. Near equilibrium, the average activity in the conditioned transition ensemble vanishes due to time-reversal symmetry. In cases of instantonic transitions, where the driving force varies slowly relative to the characteristic transition path time the activity is small. Even away from limiting cases where the barrier is much larger than the scale of the noise, the activity can be neglected. Consider for instance free diffusion, where rare transitions are largely noise assisted. In these sojourns over broad, diffusive regions, the activity is strictly negative. By neglecting it, a bound based on the heat alone is satisfied though weakened. Each case is considered explicitly in Sec. S2. Remarkably, this implies that the dissipation accumulated over a transition bounds the rate enhancement,

$$\ln \frac{k_\lambda}{k_0} \leq \frac{\beta}{2} \langle Q \rangle_\lambda \,, \tag{2.11}$$

which is our main result. Identifying the system under finite λ as a nonequilibrium system, and its absence as an equilibrium one, we identify Eq. 2.11 as a more precise statement of Eq. 2.1. We note, however that the rate enhancement relation is general for any two transition path ensembles.

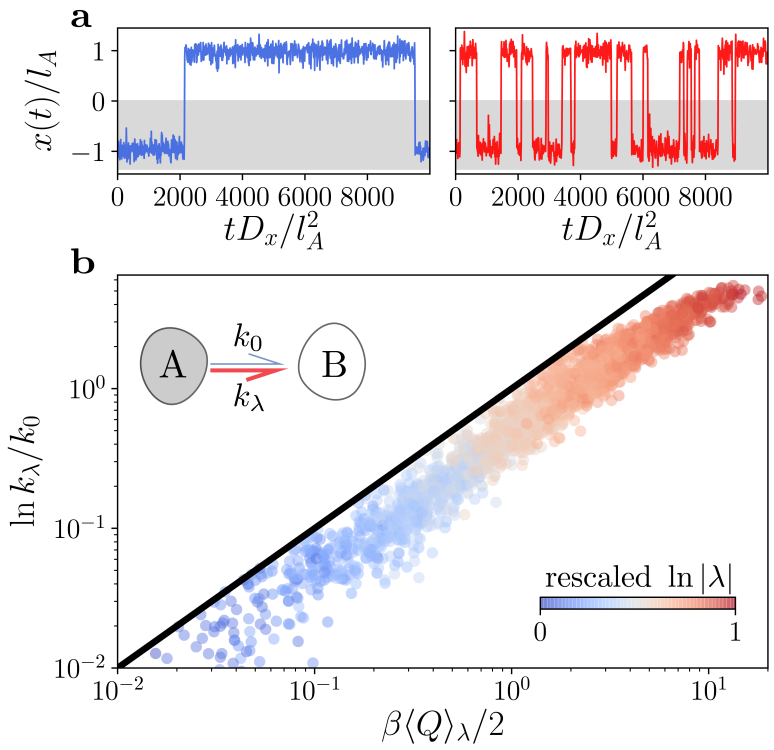

Figure 2.1: Nonequilibrium driving enhances transition rates. **a**, Trajectories of a two-state system in equilibrium (left) and the same process driven away from equilibrium (right) by a time-dependent external force. **b**, Rate enhancement for different asymmetric two-state systems as a function of dissipated heat, each driven by a randomly chosen combination of deterministic and stochastic external forces. The bound in Eq. 2.11 is shown as a black line and points are colored according to the magnitude of the force, low (blue) to high (red).

2.2 Autonomous and deterministic forcing

To illustrate the robustness of our dissipative bound, we first consider the overdamped dynamics of a particle in a one-dimensional asymmetric potential subject to both external time dependent and nonthermal forces. Specifically, the equation of motion for the position of the particle, x, is taken as $\gamma\dot{x} = -\partial_x V(x) + \lambda(t) + \sqrt{2k_\mathrm{B}T\gamma}\eta_x$ where γ is the friction due to the surrounding medium, imposing a diffusion constant $D_x = k_\mathrm{B}T/\gamma$, and η_x is a Gaussian random variable with $\langle\eta_x(t)\rangle = 0$ and $\langle\eta_x(t)\eta_x(t')\rangle = \delta(t-t')$. The static external potential consists of two quartic states, $V(x) = V_A(x)\Theta(-x) + V_B(x)\Theta(x)$. Each basin $V_i(x) = (\Delta V_i x^2/(2l_i^2))(x^2/(2l_i^2) - 1)$, $i \in \{A, B\}$, is characterized by the distance of its minimum $(|x|, |y|) = (l_i, \Delta V_i)$ to the origin, where the states are joined by $\Theta(x)$, the Heaviside function. $V(x)$ supports two metastable states with a barrier between them if

$\beta\Delta V_i > 1$ for both $i = A$ and B. Figure. 2.6a), where $x(t)$ exhibits fluctuations concentrated around two regions of the potential, with few, fleeting transitions between them, manifests this metastability.

We drive the system out of equilibrium according to a time-dependent protocol $\lambda(t) = f[p\cos(\omega t) + (1-p)\cos(\theta(t))]$ with maximum amplitude f partitioned, $p \in [0,1]$, into deterministic and autonomous components. The deterministic portion of the driving is periodic with frequency ω, whereas the autonomous piece is determined by an additional nonthermal process, $\dot{\theta}(t) = \sqrt{2D_\theta}\eta_\theta$, with diffusion constant D_θ, and delta-correlated white noise, $\langle\eta_\theta(t)\rangle = 0$ and $\langle\eta_\theta(t)\eta_\theta(t')\rangle = \delta(t-t')$.

Considering transitions that take the particle from one side of the potential to the other, we define $h_A = \Theta(x + l_A/\sqrt{3})$ and $h_B = \Theta(x - l_B/\sqrt{3})$, which correspond to the locations of the maximum force opposing the transition in the absence of $\lambda(t)$. The dissipated excess heat can be computed from

$$Q(t) = \int_0^t dt'\lambda(t')\dot{x}(t'), \tag{2.12}$$

and its mean estimated within the nonequilibrium steady-state by integrating the dissipation rate over reactive trajectories of length t, given by the typical transition path time as discussed in *Methods* and Sec. S3. Given a suitable separation between inverse rate and relaxation time, Q is often insensitive to the precise value of t. We reiterate that Q differs from the total heat by the conservative boundary portion which appears in both the equilibrium and driven actions, and therefore does not play a role in the rate enchancement.

Figure 2.6b) shows the results of 3000 randomly constructed models, where $\Delta V_i, l_i$, p, f, and D_θ, were chosen uniformly over a wide range of parameters, as detailed in *Methods*. For each model, k_λ, k_0 and $\langle Q\rangle_\lambda$ have been independently evaluated, and Fig. 2.6b) demonstrates that the bound holds across the broad parameter-space. Points are colored blue to red in increasing magnitude of the driving force f, showing that the protocol most efficiently amplifies the equilibrium rate when $\beta\langle Q\rangle_\lambda < 10$. Pushing past this regime, the bound becomes progressively weaker, as dissipation increases, but rate enhancement reaches a plateau. This corresponds to a limit where driving is large enough to degrade the assumption that basin A is metastable.

In order to understand the physical processes that determine whether or not the bound is saturated, we focus on two cases of the model presented above. First, we set $p = 0$, which corresponds to an active Brownian particle in an external potential. Active Brownian particles provide a canonical realization of how autonomous athermal noise can drive novel steady-states without simply imparting an effective temperature[21]. These self-propelled agents exhibit dynamical symmetry breaking and collective motion [22, 23], and previous studies have shown that the escape of active particles from a metastable potential exhibits interesting behavior arising from an interplay between the driving force, persistence time statistics, and the shape of the potential[24, 25]. For simplicity we take a symmetric potential, with $l_A = l_B = 1$ and $\beta\Delta V_A = \beta\Delta V_B = 10$, setting $\gamma = 1$ and $k_\mathrm{B}T = 1/2$.

Figure 2.5a) shows the dependence of the rate enhancement on the rotational diffusivity,

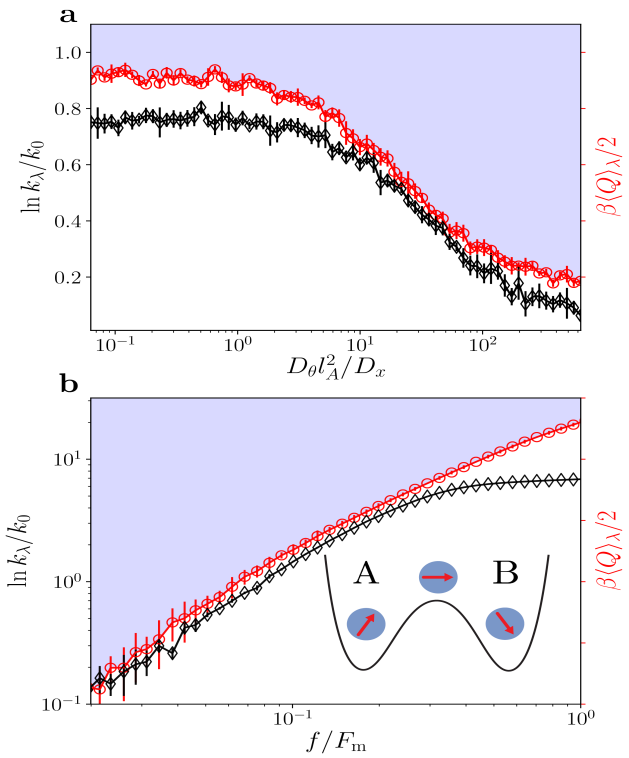

Figure 2.2: Rate enhancement for an active Brownian particle. **a**, Rate enhancement as a function of rotational diffusivity constant $D_\theta l_A^2/D_x$ for $f = 1$. **b**, Rate enhancement as a function of the magnitude of active driving relative to the maximum force opposing the transition in equilibrium f/F_m for $D_\theta = 1/2$. In both, the rate enhancement (black diamonds) is bounded by the dissipated heat (red circles).

D_θ at fixed $f = 1$. The small D_θ limit is a quasi-stationary regime corresponding to an equilibrium system with an additional linear force added to the potential.

Increasing the rotational diffusivity decreases the persistence of the driving, and as a consequence $\langle Q\rangle_\lambda$ and k_λ/k_0 fall off in this limit. In the large D_θ limit, the system is effectively in equilibrium at an elevated temperature, as λ averages to 0. Lower rate enhancement and little dissipation are observed across this range of D_θ, and our bound is uniformly close. These results are consistent with a recent study in which an effective potential approach was used to derive k_λ/k_0 for an active Ornstein–Uhlenbeck process in a cubic well[26]. As shown in Fig. 2.5b), our bound is closest to the true rate enhancement when driving is small compared to the maximum force needed to surmount the barrier in equilibrium, $f < F_{\rm m}$ where $F_m = 8\Delta V_A/3\sqrt{3l_A}$. In that regime, the heat and rate enhancement both scale with f^2, as predicted by linear response theory. When the protocol and gradient forces are comparable, the transition ceases being a rare event and further increasing f has little effect

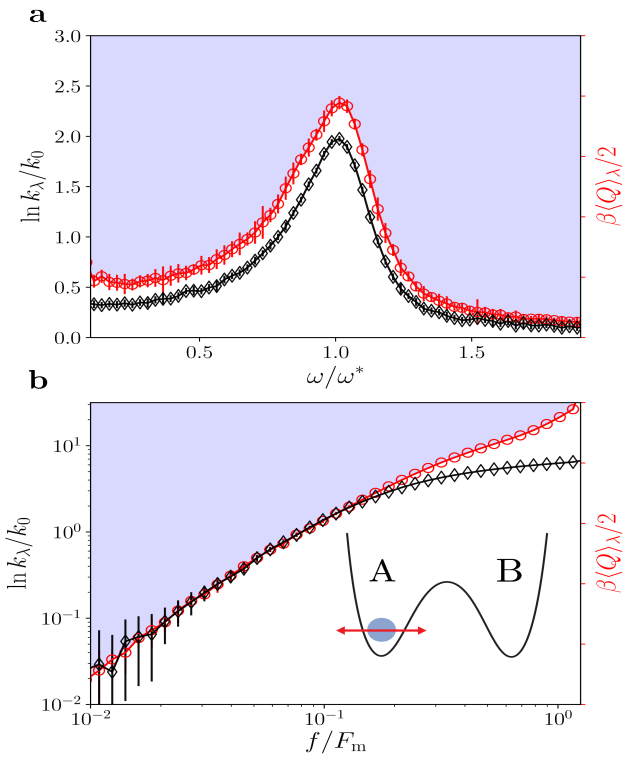

Figure 2.3: Rate enhancement in the Duffing oscillator. **a**, Rate enhancement as a function of the magnitude of driving frequency relative to the natural frequency of the equilibrium system, ω/ω^*, for $f = 1.4$. **b**, Rate enhancement as a function of the magnitude of function of the periodic driving relative to the maximum force opposing the transition in equilibrium f/F_m for $\omega/\omega^* = 1$ at $f = 1.4$. In both, the rate enhancement (black diamonds) is bounded by the dissipated heat (red circles).

on the rate, but increases the heat.

As a second test case, we consider underdamped dynamics with time-periodic driving, $p = 1$. This model, known as the Duffing oscillator,[27] is the simplest model of a stochastic pump and one whose nonequilibrium behavior is marked by significant nonlinearity. As an underdamped process, its barrier crossing behavior is determined both by spatial as well as energy diffusion, in which both position and velocity correlations play a role. The equation of motion is given by $m\ddot{x} = -\gamma\dot{x} - \partial_x V(x) + \lambda(t) + \sqrt{2k_{\rm B}T}\eta_x$, where the mass m reflects the change to underdamped dynamics. Again we take a symmetric potential, $l_A = l_B = 1$, now with $\beta\Delta V_A = \beta\Delta V_B = 14$ and $m = \beta = \gamma = D_x = 1$.

Figure 2.4a) shows that for a moderate force, $f/F_{\rm m} \approx 0.13$, there is an optimal driving frequency, denoted here as ω^*, which greatly enhances the transition rate. This phenomenon is known as stochastic resonance[27]. For slow driving, $\omega \ll \omega^*$, the particle typically

makes a transition before the external force reaches its maximum. For $\omega \gg \omega^*$ driving is inefficient, and on average requires multiple forcing cycles before presenting a chance to cross the barrier with the help of a positive force within the time of a typical transition. The approximate shape of the rate enhancement profile is Lorentzian, a trait inherited from the absorption lineshape of an underdamped harmonic oscillator. In this case, the resonant frequency coincides with the curvature of the equilibrium double-well potential driven with a quasi-static force, $\omega^* \approx \sqrt{8(\Delta V - f) - \gamma^2/4}$[28]. We find near saturation of the bound throughout a wide range of frequencies and across even such nonlinear behavior as stochastic resonance. In Fig. 2.4b), we plot $k_{\rm neq}/k_{\rm eq}$ and $\beta\langle Q\rangle_\lambda/2$ against the driving amplitude relative to $F_{\rm m}$. As in the other examples, the bound on rate enhancement is tight so long as the metastability of state A is preserved.

2.3 Additional bounds on ratios and differences of transition rates

Our main results were obtained under the assumption that the probability of starting in A does not change much under the influence of driving $Z_A(\lambda) \approx Z_A(0)$. In this section, we derive a more general bound that relaxes this assumption, and discuss in what cases we expect initial conditions to play a significant role. The factor $Z_A(0)/Z_A(\lambda)$ in question is the ratio of single-time probabilities and has played an important role in the development of Monte-Carlo sampling on the space of driving protocols.[29] As with the transition path partition function, it is still possible to express such a ratio in terms of the moment generating function of the relative action

$$\frac{Z_A(0)}{Z_A(\lambda)} = \left\langle \frac{\rho_0}{\rho_\lambda} e^{-\beta\Delta U_\lambda} \right\rangle_{A,\lambda} = \left\langle \frac{\rho_\lambda}{\rho_0} e^{\beta\Delta U_\lambda} \right\rangle_{A,0}^{-1}, \tag{2.13}$$

which implies

$$-\beta\langle\Delta U_\lambda\rangle_{A,\lambda} - D_{\rm KL}(\rho_\lambda||\rho_0) \leq \ln\frac{Z_A(0)}{Z_A(\lambda)} \leq -\beta\langle\Delta U_\lambda\rangle_{A,0} + D_{\rm KL}(\rho_\lambda||\rho_0) \tag{2.14}$$

where the paths of any length are conditioned to start in state A. After restoring initial conditions, Eq. 9 of the main text then implies

$$\beta(\langle\Delta U_\lambda\rangle_0 - \langle\Delta U_\lambda\rangle_{A,\lambda}) - J(\rho_\lambda||\rho_0) \leq \ln\frac{k_\lambda}{k_0} \leq \beta(\langle\Delta U_\lambda\rangle_\lambda - \langle\Delta U_\lambda\rangle_{A,0}) + J(\rho_\lambda||\rho_0), \tag{2.15}$$

which is our most general result. Initial conditions $\rho_\lambda(x_0)$ and $\rho_0(x_0)$ enter in through a symmetric variant of the Kullback–Leibler divergence known as the Jeffreys divergence,

$$J(\rho_\lambda||\rho_0) = D_{KL}(\rho_\lambda||\rho_0) + D_{KL}(\rho_0||\rho_\lambda) = \int dx_0(\rho_\lambda - \rho_0)\ln\frac{\rho_\lambda}{\rho_0} \geq 0, \tag{2.16}$$

and we will discuss shortly under what conditions it can be neglected. Even when its most likely state changes, if A remains metastable, the system will start in an effective stationary state with flux and force that both vanish on average, leaving the negative, λ^2 piece of the dynamical activity as the only portion of $\langle \Delta U_\lambda \rangle_{A,0}$ left. However, this piece does not depend on the trajectory since under our assumptions, the external force is independent of the system configuration, so it cancels in the relative action with an identical contribution from the driven ensemble. We can therefore deduce that only the trajectory dependent portions of $\langle \Delta U_\lambda \rangle_\lambda$, and $J(\rho_\lambda || \rho_\lambda)$ are likely to contribute to our upper bound.

To understand the influence of initial conditions, we begin by noting that they cease to enter into the bound when the initial configuration is deterministic or the system is prepared according to the same distribution, either ρ_λ or ρ_0. Next, that ρ_λ and ρ_0 are equivalent in linear response means $J(\rho_\lambda || \rho_\lambda)$ vanishes in this regime. Thus, if making the transition is correlated with the initial distribution in state A, then whenever that correlation remains relatively invariant under driving, initial conditions can be neglected. Perhaps most relevant to the present work and the study of rare events is the situation when initial conditions becoming uncorrelated with the 'bulk' of the transition path, and are forgotten. The reason why loosing memory of initial conditions is far from exotic for such reactive trajectories stems directly from the fact that state A is assumed, by definition of the rate constant, to be metastable, harboring a local minima of the free energy landscape. The initial portion of a typical rare event transition path from A to B will therefore relax exponentially quickly in this local basin and loose memory over a short timescale τ_A determined by the curvature of the basin. Since the rate constant is defined for times that are order τ_A, as detailed in the paragraph following Eq. 5 in the main text, the brunt of reactive paths are marked by boundary portions of rapid relaxation on either side of the transition region. When state A is suitably defined so that this memory loss occurs in both the driven and reference ensembles, all that matters is that the particle starts somewhere in state A, which is guaranteed by the transition portion of the path probability. In other words, knowledge about the character of state A is contained directly in $\langle \Delta U_\lambda \rangle_\lambda$ when both rate constants are defined with a common observation time, which is the topic of Sec. S3. When the driving is so extreme that a common observation time ensuring relaxation before escaping state A cannot be obtained, the rate constant effectively looses meaning. Such atypical cases can be addressed taking into account initial conformational distributions, or preferably, redefining A so that k_0 and k_λ are comparable by memory-less transition paths of the same length.

The recently proposed dissipation-time uncertainty relation [30] claimed that the rate of steady state entropy production bounds the difference of forward to backward, denoted here with a tilde, transition rates, out of and into a metastable state A, from above

$$k_\lambda - \tilde{k}_\lambda \leq \beta \langle \dot{Q} \rangle_\lambda. \tag{2.17}$$

In this section, we follow the same line of thinking, but use the definition of ΔU in place of the traditional fluctuation theorem. For a *two-state system*, the indicator functions (Main: Eq. 2) defining states A and B are simply related by $h_A(t) + h_B(t) = 1$. In this case, the

cumulative probability $p(t)$ that no transitions occur up to time t, the survival probability, is related to the probability that at least one occurs in the same way:

$$p_\lambda(t) = \frac{Z_{AA}(\lambda,t)}{Z_A} = 1 - \frac{Z_{AB}(\lambda,t)}{Z_A}. \tag{2.18}$$

From Eq. 1. in the main text, this relation implies that the transition rate is given by $k_\lambda = -d\ln p_\lambda(t)/dt$. Bounding $p_\lambda(t)/p_0(t)$ from above and below in the AA path ensemble, and subsequently using the fact that $h_A(t) = 1 - h_B(t)$ to convert to the AB ensemble yields the envelope presented in Eq. 13. For example,

$$\frac{p_0}{p_\lambda} = \frac{Z_{AA}(0,t)}{Z_{AA}(\lambda,t)} = \langle e^{-\beta\Delta U}\rangle_{AA,\lambda} \geq e^{-\beta\langle\Delta U\rangle_{AA,\lambda}} = e^{-\beta(1-\langle\Delta U\rangle_{AB,\lambda})} \implies k_\lambda - k_0 \geq \beta\langle\Delta\dot{U}_\lambda\rangle_{AB,\lambda} \tag{2.19}$$

yields an upper bound, where, to be explicit, we label averages by which path ensemble they belong to (AA and AB). Again, bounds in Eq. 13 assume the meta-stability of the A state is sufficiently preserved, though this can be relaxed by following the procedure laid out in SM section 3. We leave it to future work to delve into the utility of bounding the change of the rate $k_\lambda - k_0$ in this fashion.

2.4 Instantonic, slow driving and near equilibrium limits

In order to understand when the contribution to the dynamical activity to the rate enhancement bound can be neglected, we consider both simplified limiting cases that are analytically tractable, as well as additional numerical experiments on the systems considered in the main text. In general, the dissipative bound is valid in cases where the activity can be neglected due to its size or if it can be shown to be strictly negative. The former case can be argued to occur near-equilibrium or when specific spatial symmetries exist that result in its average being small. The latter case occurs when a particle traverses both large, narrow barriers as well as broad diffusive barriers.

Limiting cases of high and broad diffusive barriers

For simplicity and concreteness, we consider an overdamped particle in a $1d$ symmetric double well potential. The equation of motion is analogous to that considered in the first two examples in the main text,

$$\gamma\dot{x}(t) = F_\lambda(x) + \sqrt{2k_\mathrm{B}T\gamma}\eta, \qquad \langle\eta(t)\rangle = 0 \qquad \langle\eta(t)\eta(t')\rangle = \delta(t-t'), \tag{2.20}$$

where $F_\lambda(x) = -V'(x) + \lambda(t)$ is the total force, including the gradient and non-gradient contributions. From the equation of motion, it is straight-forward to show that the conditioned

transition probability takes the form $\ln P_\lambda[x(t)|x(0)] = \beta U_\lambda[x(t)] + C$, where C is a constant and the Onsager-Machlup path-action is

$$U_\lambda = -\frac{1}{4\gamma} dt' \, \{\gamma \dot{x}(t') - F_\lambda[x(t')]\}^2 \tag{2.21}$$

for an ensemble with finite nonequilibrium driving and path of length t.

In the limit that the barrier separating two metastable states is large, such that transitions between them are rare, the rate can be computed by extremizing the path action, $\delta U_\lambda / \delta x = 0$. Preforming the functional differentiation, the instantonic trajectory satisfies the second order differential equation

$$\gamma^2 \ddot{x} - \dot{F}_\lambda[x(t')] - F_\lambda[x(t)] \partial_x F_\lambda[x(t)] = 0, \tag{2.22}$$

which can, in principle, be solved subject to the boundary conditions of starting in state A and ending in state B. Here we will consider both states being defined at specific points $x = x_A$ and $x = x_B$ for states A and B, respectfully. However, solving this equation is difficult for nonconservative potentials and time dependent driving forces. We consider the case of a large, sharply peaked barrier, such that the maximum force due to the potential, F_m, is much larger than the magnitude of the applied driving force, $f = \max |\lambda(t)| \ll F_m$. In this limit, the instanton equation simplifies to

$$\gamma^2 \ddot{x} \approx V'(x) V''(x) \tag{2.23}$$

which results in the trajectory traced out by the particle in equilibrium. Its first integral is a constant of motion, yielding

$$\gamma^2 \dot{x}^2 = [V'(x)]^2, \tag{2.24}$$

that provides both branches of the instanton, or extremal path.[19] The positive root yields the trajectory beginning at x_A and ending at the maximum of the potential separating the two metastable states, $x = x_m$, and results in a positive contribution to the action. The negative root yields the trajectory beginning at x_m and ending at x_B, with zero action. The resultant total action is approximated in this limit as

$$U_\lambda \approx \int_0^{t/2} dt' \, \dot{x}(t') F_\lambda[x(t')] = \int_0^{t/2} dt' \, \dot{x}(t') \lambda(t') - \int_{x_A}^{x^\ddagger} dx \, V'(x) \tag{2.25}$$

where consistent with the assumption of an equilibrium trajectory minimizing action, we have neglected terms of order $O(\lambda^2)$ that are strictly negative and invoked time reversal symmetry to set the domain of the integrals. The first term is the equilibrium change in energy, while the second is the heat. Using this result, we can compute the averaged relative action between equilibrium and nonequilibrium path ensembles

$$\langle \Delta U_\lambda \rangle_\lambda \approx \beta \, \langle Q \rangle_{AB}^{\text{neq}} / 2 \tag{2.26}$$

which has no contributions from the activity and coincides with our main result. This calculation clarifies a relevant linear response limit in which the dissipative bound is tight.

It is one in which the typical reactive trajectory in and out of equilibrium are similar, following a gradient path, which occurs in cases where the conservative forces experienced during a transition are large relative to the nonequilibrium driving forces, $f \ll F_m$.

For a second example, we consider a barrier that is locally parabolic in the vicinity of its maximum,

$$V(x) \approx -kx^2/2 + V_0$$

where k denotes its local curvature and V_0 is the offset from the minimum in the A state. In the limit that $\beta V_0 \gg 1$, we can assume that the majority of the action required to overcome the barrier is localized to the region around the maximum. In such a case, the stochastic action in the presence of the external force is extremized by the solution of

$$\gamma^2 \ddot{x}(t) = k^2 x(t) + k\lambda(t) + \gamma\dot{\lambda}(t)\,, \tag{2.27}$$

for the instantonic trajectory subject to boundary conditions which we take to be symmetric about the maximum, $x(0) = -x_0$ and $x(t^*) = x_0$. It is convienent to introduce the charactoristic relaxation time, $\tau = \gamma/k$. This linear ordinary differential equation can be solved by the method of Laplace transforms, yielding

$$x(t) = -x_0 \cosh(t/\tau) + \frac{x_0 + x_0 \cosh(t^*/\tau)}{\sinh(t^*/\tau)} \sinh(t/\tau) + f(t)/\gamma - \frac{f(t^*)/\gamma}{\sinh(t^*/\tau)} \sinh(t/\tau) \tag{2.28}$$

where $f(t)$ is the convolution of the external force with the Green's function,

$$f(t) = \int_0^t dt' \lambda(t') e^{(t-t')/\tau} \tag{2.29}$$

and acts as an inhomogenious source. From the definitions in the main text, the heat and activity are given by

$$Q = \int_0^{t^*} dt\, \dot{x}(t)\lambda(t) \qquad \Gamma = -\frac{1}{2\gamma} \int_0^{t^*} dt\, \lambda(t) \left[2kx(t) + \lambda(t)\right] \tag{2.30}$$

whose averages are computed within the instantonic trajectory. The heat has a familiar form. The activity includes a contributions due to the product of the gradient force and the external force, and a contribution due to the external force squared. While the latter is strictly negative, path independent, and can be dropped while still satisfying the bound, the former may be positive or negative.

