

Article

Quantifying configuration-sampling error in Langevin simulations of complex molecular systems

Josh Fass ^{1,6}, David A. Sivak ², Gavin E. Crooks ³, Kyle A. Beauchamp ⁴, Benedict Leimkuhler ⁵, John D. Chodera ^{6*}

¹ Tri-Institutional PhD Program in Computational Biology & Medicine, New York, NY 10065; josh.fass@choderlab.org

² Department of Physics, Simon Fraser University, Burnaby, BC V5A 1S6, Canada; dsivak@sfsu.ca, <http://www.sfu.ca/physics/sivakgroup.html>

³ Rigetti Computing, Berkeley, CA 94710; gec@threeplusone.com

⁴ Counsyl, South San Francisco, CA 94080; kyleabeauchamp@gmail.com

⁵ School of Mathematics and Maxwell Institute of Mathematical Sciences, James Clerk Maxwell Building, Kings Buildings, University of Edinburgh, Edinburgh, EH9 3FD, UK; B.Leimkuhler@ed.ac.uk, <http://kac.maths.ed.ac.uk/~bl>

⁶ Computational and Systems Biology Program, Sloan Kettering Institute, Memorial Sloan Kettering Cancer Center, New York, NY 10065; john.chodera@choderlab.org, <http://choderlab.org>

* Correspondence: john.chodera@choderlab.org

Version July 22, 2018 submitted to Entropy

1 **Abstract:** While Langevin integrators are popular in the study of equilibrium properties of complex
2 systems, it is challenging to estimate the timestep-induced discretization error: the degree to which
3 the sampled phase-space or configuration-space probability density departs from the desired target
4 density due to the use of a finite integration timestep. In [1], Sivak *et al.* introduced a convenient
5 approach to approximating a natural measure of error between the sampled density and the target
6 equilibrium density, the KL divergence, in *phase space*, but did not specifically address the issue of
7 *configuration-space properties*, which are much more commonly of interest in molecular simulations.
8 Here, we introduce a variant of this near-equilibrium estimator capable of measuring the error in
9 the configuration-space marginal density, validating it against a complex but exact nested Monte
10 Carlo estimator to show that it reproduces the KL divergence with high fidelity. To illustrate its utility,
11 we employ this new near-equilibrium estimator to assess a claim that a recently proposed Langevin
12 integrator introduces extremely small configuration-space density errors up to the stability limit at
13 no extra computational expense. Finally, we show how this approach to quantifying sampling bias
14 can be applied to a wide variety of stochastic integrators by following a straightforward procedure to
15 compute the appropriate shadow work, and describe how it can be extended to quantify the error in
16 arbitrary marginal or conditional distributions of interest.

17 **Keywords:** Langevin dynamics; Langevin integrators; KL divergence; nonequilibrium free energy;
18 molecular dynamics integrators; integrator error; sampling error; BAOAB; VRORV; OBABO; OVRVO;
19 VVVR; Bussi-Parrinello; shadow work; integrator error

20 **Contents**

21 1 Introduction	2
22 2 Numerical discretization methods and timestep-dependent bias	4

23	3 Estimators for KL divergence and the configurational KL divergence	7
24	3.1 Near-equilibrium estimators for KL divergence	7
25	3.2 A simple modification to the near-equilibrium estimator can compute KL divergence in configuration space	10
27	3.3 Comparison of phase-space error for different integrators	11
28	3.4 Comparison of configurational KL divergence for different integrators	11
29	3.5 Influence of the collision rate	13
30	3.6 Comparison with reference methods validates the near-equilibrium estimate	13
31	3.6.1 Practical lower bound from nested Monte Carlo	15
32	3.6.2 Practical upper bound from Jensen's inequality	16
33	3.6.3 Sandwiching the KL divergence to validate the near-equilibrium estimate	17
34	4 Relation to GHMC acceptance rates	17
35	5 Discussion	20
36	5.1 Future directions	20
37	6 Detailed methods	21
38	6.1 One-dimensional model system: Double well	21
39	6.2 Model molecular mechanics system: A harmonically restrained water cluster	21
40	6.3 Caching equilibrium samples	22
41	6.4 Computing shadow work for symmetric Strang splittings	23
42	6.5 Computation of shadow work for OVRVO	24
43	6.6 Variance-controlled adaptive estimator for KL divergence	25
44	A Statistics of shadow work distributions	26
45	B Log-scale plots	28
46	C Further comments on the collision rate	28
47	References	30

48 1. Introduction

49 Langevin dynamics [2] is a system of stochastic differential equations which describes the behavior
 50 of condensed phase systems subject to random weak collisions with fictitious bath particles at thermal
 51 equilibrium. In this article we are concerned with the efficient numerical simulation of the Langevin
 52 dynamics system. The equations governing the i th atom of an N -body Langevin system are

$$\dot{\mathbf{x}}_i = \mathbf{v}_i, \quad (1)$$

$$\dot{\mathbf{v}}_i = -m_i^{-1} \nabla_{\mathbf{x}_i} U(\mathbf{x}) - \gamma \mathbf{v}_i + \sigma m_i^{-1/2} \dot{\mathbf{W}}_i. \quad (2)$$

53 Here, \mathbf{x}_i and \mathbf{v}_i denote the position vector and velocity vector of the i th particle of the system (typically
 54 each is a vector in \mathbb{R}^3), m_i is the particle mass, and $U(\mathbf{x})$ is the total potential energy, assumed to be
 55 a function of all the coordinates, $\mathbf{x} = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N)$. The constant $\sigma^2 \equiv 2k_B T \gamma$ quantifies the rate
 56 of heat exchange with the bath, where $k_B T$ denotes the thermal energy per degree of freedom, γ is
 57 the collision rate (with dimensions of inverse time), and $\dot{\mathbf{W}}_i(t)$ is a standard 3-dimensional Wiener
 58 process [3,4]. It is easily shown that Langevin dynamics preserves the canonical distribution with
 59 stationary density

$$\pi(\mathbf{x}, \mathbf{v}) \propto e^{-\beta E(\mathbf{x}, \mathbf{v})} \propto e^{-\beta U(\mathbf{x})} e^{-\beta K(\mathbf{v})}, \quad (3)$$

60 where $\beta \equiv (k_B T)^{-1}$ is the inverse temperature, $E(\mathbf{x}, \mathbf{v})$ is a separable energy function $E(\mathbf{x}, \mathbf{v}) =$
 61 $U(\mathbf{x}) + K(\mathbf{v})$, and $K(\mathbf{v}) = \mathbf{v}^T M \mathbf{v} / 2$ is the kinetic energy in which M is a diagonal matrix constructed
 62 from the masses. We assume that the system (1)-(2) is ergodic, meaning that the canonical density is
 63 the unique stationary distribution; almost all stochastic paths consistent with (1)-(2) will sample the
 64 canonical distribution with density (3).

65 On a computer, approximating the solution of equations (1)-(2) requires discretizing these
 66 equations in time to produce a *finite-timestep Langevin integrator* which can be iterated to compute
 67 equilibrium or dynamical properties [5]. A wide variety of schemes have been proposed for this
 68 discretization [6–15]. For compact presentation of integration methods, we recast the equations (1)-(2)
 69 in vectorial form:

$$\begin{aligned}\dot{\mathbf{x}} &= \mathbf{v}, \\ \dot{\mathbf{v}} &= M^{-1} \nabla U(\mathbf{x}) - \gamma \mathbf{v} + \sigma M^{-1/2} \mathbf{W}.\end{aligned}\quad (4)$$

70 Popular numerical methods for integrating (4) can then be viewed as defining Markov chains on
 71 the phase space (\mathbf{x}, \mathbf{v}) , with $(\mathbf{x}_{k+1}, \mathbf{v}_{k+1})$ defined in relation to $(\mathbf{x}_k, \mathbf{v}_k)$ where the subindex k from this
 72 point forward should be taken to be a timestep index. Assuming ergodicity, these schemes provide a
 73 practical means of sampling the equilibrium distribution.

74 Until now, assessing whether specific integrators sample the true equilibrium density with greater
 75 fidelity than others in specific settings has relied on computing low-dimensional marginal distributions
 76 of specific observables perceived to be sensitive to configuration-space sampling errors [16], such as
 77 radial distribution functions, marginal distributions of internal coordinates [17], or the configurational
 78 temperature [14,18,19]. While it is clear that some observables are more sensitive to errors in
 79 configuration-space density than others (Figure 3), and the error in the observables of interest is
 80 paramount for a particular application, this highlights the risk of using the error in a single physical
 81 property as a surrogate for judging integrator quality, as the error in other observables of interest may
 82 be large despite small error in the test observable.

83 To evaluate numerical Langevin integrators, there would be great utility in a *computable, universal*
 84 *measure* of the bias they introduce in specific concrete settings, such that low error in this measure
 85 ensures low error in all observables of interest. There is currently no computable measure of the total
 86 configuration-sampling bias for concrete choices of integrator parameters and target system.

87 Controlling the magnitude of the integrator-induced sampling bias is crucial when computing
 88 quantitative predictions from simulation. Following [20], we assume that the numerical method
 89 samples an associated probability density $\rho(\mathbf{x}, \mathbf{v})$ which differs from the exact canonical probability
 90 density $\pi(\mathbf{x}, \mathbf{v})$ of (3). Because ρ does not have a closed-form, easily computable expression¹, it is
 91 difficult to quantify the error introduced by a given choice of timestep or integrator. We will show
 92 how, for a particularly useful measure of error, we can circumvent this problem and develop a simple,
 93 effective approach to measuring error in complex molecular systems. Our focus in the sequel is on
 94 the error in configurational averages. We therefore introduce the notation ρ_x and π_x to indicate the
 95 position-dependent configurational marginal densities of ρ and π .

96 KL divergence as a natural measure of sampling bias

An ideal measure of the discrepancy between the sampled distribution ρ_x and the equilibrium
 distribution π_x should be “universal” in the sense that driving that measure to zero implies that error

¹ The concept of a *shadow Hamiltonian* has been used to embed this density in a canonical density context, but the shadow Hamiltonian cannot be directly computed, though some approaches to approximate it via expansion (generally requiring higher-order derivatives than gradients) have been proposed—see [21], Chapter 3 of [4], and references therein.

in any expectation also goes to zero. It should also be defined for all densities, and not rely on a system-specific choice of observables. One such measure is the Kullback-Leibler (KL) divergence,

$$\mathcal{D}_{\text{KL}}(p\|q) \equiv \int d\mathbf{x} p(\mathbf{x}) \ln \left(\frac{p(\mathbf{x})}{q(\mathbf{x})} \right). \quad (5)$$

97 The KL divergence is defined and non-negative for any pair of distributions on the same support, and
 98 $\mathcal{D}_{\text{KL}}(p\|q) = 0$ if and only if $p = q$ almost everywhere.

99 In [1], Sivak and colleagues demonstrated how to approximate the KL divergence of the sampled
 100 distribution ρ from the target distribution π over the full phase-space distribution, $\mathcal{D}_{\text{KL}}(\rho\|\pi)$, in terms
 101 of a work-like quantity—the *shadow work*—that is readily computable for a large family of Langevin
 102 integrators (Figure 1a). This estimator depends only on the ability to draw samples from π and to
 103 measure a suitable work-like quantity. This method was applied in [1] to measure the phase-space
 104 sampling bias introduced by a particular Langevin integrator (OVRVO) on periodic boxes of TIP3P
 105 water [22].

106 Since the velocity marginal is not generally of interest, and since some integrators are thought to
 107 preserve the configuration marginal of the target distribution with higher fidelity than the phase-space
 108 joint distribution, we sought to generalize the technique to estimate the KL divergence of the sampled
 109 configuration-space marginal, $\mathcal{D}_{\text{KL}}(\rho_x\|\pi_x)$. Below, we show how a simple modification of the estimator
 110 described in [1] can achieve this goal, and illustrate how this provides a useful tool for measuring the
 111 integrator-induced error in configuration-space densities for real molecular systems.

112 2. Numerical discretization methods and timestep-dependent bias

There are several possible ways to discretize Langevin dynamics. A flexible approach to this task is via operator splitting, where the Langevin system is split into components, for example,

$$\begin{bmatrix} \dot{\mathbf{x}} \\ \dot{\mathbf{v}} \end{bmatrix} = \underbrace{\begin{bmatrix} \mathbf{v} \\ 0 \end{bmatrix}}_{\text{R}} + \underbrace{\begin{bmatrix} 0 \\ -M^{-1}\nabla U(\mathbf{x}) \end{bmatrix}}_{\text{V}} + \underbrace{\begin{bmatrix} 0 \\ -\gamma \mathbf{v} + \sqrt{2\gamma}(\beta M)^{-1/2}\mathbf{W} \end{bmatrix}}_{\text{O}} \quad (6)$$

113 where each component can be solved “exactly” (in the sense of distributions) for a small time increment.
 114 The label O indicates that the corresponding part of the splitting has the form of an Ornstein-Uhlenbeck
 115 process. The labels of the deterministic parts, R and V, have been chosen to reflect the deterministic
 116 updates of position and velocity, respectively; this notation has been used in some previous articles [23,
 117 24]. Note that in the articles of Leimkuhler and Matthews [4,14], the Langevin equations are cast
 118 in position-momentum form instead of position-velocity form and the labels A, B are then used to
 119 indicate the deterministic updates of positions and momenta, respectively. The choices of components
 120 of the splitting we use here are not the only options. For example the Stochastic Position Verlet method
 121 [25] groups together both of the contributions V and O. Other splitting-based methods for Langevin
 122 dynamics are discussed for example in [26–28].

123 Once a splitting is defined, the propagator $e^{\mathcal{L}\Delta t}$ can be approximated as a Trotter factorization, i.e.
 124 a product of the propagators corresponding to the individual components, defining thus a family of
 125 numerical methods indexed by the string indicating the order of appearance of these individual terms.
 126 For example, we would use OVRVO to refer to the composition method,

$$e^{\mathcal{L}\Delta t} \approx e^{\mathcal{L}_{\text{OVRVO}}\Delta t} = e^{\mathcal{L}_O \frac{\Delta t}{2}} e^{\mathcal{L}_V \frac{\Delta t}{2}} e^{\mathcal{L}_R \Delta t} e^{\mathcal{L}_V \frac{\Delta t}{2}} e^{\mathcal{L}_O \frac{\Delta t}{2}}. \quad (7)$$

127 Due to the lack of commutativity of the operators, equality between the true propagator and the
 128 Strang splitting is only achieved in the limit $\Delta t \rightarrow 0$, i.e., for vanishing timestep. However, in the case
 129 of splitting methods as defined above, it is possible to analyze the error in the effective probability
 130 distribution sampled by the finite timestep method [29].

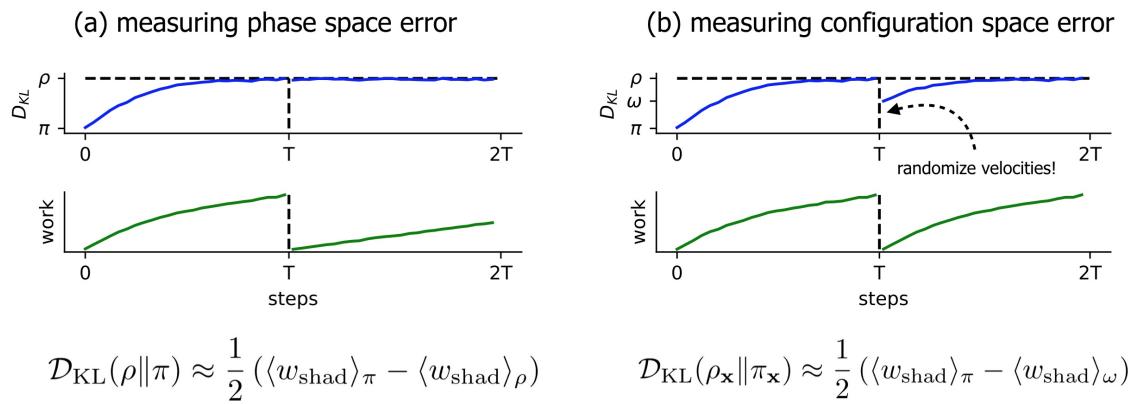


Figure 1. A simple nonequilibrium protocol allows measurement of the KL divergence in phase and configuration space close to equilibrium. Simple nonequilibrium protocols can be used in complex molecular systems to rapidly estimate—utilizing the Crooks fluctuation theorem—the KL divergence of sampled Langevin densities from equilibrium. In both panels, the x -axis is the number of steps taken so far in the length- $2T$ protocol, and $\langle w_{\text{shad}} \rangle_{\pi}$ indicates the average (reduced, unitless) shadow work accumulated over T steps of Langevin dynamics, initialized from equilibrium $((\mathbf{x}_0, \mathbf{v}_0) \sim \pi)$. **(a)** The original scheme described in Sivak *et al.* [1] to measure the KL divergence between the sampled phase-space density ρ and the equilibrium phase-space density π . $\langle w_{\text{shad}} \rangle_{\rho}$ is the average shadow work accumulated over T steps of Langevin dynamics, initialized from the integrator’s steady state $((\mathbf{x}_0, \mathbf{v}_0) \sim \rho)$. **(b)** The modified scheme introduced here to measure the KL divergence in the *configuration-space* marginal density between the marginal sampled configuration-space density ρ_x and marginal equilibrium density π_x . $\langle w_{\text{shad}} \rangle_{\omega}$ is the average shadow work accumulated over T steps of Langevin dynamics, where the initial configuration is drawn from the integrator’s steady state, and the initial velocities are drawn from equilibrium $(\mathbf{x}_0 \sim \rho_x, \mathbf{v}_0 \sim \pi(\mathbf{v}|\mathbf{x}_0))$. We denote this distribution $\omega(\mathbf{x}, \mathbf{v}) \equiv \rho_x(\mathbf{x})\pi(\mathbf{v}|\mathbf{x})$. The **top row** schematically illustrates “distance from equilibrium”, with y -axis ticks for $\mathcal{D}_{\text{KL}}(\pi\|\pi) = 0$, $\mathcal{D}_{\text{KL}}(\omega\|\pi) \leq \mathcal{D}_{\text{KL}}(\rho\|\pi)$. The **bottom row** illustrates the average work (here, just shadow work) accumulated throughout each protocol.

