

Travis Leadbetter, Prashant K. Purohit and Celia Reina*

From Langevin dynamics to macroscopic thermodynamic models: a general framework valid far from equilibrium

<https://doi.org/10.1515/jnet-2025-0071>

Received June 9, 2025; accepted September 26, 2025; published online October 16, 2025

Abstract: Given a particle system obeying overdamped Langevin dynamics, we demonstrate that it is always possible to construct a thermodynamically consistent macroscopic model which obeys a gradient flow with respect to its non-equilibrium free energy. To do so, we significantly extend the recent Stochastic Thermodynamics with Internal Variables (STIV) framework, a method for producing macroscopic thermodynamic models far-from-equilibrium from the underlying mesoscopic dynamics and an approximate probability density of states parameterized with so-called internal variables. Though originally explored for Gaussian probability distributions, we here allow for an arbitrary choice of the approximate probability density while retaining a gradient flow dynamics. This greatly extends its range of applicability and automatically ensures consistency with the second law of thermodynamics, without the need for secondary verification. We demonstrate numerical convergence, in the limit of increasing internal variables, to the true probability density of states for both a multi-modal relaxation problem, a protein diffusing on a strand of DNA, and for an externally driven particle in a periodic landscape. Finally, we provide a reformulation of STIV with the quasi-equilibrium approximations in terms of the averages of observables of the mesostate, and show that these, too, obey a gradient flow.

Keywords: stochastic thermodynamics; thermodynamics with internal variables; gradient flow

1 Introduction

The field of non-equilibrium thermodynamics boldly seeks to describe the macroscopic behavior of a wide range of physical systems. Using various formulations, such as irreversible thermodynamics [1], [2], rational thermodynamics [3], [4], thermodynamics with internal variables [5], or the General Equation for Non-Equilibrium Reversible-Irreversible Coupling (GENERIC) formalism [6]–[8], complex phenomena like creep [9], damage [10]–[12], thermoplasticity [7], and fluid rheology [13], to name a few, have been modeled effectively and systematically. However, central to our understanding, and at the heart of Hilbert's sixth problem, is the ability to formally and practically connect these macroscopic models to more fundamental microscopic physics. Perhaps the two most important contributions on this front are the Mori–Zwanzig formalism and other projection operator approaches which seek to project out noisy or irrelevant aspects of a Hamiltonian system, leaving only quantities which vary over long time scales, and the utilization of large deviation principles to

*Corresponding author: Celia Reina, Department of Mechanical Engineering and Applied Mechanics, University of Pennsylvania, 220 South 33rd Street, 229 Towne Building, Philadelphia, PA, 19104, USA, E-mail: creina@seas.upenn.edu.

<https://orcid.org/0000-0001-9302-3401>

Travis Leadbetter, Graduate Group in Applied Mathematics and Computational Science, University of Pennsylvania, Philadelphia, PA, 19104, USA. <https://orcid.org/0000-0001-7389-6248>

Prashant K. Purohit, Department of Mechanical Engineering and Applied Mechanics, University of Pennsylvania, 220 South 33rd Street, 229 Towne Building, Philadelphia, PA, 19104, USA. <https://orcid.org/0000-0002-1939-907X>

 Open Access. © 2025 the author(s), published by De Gruyter.  This work is licensed under the Creative Commons Attribution 4.0 International License.

coarse-grain microscopic stochastic processes [14]–[17]. Although these two theories provide the cornerstone of our theoretical understanding, significant challenges limit their utility in producing usable models.

In a previous work [18], the authors leveraged the theory of stochastic thermodynamics [19], [20] and a variational principle due to Eyink [21] to derive thermodynamically consistent macroscopic models directly from microscopic particle systems obeying an overdamped Langevin dynamics of the form

$$dx_t = -\frac{1}{\eta} \partial_x e \, dt + \sqrt{\frac{2}{\eta\beta}} \, dw_t, \quad (1)$$

where the interaction energy $e(x, \lambda)$ is potentially dependent on an external driving protocol λ , η and β are the viscosity and inverse temperature respectively (both assumed to be constant), and w_t is a standard Brownian motion. Stochastic thermodynamics provided the definitions for thermodynamic quantities both at the level of individual trajectories and at the macroscopic scale via ensemble averages. The system was then modeled using a parametrized approximate probability density of states, where the dynamics of the parameters (called internal variables) were determined using the variational method of Eyink. One obtained an approximate macroscopic thermodynamic model with internal variables by replacing the true probability density of states with its parameterized approximation in the definitions of macroscopic quantities from stochastic thermodynamics. The resulting macroscopic models exhibited the thermodynamics with internal variable structure for an arbitrary choice of approximation to the probability density of states, and hence the framework was named Stochastic Thermodynamics with Internal Variables (STIV). For the particular choice of a Gaussian approximation to the probability density of states, the dynamics of the internal variables obeyed a gradient flow with respect to the non-equilibrium free energy, $\hat{A}^{\text{neq}}$, thereby guaranteeing a non-negative rate of total entropy production. The STIV framework with a Gaussian approximation to the probability density of states has proven useful so far in providing the first ab-initio derivation of phase field models of front propagation in elastic media [22] as well as capturing the evolution and spreading of orientations of cell populations [23].

In this work, we expand the utility of the STIV framework by showing that the Gaussian approximation is not the only way to ensure the gradient flow structure, and hence a non-negative rate of total entropy production. In fact, the gradient flow structure can be achieved for any choice of approximate probability density of states. This is fitting, as it is well known that the Fokker–Planck equation associated with Langevin dynamics can be formulated as a gradient flow of the free energy $A^{\text{neq}} = \int (e + \beta^{-1} \log p) p \, dx$ with respect to the Wasserstein-2 metric [24]

$$\partial_t p = \frac{1}{\eta} \partial_{x_i} \left(p \, \partial_{x_i} \frac{\delta A^{\text{neq}}}{\delta p} \right). \quad (2)$$

Preserving the gradient flow structure offers several benefits. As already mentioned, within the STIV framework it guarantees that the approximate rate of total entropy production is non-negative and ensures thermodynamic consistency with the second law of thermodynamics. Typically, the second law must be enforced in an additional step, such as through the Coleman–Noll procedure [25]. Structure preservation has also been found both to improve the numerical accuracy of an approximation, and lead to more accurate qualitative behaviors [26]. This has proved vital to the numerical study of complex systems in many fields, including celestial mechanics [27], [28], molecular dynamics [29], control theory [30], and non-equilibrium mechanics [31].

2 Derivation

The dynamical equations for the STIV framework are derived from the variational method of Eyink [21], which provides a method for approximating the solution to differential equations arising in non-equilibrium statistical

mechanics. Starting from the Fokker–Planck equation associated with the microscopic stochastic differential equation in Eq. (1)¹

$$\partial_t p(x, t) = \mathcal{L} p(x, t) = \partial_{x_i} \left(\frac{1}{\eta} p(x, t) \partial_{x_i} e + \frac{1}{\eta \beta} \partial_{x_i} p(x, t) \right)$$

one can recast the problem in a variational form by introducing a test function ψ , and the action

$$\Gamma[\psi, p] = \int_0^\infty \int_X \psi (\partial_t - \mathcal{L}) p \, dx dt.$$

Stationary points (ψ^*, p^*) of this action are solutions to the adjoint and standard forms of the original Fokker–Planck equation respectively, and solve an infinite dimensional Hamilton equation

$$\begin{aligned} \partial_t p &= \frac{\delta}{\delta \psi} \mathcal{H}[\psi, p] \\ \partial_t \psi &= -\frac{\delta}{\delta p} \mathcal{H}[\psi, p] \end{aligned}$$

where $\mathcal{H} = \int_X \psi \mathcal{L} p \, dx$ is the Hamiltonian. The adjoint equation is trivially solvable with solution $\psi^* \equiv 1$. The equation for p is intractable in general, so to find an approximate solution one proposes a parameterized approximate probability density of states, $\hat{p}(x, \alpha)$, and a parameterized test function, $\hat{\psi}(x, \alpha)$. Writing $\hat{u} = \log(\hat{p})$ and $\langle f \rangle_{\hat{p}} = \int f(x) \hat{p}(x, \alpha) dx$, the dynamical equations for the parameters α reduce to

$$\langle [\partial \hat{\psi}, \partial \hat{u}]_{ij} \rangle_{\hat{p}} \dot{\alpha}_j = \partial_{\alpha_i} \langle \mathcal{L}^\dagger \hat{\psi} \rangle_{\hat{p}}$$

where $[\partial \hat{\psi}, \partial \hat{u}]_{ij} = \partial_{\alpha_i} \hat{\psi} \partial_{\alpha_j} \hat{u} - \partial_{\alpha_j} \hat{\psi} \partial_{\alpha_i} \hat{u}$ and $\mathcal{L}^\dagger$ is the adjoint of $\mathcal{L}$. The matrix $\langle [\partial \hat{\psi}, \partial \hat{u}]_{ij} \rangle_{\hat{p}}$ is anti-symmetric and obeys the Jacobi identity, so these dynamics retain the Hamiltonian form.