In the limit of small $\lambda(t)$, it can be verified that the first contribution to the activity vanishes, as the instanton trajectory spends equal time on the left and right side of the barrier with corresponding equal and opposite forces from the external potential. The second term enters proportional to $-\lambda^2(t)$ which can also be neglected if $\lambda(t)$ is small. This is identical to the near equilibrium case considered above. Analogously, in the limit that $\dot{\lambda} \approx 0$, such that over the transition the external force is well approximated by a constant, by symmetry the

activity will be strictly negative. Consider for concreteness a periodic force $\lambda(t) = f\cos(\omega t)$ with characteristic amplitude f and frequency ω. In the limit that $\omega\tau \ll 1$, the heat will be $\beta\langle Q\rangle_{AB}^{\rm neq} = 2fx_0$ and the activity $\langle A\rangle_{AB}^{\rm neq} = -f^2t^*/2\gamma < 0$. In the opposite limit that $\omega\tau \gg 1$, many cycles of the driving force will elapse during the instantonic trajectory and its influence on breaking the symmetry around the barrier will vanish. In the case of a monochromatic driving force, the heat is $\beta\langle Q\rangle_{AB}^{\rm neq} = f^2t^*/2\gamma$ and the activity $\langle A\rangle_{AB}^{\rm neq} = -f^2t^*/4\gamma < 0$ can be neglected in the bound. Away from these cases, the trajectory need not be symmetrically distributed around the barrier, and the activity maybe finite. However, this occurs when the driving is large, in which case the second term in the activity will dominate leading to it being negative.

Finally, in the case where the barrier is broad and flat, we can approximate the motion at the top as free diffusion conditioned a set of starting and ending points, x_A and x_B. Specifically, when $V(x) = 0$, for an arbitrary external force, the average heat and activity simplify to

$$\beta\langle Q\rangle_{AB}^{\rm neq} = \frac{(x_B - x_A)}{t^*}\int_0^{t^*} dt\,\lambda(t) + \gamma^{-1}\int_0^{t^*} dt\,\left[\lambda(t) - \bar{\lambda}(t^*)\right]^2 \quad \langle A\rangle_{AB}^{\rm neq} = -\frac{1}{2\gamma}\int_0^{t^*} dt\,\lambda^2(t)$$

where $\bar{\lambda}(t^*)$ denotes the external force averaged over the transition time t^*. The activity is manifestly negative and can be dropped from the dissipative bound. As above, it is negligible when the scale of λ is small as it scales quadratically. Taking $|x_B - x_A|$ and t^* large, the heat becomes the displacement times the average force, $\beta\langle Q\rangle_{AB}^{\rm neq} \approx (x_B - x_A)\,\bar{\lambda}$ and the activity $\langle A\rangle_{AB}^{\rm neq} \approx -t^*\bar{\lambda^2}/2\gamma$ where $\bar{\lambda^2}$ is the average squared size of the external driving.

Relative magnitudes of the heat and activity in the transition path ensemble

Here, we discuss the results of additional simulations on systems featured in the main text Figs. 2 and 3. Shown in Fig. 2.4 and 2.5 are the heat (top row), and also the activity (bottom row) *throughout* the course of the transition. Each column represents a driving *speed*: low, medium and, high, from left to right, and we fix all other parameters as in the main text. The active Brownian particle in Fig. 2.4 is driven at a speed D_θ equal to its inverse auto-correlation time, and ω is the analogous variable in the Duffing Oscillator depicted in The quantities plotted in these figures as a function of time, in units of the reactant (A) relaxation time, are averages of $\int_{-t/2}^{\tau-t/2}\dot{Q}$ and $\int_{-t/2}^{\tau-t/2}\dot{\Gamma}$, which are *not* part of the transition path ensemble until crossing occurs at time zero. For the overdamped active Brownian particle, $\dot{\Gamma} = -\gamma^{-1}\lambda(t)F_\lambda[x(t)]$ and in the underdamped Duffing oscillator $\dot{\Gamma} = \gamma^{-1}\lambda(t)(\gamma\ddot{x}(t) - F_\lambda[x(t)])$, which can be derived by constructing the symmetric part of the corresponding path action.

As the reaction proceeds from beginning to end, heat and activity are computed in the path ensemble defined by the additional constraint that the particle position coincides with the transition state, $x^\ddagger$ at time zero. That is, we average over all trajectories starting at time

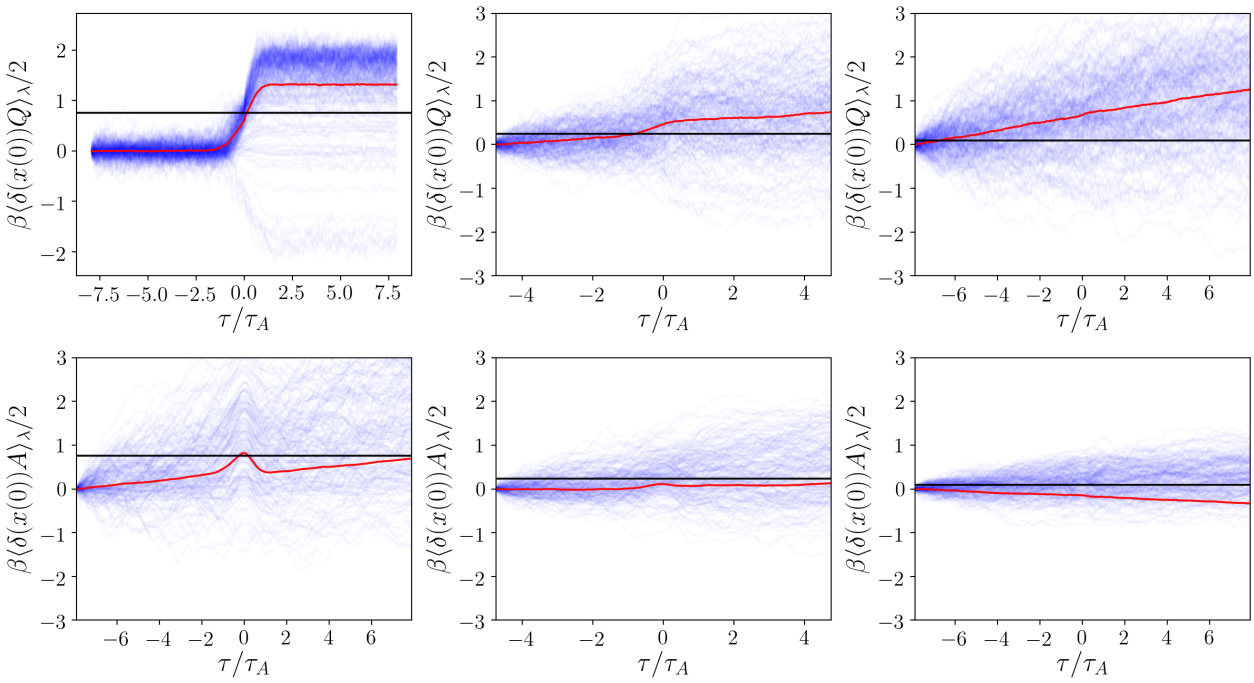

Figure 2.4: Additional examples of active barrier crossing. Rate amplification $\ln k_\lambda/k_0$ (black, from counting transitions), dissipation (red, **top row**), and dynamical activity (red, **bottom row**) along the typical reaction trajectory, in units of the reactant (state A) relaxation time τ_A. All parameters except for active diffusivity D_θ relative to that at which the enhancement is half its maximum D_θ^* are set according to Fig. 3 a). **Left**: slow driving with a long persistence time $D_\theta/D_\theta^* \approx 2 \times 10^{-3}$. **Center**: close half-maximum persistence $D_\theta/D_\theta^* \approx 1$. **Right**: low-persistence forcing $D_\theta/D_\theta^* \approx 7$.

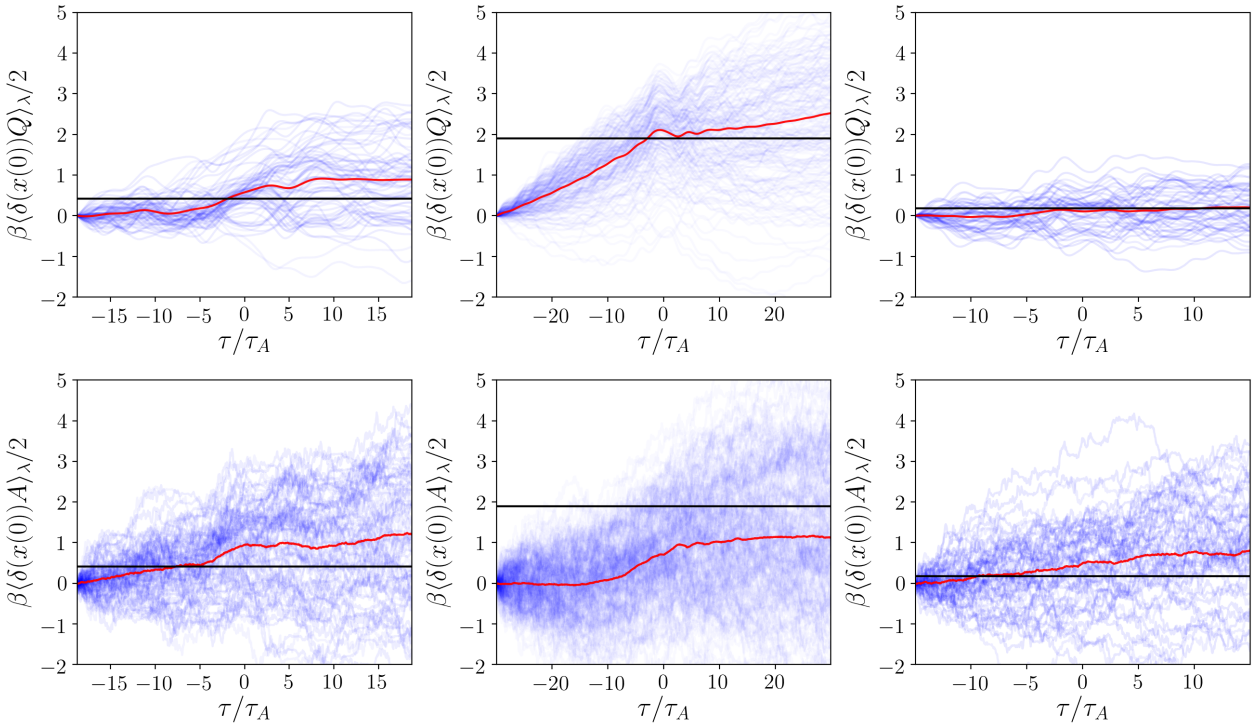

Figure 2.5: Revisiting stochastic resonance. Rate amplification $\ln k_\lambda/k_0$ (black, from counting transitions), dissipation (red, **top row**), and dynamical activity (red, **bottom row**) along the typical reaction trajectory, in units of the reactant (state A) relaxation time τ_A. Parameters apart from driving frequency ω are set according to FIG. 3 a) in the main text. **Left**: slow, quasi-adiabatic driving $\omega/\omega^* \approx 0.42$. **Center**: close to resonance $\omega/\omega^* \approx 1.04$. **Right**: very fast forcing $\omega/\omega^* \approx 1.55$.

$-t/2$ in A before crossing at time 0 and ending up in B a time $t/2$ after that. Since both the underlying system and the driving are symmetric in time in the long-time limit, it is reasonable to believe the instanton connecting A to B is also, so this scheme should sample the AB ensemble as the number of trajectories becomes large. We collect long trajectories of $\dot{Q}$ and $\dot{\Gamma}$, integrating from $-t/2$ to $\tau - t/2$ for $\tau \in (0, t)$. Specifically, we run simulations of $\sim 10^{8-9}\Delta t$ until $\sim 10^{3-4}$ reactions are observed. The majority of paths lie above $\beta \langle Q \rangle_{AB}^{\text{neq}}$ but there are a number of outlying negative heat trajectories, wherein the particle crosses the barrier while also *opposes* the driving force, which is a consequence of the integral fluctuation theorem.

In Heat (top, red) rises past the rate enhancement (black) in the immediate vicinity of this time, an observation consistent with the non-equilibrium transition-state theory we develop below. On the other hand, the activity need not serve as an upper bound, as is

the case for particles driven at the half-maximum speed and resonance frequency in in the center column of Fig. 2.4 and 2.5, respectively. Like we discussed in the section prior, the activity accumulated over a symmetric barrier should vanish by symmetry in the adiabatic limit, a signature which is approximately realized in Fig. 2.4 (bottom left). Finally, when driving varies so quickly that it couples effectively as a second bath, Fig. 2.4, the activity is negative, as would be expected for free diffusion.

2.5 Separation of timescales in the driven and equilibrium ensembles

The existence of a time-independent transition rate between two metastable states requires a separation of timescales between local relaxation within a metastable state and the characteristic time to transition between the two states. [1] The rate enhancement relation in the main text additionally requires that separation exists for both the driven and equilibrium ensembles, and further that these two intervals defined by the local relaxation and typical transition waiting times have some amount of overlap, starting at $t_{\min}$.

Crucially, the observation time t past which our bounds are defined and valid in the conditioned path ensemble must be chosen at least as large as the minimum overlap time $t > t_{\min}$. Before $t_{\min}$, heat will increase approximately linearly. An observed transition to different behavior, typically occurring around the reactant relaxation time τ_A, can be employed as a signature of when our bound is tightest in situations wherein the true rate enhancement is unknown.

Heat does not accumulate in a quasi-equilibrium state, so if both A and B remain deeply metastable, one can expect there to exist an interval of time starting around $t_{\min}$ when the particle commits to state B and $\beta \langle Q \rangle_{AB}^{\text{neq}}$ varies relatively slowly. To test these predictions, we consider the mixed driving system studied in Fig. 1. Specifically, Fig. 2.6 shows the log ratio of the rates in and out of equilibrium as a function of time for that system, together with the accumulated heat, where τ is an intermediate time between 0 and the observation time t. In all cases where there are is large barrier separating states A and B, and the applied force f is smaller than the maximum force F_m due to the potential, the ratio of rates plateaus within a time of $\mathcal{O}(1)$. Further, over the time window when the rate ratio is time independent, heat gains support at a much smaller rate than prior to this window, plateauing in the case of dual metastability in Fig. 2.6 a. Therefore, under these conditions, we find path ensemble averages in question to be somewhat insensitive to precise time at which they are taken.

2.6 Connection to transition-state theory

Our main focus is on rates calculated within the transition path ensemble, obtained by approximating the sum over time-dependent trajectories. The benefit of this approach is

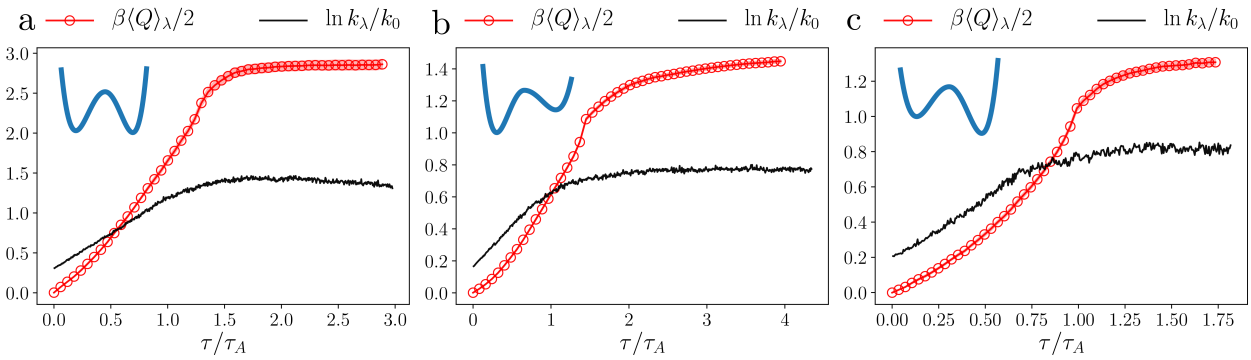

Figure 2.6: Rate amplification k_λ/k_0 (black) and dissipation (red) for three examples of heterogeneously driven two-state systems. The shape of the equilibrium potential for each system is inset (blue) in the upper-left hand corner, and time is taken in units of the relaxation time for state A (the state on the left side of each inset). **a**, transitions in an approximately symmetric double-well driven by equal parts deterministic and active forces, amplitudes summing to around half the equilibrium well-depth. **b** Escape from a basin of attraction to a less-stable state driven by mostly active forces with total max-amplitude around $0.35\Delta V_A$. **c** Excursions to a state with greater stability, forces mainly by a time-periodic protocol with maximum amplitude about $0.75\Delta V_A$.

that all that one needs is a definition of two states, A and B. Detailed mechanistic knowledge of the kinetic bottlenecks through which the transitions pass is unnecessary, alleviating the potentially hard problem of finding a good reaction coordinate, q, and relevant transition state, $q^\ddagger$, dividing metastable states A and B. However, if a pertinent reaction coordinate is known, the rate can be estimated by much simpler means using transition-state theory, which trades in the path partition function $Z_{AB}(\lambda)$ for the average flux through a transition state times the probability of being at $q^\ddagger$. [2] Generalizing the traditional result, for a nonequilibrium system,

$$k_\lambda \leq \frac{\langle|\dot{q}|\rangle_{\lambda,q^\ddagger}}{2}\frac{\rho_\lambda(q^\ddagger)}{\rho_\lambda(A)}. \tag{2.31}$$

where $\langle|\dot{q}|\rangle_{\lambda,q^\ddagger}$ is the average flux through the transition state, and $\rho_\lambda(q^\ddagger)/\rho_\lambda(A)$ is the relative probability of finding the system at the transition state relative to the total probability of finding the system in state A. In equilibrium, the equipartition theorem dictates that the average flux is independent of q, but that is not necessarily true away from equilibrium. The transition-state theory approximation is an underestimate of the true rate, as it neglects those trajectories that switch direction after already passing $q^\ddagger$, and can be made exact by computing a correction termed the transmission coefficient.

For rare, instantonic transitions, the relative probability to reach the transition state,

$\rho_\lambda(q^\ddagger)/\rho_\lambda(A)$, is by far the dominant contribution to the transition-state theory rate. In equilibrium, this term gives rise to the Arrhenius law

$$k_0 \propto \frac{\rho_0(q^\ddagger)}{\rho_0(q_A)} = e^{-\beta \Delta F} \tag{2.32}$$

where $\Delta F = F(q^\ddagger) - F(q_A)$ is the height of the free energy barrier separating $q^\ddagger$ from the most probable state $q_A \in A$. Out of equilibrium, rate estimation via transition state theory becomes a much more difficult problem, because while a time averaged stationary distribution $\rho_\lambda(q)$ still exists, one cannot in general express it in a simple close form.

However, using the Kawasaki distribution equation, we can approximate the nonequilibrium steady-state distribution in terms of a product of the equilibrium distribution and a correction dependent on the mean dissipation. As formulated by Crooks [6] for stochastic dynamics and similar to that derived by Evans and Searles[31], the Kawasaki distribution provides a relation for the probability of being in state q having started in an equilibrium distribution, and cumulant generating function of the accumulated dissipated heat. Specifically,

$$\rho_0(q) = \rho_\lambda(q) \left\langle e^{-\beta Q} \right\rangle_{\lambda,q} \tag{2.33}$$

where Q is the heat dissipated to the environment, $\rho_0(q)$ is the initial equilibrium probability of state q, $\rho_\lambda(q)$ is the nonequilibrium probability of state q, and the brackets $\langle \dots \rangle_{\lambda,q}$ denote an average under the driving force λ conditioned on ending at state q. Taking the saddle point, and applying Jensen's inequality,

$$\frac{\rho_\lambda(q)}{\rho_0(q)} \leq e^{\beta \langle Q \rangle_{\lambda,q}} \tag{2.34}$$

we arrive at a bound for the ratio of the nonequilibrium to equilibrium distributions.

The bound of the steady state distribution cannot be applied directly to the estimation of rates, but under mild assumptions it can provide an estimate similar to the result in the main text. To proceed, we first assume that the probability of being in state A is unchanged between the equilibrium and nonequilibrium ensemble, $\rho_\lambda(A) = \rho_0(A)$. From the Jarzynski equality [8] and integral fluctuation theorems[32], this is true if the protocol is cyclic, because then $\langle \exp(\beta Q) \rangle_\lambda = 1$. This can be obtained by taking the observation time long enough for the driven system to relax back to the original equilibrium distribution by the final time, since in that case, the conservative heat and the change in free energy both cancel, leaving the excess heat behind. Assuming like initial distributions is approximately valid in the limit that A is deeply metastable. Next, we assume that we can ignore the change to the flux over the transition state due to coupling to λ. Moreover, the fluctuation-response inequality[33] implies the first order change to the flux is expected to scale as $\sim \sqrt{\beta \langle Q \rangle_\lambda}$, which is subdominant to the exponential dependence from the change in the probability distribution. Finally, we assume that the transmission coefficient in and out of equilibrium are the same. This is likely a good assumption in the limit that the original

transition state theory estimate in equilibrium is tight, and the transition is instantonic such that the protocol is slowly varying relative to the typical transition path time. Under those assumptions, we find

$$\frac{k_\lambda}{k_0} \leq e^{\beta \langle Q \rangle_{\lambda, q^\ddagger}} \tag{2.35}$$

where we have applied Eq. 2.34 to the probability of reaching the transition state. Note here, as previously, the activity does not enter the approximate bound, and we have a purely mechanical relationship between the enhancement speed of a process and the energy required. Further, the dissipated heat is that accumulated in going from state A to the transition state, $q^\ddagger$. In an instantonic limit, and near equilibrium, we expect the heat accumulated in reaching the top of the barrier to be half that to reach state B, which would result in an analogous expression as the main text.

2.7 Conclusion

We have shown that separation of timescales in rare event problems causes the dissipated heat to serve as a speed limit on how much the rate can increase under nonequilibrium force. We have also formulated a more general theory in terms of the stochastic path action which yields a variational principle that can be used to infer target rates by applying external forces to driven and equilibrium systems alike.

The variational relationship in Eq. 2.11 between the transition rate and a thermodynamic quantity is reminiscent of equilibrium transition state theory, where the rate is bounded by the thermal flux times the probability that a thermal fluctuation brings the system from an initial reactant state to a rare transition state. From the Kawasaki distribution,[34] the nonequilibrium configurational distribution is related to the equilibrium one by a cumulant generating function of the excess dissipation conditioned on ending at a specific configuration. Neglecting changes to the flux generated by coupling to a nonequilibrium process, the transition state theory estimate of the rate is bounded by the average heat conditioned on ending at the transition state. Eq. 2.11 thus has the interpretation of a nonequilibrium extension to standard transition state theory, and is expected to be a good approximation to the rate enhancement in cases where the system spends little time at the top of the barrier.

While our examples have focused on simple one dimensional models, our formalism is general and can be straightforwardly applied to many-body interacting systems. One immediate consequence of higher dimensionality is that unlike in the examples explored here, applied forces need not be aligned with the direction of the most likely transition path. In such cases, we expect the efficiency with which an arbitrarily applied force enhances the rate to be suppressed relative to the largely saturated bounds we have found here, as energy may be transduced into modes not correlated with overcoming the barrier. In light of our results, a natural optimal control problem arises in which nonequilibrium protocols can be

constructed that minimize the dissipation for a given desired rate enhancement. Methods to perform such optimizations employing molecular simulation and importance sampling have been recently developed and show promise in complex systems [35, 9, 36] . Similarly, the protocols uncovered by such an optimization have the potential to lend mechanistic insight into reactions far from equilibrium, as basic concepts like free energy barriers and gradient flows cease being well defined. The effect of dissipation on rate enhancement under counter-diabatic[37] constraints, as well as in discrete-state networks and reaction-diffusion settings, remains to be seen.

Overall, our investigations build a general framework for the systematic and computationally efficient characterization of rate enhancement. Predicting how structure and external influence conspire to alter reaction rates far from equilibrium is of immediate importance in designing proteins, enzymes, small-molecule drugs, and the complex environments in which they operate. We foresee future studies in interacting colloidal and polymeric systems, both in shear and confining geometries that change dynamically in time [38]. Applications to heterogeneous growth, nucleation and jamming will also be interesting avenues to explore [39]. Time-dependent chemical potential gradients in gated release and receptor binding contexts, as well as designing interaction protocols [40] for quick and robust self-assembly [41] and pattern-formation [42] are another set of pressing examples to which our theory can apply. Local heating in ATP hydrolysis and facilitated diffusion on DNA[43], where electric fields play a pivotal role, are all more complex problems that may prove fruitful to study in this manner.

Chapter 3

Variational path sampling to learn and control rare event rates in and out of equilibrium

This work was done in a collaboration with Avishek Das and David Limmer that resulted in the publication "Direct evaluation of rare events in active matter from variational path sampling." [44]

With a theory for how transition rates change with applied nonequilibrium forces, we turn to the variational principle our theory affords to learn transition rates. We start in a 'klunky' transition path ensemble with external forces that make the transition happen always but at great energetic cost, and through iterations of sampling and gradient descent in trajectory space, we generate new forces that affect the transition at lower and lower cost, mimicking the forces a physical system would naturally feel during a spontaneous unforced transition affected only by steady state fluctuations. Transitions in solvated molecules immersed in explicit thermal and athermal active baths are of great interest in biophysics and physical chemistry. We show that our generative reinforcement learning scheme can learn rates and transition mechanisms in systems in and out of equilibrium. We study transitions of a two state passive solute in a volume excluding liquid with and without the persistent angular motion afforded by activity.

Active matter represents a broad class of systems that evolve far from equilibrium due to the local injection of energy. Like their passive analogues, transformations between distinct metastable states in active matter proceed through rare fluctuations, however their detailed balance violating dynamics renders these events difficult to study. Here, we present a simulation method for evaluating the rate and mechanism of rare events in generic nonequilibrium systems and apply it to study the conformational changes of a passive solute in an active fluid. The method employs a variational optimization of a control force that renders the

rare event a typical one, supplying an exact estimate of its rate as a ratio of path partition functions. Using this method we find that increasing activity in the active bath can enhance the rate of conformational switching of the passive solute in a manner consistent with recent bounds from stochastic thermodynamics.

3.1 Direct evaluation of rare events in active matter from variational path sampling

The constituent agents of active matter– biomolecules, colloids, or cells– autonomously consume energy to fuel their motion.[45, 46] Their resultant nonequilibrium states have non-Boltzmann phase-space densities and exhibit exotic structural and dynamical collective fluctuations, including motility-induced phase separation and swarming.[47, 48, 22, 49, 50] Within these nonequilbrium steady-states, fleeting fluctuations can free particles from external potentials,[24, 51, 25] nucleate stable phases from metastable ones,[52, 53] and assemble passive objects.[54, 55] The study of such rare dynamical events within active matter and the calculation of their associated rates is difficult. Traditional equilibrium rate theories like transition state theory and Kramer's theory require knowledge of the form of the steady-state distribution that is not in general available.[56] Further, only a few numerical methods exist that can be used to tame the exponential computational cost associated with sampling the unlikely fluctuations that lead to transitions between metastable states. Existing methods improve sampling by stratifying or branching stochastic trajectories[57, 58, 59] but do not typically employ driving forces to specifically enhance the sampling of these rare events.

Here we present a perspective and an associated numerical algorithm, termed Variational Path Sampling (VPS), for estimating transition rates in active systems using optimized time-dependent driving forces. Our approach relies on a equality between the rate of a rare event in a reference system and a ratio of path partition functions in the reference system and with a driving force that makes the rare event occur with high probability. The VPS algorithm solves a variational problem to approximate the functional form of an optimal time-dependent driving force for this estimate and is applicable to any stochastic dynamics. With VPS we investigate how driven fluids can direct motion into useful function. We apply this technique to study the rate of conformational changes of a passive dimer in a dense bath of active Brownian particles.[60, 61, 62] This model exemplifies how collective active fluctuations around passive solutes can drive self-assembly and speed up transitions between distinct metastable states. [63, 64] We find the rate to switch between the dimer's two metastable states increases dramatically with increasing activity in the bath, which we rationalize with a recent dissipation bound from stochastic thermodynamics.[13] We study the computational efficiency of rate estimation with VPS and demonstrate its advantage over existing trajectory stratification based methods like Forward Flux Sampling.[57]

We consider a system described by overdamped Brownian dynamics of the form,

$$\gamma_i \dot{\mathbf{r}}_i(t) = \mathbf{F}_i[\mathbf{r}^N(t)] + \boldsymbol{\eta}_i(t) \tag{3.1}$$

where $\dot{\mathbf{r}}_i$ is the rate of change of the i-th particle's position, γ_i is the corresponding friction coefficient, and $\mathbf{F}_i[\mathbf{r}^N(t)]$ is the sum of all conservative, nonconservative and active forces exerted on the i-th particle that depends on the full configuration of the N-particle system, $\mathbf{r}^N$. The final term, $\boldsymbol{\eta}_i(t)$, is a Gaussian white-noise with $\langle \eta_{i\alpha}(t) \rangle = 0$ and

$$\langle \eta_{i\alpha}(t)\eta_{j\beta}(t') \rangle = 2\gamma_i k_{\mathrm{B}} T \delta_{ij}\delta_{\alpha\beta}\delta(t-t') \tag{3.2}$$

for component (α, β) and $k_{\mathrm{B}}T$ is Boltzmann's constant times the temperature. In order to study the transition rate between two long-lived metastable states, denoted A and B, we define each from a given configuration using the indicator functions,

$$h_X[\mathbf{r}^N(t)] = \begin{cases} 1 & \text{if } \mathbf{r}^N(t) \in X \\ 0 & \text{else} \end{cases} , \tag{3.3}$$

for either $X = A, B$. In practice this designation requires an order parameter capable of distinguishing configurations and grouping them into these distinct metastable states like that illustrated in Fig. 4.1(a) in one dimension. Assuming there exists a separation between the time $\tau^{\ddagger}$ required to traverse the transition region between the two metastable states, and the typical waiting time for the transition, the rate k can be evaluated from the probability to observe a transition, per unit time [2]

$$k = \frac{\langle h_B(t_f) h_A(0) \rangle}{t_f \langle h_A \rangle} = t_f^{-1} \langle h_{B|A}(t_f) \rangle \,, \tag{3.4}$$

where the angular brackets denote an average over trajectories of duration $\tau^{\ddagger} < t_f \ll 1/k$ started from a steady-state distribution in A and $\langle h_{B|A}(t_f) \rangle$ denotes the conditional probability for transitioning between A and B in time t_f. When t_f is chosen to satisfy the timescale separation described above, k is independent of time.