131 The advantage of the splitting approach is that each component propagator, $e^{\mathcal{L}_R\tau}$, $e^{\mathcal{L}_V\tau}$, and $e^{\mathcal{L}_O\tau}$,
 132 has a straightforward interpretation in terms of arithmetic operations:

$$R : e^{\mathcal{L}_R\tau} : \begin{bmatrix} \Delta x \\ \Delta v \end{bmatrix} = \begin{bmatrix} v \\ 0 \end{bmatrix} \tau \quad (8)$$

$$V : e^{\mathcal{L}_V\tau} : \begin{bmatrix} \Delta x \\ \Delta v \end{bmatrix} = \begin{bmatrix} 0 \\ -M^{-1}\nabla U(x) \end{bmatrix} \tau \quad (9)$$

$$O : e^{\mathcal{L}_O\tau} : \begin{bmatrix} \Delta x \\ \Delta v \end{bmatrix} = \begin{bmatrix} 0 \\ (a(\tau) - 1)v + \sqrt{1 - a(\tau)^2} (\beta M)^{-1/2} \xi \end{bmatrix} \quad (10)$$

133 where $a(\tau) = e^{-\gamma\tau}$ and $\xi \sim \mathcal{N}(0, 1)^{3N}$ is a vector of standard normal random variables drawn for each
 134 degree of freedom in each O step.

135 By chaining the operations in the order specified by the splitting string, we can unroll them into
 136 the sequence of mathematical updates needed to implement one cycle of the integrator for a total time
 137 Δt . For VRORV, for example, translating the splitting string into the appropriate sequence of update
 138 equations in (8)–(10) produces the following equations for one complete integrator timestep:²

$$\begin{aligned} \mathbf{v}_{k+1/4} &= \mathbf{v}_k - \frac{\Delta t}{2} M^{-1} \nabla U(\mathbf{x}_k) \\ \mathbf{x}_{k+1/2} &= \mathbf{x}_k + \frac{\Delta t}{2} \mathbf{v}_{k+1/4} \\ \mathbf{v}_{k+3/4} &= a_1 \mathbf{v}_{k+1/4} + \sqrt{1 - a_1^2} (\beta M)^{-1/2} \xi_{k+1/2} \\ \mathbf{x}_{k+1} &= \mathbf{x}_{k+1/2} + \frac{\Delta t}{2} \mathbf{v}_{k+3/4} \\ \mathbf{v}_{k+1} &= \mathbf{v}_{k+3/4} - \frac{\Delta t}{2} M^{-1} \nabla U(\mathbf{x}_{k+1}) \end{aligned} \quad (11)$$

139 Due to the different naming convention they adopt, this is referred to as the “BAOAB” method in the
 140 work of Leimkuhler and Matthews [4,14].

141 As another example, the Langevin integrator of Bussi and Parrinello [12] corresponds to the
 142 splitting OVRVO; this splitting is also known as velocity Verlet with velocity randomization [24] due to
 143 its use of a velocity Verlet integrator core (substeps VRV) [30].

144 While both the VRORV and OVRVO discrete time integration schemes reduce to the same stochastic
 145 differential equations in the limit that $\Delta t \rightarrow 0$, they can behave quite differently for finite timesteps
 146 ($\Delta t > 0$), especially for timesteps of practical interest for atomistic molecular simulation.

147 *Langevin integrators introduce sampling bias that grows with the size of the timestep*

148 In many types of molecular simulations, only configuration-space properties are of interest. The
 149 configurational canonical density is defined by marginalization, viz,

$$\pi_x(\mathbf{x}) = \int d\mathbf{v} \pi(\mathbf{x}, \mathbf{v}) \propto e^{-U(\mathbf{x})}. \quad (12)$$

150 In practice, the velocities are simply discarded in computations while the positions are used to estimate
 151 configuration-dependent properties.

152 The continuous-time Langevin equations of motion (4) possess the target equilibrium density π (3)
 153 as their unique stationary density, suggesting that, at least in principle, temporal averages along

² The subscripts 1/4, 1/2, etc., have no relation to intermediate physical times, and are used solely to denote the sequence of intermediate computational steps.

154 Langevin trajectories can be used to approximate averages with respect to the equilibrium distribution.
 155 However, numerical simulations with a finite timestep $\Delta t > 0$ will generally sample a different
 156 distribution, which we will denote by $\rho(\mathbf{x}, \mathbf{v})$, which implicitly depends on timestep Δt . The
 157 discrepancy between the distributions ρ and π will grow with the size of Δt according to some
 158 power law (e.g., $\mathcal{O}(\Delta t^2)$ or $\mathcal{O}(\Delta t^4)$). For sufficiently large stepsize (typically inversely proportional
 159 to the fastest oscillatory mode present) the discretization will become unstable, but the value of the
 160 stepsize for which the bias is unacceptable may occur well below the molecular dynamics stability
 161 threshold [14].

162 Note that this phenomenon is completely separate from numerical issues introduced by
 163 finite-precision arithmetic on digital computers, which introduces roundoff error in mathematical
 164 operations; here, we presume that computations can be carried out to arbitrary precision, and analyze
 165 only the effects of time discretization. The timestep-dependent bias is also unrelated to the Monte
 166 Carlo or sampling error which is due to finite approximation of the long-term average of a given
 167 quantity.

168 In Figure 2, we illustrate a few key behaviors of this stepsize-dependent sampling bias in a
 169 simple quartic 1D system. Note that: (1) the numerically sampled distribution deviates from the target
 170 distribution, (2) this deviation increases with timestep Δt , and (3) the deviation in phase space (\mathbf{x}, \mathbf{v})
 171 may be different than the deviation in configuration space only (\mathbf{x}) .

172 It has been proposed that some integrators of Langevin dynamics (particularly the VRORV
 173 aka “BAOAB” integrator of Leimkuhler and Matthews) preserve the configuration distribution with
 174 significantly greater fidelity than other equal-cost integration algorithms [14,31,32], a property that
 175 could have significant ramifications for the efficiency of sampling given fixed computational budgets.
 176 However, as the formal arguments for this “superconvergence” property rely on a high-friction
 177 limit, it is unclear how large the friction coefficient needs to be in practice for the argument to
 178 apply. Formal descriptions of the error are typically generic, in that they do not provide guidance
 179 on precisely which Δt introduces a tolerable amount of bias for a particular system, and they do not
 180 provide a way to predict how other choices, such as mass matrix modifications (e.g., hydrogen mass
 181 repartitioning) [17,33–35], will affect the error for a system of interest.

182 3. Estimators for KL divergence and the configurational KL divergence

183 3.1. Near-equilibrium estimators for KL divergence

By using work averages, Sivak and Crooks [36] derived a near-equilibrium estimator for the KL divergence between an arbitrary distribution ρ and the equilibrium distribution π :

$$\mathcal{D}_{\text{KL}}(\rho \parallel \pi) \approx \frac{1}{2} \left(\langle w \rangle_{\pi; \tilde{\Lambda}} - \langle w \rangle_{\pi, \Lambda; \tilde{\Lambda}} \right) \quad (13)$$

$$= \frac{1}{2} \left(\langle w \rangle_{\pi; \tilde{\Lambda}} - \langle w \rangle_{\rho; \tilde{\Lambda}} \right). \quad (14)$$

184 Here $\langle \dots \rangle_{p; \tilde{\Lambda}}$ indicates an average over the dynamical ensemble produced by initialization in
 185 microstates sampled from density p and subsequent driving by protocol $\tilde{\Lambda}$ that is the time-reversal
 186 of protocol Λ . The expectation $\langle w \rangle_{\pi, \Lambda; \tilde{\Lambda}} = \langle w \rangle_{\rho; \tilde{\Lambda}}$ represents a procedure in which initial microstates
 187 are sampled from an initial density $\rho = (\pi, \Lambda)$ prepared by sampling π and applying protocol
 188 Λ ; the expectation is subsequently computed by averaging the work during the application of the
 189 time-reversed protocol $\tilde{\Lambda}$ over many realizations of this sampling process. Note that this is distinct
 190 from $\langle w \rangle_{\pi; \tilde{\Lambda}}$, the expectation where the initial sample is selected from π and the work is measured
 191 during the execution of time-reversed protocol $\tilde{\Lambda}$.

Sivak, Chodera, and Crooks [1] demonstrated how to apply this estimator when ρ is the biased (nonequilibrium) stationary distribution resulting from protocol Λ , the repeated application of a particular numerical Langevin integrator for sufficiently long to reach steady state. In particular,

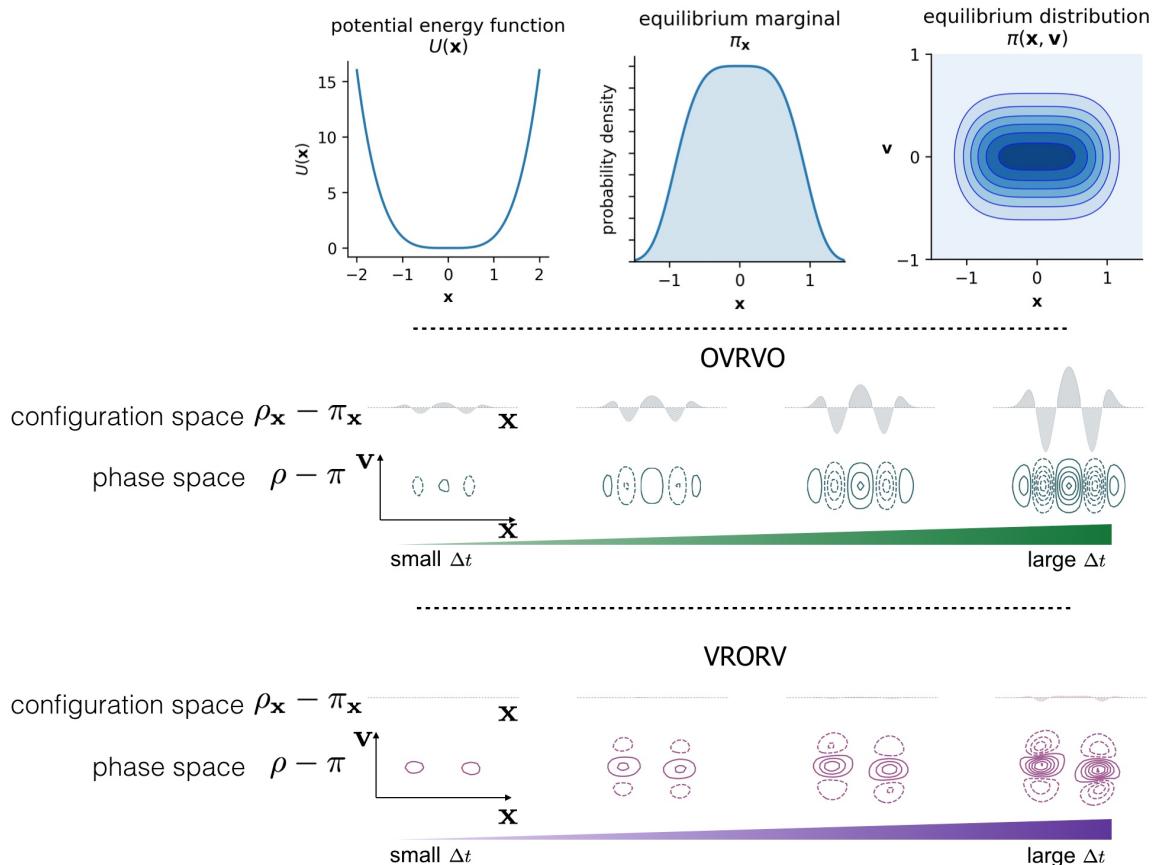


Figure 2. Comparison of Langevin integrators in terms of phase-space and marginal distributions. For a simple 1D system with the quartic potential $U(x) = x^4$, the error in sampled phase-space density ρ and its marginal density ρ_x grows as a function of timestep Δt . However, different Langevin integrators (OVRVO and VRORV shown here) derived from symmetric Strang splittings can lead to drastically different error structures in phase space, which can induce fortuitous cancellation of error in the marginal distribution under certain circumstances (VRORV), see [20]. In the **top row**, we illustrate the definition of the 1D system (**left**: the potential energy function, $U(x) = x^4$; **middle**: the equilibrium marginal density over configuration space, $\pi_x(x) \propto e^{-\beta U(x)}$; **right**: the equilibrium joint distribution over phase space $\pi(x, v)$). In the **middle row**, we illustrate the *increasing discrepancy between the sampled distribution ρ and the equilibrium distribution π* , for both the full phase-space and the marginal configuration space, as a function of timestep Δt , for the given model problem and the particular choice of the Bussi-Parinello Langevin integrator OVRVO (7). Here the difference between exact and discrete configurational measures are plotted above the contours of the phase space density, for four values of the stepsize $\Delta t = [0.43, 0.66, 0.88, 1.1]$. In the **bottom row**, we illustrate the timestep-dependent error in a similar way for another integrator VRORV (11).

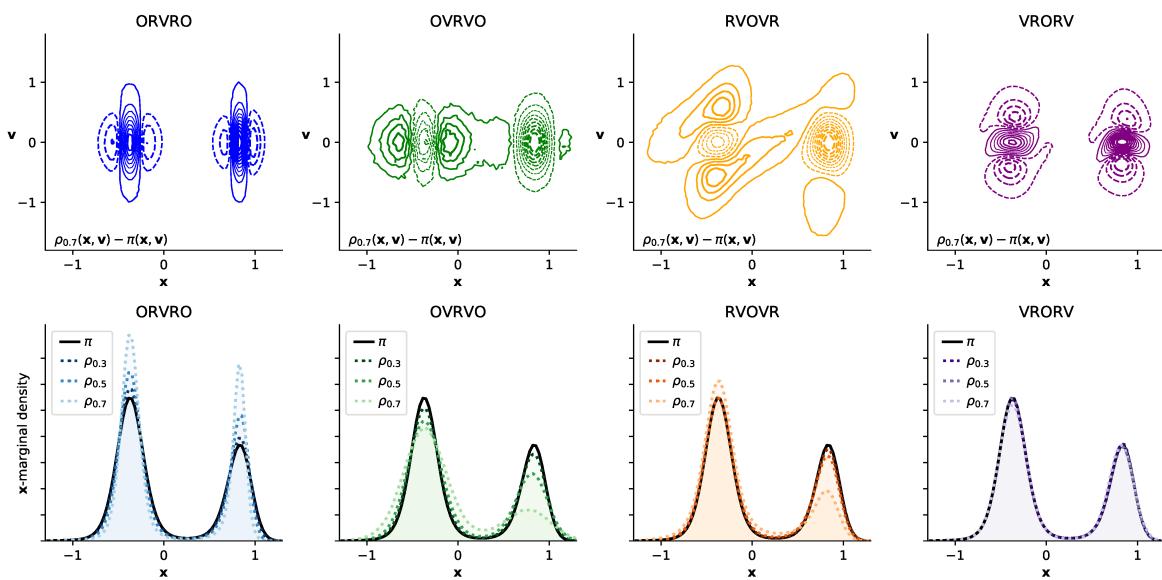


Figure 3. Different numerical integrators introduce different error structure in phase space, illustrated in a double-well system. Here, we illustrate the timestep-dependent discretization error introduced by four integrators on a 1D double-well potential [$U(x) \equiv x^6 + 2 \cos(5(x + 1))$].