In the original framework [18], the parameters were split into two sets: the variables parameterizing p denoted (with a slight abuse of notation) α and referred to as internal variables, and dummy variables γ of the same dimension as α . The approximate density is assumed to only depend on α while $\hat{\psi}$ was defined as $\hat{\psi} = 1 + \gamma_i \partial_{\alpha_i} \hat{s}$, where $\hat{s} = -k_B \hat{u} = -k_B \log(\hat{p})$ is the approximate microscopic entropy, and k_B is the Boltzmann constant. With this particular choice, the dynamical equations reduce to

$$\langle \partial_{\alpha_i} \hat{s} \partial_{\alpha_j} \hat{s} \rangle_{\hat{p}} \dot{\alpha}_j = -k_B \langle \mathcal{L}^\dagger \partial_{\alpha_i} \hat{s} \rangle_{\hat{p}}, \quad \gamma(t) \equiv 0.$$

This “linear ansatz” using the α -gradient of the microscopic entropy in the test function $\hat{\psi}$ produces a specific case of the moment-closure equations called “entropic matching” [32], and ultimately leads to irreversible dynamics for α .

Once the dynamics for the internal variables are obtained, one can approximate macroscopic thermodynamic quantities defined via stochastic thermodynamics by substituting the approximate probability density of states for the true solution. For example, the macroscopic non-equilibrium free energy defined as $A^{\text{neq}} = \langle e \rangle_p - T \langle s \rangle_p$ becomes $\hat{A}^{\text{neq}} = \langle e \rangle_{\hat{p}} - T \langle \hat{s} \rangle_{\hat{p}}$, from which one can approximate both the external force $F^{\text{ex}} = \langle \partial_\lambda e \rangle_p$ by $\hat{F}^{\text{ex}} = \partial_\lambda \hat{A}^{\text{neq}}$ and the rate of total entropy production $\frac{d}{dt} S^{\text{tot}}$ by $\frac{d}{dt} \hat{S}^{\text{tot}} = \frac{1}{T} \left(\hat{F}^{\text{ex}} \dot{\lambda} - \frac{d}{dt} \hat{A}^{\text{neq}} \right) = -\frac{1}{T} \partial_{\alpha_i} \hat{A}^{\text{neq}} \dot{\alpha}_i$.

Previously, the STIV framework was effectively used to study the dynamics and thermodynamics of phase front propagation in an elastic medium, where one could derive both the continuum limit [18], and an equivalent phase field model directly from microscopic physics [22]. This previous implementation, however, was limited to a multivariate Gaussian approximation to the probability density of states $\hat{p}$ as this was the only form in which the gradient flow structure, and hence non-decreasing total entropy production, could be ensured at the time. The key findings of this work are a number of options for preserving the gradient flow structure of the underlying Fokker–Planck equation which are flexible enough to approximate any probability density to arbitrary

¹ We shall use the Einstein summation convention, summing all repeated indices, but without regard for raised and lowered indices.

Table 1: Key features of the closure methods.^a

Method name	Operator method	Quasi-equilibrium	Exponential family
Probability density of states $\hat{p}$	Arbitrary	$\exp(-\beta e + \alpha \phi - f)$	$\exp(\alpha \phi - f)$
Test functions $\hat{\psi}$	$1 + \gamma \left(\mathcal{L}_0^\dagger \right)^{-1} \partial_\alpha \hat{u}$	$1 + \gamma \partial_\alpha \hat{u}$	$1 + \gamma \partial_\alpha \hat{u}$
Dynamical equations	$\mathbb{F} \dot{\alpha} = -\frac{1}{\eta} \partial_\alpha \hat{A}^{\text{neq}}$	$\mathbb{C} \dot{\alpha} = -\frac{1}{\eta} \mathbb{M} \mathbb{C}^{-1} \partial_\alpha \hat{A}^{\text{neq}}$	$\mathbb{C} \dot{\alpha} = -\frac{1}{\eta} \mathbb{M} \mathbb{C}^{-1} \partial_\alpha \hat{A}^{\text{neq}}$
Dynamics with external driving	$\mathbb{F} \dot{\alpha} = -\frac{1}{\eta} \partial_\alpha \hat{A}^{\text{neq}}$	$\mathbb{C} \dot{\alpha} + \langle \bar{\phi} \partial_\alpha \hat{u} \rangle_{\hat{p}} \dot{\beta} = -\frac{1}{\eta} \mathbb{M} \mathbb{C}^{-1} \partial_\alpha \hat{A}^{\text{neq}}$	$\mathbb{C} \dot{\alpha} = -\frac{1}{\eta} \mathbb{M} \mathbb{C}^{-1} \partial_\alpha \hat{A}^{\text{neq}}$
Dual dynamics	Not studied here	$\frac{d}{dt} \hat{\Phi} = -\frac{1}{\eta} \mathbb{M} \partial_\Phi \hat{A}^{\text{neq}}$	$\frac{d}{dt} \hat{\Phi} = -\frac{1}{\eta} \mathbb{M} \partial_\Phi \hat{A}^{\text{neq}}$
Entropy production $T \frac{d}{dt} \hat{S}^{\text{tot}}$	$\eta \mathbb{F}_{ij} \dot{\alpha}_i \dot{\alpha}_j$	$\frac{1}{\eta \beta^2} \mathbb{M}_{ij} \alpha_i \alpha_j$	$\eta \mathbb{M}_{ij}^{-1} \frac{d}{dt} \hat{\Phi}_i \frac{d}{dt} \hat{\Phi}_j$

^aDefinition of terms for Table 1: $\mathbb{C}_{ij} = \langle \partial_{\alpha_i} \hat{u} \partial_{\alpha_j} \hat{u} \rangle_{\hat{p}} = \text{Cov}_{\hat{p}}(\phi_i, \phi_j)$, $\mathbb{M}_{ij} = \langle \partial_{x_k} \partial_{\alpha_i} \hat{u} \partial_{x_k} \partial_{\alpha_j} \hat{u} \rangle_{\hat{p}} = \langle \partial_{x_k} \phi_i \partial_{x_k} \phi_j \rangle_{\hat{p}}$,

$\mathcal{L}_0^\dagger g = \partial_{x_i} \hat{u} \partial_{x_i} g + \partial_{x_i} \partial_{\alpha_i} g$, $\mathbb{F}_{ij} = \langle \partial_{\alpha_i} \hat{u} \left(\mathcal{L}_0^\dagger \right)^{-1} \partial_{\alpha_j} \hat{u} \rangle_{\hat{p}}$, $\bar{\phi} = \phi - \hat{\Phi}$, $\hat{\Phi} = \langle \phi \rangle_{\hat{p}}$. ^bThis equation holds with and without external driving.

^cThese dynamics are not given by moment closure, but approximate it for large number of observables.

accuracy in the limit of increasing internal variables. The first method, which we call the “operator method”, utilizes an alternative form for the approximate test function that allows for a completely unconstrained choice of approximate probability density of states. The second uses the original approximate test functions and entropic matching moment-closure equations, but requires one to restrict the approximate probability density of states using a “quasi-equilibrium” distributions [33]. The final method forgoes the need to include the Boltzmann factor $\exp(-\beta e(x, \lambda))$ in the approximate probability density of states, which could reduce computational costs in some cases, but the internal variable dynamics given by entropic matching only approximate a gradient flow. We label this method as “exponential family” throughout. A detailed analysis of each of these three methods is provided next, and the key elements and results for each of them are summarized in Table 1 as reference.

2.1 The operator method

We now motivate an alternative choice for the approximate test function $\hat{\psi}$ to show how the STIV method can produce gradient flow equations for an arbitrary approximate probability density of states. Again, we fix an approximate density $\hat{p}(x, \alpha)$, and let $\hat{u} = \log(\hat{p})$. We then assume that $\chi(x, \alpha)$ is a vector valued function with the same dimension as α , and define $\hat{\psi}(x, \alpha, \gamma) = 1 + \gamma_i \chi_i(x, \alpha)$. Following the same steps as above, the dynamical equations induced by the variational method become

$$\langle \chi_i \partial_{\alpha_j} \hat{u} \rangle_{\hat{p}} \dot{\alpha}_j = \langle \mathcal{L}^\dagger \chi_i \rangle_{\hat{p}}, \quad \gamma(t) = 0.$$

For convenience, we define $\mathbb{F}_{ij} \equiv \langle \chi_i \partial_{\alpha_j} \hat{u} \rangle_{\hat{p}}$ and $\mathbb{C}_{ij} \equiv \langle \partial_{\alpha_i} \hat{u} \partial_{\alpha_j} \hat{u} \rangle_{\hat{p}}$.