If the transition is rare, most short trajectories are nonreactive leading to difficulties in estimating the rate directly. Instead of trying to evaluate the small transition probability through stratification as other existing methods do,[57, 58] we instead optimize a time-dependent driving force $\boldsymbol{\lambda}(\mathbf{r}^N, t)$ that constrains the transition to occur, and evaluate the probability cost associated with adding that force to the original dynamics. For a general time-dependent force $\boldsymbol{\lambda}$, using the Onsager-Machlup form for the probabilities of stochastic trajectories,[4] the rate expression in Eq. 3.4 can be rewritten as[13]

$$k = t_f^{-1} \left\langle e^{-\Delta U_{\boldsymbol{\lambda}}} \right\rangle_{B|A,\boldsymbol{\lambda}} , \tag{3.5}$$

where $\langle \rangle_{B|A,\boldsymbol{\lambda}}$ denotes a conditioned average computed in presence of the additional force. This relation holds for forces $\boldsymbol{\lambda}$ that affect the transition to occur with probability 1, such that the rate in the driven ensemble is $1/t_f$. The average is of the exponential of the change in the path action, $\Delta U_{\boldsymbol{\lambda}}$,

$$\Delta U_{\boldsymbol{\lambda}}[\mathbf{X}] = -\int_0^{t_f} dt \sum_i \frac{[\boldsymbol{\lambda}_i^2 - 2\boldsymbol{\lambda}_i \cdot (\gamma_i \dot{\mathbf{r}}_i - \mathbf{F}_i)]}{4\gamma_i k_{\mathrm{B}} T} \,, \tag{3.6}$$

between trajectories generated with the added force and in its absence. The path action and all other stochastic integrals are evaluated in the Ito convention.

Equation 3.5 is a direct estimator for a rate employing an auxiliary control system, but it only becomes useful when the protocol $\boldsymbol{\lambda}(\mathbf{r}^N, t)$ generates trajectories in a manner equivalent to the unbiased reactive trajectory distribution. This is because the expectation can be viewed as an overlap between the two reactive path distributions, and without significant overlap the exponential average is difficult to estimate. We express the optimality of $\boldsymbol{\lambda}$ using Jensen's inequality after taking the logarithm of Eq. 3.5 to obtain a variational principle,

$$\ln k \geq -\ln t_f - \langle \Delta U_{\boldsymbol{\lambda}} \rangle_{B|A,\boldsymbol{\lambda}} \,. \tag{3.7}$$

If the average change in conditioned path action $\langle \Delta U_{\boldsymbol{\lambda}} \rangle_{B|A,\boldsymbol{\lambda}}$ is minimized over all possible functional forms of $\boldsymbol{\lambda}$, the rate can be obtained directly as a simple ensemble average of $\Delta U_{\boldsymbol{\lambda}^*}$ at the minimizer $\boldsymbol{\lambda} = \boldsymbol{\lambda}^*$.

The optimal control force $\boldsymbol{\lambda}^*$ that saturates Eq. 3.7 is unique and given by the solution of the backward Kolmogorov equation[65, 66, 67] as detailed in the Supporting Material (SM). Specifically, the optimal force is $2k_{\mathrm{B}}T$ times the gradient of the logarithm of the commitor probability[68] of ending in state B at t_f. A schematic illustration of the optimal effective time-dependent potential $V_t(R)$ added to a double well potential is illustrated in Fig. 4.1(a). The resultant force gradually destabilizes the reactant well to ensure the transition almost surely within the short duration t_f. Viewed in the backwards direction of time, the potential follows the negative logarithm of the relaxation of an initially localized distribution in B to its steady-state. The force is thus optimal in the sense that reactive trajectories, like those in Fig. 4.1(b), generated with it are drawn from the reference path ensemble with the correct statistical weights. Generically, $\boldsymbol{\lambda}^*(\mathbf{r}^N, t)$ is a function of all particle coordinates, so it is not typically tractable to compute. We demonstrate here that one- and two-body representations of $\boldsymbol{\lambda}$ can be sufficiently close to optimal as to estimate the rate accurately even in cases where the rare event is collective, similar to related observations in large-deviation sampling.[69, 70, 71, 9]

We study the accuracy and utility of this formalism in a system comprised of an active bath and a passive dimer that can undergo conformational changes between two metastable states. All particles interact pairwise via a Weeks-Chandler-Andersen (WCA) repulsive potential[72]

$$V_{\mathrm{WCA}}(\mathbf{r}) = \left\{ 4\epsilon \left[\left(\frac{\sigma}{r}\right)^{12} - \left(\frac{\sigma}{r}\right)^{6} \right] + \epsilon \right\} \Theta(r_{\mathrm{WCA}} - r) \tag{3.8}$$

with energy scale ϵ, and particle diameter σ, truncated at $r_{\mathrm{WCA}} \equiv 2^{1/6}\sigma$ with the Heaviside function Θ. Active particles experience an additional self-propulsion force of magnitude v_0, $\mathbf{F}_i^a(t) = v_0 \mathbf{e}[\theta_i(t)]$ where the director is $\mathbf{e}(\theta_i) = (\cos\theta_i, \sin\theta_i)$ and θ_i obeys $\dot{\theta}_i(t) = \xi_i(t)$ with,

$$\langle \xi_i(t) \rangle = 0, \quad \langle \xi_i(t)\xi_j(t') \rangle = 2D_\theta \delta_{ij} \delta(t - t') \tag{3.9}$$

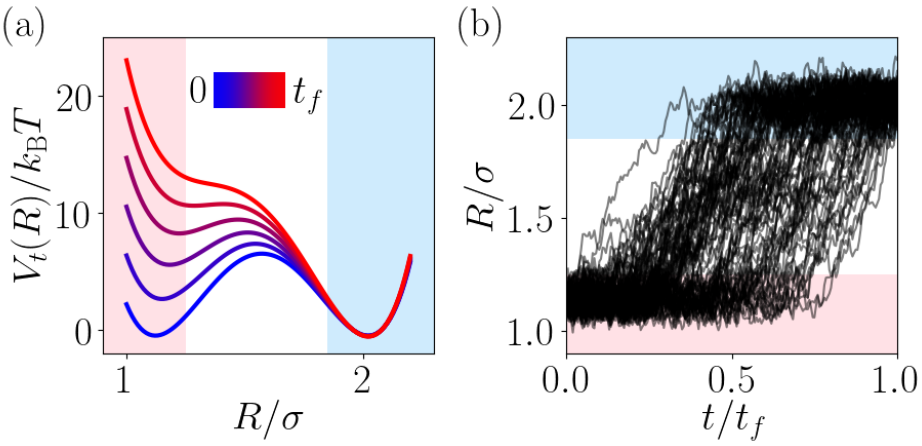

Figure 3.1: Reactive trajectories with VPS. (a)Schematic representation of the total optimal time-dependent potential in an isolated passive dimer as t goes from 0 to t_f. Shaded regions are the compact (A, pink) and extended (B, light blue) states. (b) Unbiased reactive trajectories generated with $\boldsymbol{\lambda}(R,t)$.

for angular diffusion constant D_θ. Passive solutes separated by distance R are bound by a double-well potential

$$V_{\mathrm{dw}}(R) = \Delta V \left[1 - (R - r_{\mathrm{WCA}} - w)^2/w^2\right]^2 \tag{3.10}$$

with an energy barrier of height ΔV between the compact and extended states at $R = r_{\mathrm{WCA}}$ and $R = r_{\mathrm{WCA}} + 2w$ respectively.[73] We study the transition rates between these states, employing indicator functions $h_A(t) = \Theta(R_A - R)$ and $h_B(t) = \Theta(R - R_B)$ for $R_A = 1.25\sigma$ and $R_B = 1.85\sigma$. Conformation transitions like these in dense fluids are collective in origin[73] and serve as a sensitive probe of the bath.

The VPS algorithm estimates an optimal force using a low-rank ansatz by iteratively solving the variational problem in Eq. 3.7, and uses this force to directly obtain a rate estimate. For computing the rate of isomerization of the passive dimer, we approximate $\boldsymbol{\lambda}^*$ with a time-dependent interaction along the dimer bond vector $\mathbf{R}$, expressed as a sum of Gaussians

$$\boldsymbol{\lambda}(\mathbf{R},t) = \hat{\mathbf{R}} \sum_{p,q=1}^{M_R,M_t} c_{pq}^{(i)} e^{-\frac{(R-\mu_{R,p})^2}{2\nu_R^2} - \frac{(t-\mu_{t,q})^2}{2\nu_t^2}} \tag{3.11}$$

where $c_{pq}^{(1)} = -c_{pq}^{(2)}$ are variational parameters to be tuned, and the locations and widths $\mu_{R,p}$, $\mu_{t,q}$, ν_R and ν_t are held fixed. To impose the conditioning while minimizing $\langle \Delta U_{\boldsymbol{\lambda}} \rangle_{B|A,\boldsymbol{\lambda}}$, we use a Lagrange multiplier s to construct a loss function $\Omega_{\boldsymbol{\lambda}} = \langle \Delta U_{\boldsymbol{\lambda}} \rangle_{\boldsymbol{\lambda}} + s(\langle h_{B|A} \rangle_{\boldsymbol{\lambda}} - 1)$. For a general force that does not ensure the transition with unit probability, there is a

multiplicative contribution to the estimate of the rate in Eq. 3.5 from $\langle h_{B|A}\rangle_{\boldsymbol{\lambda}}$, which for most optimized forces is negligible.

The optimization problem maps onto the computation of a cumulant generating function for the statistics of the indicator $h_B(t_f)$ studied previously,[67, 74] with the short trajectories starting from a steady-state distribution in the initial state. As such we can employ generalizations of recent reinforcement learning procedures to efficiently estimate the gradients of the loss function with respect to the variational parameters.[75] Specifically, we modify the Monte-Carlo Value Baseline (MCVB) algorithm[74] which performs a stochastic gradient descent to optimize $c_{pq}^{(i)}$. We add two preconditioning steps over the MCVB algorithm. First, we generate an initial reactive trajectory using a routine reminiscent of well-tempered metadynamics.[76] Then we symmetrize the learned force to ensure time translational invariance of the transition paths. We denote this preconditioning algorithm MCVB-T. Further information is available in the SM.

We first illustrate the systematic convergence of VPS by estimating the isomerization rate of an isolated passive dimer. Such a simplified system allows us to compare to numerically exact results, and study convergence of the force ansatz in the complete basis limit, where $M_R, M_t \to \infty$ and the Gaussians cover the thermally sampled region in R and t. For this simple system, we take $k_{\mathrm{B}}T = \gamma = \sigma = \epsilon = 1$, $w = 0.25\sigma$, with diffusive timescale $\tau = \sigma^2\gamma/k_{\mathrm{B}}T$. We simulate the one-dimensional version of Eq. 5.7 along R, with $V_{\mathrm{dw}}(R)$ only. For simplicity we define state A by the initial condition $R(0) = r_{\mathrm{WCA}}$, and state B via $R_B = 1.45\sigma$. To provide a steady-state value in Eq. 3.4[74, 77], we use an Euler method and take in this example $t_f = \gamma w\sigma/\sqrt{8k_{\mathrm{B}}T\Delta V}$. We choose $\mu_{R,p}$ and $\mu_{t,q}$ evenly distributed in $R/\sigma \in [0.9, 1.77]$ and $t \in [0, t_f]$, respectively, and ν_R, ν_t to be half the distance between Gaussian centers. We consider basis sizes $M_R = M_t = 2 - 40$, each optimized independently and used to sample $\sim 10^5$ transition paths.

Figure 4.2(a) illustrates a typical learning curve for the control force, showing convergence of the variational rate bound towards the numerically exact rate. The variational estimate requires a basis of $M_R, M_t > 40$ to approach the rate to within the statistical uncertainty of the estimate, however alternative estimates with small basis sets can be refined using a cumulant expansion approximation to Eq. 3.5. Specifically, truncating the exact exponential relation at the ℓth cumulant as

$$\ln k \approx -\ln t_f + \sum_{n=1}^{\ell} \frac{1}{n!} \frac{d^n \ln \left\langle e^{-\Delta U_{\boldsymbol{\lambda}}} \right\rangle_{B|A,\boldsymbol{\lambda}}}{d\Delta U_{\boldsymbol{\lambda}}^n} \tag{3.12}$$

provides an approximation to the rate that converges in the limit that ℓ is large. Figure 4.2(b) illustrates this convergence, where we find that even coarse-representations of the control force can yield close estimates of the rate with only the first few cumulants, illustrating a tradeoff between basis set completeness and statistical efficiency. Sweeping across a wide range of barrier heights in Fig. 4.2(c), we find excellent agreement between the log-rate from brute force simulations and a truncation of the cumulant expansion to $\ell = 2$ using $M_R = 80$ and $M_t = 30$.

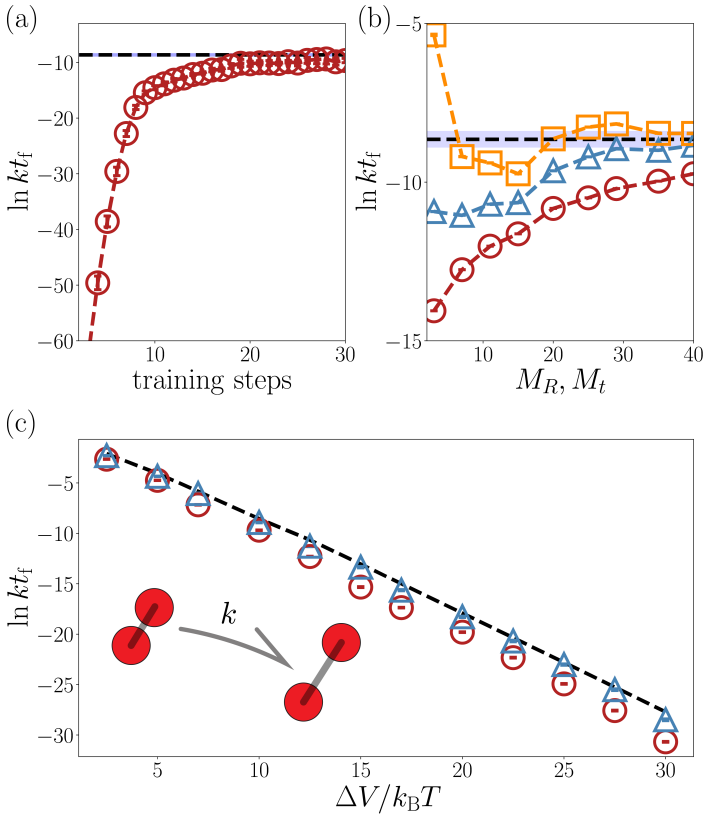

Figure 3.2: Convergence of isomerization rates for an isolated passive dimer. (a) Learning curve for $\Delta V = 10k_BT$ and $M_R, M_t = 20$. (b) Convergence of the variational rate estimate (circles) and cumulant corrections for $\ell = 2$ (triangles) and $\ell = 4$ (squares) with basis size as compared to the numerically exact answer (dashed line). (c) Variational (circles) and $\ell = 2$ (triangles) estimate of the rate compared to the exact value (dashed line) with increasing barrier height.

We next compute the isomerization rate with VPS when the dimer is immersed in an explicit solvent of active Brownian particles with $N = 80$ and a total density of $0.6/\sigma^2$. The dimer particles have a friction $\gamma_{\rm d} = 2\gamma$ and the solvent particles have $\gamma_{\rm s} = 4\gamma$. We take $\gamma = \sigma = \epsilon = 1$, $k_{\rm B}T = 0.5$, $\Delta V = 7k_{\rm B}T$, $\tau = \sigma^2\gamma/2k_{\rm B}T = 1$, $D_\theta = 1/\tau$ and timestep $10^{-5}\tau$. We also change $w = 0.45\sigma$ such that the collisional cross-section of the dimer is large. Collisions with active particles transduce energy along the dimer bond and we study the change in the isomerization rate as the bath activity $v_0\sigma/k_{\rm B}T$ is varied from 0 to 18. We use a basis size of $M_R = M_t = 50$ distributed between $R/\sigma \in [0.9, 2.3]$ and $t \in [0, t_f]$ where $t_f = 0.2\tau$. The optimization starts by learning forces $\boldsymbol{\lambda}(\mathbf{R}, t)$ for the isolated dimer with WCA interactions between monomers, followed by the MCVB-T algorithm. Then, $\boldsymbol{\lambda}(\mathbf{R}, t)$ is optimized in the presence of the bath for $v_0 = 0$ and higher values of v_0 are initialized from converged forces at the previous v_0.

The rate is a strong function of activity, increasing twenty-fold over the range of v_0's considered. While the variational rate estimate from Eq. 3.7 is closest for the passive bath, it weakens with increasing v_0, indicating a growing importance of solvent degrees of freedom in the optimal control force. With converged forces at each v_0, we run 10^6 trajectories of length t_f to compute k from Eq. 3.5. This estimate correctly predicts the suppression of k due to passive solvation and can be converged statistically for $v_0\sigma/k_{\rm B}T < 9$, which is supported by direct rate estimates from unbiased simulations in Fig. 4.3(a). Above $v_0\sigma/k_{\rm B}T = 9$, the optimized force is not close enough to $\boldsymbol{\lambda}^*$ to estimate k directly through the exponential average or a low order cumulant expansion.

Provided we have access to the transition path ensemble from direct unbiased simulations or methods like Transition Path Sampling[1, 78, 79] we can supplement the estimate of k using histogram reweighting.[80] k satisfies a reweighting relation of the form,

$$k = \frac{e^{-\Delta U_{\boldsymbol{\lambda}}} P_{B|A,\boldsymbol{\lambda}}(\Delta U_{\boldsymbol{\lambda}})}{t_f P_{B|A,0}(\Delta U_{\boldsymbol{\lambda}})} \tag{3.13}$$

where we have defined $P_{B|A,\boldsymbol{\lambda}}(\Delta U_{\boldsymbol{\lambda}}) = \langle \delta(\Delta U_{\boldsymbol{\lambda}}[\mathbf{X}] - \Delta U_{\boldsymbol{\lambda}}) \rangle_{B|A,\boldsymbol{\lambda}}$ and similarly for its undriven counterpart $\boldsymbol{\lambda} = 0$. We evaluate k with this estimator by sampling 10^4 driven and only 6-100 unbiased reactive paths, using the Bennett Acceptance Ratio[81] to evaluate the ratio of probabilities. Compared with the brute-force estimate in Fig. 4.3(a), we find this reweighting predicts k accurately across all values of v_0 with significantly higher statistical efficiency then a brute force calculation, which validates the accuracy and utility of the control forces. We have compared the VPS rate estimates in the SM, using either Eqs. 3.5 and 3.13, to the Rosenbluth variant of Forward Flux Sampling[57], and find that VPS is statistically more efficient and converges more quickly with the number of reactive trajectories.

Access to an ensemble of transition paths in this active system gains us mechanistic insight into the process. The rate enhancement observed for the compact to extended state transition of the passive dimer with bath activity can be understood using recent results from stochastic thermodynamics. Specifically the rate enhancement achievable by coupling a reactive mode to a nonequilibrium driving force is bounded from above by the heat dissipated

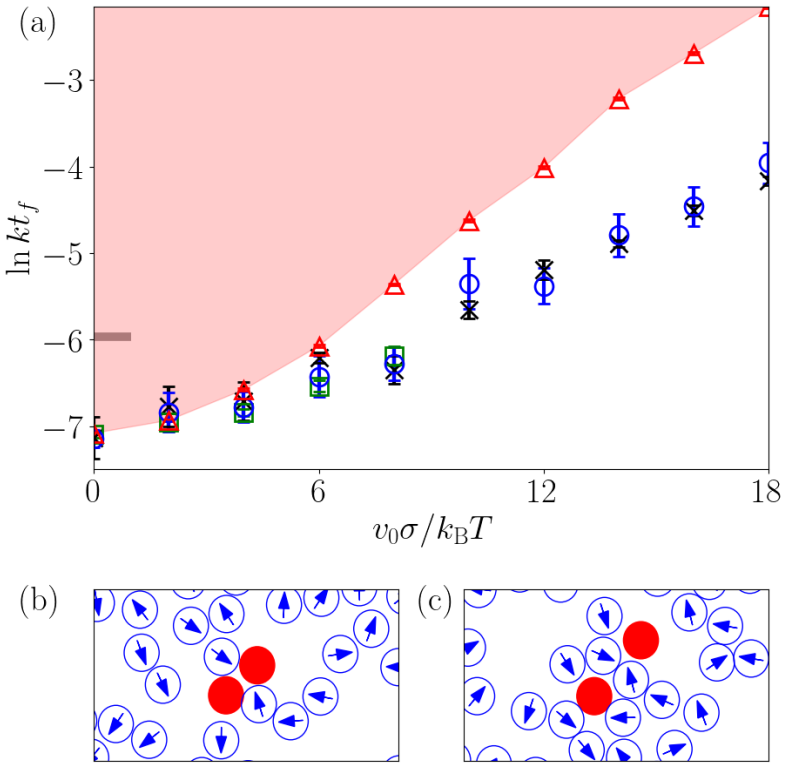

Figure 3.3: Rate enhancement of isomerization in an active fluid. (a) Change in the rate as estimated from direct unbiased simulations (crosses), from exponential estimate (squares), and from histogram reweighting (circles). The excess dissipated heat (triangles) bounds the rate enhancement achievable demarked by the red shaded region. The thick tick mark on the left denotes the rate for the isolated dimer. (b) and (c) Typical snapshots of reactive trajectories of the active bath (blue) and passive dimer (red), at $t = 0$ and $t = t_f$.

over the course of the transition.[13] In this case the nonequilibrium driving is afforded by the interactions between the dimer and the active bath, so the bound takes the form

$$\ln k \leq \ln k_0 + \frac{1}{2k_BT}\langle Q - Q_0\rangle_{B|A} \tag{3.14}$$

where k_0 is the rate at $v_0 = 0$ and $\langle Q - Q_0\rangle_{B|A}$ is the dissipative heat less its average at $v_0 = 0$ given by

$$Q = \int_0^{t_f} dt \sum_{i\in \mathrm{d}} \sum_{j\in \mathrm{s}} (\dot{\mathbf{r}}_i - \dot{\mathbf{r}}_j) \cdot \mathbf{F}_{\mathrm{WCA}}(\mathbf{r}_{ij}) \tag{3.15}$$

which is a sum of the total force from the WCA potential of the solvent particles (s) on the dimer (d) times the difference in their velocities in an ensemble at fixed v_0 (SM). This

bound is verified in Fig. 4.3(a) for all v_0, and saturated at small v_0. The specific mechanism of energy transfer from bath to dimer that promotes transitions is clarified by examining reactive trajectories driven by the biasing force and are typical, after removal of the bias from the incomplete basis set. Figures 4.3 (b) and (c) show typical snapshots of the solvated dimer at the start and end of the reaction. Energy transfer results from active particles accumulating around the dimer, and preferentially in its cross-section, pushing it apart into an extended state. This mechanism of action is reminiscent of how nonequilibrium agents collect in the corners of mesoscopic gears to power their directed rotation.[63, 64] At low v_0, we find the driven isomerization process is efficient, while deviation from the bound at large v_0 demonstrates that energy is additionally funneled into non-reactive modes. Further studies showing the unbiased nature of the VPS-sampled transition path ensemble in terms of duration and distribution of transition paths, and quantification of the changing solvation environment with v_0 are provided in the SM.

3.2 Doob's optimal force saturates the rate bound

For ease of notation, we consider states A and B being specific phase space points, $\mathbf{r}_A^N$ and $\mathbf{r}_B^N$, respectively. Let $\boldsymbol{\lambda}^*(\mathbf{r}^N, t) = 2k_\mathrm{B}T\boldsymbol{\nabla}\Phi$, where $\Phi(\mathbf{r}_B^N, t_f|\mathbf{r}^N, t)$ is the log of the probability to end at a single target configuration, $\mathbf{r}_B^N$, conditioned on being at $\mathbf{r}^N$ at time t. This conditioned probability satisfies the logarithmic transform of the backward Kolmogorov equation,[82, 65, 66, 67, 83]

$$\partial_t\Phi + \sum_i \Big[D_i(\boldsymbol{\nabla}_i\Phi)^2 + \boldsymbol{\nabla}_i\Phi\cdot\mathbf{F}_i/\gamma_i + D_i\boldsymbol{\nabla}_i^2\Phi \Big] = 0 \tag{3.16}$$

where $D_i = k_\mathrm{B}T/\gamma_i$ is the diffusion constant and the gradients act on $\mathbf{r}_i$. The optimal force achieves the reactive transition by construction, rendering $\langle h_{B|A}\rangle_{\boldsymbol{\lambda}^*} = 1$. The change in path action with the optimal force is

$$\Delta U_{\boldsymbol{\lambda}^*}[\mathbf{X}] = -\int_0^{t_f} dt \sum_i D_i(\boldsymbol{\nabla}_i\Phi)^2 + \boldsymbol{\nabla}_i\Phi\cdot\mathbf{F}_i/\gamma_i - \boldsymbol{\nabla}_i\Phi\cdot\dot{\mathbf{r}}_i \tag{3.17}$$

Using Ito's Lemma for the total time derivative of Φ to eliminate the final term,

$$\Delta U_{\boldsymbol{\lambda}^*}[\mathbf{X}] = -\int_0^{t_f} dt \left\{ \sum_i \left[D_i(\boldsymbol{\nabla}_i\Phi)^2 + \boldsymbol{\nabla}_i\Phi\cdot\mathbf{F}_i/\gamma_i + D_i\boldsymbol{\nabla}_i^2\Phi \right] - \dot{\Phi} + \partial_t\Phi \right\} \tag{3.18}$$

and substituting the backward Kolmogorov equation

$$\Delta U_{\boldsymbol{\lambda}^*}[\mathbf{X}] = \int_0^{t_f} dt\, \dot{\Phi} = \Phi(\mathbf{r}_B^N, t_f|\mathbf{r}_B^N, t_f) - \Phi(\mathbf{r}_B^N, t_f|\mathbf{r}_A^N, 0) \tag{3.19}$$

the change of action can be evaluated exactly. The boundary terms from the exact integration are 0 and the log of the transition probability between A and B in the reference system. The

latter can be identified with $\ln\langle h_{B|A}\rangle$, resulting in the equality in Eq. 7 using the definition of the rate k in Eq. 4. This reasoning extends linearly to cases where A and B are collections of configurations.