The **top row** of 2D contour plots illustrates the difference between the phase-space density $\rho(x, v)$ sampled at the maximum timestep considered ($\Delta t = 0.7$, close to the stability limit) and the equilibrium density $\pi(x, v)$; solid lines indicate positive contours, while dashed lines indicate negative contours. The **bottom row** of 1D density plots shows timestep-dependent perturbation in the sampled marginal distribution in configuration space, ρ_x , with the equilibrium distribution π_x depicted as a solid black line. The sampled marginal distributions ρ_x are shown for increasingly large timestep, denoted $\rho_{\Delta t}$, depicted by increasingly light dotted lines, for $\Delta t = 0.3, 0.5, 0.7$ (arbitrary units). Inspecting the contour plots suggests that some integrator splittings (especially VRORV) induce error that fortuitously “cancels out” when the density is marginalized by integrating over v , while the error in other integrator splittings (ORVRO, OVRVO) constructively sums to amplify the error in configuration space.

because the numerical Langevin integrator is symmetric (so $\tilde{\Lambda} = \Lambda$) and because the time-independent Hamiltonian produces no explicit protocol work (so $w = w_{\text{shad}}$), in this case the KL divergence is approximately

$$\mathcal{D}_{\text{KL}}(\rho\|\pi) \approx \frac{1}{2} (\langle w \rangle_{\pi;\Lambda} - \langle w \rangle_{\rho;\Lambda}) \quad (15)$$

$$= \frac{1}{2} (\langle w_{\text{shad}} \rangle_{\pi;\Lambda} - \langle w_{\text{shad}} \rangle_{\rho;\Lambda}) , \quad (16)$$

the halved difference of two work averages: the work $\langle w_{\text{shad}} \rangle_{\pi}$ required to drive from equilibrium π into the steady state ρ , and the steady-state work $\langle w_{\text{shad}} \rangle_{\rho}$ expended over the same length of time, but starting in ρ (Figure 5a)³.

3.2. *A simple modification to the near-equilibrium estimator can compute KL divergence in configuration space*

In this study, we are especially interested in the configuration-space marginal distribution $\rho_{\mathbf{x}} \equiv \int d\mathbf{v} \rho(\mathbf{x}, \mathbf{v})$. The KL divergence $\mathcal{D}_{\text{KL}}(\rho_{\mathbf{x}}\|\pi_{\mathbf{x}})$ between the respective configuration-space marginal distributions $\rho_{\mathbf{x}}$ and $\pi_{\mathbf{x}} \equiv \int d\mathbf{v} \pi(\mathbf{x}, \mathbf{v})$ equals $\mathcal{D}_{\text{KL}}(\omega\|\pi)$, the KL divergence between the full equilibrium distribution π and the distribution

$$\omega(\mathbf{x}, \mathbf{v}) \equiv \rho_{\mathbf{x}}(\mathbf{x})\pi(\mathbf{v}|\mathbf{x}) , \quad (17)$$

that differs from π only in its \mathbf{x} -marginal:

$$\mathcal{D}_{\text{KL}}(\omega\|\pi) = \int d\mathbf{x} d\mathbf{v} \omega(\mathbf{x}, \mathbf{v}) \ln \left[\frac{\omega(\mathbf{x}, \mathbf{v})}{\pi(\mathbf{x}, \mathbf{v})} \right] \quad (18)$$

$$= \int d\mathbf{x} d\mathbf{v} \rho_{\mathbf{x}}(\mathbf{x})\pi(\mathbf{v}|\mathbf{x}) \ln \left[\frac{\rho_{\mathbf{x}}(\mathbf{x})\pi(\mathbf{v}|\mathbf{x})}{\pi_{\mathbf{x}}(\mathbf{x})\pi(\mathbf{v}|\mathbf{x})} \right] \quad (19)$$

$$= \int d\mathbf{x} \rho_{\mathbf{x}}(\mathbf{x}) \left[\int d\mathbf{v} \pi(\mathbf{v}|\mathbf{x}) \right] \ln \left[\frac{\rho_{\mathbf{x}}(\mathbf{x})}{\pi_{\mathbf{x}}(\mathbf{x})} \right] \quad (20)$$

$$= \int d\mathbf{x} \rho_{\mathbf{x}}(\mathbf{x}) \ln \left[\frac{\rho_{\mathbf{x}}(\mathbf{x})}{\pi_{\mathbf{x}}(\mathbf{x})} \right] \quad (21)$$

$$= \mathcal{D}_{\text{KL}}(\rho_{\mathbf{x}}\|\pi_{\mathbf{x}}) \quad (22)$$

This distribution ω is reached via an augmented protocol Λ' consisting of the original protocol Λ (repeated application of the numerical Langevin integrator) followed by final randomization of velocities according to the equilibrium conditional distribution $\pi(\mathbf{v}|\mathbf{x})$. Thus we construct an

³ In the limit that $\Delta t \rightarrow 0$, this relation is exact, since $\mathcal{D}_{\text{KL}}(\rho\|\pi) \rightarrow 0$ and both $\langle w_{\text{shad}} \rangle_{\pi;\Lambda} \rightarrow 0$ and $\langle w_{\text{shad}} \rangle_{\rho;\Lambda} \rightarrow 0$.

analogous near-equilibrium estimator of $\mathcal{D}_{\text{KL}}(\omega \parallel \pi)$ based on average the average shadow work w_{shad} accumulated by trajectories initiated from the respective distributions:

$$\mathcal{D}_{\text{KL}}(\rho_x \parallel \pi_x) = \mathcal{D}_{\text{KL}}(\omega \parallel \pi) \quad (23)$$

$$\approx \frac{1}{2} \left(\langle w_{\text{shad}} \rangle_{\pi; \tilde{\Lambda}'} - \langle w_{\text{shad}} \rangle_{\pi, \Lambda'; \tilde{\Lambda}'} \right) \quad (24)$$

$$= \frac{1}{2} \left(\langle w_{\text{shad}} \rangle_{\pi; \tilde{\Lambda}'} - \langle w_{\text{shad}} \rangle_{\omega; \tilde{\Lambda}'} \right) \quad (25)$$

$$= \frac{1}{2} \left(\langle w_{\text{shad}} \rangle_{\pi; \tilde{\Lambda}} - \langle w_{\text{shad}} \rangle_{\omega; \tilde{\Lambda}} \right) \quad (26)$$

$$= \frac{1}{2} \left(\langle w_{\text{shad}} \rangle_{\pi; \Lambda} - \langle w_{\text{shad}} \rangle_{\omega; \Lambda} \right) \quad (27)$$

$$= \frac{1}{2} \left(\langle w_{\text{shad}} \rangle_{\pi} - \langle w_{\text{shad}} \rangle_{\omega} \right). \quad (28)$$

Equation (25) follows from (26) because applying the time-reversed protocol $\tilde{\Lambda}'$ (velocity randomization followed by repeated application of the numerical Langevin integrator) to π or to ω is equivalent to applying $\tilde{\Lambda}$: π and ω already have velocity distributions randomized according to $\pi(\mathbf{v}|\mathbf{x})$, and the velocity randomization step does no work. In equation (28), we suppress the explicit protocol dependence, since henceforth all work averages will be averaged over the same protocol Λ .

Compared to the full-distribution KL divergence, the configuration-only KL divergence replaces the second work average $\langle w_{\text{shad}} \rangle_{\rho}$ with $\langle w_{\text{shad}} \rangle_{\omega}$, computing the expectation of the shadow work over a modified initial density constructed from the nonequilibrium steady-state configuration distribution but with Maxwell-Boltzmann velocity distribution $\pi(\mathbf{v}|\mathbf{x})$. Practically, this corresponds to drawing samples from the nonequilibrium steady-state ρ and replacing the velocities \mathbf{v} with an i.i.d. sample from $\pi(\mathbf{v}|\mathbf{x})$. The modified procedure is depicted schematically in Fig. 1b.

3.3. Comparison of phase-space error for different integrators

We first applied the original near-equilibrium method of Sivak *et al.* [1] to measure the timestep-dependent phase-space error introduced by four common Langevin integrators on a molecular mechanics model system (a cluster of 20 TIP3P waters [22] in a harmonic restraining potential), for a range of timesteps Δt between 0.1 and 8.0 femtoseconds (0.1 fs, 0.5 fs, 1.0 fs, ..., 7.5 fs, 8.0 fs). As illustrated in Figure 5a, while this approach can resolve statistically significant differences among schemes, none of the four integrator splitting schemes offers a considerable reduction in phase space error. This may be unsurprising, as none of these integrator schemes, when Metropolized, are known to provide a significant reduction in acceptance rates, which depend in some manner on the induced phase-space sampling error (Figure 8a; see Section 4 for more details about the relationship between Metropolized acceptance rates and KL divergence). Consistent with the results of [1], the phase-space error appears to scale approximately as $\mathcal{O}(\Delta t^4)$.

3.4. Comparison of configurational KL divergence for different integrators

Figure 5b shows that the measured KL divergence between the configuration-space marginals ρ_x and π_x can be drastically different among the four integrator schemes, and in some cases grow much more slowly than the associated phase-space sampling error (Fig. 5a). In particular, for VRORV, the error in the x -marginal is very nearly zero for the entire range of feasible timesteps, and it can be run at $\Delta t \approx 6$ fs while introducing the same amount of configuration error as other methods at $\Delta t \approx 2$ fs. This is consistent with prior findings [4,14], which showed the VRORV scheme introduces very little error in the average potential energy and multiple other system-specific observables.

We also note that for the OVRVO scheme, $\mathcal{D}_{\text{KL}}(\rho_x \parallel \pi_x) \approx \mathcal{D}_{\text{KL}}(\rho \parallel \pi)$ over the range of measured timesteps (Fig. 5), consistent with prior findings that estimates of $\mathcal{D}_{\text{KL}}(\rho \parallel \pi)$ tracked well with several measures of configuration-sampling error [1].

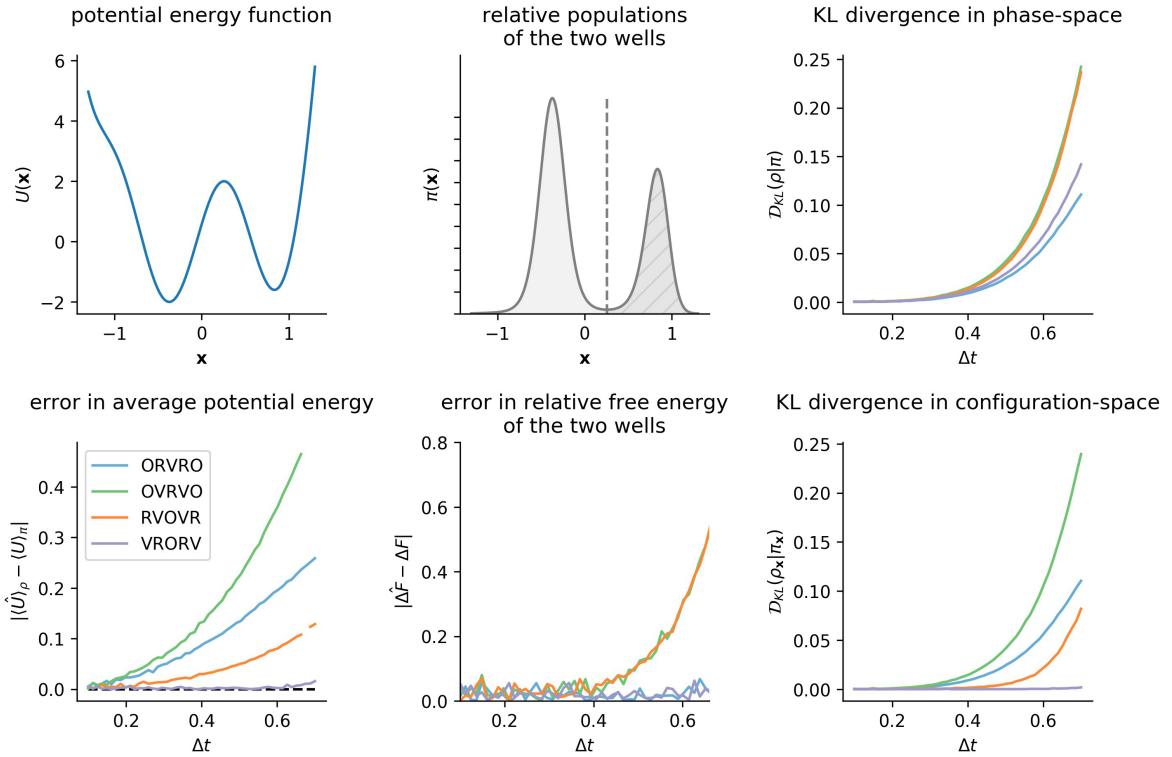


Figure 4. KL divergence is a natural measure of sampling error, although system-specific observables display different sensitivities to sampling error. Even for the simple double-well potential considered in Figure 3, configuration-space properties display different sensitivities to sampling error, motivating the use of a “universal” error measure, such as the KL divergence. The **top left** panel illustrates the double-well potential energy function from Figure 3, and the **top center** panel shows the resulting marginal equilibrium density, π_x , at $\beta = 1$. The **bottom left** panel shows, as a function of Δt , growth in the magnitude of the error in average potential energy, $|\langle U \rangle_\rho - \langle U \rangle_\pi|$, which has been used previously as a sensitive measure of sampling error [14]. The **bottom center** panel shows the error in the apparent free energy difference between the two wells as a function of Δt . Note that the timestep-dependent behavior of these two observables imply different rankings of integrator fidelity that may mislead one into believing error in *all* observables remains low with increasing timestep. However, as is clear here, just because an integrator introduces low timestep-dependent error in one observable does not mean that the method will introduce low error in another observable: for example, OVRVO preserves the well populations as accurately as VRORV, but introduces much larger errors in the average potential energy. The **right column** summarizes the growth in timestep-dependent error, as measured by the KL divergence. While all four integrators introduce comparable levels of Δt -dependent error in the phase-space distribution, they induce dramatically different magnitudes of error in the configuration-space marginal.

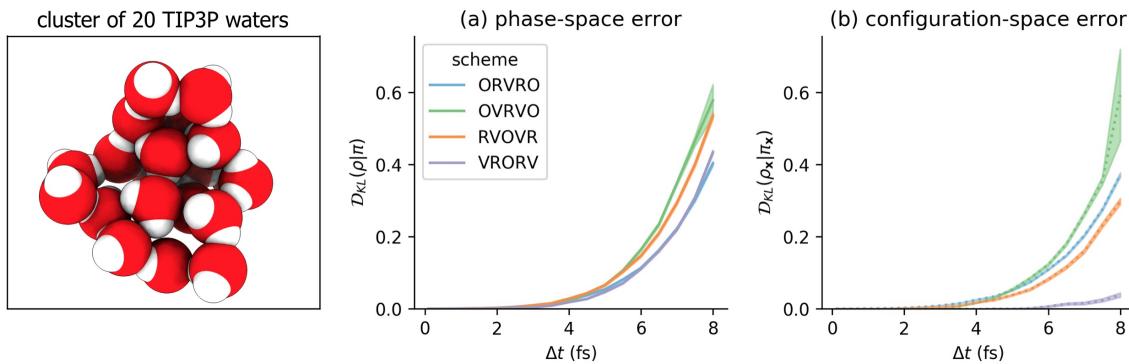


Figure 5. Using the near-equilibrium approximation, some numerical methods introduce far less configuration-space bias in molecular mechanics models than others. The results here are reported for a small cluster of rigid TIP3P waters, described in more detail in the Detailed Methods section, and illustrated in the leftmost panel. On the x -axis is the timestep Δt , measured in femtoseconds (fs). On the y -axis is the estimated KL divergence D_{KL} . (a) The error over the joint distribution on $D_{KL}(\rho \parallel \pi)$. (b) The error over the configuration-space marginal $D_{KL}(\rho_x \parallel \pi_x)$. Each colored curve corresponds to a numerical scheme for Langevin dynamics. The shaded region is the mean $\pm 95\%$ confidence interval.