We have not specified new test functions yet, but to do so we examine the approximate rate of total entropy production, $T \dot{\hat{S}}^{\text{tot}} = -\dot{\alpha}_i \partial_{\alpha_i} \hat{A}^{\text{neq}}$. We recall that $\hat{A}^{\text{neq}} = \hat{E} - T \hat{S} = \langle e \rangle_{\hat{p}} + \langle \hat{u} / \beta \rangle_{\hat{p}}$, and since $0 = \partial_{\alpha_i} \langle 1 \rangle_{\hat{p}} = \langle \partial_{\alpha_i} \hat{u} \rangle_{\hat{p}}$ and e does not depend on α , we have $\partial_{\alpha_i} \hat{A}^{\text{neq}} = \langle (e + \hat{u} / \beta) \partial_{\alpha_i} \hat{u} \rangle_{\hat{p}}$. On the other hand, repeated integration by parts reveals that

$$\begin{aligned}
 \langle \mathcal{L}^\dagger \chi_i \rangle_{\hat{p}} &= \left\langle \left(-\frac{1}{\eta} \partial_{x_k} e \partial_{x_k} \chi_i + d \partial_{x_i} \partial_{x_i} \chi_i \right) \right\rangle_{\hat{p}} \\
 &= -\frac{1}{\eta} \langle \partial_{x_k} (e + \hat{u} / \beta) \partial_{x_k} \chi_i \rangle_{\hat{p}} \\
 &= \frac{1}{\eta} \langle (e + \hat{u} / \beta) (\partial_{x_k} \hat{u} \partial_{x_k} \chi_i + \partial_{x_k} \partial_{x_k} \chi_i) \rangle_{\hat{p}} \\
 &\equiv \frac{1}{\eta} \langle (e + \hat{u} / \beta) \mathcal{L}_0^\dagger \chi_i \rangle_{\hat{p}}.
 \end{aligned}$$

This new operator $\mathcal{L}_{\hat{u}}^\dagger g = \partial_{x_k} \hat{u} \partial_{x_k} g + \partial_{x_i} \partial_{x_i} g$ is also the adjoint of a Fokker Planck operator which only depends on $\hat{u}$. The final line resembles the equation for $\partial_{\alpha_i} \hat{A}^{\text{neq}}$, and suggests that one should choose test functions which obey

$$\mathcal{L}_{\hat{u}}^\dagger \chi_i = -\partial_{\alpha_i} \hat{u}. \quad (3)$$

Assuming one can solve this differential equation, the dynamical equations become a gradient flow

$$\mathbb{F}_{ij} \dot{\alpha}_j = -\frac{1}{\eta} \partial_{\alpha_i} \hat{A}^{\text{neq}}, \quad (4)$$

since $\mathbb{F}_{ij}$ is positive semi-definite as

$$\begin{aligned} \mathbb{F}_{ij} &= \left\langle \chi_i \partial_{\alpha_j} \hat{u} \right\rangle_{\hat{p}} \\ &= -\left\langle \chi_i \mathcal{L}_{\hat{u}}^\dagger \chi_j \right\rangle_{\hat{p}} \\ &= -\int \chi_i \partial_{x_k} (\hat{p} \partial_{x_k} \chi_j) dx \\ &= \left\langle \partial_{x_k} \chi_i \partial_{x_k} \chi_j \right\rangle_{\hat{p}}. \end{aligned}$$

Between the second and third lines, we have made use of the fact that $\hat{p} \mathcal{L}_{\hat{u}}^\dagger \chi_j = \partial_{x_k} (\hat{p} \partial_{x_k} \chi_j)$. With the gradient flow structure comes the guarantee of non-negative rate of total entropy production

$$T \dot{S}^{\text{tot}} = \eta \mathbb{F}_{ij} \dot{\alpha}_i \dot{\alpha}_j \geq 0.$$

This strategy allows for a completely general choice of approximate probability density of states $\hat{p}(x, \alpha)$. Solving Eq. (3), though, is challenging in general. However, following a well known trick [34], one can define a “Hamiltonian” conjugation of $\mathcal{L}_{\hat{u}}^\dagger$ by

$$-\mathcal{H}_{\hat{u}} g = \sqrt{\hat{p}} \mathcal{L}_{\hat{u}}^\dagger \left(\frac{g}{\sqrt{\hat{p}}} \right) = \partial_x^2 g - v(x, \alpha) g$$

where $v = \partial_x^2 \hat{u} / 2 + |\partial_x \hat{u}|^2 / 4$. Unlike $\mathcal{L}_{\hat{u}}^\dagger$, $\mathcal{H}_{\hat{u}}$ is self-adjoint with respect to the Lebesgue measure. Note, also that $\mathcal{H}_{\hat{u}}$ defines a Schrödinger operator, and it is immediately clear that if (λ^*, g^*) is an eigenvalue-eigenvector pair for $\mathcal{H}_{\hat{u}}$, then $(-\lambda^*, g^* / \sqrt{\hat{p}})$ is an eigenvalue-eigenvector pair for $\mathcal{L}_{\hat{u}}^\dagger$. Hence, one can take advantage of existing eigenvalue solvers for quantum systems in order to compute the necessary test functions.

2.2 Quasi-equilibrium distribution

The primary advantage of the previous method for producing a gradient flow dynamics is that it allows for an arbitrary choice of approximate probability density of states. The disadvantage is that the test functions must be computed as solutions to the differential equation in Eq. (3). Alternatively, we show now that if one restricts the approximate probability density of states to a specific class known as quasi-equilibrium densities (a specific form of exponential family) [33], then one may use the original test functions $\chi_i = \partial_{\alpha_i} \hat{u}$ and still achieve a gradient flow (here, we have replaced $\hat{s}$ by $\hat{u}$ for notational simplicity). This is only a minor restriction as any continuous probability distribution can be approximated by an exponential family or quasi-equilibrium distribution to arbitrary accuracy (see Appendix C for one such proof).

Assume that the approximate probability density of states takes a quasi-equilibrium form. That is

$$\hat{p}(x, \alpha, \beta) = \exp(-\beta e(x) + \alpha_i \phi_i(x) - f(\alpha, \beta)), \quad (5)$$

for some set of “observables” $\phi_i(x)$ and internal variables $\alpha_i(t)$. Here, β (taken as a constant) and the internal variables α , form the natural parameters of the exponential family, while the energy e and observables ϕ form the “sufficient statistics” [35]. It will be helpful to define $\hat{\Phi}_i$ as the $\hat{p}$ averages of the observables, i.e., $\hat{\Phi}_i \equiv \langle \phi_i \rangle_{\hat{p}}$.

To see the gradient flow structure, note first that in this case

$$\langle \chi_i \partial_{\alpha_j} \hat{u} \rangle_{\hat{p}} = \langle (\phi_i - \hat{\Phi}_i)(\phi_j - \hat{\Phi}_j) \rangle_{\hat{p}} = \mathbb{C}_{ij}.$$

Second, since $e + \hat{u}/\beta = (\alpha_i \phi_i - f(\alpha, \beta))/\beta$, we have

$$\partial_{\alpha_i} \hat{A}^{\text{neq}} = \langle (e + \hat{u}/\beta) \partial_{\alpha_i} \hat{u} \rangle_{\hat{p}} = \frac{\mathbb{C}_{ij} \alpha_j}{\beta}.$$

Finally, again by integrating by parts

$$\begin{aligned} \langle \mathcal{L}^\dagger \chi_i \rangle_{\hat{p}} &= -\frac{1}{\eta} \langle \partial_{x_k} (e + \hat{u}/\beta) \partial_{x_k} \chi_i \rangle_{\hat{p}} \\ &= -\frac{\langle \partial_{x_k} \phi_i \partial_{x_k} \phi_j \rangle_{\hat{p}} \alpha_j}{\eta \beta} \\ &\equiv -\frac{1}{\eta \beta} \mathbb{M}_{ij} \alpha_j. \end{aligned}$$

As with F in the earlier discussion, $\mathbb{M}_{ij} = \langle \partial_{x_k} \phi_i \partial_{x_k} \phi_j \rangle_{\hat{p}}$ is positive semi-definite. Putting the pieces together, we get

$$\mathbb{C}_{ij} \dot{\alpha}_j = -\frac{1}{\eta \beta} \mathbb{M}_{ij} \alpha_j \quad (6)$$

or, by inserting the identity as $\delta_{jk} = \mathbb{C}_{jl}^{-1} \mathbb{C}_{lk}$ between $\mathbb{M}$ and α (we assume $\mathbb{C}$ is invertible),

$$\dot{\alpha}_i = -\frac{1}{\eta} \mathbb{C}_{ij}^{-1} \mathbb{M}_{jk} \mathbb{C}_{kl}^{-1} \partial_{\alpha_i} \hat{A}^{\text{neq}}. \quad (7)$$

Since $\mathbb{M}$ is positive semi-definite, and $\mathbb{C}$ is symmetric, $\mathbb{C}^{-1} \mathbb{M} \mathbb{C}^{-1}$ is positive semi-definite and the dynamics obey a gradient flow.