3.3 Generative reinforcement learning protocol for learning optimal force

We optimize the time-dependent control force $\boldsymbol{\lambda}(\mathbf{R}, t)$ by minimizing $\langle \Delta U_{\boldsymbol{\lambda}} \rangle_{B|A,\boldsymbol{\lambda}}$ over variational parameters $c_{pq}^{(i)}$ using Lagrange multiplier s to impose the $B|A$ conditioning. If s is chosen to be more negative than an approximately estimated threshold value $s^* = \ln[\langle h_{B|A}\rangle/(1 - \langle h_{B|A}\rangle)]$, the optimized forces drive the reaction with unit probability and s need not be individually optimized. For a rare transition, any choice of s with a magnitude an order or more larger than the energy barrier height will robustly provide forces that always satisfy the $B|A$ conditioning.[74]

For optimization we use an extension of a reinforcement learning algorithm called Monte Carlo Value Baseline (MCVB).[74] This algorithm computes the correlation of the gradient of the log of the trajectory probability, called Malliavin weights,[9, 84] with the instantaneous change in $\Delta U_{\boldsymbol{\lambda}}$ and h_B over the course of the trajectory. These yield the exact gradients of the loss-function $\Omega_{\boldsymbol{\lambda}}$ with respect to the tunable parameters, with which a stochastic gradient descent is performed. The MCVB algorithm simultaneously learns the driving force and a corresponding value function, $\mathcal{V}(\mathbf{R}, t) = \langle \Delta U_{\boldsymbol{\lambda},t} + s h_{B|A} \rangle_{\boldsymbol{\lambda}|R(t)=R}$, where $\Delta U_{\boldsymbol{\lambda},t}$ contains the integrated action difference only within $[t, t_f]$ and the expectation is conditioned on starting from R at time t. The value function greatly reduces the the variance of the gradients at zero cost, allowing better convergence. Our modification to this algorithm, refered to as MCVB-T, is a preconditioning step that enforces time translational symmetry for the log of the bridge probability, $\Phi(\mathbf{r}_f^N, t_f|\mathbf{r}^N, t) = \Phi(\mathbf{r}_f^N, t_f - t|\mathbf{r}^N, 0)$, as the reference forces are not explicit functions of time. This is achieved by randomly choosing a $t_{\text{mid}} \in [0, t_f]$ for every trajectory used for averaging the force gradient, and applying the force $\boldsymbol{\lambda}(\mathbf{R}, t \in [t_{\text{mid}}, t_f])$ on it only for a duration $[t_{\text{mid}}, t_f]$. This ensures that trajectories undergoing the transition at late times are accounted for while training the force.

Details of the MCVB-T algorithm are available in the pseudocode in Algorithm 1. The set of Gaussian coefficients parametrizing the force and the value function are denoted in short by χ and ψ respectively. The MCVB algorithm is a special case of MCVB-T with fixed $t_{\text{mid}} = 0$.

Figure 3.4 illustrates all learning curves that led to the results in Fig. 3, plotting $\ln \overline{h_{B|A}} - \overline{\Delta U_{\boldsymbol{\lambda}}}\big|_{B|A}$ as a function of training steps. The averages of this estimator and the gradients are computed over 40 trajectories simulated at each training step. The trajectories are initialized with coordinates randomly chosen from a collection of 10000 steady-state dimer and bath configurations in state A, collected once in every 0.1τ time units from a long trajectory without any driving forces. We first learn optimal forces in the absence of the

Algorithm 1 Monte-Carlo Value Baseline with Time-translation invariance (MCVB-T)

inputs Gaussian coefficients for a general force $\boldsymbol{\lambda}_\chi(\mathbf{r}^N, t)$ and value function $\mathcal{V}_\psi(\mathbf{r}^N, t)$
parameters learning rates α^χ, α^ψ; total optimization steps I; trajectory length t_f consisting of J timesteps of duration Δt each; number of trajectories N
initialize choose initial weights χ and ψ, define iteration variables i and j, force and value function gradients δ_P, δ_V, define functional form for stepwise increments (rewards) ξ to the loss-function $\Delta U_{\boldsymbol{\lambda}} + sh_{B|A}$
$i \leftarrow 0$
repeat
    Generate trajectories $[\mathbf{X}(t)]$ with first-order Euler propagation starting from uncorrelated steady-state configurations in state A. Every trajectory starts experiencing the force $\boldsymbol{\lambda}$ from a random time t_{mid} which is sampled uniformly from $[0, t_f]$. Configurations, times, noises (with variance $2\boldsymbol{\gamma} k_{\text{B}} T \Delta t$), Malliavin weights, integral of value function gradients, and rewards are denoted by $\mathbf{r}_j^N, t_j, \boldsymbol{\eta}_j, y_\chi(t_j), z_\psi(t_j)$ and $\xi(t_j) = \xi_j$ respectively.
    $j \leftarrow 0$
    $\delta_P \leftarrow 0$
    $\delta_V \leftarrow 0$
    $y_\chi(t_0) \leftarrow 0$
    $z_\psi(t_0) \leftarrow 0$
    repeat
        $\dot{y}_\chi(t_j) \leftarrow \boldsymbol{\eta}_j \cdot \nabla_\chi \boldsymbol{\lambda}_\chi(\mathbf{r}_j^N, t_j)/2k_{\text{B}}T\Delta t$
        $y_\chi(t_{j+1}) \leftarrow y_\chi(t_j) + \Delta t \dot{y}_\chi(t_j)$
        $\dot{z}_\psi(t_j) \leftarrow \nabla_\psi \mathcal{V}_\psi(\mathbf{r}_j^N, t_j)$
        $z_\psi(t_{j+1}) \leftarrow z_\psi(t_j) + \Delta t \dot{z}_\psi(t_j)$
        $\delta_P \leftarrow \delta_P + \xi_j y_\chi(t_{j+1}) - \mathcal{V}_\psi(\mathbf{r}_j^N, t_j)\dot{y}_\chi(t_j))$
        $\delta_V \leftarrow \delta_V + \xi_j z_\psi(t_{j+1}) - \mathcal{V}_\psi(\mathbf{r}_j^N, t_j)\dot{z}_\psi(t_j)$
        $j \leftarrow j + 1$
    until $j = J$
    average δ_P,δ_V over N trajectories to get $\overline{\delta}_P$, $\overline{\delta}_V$
    $\chi \leftarrow \chi - \alpha^\chi \overline{\delta}_P$
    $\psi \leftarrow \psi + \alpha^\psi \overline{\delta}_V$
    $i \leftarrow i + 1$
until $i = I$

explicit bath in Fig. 3.4(a), and then optimize these forces in the presence of the bath in Fig. 3.4(b). We start our optimization by first finding an arbitrary force that ensures the transition with a nonzero probability. We learn initial parameters $c_{pq}^{(i)}$ with a routine similar to well-tempered metadynamics.[76] Starting with $c_{pq}^{(i)} = 0$, at fixed frequency we add $c_{pq}^{(i)} \mapsto c_{pq}^{(i)} + \tau_m \omega T_m / [T_m + \omega \mathcal{N}(R,t)]$, where $\mathcal{N}(R,t)$ is a running histogram of order parameter R up to the current time t, and hyperparameters T_m, ω and τ_m determine how quickly the force landscape is filled. The blue learning curve in Fig. 3.4(a) refers to 100 steps of metadynamics run with $\tau_m = 10t/M_t$, $\omega = 4000$ and $T_m = 9000$. We find that the force solely from metadynamics is highly suboptimal compared to the rate bound, indicated by the black dashed line. Starting with a $\boldsymbol{\lambda}$ averaged over all metadynamics steps and with $\mathcal{V} = 0$, we next optimize both sets of parameters with MCVB-T and then MCVB each over 1000 steps with learning rates $\alpha^\chi = 40, \alpha^\psi = 200$ and $s = -100$. We find that the variational estimate converges tightly to the exact rate bound.

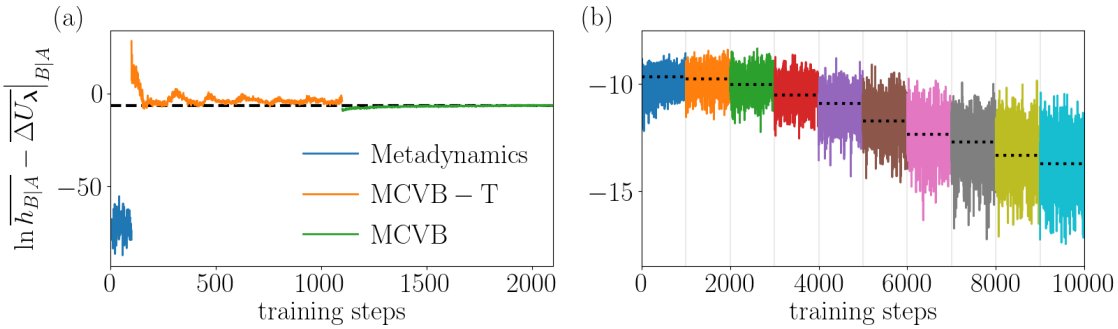

Figure 3.4: Learning curves for variational bound. (a) Optimization for the isolated dimer with 100 steps of well-tempered metadynamics(blue), 1000 steps MCVB-T(orange) and 1000 steps MCVB(green). Black dashed line is $\ln kt_f$ for the isolated dimer. (b) Learning curves for 1000 steps each, in presence of the explicit bath with $v_0\sigma/k_\mathrm{B}T = 0, 2, 4, ..., 18$. Black dotted lines denote the corresponding converged values.

Next we use the converged $\boldsymbol{\lambda}$ and $\mathcal{V}$ to start the optimization in presence of the bath, as illustrated in Figure 3.4(b). We successively optimize for each $v_0\sigma/k_\mathrm{B}T \in \{0, 2, 4, ..., 18\}$ starting from the previous converged result, each over 1000 steps. Each time we choose $(\alpha^\chi, \alpha^\psi) = (0, 200)$ for the first 200 steps and $(40, 200)$ for the remaining 800 steps. Learning the value function before starting to change the force in this way avoids a brief period of divergence at the beginning of each optimization run.[74] The results are robust towards changing the learning rates as long as α^ψ is kept about 5-10 times of α^χ, such that the value function is always approximately accurate whenever the force is being changed.

Results in Fig. 2 were also obtained similar to this protocol, but with no value function. For Fig. 2(c), the initialization parameters τ_m, ω and T_m are chosen at each barrier height so that at least half of the transitions are reactive.

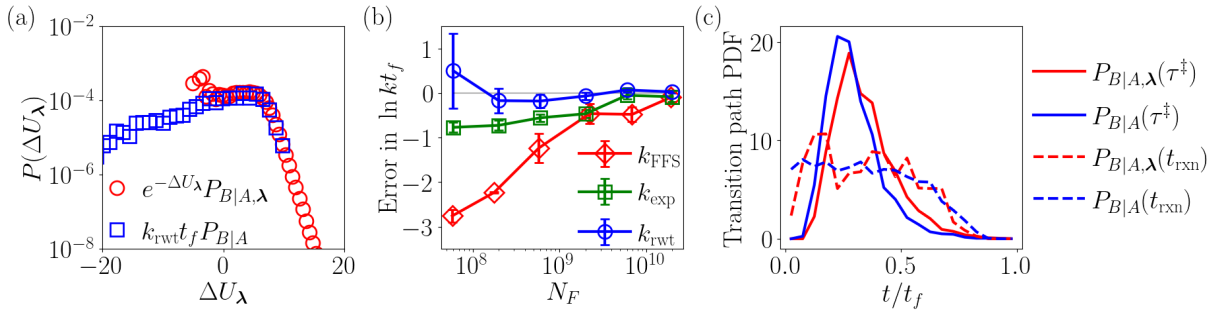

Figure 3.5: Unbiased rates, statistical convergence and transition path ensemble with VPS for $v_0\sigma/k_{\rm B}T = 8$

. (a) Overlap of the incomplete tilted biased and the unbiased distributions, with the scaling coefficient computed from Bennett Acceptance Ratio. (b) Errors in the rate estimates from Equations 5 ($k_{\rm exp}$) and 13 ($k_{\rm rwt}$) and from Forward Flux Sampling ($k_{\rm FFS}$)as the amount of total simulation timesteps N_F is varied. (c) Probability Density Functions (PDF) of transition path times and reactive escape times in the transition path ensemble, computed from the driven trajectories and unbiased reactive trajectories.

3.4 Unbiased reactive events from VPS

We use Eq. 5 and 13 in the main text to obtain rate estimates from direct simulations using the low-rank optimized force $\boldsymbol{\lambda}$. For the passive dimer in an active bath with $v_0\sigma/k_{\rm B}T = 8$, we have illustrated in Figure 3.5(a) overlap of the driven distribution $P_{B|A,\boldsymbol{\lambda}}(\Delta U_{\boldsymbol{\lambda}})$ with the unbiased distribution $P_{B|A}(\Delta U_{\boldsymbol{\lambda}})$ after tilting to correct the systematic error. The scaling constant $k_{\rm rwt}$, which is our estimate for the rate k, has been evaluated from Bennett Acceptance Ratio method[81] by using the tilting exponent as the reduced potential. This overlap is observable only when the driving force $\boldsymbol{\lambda}$ is near-optimal. If the tilted distribution does not contain enough statistics to represent the unbiased distribution, the estimate $k_{\rm exp}$ from Eq. 5 given by the area under the tilted distribution will underestimate the rate. If the basis set is complete and the exact optimal force $\boldsymbol{\lambda}^*$ can be obtained, $\Delta U_{\boldsymbol{\lambda}^*}$ will follow a Dirac delta distribution $P_{B|A,\boldsymbol{\lambda}^*}(\Delta U_{\boldsymbol{\lambda}^*}) = \delta(\Delta U_{\boldsymbol{\lambda}^*} + \ln kt_f)$, and the first cumulant will be sufficient to describe the log of the average of the exponential. This is also evident in Eq. 3.19 where $\boldsymbol{\lambda} = \boldsymbol{\lambda}^*$ makes $\Delta U_{\boldsymbol{\lambda}}[\mathbf{X}]$ trajectory independent. In that case, all three estimates of k from Eqs. 5, 7 and 13 will be equal and the unbiased reactive events will be entirely force-assisted rather than being driven by thermal fluctuations.

Figure 3.5(b) shows the systematic and statistical errors in $\ln kt_f$ calculated as VPS estimates $k_{\rm exp}$ and $k_{\rm rwt}$ from Equations 5 and 13 respectively in the main text. We have computed the errors by comparing to direct unbiased simulation, as the number N_w of uncorrelated trajectories of duration t_f is varied, expressed through the total number of

simulation timesteps $N_F = N_w t_f/\delta t$ where δt is a single timestep. Given the optimized driving force, $k_{\rm exp}$ is computed by averaging over N_w trajectories and $k_{\rm rwt}$ is obtained by reweighting $P_{B|A,\boldsymbol{\lambda}}(\Delta U_{\boldsymbol{\lambda}})$ and $P_{B|A}(\Delta U_{\boldsymbol{\lambda}})$ each computed with $N_w/2$ total trajectories, of which only a fraction are reactive without the driving force. At small N_F, $k_{\rm exp}$ systematically underestimates the rate due to the full area under $e^{-\Delta U_{\boldsymbol{\lambda}}} P_{B|A}(\Delta U_{\boldsymbol{\lambda}})$ not being accessible because of incomplete overlap, making $k_{\rm exp}$ formally unbiased but statistically biased. This error disappears with large N_F. However, the full rate can still be successfully obtained by comparing segments of incomplete distributions. Thus even when the undriven trajectory ensemble has ≤ 10 reactive trajectories at smaller values of N_F, $k_{\rm rwt}$ incurs much less error and provides a rate estimate that is both formally and statistically unbiased.

Figure 3.5(c) demonstrates convergence of the transition path ensemble obtained from direct simulations with the optimized forces even before the tilting correction. $P_{B|A,\boldsymbol{\lambda}}(\tau^\ddagger)$ and $P_{B|A}(\tau^\ddagger)$ are distributions of the transition path time $\tau^\ddagger$ measured as the time after leaving state A and before reaching state B without returning to A. $P_{B|A}(\tau^\ddagger)$ is from 2000 reactive trajectories obtained from 10^6 total unbiased simulated trajectories, while $P_{B|A,\boldsymbol{\lambda}}(\tau^\ddagger)$ is from a total of 2000 driven trajectories all of which were reactive. We find convergence in the distribution of transition path times signifying direct access to the nearly unbiased transition path ensemble by using the optimal force. Similarly we compare the distribution of the start time of the reaction $t_{\rm rxn} \in [0, t_f]$ measured as the time the trajectory last leaves A before arriving in B. We again find convergence in the driven ensemble compared to the unbiased reactive ensemble indicating the forces $\boldsymbol{\lambda}(\mathbf{R}, t)$ are near-optimal at all values of t.

The directly evaluated rates without additional forces used to compare VPS estimates have in most cases been computed from 5 trajectories, each of duration $10^4\tau$ with τ being the diffusive timescale. We compute k using Eq. 4 by evaluating the expectation with a rolling window over the trajectories after relaxing to a steady-state. We deviate from this protocol only for Figure 2(c), where the barrier heights are too large to estimate the rate from direct simulations. Here we use a numerically exact escape rate obtained from Kramer's theory.[85]

3.5 Comparison with Forward Flux Sampling

In Figure 3.5(b) we have compared the numerical cost of VPS at an active self-propulsion $v_0\sigma/k_{\rm B}T = 8$ with that of a Rosenbluth-like variant of Forward Flux Sampling (RB-FFS).[57] Starting from an ensemble of steady-state configurations in A, RB-FFS uses multiple interfaces between A and B to sequentially generate the transition path ensemble and compute the nonequilibrium reaction rate without an additional driving force.[86, 87] The transition paths generated from RB-FFS are unbranched and each has an associated weight as part of the transition path ensemble, analogous to VPS, from which the rate is estimated.

In order to apply RB-FFS, we define the interfaces along R as $R/\sigma \in \{1.25, 1.29, 1.33, 1.38, 1.43, 1.50, 1.57, 1.65, 1.72, 1.77, 1.81, 1.85\}$ with the first and the last interfaces corresponding to R_A and R_B respectively. We start RB-FFS trajectories from the same ensemble of steady-state configurations in A that we have used for VPS, and record the configurations

whenever the trajectories cross R_A from the A side. Every time a trajectory reaches B, we replace it in A at a random steady-state configuration. From each of $\mathcal{M}_0$ recorded configurations located at R_A, we generate reactive paths by shooting $\mathcal{M} = 100$ trajectories from each interface sequentially and randomly choosing one out of those that reach the next interface instead of coming back to A. The rate estimate is given by product of the forward flux of crossing R_A and the conditional probabilities of reaching subsequent interfaces, computed from an average over weighted reactive trajectories from the RB-FFS simulation[57]. We have varied $\mathcal{M}_0$ between 20 and 7000 to study the convergence of the RB-FFS rate estimate k_{FFS} as a function of the total number of simulation timesteps N_F, as shown in Figure 3.5(b). Statistical errors have been estimated over 3 independent parallel runs of the entire RB-FFS procedure.

We find that the VPS rate estimates k_{exp} and k_{rwt} incur much smaller errors than k_{FFS} at small N_F, though at large N_F all estimates converge to the same rate. Specifically, k_{rwt} converges to the true rate fastest among the three estimates, and k_{exp} incurs much smaller systematic errors than k_{FFS} even before convergence. This demonstrates that the use of the optimized force in a simple low-dimensional basis in VPS reduces the computational cost of estimating the exact rate by an order of magnitude or more compared to a trajectory stratification based method like RB-FFS. Further, we find that in RB-FFS, obtaining sufficient statistics given by a large $\mathcal{M}_0$ required the use of a long serial simulation to converge the flux of crossing the first interface at R_A. Parallelization of the $\mathcal{M}$ trajectory segments starting from each interface scaled poorly due to a broad distribution of durations over the trajectory segments, each of which must continue till they reach either the next interface, or A. Since shooting from the next interface can only start after the slowest of the previous trajectory segment has concluded, parallel implementations of RB-FFS scaled poorly and required overall a very long serial simulation. Our attempts to parallelize RB-FFS in an alternate fashion by reducing the serial configurations $\mathcal{M}_0$ worsened the systematic error in k_{FFS} even when averaged over fully independent RB-FFS implementations over many parallel threads. In contrast, every step of VPS is trivially parallelizable because of all trajectories being of the same duration t_f, which corresponds to $N_F = 2 \times 10^4$. As a result, a parallel implementation of VPS reduced its cost linearly and the overall computation required much shorter serial simulations.

3.6 Dissipative rate bounds

Stochastic thermodynamics provides a fundamental speed-limit on the enhancement achievable of the rate k of a rare nonequilibrium transition over a reference equilibrium rate k_0 in terms of the excess heat dissipation in the reactive path ensemble,[13]

$$\ln k \leq \ln k_0 + \frac{1}{2k_{\mathrm{B}}T}\langle Q \rangle_{B|A} \tag{3.20}$$

where $\langle Q \rangle_{B|A}$ is defined as the time-reversal asymmetric contribution from the change in path action between the equilibrium reference and the nonequilibrium system in which it

is measured. This bound holds under mild assumptions of instantonic or diffusive transitions and follows from a similar change of measure as leads to Eq. 7, with the additional observation of the time-reversal symmetric contribution of the change in path action being negligible near equilibrium. Here we show how to arrive at Eq. 14 employing this bound. As in the main text, we assume a separation of timescales between local relaxation and typical transitions so that the rate problem is well-posed.

We consider the rate enhancement afforded by coupling the dimer to an active solvent over the equilibrium passive bath isomerization rate. However, if we simply compute the excess heat dissipated as resulting from the time reversal asymmetric change in path action in turning v_0 from 0 to some finite value, the heat will be extensive in the number of solvent degrees of freedom and thus not have a well-defined thermodynamic limit. To mitigate this, we note that the isomerization rate of the dimer would be independent of v_0 if the dimer and solvent did not interact. Denoting k^{ni} and k_0^{ni} the rates of isomerization when the dimer is uncoupled to the solvent at finite or zero v_0, respectively, then $k^{\mathrm{ni}} = k_0^{\mathrm{ni}}$ and

$$\ln \frac{k}{k_0} = \ln \frac{k}{k^{\mathrm{ni}}} \frac{k_0^{\mathrm{ni}}}{k_0} \leq \frac{1}{2k_{\mathrm{B}}T} \left(\langle Q \rangle_{B|A} - \langle Q \rangle_{B|A,0} \right) \tag{3.21}$$

where $\langle Q \rangle_{B|A}$ is the excess dissipation resulting from turning on interactions between the dimer and solvent at finite v_0, and $\langle Q \rangle_{B|A,0} = Q_0$ results from turning on interactions between the dimer and solvent at $v_0 = 0$. The inequality is preserved even though a difference of heats is taken since the second ratio of rates k_0^{ni}/k_0 are both evaluated at equilibrium and thus the symmetric part of the action is zero. This second heat subtracts out the dissipation that is uncorrelated with the isomerization, and the remaining excess dissipation is left finite even when the number of solvent particles is large, so long as the dimer is correlated with the solvent over a finite distance.

The full path action for a system at finite v_0 in the presence of dimer-solvent interactions is

$$\begin{aligned} U_{v_0}(t_f) = -\frac{1}{4k_B T} \int_0^{t_f} dt\, \gamma_{\mathrm{d}}^{-1} \sum_{i \in \mathrm{d}} \left[\gamma_{\mathrm{d}} \dot{r}_i + \nabla_i V_{\mathrm{dw}} - \sum_{j \in \mathrm{d}} \mathbf{F}_{\mathrm{WCA}}(\mathbf{r}_{ij}) - \sum_{j \in \mathrm{s}} \mathbf{F}_{\mathrm{WCA}}(\mathbf{r}_{ij}) \right]^2 \\ + \gamma_{\mathrm{s}}^{-1} \sum_{i \in \mathrm{s}} \left[\gamma_{\mathrm{s}} \dot{r}_i - v_0 \mathbf{e}[\theta_i] - \sum_{j \in \mathrm{d}} \mathbf{F}_{\mathrm{WCA}}(\mathbf{r}_{ij}) - \sum_{j \in \mathrm{s}} \mathbf{F}_{\mathrm{WCA}}(\mathbf{r}_{ij}) \right]^2 \\ - \frac{1}{4D_\theta} \int_0^{t_f} dt \sum_{i \in \mathrm{s}} \dot{\theta}_i^2 \end{aligned} \tag{3.22}$$

and using the convention that v_0 is invariant under time-reversal,[49] the dissipated heat associated with turning on interactions between the solvent and dimer is

$$Q(t_f) = \int_0^{t_f} dt \left[\sum_{i \in \mathrm{d}} \sum_{j \in \mathrm{d}} \dot{\boldsymbol{r}}_i \mathbf{F}_{\mathrm{WCA}}(\mathbf{r}_{ij}) + \sum_{j \in \mathrm{s}} \sum_{i \in \mathrm{d}} \dot{\boldsymbol{r}}_j \mathbf{F}_{\mathrm{WCA}}(\mathbf{r}_{ji}) \right] \tag{3.23}$$

and is the same if $v_0 = 0$ or is nonzero. Since $\mathbf{F}_{\rm WCA}(\mathbf{r}_{ji}) = -\mathbf{F}_{\rm WCA}(\mathbf{r}_{ij})$ we find the definition of the Q in Eq. 15.

We measure $\langle Q \rangle_{B|A}(t_f)$ by averaging Eq. 3.23 over reactive trajectories of length $t_f = .2\tau$ sampled from long, 2×10^8 time-step, simulations in the nonequilibrium steady state at fixed v_0, with all other parameters as in the main text. Assuming transitions are uncorrelated, we compile Q samples from $24 - 96$ independent simulations, and use this data to calculate a mean and standard error, as depicted in Fig. 3 (red triangles).

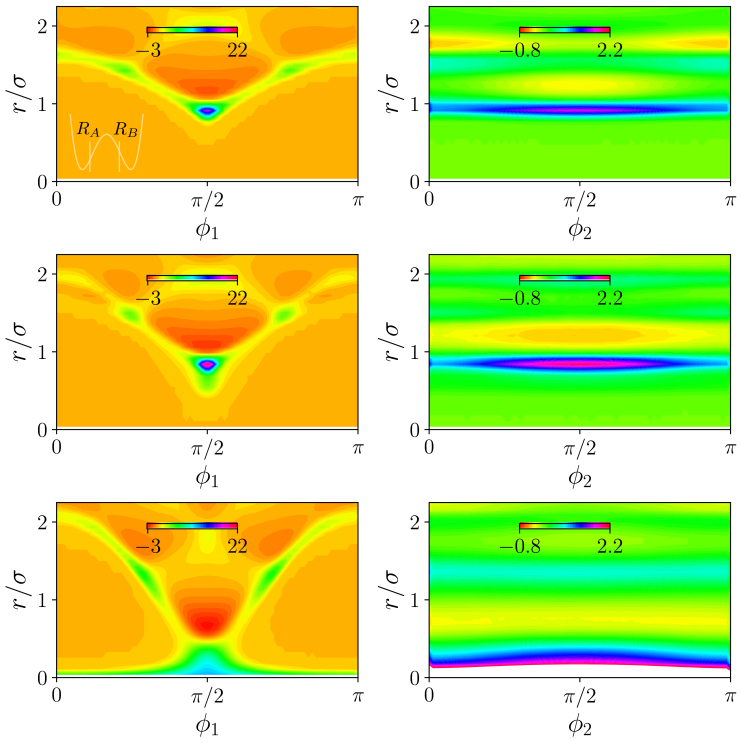

Figure 3.6: Solvation structure of the dimer by the active bath. Difference in pair distributions $\Delta g_X(r, \phi_1)$ (Left) and $\Delta g_X(r, \phi_2)$ (Right). Configurations are conditioned such that the bath is sampled with the dimer in the collapsed state $\Delta g_A(r, \phi_{1,2})$ (Top), the transition region $\Delta g_{AB}(r, \phi_{1,2})$ (Middle) or the extended state $\Delta g_B(r, \phi_{1,2})$ (Bottom).