231 3.5. Influence of the collision rate

232 The results reported above are for a single, relatively weak collision rate of 1 ps^{-1} . This collision
 233 rate was selected, following prior work [14], because it is low enough that barrier-crossing and
 234 conformational exploration should be fast, but large enough that low errors are observed empirically
 235 for configuration-space averages. As the formal “superconvergence” properties of VRORV were
 236 derived in the high-friction limit [14,31,32], it is worth considering how robustly the splitting VRORV
 237 introduces low configuration-space error at various collision rates. Additionally, as the collision rate
 238 goes to zero, we would expect the differences between pairs of the schemes that are equivalent upon
 239 removal of the O step (such as OVRVO and VRORV) to become smaller as the collision-rate goes to
 240 zero.

241 In Figure 6, we report near-equilibrium estimates of D_{KL} (as in Figure 5) over a range of collision
 242 rates spanning γ from 0.1 – 100 ps^{-1} . Strikingly, the configuration-space error introduced by VRORV
 243 remains low over this entire range, while other integrators (such as OVRVO) display a significant
 244 sensitivity to collision rate. In general, increasing collision rate increases phase- and configuration-space
 245 error for all integrators for which differences can statistically be differentiated, though the effect is
 246 modest over several orders of magnitude relative to the sensitivity with respect to timestep. For each
 247 condition, 50000 protocol samples were collected (one half the number of samples used in Figure 5), and
 248 the protocol length was 2000 steps (twice the protocol length used in Figure 5). Figure A7 demonstrates
 249 that this result is robust to protocol length for all collision rates considered.

250 3.6. Comparison with reference methods validates the near-equilibrium estimate

251 The accuracy of the near-equilibrium approximation introduced by Sivak *et al.* is largely
 252 unexplored. While the near-equilibrium approximation is computationally and statistically appealing,
 253 it is important to validate the accuracy of the approximation over the practical timestep Δt range
 254 of relevance to molecular simulation. In particular, it is unknown whether the near-equilibrium
 255 approximation produces an over-estimate or under-estimate of the KL divergence, or how accurate the
 256 approximation is for high-dimensional systems. Further, it is unknown whether any bias introduced
 257 by the approximation is uniform across different numerical methods for Langevin dynamics.

258 How well does the near-equilibrium estimator approximate the true KL divergence of relevant
 259 timestep ranges? The task of validating the near-equilibrium approximation is numerically challenging,

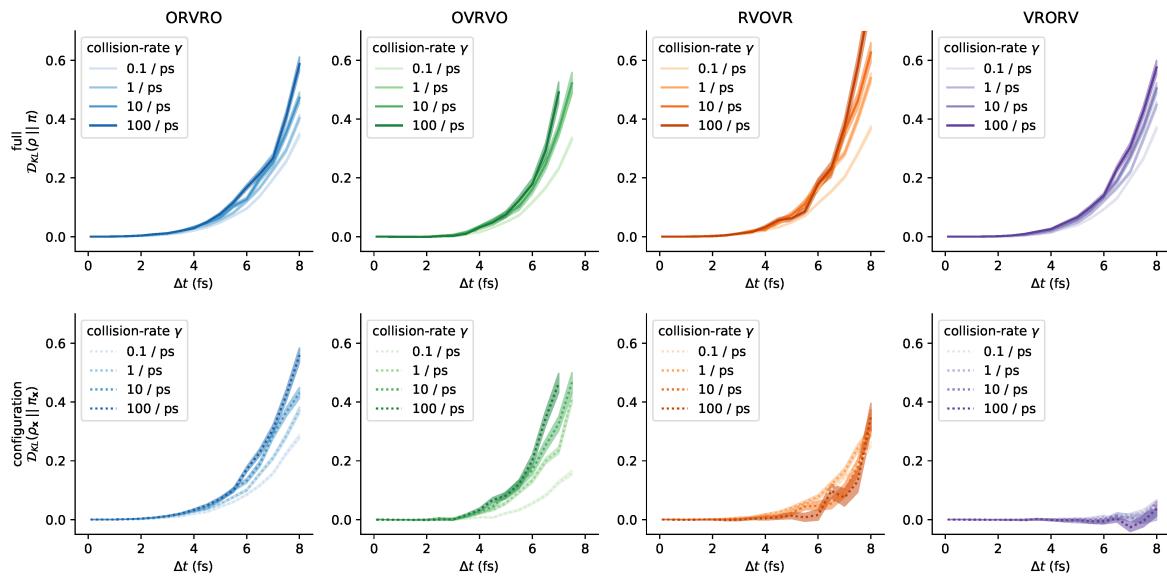


Figure 6. The choice of collision rate influences sampling bias. As we vary the collision rate γ over a few orders of magnitude, the resulting measured KL divergence responds in different ways for the different schemes. The phase-space bias appears to increase with increasing collision rate for all schemes. The configuration-space bias for OVRVO and ORVRO appears to *increase* with increasing collision rate, but the configuration-space bias for RVOVR appears to *decrease* with increasing collision rate. The anomalous low configuration-space error for VRORV is observed across all collision rates tested. The non-monotonic curves in the $\gamma = 100 \text{ ps}^{-1}$ condition are expected to be due to finite-sampling error, and are expected to be attenuated at a larger number of protocol samples. (Note that one condition is omitted from these plots for clarity: estimates of \mathcal{D}_{KL} for OVRVO at $\Delta t = 8 \text{ fs}$. At that timestep, the variance of the resulting \mathcal{D}_{KL} estimates for this scheme were much larger than for the other schemes.) See Figure A6 for a comparison grouped by collision rate, rather than by integrator.

260 since we are unaware of exact estimators for $\mathcal{D}_{\text{KL}}(\rho\|\pi)$ that remain tractable in high dimensions⁴. In
 261 the case of simple fluids, approximate methods are available that express the KL divergence in terms
 262 of a series of N -body correlations (as in [38]), typically truncating to two-body correlation functions
 263 (*i.e.*, comparing the radial distribution functions). However, in general we do not know the effect of
 264 truncating the expansion, since the successive terms in the series do not necessarily have decreasing
 265 magnitude.

266 To validate the near-equilibrium estimate of the KL-divergence, we attempt to “sandwich”
 267 it between two reference estimates that are unlikely to substantially over- or under-estimate the
 268 KL-divergence, and verify that the proposed method is consistent. In Section 3.6.1, we derive
 269 an asymptotically exact nested Monte Carlo method, that is computationally inefficient and an
 270 under-estimate in practice. In Section 3.6.2, we note that the results from the nested Monte Carlo
 271 method can be reprocessed to yield an over-estimate. In Section 3.6.3, we compute both, and compare
 272 with the near-equilibrium estimator.

273 3.6.1. Practical lower bound from nested Monte Carlo

274 First, we will derive an exact expression for the KL divergence between ρ and π in terms of
 275 quantities that we can measure, then discuss practical challenges that arise when using this expression,
 276 and under what conditions it becomes impractical. We start by writing the KL divergence as an
 277 expectation over ρ (30), since we cannot evaluate $\rho(\mathbf{x}, \mathbf{v})$ pointwise, but we can draw samples $(\mathbf{x}, \mathbf{v}) \sim \rho$.

$$\mathcal{D}_{\text{KL}}(\rho\|\pi) = \int d\mathbf{x} d\mathbf{v} \rho(\mathbf{x}, \mathbf{v}) \ln \left[\frac{\rho(\mathbf{x}, \mathbf{v})}{\pi(\mathbf{x}, \mathbf{v})} \right] \quad (29)$$

$$= \left\langle \ln \left[\frac{\rho(\mathbf{x}, \mathbf{v})}{\pi(\mathbf{x}, \mathbf{v})} \right] \right\rangle_{\rho} \quad (30)$$

$$= \left\langle \ln \left[\frac{\pi(\mathbf{x}, \mathbf{v}) \langle e^{-w} \rangle_{\mathbf{x}, \mathbf{v}; \tilde{\Lambda}}}{\pi(\mathbf{x}, \mathbf{v})} \right] \right\rangle_{\rho} \quad (31)$$

$$= \langle \ln \langle e^{-w} \rangle_{\mathbf{x}, \mathbf{v}; \Lambda} \rangle_{\rho} \quad (32)$$

$$\approx \frac{1}{N} \sum_{i=1}^N \left(\ln \sum_{j=1}^M \frac{1}{M} e^{-w_{ij}} \right) \quad (33)$$

278 We note that the inner ratio of nonequilibrium steady-state to equilibrium densities, $\rho(\mathbf{x}, \mathbf{v}) / \pi(\mathbf{x}, \mathbf{v})$,
 279 can be expressed in terms of $\langle e^{-w} \rangle_{\mathbf{x}, \mathbf{v}; \tilde{\Lambda}}$ —the average of exponentiated nonequilibrium work measured
 280 under the application of the time-reversed protocol $\tilde{\Lambda}$ starting from (\mathbf{x}, \mathbf{v}) (31). Λ denotes the protocol
 281 used to generate ρ from a sample from π —in this case, T applications of the Langevin integrator step
 282 kernel; $\tilde{\Lambda}$ denotes the time-reverse of this protocol. Since the protocol we apply to generate ρ from π is
 283 time-symmetric for integrators derived from symmetric Strang splittings⁵, we can substitute $\Lambda = \tilde{\Lambda}$. In
 284 the final step, we substitute a simple Monte Carlo estimator of that average, in terms of work samples
 285 w_{ij} , where w_{ij} is the j th reduced (unitless) work measurement collected from initial condition i . Here,
 286 N is the number of initial conditions sampled (*i.e.*, the number of “outer-loop” samples), and M is the
 287 number of work samples (*i.e.*, the number of “inner-loop” samples) collected at each initial condition
 288 $(\mathbf{x}_i, \mathbf{v}_i) \sim \rho$.

⁴ Spatial discretization or density estimation are infeasible, due to curse of dimensionality. There are direct estimators of the KL divergence based on Euclidean nearest-neighbor distances that perform well in some high-dimensional settings (*e.g.*, [37]), but Euclidean distance is an unsuitable metric on molecular configurations.

⁵ Note that applying this methodology to non-symmetric integrators (where the sequence of operations for the integrator its time-reverse are not identical) would require modifications to this scheme, as well as the manner in which shadow work is computed.

289 Related nested estimators have been proposed in the literature. Notably, an estimator for the
 290 nonequilibrium entropy $S(\rho) \equiv -\int d\mathbf{x} \rho(\mathbf{x}) \ln \rho(\mathbf{x})$ in terms of work averages is given in equation 18
 291 of [39], and this can in turn be used to estimate the desired KL divergence if we also have suitable
 292 estimates for the equilibrium entropy $S(\pi)$.

293 The required work values w_{ij} can be easily computed from simulations. To sample an initial
 294 condition $(\mathbf{x}_i, \mathbf{v}_i)$ from ρ , we simply run the Langevin integrator of interest for a sufficient number of
 295 steps to sample a new uncorrelated configuration from the nonequilibrium steady-state sampled by
 296 the integrator. To compute the work accumulated from a given starting condition, we use the notion of
 297 *shadow work* [23]. For numerical methods constructed from symmetric Strang splittings involving the R,
 298 V, and O operations described above, we simply need to compute the sum of the total energy changes
 299 during the deterministic substeps (*i.e.*, the potential energy change during deterministic updates of
 300 the position variables, and the kinetic energy change during deterministic updates of the momentum
 301 variables). For convenience, we use reduced (unitless) energies and work values throughout, where
 302 factors of $k_B T$ have been removed, without loss of generality. See Detailed Methods (Section 6.4) for a
 303 detailed description on how shadow work can be computed from this family of Langevin integrators
 304 in general.

305 Like the near-equilibrium scheme, this nested scheme can be modified analogously to
 306 measure the configuration-space error in isolation, by initializing instead from the distribution
 307 $\omega(\mathbf{x}, \mathbf{v}) \equiv \rho_{\mathbf{x}}(\mathbf{x}) \pi(\mathbf{v}|\mathbf{x})$, allowing us to compute $\mathcal{D}_{\text{KL}}(\omega\|\pi)$, a quantity that is identical to $\mathcal{D}_{\text{KL}}(\rho_{\mathbf{x}}\|\pi_{\mathbf{x}})$
 308 (see 3.2). Specifically, to measure the full KL divergence, we sample initial conditions from the Langevin
 309 integrator's steady state: $(\mathbf{x}_i, \mathbf{v}_i) \sim \rho$. To measure configuration-space-only KL divergence, we draw
 310 initial configuration from the integrator's steady state, and velocities from equilibrium: $\mathbf{x}_i \sim \rho_{\mathbf{x}}$,
 311 $\mathbf{v}_i \sim \pi(\mathbf{v}_i|\mathbf{x}_i)$. Note that, for constrained systems, $\pi(\mathbf{v}|\mathbf{x})$ is not independent of \mathbf{x} , and care must be
 312 taken to eliminate velocity components along constrained degrees of freedom before measuring the
 313 contribution of the integrator substep to the shadow work (see Detailed Methods).

314 We note that the nested plug-in Monte Carlo estimator of the KL divergence is asymptotically exact
 315 only when both N (the number of “outer-loop” samples) and M (the number of “inner-loop” samples)
 316 go to infinity. For a finite number of samples, the nested estimator will produce an under-estimate
 317 of the KL divergence. To see this, note that the nested estimator is a simple average of many likely
 318 underestimates, so it should itself be an underestimate. More specifically, although the outer-loop
 319 expectation can be approximated without bias (since samples can be drawn from ρ), each of the
 320 inner-loop expectations is an exponential average ($\log\langle\exp(-w)\rangle$), which we will underestimate when
 321 we plug in a finite sample of w 's (for the same reasons that the EXP estimator for free energies is biased
 322 [40] – To leading order, that under-estimate is related to the variance of w [40], which here grows
 323 rapidly with Δt as shown in Figure A2).

324 In practice, we use a simple adaptive scheme (described in detail in Section 6.6) that draws inner-
 325 and outer-loop samples until uncertainty thresholds are met, which should minimize the magnitude
 326 of this bias.

327 3.6.2. Practical upper bound from Jensen's inequality

328 A simple, but practically useful, upper bound for the KL divergence can be obtained from the
 329 application of Jensen's inequality, $\langle \ln x \rangle \leq \ln \langle x \rangle$, to (32):

$$\mathcal{D}_{\text{KL}}(\rho\|\pi) = \langle \ln \langle e^{-w} \rangle_{\mathbf{x}, \mathbf{v}; \Lambda} \rangle_{\rho} \quad (34)$$

$$\leq \ln \langle \langle e^{-w} \rangle_{\mathbf{x}, \mathbf{v}; \Lambda} \rangle_{\rho} \quad (35)$$

$$= \ln \langle e^{-w} \rangle_{\rho; \Lambda}. \quad (36)$$

330 The analogous inequality for the configuration-space marginal is

$$\mathcal{D}_{\text{KL}}(\rho_{\mathbf{x}}\|\pi_{\mathbf{x}}) = \mathcal{D}_{\text{KL}}(\omega\|\pi) \quad (37)$$

$$= \langle \ln \langle e^{-w} \rangle_{\mathbf{x}, \mathbf{v}; \Lambda} \rangle_{\omega} \quad (38)$$

$$\leq \ln \langle \langle e^{-w} \rangle_{\mathbf{x}, \mathbf{v}; \Lambda} \rangle_{\omega} \quad (39)$$

$$= \ln \langle e^{-w} \rangle_{\omega; \Lambda}. \quad (40)$$

331 This provides a particularly convenient upper bound on the KL-divergence, since it can be computed
 332 by reprocessing work samples collected from the nested Monte Carlo scheme⁶.

333 3.6.3. Sandwiching the KL divergence to validate the near-equilibrium estimate

334 We compared these three estimates of the KL divergence on the molecular mechanics system
 335 introduced in Figure 5, and confirmed that the near equilibrium estimate falls between the likely over-
 336 and under-estimate for all four integrator schemes, over a range of feasible timesteps (Figure 7, and in
 337 log-scale in Figure A5). We conclude that the near-equilibrium approximation is empirically reliable
 338 for measuring integrator bias on molecular mechanics models for practical timesteps.

339 **4. Relation to GHMC acceptance rates**

340 What is the relationship between the bias introduced by an integrator at steady state, and the
 341 acceptance rate of the corresponding Metropolized integrator? Specifically, why not construct a
 342 Metropolized version of VRORV to guarantee samples are drawn appropriately from the equilibrium
 343 target density $\pi(\mathbf{x})$? Following [3,41], we can construct an exact MCMC method that uses one or more
 344 steps of Langevin dynamics as a proposal, by using

$$\alpha \equiv \min\{1, e^{-w_{\text{shad}}}\} \quad (41)$$

345 as the acceptance criterion. The resulting method is called *generalized hybrid Monte Carlo* (GHMC),
 346 and eliminates time-discretization error at the cost of increasing sample autocorrelation (see also the
 347 comments in Section 6.3). A natural question arises: if an uncorrected Langevin integrator introduces
 348 low configuration-space error, is the rejection rate of the corresponding Metropolis-corrected method
 349 also low?