2.3 Exponential family

Using a quasi-equilibrium distribution which includes the total energy leads to gradient flow dynamics exactly. However, one may not always wish to include the energy in the approximate density of states, for example, if it would make computing the normalizing function more challenging. If one removes the energy and uses an exponential family of the form

$$\hat{p}(x, \alpha) = \exp(\alpha_i \phi_i - f(\alpha)) \quad (8)$$

the dynamics are not always a gradient flow. Specifically, they become

$$\begin{aligned} \mathbb{C} \dot{\alpha} &= -\frac{1}{\eta} \mathbb{M} \mathbb{C}^{-1} \partial_{\alpha} \hat{A}^{\text{neq}} \\ &\quad - \frac{1}{\eta} \left(\langle \partial_{x_k} e \partial_{x_k} \phi \rangle_{\hat{p}} - \mathbb{M} \mathbb{C}^{-1} \langle e(\phi - \hat{\Phi}) \rangle_{\hat{p}} \right). \end{aligned} \quad (9)$$

The second term on the right hand side only depends on the (non-constant) part of the energy which is $\hat{p}$ -orthogonal to the span of the ϕ . Mathematically, if we can write $e(x) = \gamma_j(\phi_j(x) - \hat{\Phi}_j) + e_0 + e^\perp(x)$ where $\langle e^\perp(\phi_j - \hat{\Phi}_j) \rangle_{\hat{p}} = 0$ for all j , then the second term becomes

$$\left(\langle \partial_{x_k} e \partial_{x_k} \phi \rangle_{\hat{p}} - \mathbb{M} \mathbb{C}^{-1} \langle e(\phi - \hat{\Phi}) \rangle_{\hat{p}} \right) = \langle \partial_{x_k} e^\perp \partial_{x_k} \phi \rangle_{\hat{p}}.$$

If the observables are chosen to produce a highly flexible approximation, for example by using machine learning techniques or a basis expansion, it is reasonable to assume that this term is negligible and therefore one

can set it to zero, giving a gradient flow structure. In this case, $\partial_\alpha \hat{A}^{\text{neq}} = \mathbb{C}\alpha/\beta + \left\langle e(\phi - \hat{\Phi}) \right\rangle_{\hat{p}}$, and again the dynamical equations are given by Eq. (7). To minimize confusion, we shall refer to an approximate probability density of states of the form $\hat{p} = \exp(\alpha_i \phi_i - f(\alpha))$ which does not include the energy as an exponential family approximation. When $\hat{p} = \exp(-\beta e + \alpha_i \phi_i - f(\alpha, \beta))$ explicitly includes the energy, we shall refer to this as a quasi-equilibrium approximation (although it is also an exponential family).

2.4 Time dependent external driving

Now we consider the case when the total energy is time dependent through some external driving protocol, $e = e(x, \lambda(t))$. In the operator formulation, with test functions $\mathcal{L}_u \chi_i = \partial_{\alpha_i} \hat{u}$, nothing needs change to maintain the original STIV definitions. Since the approximate probability density of states can still be chosen to be independent of λ (for fixed α), the approximate external force remains $\hat{F}^{\text{ex}} = \partial_\lambda \hat{A}^{\text{neq}}$ [18]. In the quasi-equilibrium formulation, $\hat{p}$ explicitly depends on λ through e and so it is no longer true that $\hat{F}^{\text{ex}} = \partial_\lambda \hat{A}^{\text{neq}}$. Moreover, when deriving the dynamical equations from the variational method [21], it was assumed that all of the time dependence of $\hat{p}$ was captured through the internal variables α . When $\hat{p}$ takes the quasi-equilibrium form this is not the case since the energy is now time dependent. This is fixed through the proper modification of the dynamical equations given in Eq. (6) (which can be found in Appendix A). The resulting dynamical equations are

$$\mathbb{C}_{ij} \dot{\alpha}_j + \left\langle (\phi_i - \hat{\Phi}_i) \partial_\lambda \hat{u} \right\rangle_{\hat{p}} \dot{\lambda} = -\frac{1}{\eta\beta} \mathbb{M}_{ik} \alpha_k \quad (10)$$

and the external force becomes

$$\hat{F}^{\text{ex}} = \partial_\lambda \left(\hat{A}^{\text{neq}} - \frac{\alpha_i \hat{\Phi}_i}{\beta} \right) = -\frac{1}{\beta} \partial_\lambda f(\alpha, \beta, \lambda). \quad (11)$$

The equation for the rate of total entropy production is also impacted, but it still remains non-negative

$$T \frac{d}{dt} \hat{S}^{\text{tot}} = \frac{1}{\eta\beta^2} \alpha_i \mathbb{M}_{ij} \alpha_j \quad (12)$$

(see the Appendix B for the derivation of both the approximate external force and approximate rate of total entropy production).

The new dynamical equations appear to break the gradient flow structure. However, in the following section, we show that the gradient flow structure remains intact when the equations are recast in terms of the dynamics of the averages of observables $\hat{\Phi}$.

2.5 Dual formulation

So far, the dynamics of the system have been tracked through the internal variables α . However, in traditional internal variable formulations, the internal variables are generally assumed to be averages of microscopic sub-structures which are not captured in equilibrium. Thus, it seems more natural to consider the averages of the observables, that is $\hat{\Phi}_i = \langle \phi_i \rangle_{\hat{p}}$ as independent variables, and develop the system's dynamics in terms of their behavior. Fortunately, this is readily achievable through a simple change of variables. Assuming the map from the internal variables to the averages of the observables, $\alpha \rightarrow \hat{\Phi}$, is invertible, we can define the non-equilibrium free energy as a function of $\hat{\Phi}$. Using the chain rule, we find that $\partial_\alpha \hat{A}^{\text{neq}} = \mathbb{C} \partial_\Phi \hat{A}^{\text{neq}}$. Combining this with

$$\begin{aligned} \frac{d}{dt} \hat{\Phi}_i &= \left\langle \phi_i \partial_{\alpha_j} u \right\rangle_{\hat{p}} \dot{\alpha}_j + \left\langle \phi_i \partial_\lambda \hat{u} \right\rangle_{\hat{p}} \dot{\lambda} \\ &= \mathbb{C}_{ij} \dot{\alpha}_j + \left\langle (\phi_i - \hat{\Phi}_i) \partial_\lambda \hat{u} \right\rangle_{\hat{p}} \dot{\lambda}, \end{aligned}$$

which reproduces the left hand side of Eq. (10) (or Eqs. (6) and (9) when $\hat{p}$ doesn't depend on λ), gives the equivalent dynamics for $\hat{\Phi}$

$$\frac{d}{dt}\hat{\Phi}_i = -\frac{1}{\eta}\mathbb{M}_{ij}\partial_{\Phi_j}\hat{A}^{\text{neq}}$$

which again have the gradient flow structure. This holds true for both the quasi-equilibrium and exponential family approximations. Moreover, when $\hat{p}$ is independent of λ (either the exponential family approximation or the quasi-equilibrium approximation with no external driving), the form of the rate of total entropy production is invariant under the change of variables [22]

$$\begin{aligned} T\frac{d}{dt}\hat{S}^{\text{tot}} &= -\partial_{\alpha_i}\hat{A}^{\text{neq}}(\alpha_i, \beta)\dot{\alpha}_i \\ &= -\partial_{\Phi_i}\hat{A}^{\text{neq}}(\hat{\Phi}, \beta)\frac{d\hat{\Phi}_i}{dt} \\ &= \eta\frac{d\hat{\Phi}_i}{dt}\mathbb{M}_{ij}^{-1}\frac{d\hat{\Phi}_j}{dt}. \end{aligned}$$

In this form, it becomes clear that $\mathbb{M}_{ij} = \langle \partial_x \phi_i \partial_x \phi_j \rangle_{\hat{p}}$ is an approximation to the true average dissipation matrix of the observables $\phi_i(x_t)$ along fluctuating trajectories. Mathematically, since the governing Langevin equation is given by

$$dx_t = -\partial_x e / \eta dt + \sqrt{2/\eta\beta} dw_t,$$

we know by Itô's formula that

$$d\phi_i(x_t) = \left(-\frac{1}{\eta}\partial_x e \partial_x \phi_i + \frac{1}{\eta\beta}\Delta_x \phi_i \right) dt + \sqrt{\frac{2}{\eta\beta}}\partial_x \phi_i dw_t$$

and hence the diffusion matrix of $\phi(x_t)$ is $\mathcal{M}_{ij}(x_t) \propto \partial_x \phi_i \partial_x \phi_j$. When $\hat{p}$ depends on λ , the approximate rate of total entropy production cannot be written as a quadratic form in $\frac{d}{dt}\hat{\Phi}$ but is still given by Eq. (12).

It is also interesting to note that in the case of the quasi-equilibrium approximation, the non-equilibrium free energy is the convex dual of the normalizing function

$$f(\alpha, \beta, \lambda) = \log \left(\int \exp(-\beta e + \alpha \phi) dx \right).$$

f is convex in α since $\partial_{\alpha_i} \partial_{\alpha_j} f = \langle (\phi_i - \hat{\Phi}_i)(\phi_j - \hat{\Phi}_j) \rangle_{\hat{p}} = \mathbb{C}_{ij}$ is positive semi-definite and we know

$$\hat{\Phi}_i = \langle \phi_i \rangle_{\hat{p}} = \partial_{\alpha_i} f.$$

Using the Legendre transform, we can define the convex dual to f as $f^*(\Phi, \beta, \lambda) = \sup_{\alpha} (\alpha \Phi - f(\alpha, \beta, \lambda))$ so that

$$\alpha_i = \partial_{\Phi_i} f^*(\Phi = \hat{\Phi}, \beta, \lambda).$$

Looking at the definition for $\hat{A}^{\text{neq}}(\alpha, \beta, \lambda) = \langle e + \hat{u}/\beta \rangle_{\hat{p}} = \langle (\alpha_i \phi_i - f(\alpha, \beta, \lambda)) \rangle_{\hat{p}} / \beta = (\alpha_i \hat{\Phi}_i - f(\alpha, \beta, \lambda)) / \beta$, it's clear that $\hat{A}^{\text{neq}}(\Phi, \beta, \lambda) = f^*(\Phi, \beta, \lambda) / \beta$. Since $\hat{A}^{\text{neq}}(\alpha, \beta, \lambda) \neq f(\alpha, \beta, \lambda)$, the non-equilibrium free energies driving the dynamics for α and $\hat{\Phi}$, that is $\hat{A}^{\text{neq}}(\alpha)$ and $\hat{A}^{\text{neq}}(\Phi)$, are not related through a Legendre transform as is typically the case when one changes between dual variables in classical thermodynamics.