3.7 Nonequilibrium solvation structure

We have studied how the solvation structure around the dimer evolves with activity. We compute the two-dimensional pair distribution functions for the position of the solvent around

the dimer bond within a conditioned steady state ensemble average

$$g_X(r,\phi_1) = \frac{1}{\rho_{\rm s}\rho_{\rm d}} \frac{\langle \sum_{i\in {\rm s}} \delta(\mathbf{r}^{\rm d}_{\rm cm})\delta(r-|\mathbf{r}_i-\mathbf{r}^{\rm d}_{\rm cm}|)\delta(\phi_1-\arccos(\mathbf{r}_i\cdot\mathbf{R}))h_X(R)\rangle}{\langle h_X(R)\rangle}, \tag{3.24}$$

with the center of the dimer bond $\mathbf{r}^{\rm d}_{\rm cm} = (\mathbf{r}_1+\mathbf{r}_2)/2$ as a reference. Here, r is the radial distance between $\mathbf{r}^{\rm d}_{\rm cm}$ and surrounding bath particles, which make an angle ϕ_1 with the bond vector $\mathbf{R} = \mathbf{r}_1 - \mathbf{r}_2$. The indicator function $h_X(R)$ restricts configurations where the bond length R falls into state X. Similarly, we probe the orientation of active solute particles around the dimer bond vector with the pair distribution

$$g_X(r,\phi_2) = \frac{1}{\rho_{\rm s}\rho_{\rm d}} \frac{\langle \sum_{i\in {\rm s}} \delta(\mathbf{r}^{\rm d}_{\rm cm})\delta(r-|\mathbf{r}_i-\mathbf{r}^{\rm d}_{\rm cm}|)\delta(\phi_1-\arccos(\mathbf{e}[\theta_i]\cdot\mathbf{R}))h_X(R)\rangle}{\langle h_X(R)\rangle}. \tag{3.25}$$

where ϕ_2 is the angle between a bath director and the dimer bond, and $\rho_{\rm s}\sigma^2 = 0.6$ is the density of the solvent, and $\rho_{\rm d}\sigma^2 = 0.008$ is the density of the dimers. To compute $g_X(r,\phi_i)$ we average over configurations sampled from 24 simulations each with a length of 2×10^8 time-steps. In the left column of Fig. 3.6, ϕ_1 denotes the angle between dimer bond center, and bath particle position, whereas ϕ_2 in the right column of Fig. 3.6 is the angle between dimer bond center the bath particle director.

In Fig. 3.6, we consider the change in the pair distributions $\Delta g_X(r,\phi_{i=1,2}) = g_X(r,\phi_{i=1,2},v_0 = 9) - g_X(r,\phi_{i=1,2},v_0=0)$ in an active bath with $v_0 = 9$ and its equilibrium counterpart at $v_0 = 0$. The different rows impose different conditions for the dimer bond distance $R = |\mathbf{R}|$ to be either primarily in state $X = A$ (top row), with $R < 1.55\sigma$, in the transition region $X = AB$ between states $1.55\sigma < R < 1.65\sigma$ (center row), or mostly in state $X = B$ (bottom row), with $R > 1.65\sigma$.

These results demonstrate that the rate enhancement is correlated with active particles dynamically wedging within the cross section of the dimer, pushing it apart into an extended state. The left column of Fig. 3.6 demonstrates that activity greatly enhances the packing of bath particles between the two bonded dimer particles, while the right column illustrates that bath particles preferentially orient perpendicular to the bond vector once far enough within the cross section. In state A, the active nature of the bath causes particles to push the dimer apart along R, as evidenced by the depletion for $\phi_2 = \pi/2$ and $r/\sigma > 1$ in the top-right panel of Fig. 3.6. The transition region, center-left, shows a significantly higher peak in the radial distribution function around $\phi_2 = \pi/2$, marking a decrease in the height of the effective free energy barrier along R. This analysis also illustrates the mechanism of increased stability in the active dimer extended state. Namely, Fig. 3.6 bottom-left shows that the driven bath particles act to inhibit the extended state from closing.

3.8 Conclusion

In conclusion, we developed a novel formalism and corresponding algorithm termed Variational Path Sampling to compute rate constants in nonequilibrium systems by optimally

driving the systems to transition between metastable states. VPS can be used to compute rates in arbitrary stochastic systems and extends the use of optimal control forces in large deviation sampling to transient rare events.[9, 67, 69, 88, 70] VPS complements trajectory-level importance sampling methods by generating the rare reactive event through a time-series of driving forces instead of a sequence of rare noise histories. We expect this approach to find broad use in rate computations for rare events in dissipative systems throughout the physical sciences and across scales.

We have shown that forces that generate unbiased transition paths in complex noisy systems in and out of equilibrium can be iteratively approximated with VPS. VPS can be used to sample with ease unbiased transitions and estimate rates which would otherwise be inaccessible to brute-forced simulations. In the future, we envision VPS will be used to investigate rare events and map their mechanisms in more complex biophysical simulations.

Chapter 4

Inferring equilibrium transition rates from nonequilibrium protocols: a general perspective on Bell's law

This work was done in a collaboration with David Limmer that resulted in the publication "Inferring equilibrium transition rates from nonequilibrium protocols." [89]

In chapter 2, we saw that there exists a fundamental relationship between the speed at which a driven system transitions states and the dissipation due to the driving, both through a path-action a transition-state theory view of the rate process. Here, we extend our transition-state theory picture to the case of an absorbing transition state and leverage our theory to infer intrinsic rate constants from driven trajectories in the context of single-molecule pulling experiments.

Extracting thermodynamic information from molecular systems through nonequilibrium processes was made possible with the revelation of Jarzynski's equality. [8, 34, 90, 91] However, inference of kinetic information, such as the intrinsic rate of molecular transitions, has remained more elusive.[92] Although useful ways of extracting transition rates from single molecule force data exist, they often rely on fitting to phenomenological expressions[93] or specifying a low dimensional model of the underlying conformational landscape. [94, 95] Such theories also typically assume that the driving forces are quasi-adiabatic, so that the molecule is assumed to locally equilibrate with the experimental forces imposed on it before a transition occurs. Here, we report that a molecule's bare equilibrium transition rate can be inferred from the statistics of the excess heat released during a nonequilibrium protocol. This result derives from expressions from stochastic thermodynamics[3] and an extension of transition state theory into nonequilibrium regimes.

One of the most common methods for extracting rate information from nonequilibrium

experiments and simulations employs Bell's law.[93] Bell's phenomenological rate law postulates that the speed of a molecular transition is accelerated with an applied external force by a factor that varies exponentially

$$k_\lambda \approx k_0 e^{\beta \lambda^\dagger x^\dagger}, \tag{4.1}$$

where k_λ is the rate in the presence of the external force λ, k_0 is the equilibrium rate, $x^\dagger$ is the distance along the forcing direction between the reactant state and a putative transition state, $\lambda^\dagger$ is the rupture force at the transition state and β is the inverse temperature times Boltzmann's constant. Evans and Richie showed that such a form emerges from Kramers theory for specific model potentials, within a high friction limit.[96] Dudko, Hummer and Szabo [97] developed alternative rate expressions from Kramers theory, as well as expressions for other experimentally observable quantities, and more recently, introduced a model-free method of estimating the force-dependent transition rate using statistics from the rupture force distribution.[95] Using a nonequilibrium response relation for reaction rates,[13] we provide a general perspective on the origin of Bell's law. We explore subsequent generalizations in a number of molecular systems with increasing complexity, and study the utility of model-independent rate estimates that depend only on the statistics of the dissipation, a thermodynamic quantity that is measurable experimentally.

4.1 Bell-type rate expressions from stochastic thermodynamics conditioned on absorbing boundaries

To start, we demonstrate how Bell's law can be understood as a consequence of two distinct approximations, a transition state theory approximation and a near equilibrium approximation. Within transition state theory, the rate of a transition between two metastable states is

$$k_\lambda = \nu p_\lambda(x^\dagger) \tag{4.2}$$

where ν is the probability flux through $x^\dagger$ and $p_\lambda(x^\dagger)$ the probability to reach a transition state, or dividing surface in phase space, starting from a metastable state.[98] This expression is valid for any value of λ, but is an approximation to the rate because ν in principle depends on $x^\dagger$, and errors associated with this approximation can be minimized with a judicious choice of dividing surface.[2] In equilibrium, the probability to reach a transition state is given by $p_0(x^\dagger) \propto \exp[-\beta F(x^\dagger)]$, where $F(x^\dagger)$ is the free energy to reach $x^\dagger$, resulting in the expected Arrhenius temperature dependence.

While transition state theory is still applicable to systems away from equilibrium,[19, 99, 13] the likelihood of reaching the transition state is not generally known, rendering it difficult to employ. For a system initially in equilibrium and acted upon by an external force, the nonequilibrium distribution is encoded in the Kawasaki relation,[100, 101, 6, 102]

$$p_0(\mathbf{r}) = p_\lambda(\mathbf{r}) \left\langle e^{-\beta Q} \right\rangle_{\lambda, \mathbf{r}} \tag{4.3}$$

where $p_0(\mathbf{r})$ is the initial equilibrium probability of full configuration $\mathbf{r}$, $p_\lambda(\mathbf{r})$ is the nonequilibrium probability, and the brackets $\langle \dots \rangle_{\lambda,\mathbf{x}}$ denote a trajectory ensemble average evolved under the driving force λ conditioned on ending at $\mathbf{r}$. The likelihood of the transition state, $p_\lambda(x^\dagger)$, is the marginalization of the full configurational probability onto the reaction coordinate. For any trajectory ending in rupture at time $t_{\rm rup}$, the excess heat dissipated to the environment, Q, over that from the conservative force is given by

$$Q(t_{\rm rup}) = \int_0^{t_{\rm rup}} dt\, \dot{\mathbf{r}}(t) \cdot \boldsymbol{\lambda}(\mathbf{r}, t), \tag{4.4}$$

which for a single molecule pulling experiment could be inferred from the force-extension curve.

In general, an applied time dependent force can change the mechanism of the transition. Even if the mechanism is conserved, both the height of the relevant barrier as well as the location of the transition state can be altered.[97] However, if the barrier to transitioning is large, we expect that the dominant change in the rate under an additional force is due to modulation of the transition state probability $p_\lambda(x^\dagger)$, leaving the location and flux through the transition state unchanged. Under these assumptions, we employ the Kawasaki relationship, Eq. 4.3, together with transition state theory, Eq. 4.2, to relate the transition rates in the presence and absence of λ,

$$k_0 \approx k_\lambda \left\langle e^{-\beta Q} \right\rangle_{\lambda, x^\dagger} , \tag{4.5}$$

to the statistics of the dissipated heat. In this limit, the Kawasaki response relation connects the transition rate amplification and the distribution of excess heat dissipated to the bath. A similar rate enhancement relation has been derived from path ensemble techniques.[13] While under the assumption of a conserved transition state this is the best estimate of the rate enhancement, converging the exponential average requires significant data. The equilibrium rate can be further approximated under an assumption of small applied force with a cumulant expansion. Expanding the logarithm of the relative rate for small values of the βQ,

$$\ln k_0 \approx \ln k_\lambda - \beta \langle Q \rangle + \frac{\beta^2}{2} \langle \delta Q^2 \rangle, \tag{4.6}$$

which is our first main result. For simplicity of notation we drop the explicit condition on the heat averages. Equations 4.5 and 4.6 imply that by measuring the rate of a rare event, as determined by a mean first passage time $k_\lambda = 1/\langle t_{\rm rup} \rangle$ to $x^\dagger$, and accompanying heat statistics in a driven system, we can infer the rate of a rare event in thermal equilibrium. These results can be considered as a nonlinear response theory for the rate,[103] in which frenetic contributions are neglected.[5] Equation 4.6 is similar to higher order corrections to Bell's law valid for constant applied forces,[104, 105] generalized to arbitrary protocols.

Near-equilibrium, $\beta \langle Q \rangle \lesssim 1$, and for slow driving forces, $\dot{\lambda} \approx 0$, the heat dissipated until crossing the transition state $x^\dagger$ is well approximated by $\langle Q \rangle_{\lambda, x^\dagger} \approx \lambda^\dagger x^\dagger$, where we have included only terms first order in λ and integrated Eq. 4.4 by parts. Such an approximation is good when the transition remains activated, so the transition path time is much shorter

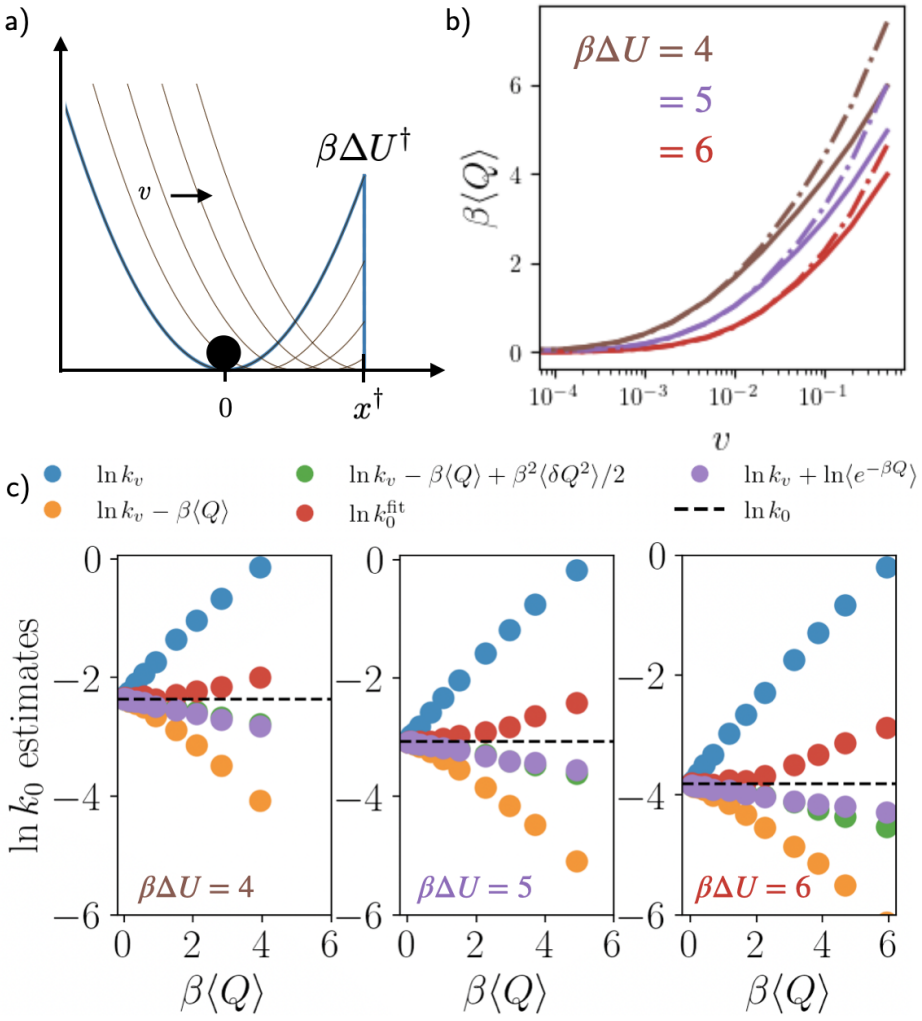

Figure 4.1: Pulling on a harmonic molecule. (a) A cartoon of the time dependent potential, initial barrier height $\beta\Delta U^\dagger$, and absorbing boundary condition, $x^\dagger$. (b) The heat dissipated gives Bell's law in the small loading rate, long rupture time limit. Dashed curves are Bell's law and the solid curves the excess heat for three potential stiffnesses. (c) Estimates of the equilibrium escape rate for increasing barrier heights and a range of pulling velocities, each with a typical heat as in (b).

than the time over which $\lambda(t)$ varies. Substituting this approximation for the average heat into Eq. 4.6 and neglecting higher order heat statistics, we find Bell's law.

4.2 Simple Harmonic Molecule

To understand the various rate inference expressions, we consider a hierarchy of models with increasing complexity. In each, we apply a simple force ramp, with constant velocity v so that $\lambda \propto v$, and we measure the rate under this driving protocol, $k_\lambda \to k_v$, as a mean first passage time to an absorbing boundary condition. The first model we consider is a simple overdamped particle trapped in a harmonic potential in one dimension, $\mathbf{r} \to x$. The equation of motion is

$$\dot{x} = \mu F(x) + \mu\lambda + \sqrt{2D}\eta \tag{4.7}$$

where the mobility μ and diffusivity D satisfy an Einstein relation $\beta D = \mu$, F is a conservative force, and η is a Gaussian random variable with $\langle\eta\rangle = 0$ and $\langle\eta(t)\eta(t')\rangle = \delta(t-t')$. The conservative force is $F = -a\kappa x$, and the particle is pulled with a linear ramp at loading rate v, such that $\lambda = a\kappa v t$ as depicted in Fig. 4.1a. We fix $\kappa = x^\dagger = \beta = \mu = 1$ and vary $v = \{0, .., 0.5\}$ and $a = \{8, 10, 12\}$. Simulations are run from an initial condition at the origin and stopped upon crossing $x^\dagger$ at $t = t_{\rm rup}$, where $k_v = 1/\langle t_{\rm rup}\rangle$. Time is measured in units of $\tau = \kappa/\mu$. Associated with the location of the absorbing boundary condition is an increased potential energy, equal to $\Delta U^\dagger = a\kappa x^{\dagger 2}/2$. We use a time step $t = 10^{-2}\tau$ and average over 10^4 trajectories.

In Fig. 4.1b) we verify the relationship between the heat and the argument of the exponential in Bell's law. Under a constant velocity force to lowest order in v, $Q \approx \kappa v t_{\rm rup} x^\ddagger$, which is plotted against the full expression for Q. As expected, for small v, in which $\beta\langle Q\rangle \lesssim 1$, both estimates agree. Note that $\langle Q\rangle$ does not vary linearly with v because $t_{\rm rup}$ depends implicitly on loading rate. Figure 4.1c) illustrates the convergence of the equilibrium rate employing different estimators, for a range of pulling velocities and barrier heights as $\langle Q\rangle \to 0$. The Bell's law like rate estimate, correcting the rate with just the mean heat, approaches the true equilibrium rate $\ln k_0$ from below. Within the validity of transition state theory, this behavior can be understood as a consequence of Jensen's inequality, applied to Eq. 4.5. Including additional dissipation statistics as in the full expression in Eq. 4.5 yields a faster convergence to the equilibrium rate for all barrier heights considered. Significant deviations from Eq. 4.5 occur when the dissipation is comparable to the size of the energy barrier, in which case the barrier is degraded enough for the event to no longer be rare. The agreement between the full exponential average of the heat and its second order expansion is a result of the linearity of the model studied.

From Bell's law, methods exist to extract the equilibrium rate from a forced measurement or simulation. A particularly accurate means of doing this is to perform maximum likelihood estimation (MLE) on a known form of the rupture force distribution,[106, 94, 107]

$$p_R(\lambda^\ddagger) = \frac{\beta k_0^{\rm fit}}{a\kappa v}\exp\left[\beta\lambda^\ddagger x^\ddagger - \frac{k_0^{\rm fit}}{a\kappa v x^\ddagger}(e^{\beta\lambda^\ddagger x^\ddagger} - 1)\right]. \tag{4.8}$$

to extract the equilibrium rate estimate $k_0^{\rm fit}$ and transition state distance $x^\ddagger$. With k, v and a fixed, we collect many samples of $\lambda^\ddagger$, the external force immediately after crossing the

transition state, and find the MLE parameters by minimizing $-\ln p_R(\lambda^\ddagger)$ with the SciPy minimize function, with initial guesses of -10 and 1 for $\ln k_0^{\text{fit}}$ $x^\ddagger$ respectively.[108] Other state of the art rupture force distributions exist, [109] but maximizing $p_R(\lambda^\ddagger)$ over k_0^{fit} and $x^\ddagger$ is a common method of inferring k_0 from single molecule data, and requires the least fitting parameters. We find k_0^{fit} from fitting the rupture force distribution is comparable in accuracy to estimates from Eq. 4.6. For this system, all estimators that employ corrections to Bell's law are accurate to within 30% across the full range of pulling speeds studied.

4.3 Pulling in multiple dimensions

We now consider a simple model often adopted in force spectroscopy studies to understand the role of flexible linkers. Specifically, we consider overdamped motion in 2 spatial dimensions,

$$\dot{\mathbf{r}} = -\boldsymbol{\mu}\nabla U(\mathbf{r}) + \boldsymbol{\mu}\lambda + \sqrt{2\mathbf{D}}\boldsymbol{\eta} \tag{4.9}$$

where $\mathbf{r} = \{q, x\}$, and q is envisioned as a measured extension that is coupled to the true molecular extension, x, through a potential

$$U(\mathbf{r}) = \Delta U x^2(x-2)^2 + \kappa_l(x-q)^2/2 + \kappa_s\mathbf{r}^2/2 \tag{4.10}$$

where ΔU denotes the height of the barrier in the molecular coordinate, while κ_l and κ_s are stiffnesses associated with the linker and trap, respectively.[110, 111, 112, 113, 114] In this simplest multi-dimensional model of a single molecule pulling experiment, the molecule undergoes diffusion in the 2-d landscape as shown in Fig. 4.2a. We perform force-ramp simulations with an added force $\boldsymbol{\lambda} = \kappa_s\mathbf{e}_\Theta vt$, parameterized by a pulling vector $\mathbf{e}_\theta = \{\cos(\Theta), \sin(\Theta)\}$ determined by the angle Θ relative to the q axis. The rate of heat dissipation is computed from $\boldsymbol{\lambda}\cdot\dot{\mathbf{r}}$. We fix $\beta = \mu_q = 1$, $x^\dagger = 1.5$, $\kappa_l = \kappa_s = 5$, $\beta\Delta U = 5$ and vary $v = \{0, .., 0.5\}$ and Θ. As before we estimate averages from 10^4 trajectories with an absorbing boundary condition at $x^\dagger$.

We first consider the experimentally relevant case of $\Theta = 0$, where the q direction is slow, $\mu_q/\mu_x = 1/20$ – the opposite case, when x is slow, being treatable analytically. [115] Note that when $\Theta = 0$, the heat is measured along the q direction only, and is therefore experimentally accessible. Under these conditions, shown in Fig. 4.2b, we observe that similar to the harmonic system, for small loading rates the dissipative second-cumulant estimate of Eq. 4.6 converges faster to the exact equilibrium rate $\ln k_0$, than either the bare driven rate or $\ln k_0 \approx \ln k_v - \beta\langle Q\rangle$ as $\langle Q\rangle \to 0$. As before, the estimate from the mean dissipation converges from below, however in this case the inclusion of the second cumulant results in convergence from above. The fit from the rupture force distribution is found to perform better than Eq. 4.6, but is comparable to the exponential average from Eq. 4.5, even for higher v. This is a consequence of the nonlinear system considered in this example, where the non-Gaussian heat statistics manifest an enhanced variance and nonvanishing higher order cumulants. Estimates from both Eq. 4.5 and Eq.4.8 provide accuracies within

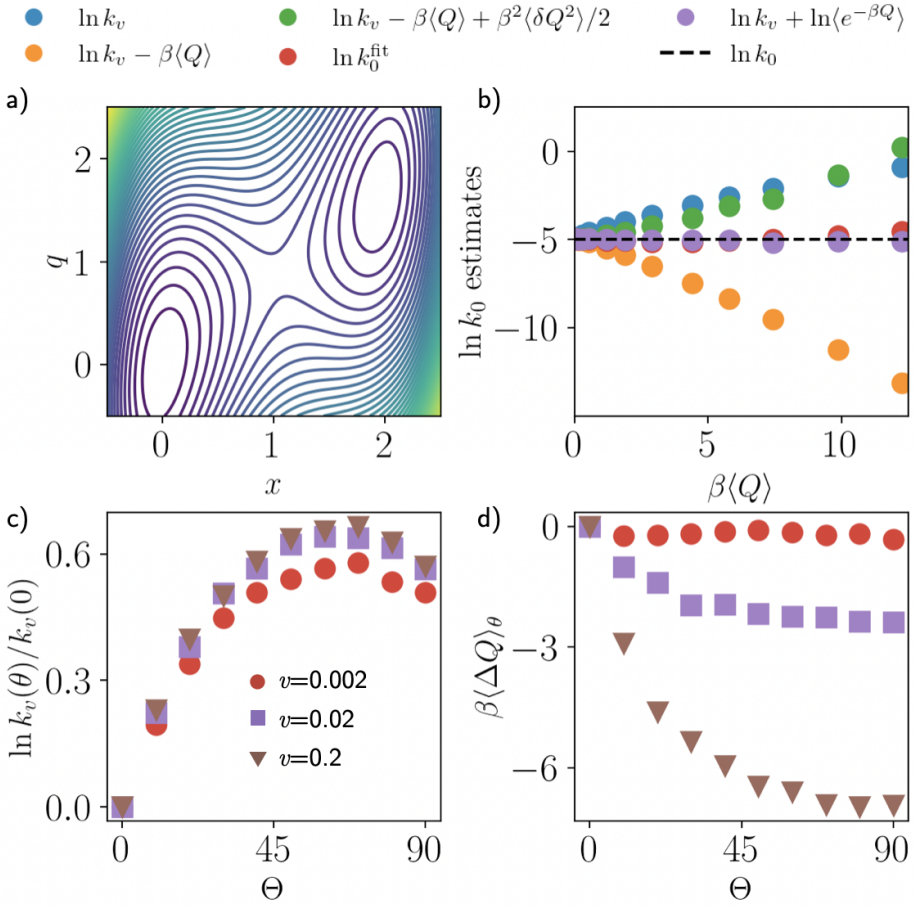

Figure 4.2: Protocol dependence. (a) The potential energy surface in the molecular x and observed q coordinates with absorbing boundary placed at $x^\dagger = 1.5$. (b) Rate estimates labeled as in Fig. 1c. (c) Modulation of the nonequilibrium rate with pulling velocity and angle relative to the q axis where the reference rate is taken as pulling along the q direction. (d) Modulation of the dissipated heat with pulling velocity and angle relative to the q axis labeled as in (c) and reference analogously.

5% across the full range of pulling velocities, as pulling orthogonally to x does not degrade the barrier, allowing for transition state theory to remain accurate.

In order to understand the protocol dependence of our rate inference, we now imagine that both pulling directions are accessible. While this is not typically experimentally practical, such a study provides insight into the convergence properties of the estimators. We consider the case $\mu_q/\mu_x = 1$ and vary Θ and v in Figs. 4.2c and 4.2d. As expected, pulling along coordinate x, $\Theta = 90^\mathrm{o}$, results in a faster process relative to $\Theta = 0^\mathrm{o}$, and one which dissipates less heat as quantified by $\Delta\langle Q\rangle_\Theta = \langle Q\rangle_\Theta - \langle Q\rangle_0$. The rate within the driven dynamics is

maximized near $\Theta = 60^{\circ}$, which is in reasonable agreement with the optimal transition state predicted by multidimensional transition state theory.[116] For small v, near equilibrium, the dissipated heat approaches zero independent of Θ as expected for a quasi-reversible process. These results suggest that in systems with multiple spatial dimensions, pulling along any direction may be sufficient to estimate k_0, but geometries that minimize the heat until rupture may lead to faster convergence of Eq. 4.5 and Eq. 4.6.