350 To answer this question, we estimated the GHMC acceptance rate at all conditions for which we
 351 have estimated steady-state \mathcal{D}_{KL} . Given a collection of equilibrium samples (described in Section 6.3),
 352 we can efficiently estimate the acceptance rate of an MCMC proposal by taking the sample average of
 353 the acceptance ratio α over proposals originating from equilibrium, $(\mathbf{x}_0, \mathbf{v}_0) \sim \pi$.

354 We compared the GHMC acceptance rate to the histogram-based \mathcal{D}_{KL} estimates for the 1D
 355 double-well system in Figure 8. There does not appear to be a consistent relationship between \mathcal{D}_{KL}
 356 and acceptance rate across the four schemes. Notably, the GHMC rejection rate can be extremely
 357 “conservative” for splittings such as VRORV.

358 Next, we compared the GHMC rejection rate with the near-equilibrium \mathcal{D}_{KL} estimates for the
 359 water cluster considered in Figure 9. A similar pattern is recapitulated in this molecular mechanics
 360 model as in the 1D system—there is not a consistent relationship between configuration-space bias
 361 introduced by a Langevin integrator and the rejection rate of its corresponding GHMC method.

362 This complicates the decision of whether to Metropolize or not. As noted in Section 6.3, incurring
 363 even a small rejection rate in GHMC can have a large effect on statistical efficiency, due to the effect of

6 Note that when we approximate this bound with a finite number of samples, we will underestimate it for the reasons mentioned in Section 3.6.1. However, since we are pooling all work samples, the magnitude of this underestimate should be much smaller than for the inner-loop underestimates in Section 3.6.1 above, and we expect the magnitude of this bias to be negligible compared with the effect of invoking Jensen’s inequality.

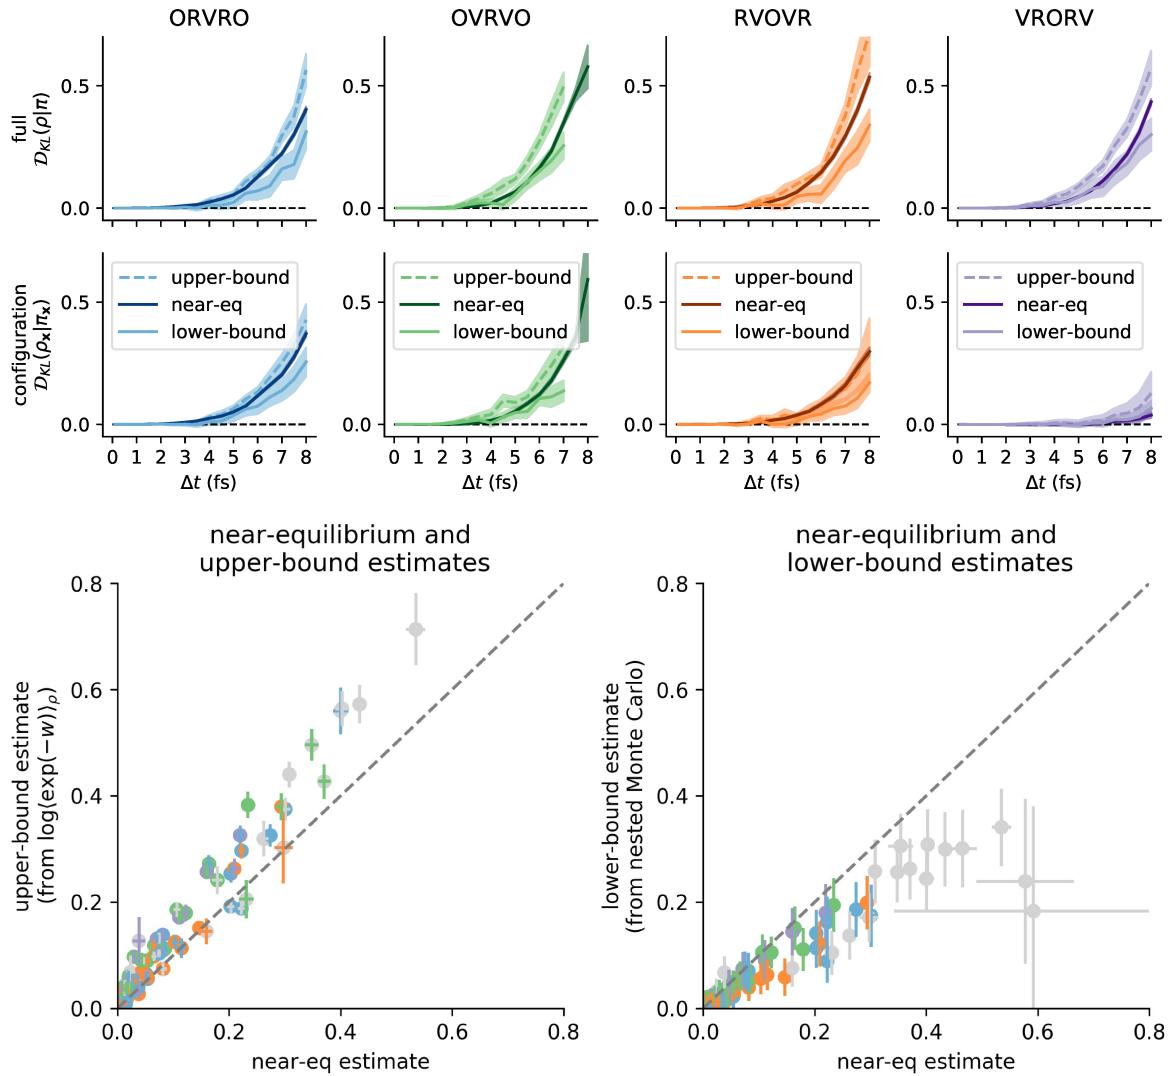


Figure 7. The near-equilibrium estimator is consistent with reference estimators for a practical range of Δt . We compared the near-equilibrium estimates reported in Figure 5 for the water cluster against a likely under-estimate and a likely over-estimate of the \mathcal{D}_{KL} . In the **top row**, we validate near-equilibrium estimates of the KL divergence on the full state space (x, v). In the **bottom row**, we validate near-equilibrium estimates of the KL divergence on configuration space (x) alone. Each column corresponds to a numerical method for Langevin dynamics. The **darker band** in each plot corresponds to the near-equilibrium estimate $\pm 95\%$ confidence intervals from asymptotic uncertainty estimate (details in section 3.1). The **lighter band with a solid line** corresponds to the nested Monte Carlo estimate $\pm 95\%$ confidence intervals from bootstrapping (details in section 6.6). The **lighter band with a dotted line** corresponds to the exponential average estimate $\pm 95\%$ confidence intervals from bootstrapping (details in section 6.6). Log-scale versions of these plots are provided in the appendix also, A5. In the **lower two panels**, we summarize these results by plotting all near-equilibrium estimates vs. all exponential-average estimates (left) and all near-equilibrium estimates vs. all nested Monte Carlo estimates (right). The colored dots and bars correspond to the means \pm uncertainties used in the earlier panels. The dashed diagonal line shows parity. Grey error dots and error bars correspond to conditions where the nested Monte Carlo estimate reached the computational budget (5×10^4 inner-loop samples) but failed to reach the inner-loop uncertainty threshold, and is thus more biased. See Section 6.6 for additional details.

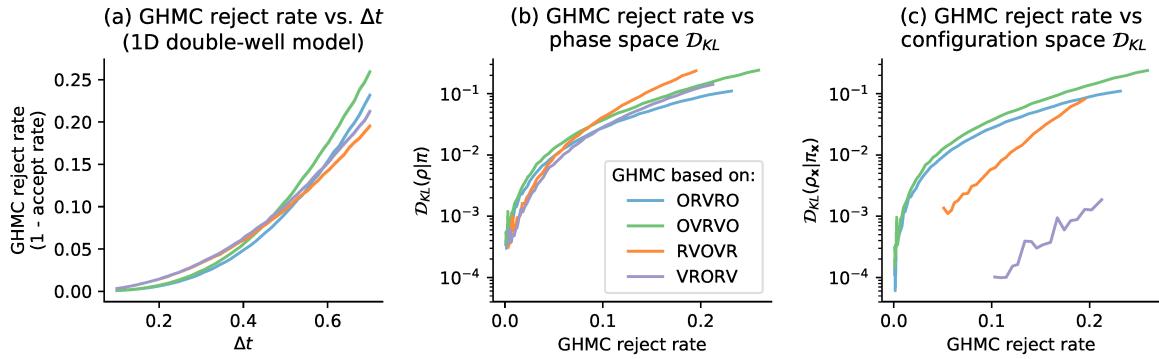


Figure 8. No consistent relationship between the GHMC acceptance rate and the steady-state bias is apparent for a 1D system. Since the GHMC rejection rate grows similarly with Δt across all four schemes, but the configuration-space KL divergence does not, the GHMC rejection rate can be overly “conservative” for some splittings. Panel (a) shows the growth in the GHMC rejection rate as a function of timestep Δt , for the 1D double-well model considered in Figures 3 and 4. On the x-axis is an evenly spaced grid of 50 timesteps between 0.1 and 0.7. On the y-axis is the estimated rejection rate, which is based on a sample average of the GHMC acceptance criterion. The shaded region is the mean \pm 95% confidence interval. Panel (b) compares the GHMC rejection rate vs. the *phase-space* bias at steady state, over the range of timesteps plotted in panel (a). The y-axis is KL divergence between the phase-space histograms, plotted on a log-scale. Panel (c) compares the GHMC rejection rate vs. the *configuration-space* bias at steady state, over the range of timesteps plotted in panels (a), (b). The y-axis is the KL divergence between the configuration-space histograms, plotted on a log-scale. Note that in panel (c), we have truncated the leftmost parts of the curves for RVOVR and VRORV rejection rates less than 0.05 and 0.1, respectively, due to noise in histogram estimates of very small $D_{KL}(\rho_x \parallel \pi_x)$.

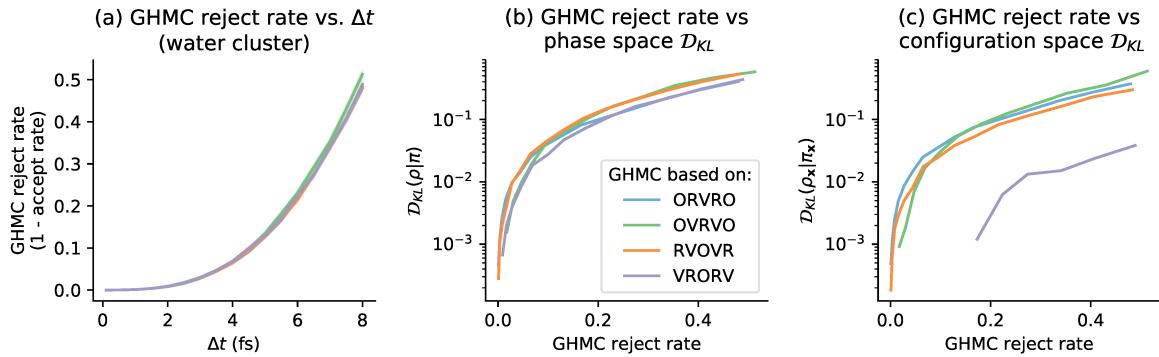


Figure 9. Near-equilibrium measurements recapitulate the relationship between steady-state D_{KL} and GHMC acceptance rate for the water cluster test system. Panel (a) shows the growth in the GHMC rejection rate (1 minus the acceptance rate) as a function of timestep Δt (in femtoseconds), for the water cluster test system illustrated in Figure 5. On the x-axis are timesteps [0.1 fs, 0.5 fs, 1.0 fs, ... 7.5 fs, 8 fs]. On the y-axis is the estimated rejection rate, which is based on a sample average of the GHMC acceptance criterion, over 10000 proposals per condition. The shaded region is the mean \pm 95% confidence interval. Panel (b) compares the GHMC rejection rate vs. the *phase-space* bias at steady state, over the range of timesteps plotted in panel (a). The y-axis is the KL divergence between the phase-space distributions as measured by the near-equilibrium estimate, plotted on a log-scale. Panel (c) compares the GHMC rejection rate vs. the *configuration-space* bias at steady state, over the range of timesteps plotted in panels (a), (b). Note that in panels (b) and (c), we have truncated at $D_{KL} \leq 10^{-4}$, due to noise in near-equilibrium estimates of very small D_{KL} .

364 momentum flipping. An open challenge is to construct Metropolis criteria for GHMC that might be
 365 less “wasteful” for Langevin splittings that introduce low configuration-space bias. One possibility is
 366 to use a multi-proposal Metropolization procedure to correct estimates of equilibrium expectations, as
 367 done in [42].

368 **5. Discussion**

369 We have introduced and validated a work-based estimator of the KL divergence over the
 370 configuration-space marginal sampled by Langevin dynamics. We demonstrated that we could
 371 use this estimator to measure differences between the timestep-dependent configuration-sampling
 372 error introduced by each member of a small class of Langevin integrators on molecular mechanics
 373 models. Work-based estimators are especially attractive for biomolecular systems, since expectations
 374 over work distributions are often tractable when other approaches are not.

375 Reliable estimates of KL divergence using the work-based estimator considered here require
 376 knowledge of the time to reach nonequilibrium steady state. This near-equilibrium approach requires
 377 that the user select a trajectory length T sufficiently large to reach the nonequilibrium steady state, or
 378 else the KL divergence estimate could be substantially biased. Opposing the choice of large T is the
 379 variance of the estimate, since the contribution of the steady-state work to the variance of the estimate
 380 grows as T . Taken together, this suggests the smallest time T that produces unbiased estimates is
 381 optimal. In our calculations, it was sufficient to use a protocol that was twice the average collision
 382 time, but in general this choice should be validated for the system under study. One way to do this, for
 383 example, is to perform the same computation for T and $2T$ and ensure estimates are concordant to
 384 within statistical error.

385 Generating equilibrium samples from $\pi(\mathbf{x}, \mathbf{v})$ can be difficult for large, complex molecular systems.
 386 The near-equilibrium method requires access to a large number of independent samples from the
 387 equilibrium distribution of the system under study. In this work, we used extra-chance HMC [41,43]
 388 to construct a large cache of independent equilibrium samples, amortizing the cost of equilibrium
 389 sampling across the many integrator variants (described in Section 6.3). In cases where we would
 390 like to compare a large number of integrator variants on the same system, this can be an acceptable
 391 cost, but in other cases, it may be prohibitive. It is unknown whether we can relax this requirement
 392 in practice, and use only samples in “local equilibrium,” since integrator error may be dominated by
 393 local features of the energy landscape. An alternative would be that, if the primary contributions to
 394 the dissipation process that drive integration errors arise from high-frequency motions, adding a weak
 395 restraint to parts of the system under study (such as a biological macromolecule) may also allow rapid
 396 assessment of integrator schemes in a region of configuration space where estimates can be easily
 397 converged. We have not yet tested this hypothesis, but if it is possible to relax the requirement of i.i.d.
 398 equilibrium samples and retain accurate estimates, then the method will be much cheaper to apply in
 399 difficult settings.

400 *5.1. Future directions*

401 The validation of the near-equilibrium estimate makes it possible to apply the technique to a
 402 systematic comparison of sampling bias in integrator variants and biomolecular systems. Although
 403 we considered only four Langevin integrators here, this approach can be applied to any stochastic
 404 integrator for which the relative path action can be computed (see [44] for examples of how to compute
 405 the relative action for stochastic integrators not necessarily derived from operator splitting).

406 Independently, the work-based estimate for $\ln[\rho_{\mathbf{x}}(\mathbf{x})/\pi(\mathbf{x})]$ we used in the expensive lower bound
 407 (Section 3.6) could be useful for other analyses. For example, an estimate of $\ln[\rho_{\mathbf{x}}(\mathbf{x})/\pi(\mathbf{x})]$ could
 408 be used to interpret what features of \mathbf{x} are most distorted by integrator bias, *e.g.*, by checking which
 409 features of \mathbf{x} are most predictive of extreme values of $\ln[\rho_{\mathbf{x}}(\mathbf{x})/\pi(\mathbf{x})]$.