3 An information theoretic view

Before moving on to the example implementations of the STIV framework, we highlight an information theoretic perspective on the dynamical equations for the internal variables based on minimizing the Kullback–Leibler (KL) divergence between the approximate and true probability density of states. Following Bronstein and Koepl

[32], we assume that the true probability density of states is governed by the equation $\partial_t p = \mathcal{L} p$, and try to minimize the KL-divergence

$$\mathbb{D}_{\text{KL}}(\alpha) \equiv \mathbb{D}_{\text{KL}}[\hat{p} \| p] = \int \hat{p} \log\left(\frac{\hat{p}}{p}\right) dx.$$

Differentiating in α gives

$$0 = \int \log\left(\frac{\hat{p}}{p}\right) \partial_{\alpha_i} \hat{p} dx.$$

We would like this to be true for all time, and hence we differentiate in time to get

$$\begin{aligned} 0 = & \int \left\{ \log\left(\frac{\hat{p}}{p}\right) \partial_{\alpha_i} \partial_{\alpha_j} \hat{p} + \hat{p} \partial_{\alpha_i} \hat{u} \partial_{\alpha_j} \hat{u} \right\} \dot{\alpha}_j dx \\ & + \int \left\{ \log\left(\frac{\hat{p}}{p}\right) \partial_{\alpha_i} \partial_{\lambda} \hat{p} + \hat{p} \partial_{\alpha_i} \hat{u} \partial_{\lambda} \hat{u} \right\} \dot{\lambda} - \int \partial_{\alpha_i} \hat{u} \frac{\mathcal{L} p}{p} \hat{p} dx \end{aligned}$$

or in terms of expectations with respect to $\hat{p}$,

$$\begin{aligned} \left\langle \partial_{\alpha_i} \hat{u} \frac{\mathcal{L} p}{p} \right\rangle_{\hat{p}} &= \left\langle \log\left(\frac{\hat{p}}{p}\right) \frac{\partial_{\alpha_i} \partial_{\alpha_j} \hat{p}}{\hat{p}} \right\rangle_{\hat{p}} \dot{\alpha}_j + \mathbb{C}_{ij} \dot{\alpha}_j \\ &+ \left\langle \log\left(\frac{\hat{p}}{p}\right) \frac{\partial_{\alpha_i} \partial_{\lambda} \hat{p}}{\hat{p}} \right\rangle_{\hat{p}} \dot{\lambda} + \langle \partial_{\alpha_i} \hat{u} \partial_{\lambda} \hat{u} \rangle_{\hat{p}} \dot{\lambda}. \end{aligned}$$

Since this equation contains p , it is unlikely that one could simplify it without significant approximation. Moreover, the best approximation one has for p at any given time within our framework is $\hat{p}$. If we make this substitution in the previous equation, we achieve

$$\mathbb{C}_{ij} \dot{\alpha}_j + \langle \partial_{\alpha_i} \hat{u} \partial_{\lambda} \hat{u} \rangle_{\hat{p}} \dot{\lambda} = \langle \mathcal{L}^\dagger \partial_{\alpha_i} \hat{u} \rangle_{\hat{p}}$$

which is exactly the original dynamical equations given by the variational method of Eyink. Thus, these dynamics minimize the KL-divergence between the true probability density of states and the approximate probability density of states at each instance in time.

4 Examples

As a demonstration of the enhanced flexibility of the modified STIV framework presented above, we compare the results of STIV using an exponential family and quasi-equilibrium approximation to the true probability density of states for two example systems. The first example is that of a colloidal protein diffusing along a string of DNA [36], [37]. This system is approximately described by a single particle diffusing in a one dimensional potential consisting of the product of a sinusoidal function and a decaying exponential. Mathematically,

$$u(x) = c_0 \exp(-c_1 x) \sin(c_2 x + c_3) + \text{bounding box}. \quad (13)$$

The parameters $c_0, \dots, c_3$ were chosen to fit the potential shown in Figure 6 of [37], and the bounding box is specified to hold the particle within 60 base pairs of deepest minima (the values of the parameters and the bounding box are detailed in Appendix D). This example allows us to study the approximation error during a relaxation towards equilibrium. The second example includes time dependent driving. A colloidal particle is assumed to diffuse in a sinusoidal landscape while coupled to an external driving protocol via a harmonic potential, a prototypical model of an optical trap. The external driving moves in the positive x axis with constant velocity v . Mathematically,

$$e(x, \lambda) = -\cos\left(\frac{2\pi x}{l}\right) + \frac{k}{2}(x - \lambda)^2, \quad (14)$$

with $\lambda(t) = vt$ (see Appendix D for parameter details). Here, we can evaluate the accumulation of error in time since this system never reaches a steady state.

For these examples, we compare the approximation accuracy of the quasi-equilibrium approximate density as in Eq. (5), and a traditional exponential family without the energy as in Eq. (8). Since both examples are one dimensional, we chose sinusoidal Fourier expansion functions as the observables $\phi_n = \sin(n\pi(x - x_0)/L)$ where x_0 and L are chosen to fit the bounding box in the diffusing protein example or the spatial extent of the simulation in the external driving example. Due to the computational complexity of solving the differential equation $\mathcal{L}_h^\dagger \chi_i = \partial_{\alpha_i} \hat{u}$ to find χ_i at each step, we choose not to implement the operator method. It is likely that this method will only become feasible for specific choices of observables for which the differential equation can be inverted exactly (as happens to be the case for a univariate Gaussian).

Figure 1 shows results for the colloidal protein example, and Figure 2 shows those of the time dependent driving example. Three different implementations of the STIV framework are highlighted throughout. The first, labeled “ $\hat{p} \propto \exp(-\beta e + \alpha \phi)$ ” and colored green, uses the quasi-equilibrium probability density of states and the dynamics obtained in Eq. (6). The second, labeled “ $\hat{p} \propto \exp(\alpha \phi)$ Moment Closure” and colored yellow, uses just the exponential family approximation and the dynamics obtained from the variational method of Eyink, Eq. (9). Finally, the method labeled “ $\hat{p} \propto \exp(\alpha \phi)$ Gradient Flow” and colored pink also uses the exponential family approximation but the second term on the right hand side of Eq. (9) (which is expected to vanish as the number of observables increases) is set to zero to obtain a gradient flow. The convergence of each method to the ground truth probability density of states is shown in panel A. The metric used is the time averaged total variation distance between the ground truth and the approximation




Figure 1: The convergence of STIV to the true probability density of states in a model of protein diffusion on DNA. Panel A shows the time averaged total variation distance between each STIV model (the quasi-equilibrium approximation in green, the exponential family approximation without the energy in yellow, and the exponential family approximation with modified dynamics to produce a gradient flow in pink) and the true probability density of states (obtained via a traditional solver) as a function of the number of internal variables. Panel B depicts the energy landscape. Panels C and D show the approximate probability density of states for each model for 8 and 32 internal variables respectively and the true probability density of states (in dark blue) at the same fixed time point (midway between the initial time and fully relaxed in equilibrium).




Figure 2: Investigation of the error accumulation of the STIV models for a single particle in a sinusoidal potential driven by a linear external force. Panel A shows the total variation distance between each model (same colors as Figure 1) and the true probability density of states when 64 internal variables are used. Panel B depicts the sinusoidal landscape and the quadratic potential associated with the external driving from an optical trap (at $t = 0.79 t_{\text{final}}$). Panels C and D show the approximate probability density of states for each model using 64 internal variables and the true probability density of states (in dark blue) for two different points in time. ($t_{\text{final}} = 10$).

$$d(p, \hat{p}) = \frac{1}{2T} \int_0^T \int_{\mathbb{R}} |p(x, t) - \hat{p}(x, \alpha(t))| dx dt.$$

Although previously we have made use of the KL divergence to interpret the choice of dynamical equations for the internal variables, here we use the total variation distance to measure the error in the approximation as the numerical computation of the KL divergence tended to be poorly behaved. Fortunately, the total variation distance is bounded above by the KL divergence by [38]–[40]

$$d(p, q) \leq \min \left(\sqrt{\frac{1}{2} \mathbb{D}_{\text{KL}}(p \| q)}, \sqrt{1 - \exp(-\mathbb{D}_{\text{KL}}(p \| q))} \right).$$

In panel B of Figure 1, we show the potential energy governing the system, along with a cartoon of the model setup. In the remaining panels C and D, snapshots of the densities at an intermediate time for two different numbers of observables are shown. In these frames, the ground truth probability density of states is shown in dark blue. At this time point, the probability density of states has not fully relaxed from the initial Gaussian distribution, and is still far from equilibrium. All methods converge to the ground truth solution, however, the quasi-equilibrium distribution outperforms the two exponential family models by an order of magnitude for 32 and 64 observables. As is expected, the interaction energy is a good observable and provides important information about the probability density of states, particularly in a problem concerning relaxation to the equilibrium state. Interestingly, the exponential family with gradient flow dynamics outperforms the moment-closure dynamics for small numbers of internal variables. As expected, this gap decreases for 64 internal variables once the approximate density has become highly flexible. Computationally, each of the three models require about the same run time (on a standard laptop computer), and run faster than computing the ground truth pde for few observables (4–32), but become slower for large numbers of internal variables (64) due to the need to solve linear systems involving the covariance matrix $\mathbb{C}_{ij}$.