4.4 Pulling a semi-flexible polymer

As an example of a non-linear many-particle system, we consider stretching a semi-flexible polymer with reactive ends[117] that attract each other with a strong, short-ranged potential. The configuration of the polymer consists of $3N$ dimensions, $\mathbf{r} = \{\mathbf{r}_1, \ldots, \mathbf{r}_N\}$, and evolves with underdamped Langevin dynamics

$$m\ddot{\mathbf{r}}_i = -\mu^{-1}\dot{\mathbf{r}}_i - \boldsymbol{\nabla}_i U(\mathbf{r}) + \boldsymbol{\lambda} + \mu^{-1}\sqrt{2D}\boldsymbol{\eta_i}. \tag{4.11}$$

where m is the mass of a monomer. The monomers interact through a potential that consists of $U(\mathbf{r}) = U_{\mathrm{b}}(\mathbf{r}) + U_{\mathrm{a}}(\mathbf{r}) + U_{\mathrm{nb}}(\mathbf{r})$ where $U_{\mathrm{b}}(\mathbf{r})$ is a harmonic bond potential between adjacent monomers

$$U_{\mathrm{b}}(\mathbf{r}) = \frac{\kappa_{\mathrm{b}}}{2}\sum_{i=1}^{N-1}\left(\mathbf{r}_{i+1} - \mathbf{r}_i\right)^2 \tag{4.12}$$

with stiffness κ_{b} and $U_{\mathrm{a}}(\mathbf{r})$ is a harmonic angular potential that penalizes bending

$$U_{\mathrm{a}}(\mathbf{r}) = \frac{\kappa_{\mathrm{a}}}{2}\sum_{i=2}^{N-1}\left(\frac{\mathbf{r}_{i+1,i}\cdot\mathbf{r}_{i,i-1}}{r_{i+1,i}r_{i,i-1}} - 1\right)^2 \tag{4.13}$$

where $\mathbf{r}_{i,i-1} = \mathbf{r}_i - \mathbf{r}_{i-1}$ is the vector between two adjacent monomers, $r_{i,i-1}$ denotes its magnitude, and the potential has a stiffness κ_{a}. The nonbonding potential $U_{\mathrm{nb}}(\mathbf{r})$ has a short ranged form,

$$U_{\mathrm{nb}}(\mathbf{r}) = \frac{1}{2}\sum_{i\neq j=1}^{N}\epsilon_{ij}\left[5\left(\frac{\sigma}{r_{i,j}}\right)^{12} - 6\left(\frac{\sigma}{r_{i,j}}\right)^{10}\right] \tag{4.14}$$

where σ sets the characteristic size of a monomer and ϵ_{ij} the interaction strength between monomer i and j. In order to model a reactive linker, we set the two monomers on each end to interact more strongly than the monomers in the interior of the polymer, and the cross interactions, between the linker and interior monomers, are neglected. We hold the 1st monomer fixed at the origin, so the end to end vector is $\mathbf{R_{EE}} = \mathbf{r_N}$, and pull the two ends apart with $\boldsymbol{\lambda} = -\kappa(\mathbf{R_{EE}}\cdot\hat{\mathbf{x}} - vt)$, where $\hat{\mathbf{x}}$ is the unit vector in the x direction. Characteristic snapshots are shown in Fig. 4.3a). We adopt a unit system with $\beta = m = \sigma = 1$, with a natural time-scale $\tau = 1/\sqrt{\beta m\sigma^2}$. We set $\mu = 5$, $\kappa_{\mathrm{b}} = 110$, $\kappa_{\mathrm{a}} = 4.5$, $\kappa = 5.5$, end monomer

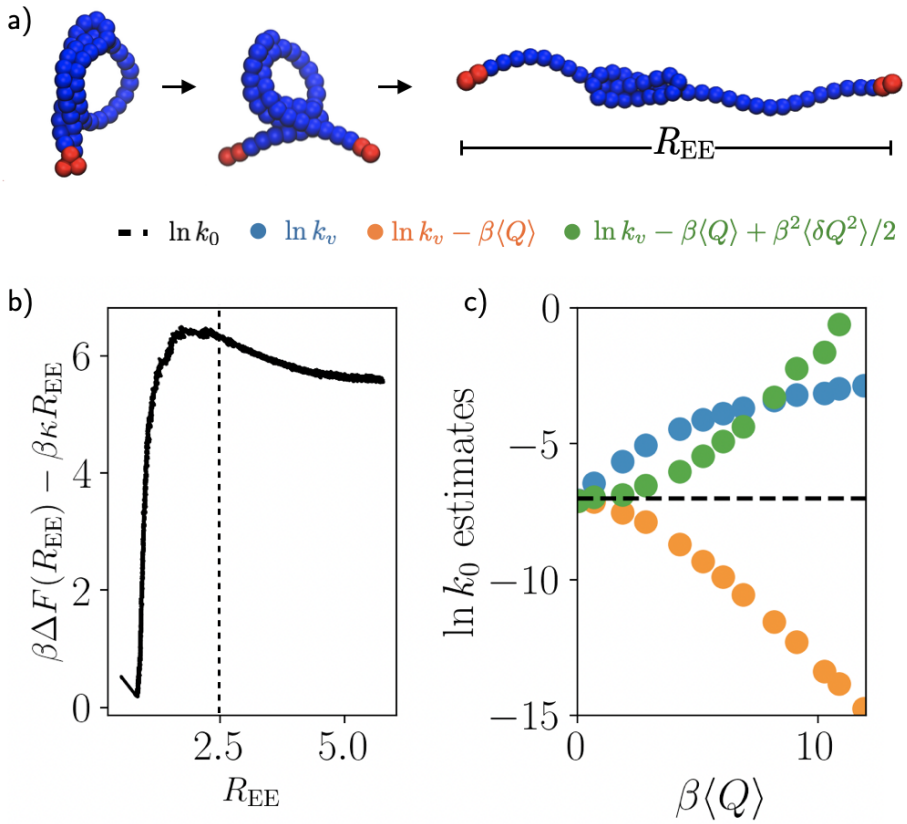

Figure 4.3: (a) A semi-flexible polymer with reactive ends, highlighted in red, mechanically unfolded to a large end to end distance during a force ramp experiment. (b) The free energy of the untethered polymer as a function of the end to end distance. The black dashed line denotes the absorbing boundary condition. (c) Rate estimates, labeled as in Fig 1c, as a function of heat.

interactions $\epsilon_{1,N} = 4.5$, and interior monomer interactions $\epsilon_{3,N-3} = 1.7$, $N = 50$. We employ a time step of $10^{-2}\tau$ and estimate averages from 500 trajectories.

Under these conditions, the semi-flexible polymer permits two types of conformations– a folded state for small $R_{EE} = |\mathbf{R_{EE}}|$ where the linker monomers are bound, and an unfolded state for large R_{EE} where the linker monomers do not interact strongly. To differentiate between these two regimes, we first computed the work to pull R_{EE} reversibly, denoted by $\beta\Delta F(R_{EE})$. This is shown in Fig. 4.3b), and evaluated using the Jarzynski equality[8] within a steered Brownian dynamics framework.[118] For the free energy calculation we employed the same constant v protocols, including all of the simulation data shown in Fig. 4.3c). The free energy exhibits a deep minimum for small end to end distances and a second shallow basin for large end to end distances. Under a small additional load, $-\kappa R_{EE}$, the two basins

are separated by a barrier in the free energy. Using this biased free energy we set an absorbing boundary condition for our pulling calculation to ${R_{\rm EE}}^\dagger = 2.5\sigma$.

We pulled the polymer at loading rates over the interval $v = \{0.001 - 0.2\}\sigma/\tau$ and measured the dissipative heat and first passage time to ${R_{\rm EE}}^\dagger$. Shown in Fig. 4.3c), we find that convergence to the true equilibrium rate is fastest using Eq. 4.6, and as in the 2d model the estimate converges from above upon the approach to equilibrium. As in both previous models, the incorporation of the variance of the heat provides an accurate estimate of the equilibrium over a range of heats that is comparable to a significant fraction of the native barrier, in this case as long as $\beta\langle Q\rangle < 3$. The first order estimate converges to the true equilibrium rate slowly from below. Over the range of pulling velocities considered, if either the first-order estimate or the bare rate were fit and linearly extrapolated to $v = 0$, both would overestimate the equilibrium rate by about an order of magnitude.

4.5 Conclusion

Our results demonstrate an underlying stochastic thermodynamic basis for Bell's law under nonequilibrium driving and a useful means for going beyond it to infer equilibrium transition rates. Within the context of single molecule force ramp experiments, we have demonstrated a robust way to infer unfolding rates using the statistics of the heat distribution, conditioned on ending at an absorbing transition state. While within transition state theory, the full exponential estimator is most accurate, we suspect that in general its convergence will be cumbersome, and the perturbative expansions illustrated here will provide an intermediate means of rate estimation. In the future, this response method may be used to study the rare kinetics of more detailed protein models, protein unfolding in optical tweezing and atomic force microscopy experiments, and other rare molecular transitions that can be sped up by applied force.[119, 120, 121, 122] The nonequilibrium thermodynamic framework developed here works not only with constant velocity force ramps, but could be used with more complex protocols. Indeed, protocols can be optimized to allow for rate inferences,[44] or could be used as a theoretical framework for understanding other approximate methods that use driven dynamics to infer rates with applied force.[123, 124, 125, 126]

We have shown how forces influence transition rates between states, but it is often the case in experimental settings that we want to use transition information to learn about the nature of hidden states. In what is to come, we will explore how single molecule unfolding trajectories can be used to infer information about hidden intermediate states on the route to folding and unfolding.

Chapter 5

The hidden dimensions of single molecules

The work in this chapter was done in collaboration with Robert Sosa, JeongHoon Kim, Alex Tong and Carlos Bustamante.

Single molecule force spectroscopy, and in particular optical tweezing, has changed the way we study molecular function, shifting the paradigm from ensemble bulk assays to trajectories of individual molecules. [127, 119] In optical tweezing experiments, a single molecule, a protein, is tethered via DNA linkers on either end to two micro-scale beads trapped by laser light. By manipulating the lasers, the beads can be pulled apart while the distance between them and the restoring force exerted by the molecule are measured. Pulling on the protein until the bond network holding it together in its native folded structure breaks allows for the exploration of how the protein of interest unfolds under physiological forces and refolds in their absence. [128, 129, 130] Here, we probe the intermediate states of a globular protein on its unfolding path to unravel how its secondary structure coordinates to allow it to properly fold.

5.1 Foldon theory of protein folding

The traditional model of protein folding is that of the entropic funnel. [131] In this 'funnel' model, the unfolded protein starts in a state of high conformational entropy, which decreases to a global minimum as the protein travels along any one of an infinite number of paths to its native folded state. This high number of paths allows the protein to initiate folding from any section along its linear sequence of amino acids and proceed with transitions of any other section of its sequence after that.

In recent years, a refinement of the funnel model called the 'foldon' model of protein folding has been proposed and verified in bulk assays. [132, 133] The foldon model differs

from the classic funnel model in that it posits that there exist at most a handful of probable paths the protein traverses from its unfolded to native states. Instead of starting at any location along its peptide sequence, in the foldon model, folding begins at the location of a particular foldon, a $\sim 20-30$ amino acid long section comprised of a few secondary structural elements (alpha helices and beta strands). Folding of the first foldon triggers a cascade in which each subsequent foldon acts as a template to nucleate folding of the next foldon even at long ranges along the sequence. Critically if one foldon is knocked out, its inability to fold disrupts the cascade, hindering the folding of the rest of the foldons. [134] In this way, a protein in the foldon model travels on one of few deep valleys in the entropic funnel, along a distinct set of well defined intermediates on its way to its native state.

5.2 Probing unfolding intermediates in a globular protein

We aim to dissect the (un)folding pathway of a small highly stable single domain globular protein called Ross. [135, 136] Ross gets its name from its Rossmann fold, which is found in most DNA binding domains and consists of a four-strand beta sheet wedged between four alpha helices, as shown in Fig 5.1 a. To study (un)folding paths, we locate and examine the intermediates along the transition path of the protein as it is forced open during force ramp experiments, in which the optically trapped beads are pulled apart to a large fixed extension and brought back together at constant speed. [137]

The raw data is typically visualized with a force-extension curve, as shown in Fig. 5.2 a. To transform from extension to protein contour length L_c, we need a model that links force F and extension x. The Marko-Siggia model

$$\beta F(x)P = \frac{1}{4}\left(1 - \frac{x}{L} + \frac{F}{K}\right)^{-2} - \frac{1}{4} + \frac{x}{L} - \frac{F}{K} \tag{5.1}$$

works well at describing such a relationship in the context of stretching a worm-like chain such as DNA. [138] This model relies on two material fitting parameters: the persistence length of the polymer P and its stretch modulus K. In the case that the polymer is highly extended under large force, Eq. 5.1 can be inverted [139] to give

$$\frac{x}{L} = x^{WLC}(F) = 1 - \frac{1}{2\sqrt{\beta F P}} + \frac{F}{K}. \tag{5.2}$$

When the protein is folded, the force extension curve reflects the stretching of bare DNA and can be fit to Eq. 5.2, $x^{WLC}_{folded} = x^{WLC}_{DNA}$ with parameters L_{DNA}, P_{DNA} and K. The unfolded extension is that of the DNA and the protein combined

$$x_{unfolded}(t) = L_{DNA}x^{WLC}_{DNA}(P_{DNA}, K) + L_p x^{WLC}_p(P_p), \tag{5.3}$$

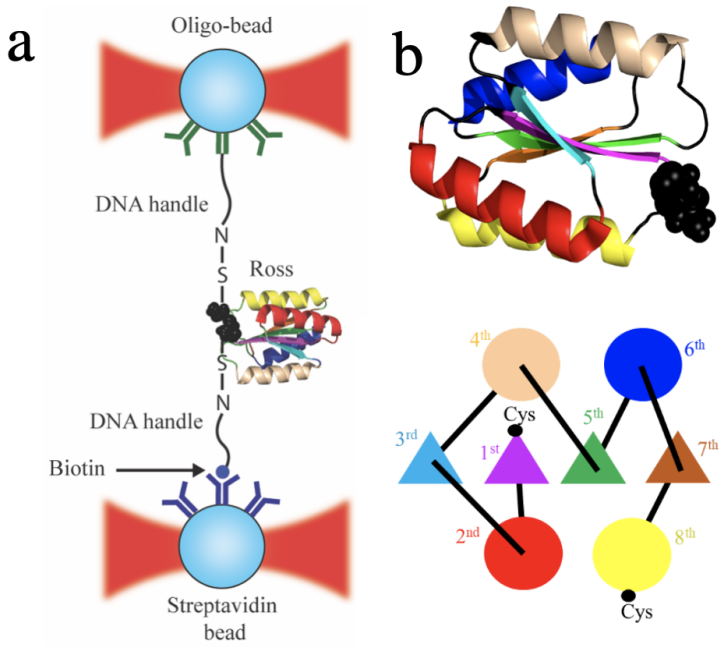

Figure 5.1: (a) Optical tweezers set-up: a single molecule tethered to DNA handles is trapped between two micro-scale beads held in place by laser light. The molecule is unfolded by force as the lasers are separated and the beads pulled apart. (b) The synthetically designed Ross protein. Atomic structure from the protein data base (top) and a visualization of the Rossmann fold secondary structure (bottom). Beta strands are shown as triangles and alpha helices as circles. We refer to secondary structures as numbers 1 (purple inner beta strand) to 8 (yellow outer alpha helix).

where we take the $K_p \to \infty$ limit as the protein is non-extensible. Once the five free parameters L_{DNA}, P_{DNA}, K, L_p and P_p are fit, we have described the typical force extension relationship for the folded and unfolded states, shown as blue and red dashed curves in Fig. 5.2 b. Along the transition path that connects the unfolded state to the folded state, the contour length of the protein varies $L_p \to L_c(t)$ and can be solved for explicitly from Eq. 5.4 as

$$L_c(t) = \frac{x_{unfolded}(t) - L_{DNA} x_{DNA}^{WLC}(F)}{x_p^{WLC}(F)}. \tag{5.4}$$

The contour length along the transition path corresponding to the rip in Fig. 5.2 a is shown in Fig. 5.2 b and exhibits two well defined intermediate states, which we will call I1 and I2. We separate the folded and unfolded portions of the force extension data by the temporal location of the transition, found by applying a step-finding algorithm [140] to the (folded) contour length $x/x^{WLC}(F)$, where x and x^{WLC} include the entire extension and force time-

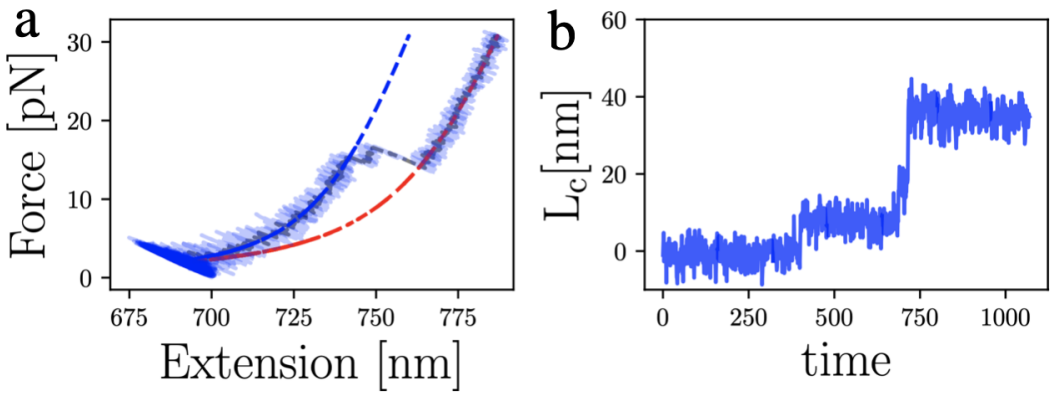

Figure 5.2: Analysis of a single molecule unfolding trajectory. (a) A force extension curve. Fits to wormlike chains, DNA and combined DNA and protein, are shown as dashed blue and red curves, respectively. (b) Contour length of the protein as a function of time along the transition path from folded to unfolded states, from transforming force-extension data with worm-like chain polymer theory.

series, respectively, and x^{WLC} is evaluated with parameters $P = 50nm$ and $K = 900pN^{-1}$ typical of DNA. [141]

Each cycle of pulling and relaxing is analysed in this way to produce a contour length transition path, which we histogram to arrive at the transition path probability $p(L_c)$. [142, 143] The transition path probability for wild-type ROSS is shown in Fig. 5.3 (blue).

5.3 Dynamics of structure based protein models

The dynamics of proteins with known atomic structure can be studied on long timescales with molecular dynamics of structure-based models (SBMs). The most simple course grained SBMs work at the resolution of single amino acids, the alpha-carbons that comprise the protein backbone, and posit that particles experience vibrations around equilibrium bonds and angles, a torsional potential, attraction between native contacts which take the place of a hydrogen bond network in the protein and steric repulsion between particles not in native contact with one another. To simulate ROSS, we employ SMOG to generate the course grained coordinates and native contacts from the all atom equilibrium configuration of the protein.[144] We then feed the coordinates and contacts to the open source molecular dynamics software SMBopenmm from which we can run customized simulations.[145] We run over-damped Brownian dynamics

$$\dot{\mathbf{x}}_i(t) = \mu \mathbf{F}_i[\mathbf{x}(t), t] + \boldsymbol{\eta}_i(t) \tag{5.5}$$

$$\langle \eta_{i\alpha}(t)\eta_{j\beta}(t')\rangle = 2\mu k_{\mathrm{B}} T \delta_{ij}\delta_{\alpha\beta}\delta(t-t') \tag{5.6}$$

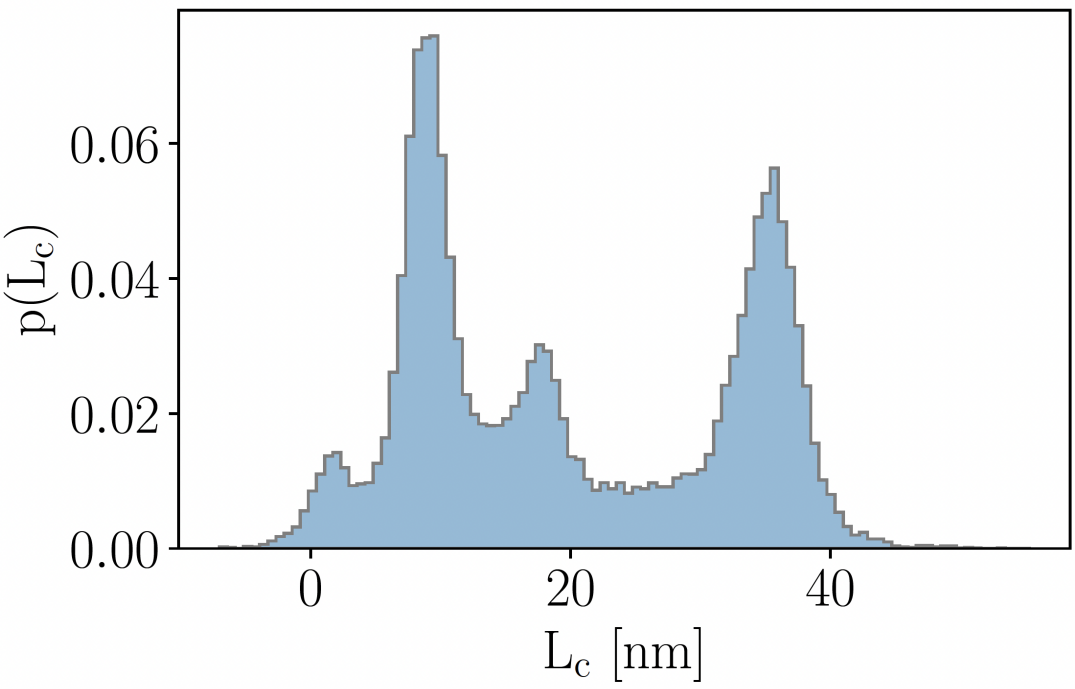

Figure 5.3: Unfolding transition path distributions of wild-type Ross

$$\mathbf{F}_i = -\nabla_{\mathbf{x}_i} V(\mathbf{x}) - \delta_{\alpha,1}\delta_{i,1}k(\mathbf{x}_{11} + vt) - \delta_{\alpha,1}\delta_{i,N}k(\mathbf{x}_{N1} - vt) \tag{5.7}$$

to pull the protein apart from either end along the x-axis. As in the SMBopenmm documentation, the total potential is

$$V(\mathbf{x}) = \sum_{bonds} V_{bond} + \sum_{angles} V_{angle} + \sum_{torsions} V_{torsion} + \sum_{contacts} V_{contact} + \sum_{non-contacts} V_{non-contact}, \tag{5.8}$$

where

$$V_{bond} = \frac{k_b}{2}(r - r_0)^2, \tag{5.9}$$

$$V_{angle} = \frac{k_a}{2}(\theta - \theta_0)^2, \tag{5.10}$$

$$V_{torsion} = k_t(1 - \cos(\phi - \phi_0)) - \frac{1}{2}(1 - \cos(3(\phi - \phi_0))), \tag{5.11}$$

$$V_{contact} = \epsilon_c \left(5 \left(\frac{\sigma_{ij}}{r_{ij}} \right)^{12} - 6 \left(\frac{\sigma_{ij}}{r_{ij}} \right)^{10} \right), \tag{5.12}$$

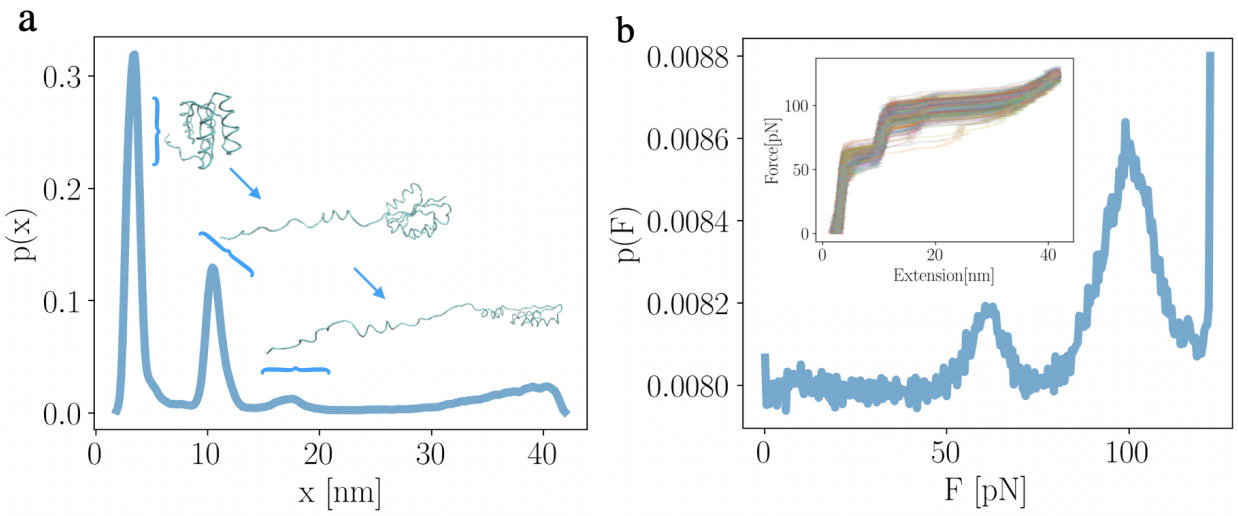

Figure 5.4: Ross extension (a) and force (b) transition path distributions exhibit two intermediate states, as observed in experiments. Typical conformations in each intermediate shown as insets (a). Force extension traces inset in (b).

and

$$V_{non-contact} = \epsilon_{nc} \left(\frac{\sigma_{nc}}{r_{ij}} \right)^{12}. \tag{5.13}$$

Here r_0 is the equilibrium bond distance between adjacent particles, θ_0 is equilibrium angle between three particles, ϕ_0 is the equilibrium dihedral angle between 4 particles, and σ_{ij} is the equilibrium contact distance between particles i and j that are in contact in the native state. The parameters r_0, θ_0 and ϕ_0, as well as the contacts ij are set by the reference native structure of the protein, in this case from the protein data base. The parameters k_b, k_a and k_t are the bond and angle stiffnesses, and ϵ_c and ϵ_{nc} the repulsive and attractive energy scales that determine the interaction between particles not in native contact and those that are in native contact, respectively. The parameters are set to predetermined values: $k_b = 2 \times 10^4 kJ/(mol \times nm^2)$, $k_a = 40 kJ/(mol \times rad^2)$, $k_t = 1 kJ/(mol)$, $\epsilon_c = \epsilon_{nc} = 1 kJ/mol$, $\sigma_{nc} = 0.4nm$. For our simulations, we set the temperature to $T = 300k$, take a time step of $\Delta t = 0.5fs$, a relatively weak trap spring constant $k = 0.1 kJ/(mol)$ and a fast pulling speed $v = 2 \times 10^1 2nm/s$, so that the center of the traps are displaced by $60nm$ over 10^4 time steps. Though v is extremely large, with such small trap spring constants the protein takes $\sim 10^8 \Delta t = 100ns$ to rupture.

We perform constant speed force ramp experiments with harmonic traps, beginning each trajectory in the equilibrium native state ensemble and ending once the protein has been pulled apart to a predefined length, $40nm$. We histogram these transition paths, as shown

in Fig. 5.4, and confirm the presence of the two intermediates we found in experiments.

We examine the simulated trajectories and show representative snapshots of the observed intermediates in Fig. 5.4. Specifically, we see that the first intermediate corresponds to the tethered 8th secondary structure, the outside alpha helix shown in yellow in Fig. 5.1, unfolding. We next observe that the second, shorter lived intermediate corresponds to the 7th and 6th secondary structures, an outside beta strand (brown) and the alpha helix next to it (blue), unfolding. From here our simulations suggest that the rest of the secondary structures comprising ROSS unfold cooperatively, all at once. This unfolding pathway seems plausible considering that the tension on the protein is applied to the 1st and 8th secondary structures. This would perhaps lead one to guess the 1st or 8th secondary structures would unravel first, but since the 1st secondary structure is part of a beta sheet in the core of the protein, it would have to be sheared to unfold first, which requires significantly more applied force than ripping an outside alpha helix off the protein.

5.4 Secondary structure contact indicators

A proteins' native contacts to large part determine the paths it takes both to collapse into its native state and be pulled out of it by force. [146, 147, 148, 149] To quantitatively track the movement of individual secondary structures within the protein, we compute a switching function which reports on the percent contact for each amino acid pair in the protein. The switching function

$$f(r_{ij}) = \frac{1}{1 + e^{-\delta(r_{ij} - r^*)}} \tag{5.14}$$

$f(r_{ij})$ interpolates between $f(r_{ij}) \approx 1$ when amino acid i is within a radius $r_{ij} < 1nm$ from amino acid j and $f(r_{ij}) \approx 0$ when $r_{ij} > 2.5nm$ for $\delta = 5nm^{-1}$ and $r^* = 1.7nm$.

We then calculate what percent of contact each secondary structure σ_1 is with every other secondary structure σ_2 as

$$\%_c(\sigma_1, \sigma_2) = \frac{\sum_{i \in \sigma_1} \sum_{j \in \sigma_2} f(r_{ij})}{\sum_{i \in \sigma_1} \sum_{j \in \sigma_2} 1}. \tag{5.15}$$

In Fig. 5.5 we show heat maps of Eq. 5.16 at the beginning (left) of a force ramp trajectory when the protein is folded in its native structure, and towards the end of the force ramp when most of the native contacts have been broken. It is our hope that by visualizing how the secondary structural contacts change in time as the protein is unfolded at low enough pulling speed, we will be able to detect groups of secondary structures that loose or maintain contact together. These secondary structural groups would be our purported foldons.

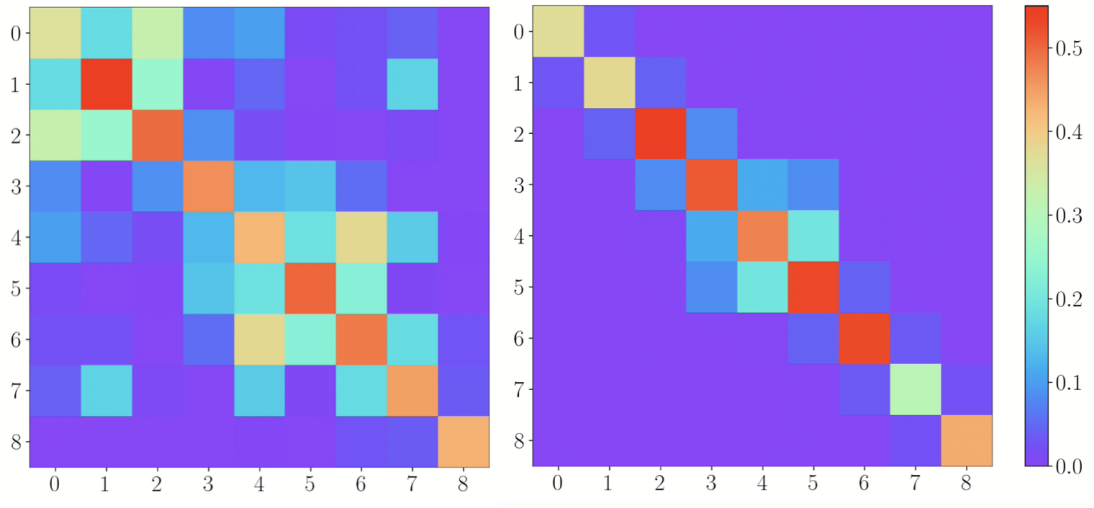

Figure 5.5: Secondary structural percent-contact maps at the beginning (left) and close to the end (right) of a fast force ramp trajectory.