410 We also note that nothing about the derivation is specific to the partition between configuration
 411 degrees of freedom and velocities. We could also use this method to measure the KL divergence over

any subset S of the state variables $\mathbf{z} = (\mathbf{x}, \mathbf{v})$, provided we can sample from the conditional distribution for the complementary subset S' of the state variables: $\pi(\mathbf{z}_{S'} | \mathbf{z}_S)$. To measure KL divergence over the configuration variables, we need only sample from the conditional distribution of velocities given positions, which is typically tractable. Provided that the required conditional distribution is tractable, this method could also prove useful in contexts other than measuring integrator error.

Finally, in this study we have only considered sampling the canonical ensemble (NVT; constant temperature, particle number, and *volume*), but the isothermal-isobaric ensemble (NpT; constant temperature, particle number, and *pressure*) also has wide practical relevance. In OpenMM, the isothermal-isobaric ensemble is simulated by alternating between sampling the canonical ensemble (using thermostatted dynamics, such as Langevin dynamics), and periodically sampling the volume using a molecular-scaling Monte Carlo barostat. For sufficiently infrequent barostat proposals, the bias introduced by interaction between the finite-timestep Langevin integrator and the volume perturbations is expected to be minimal, so applying the proposed near-equilibrium method is expected to approximate well the overall sampling bias⁷.

6. Detailed methods

All code used in this paper, along with a manifest of all conda-installable prerequisites and version numbers needed to run the code, is available at <https://github.com/choderalab/integrator-benchmark> under the permissive [Open Source Initiative](#) approved MIT license.

A byproduct of this work is a flexible implementation of Langevin integrators derived from operator splitting for the GPU-accelerated OpenMM molecular simulation framework [45], also available under the MIT license in the `openmmtools` library: <https://github.com/choderalab/openmmtools>. This implementation allows the user to specify a Langevin integrator using a splitting string (like OVRVO) and can automatically compute shadow work for each splitting.

6.1. One-dimensional model system: Double well

For illustration and to have a model system where the exact \mathcal{D}_{KL} was readily computable using histograms, we constructed and analyzed a double-well model in 1D. The potential energy function of this model is $U(x) \equiv x^6 + 2 \cos(5(x+1))$, illustrated in Figure 4. We implemented the four Langevin schemes under study using Numba 0.35.0 [46] for use with 1D toy models. We used a temperature of $\beta = 1$, a collision rate of $\gamma = 10$, and a mass of $m = 10$. For these conditions, we found a maximum stable timestep of approximately $\Delta t = 0.7$. Histogram-based estimates of the configuration-space density and phase-space density used 100 bins per dimension, where the bin edges were set by bounding box of a trial run at the maximum Δt . The equilibrium density of each bin was computed using numerical quadrature (using the trapezoidal rule, `numpy.trapz`). The KL divergence between a given ρ and π was then computed using `scipy.stats.entropy` on the histogram representation.

6.2. Model molecular mechanics system: A harmonically restrained water cluster

As noted, the exact Monte Carlo method involves exponential work averages, resulting in a statistical inefficiency that grows rapidly with both the size of the system and the distance from equilibrium. Since we are interested in identifying whether the near-equilibrium approximation breaks down over the timestep Δt range of interest to molecular simulations, it is important to be able to compute a reliable estimate of \mathcal{D}_{KL} far from equilibrium. Thus, we aim to select the smallest system we think will be representative of the geometry of molecular mechanics models generally in order to allow the exact estimate to be computable with reasonable computing resources.

⁷ In [23], the error in phase space was measured for an ensemble of constant-volume Langevin trajectories with initial conditions drawn from the isothermal-isobaric ensemble.

454 To compare the proposed method with a reference estimator, we needed to select a test system
 455 which met the following criteria:

- 456 1. The test system must have **interactions typical of solvated molecular mechanics models**, so
 457 that we would have some justification for generalizing from the results. This rules out 1D
 458 systems, for example, and prompted us to search for systems that were not alanine dipeptide in
 459 vacuum.
- 460 2. The test system must have **sufficiently few degrees of freedom that the nested Monte Carlo**
 461 **estimator remains feasible**. Because the nested estimator requires converging many exponential
 462 averages, the cost of achieving a fixed level of precision grows dramatically with the standard
 463 deviation of the steady-state shadow work distribution. The width of this distribution is
 464 extensive in system size. Empirically, this ruled out using the first water box we had tried
 465 (with approximately 500 rigid TIP3P waters [22], with 3000 degrees of freedom). Practically,
 466 there was also a limit to how small it is possible to make a water system with periodic boundary
 467 conditions in OpenMM (about 100 waters, or 600 degrees of freedom), which was also infeasible.
- 468 3. The test system must have **enough disordered degrees of freedom that the behavior of work**
 469 **averages is typical of larger systems**. This was motivated by our observation that it was
 470 paradoxically much easier to converge estimates for large disordered systems than it was to
 471 converge estimates for the 1D toy system.

472 To construct a test system that met all of those criteria, we used a `WaterCluster` test system, which
 473 comprises 20 rigid TIP3P waters weakly confined in a central harmonic restraining potential with force
 474 constant $K = 1 \text{ kJ/mol/nm}^2$ applied to all atoms. This test system is available in version 0.14.0 of
 475 the `openmmtools` package [47]. Simulations were performed in double-precision using the `Reference`
 476 platform in OpenMM 7.2 [48] to minimize the potential for introducing significant round-off error due
 477 to finite floating point precision.

478 6.3. Caching equilibrium samples

479 To enable this study, we attempted to amortize the cost of collecting i.i.d. samples from each test
 480 system's equilibrium distribution π and various integrator-and- Δt -specific distributions ρ . Since there
 481 are many different distributions ρ , and all are relatively small perturbations of π , we invest initial
 482 effort into sampling π exhaustively, and then we draw samples from each integrator-specific ρ by
 483 running the integrator of interest from initial conditions $(\mathbf{x}_0, \mathbf{v}_0) \sim \pi$.

484 For each test system, we pre-computed a large collection of $K = 1000$ equilibrium samples

$$\pi_{\mathbf{x}}(\mathbf{x}) \approx \pi^{\text{cache}}(\mathbf{x}) \equiv \frac{1}{K} \sum_{k=1}^K \delta(\mathbf{x} - \mathbf{x}^{(k)}) \quad (42)$$

485 using *extra-chance Hamiltonian Monte Carlo* (XC-HMC)[41,43], implemented as a `CustomIntegrator` in
 486 OpenMM [48]. In brief, XC-HMC is a strategy to reduce the adverse effects of momentum flipping on
 487 sampling autocorrelation times in GHMC. GHMC uses Langevin integration of Hamiltonian dynamics
 488 as a proposal mechanism, and accepts or rejects each proposal according to a Metropolis criterion.
 489 Whenever the Metropolis test fails, the proposal is rejected and the momentum must be reversed;
 490 this is necessary to maintain detailed balance[3,49], but can lead to extremely large autocorrelation
 491 times when the acceptance rate is not sufficiently close to 100% (see [43] for empirical examples and
 492 further discussion). In “extra-chance” HMC [41,43], rather than immediately flipping the momentum
 493 whenever the Metropolis criterion fails, the proposal trajectory is instead extended to generate a new
 494 proposal, and another (suitably modified) Metropolis criterion is checked. For suitable choices of
 495 parameters (length of trajectory proposal, timestep, number of extra chances, length of “extra-chance”
 496 trajectories), this strategy can virtually eliminate the effect of momentum flipping, at the cost of
 497 increasing the average length of proposal trajectories.

498 Our initial experiments (on larger systems than reported here) suggested that the cost of collecting
 499 uncorrelated samples using GHMC without “extra-chances” was prohibitive, since we needed to make
 500 the timestep extremely small (around 0.1–0.25 fs) to keep the acceptance rate sufficiently near 100%
 501 that the effect of momentum flipping was acceptable. Instead, we equilibrated for approximately 1 ns
 502 (10^5 XC-HMC iterations, 10 steps per XC-HMC proposal trajectory, 15 extra-chance trajectories per
 503 iteration, 1 fs per timestep) from an energy-minimized starting structure. We then saved one sample
 504 $\mathbf{x}^{(i)}$ every 10^4 XC-HMC iterations afterwards.

505 To draw an i.i.d. sample from $\pi(\mathbf{x}, \mathbf{v})$, we draw a sample \mathbf{x} uniformly from π^{cache} , and then sample
 506 \mathbf{v} from the equilibrium distribution of velocities conditioned on $\mathbf{x}, \mathbf{v} \sim \pi(\mathbf{x}|\mathbf{v})$. (In the presence of
 507 holonomic constraints, the velocity distribution is not independent of the configuration distribution.
 508 For example, if bond lengths involving hydrogen atoms are constrained, the velocity of a hydrogen
 509 minus the velocity of the bonded heavy atom cannot have any component parallel to the bond.)

510 To draw an i.i.d. sample from ρ , we draw an i.i.d. sample from π and then simulate Langevin
 511 dynamics for a large number of steps. We tested using 1000 steps, $(1/\gamma)/\Delta t$ steps, and $(2/\gamma)/\Delta t$
 512 steps.

513 6.4. Computing shadow work for symmetric Strang splittings

514 Here, we demonstrate how to compute the appropriate shadow work for a given discrete-timestep
 515 Langevin integration scheme. We note that the log-ratio of forward and reverse conditional path
 516 probabilities under specific time-discretizations of Langevin dynamics has been calculated in many
 517 prior works, such as [10,42,50]. While there has been a great deal of confusion in the literature about
 518 how nonequilibrium work should be computed [51], fortunately, there is an unambiguous mechanical
 519 (if tedious) approach to the computation of the appropriate work-like quantity.

520 Assemble the sequence of all steps and substeps of an integrator cycle into a trajectory Z . For
 521 example, for OVRVO, we have

$$522 Z \equiv \{z_0 \xrightarrow{O} z_1 \xrightarrow{V} z_2 \xrightarrow{R} z_3 \xrightarrow{V} z_4 \xrightarrow{O} z_5\} \quad (43)$$

523 where $z_n \equiv (\mathbf{x}_n, \mathbf{v}_n)$ are phase-space points. Let \tilde{Z} denote the time-reversal of all substeps of Z with all
 velocities negated,

$$524 \tilde{Z} \equiv \{\tilde{z}_5 \xrightarrow{O} \tilde{z}_4 \xrightarrow{V} \tilde{z}_3 \xrightarrow{R} \tilde{z}_2 \xrightarrow{V} \tilde{z}_1 \xrightarrow{O} \tilde{z}_0\} \quad (44)$$

525 where $\tilde{z}_n = (\mathbf{x}_n, -\mathbf{v}_n)$ denotes negating the velocity of phase-space point z_n .

526 To compute the reduced, unitless shadow work $w[Z]$, we use the definition of work that satisfies
 527 the Crooks fluctuation theorem for starting with a sample from the target equilibrium density $\pi(z_0)$
 and taking one integrator cycle step (eqn 4 of [1]):

$$528 w[Z] = \ln \frac{P[Z|\Lambda]}{P[\tilde{Z}|\tilde{\Lambda}]} + \Delta f_{\text{eq}} \quad (45)$$

529 where the target equilibrium density is $\pi(z) = e^{f-h(z)}$, f is a log normalizing constant (dimensionless
 530 free energy), $h(z) \equiv u(\mathbf{x}) + t(\mathbf{v})$ is the reduced Hamiltonian, $u(\mathbf{x})$ the reduced potential, and $t(\mathbf{v})$
 531 the reduced kinetic energy [52]. The quantity $P[\tilde{Z}|\tilde{\Lambda}]$ denotes the probability of generating the
 532 time-reversed trajectory \tilde{Z} by starting with \tilde{z}_5 drawn from the target equilibrium density $\pi(\tilde{z}_5)$ (and
 533 not the nonequilibrium steady state) and applying the reverse sequence of integrator operations $\tilde{\Lambda}$,
 which is identical to the forward sequence of integrator operations Λ because the integrators we

534 consider here are symmetric. Since the Hamiltonian is time-independent, the free energy change
 535 $\Delta f_{\text{eq}} = 0$, and this simplifies to

$$w[Z] = \ln \frac{P[Z|\Lambda]}{P[\tilde{Z}|\tilde{\Lambda}]} = \ln \frac{\pi(z_0)}{\pi(\tilde{z}_5)} \frac{P[Z|z_0]}{P[\tilde{Z}|\tilde{z}_5]} \quad (46)$$

$$= \ln \frac{e^{f-h(z_0)}}{e^{f-h(\tilde{z}_5)}} + \ln \frac{P[Z|z_0]}{P[\tilde{Z}|\tilde{z}_5]} \quad (47)$$

$$= \Delta h[Z] + \ln \frac{P[Z|z_0]}{P[\tilde{Z}|\tilde{z}_5]} \quad (48)$$

536 Computation of the shadow work then proceeds by simple mechanical algebra by computing the log
 537 ratio of conditional path probabilities in the last term.

538 For the family of integrators considered here (symmetric Strang splittings of the propagator,
 539 composed of R, V, and O steps), the shadow work has an especially simple form:

$$w[Z] = \Delta h[Z] - \Delta q[Z] \quad (49)$$

540 where Δh is the total change in reduced Hamiltonian, and Δq is the total change in reduced heat across
 541 each of the O substeps. We note that accumulation of the shadow work during integration requires no
 542 extra force evaluations, and simply requires knowledge of the potential energy at the beginning and
 543 end of the integrator cycle as well as the changes in kinetic energy for each O substep.

544 For OVRVO, this is

$$\Delta h[Z] \equiv [u(\mathbf{x}_5) + t(\mathbf{v}_5)] - [u(\mathbf{x}_0) + t(\mathbf{v}_0)] \quad (50)$$

$$\Delta q[Z] \equiv [t(\mathbf{v}_1) - t(\mathbf{v}_0)] + [t(\mathbf{v}_5) - t(\mathbf{v}_4)] \quad (51)$$

545 We illustrate how to arrive at this result in detail for OVRVO below.

546 6.5. Computation of shadow work for OVRVO

547 For OVRVO, we can represent the forward and reverse of a single integrator cycle diagrammatically
 548 as

$$Z \equiv \{z_0 \xrightarrow{O} z_1 \xrightarrow{V} z_2 \xrightarrow{R} z_3 \xrightarrow{V} z_4 \xrightarrow{O} z_5\} \quad (52)$$

$$\tilde{Z} \equiv \{\tilde{z}_5 \xrightarrow{O} \tilde{z}_4 \xrightarrow{V} \tilde{z}_3 \xrightarrow{R} \tilde{z}_2 \xrightarrow{V} \tilde{z}_1 \xrightarrow{O} \tilde{z}_0\} \quad (53)$$

549 where $\tilde{z}_n = (\mathbf{x}_n, -\mathbf{v}_n)$ denotes negating the velocity of phase-space point z_n .

550 To compute the conditional path probability $P[Z|z_0]$, we write a transition probability density
 551 kernel for each substep:

$$P[Z|z_0] = K_O(z_0, z_1) K_V(z_1, z_2) K_R(z_2, z_3) K_V(z_3, z_4) K_O(z_4, z_5) \quad (54)$$

552 We can write the log ratio of conditional path probabilities as

$$\ln \frac{P[Z|z_0]}{P[\tilde{Z}|\tilde{z}_5]} = \ln \frac{K_O(z_0, z_1)}{K_O(\tilde{z}_1, \tilde{z}_0)} \frac{K_V(z_1, z_2)}{K_V(\tilde{z}_2, \tilde{z}_1)} \frac{K_R(z_2, z_3)}{K_R(\tilde{z}_3, \tilde{z}_2)} \frac{K_V(z_3, z_4)}{K_V(\tilde{z}_4, \tilde{z}_3)} \frac{K_O(z_4, z_5)}{K_O(\tilde{z}_5, \tilde{z}_4)} \quad (55)$$

553 The probability kernels K_V and K_R are both deterministic, so as long as we are considering a trajectory
 554 Z and its time-reverse \tilde{Z} generated by a symmetric integrator splitting, the ratios involving these
 555 kernels are unity.