Figure 2 shows the results for the external driving example. Panel A, again, shows the total variation distance from the ground truth for each model, but now as a function of time over the whole course of the simulation, and for the approximations with 16 and 64 internal variables. After an initial increase, the error of each of the methods is relatively stable in time. This mirrors the observation made in previous implementations of the STIV framework, where the approximation seemed to recover despite passing through regions in time where the approximation was quite poor [18]. Panel B shows the sinusoidal potential in blue and the external driving potential due to an optical trap in pink, along with a cartoon of the model setup. The remaining two panels (C and D) compare the approximate densities with 64 internal variables to the ground truth at two representative points in time. Although not shown here, it is worth noting that when fewer than 64 internal variables are used, including the energy greatly improves the approximation in this example. The external driving is slow enough that the density remains close to equilibrium at all times, and hence the energy is a very good observable.

5 Conclusions

We have shown that given a microscopic particle system governed by overdamped Langevin dynamics it is always possible to construct a thermodynamically consistent macroscopic (ensemble averaged) model with a gradient flow evolution for the internal variables. One strategy, which we've called the operator method, applicable for an arbitrary choice of approximate probability density of states, is presented but not pursued due to computational challenges of a general implementation. Alternative methods, restricted to quasi-equilibrium or exponential family approximations, can be implemented relatively simply. These latter models can either be formulated in terms of averaged observables or their dual internal variables, and can approximate continuous probability distributions with arbitrary accuracy.

In this work, we have shown the flexibility of this framework by considering Fourier basis functions as observables. Such basis functions provide an essentially system agnostic model, as the only way the approximate probability density of states depends on the system of study is through the system size. Since the Fourier basis functions can approximate any continuous function, in principle, one could use them to model systems of arbitrary complexity. However, this becomes impracticable both because the number of observables needed for a reasonable approximation would scale exponentially in the number of degrees of freedom in the system, and because the Fourier functions provide relatively little insight into the underlying structure of the system. Alternatively, if one can capture the salient features of the system using observables especially chosen for the problem, one can expect the STIV framework to produce accurate predictions with relatively little computational cost, as was the case in the previous works [18], [22]. The true utility of internal variable models, therefore, comes from the intelligent choice of a few highly descriptive observables. Thus, a key task for continuing work is that of finding descriptive observables for a given microscopic system (see [41], [42] for emerging efforts in this direction).

Research ethics: Not applicable.

Informed consent: Not applicable.

Author contributions: All authors have accepted responsibility for the entire content of this manuscript and approved its submission. Travis Leadbetter: Conceptualization, Methodology, Software, Formal Analysis, Writing – Original Draft, Writing – Review & Editing. Prashant K. Purohit: Conceptualization, Writing – Review & Editing, Supervision. Celia Reina: Conceptualization, Methodology, Writing – Review & Editing, Supervision.

Use of Large Language Models, AI and Machine Learning Tools: None declared.

Conflict of interest: All authors state no conflict of interest.

Research funding: T. L. acknowledges funding support from the University of Pennsylvania Applied Mathematics and Computational Science graduate group. C. R. gratefully acknowledges support from NSF CAREER Award (CMMI-2047506). P. K. P. acknowledges partial support from NSF Grant DMR 2212162.

Data availability: Source code used to generate the data and figures displayed in this work is freely available via the GitHub repository celiareina/STIV-GradFlow.

Appendix A: The variational method of Eyink with external driving

In the original derivation of the variational method of Eyink [21], when the true solution to the Fokker–Planck equation $\partial_t p(x, t) = \mathcal{L} p(x, t)$ is approximated by a parameterized density of states $\hat{p}(x, \alpha)$, it was assumed that all of the time dependence of $\hat{p}$ is captured by the parameters $\alpha(t)$. However, if we wish to implement a quasi-equilibrium approximation $\hat{p} = \exp(-\beta e + \alpha_i(t)\phi_i(x) - f(\alpha, \beta))$ and the energy depends on a time-dependent external driving protocol $e = e(x, \lambda(t))$, we must modify the variational method since the time dependence assumption has been broken.

Assume for fixed external driving protocol $\lambda(t)$, we approximate the density of states via $\hat{p}(x, \alpha, \lambda)$ (here β is assumed to be fixed), and approximate test function $\hat{\psi}(x, \alpha)$. The variational action is $S(\alpha) = \int_0^\infty \int_X \hat{\psi}(\partial_t - \mathcal{L})\hat{p} dx dt$. Stationary points of the action are given by

$$\begin{aligned} 0 &= \partial_{\alpha_i} S(\alpha) \\ &= \int_0^\infty \int_X \{ \partial_{\alpha_i} \hat{\psi} \partial_t \hat{p} - \partial_{\alpha_i} \hat{p} \partial_t \hat{\psi} \} dx dt - \partial_{\alpha_i} \int_0^\infty \int_X \hat{\psi} \mathcal{L} \hat{p} dx dt \\ &= \int_0^\infty \int_X \left\{ \partial_{\alpha_i} \hat{\psi} \partial_{\alpha_j} \hat{p} \dot{\alpha}_j + \partial_{\alpha_i} \hat{\psi} \partial_\lambda \hat{p} \dot{\lambda} - \partial_{\alpha_i} \hat{p} \partial_{\alpha_j} \hat{\psi} \dot{\alpha}_j \right\} \\ &\quad - \int_0^\infty \partial_{\alpha_i} \langle \mathcal{L}^\dagger \hat{\psi} \rangle_{\hat{p}} dt \\ &= \int_0^\infty \left\{ \langle [\partial \hat{\psi}, \partial \hat{u}]_{ij} \rangle_{\hat{p}} \dot{\alpha}_j + \langle \partial_{\alpha_i} \hat{\psi} \partial_\lambda \hat{u} \rangle_{\hat{p}} \dot{\lambda} - \partial_{\alpha_i} \langle \mathcal{L}^\dagger \hat{\psi} \rangle_{\hat{p}} \right\} dt. \end{aligned}$$

The proper dynamics are recovered by setting the integrand equal to zero.

We now assume the “linear ansatz” for the test function by introducing the paired set of parameters (α, γ) and write $\hat{\psi} = 1 + \gamma_i \partial_{\alpha_i} \hat{u}$. Plugging this in to the dynamical equation above leads to two coupled equations

$$\begin{aligned} \gamma_k \partial_{\alpha_i} \langle \mathcal{L}^\dagger \partial_{\alpha_k} \hat{u} \rangle_{\hat{p}} &= \gamma_k \left\langle \partial_{\alpha_i} \partial_{\alpha_k} \hat{u} \partial_{\alpha_j} \hat{u} - \partial_{\alpha_j} \partial_{\alpha_k} \hat{u} \partial_{\alpha_i} \hat{u} \right\rangle_{\hat{p}} \dot{\alpha}_j \\ &\quad - \left\langle \partial_{\alpha_i} \hat{u} \partial_{\alpha_j} \hat{u} \right\rangle_{\hat{p}} \dot{\gamma}_j + \gamma_k \left\langle \partial_{\alpha_i} \partial_{\alpha_k} \hat{u} \partial_\lambda \hat{u} \right\rangle_{\hat{p}} \dot{\lambda} \\ \langle \mathcal{L}^\dagger \partial_{\alpha_i} \hat{u} \rangle_{\hat{p}} &= \left\langle \partial_{\alpha_i} \hat{u} \partial_{\alpha_j} \hat{u} \right\rangle_{\hat{p}} \dot{\alpha}_j + \langle \partial_{\alpha_i} \hat{u} \partial_\lambda \hat{u} \rangle_{\hat{p}} \dot{\lambda}. \end{aligned}$$

The first equation is trivially solved by $\gamma_i \equiv 0$, whereas the second equation is independent of γ_i and gives the resulting dynamics we need.

Appendix B: Approximate thermodynamic quantities for the quasi-equilibrium approximation of a system with external driving

Since the inclusion of the energy in the approximate probability density of states breaks the assumption that $\hat{p}$ be independent of λ , the form of the approximate thermodynamic quantities must change as a result. Recall that the quasi-equilibrium approximation takes the form $\hat{p} = \exp(-\beta e(x, \lambda) + \alpha_i \phi_i(x) - f(\alpha, \beta, \lambda))$ so that $\partial_{\alpha_i} \hat{u} = \phi_i - \hat{\Phi}_i$, $\hat{u}^{\text{neq}} = e + \hat{u}/\beta = (\alpha_i \phi_i - f)/\beta$ and, following the same steps as in the main text, $\langle \mathcal{L}^\dagger \partial_{\alpha_i} \hat{u} \rangle_{\hat{p}} = -\frac{1}{\eta\beta} \mathbb{M}_{ij} \alpha_j$ where $\mathbb{M}_{ij} = \langle \partial_x \phi_i \partial_x \phi_j \rangle_{\hat{p}}$. Plugging this into the dynamical equations derived from the previous section leads to

$$\mathbb{C}_{ij}\dot{\alpha}_j + \left\langle (\phi_i - \hat{\Phi}_i)\partial_\lambda \hat{u} \right\rangle_{\hat{p}} \dot{\lambda} = -\frac{1}{\eta\beta} \mathbb{M}_{ij}\alpha_j$$

where $\mathbb{C}_{ij} = \left\langle (\phi_i - \hat{\Phi}_i)(\phi_j - \hat{\Phi}_j) \right\rangle_{\hat{p}}$. The approximate external force is defined as $\hat{F}^{\text{ex}} \equiv \langle \partial_\lambda e \rangle_{\hat{p}}$ which is no longer simply $\partial_\lambda \hat{A}^{\text{neq}}$. Now

$$\begin{aligned} \partial_\lambda \hat{A}^{\text{neq}} &= \langle \partial_\lambda e \rangle_{\hat{p}} + \langle (e + \hat{u}/\beta) \partial_\lambda \hat{u} \rangle_{\hat{p}} \\ &= \hat{F}^{\text{ex}} + \frac{1}{\beta} \langle \alpha_i \phi_i \partial_\lambda \hat{u} \rangle_{\hat{p}} \\ &= \hat{F}^{\text{ex}} + \frac{1}{\beta} \partial_\lambda (\alpha_i \hat{\Phi}_i) \end{aligned}$$

which gives the first part of Eq. (11) of the main text. Since $\hat{A}^{\text{neq}} = (\alpha_i \hat{\Phi}_i - f)/\beta$, $\hat{A}^{\text{neq}} - \alpha_i \hat{\Phi}_i/\beta = -f/\beta$ and we recover the second part.