5.5 Work propagation during forced extension reveals secondary structure dynamics

In order to visualise unfolding of different parts of Ross along a single dimensional coordinate, we studied how work propagates through the alpha-carbon bonds connecting the back bone of the protein under increasing tension. To compute the work transduced through the bonds, for each bond i in the protein,

$$\beta \dot{W}_i(t) = \frac{f_i(t) + f_i(t + \Delta t)}{2}(r_i(t + \Delta t) - r_i(t)) \tag{5.16}$$

with distance r_i, we measure how it extends or shortens over a single timestep as well as the net force f_i along the bond. We average trajectories $\dot{W}_i(t)$ as the protein is unfolded by force to visualize how dissipation travels along the backbone in real time, with the expectation that as a portion of the protein is unfolded and its constituent bonds stretched, the large restoring forces should increase the power delivered through those bonds.

This is illustrated in Fig. 5.6, where we show typical snapshots of $\beta \dot{W}_i(t)$ (top) and the protein conformation (bottom) as it is pulled open from left to right. The vertical dashed lines in Fig. 5.6 mark the boundaries between the secondary structures of the protein. As the protein begins to unfold (left) from the 8th helix, the corresponding dissipation along its alpha carbon bonds begins to increase. As those 8th helix bonds are increasingly stretched while the adjacent secondary structure remains primarily folded, the power through them increases sharply (middle). At the same time (middle), at the extremely high loading rate at which these force ramp experiments were performed, the 1st (beta strand) and second

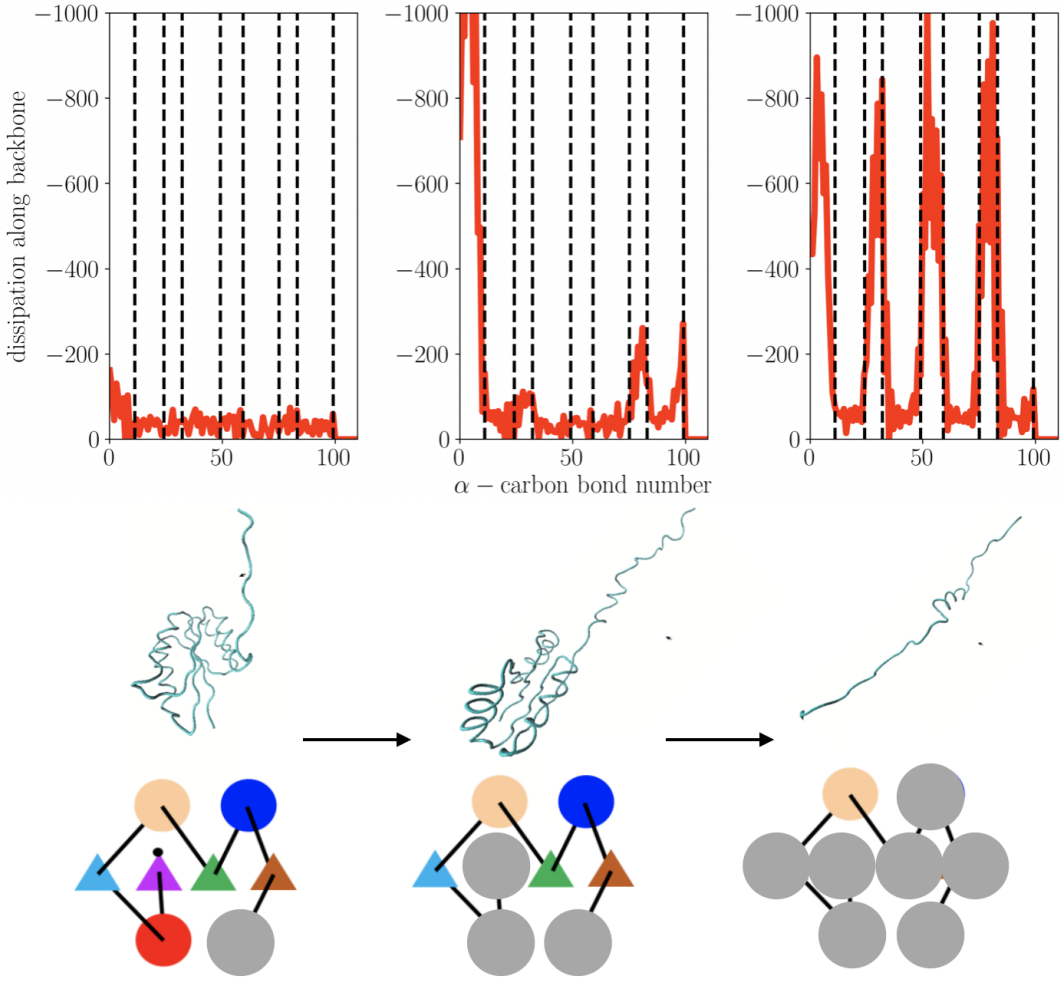

Figure 5.6: Top: dissipation profiles along the alpha-carbon backbone bonds as the protein is pulled open quickly (left to right). Borders between secondary structures shown as black dashed lines. Middle: representative snapshots of the protein at simulation times corresponding to Top. Bottom: secondary structures are shaded in grey as they are unraveled.

(alpha helix) secondary structures are pried from the center of the protein. Finally, when the protein is fully unfolded (right), the alpha-carbon dissipation profile exhibits a standing wave pattern, where the peaks and troughs are marked by alternating secondary structures. The alpha helices experience an increased dissipation as their bonds undergo higher restoring forces.

We have shown that the dissipation transduced along the backbone of a protein as it is pulled open can be leveraged as a coarse grained single-dimensional reporter of how each secondary structure unfolds in real time. We expect this easy to compute observable will

find future use in studying the mechanical unfolding process of more complex proteins.

5.6 Conclusion

We have studied single molecule trajectories from optical tweezing experiments and coarse-grained molecular simulations to infer information about intermediate states on the path a small protein takes as it folds and unfolds. We developed an experimental data analysis pipeline and explored new computational coordinates that could be used to investigate the temporal order of protein folding in simulations.

Chapter 6

Minimum variance free energy estimation

The work in this chapter was done in collaboration with Adrianne Zhong.

Given two potential energy landscapes $V_A(\mathbf{x})$ and $V_B(\mathbf{x})$, equilibrium probability distributions are defined as $\rho_A(\mathbf{x}) = Z_A^{-1}\exp(-\beta\, V_A(\mathbf{x}))$ and $\rho_B(\mathbf{x}) = Z_B^{-1}\exp(-\beta\, V_B(\mathbf{x}))$ respectively, with inverse temperature $\beta = (k_{\mathrm{B}}T)^{-1}$ and partition functions $Z_A = \int d\mathbf{x}\exp(-\beta\, V_A(\mathbf{x}))$ and $Z_B = \int d\mathbf{x}\exp(-\beta\, V_B(\mathbf{x}))$. The free energy difference $\Delta F = \beta^{-1}\ln(Z_A/Z_B)$ encodes, through detailed balance, how forward and backward transition rates are related in equilibrium.

6.1 Nonequilibrium free energy estimation

There are many nonequilibrium estimators for equilibrium free energy differences. One of the simplest uses the microscopic fluctuation theorem[6]

$$P_F(W_F) = P_R(-W_R)e^{-\beta(W-\Delta F)} \tag{6.1}$$

to find the work value at which the forward and backward work distributions overlap, but poor overlap between distributions in the fast switching regime can prohibit its use. The maximum likelihood estimator for the free energy difference $\Delta F = \beta^{-1}\ln(Z_A/Z_B)$ uses the Bennett acceptance ratio (BAR) and is given implicitly as the value ΔF that satisfies [150, 151, 152]

$$h(\Delta F) = \left\langle \frac{1}{1+e^{\beta(W-\Delta F)}} \right\rangle_F - \left\langle \frac{1}{1+e^{\beta(W+\Delta F)}} \right\rangle_R = 0. \tag{6.2}$$

where W is the work done over a stochastic trajectory transforming between $V_A(\mathbf{x})$ and $V_B(\mathbf{x})$ in finite time. Here, the brackets denote ensemble averages from forward (F) and

reverse (R) measurements respectively, and we have assumed for simplicity an equal number of forward and reverse samples $n_F = n_R = n_{\text{tot}}/2$ are made, though it is straightforward to extend to $n_F \neq n_R$. In the forward process, trajectories are initialized with the equilibrium distribution $\rho_A(\mathbf{x})$ and the time-varying potential $V^F(\mathbf{x}, t)$ begins at $V_A(\mathbf{x})$ at $t = 0$ and ends at $V_B(\mathbf{x})$ at finite time $t = t_f$; while in the reverse they are initialized with $\rho_B(\mathbf{x})$ and the time-varying potential is the time-reversal of the forward process $V^R(\mathbf{x}, t) = V^F(\mathbf{x}, t_f - t)$.

Given finite samples, the average information per sample $\iota(\Delta F) = I(\Delta F)/n_{\text{tot}}$ is given by

$$\iota(\Delta F) = \frac{\beta^2}{4}\left[\left\langle \frac{1}{1+\cosh(\beta(W-\Delta F))}\right\rangle_F + \left\langle \frac{1}{1+\cosh(\beta(W+\Delta F))}\right\rangle_R\right], \tag{6.3}$$

and related to the estimator variance by $\sigma^2 = n_{\text{tot}}^{-1}(\iota^{-1} - 4\beta^{-2})$ – protocols that give yield ι correspond to low estimator variance σ^2 and are more efficient, i.e. on average requiring fewer samples to obtain a desired level of estimator error. [151]

In general the variance of the BAR estimator depends on the particular time-dependent potential energy $V^F(\mathbf{x}, t)$ that transforms from $V_A(\mathbf{x})$ and $V_B(\mathbf{x})$, which we denote as BAR estimator protocol, with different protocols yielding different values of the information-per-measurement in Eq. (6.3). In this paper, we study optimal protocols, i.e. those that maximize ι and thus minimize σ^2.

To make finite-time work measurements, we consider overdamped Langevin dynamics, though it is possible to extend our work to the underdamped regime. We generate stochastic trajectories $\mathbf{x}(t)$ for time $t \in [0, t_f]$ by integrating the equations of motion

$$\begin{aligned} \gamma\dot{\mathbf{x}} &= -\nabla V_\theta(\mathbf{x}, t) + \sqrt{2\gamma k_{\text{B}}T}\boldsymbol{\eta}(t) \\ \mathbf{x}(0) &\sim \rho_0(\cdot) \end{aligned} \tag{6.4}$$

where γ is the friction coefficient, ∇ denotes the spatial gradient, $V_\theta(\mathbf{x}, t)$ is a time-varying potential energy parameterized by some set of scalar parameters $\theta \in \mathbb{R}^M$, and each component of $\boldsymbol{\eta}(t)$ is an independent instance of delta-correlated white noise with unit variance; and $\rho_0(\cdot)$ is the probability distribution from which $\mathbf{x}(0)$ is drawn.

For an individual trajectory $\mathbf{x}(t)$ generated from the above dynamics, the work up to time t_f is given by

$$W_\theta[\mathbf{x}(t)] := \int_0^{t_f} \left(\frac{\partial U_\theta}{\partial t}\right) dt \tag{6.5}$$

integrated along the trajectory $\mathbf{x}(t)$.

6.2 Protocol optimization

Forward protocols have trajectories initialized from the equilibrium distribution $\rho_0^F(\mathbf{x}) = \rho_A(\mathbf{x})$ and the dynamics are under the forward protocol $V_\theta^F(\mathbf{x}, t)$ satisfy $V_\theta^F(\mathbf{x}, 0) = V_A(\mathbf{x})$

and $V_\theta^F(\mathbf{x}, t_f) = V_B(\mathbf{x})$. We consider affine control protocols

$$V_\theta^F(\mathbf{x}, t) = \begin{cases} V_A(\mathbf{x}) & \text{for } t = 0 \\ \sum_n \lambda_\theta^n(t)\, V_n(\mathbf{x}) & \text{for } t \in (0, t_f) \\ V_B(\mathbf{x}) & \text{for } t = t_f, \end{cases} \tag{6.6}$$

where $\{V_n(\mathbf{x}), \}$ is a set of control potentials, each with a control $\lambda_\theta^n(t)$ that is parameterized as a linear combination of orthogonal Legendre polynomials

$$\lambda_\theta^n(t) = \sum_{m=0}^{m_{\max}} \theta^{nm} P_m\left(\frac{2t}{t_f} - 1\right), \tag{6.7}$$

with argument scaled to be in the range $[-1, 1]$, though in theory any other temporal basis functions may be used. To allow for jump discontinuities at the start and end of the protocols commonly observed in optimal protocols, we do not impose the constraints $\lim_{t\to 0^+} V_\theta^F(\mathbf{x}, t) = V_A(\mathbf{x})$ and $\lim_{t\to t_f^-} V_\theta^F(\mathbf{x}, t) = V_B(\mathbf{x})$.

Reverse protocols on the other hand have trajectories initialized from $\rho_0^R(\mathbf{x}) = \rho_B(\mathbf{x})$, given by the time-reversal of the forward protocol $V_\theta^R(\mathbf{x}, t) = V_\theta^F(\mathbf{x}, t_f - t) = \sum_n \lambda_\theta^n(t)\, V_n(\mathbf{x})$.

There are infinitely many choices of the set of control potentials $\{U_n(\mathbf{x}), \}$, but for simplicity we consider the two-dimensional case $V_\theta^F(\mathbf{x}, t) = \lambda_\theta^A(t)\, V_A(\mathbf{x}) + \lambda_\theta^B(t)\, V_B(\mathbf{x})$. This is a generalization of the linear interpolation $V_\lambda(\mathbf{x}) = (1 - \lambda)\, V_A(\mathbf{x}) + \lambda\, V_B(\mathbf{x})$ with single-dimension control $\lambda(t)$ commonly considered. Thus, altogether our set of parameters is

$$\theta = \{\theta^{nm}, \mid n \in \{A, B\}, m \in \{0, 1, ..., m_{\max}\}\} \tag{6.8}$$

containing $2(m_{\max} + 1)$ parameters in total.

We note that both the constraint equation that defines the BAR estimator Eq. (6.2) and our objective function Eq. (6.3) are given in terms of expectations of path observables of the form $O(W, \Delta F)$, formally written as

$$\langle O \rangle = \int D[\mathbf{X}(t)]\, P_\theta[\mathbf{X}(t)]\, O(W_\theta[\mathbf{x}(t)], \Delta F). \tag{6.9}$$

where $\int D[\mathbf{X}(t)]$ a path integral, and $P_\theta[\mathbf{X}(t)]$ is the path probability. The work functional for a path Eq. (6.5) is dependent on the protocol explicitly through θ, and the probability of observing a path $\mathbf{x}(t)$ from Eq. (6.4) is

$$P_\theta[\mathbf{x}(t)] \propto p_0(\mathbf{x}(0)) \exp(U_\theta[\mathbf{x}(t)]). \tag{6.10}$$

Here, $U_\theta[\mathbf{x}(t)]$ is the Onsager-Machlup action

$$U_\theta[\mathbf{x}(t)] = -\int_0^{t_f} \frac{|\gamma \dot{\mathbf{x}}(t') + \nabla V_\theta(\mathbf{x}(t'), t')|^2}{4\gamma k_\mathrm{B} T} dt'. \tag{6.11}$$

Algorithm 2 Forward-Reverse time Malliavin weights for Free energy estimation

inputs Legendre polynomial coefficients for a general force $\boldsymbol{\lambda}_\theta(\mathbf{x}, t)$
parameters learning rate l; total optimization steps I; trajectory length t_f consisting of J timesteps of duration Δt each; number of trajectories N
initialize choose initial weights θ, define iteration variables i and j, loss-function gradients $\nabla\langle O(W, \Delta F)\rangle_\theta$
$i \leftarrow 0$
repeat
    Generate forward $[\mathbf{X}(t)]$ and reverse time $[\tilde{\mathbf{X}}(t)]$ trajectories with first-order Euler propagation starting from uncorrelated steady-state configurations in state A and compute a work values for each trajectory. Configurations, times, noises (with variance $2\boldsymbol{\gamma} k_{\mathrm{B}} T \Delta t$) and Malliavin weights are denoted $\mathbf{x}_j, t_j, \boldsymbol{\eta}_j, y_\theta^F(t_j)$ and $y_\theta^R(t_j)$, respectively. Time reversal of configurations, times and noises is denoted by a tilde.
    $j \leftarrow 0$
    $\nabla\langle O\rangle_\theta \leftarrow 0$
    $y_\theta^F(t_0) \leftarrow 0$
    $y_\theta^R(t_J) \leftarrow 0$ For each trajectory
    repeat
        $y_\theta^R(\tilde{t}_{j+1}) \leftarrow y_\theta^R(\tilde{t}_j) + \tilde{\boldsymbol{\eta}}_j \cdot \nabla_\theta \boldsymbol{\lambda}_\theta(\tilde{\mathbf{x}}_j, \tilde{t}_j)/2k_{\mathrm{B}}T$
        $y_\theta^F(t_{j+1}) \leftarrow y_\theta^F(t_j) + \boldsymbol{\eta}_j \cdot \nabla_\theta \boldsymbol{\lambda}_\theta(\mathbf{x}_j, t_j)/2k_{\mathrm{B}}T$
        $j \leftarrow j + 1$
    until $j = J$
    $\Delta F \leftarrow \mathrm{argmin}_{\Delta F}\langle(1 + e^{\beta(W-\Delta F)})^{-1}\rangle_F - \langle(1 + e^{\beta(W+\Delta F)})^{-1}\rangle_R$
    average y_θ^F and y_θ^R over N trajectories to get $\bar{y}_\theta^F$ and $\bar{y}_\theta^R$
    $\nabla\langle O\rangle_\theta \leftarrow \langle\nabla O\rangle_\theta + \langle O(y_\theta^F - \bar{y}_\theta^F + y_\theta^R - \bar{y}_\theta^R)\rangle_\theta$
    note the above averages will include sums over forward and reverse trajectories
    $\theta \leftarrow \theta - l\nabla\langle O\rangle_\theta$
    $i \leftarrow i + 1$
until $i = I$

The partial derivative of expectations in Eq. (6.9) with respect to a parameter θ^μ is

$$\frac{\partial\langle O\rangle}{\partial\theta^\mu} = \left\langle \frac{\partial O}{\partial W}\frac{\partial W_\theta}{\partial\theta^\mu} + O\,\frac{\partial \ln \mathcal{P}_\theta}{\partial\theta^\mu}\right\rangle, \tag{6.12}$$

6.3 Forward-Reverse time Algorithm

The path probability for observable expectations (6.9) are protocol-dependent, i.e. θ-dependent, the BAR estimator Eq. (6.2) and information-per-sample Eq. (6.3) are functions of both the

estimated free energy and the protocol parameters $h(\Delta F, \theta)$ and $\iota(\Delta F, \theta)$.

For each stochastic trajectory $\mathbf{x}(t)$ generated through Eq. (6.4) with the forward/reverse protocol, we keep track of additional dynamical variables, Malliavin weights, $q_\mu^{F/R}(t)$ with $\mu \in \{(n, m)\}$.

Keeping track of $q_\mu(t)$ incurs a trivial amount of additional cost in memory and computation time: the most expensive part of simulating Langevin dynamics Eq. (6.4) for a multi-dimensional system is the evaluation of the energy and gradient forces, $V_n(\mathbf{x})$ and $-\nabla V_n(\mathbf{x})$ for each control potential n. In total, $2M$ additional variables are needed. As illustrated in Algorithm 2, we first calculate Malliavin weights for all trajectories, forward and reverse. Next we calculate the BAR estimate of the free energy with work values for all trajectories and use this to calculate the observable $O(W, \Delta F)$ and its trajectory gradients $\nabla\langle O\rangle_\theta$, and propogate the parameters forward iteration by iteration.

6.4 Biased symmetric double well

We consider a protocol that switches the bias on a symmetric double well with barrier height ΔV from one side to another, as shown in Fig. 6.1 a. The ground truth free energy difference in this system is $\Delta F = 0$ by the underlying symmetry of the system. The equation of motion is

$$\gamma\dot{x} = -\partial_x V(x,t) + \sqrt{2\gamma k_\mathrm{B} T}\eta \tag{6.13}$$

with white noise $\langle\eta\rangle = 0$, $\langle\eta(t)\eta(t')\rangle = \delta(t - t')$ and potential

$$V(x,t) = \Delta V/4(x - x_A)^2(x - x_B)^2 - \Delta V\lambda(t)x,\ \lambda_A(0) = -\lambda_B(t_f) = -1, \tag{6.14}$$

where we fix $\beta\Delta V = 16$ and $x_B = -x_A = \beta = \gamma = D = 1$. We optimize the control force $\lambda(t)$ over 100 iterations of gradient descent each consisting of 2000 forward trajectories and 2000 reverse time trajectories, using a basis set of 15 Legendre polynomials to represent $\lambda(t)$. Taking time in units of $\tau = x_A^2/D = 1$ we study a switching time $t_f = 5\tau$ with time step $\Delta t = \tau/100$. We use an Euler-Maruyama discretization of the equations of motion. Initial conditions for forward and backward trajectories were sampled from long equilibrium simulations under potentials $V_A = \Delta V/4(x - x_A)^2(x - x_B)^2 - \Delta V x$ and $V_B = \Delta V/4(x - x_A)^2(x - x_B)^2 + \Delta V x$. At each iteration we compute the BAR estimator free energy difference ΔF by minimizing $h(\Delta F)$ over our current and previous work samples with the SciPy minimize function[108], and the BAR variance Eq. 6.3 is calculated with work samples and ΔF. We use a learning rate of 0.1 to propagate the control parameter gradients from iteration to iteration. We initialize our coefficients to a linear ramp $2 * t/t_f - 1$ that interpolates between $\lambda_A(0) = -1$ and $\lambda_B(t_f) = 1$. The details of how we update the Malliavin weights with forward and backward trajectory data is provided in Alg. 2.

The optimized control force $\lambda(t)$ is shown in Fig. 6.1; it is approximately symmetric in time about the midpoint and qualitatively agrees with the minimum work protocol of Sivak and Crooks, with a plateau bridging two linearly increasing portions.[153, 154] On iteration

0, before the optimization, the distributions of forward and reverse time work measurements in Fig. 6.1 c have little overlap at the value of the free energy $\Delta F = 0$. After 100 iterations of optimization however, after the BAR variance has been sufficiently minimized, the work distributions in Fig. 6.1 d have been transformed to exhibit significant overlap at $W = 0$. In Fig. 6.1 e we show the BAR variance per sample, both cumulatively using all previous trajectories in red and per iteration using only trajectories generated with the current control force. We use the same color scheme in Fig. 6.1 to illustrate how the free energy estimate becomes better as training progresses, both cumulatively and, although more noisily, per iteration.

6.5 Conclusion

Whether by Malliavin weights, neural networks or some other stochastic optimization protocol, we expect BAR variance minimization, and work minimization, to be of use in estimating free energy differences in noisy physical systems, such as the free energy differences between the bound and unbound states of enzymes and substrates, and drug targets and small molecules, as well as the folded and unfolded states of proteins. When the control parameter is an order parameter, such as is the case in the mechanical extension of proteins, free energy profiles (potentials of mean force) can be constructed with our method.

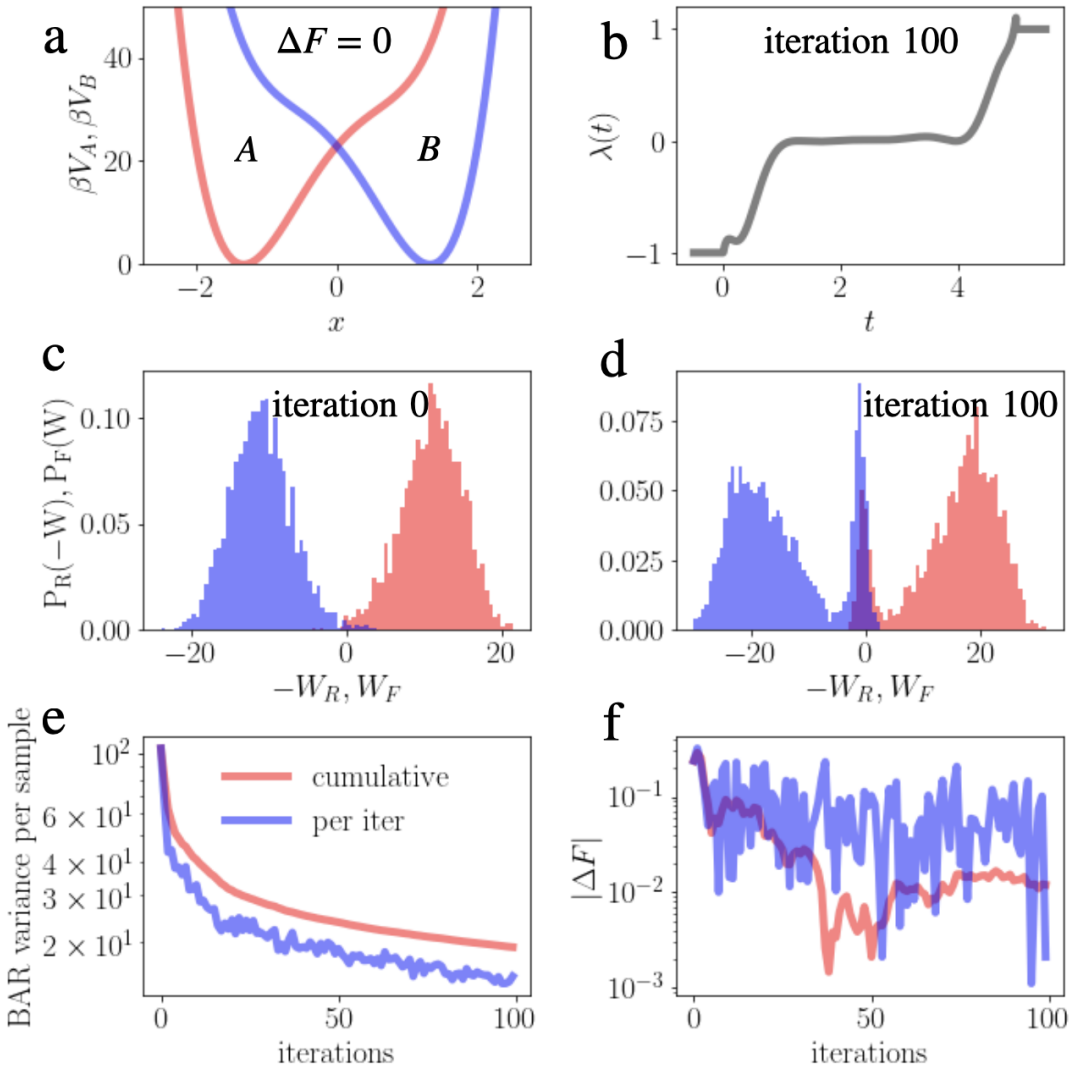

Figure 6.1: (a) A biased double well A, red, is transformed to B, blue, using a finite time control protocol with net free energy difference $\Delta F = 0$. (b) The optimized control protocol linearly increases until the double well is symmetric, pauses, and linearly increases again. (c, d) Work distributions before and after optimization. Before (c) there is little overlap at the free energy value; after optimization (d) the overlap is significant. (e) Cumulative and per iteration BAR variance per sample across iterations. (f) Cumulative and per iteration free energy estimate across iterations.

Bibliography

[1] Peter G Bolhuis et al. "Transition path sampling: Throwing ropes over rough mountain passes, in the dark". In: *Annual review of physical chemistry* 53.1 (2002), pp. 291–318.