556 To compute the ratios involving K_O kernels, we note that $K_O(z_0, z_1)$ perturbs the velocity
 557 according to the update equation

$$\mathbf{v}_1 = a_2 \mathbf{v}_0 + \sqrt{1 - a_2^2} (\beta M)^{-1/2} \xi \quad (56)$$

558 where ξ is a random variate drawn from the unit normal density, which allows us to solve for the
 559 random variate required to propagate from z_0 to z_1 ,

$$\xi = (1 - a_2^2)^{-1/2} (\beta M)^{+1/2} (\mathbf{v}_1 - a_2 \mathbf{v}_0) \quad (57)$$

$$\xi' = (1 - a_2^2)^{-1/2} (\beta M)^{+1/2} (\tilde{\mathbf{v}}_0 - a_2 \tilde{\mathbf{v}}_1) \quad (58)$$

560 where the probability density is given by

$$p(\xi) \propto e^{-\frac{1}{2}|\xi|^2}. \quad (59)$$

561 We can then rewrite the log ratio of O kernels as

$$\ln \frac{K_O(z_0, z_1)}{K_O(z_1, z_0)} = \ln \frac{p(\xi)}{p(\xi')} \quad (60)$$

$$= -\frac{\beta M}{2(1 - a_2)^2} \left[(|\mathbf{v}_1|^2 - 2a_2 \mathbf{v}_1 \cdot \mathbf{v}_0 + a_2^2 |\mathbf{v}_0|^2) - (|\mathbf{v}_0|^2 - 2a_2 \mathbf{v}_0 \cdot \mathbf{v}_1 + a_2^2 |\mathbf{v}_1|^2) \right] \quad (61)$$

$$= -\frac{\beta M}{2(1 - a_2)^2} (1 - a_2)^2 (|\mathbf{v}_1|^2 - |\mathbf{v}_0|^2) \quad (62)$$

$$= -[t(\mathbf{v}_1) - t(\mathbf{v}_0)] \quad (63)$$

562 Combining this with (49), this provides the overall work as

$$w[Z] = [u(\mathbf{x}_5) - u(\mathbf{x}_0)] + [t(\mathbf{v}_5) - t(\mathbf{v}_0)] - [t(\mathbf{v}_1) - t(\mathbf{v}_0)] - [t(\mathbf{v}_5) - t(\mathbf{v}_4)] \quad (64)$$

563 6.6. Variance-controlled adaptive estimator for KL divergence

564 As noted in Section 3.6, the nested Monte Carlo estimator we use as an expensive, but in principle
 565 asymptotically exact, estimate of the KL divergence requires converging a separate exponential average
 566 $\ln \langle e^{-w} \rangle_{\mathbf{x}, \mathbf{v}; \Lambda}$ for every sample $(\mathbf{x}, \mathbf{v}) \sim \rho$ or ω . It is obviously impossible to compute this exactly; any
 567 practical approach employing finite computational resources can only estimate this quantity to some
 568 finite statistical precision, and even then, the logarithm will induce some bias in the computed estimate
 569 for finite sample sizes. Here, we take the approach of setting a sensible target statistical error for this
 570 inner estimate, arbitrarily selecting 0.01, since we would like to resolve features in \mathcal{D}_{KL} larger than this
 571 magnitude. Notably, the difficulty in achieving this threshold increases exponentially as the width of
 572 the sampled distribution $p(w)$ increases with increasing timestep Δt .

573 To determine the number of inner-loop samples required to meet this statistical error threshold, we
 574 periodically compute an estimate of $\sigma_{\text{inner}}(\mathbf{x}, \mathbf{v}) = \text{stddev}(\ln \langle e^{-w} \rangle_{\mathbf{x}, \mathbf{v}; \Lambda})$ from the available samples
 575 using standard first-order Taylor series propagation of uncertainty. We compare the estimated standard
 576 deviation with the user-defined threshold, and continue to draw samples until we meet the threshold
 577 or exceed the sample budget of 5×10^4 samples⁸. The scaling of computational effort with Δt is shown
 578 in Figure A4.

579 Choosing the inner-loop threshold is subtle. If the inner-loop threshold is chosen too large, then
 580 the resulting estimate will be very biased (the whole procedure only becomes unbiased in the limit that

8 This limited the maximum CPU time spent collecting an “outer-loop” sample to approximately 2 hours—conditions where this limit was met are colored grey in the lower panel of Figure 7.

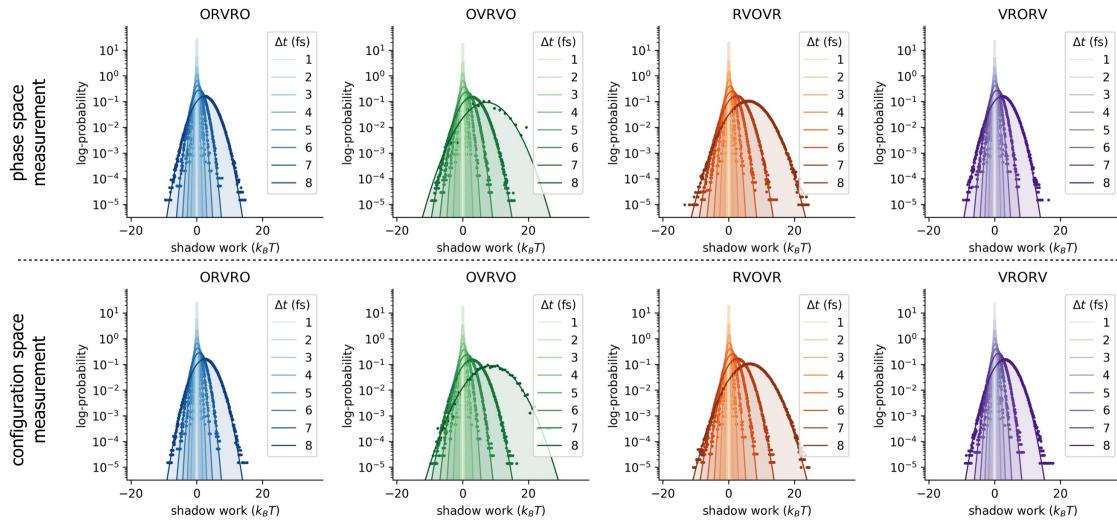


Figure A1. Shadow work distributions for the water cluster are approximately Gaussian for all integrators examined. In all panels, solid lines and shaded regions denote Gaussian fits, while dots denote histogram estimates. The **top row** depicts work distributions where initial conditions are sampled from the nonequilibrium steady-state induced by the corresponding integrator and timestep $[(\mathbf{x}, \mathbf{v}) \sim \rho]$; these shadow work values are used to measure phase-space error in the near-equilibrium estimates of \mathcal{D}_{KL} . The **bottom row** depicts work distributions where initial conditions are sampled from the ω ensemble $[\mathbf{x} \sim \rho_{\mathbf{x}}, \mathbf{v} \sim \pi(\mathbf{v}|\mathbf{x})]$; these work values are used to estimate configuration-space error in the near-equilibrium estimates of \mathcal{D}_{KL} .

581 $\sigma_{\text{inner}} \rightarrow 0$). If the threshold is chosen too small, then the computational effort becomes prohibitive.
 582 Controlling the *precision* of the inner-loop estimates should also be expected to control their *bias*, since
 583 the bias of the inner-loop estimates is approximately $\sigma_{\text{inner}}^2 / 2M$ (see Section II.B, eqn. 8 in [40]), in the
 584 direction of *under-estimating* the \mathcal{D}_{KL} .

585 To compute and report uncertainty in Figure 7, we use bootstrap resampling, rather than Taylor
 586 propagation. The data for each condition is a jagged array of sampled work values, where each row
 587 represents an “outer-loop sample” (*i.e.*, a different initial condition (\mathbf{x}, \mathbf{v}) sampled from ρ (or ω)), and
 588 the length of each row is variable, reflecting the number of “inner-loop” samples required to estimate
 589 $\ln [\rho(\mathbf{x}, \mathbf{v}) / \pi(\mathbf{x}, \mathbf{v})]$ (or $\ln [\omega(\mathbf{x}, \mathbf{v}) / \pi(\mathbf{x}, \mathbf{v})]$) to the desired precision. To generate a single bootstrap
 590 sample, we resample first the rows uniformly with replacement, and then, within each row, resample
 591 the columns uniformly with replacement. The error bands for the “exact” estimator in Figure 7 are
 592 computed from 100 bootstrap samples per condition.

593 Appendix A Statistics of shadow work distributions

594 The exact expression for the KL divergence used for validation in Section 3.6 requires estimating
 595 the expectation of e^{-w} averaged over $p(w)$. Work distributions for various integrators and timesteps
 596 are plotted in Figure A1, and appear to be approximately Gaussian, as can be seen by comparison with
 597 Gaussian fits (solid lines). As expected, the width of the work distribution increases with increasing
 598 timestep, which can be seen more clearly in Figure A2, which plots the standard deviation of the work
 599 distribution as a function of timestep.

600 While the near-equilibrium estimate will find the difficulty of reaching estimates of a given
 601 statistical precision grows linearly with the variance in $p(w)$, the exact estimator must converge
 602 expectations of the *exponentiated* shadow work, which becomes exponentially difficult with increasing
 603 variance. This effect is illustrated in Figure A3, where we plot e^{-w} along with the Gaussian fits to these
 604 work distributions.

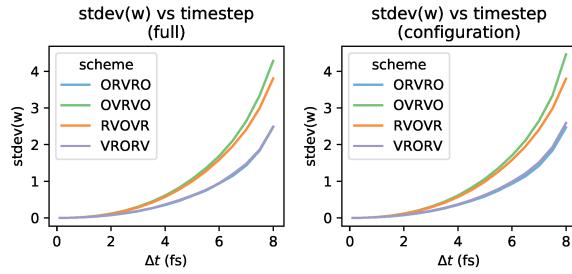


Figure A2. Standard deviation of water cluster shadow work distribution grows with time step. The left panel summarizes the top row of Figure A1 (shadow work distributions for trajectories initialized in π), and the right panel summarizes the bottom row of Figure A1 (trajectories initialized in ω).

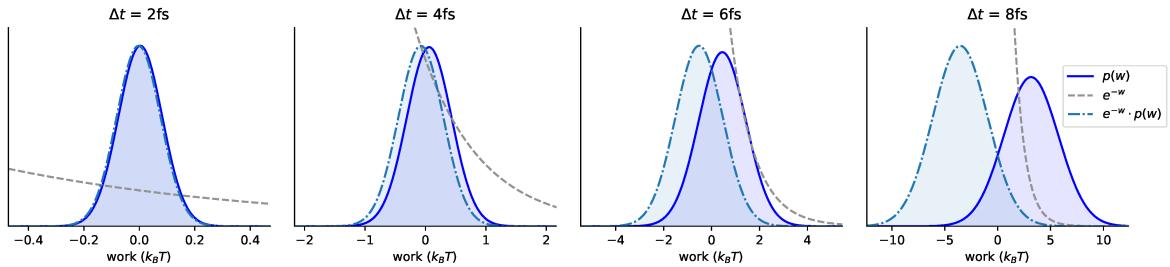


Figure A3. Exponential averages with respect to the shadow work distribution become increasingly difficult with increasing timestep. It becomes increasingly difficult to estimate the expectation of e^{-w} with respect to Gaussian fits to work distributions $p(w)$ for the water cluster, as the timestep Δt increases. The four panels increase in Δt from left to right: note the changing x -axis scales. The **solid line** is the shadow work distribution $p(w)$ measured at each timestep. The **dashed line** is e^{-w} . The **dash-dotted line** is $e^{-w} \cdot p(w)$.

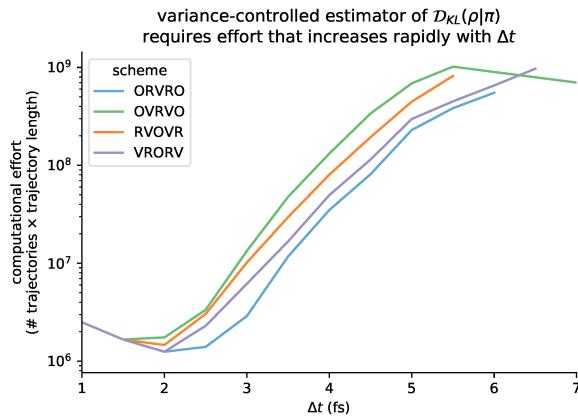


Figure A4. The computational effort required to reach a fixed uncertainty threshold depends sharply on Δt . The total computational effort is determined by the total number of trajectories sampled in that condition (i.e. $\sum_{i=1}^N M_i$), multiplied by the length of each trajectory. Note that the curves are not monotonic, since the number of required trajectories increases superlinearly with Δt , but the number of timesteps in each trajectory decreases linearly with Δt .

605 This is reflected in the computational effort required by the variance-controlled estimator,
 606 illustrated in Figure A4.

607 Appendix B Log-scale plots

608 To better visualize agreement between near-eq and nested estimates at small timesteps, we also
 609 show a version of Figure 7 with a log-scale y -axis in Figure A5.

610 Appendix C Further comments on the collision rate

611 If we are interested only in configuration-sampling, the collision rate is a free parameter. However,
 612 the collision rate affects the sampled kinetics. Depending on the application, we may choose a
 613 value of the collision rate that reproduces kinetic observables such as diffusion constants, or we
 614 may choose a value of the collision rate that minimizes correlation times. In the over-damped limit
 615 (large γ) conformational search is significantly hampered, as the trajectories become diffusive. In the
 616 under-damped limit (small γ) trajectories are not random walks, but mixing between energy levels is
 617 slower. (At the extreme, $\gamma = 0$, Langevin dynamics reduces to Hamiltonian dynamics, and explores a
 618 single constant-energy ensemble.)

619 We do not provide any guidance here on the choice of the collision rate, but we investigated how
 620 sensitive the measured D_{KL} is to the choice of γ for each integrator. To supplement the figure in the
 621 main text (Figure 6), we plot the same information again, but grouped by γ rather than by scheme in
 622 Figure A6.

623 As noted above, the accuracy of the proposed estimator hinges on the assumption that the system
 624 has been driven into steady-state during the finite-length protocol. To validate that protocols of length
 625 2000 steps were sufficient, we compare results for protocols that are 1000 steps long and 2000 steps
 626 long, and note that they produce very similar estimates for timestep-dependent error for the range
 627 of collision rates examined. The results reported in Figure 6 and Figure A6 were performed using
 628 a protocol of length $T = 2000$ steps, two times longer than the protocols used in Figure 5. We also
 629 computed near-equilibrium estimates over the same range of collision rates using a shorter protocol
 630 of length of $T = 1000$ steps Figure A7. Note that, for computational convenience, we have used
 631 substantially fewer protocol samples to compute the estimates in Figure A7 than were used to compute
 632 the results reported in the main text (10 000 samples, rather than 50 000 samples).

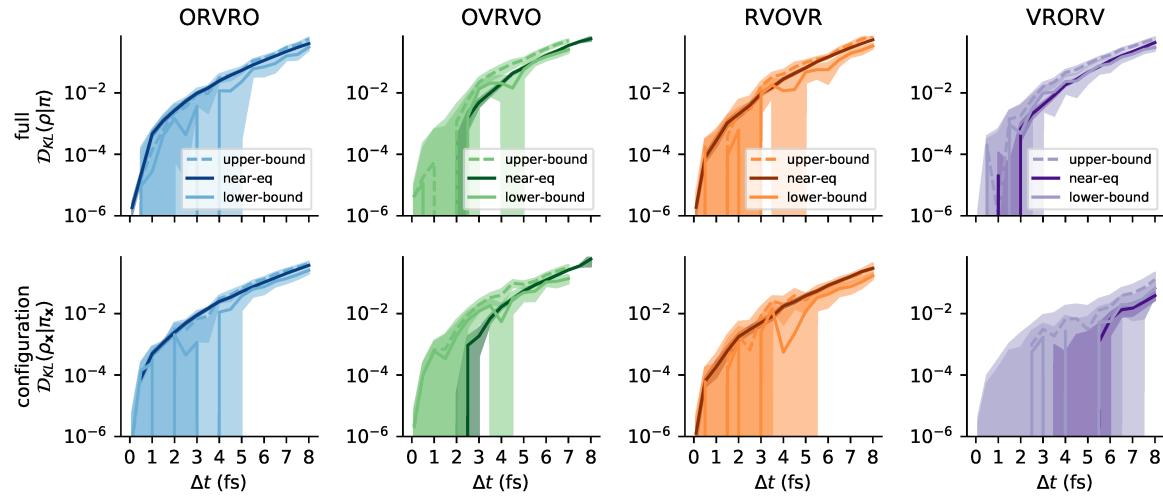


Figure A5. Agreement between near-equilibrium and reference methods may be better visualized at small timesteps on a log-scale. This is a repeat of Figure 7, but with a log-scale y -axis. Vertical artifacts are caused by the estimate becoming transiently negative, which cannot be represented on the logarithmic scale.