To identify the form of the entropy production, we return to its definition from stochastic thermodynamics and its approximation via STIV. The incremental entropy production is the difference in the incremental work minus the incremental free energy, $Tds^{\text{tot}} = dW - dA^{\text{neq}}$. Approximating these terms at the ensemble average level leads to the required relation $T \frac{d}{dt} \hat{S}^{\text{tot}} = \frac{d}{dt} \hat{W} - \frac{d}{dt} \hat{A}^{\text{neq}}$. The work rate is identified at $\frac{d}{dt} \hat{W} = \hat{F}^{\text{ex}} \dot{\lambda}$ and the rate of change in the non-equilibrium free energy writes $\frac{d}{dt} \hat{A}^{\text{neq}} = \partial_\lambda \hat{A}^{\text{neq}} \dot{\lambda} + \partial_\alpha \hat{A}^{\text{neq}} \dot{\alpha}$. Inserting $\partial_\lambda \hat{A}^{\text{neq}} = \hat{F}^{\text{ex}} + \alpha_i \langle (\phi_i - \hat{\Phi}_i) \partial_\lambda \hat{u} \rangle_{\hat{p}} / \beta$ into the equation for $T \frac{d}{dt} \hat{S}^{\text{tot}}$ leaves (note $\partial_\lambda \hat{u}$ has mean zero due to normalization)

$$T \frac{d}{dt} \hat{S}^{\text{tot}} = -\dot{\alpha} \partial_\alpha \hat{A}^{\text{neq}} - \frac{\alpha_i}{\beta} \langle (\phi_i - \hat{\Phi}_i) \partial_\lambda \hat{u} \rangle_{\hat{p}} \dot{\lambda}.$$

Since $\partial_\alpha \hat{A}^{\text{neq}} = \mathbb{C}_{ij} \alpha_j / \beta$, we can multiply the whole dynamical equation by $-\alpha_i / \beta$ to see

$$T \frac{d}{dt} \hat{S}^{\text{tot}} = \frac{\alpha_i \mathbb{M}_{ij} \alpha_j}{\eta \beta^2}$$

giving Eq. (12) in the main text.

Appendix C: Proof that exponential families and quasi-equilibrium distributions can approximate any continuous probability density on $\mathbb{R}^n$ in total variation distance

Let p be a continuous probability density on $\mathbb{R}^n$ and let $u = \log(p)$. Since p is integrable, for any fixed $0 < \epsilon < 1$, we can find a compact subset $K \subset \mathbb{R}^n$ such that $K^c = \mathbb{R}^n \setminus K$ obeys $\int_{K^c} p(x) dx < \epsilon/2$. Let $K' = K \cap \{p \geq \frac{\epsilon}{2\text{Vol}(K)}\}$. K' is compact since p is continuous, and we know

$$\begin{aligned} \int_{K'^c} p(x) dx &= \int_{K^c} p(x) dx + \int_{K \cap \{p < \epsilon/2\text{Vol}(K)\}} p(x) dx \\ &< \epsilon/2 + \epsilon/2 = \epsilon. \end{aligned}$$

u is bounded on K' , so by the Stone–Weierstrass theorem, we can find a polynomial $q(x)$ such that $\|u - q\|_{L^\infty(K')} = \sup_{x \in K'} |u(x) - q(x)| < \epsilon$. As $e^{q(x)} = e^{u(x)} e^{q(x) - u(x)} \leq e^{u(x)} e^\epsilon$, $Z_q \equiv \int_{K'} e^{q(x)} dx < \infty$ and we can define $\hat{p} = \mathbb{1}_{x \in K'} \frac{e^{q(x)}}{Z_q}$. In fact, we have $e^{-\epsilon}(1 - \epsilon) \leq Z_q \leq e^\epsilon$. Now

$$\begin{aligned}
d_{TV}(p, \hat{p}) &= \frac{1}{2} \int |p(x) - \hat{p}(x)| dx \\
&= \frac{1}{2} \int_{K'} |p(x) - \hat{p}(x)| dx + \frac{1}{2} \int_{K'^c} p(x) dx \\
&\leq \frac{1}{2} \int_{K'} e^u \left| 1 - \frac{e^{q-u}}{Z_q} \right| dx + \frac{\epsilon}{2} \\
&\leq \frac{1}{2} \|1 - \frac{e^{q-u}}{Z_q}\|_{L^\infty(K')} \int_{K'} p dx + \frac{\epsilon}{2} \\
&\leq \frac{1}{2} (e^{2\epsilon} - 1) + \epsilon \rightarrow 0
\end{aligned}$$

as $\epsilon \rightarrow 0$ since $\|1 - \frac{e^{q-u}}{Z_q}\|_{L^\infty(K')} \leq \max\left(\frac{e^{2\epsilon}}{1-\epsilon} - 1, 1 - e^{-2\epsilon}\right) = \frac{e^{2\epsilon}}{1-\epsilon} - 1$.

An analogous proof shows that the same is true for quasi-equilibrium distributions. Assuming that the energy is continuous as well, we simply choose K' as above, q such that $\|u + \beta e - q\|_{L^\infty(K')} < \epsilon$, define $Z_q = \int_{K'} e^{-\beta e + q} dx$. And let $\hat{p} = \mathbb{1}_{x \in K'} \frac{e^{-\beta e + q}}{Z_q}$. The same bounds hold for Z_q so we again have

$$d_{TV}(p, \hat{p}) \leq \frac{1}{2} (e^{2\epsilon} - 1) + \epsilon \rightarrow 0.$$

Appendix D: Details of the numerical implementation of STIV for examples

As discussed in the main text, the energy used to study relaxation to equilibrium comes from a model of the energy landscape of a protein diffusing on a strand of DNA. The original experimental data was published by [36], and the particular choice of model energy landscape was taken from [37]. The parameters used in Eq. (13) are $c_0 = -0.87817$, $c_1 = 0.04736$, $c_2 = 0.60955$, $c_3 = 1.25408$. A bounding box of the form (where x has units of base pairs)

$$\text{Bounding Box}(x) = 100 e^{-x-40} \mathbb{1}_{(x < -25)} + \left(\frac{(x+10)}{30}\right)^{100}$$

was used to constrain the simulated proteins to $x \in [-40, 20]$ (that is, the initial range of the model energy landscape in [37]) and to insure that the minima at $x \approx -30$ was the left most minima. This range is translated to $[0, 60]$ in the video in Appendix E. Since the energy has units of $k_B T$, we used $\beta = 1$. We chose $\eta = 0.83$ which is artificially large to accelerate all simulations.

For the external driving example, we chose parameters $l = 1.0$, $k = 1.6287$, $\beta = 1.0$, and $\eta = 0.8298$. The external protocol was chosen as $\lambda(t) = 0.5t$.

All models (both STIV and the ground truth) were implemented on a discrete grid of 2,000 equally spaced points (between $[-39, 21.1]$ for the relaxation example and $[-2, 8]$ for the external driving example). The ground truth was obtained from discretizing the true pde, $\partial_t p = \mathcal{L} p$, and solving the resulting ode using the Dormand–Prince adaptive Runge–Kutta scheme with variable step size to ensure the relative error remained less than 0.001. For the STIV models, the energy and observables were computed on the grid points $e_i = e(x_i) \in \mathbb{R}^{2000}$ and $\phi_{ij} = \phi_j(x_i) \in \mathbb{R}^{2000 \times d}$ where d is the number of internal variables used. $\hat{p}_i = \hat{p}(x_i, \alpha)$ was then approximated as

$$\hat{p}_i = \frac{1}{\Delta x} \exp(-\beta e_i + \phi_{ij} \alpha_j - f(\alpha, \beta))$$

with

$$f(\alpha, \beta) = \log \left(\sum_{i=1}^{2,000} \exp(-\beta e_i + \phi_{ij} \alpha_j) \right).$$

Any expectations in the dynamical equations were then approximated using the discrete approximate density of states

$$\langle g \rangle_{\hat{p}} \approx \sum_{i=1}^{2000} g(x_i) \hat{p}_i \Delta x.$$

The exponential family models which did not include the energy were obtained the same way but with β fixed to zero.