[2] David Chandler. "Statistical mechanics of isomerization dynamics in liquids and the transition state approximation". In: *The Journal of Chemical Physics* 68.6 (1978), pp. 2959–2970.

[3] Udo Seifert. "Stochastic thermodynamics, fluctuation theorems and molecular machines". In: *Reports on progress in physics* 75.12 (2012), p. 126001.

[4] L. Onsager and S. Machlup. "Fluctuations and Irreversible Processes." In: *Phys. Rev.* 91 (1953), p. 1505.

[5] Chloe Ya Gao and David T Limmer. "Nonlinear transport coefficients from large deviation functions". In: *The Journal of chemical physics* 151.1 (2019), p. 014101.

[6] Gavin E Crooks. "Path-ensemble averages in systems driven far from equilibrium". In: *Physical review E* 61.3 (2000), p. 2361.

[7] Christian Maes. "Frenesy: Time-symmetric dynamical activity in nonequilibria". In: *Physics Reports* 850 (2020), pp. 1–33.

[8] Christopher Jarzynski. "Nonequilibrium equality for free energy differences". In: *Physical Review Letters* 78.14 (1997), p. 2690.

[9] A. Das and D. T. Limmer. "Variational control forces for enhanced sampling of nonequilibrium molecular dynamics simulations." In: *Journal of Chemical Physics* 151 (2019), p. 244123. DOI: https://doi.org/10.1063/1.5128956.

[10] Patrick B Warren and Rosalind J Allen. "Malliavin weight sampling for computing sensitivity coefficients in brownian dynamics simulations". In: *Physical review letters* 109.25 (2012), p. 250601.

[11] Patrick B Warren and Rosalind J Allen. "Malliavin weight sampling: a practical guide". In: *Entropy* 16.1 (2013), pp. 221–232.

[12] Avishek Das et al. "Reinforcement learning of rare diffusive dynamics". In: *The Journal of Chemical Physics* 155.13 (2021), p. 134105.

[13] Benjamin Kuznets-Speck and David T. Limmer. In: *Proceedings of the National Academy of Sciences* 118.8 (2021).

[14] Carlos Bustamante et al. "Mechanical processes in biochemistry". In: *Annual review of biochemistry* 73.1 (2004), pp. 705–748.

[15] Mauro L Mugnai et al. "Theoretical perspectives on biological machines". In: *Reviews of Modern Physics* 92.2 (2020), p. 025001.

[16] Alireza Mashaghi et al. "Chaperone action at the single-molecule level". In: *Chemical reviews* 114.1 (2014), pp. 660–676.

[17] Yu Ling Wu et al. "Melting and crystallization of colloidal hard-sphere suspensions under shear". In: *Proceedings of the National Academy of Sciences* 106.26 (2009), pp. 10564–10569.

[18] H.A. Kramers. "Brownian motion in a field of force and the diffusion model of chemical reactions". In: *Physica* 7 (1940), p. 284.

[19] Freddy Bouchet and Julien Reygner. "Generalisation of the Eyring–Kramers transition rate formula to irreversible diffusion processes". In: *Annales Henri Poincaré*. Vol. 17. 12. Springer. 2016, pp. 3499–3532.

[20] Jordan M Horowitz and Todd R Gingrich. "Thermodynamic uncertainty relations constrain non-equilibrium fluctuations". In: *Nature Physics* 16.1 (2020), pp. 15–20.

[21] Michael E Cates and Julien Tailleur. "When are active Brownian particles and run-and-tumble particles equivalent? Consequences for motility-induced phase separation". In: *Europhysics Letters* 101.2 (2013), p. 20010.

[22] Takahiro Nemoto et al. "Optimizing active work: Dynamical phase transitions, collective motion, and jamming". In: *Physical Review E* 99.2 (2019), p. 022605.

[23] Trevor GrandPre and David T Limmer. "Current fluctuations of interacting active Brownian particles". In: *Physical Review E* 98.6 (2018), p. 060601.

[24] Eric Woillez et al. "Activated escape of a self-propelled particle from a metastable state". In: *Physical review letters* 122.25 (2019), p. 258001.

[25] Eric Woillez, Yariv Kafri, and Vivien Lecomte. "Nonlocal stationary probability distributions and escape rates for an active Ornstein–Uhlenbeck particle". In: *Journal of Statistical Mechanics: Theory and Experiment* 2020.6 (2020), p. 063204.

[26] Abhinav Sharma, René Wittmann, and Joseph M Brader. "Escape rate of active particles in the effective equilibrium approach". In: *Physical Review E* 95.1 (2017), p. 012115.

[27] Luca Gammaitoni et al. "Stochastic resonance". In: *Reviews of modern physics* 70.1 (1998), p. 223.

[28] Peter Jung. "Periodically driven stochastic systems". In: *Physics Reports* 234.4-5 (1993), pp. 175–295.

[29] Alex Albaugh and Todd R Gingrich. "Estimating reciprocal partition functions to enable design space sampling". In: *The Journal of Chemical Physics* 153.20 (2020), p. 204102.

[30] Gianmaria Falasco and Massimiliano Esposito. "Dissipation-time uncertainty relation". In: *Physical Review Letters* 125.12 (2020), p. 120604.

[31] Denis J Evans and Debra J Searles. "Equilibrium microstates which generate second law violating steady states". In: *Physical Review E* 50.2 (1994), p. 1645.

[32] Giovanni Gallavotti and Ezechiel Godert David Cohen. "Dynamical ensembles in nonequilibrium statistical mechanics". In: *Physical review letters* 74.14 (1995), p. 2694.

[33] Andreas Dechant and Shin-ichi Sasa. "Fluctuation–response inequality out of equilibrium". In: *Proceedings of the National Academy of Sciences* 117.12 (2020), pp. 6430–6436.

[34] Gavin E Crooks. "Entropy production fluctuation theorem and the nonequilibrium work relation for free energy differences". In: *Physical Review E* 60.3 (1999), p. 2721.

[35] R. Chetrite and H. Touchette. "Nonequilibrium Markov Processes Conditioned on Large Deviations." In: *Ann. Inst. Henri Poincaré* 16 (2015), p. 2005.

[36] Todd R Gingrich et al. "Near-optimal protocols in complex nonequilibrium transformations". In: *Proceedings of the National Academy of Sciences* 113.37 (2016), pp. 10263–10268.

[37] Shamreen Iram et al. "Controlling the speed and trajectory of evolution with counterdiabatic driving". In: *Nat. Phys.* 17.1 (2021), pp. 135–142. DOI: https://doi.org/10.1038/s41567-020-0989-3.

[38] S. W. Schneider et al. "Shear-induced unfolding triggers adhesion of von Willebrand factor fibers". In: *Proc. Natl. Acad. Sci.* 104.19 (2007), p. 7899. DOI: 10.1073/pnas.0608422104.

[39] M Delarue et al. "Self-driven jamming in growing microbial populations". In: *Nat. Phys.* 12 (2016), p. 762.

[40] W B Rogers, W M Shih, and V N Manoharan. "Using DNA to program the self-assembly of colloidal nanoparticles and microparticles". In: *Nat. Rev. Mater.* 1 (2016), p. 16008.

[41] A Aubret et al. "Targeted assembly and synchronization of self-spinning microgears". In: *Nat. Phys.* 14 (2018), p. 1114.

[42] J Denk et al. "MinE conformational switching confers robustness on self-organized Min protein patterns". In: *Proc. Natl. Acad. Sci.* 115 (2018), p. 4553.

[43] J Liu et al. "Cascading MutS and MutL sliding clamps control DNA diffusion to activate mismatch repair". In: *Nature* 539 (2016), p. 583.

[44] Avishek Das, Benjamin Kuznets-Speck, and David T Limmer. "Direct evaluation of rare events in active matter from variational path sampling". In: *Physical Review Letters* 128.2 (2022), p. 028005.

[45] M Cristina Marchetti et al. "Hydrodynamics of soft active matter". In: *Reviews of Modern Physics* 85.3 (2013), p. 1143.

[46] M Reza Shaebani et al. "Computational models for active matter". In: *Nature Reviews Physics* 2.4 (2020), pp. 181–199.

[47] Michael E Cates and Julien Tailleur. "Motility-induced phase separation". In: *Annual Reviews of Condensed Matter Physics* 6.1 (2015), pp. 219–244.

[48] Thomas Speck. "Collective behavior of active Brownian particles: From microscopic clustering to macroscopic phase separation". In: *The European Physical Journal Special Topics* 225.11 (2016), pp. 2287–2299.

[49] Trevor GrandPre et al. "Entropy production fluctuations encode collective behavior in active matter". In: *Physical Review E* 103.1 (2021), p. 012613.

[50] Gerhard Gompper et al. "The 2020 motile active matter roadmap". In: *Journal of Physics: Condensed Matter* 32.19 (2020), p. 193001.

[51] Andrei Militaru et al. "Escape dynamics of active particles in multistable potentials". In: *Nature Communications* 12.1 (2021), pp. 1–6.

[52] Ahmad K Omar et al. "Phase Diagram of Active Brownian Spheres: Crystallization and the Metastability of Motility-Induced Phase Separation". In: *Physical Review Letters* 126.18 (2021), p. 188002.

[53] David Richard, Hartmut Löwen, and Thomas Speck. "Nucleation pathway and kinetics of phase-separating active Brownian particles". In: *Soft Matter* 12.24 (2016), pp. 5257–5264.

[54] Joakim Stenhammar et al. "Activity-induced phase separation and self-assembly in mixtures of active and passive particles". In: *Physical review letters* 114.1 (2015), p. 018301.

[55] Stewart A Mallory, Chantal Valeriani, and Angelo Cacciuto. "An active approach to colloidal self-assembly". In: *Annual review of physical chemistry* 69 (2018), pp. 59–79.

[56] Abraham Nitzan. *Chemical dynamics in condensed phases: relaxation, transfer and reactions in condensed molecular systems*. Oxford university press, 2006.

[57] Rosalind J Allen, Chantal Valeriani, and Pieter Rein Ten Wolde. "Forward flux sampling for rare event simulations". In: *Journal of physics: Condensed matter* 21.46 (2009), p. 463102.

[58] Aryeh Warmflash, Prabhakar Bhimalapuram, and Aaron R Dinner. "Umbrella sampling for nonequilibrium processes". In: *The Journal of chemical physics* 127.15 (2007), p. 114109.

[59] Frédéric Cérou and Arnaud Guyader. "Adaptive multilevel splitting for rare event analysis". In: *Stochastic Analysis and Applications* 25.2 (2007), pp. 417–443.

[60] Yaouen Fily and M Cristina Marchetti. "Athermal phase separation of self-propelled particles with no alignment". In: *Physical review letters* 108.23 (2012), p. 235702.

[61] Gabriel S Redner, Michael F Hagan, and Aparna Baskaran. "Structure and dynamics of a phase-separating active colloidal fluid". In: *Physical review letters* 110.5 (2013), p. 055701.

[62] Julian Bialké, Hartmut Löwen, and Thomas Speck. "Microscopic theory for the phase separation of self-propelled repulsive disks". In: *Europhysics Letters* 103.3 (2013), p. 30008.

[63] SA Mallory, ML Bowers, and A Cacciuto. "Universal reshaping of arrested colloidal gels via active doping". In: *The Journal of Chemical Physics* 153.8 (2020), p. 084901.

[64] Andrey Sokolov et al. "Swimming bacteria power microscopic gears". In: *Proceedings of the National Academy of Sciences* 107.3 (2010), pp. 969–974.

[65] J Szavits-Nossan and M R Evans. "Inequivalence of nonequilibrium path ensembles: the example of stochastic bridges." In: *Journal of Statistical Mechanics* 2015 (2015), P12008. DOI: https://doi.org/10.1088/1742-5468/2015/12/P12008.

[66] R. Chetrite and H. Touchette. "Nonequilibrium Markov Processes Conditioned on Large Deviations". In: *Annales Henri Poincaré* 16 (2015), p. 2005. DOI: https://doi.org/10.1007/s00023-014-0375-8.

[67] R. Chetrite and H. Touchette. "Variational and optimal control representations of conditioned and driven processes." In: *Journal of Statistical Mechanics* 2015 (2015), P12001. DOI: https://doi.org/10.1088/1742-5468/2015/12/P12001.

[68] Eric Vanden-Eijnden et al. "Transition-path theory and path-finding algorithms for the study of rare events." In: *Annual review of physical chemistry* 61 (2010), pp. 391–420.

[69] Ushnish Ray, Garnet Kin-Lic Chan, and David T Limmer. "Exact fluctuations of nonequilibrium steady states from approximate auxiliary dynamics". In: *Physical review letters* 120.21 (2018), p. 210602.

[70] Takahiro Nemoto et al. "Population-dynamics method with a multicanonical feedback control". In: *Physical Review E* 93.6 (2016), p. 062123.

[71] Daniel Jacobson and Stephen Whitelam. "Direct evaluation of dynamical large-deviation rate functions using a variational ansatz". In: *Physical Review E* 100.5 (2019), p. 052139.

[72] John D Weeks, David Chandler, and Hans C Andersen. "Role of repulsive forces in determining the equilibrium structure of simple liquids". In: *The Journal of chemical physics* 54.12 (1971), pp. 5237–5247.

[73] Christoph Dellago, Peter G Bolhuis, and David Chandler. "On the calculation of reaction rate constants in the transition path ensemble". In: *The Journal of chemical physics* 110.14 (1999), pp. 6617–6625.

[74] A. Das et al. "Reinforcement learning of rare diffusive dynamics". In: *arXiv:2105.04321* (2021).

[75] Dominic C Rose, Jamie F Mair, and Juan P Garrahan. "A reinforcement learning approach to rare trajectory sampling". In: *New Journal of Physics* 23.1 (2021), p. 013013.

[76] Alessandro Barducci, Giovanni Bussi, and Michele Parrinello. "Well-tempered metadynamics a smoothly converging and tunable free-energy method". In: *Physical Review Letters* 100.2 (2008), p. 020603.

[77] Marc Delarue, Patrice Koehl, and Henri Orland. "Ab initio sampling of transition paths by conditioned Langevin dynamics". In: *The Journal of chemical physics* 147.15 (2017), p. 152703.

[78] Ushnish Ray, Garnet Kin-Lic Chan, and David T Limmer. "Importance sampling large deviations in nonequilibrium steady states. I". In: *Journal of chemical physics* 148.12 (2018), p. 124120.

[79] P Buijsman and PG Bolhuis. "Transition path sampling for non-equilibrium dynamics without predefined reaction coordinates". In: *The Journal of chemical physics* 152.4 (2020), p. 044108.

[80] Daan Frenkel and Berend Smit. *Understanding molecular simulation: from algorithms to applications*. Vol. 1. Elsevier, 2001.

[81] Michael R Shirts and John D Chodera. "Statistically optimal analysis of samples from multiple equilibrium states". In: *Journal of chemical physics* 129.12 (2008), p. 124105.

[82] Satya N Majumdar and Henri Orland. "Effective Langevin equations for constrained stochastic processes". In: *Journal of Statistical Mechanics: Theory and Experiment* 2015.6 (2015), P06039.

[83] Federico Carollo et al. "Making rare events typical in Markovian open quantum systems". In: *Physical Review A* 98.1 (2018), p. 010103.

[84] Patrick Warren and Rosalind Allen. "Malliavin Weight Sampling: A Practical Guide". In: *Entropy* 16.1 (Dec. 2013), pp. 221–232. ISSN: 1099-4300. DOI: 10.3390/e16010221. URL: http://dx.doi.org/10.3390/e16010221.

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

[86] Rosalind J Allen, Patrick B Warren, and Pieter Rein Ten Wolde. "Sampling rare switching events in biochemical networks". In: *Physical review letters* 94.1 (2005), p. 018104.

[87] Rosalind J Allen, Daan Frenkel, and Pieter Rein ten Wolde. "Simulating rare events in equilibrium or nonequilibrium stochastic systems". In: *The Journal of chemical physics* 124.2 (2006), p. 024102.

[88] Jakub Dolezal and Robert L Jack. "Large deviations and optimal control forces for hard particles in one dimension". In: *Journal of Statistical Mechanics: Theory and Experiment* 2019.12 (2019), p. 123208.

[89] Benjamin Kuznets-Speck and David T Limmer. "Inferring equilibrium transition rates from nonequilibrium protocols". In: *arXiv preprint arXiv:2208.05585* (2022).

[90] Gerhard Hummer and Attila Szabo. "Free energy reconstruction from nonequilibrium single-molecule pulling experiments". In: *Proceedings of the National Academy of Sciences* 98.7 (2001), pp. 3658–3661.

[91] Gerhard Hummer and Attila Szabo. "Free energy surfaces from single-molecule force spectroscopy". In: *Accounts of chemical research* 38.7 (2005), pp. 504–513.

[92] Megan L Hughes and Lorna Dougan. "The physics of pulling polyproteins: a review of single molecule force spectroscopy using the AFM to study protein unfolding". In: *Reports on Progress in Physics* 79.7 (2016), p. 076601.

[93] George I Bell. "Models for the specific adhesion of cells to cells: a theoretical framework for adhesion mediated by reversible bonds between cell surface molecules." In: *Science* 200.4342 (1978), pp. 618–627.

[94] Gerhard Hummer and Attila Szabo. "Kinetics from nonequilibrium single-molecule pulling experiments". In: *Biophysical journal* 85.1 (2003), pp. 5–15.

[95] Olga K Dudko, Gerhard Hummer, and Attila Szabo. "Theory, analysis, and interpretation of single-molecule force spectroscopy experiments". In: *Proceedings of the National Academy of Sciences* 105.41 (2008), pp. 15755–15760.

[96] Evan Evans and Ken Ritchie. "Dynamic strength of molecular adhesion bonds". In: *Biophysical journal* 72.4 (1997), pp. 1541–1555.

[97] Olga K Dudko, Gerhard Hummer, and Attila Szabo. "Intrinsic rates and activation free energies from single-molecule pulling experiments". In: *Physical review letters* 96.10 (2006), p. 108101.

[98] Baron Peters. *Reaction rate theory and rare events*. Elsevier, 2017.

[99] Suman Kumar Banik, Jyotipratim Ray Chaudhuri, and Deb Shankar Ray. "The generalized Kramers theory for nonequilibrium open one-dimensional systems". In: *The Journal of Chemical Physics* 112.19 (2000), pp. 8330–8337.

[100] Tomoji Yamada and Kyozi Kawasaki. "Nonlinear effects in the shear viscosity of critical mixtures". In: *Progress of Theoretical Physics* 38.5 (1967), pp. 1031–1051.

[101] Gary P Morriss and Denis J Evans. "Isothermal response theory". In: *Molecular Physics* 54.3 (1985), pp. 629–636.

[102] Dominika Lesnicki et al. "On the molecular correlations that result in field-dependent conductivities in electrolyte solutions". In: *The Journal of chemical physics* 155.1 (2021), p. 014507.

[103] David A Sivak and Gavin E Crooks. "Near-equilibrium measurements of nonequilibrium free energy". In: *Physical Review Letters* 108.15 (2012), p. 150601.

[104] Masoud Abkenar, Thomas H Gray, and Alessio Zaccone. "Dissociation rates from single-molecule pulling experiments under large thermal fluctuations or large applied force". In: *Physical Review E* 95.4 (2017), p. 042413.

[105] Sai Sriharsha M Konda et al. "Chemical reactions modulated by mechanical stress: extended Bell theory". In: *The Journal of chemical physics* 135.16 (2011), p. 164103.

[106] Sergei Izrailev et al. "Molecular dynamics study of unbinding of the avidin-biotin complex". In: *Biophysical journal* 72.4 (1997), pp. 1568–1581.

[107] Sebastian Getfert and Peter Reimann. "Optimal evaluation of single-molecule force spectroscopy experiments". In: *Physical Review E* 76.5 (2007), p. 052901.

[108] Pauli Virtanen et al. "SciPy 1.0: fundamental algorithms for scientific computing in Python". In: *Nature methods* 17.3 (2020), pp. 261–272.

[109] Sudeep Adhikari and KSD Beach. "Reliable extraction of energy landscape properties from critical force distributions". In: *Physical Review Research* 2.2 (2020), p. 023276.

[110] Michael Hinczewski et al. "From mechanical folding trajectories to intrinsic energy landscapes of biopolymers". In: *Proceedings of the National Academy of Sciences* 110.12 (2013), pp. 4500–4505.

[111] Pilar Cossio, Gerhard Hummer, and Attila Szabo. "On artifacts in single-molecule force spectroscopy". In: *Proceedings of the National Academy of Sciences* 112.46 (2015), pp. 14248–14253.

[112] Alexander M Berezhkovskii et al. "Multidimensional reaction rate theory with anisotropic diffusion". In: *The Journal of chemical physics* 141.20 (2014), 11B616_1.

[113] Ajay P Manuel, John Lambert, and Michael T Woodside. "Reconstructing folding energy landscapes from splitting probability analysis of single-molecule trajectories". In: *Proceedings of the National Academy of Sciences* 112.23 (2015), pp. 7183–7188.

[114] Krishna Neupane and Michael T Woodside. "Quantifying instrumental artifacts in folding kinetics measured by single-molecule force spectroscopy". In: *Biophysical journal* 111.2 (2016), pp. 283–286.

[115] Gerhard Hummer and Attila Szabo. "Free energy profiles from single-molecule pulling experiments". In: *Proceedings of the National Academy of Sciences* 107.50 (2010), pp. 21441–21446.

[116] Alexander Berezhkovskii and Attila Szabo. "One-dimensional reaction coordinates for diffusive activated rate processes in many dimensions". In: *The Journal of chemical physics* 122.1 (2005), p. 014503.

[117] Jing Wu et al. "The role of solvent quality and chain stiffness on the end-to-end contact kinetics of semiflexible polymers". In: *The Journal of Chemical Physics* 149.23 (2018), p. 234903.

[118] Sanghyun Park and Klaus Schulten. "Calculating potentials of mean force from steered molecular dynamics simulations". In: *The Journal of chemical physics* 120.13 (2004), pp. 5946–5961.

[119] Carlos J Bustamante et al. "Optical tweezers in single-molecule biophysics". In: *Nature Reviews Methods Primers* 1.1 (2021), pp. 1–29.

[120] Rafayel Petrosyan, Abhishek Narayan, and Michael T Woodside. "Single-molecule force spectroscopy of protein folding". In: *Journal of Molecular Biology* 433.20 (2021), p. 167207.

[121] Bin Wang and Olga K Dudko. "A theory of synaptic transmission". In: *Elife* 10 (2021), e73585.

[122] Yaojun Zhang and Olga K Dudko. "Statistical mechanics of viral entry". In: *Physical review letters* 114.1 (2015), p. 018104.

[123] Arthur F Voter. "Hyperdynamics: Accelerated molecular dynamics of infrequent events". In: *Physical Review Letters* 78.20 (1997), p. 3908.

[124] Willmor J Peña Ccoa and Glen M Hocky. "Assessing models of force-dependent unbinding rates via infrequent metadynamics". In: *The Journal of Chemical Physics* 156.12 (2022), p. 125102.

[125] Salman A Khan, Bradley M Dickson, and Baron Peters. "How fluxional reactants limit the accuracy/efficiency of infrequent metadynamics". In: *The Journal of Chemical Physics* 153.5 (2020), p. 054125.

[126] Pratyush Tiwary and Michele Parrinello. "From metadynamics to dynamics". In: *Physical review letters* 111.23 (2013), p. 230602.

[127] Keir C Neuman and Attila Nagy. "Single-molecule force spectroscopy: optical tweezers, magnetic tweezers and atomic force microscopy". In: *Nature methods* 5.6 (2008), pp. 491–505.

[128] Carlos Bustamante et al. "Single-molecule studies of protein folding with optical tweezers". In: *Annual review of biochemistry* 89 (2020), pp. 443–470.

[129] Bharat Jagannathan and Susan Marqusee. "Protein folding and unfolding under force". In: *Biopolymers* 99.11 (2013), pp. 860–869.

[130] Dustin B Ritchie and Michael T Woodside. "Probing the structural dynamics of proteins and nucleic acids with optical tweezers". In: *Current opinion in structural biology* 34 (2015), pp. 43–51.

[131] Ken A Dill and Justin L MacCallum. "The protein-folding problem, 50 years on". In: *science* 338.6110 (2012), pp. 1042–1046.

[132] S Walter Englander and Leland Mayne. "The nature of protein folding pathways". In: *Proceedings of the National Academy of Sciences* 111.45 (2014), pp. 15873–15880.

[133] S Walter Englander and Leland Mayne. "The case for defined protein folding pathways". In: *Proceedings of the National Academy of Sciences* 114.31 (2017), pp. 8253–8258.

[134] Xiang Ye et al. "Folding of maltose binding protein outside of and in GroEL". In: *Proceedings of the National Academy of Sciences* 115.3 (2018), pp. 519–524.

[135] Xingjie Pan et al. "Expanding the space of protein geometries by computational design of de novo fold families". In: *Science* 369.6507 (2020), pp. 1132–1136.

[136] Kirill E Medvedev, Lisa N Kinch, and Nick V Grishin. "Functional and evolutionary analysis of viral proteins containing a Rossmann-like fold". In: *Protein Science* 27.8 (2018), pp. 1450–1463.

[137] Alexander Mehlich et al. "Slow transition path times reveal a complex folding barrier in a designed protein". In: *Frontiers in Chemistry* 8 (2020), p. 587824.

[138] John F Marko and Eric D Siggia. "Stretching dna". In: *Macromolecules* 28.26 (1995), pp. 8759–8770.

[139] Theo Odijk. "Stiff chains and filaments under tension". In: *Macromolecules* 28.20 (1995), pp. 7016–7018.

[140] Bennett Kalafut and Koen Visscher. "An objective, model-independent method for detection of non-uniform steps in noisy signals". In: *Computer physics communications* 179.10 (2008), pp. 716–723.

[141] Michelle D Wang et al. "Stretching DNA with optical tweezers". In: *Biophysical journal* 72.3 (1997), pp. 1335–1346.

[142] Krishna Neupane et al. "Transition path times for nucleic acid folding determined from energy-landscape analysis of single-molecule trajectories". In: *Physical review letters* 109.6 (2012), p. 068102.

[143] Krishna Neupane et al. "Transition-path probability as a test of reaction-coordinate quality reveals DNA hairpin folding is a one-dimensional diffusive process". In: *The journal of physical chemistry letters* 6.6 (2015), pp. 1005–1010.

[144] Jeffrey K Noel et al. "SMOG 2: a versatile software package for generating structure-based models". In: *PLoS computational biology* 12.3 (2016), e1004794.

[145] Martin Floor et al. "SBMOpenMM: a builder of structure-based models for OpenMM". In: *Journal of Chemical Information and Modeling* 61.7 (2021), pp. 3166–3171.

[146] Dmitri K Klimov and D Thirumalai. "Native topology determines force-induced unfolding pathways in globular proteins". In: *Proceedings of the National Academy of Sciences* 97.13 (2000), pp. 7254–7259.

[147] Eric R Henry, Robert B Best, and William A Eaton. "Comparing a simple theoretical model for protein folding with all-atom molecular dynamics simulations". In: *Proceedings of the National Academy of Sciences* 110.44 (2013), pp. 17880–17885.

[148] Robert B Best, Gerhard Hummer, and William A Eaton. "Native contacts determine protein folding mechanisms in atomistic simulations". In: *Proceedings of the National Academy of Sciences* 110.44 (2013), pp. 17874–17879.

[149] Himadri S Samanta et al. "Protein collapse is encoded in the folded state architecture". In: *Soft Matter* 13.19 (2017), pp. 3622–3638.

[150] Charles H Bennett. "Efficient estimation of free energy differences from Monte Carlo data". In: *Journal of Computational Physics* 22.2 (1976), pp. 245–268.

[151] Michael R Shirts et al. "Equilibrium free energies from nonequilibrium measurements using maximum-likelihood methods". In: *Physical review letters* 91.14 (2003), p. 140601.

[152] Michael R Shirts and Vijay S Pande. "Comparison of efficiency and bias of free energies computed by exponential averaging, the Bennett acceptance ratio, and thermodynamic integration". In: *The Journal of chemical physics* 122.14 (2005), p. 144107.

[153] David A Sivak and Gavin E Crooks. "Thermodynamic geometry of minimum-dissipation driven barrier crossing". In: *Physical Review E* 94.5 (2016), p. 052106.

[154] Sara Tafoya et al. "Using a system's equilibrium behavior to reduce its energy dissipation in nonequilibrium processes". In: *Proceedings of the National Academy of Sciences* 116.13 (2019), pp. 5920–5924.