Some numerical difficulties arose while inverting the covariance matrix C_{ij} . Thus, C_{ij}^{-1} was further approximated using the pseudo-inverse C_{ij}^{+} truncating singular values (of C_{ij}) falling below 10^{-8} . Various truncation levels (10^{-2} through 10^{-15}) were tested to ensure this choice did not impact numerical accuracy.

Appendix E: Full video of examples

The video provided in the supplementary materials shows a comparison between the true probability density of states in dark blue and the approximate densities using the quasi-equilibrium method (green), the exponential family approximation with dynamics given by the variational method (equivalent to moment closure, yellow), and the exponential family approximation with dynamics given by a gradient flow (salmon). The first part of the video shows, sequentially, the approximation for the model of proteins diffusing on DNA using 4, 8, 16, 32, and 64 internal variables (the legend $|\phi| = N$ in each part shows the number of internal variables/observables). The second half of the video shows the approximation for the externally driven particle in a periodic landscape for 4, 8, 16, 32, and 64 internal variables.

References

- [1] D. Jou, J. Casas-Vázquez, and G. Lebon, “Extended irreversible thermodynamics,” in *Extended Irreversible Thermodynamics*, Springer, 1996, pp. 41–74.
- [2] G. Lebon, D. Jou, and J. Casas-Vázquez, *Understanding Non-Equilibrium Thermodynamics*, vol. 295, Berlin, Heidelberg, Springer, 2008.
- [3] B. Coleman, *Thermodynamics of Materials with Memory*, 1st ed. Vienna, Springer, 1971.
- [4] C. Truesdell, “Historical introit the origins of rational thermodynamics,” in *Rational Thermodynamics*, Springer, 1984, pp. 1–48.
- [5] G. A. Maugin and W. Muschik, “Thermodynamics with internal variables. Part i. General concepts,” *J. Non-Equilib. Thermodyn.*, vol. 19, no. 1, pp. 217–249, 1994.
- [6] M. Grmela and H. C. Öttinger, “Dynamics and thermodynamics of complex fluids. i. Development of a general formalism,” *Phys. Rev. E*, vol. 56, no. 6, pp. 6620–6632, 1997.
- [7] A. Mielke, “Formulation of thermoelastic dissipative material behavior using generic,” *Continuum Mech. Thermodyn.*, vol. 23, no. 3, pp. 233–256, 2011.
- [8] M. Pavelka, V. Klika, and M. Grmela, *Multiscale Thermo-Dynamics: Introduction to GENERIC*, Berlin/Boston, Walter de Gruyter GmbH & Co KG, 2018.
- [9] F. Garofalo, “An empirical relation defining the stress dependence to minimum creep rate in metals,” *Trans. Metall. Soc. AIME*, vol. 227, no. 1, p. 351, 1963.
- [10] J.-L. Chaboche, “Continuous damage mechanics — a tool to describe phenomena before crack initiation,” *Nucl. Eng. Des.*, vol. 64, no. 2, pp. 233–247, 1981.
- [11] J. Lemaitre, “A continuous damage mechanics model for ductile fracture,” *J. Eng. Mater. Technol.*, vol. 107, no. 1, pp. 83–89, 1985.
- [12] Z. P. Bazant, “Mechanics of distributed cracking,” *Appl. Mech. Rev.*, vol. 39, no. 5, pp. 675–705, 1986.
- [13] H. C. Öttinger, *Beyond Equilibrium Thermodynamics*, vol. 5, Hoboken, Wiley-Interscience, 2005.
- [14] R. Zwanzig, *Nonequilibrium Statistical Mechanics*, New York, Oxford University Press, 2001.
- [15] H. C. Öttinger, “General projection operator formalism for the dynamics and thermodynamics of complex fluids,” *Phys. Rev. E*, vol. 57, no. 2, pp. 1416–1420, 1998.
- [16] A. Montefusco, M. A. Peletier, and H. C. Öttinger, “A framework of nonequilibrium statistical mechanics. ii. Coarse-graining,” *J. Non-Equilib. Thermodyn.*, vol. 46, no. 1, pp. 15–33, 2021.
- [17] A. Mielke, M. A. Peletier, and J. Zimmer, “Deriving a generic system from a hamiltonian system,” *Arch. Ration. Mech. Anal.*, vol. 249, no. 1, 2025, Art. no. 62.

- [18] T. Leadbetter, P. K. Purohit, and C. Reina, “A statistical mechanics framework for constructing nonequilibrium thermodynamic models,” *PNAS Nexus*, vol. 2, no. 12, 2023, Art. no. pgad417.
- [19] U. Seifert, “Stochastic thermodynamics, fluctuation theorems and molecular machines,” *Rep. Prog. Phys.*, vol. 75, no. 12, 2012, Art. no. 126001.
- [20] L. Peliti and S. Pigolotti, *Stochastic Thermodynamics: An Introduction*, Princeton, Princeton University Press, 2021.
- [21] G. L. Eyink, “Action principle in nonequilibrium statistical dynamics,” *Phys. Rev. E*, vol. 54, no. 4, pp. 3419–3435, 1996.
- [22] T. Leadbetter, P. K. Purohit, and C. Reina, “A statistical mechanics derivation and implementation of non-conservative phase field models for front propagation in elastic media,” *J. Mech. Phys. Solids*, vol. 204, no. 1, p. 2025, <https://doi.org/10.1016/j.jmps.2025.106240>.
- [23] R. Abeyaratne, S. Dharmaravan, G. Saccomandi, and G. Tomassetti, “Using stochastic thermodynamics with internal variables to capture orientational spreading in cell populations undergoing cyclic stretch,” *arXiv preprint arXiv:2507.15694*, 2025.
- [24] R. Jordan, D. Kinderlehrer, and F. Otto, “The variational formulation of the Fokker–Planck equation,” *SIAM J. Math. Anal.*, vol. 29, no. 1, pp. 1–17, 1998.
- [25] B. Coleman and W. Noll, “The thermodynamics of elastic materials with heat conduction and viscosity,” *Arch. Ration. Mech. Anal.*, vol. 13, no. 1, pp. 167–178, 1963.
- [26] E. Hairer, M. Hochbruck, A. Iserles, and C. Lubich, “Geometric numerical integration,” *Oberwolfach Rep.*, vol. 3, no. 1, pp. 805–882, 2006.
- [27] B. Gladman, M. Duncan, and J. Candy, “Symplectic integrators for long-term integrations in celestial mechanics,” *Celest. Mech. Dyn. Astron.*, vol. 52, no. 3, pp. 221–240, 1991.
- [28] H. Yoshida, “Recent progress in the theory and application of symplectic integrators,” *Celest. Mech. Dyn. Astron.*, vol. 56, nos. 1–2, pp. 27–43, 1993.
- [29] B. Leimkuhler and S. Reich, *Simulating Hamiltonian Dynamics, Number 14*, Cambridge, Cambridge University Press, 2004.
- [30] M. Kobilarov, K. Crane, and M. Desbrun, “Lie group integrators for animation and control of vehicles,” *ACM Trans. Graph. (TOG)*, vol. 28, no. 2, pp. 1–14, 2009.
- [31] H. C. Öttinger, “GENERIC integrators: structure preserving time integration for thermodynamic systems,” *J. Non-Equilib. Thermodyn.*, vol. 43, no. 2, pp. 89–100, 2018.
- [32] L. Bronstein and H. Koepl, “A variational approach to moment-closure approximations for the kinetics of biomolecular reaction networks,” *J. Chem. Phys.*, vol. 148, no. 1, 2018, Art. no. 014105.
- [33] B. Turkington, “An optimization principle for deriving nonequilibrium statistical models of hamiltonian dynamics,” *J. Stat. Phys.*, vol. 152, no. 1, pp. 569–597, 2013.
- [34] H. Risken, *The Fokker–Planck Equation, Volume 18 of Springer Series in Synergetics*, Berlin Heidelberg, Springer-Verlag, 1996.
- [35] E. L. Lehmann and G. Casella, *Theory of Point Estimation*, New York, Springer Science & Business Media, 2006.
- [36] S. Kim et al., “Probing allostery through dna,” *Science*, vol. 339, no. 6121, pp. 816–819, 2013.
- [37] J. Singh and P. K. Purohit, “Elasticity as the basis of allostery in dna,” *J. Phys. Chem. B*, vol. 123, no. 1, pp. 21–28, 2018.
- [38] M. S. Pinsker, *Information and Information Stability of Random Variables and Processes*, San Francisco, Holden Day, 1964.
- [39] J. Bretagnolle and C. Huber, “Estimation des densités: risque minimax,” *Z. Wahrscheinlichkeitstheor. Verw. Geb.*, vol. 47, no. 1, pp. 119–137, 1979.
- [40] C. L. Canonne, “A short note on an inequality between kl and tv ,” *arXiv preprint arXiv:2202.07198*, 2022.
- [41] B. Liu, E. Ocegueda, M. Trautner, A. M. Stuart, and K. Bhattacharya, “Learning macroscopic internal variables and history dependence from microscopic models,” *J. Mech. Phys. Solids*, vol. 178, no. 1, 2023, Art. no. 105329.
- [42] W. Qiu, S. Huang, and C. Reina, “Bridging statistical mechanics and thermodynamics away from equilibrium: a data-driven approach for learning internal variables and their dynamics,” *J. Mech. Phys. Solids*, vol. 203, 2025, Art. no. 106211.