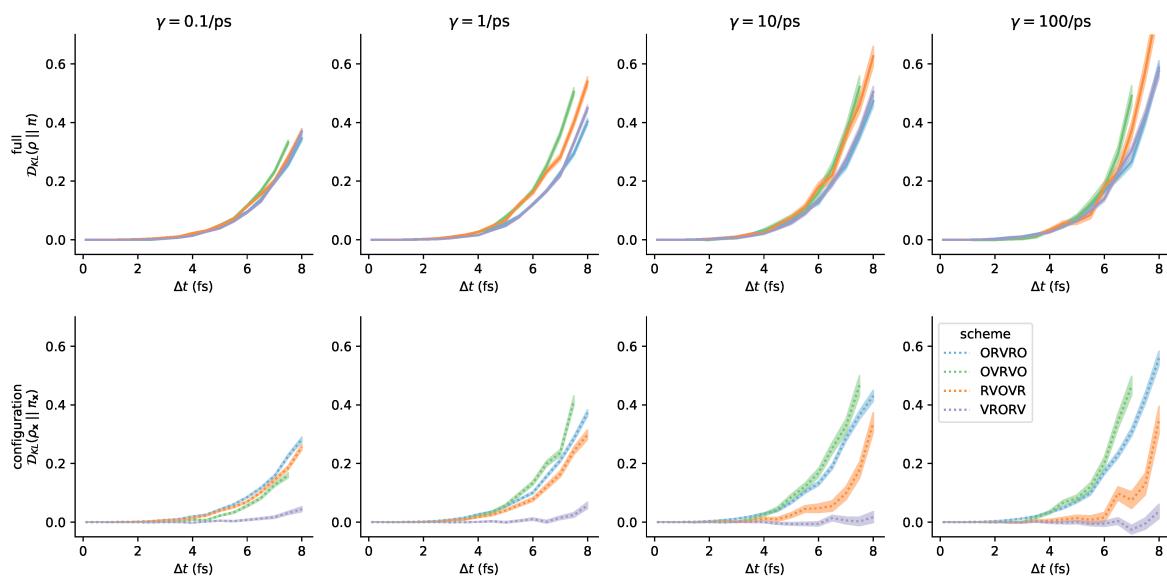


Figure A6. Integrator accuracy can be compared at different collision rates. Here we have plotted the same results as in Figure 6 but grouped by the collision rate γ rather than by scheme. Note that the ordering over schemes induced by configuration-space error can vary as a function of collision rate—*i.e.* at the lowest measured γ , the ordering from lowest to highest error is (1) VRORV, (2) OVRVO, (3) RVOVR, (4) ORVRO, but at the highest measured γ , the ordering is different.

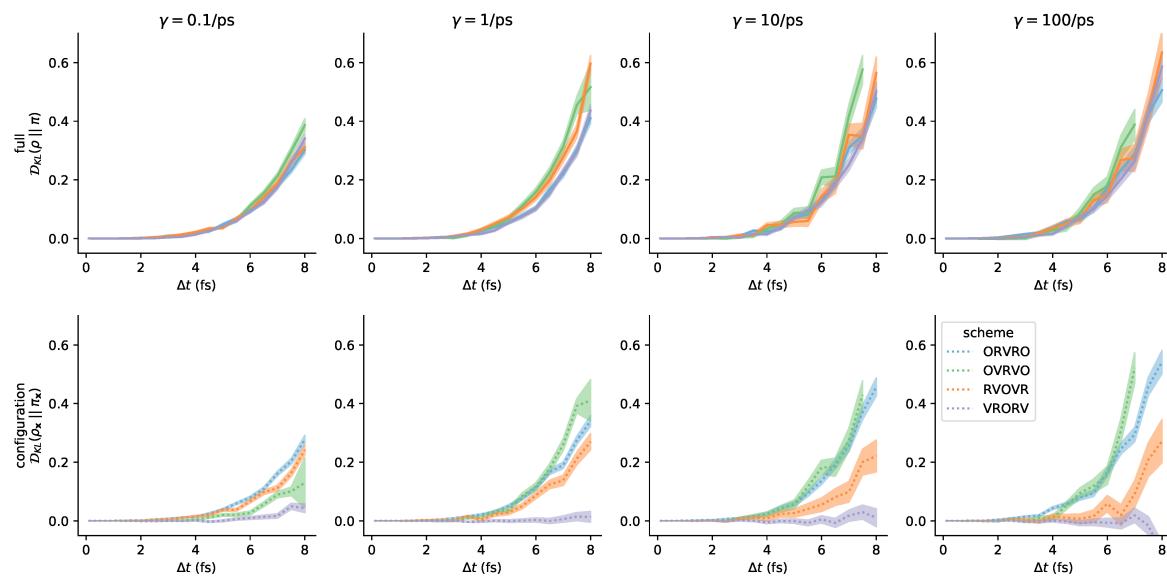


Figure A7. Similar results are obtained for 1000-step protocols as for 2000-step protocols. For the same wide array of conditions reported in Figure A6, near-equilibrium estimates generated using 5x fewer samples of 2x shorter protocols are broadly consistent with the results obtained by more exhaustive sampling.

Acknowledgments: The authors gratefully acknowledge the helpful suggestions of Grant Rotskoff (ORCID: 0000-0002-7772-5179) for his input on nonequilibrium shadow work measurement schemes; Charles Matthews (ORCID: 0000-0001-9545-1346) and Brian Radak (ORCID: 0000-0001-8628-5972) for discussions on geodesic integrators; Jason Wagoner (ORCID: 0000-0003-2359-9696) for discussions about Metropolized Langevin integrators; and members of the Chodera laboratory for input on various aspects of the implementation and theory. We thank Rodrigo Guerra (ORCID: 0000-0003-3082-0248) for helpful comments on continuous barostats. The simulations were made possible by the High Performance Computing Group at MSKCC. BJL was supported by EPSRC grant no. EP/P006175/1 and the Alan Turing Institute. JDC acknowledges support from the Sloan Kettering Institute and NIH grant P30 CA008748 and NIH grant R01 GM121505. JF acknowledges support from NSF grant CHE 1738979. DAS acknowledges support from a Natural Sciences and Engineering Research Council of Canada (NSERC) Discovery Grant and the Canada Research Chairs program. We are especially grateful to Peter Eastman (ORCID: 0000-0002-9566-9684) for implementing the *CustomIntegrator* facility within OpenMM that greatly simplifies implementation of the integrators considered here, and for helpful feedback on the manuscript. Finally, we thank the three anonymous reviewers for their many thoughtful comments that greatly improved the manuscript, for identifying relevant prior work we had missed, and for suggesting additional experiments to isolate the effect of the collision rate.

1. Sivak, D.A.; Chodera, J.D.; Crooks, G.E. Using Nonequilibrium Fluctuation Theorems to Understand and Correct Errors in Equilibrium and Nonequilibrium Simulations of Discrete Langevin Dynamics. *Physical Review X* **2013**, *3*.
2. Lemons, D.S.; Gythiel, A. Paul Langevin's 1908 paper "On the Theory of Brownian Motion" ["Sur la théorie du mouvement brownien," C. R. Acad. Sci. (Paris) 146, 530-533 (1908)]. *American Journal of Physics* **1997**, *65*, 1079–1081.
3. Lelièvre, T.; Stoltz, G.; Rousset, M. *Free Energy Computations: A Mathematical Perspective*; Imperial College Press: London ; Hackensack, N.J, 2010. OCLC: ocn244765923.
4. Leimkuhler, B.; Matthews, C. *Molecular Dynamics: With Deterministic and Stochastic Numerical Methods*; Springer: Cham, 2015. OCLC: 914391557.
5. Frenkel, D.; Smit, B. *Understanding Molecular Simulation*, 2nd ed.; Academic Press, Inc.: Orlando, FL, USA, 2001.
6. Maruyama, G. Continuous Markov Processes and Stochastic Equations. *Rendiconti del Circolo Matematico di Palermo* **1955**, *4*, 48–90.

664 7. Ermak, D.L.; Yeh, Y. Equilibrium Electrostatic Effects on the Behavior of Polyions in Solution:
665 Polyion-Mobile Ion Interaction. *Chemical Physics Letters* **1974**, *24*, 243–248.

666 8. Brünger, A.; Brooks, C.L.; Karplus, M. Stochastic Boundary Conditions for Molecular Dynamics Simulations
667 of ST2 Water. *Chemical Physics Letters* **1984**, *105*, 495–500.

668 9. Pastor, R.W.; Brooks, B.R.; Szabo, A. An Analysis of the Accuracy of Langevin and Molecular Dynamics
669 Algorithms. *Molecular Physics* **1988**, *65*, 1409–1419.

670 10. Athènes, M. A Path-Sampling Scheme for Computing Thermodynamic Properties of a Many-Body System
671 in a Generalized Ensemble. *The European Physical Journal B - Condensed Matter and Complex Systems* **2004**,
672 *38*, 651–663.

673 11. Adjanor, G.; Athènes, M.; Calvo, F. Free Energy Landscape from Path-Sampling: Application to the
674 Structural Transition in LJ38. *The European Physical Journal B - Condensed Matter and Complex Systems* **2006**,
675 *53*, 47–60.

676 12. Bussi, G.; Parrinello, M. Accurate Sampling Using Langevin Dynamics. *Physical Review E* **2007**, *75*, 056707.

677 13. Izaguirre, J.A.; Sweet, C.R.; Pande, V.S. Multiscale Dynamics of Macromolecules Using Normal Mode
678 Langevin. In *Biocomputing 2010*; World Scientific, 2009; pp. 240–251.

679 14. Leimkuhler, B.; Matthews, C. Robust and Efficient Configurational Molecular Sampling via Langevin
680 Dynamics. *The Journal of Chemical Physics* **2013**, *138*, 174102.

681 15. Leimkuhler, B.; Matthews, C. Efficient Molecular Dynamics Using Geodesic Integration and Solvent–solute
682 Splitting. *Proc. R. Soc. A* **2016**, *472*, 20160138.

683 16. Grubmüller, H.; Tavan, P. Multiple Time Step Algorithms for Molecular Dynamics Simulations of Proteins:
684 How Good Are They? *Journal of Computational Chemistry* **1998**, *19*, 1534–1552.

685 17. Hopkins, C.W.; Le Grand, S.; Walker, R.C.; Roitberg, A.E. Long-Time-Step Molecular Dynamics through
686 Hydrogen Mass Repartitioning. *Journal of Chemical Theory and Computation* **2015**, *11*, 1864–1874.

687 18. Butler, B.D.; Ayton, G.; Jepps, O.G.; Evans, D.J. Configurational Temperature: Verification of Monte Carlo
688 Simulations. *The Journal of Chemical Physics* **1998**, *109*, 6519–6522.

689 19. Leimkuhler, B.; Matthews, C. Efficient Molecular Dynamics Using Geodesic Integration and Solvent–solute
690 Splitting. *Proc. R. Soc. A* **2016**, *472*, 20160138.

691 20. Leimkuhler, B.; Matthews, C. Rational Construction of Stochastic Numerical Methods for Molecular
692 Sampling. *Applied Mathematics Research eXpress* **2013**, *2013*, 34–56.

693 21. Sweet, C.R.; Hampton, S.S.; Skeel, R.D.; Izaguirre, J.A. A Separable Shadow Hamiltonian Hybrid Monte
694 Carlo Method. *The Journal of Chemical Physics* **2009**, *131*, 174106.

695 22. Jorgensen, W.L.; Chandrasekhar, J.; Madura, J.D.; Impey, R.W.; Klein, M.L. Comparison of Simple Potential
696 Functions for Simulating Liquid Water. *J. Chem. Phys.* **1983**, *79*, 926–935.

697 23. Sivak, D.A.; Chodera, J.D.; Crooks, G.E. Using Nonequilibrium Fluctuation Theorems to Understand and
698 Correct Errors in Equilibrium and Nonequilibrium Simulations of Discrete Langevin Dynamics. *Physical
699 Review X* **2013**, *3*, 001007+.

700 24. Sivak, D.A.; Chodera, J.D.; Crooks, G.E. Time Step Rescaling Recovers Continuous-Time Dynamical
701 Properties for Discrete-Time Langevin Integration of Nonequilibrium Systems. *The Journal of Physical
702 Chemistry B* **2014**, *118*, 6466–6474.

703 25. Melchionna, S. Design of quasisymplectic propagators for Langevin dynamics. *J. Chem. Phys.* **2007**,
704 *127*, 044108.

705 26. Skeel, R.D.; Izaguirre, J.A. An impulse integrator for Langevin dynamics. *Molecular Physics* **2002**,
706 *100*, 3885–3891, [<http://www.tandfonline.com/doi/pdf/10.1080/0026897021000018321>].

707 27. Serrano, M.; De Fabritiis, G.; Espanol, P.; Coveney, P. A stochastic Trotter integration scheme for dissipative
708 particle dynamics. *J. Maths and Comp. in Sim.* **2006**, *72*, 190–194.

709 28. Thalmann, F.; Farago, J. Trotter derivation of algorithms for Brownian and dissipative particle dynamics. *J.
710 Chem. Phys.* **2007**, *127*, 124109.

711 29. Leimkuhler, B.; Matthews, C.; Stoltz, G. The Computation of Averages from Equilibrium and
712 Nonequilibrium Langevin Molecular Dynamics. *IMA Journal of Numerical Analysis* **2016**, *36*, 13–79.

713 30. Swope, W.C.; Andersen, H.C.; Berens, P.H.; Wilson, K.R. A Computer Simulation Method for the
714 Calculation of Equilibrium Constants for the Formation of Physical Clusters of Molecules: Application to
715 Small Water Clusters. *The Journal of Chemical Physics* **1982**, *76*, 637–649.

716 31. Leimkuhler, B.; Matthews, C. *Molecular Dynamics: With Deterministic and Stochastic Numerical Methods*; Springer: Cham, 2015. OCLC: 914391557.

717 32. Leimkuhler, B.; Matthews, C. Rational Construction of Stochastic Numerical Methods for Molecular Sampling. *Appl. Math. Res. Express* **2012**, [\[1203.5428\]](#).

718 33. Bennett, C.H. Mass Tensor Molecular Dynamics. *Journal of Computational Physics* **1975**, *19*, 267–279.

719 34. Pomès, R.; McCammon, J.A. Mass and Step Length Optimization for the Calculation of Equilibrium Properties by Molecular Dynamics Simulation. *Chemical Physics Letters* **1990**, *166*, 425–428.

720 35. Plecháč, P.; Rousset, M. Implicit Mass-Matrix Penalization of Hamiltonian Dynamics with Application to Exact Sampling of Stiff Systems. *Multiscale Modeling & Simulation* **2010**, *8*, 498–539.

721 36. Sivak, D.A.; Crooks, G.E. Near-Equilibrium Measurements of Nonequilibrium Free Energy. *Physical Review Letters* **2012**, *108*, 150601.

722 37. Perez-Cruz, F. Kullback-Leibler Divergence Estimation of Continuous Distributions. 2008 IEEE International Symposium on Information Theory, 2008, pp. 1666–1670.

723 38. Dhabal, D.; Nguyen, A.H.; Singh, M.; Khatua, P.; Molinero, V.; Bandyopadhyay, S.; Chakravarty, C. Excess Entropy and Crystallization in Stillinger-Weber and Lennard-Jones Fluids. *The Journal of Chemical Physics* **2015**, *143*, 164512.

724 39. Athènes, M.; Adjancor, G. Measurement of Nonequilibrium Entropy from Space-Time Thermodynamic Integration. *The Journal of Chemical Physics* **2008**, *129*, 024116.

725 40. Shirts, M.R.; Pande, V.S. Comparison of Efficiency and Bias of Free Energies Computed by Exponential Averaging, the Bennett Acceptance Ratio, and Thermodynamic Integration. *The Journal of Chemical Physics* **2005**, *122*, 144107.

726 41. Campos, C.M.; Sanz-Serna, J.M. Extra Chance Generalized Hybrid Monte Carlo. *Journal of Computational Physics* **2015**, *281*, 365–374, [\[1407.8107\]](#).

727 42. Athènes, M.; Marinica, M.C. Free Energy Reconstruction from Steered Dynamics without Post-Processing. *Journal of Computational Physics* **2010**, *229*, 7129–7146.

728 43. Wagoner, J.A.; Pande, V.S. Reducing the Effect of Metropolization on Mixing Times in Molecular Dynamics Simulations. *The Journal of Chemical Physics* **2012**, *137*, 214105, [\[1209.5944\]](#).

729 44. Nilmeier, J.P.; Crooks, G.E.; Minh, D.D.L.; Chodera, J.D. Nonequilibrium Candidate Monte Carlo Is an Efficient Tool for Equilibrium Simulation. *PNAS* **2011**, *108*, E1009–E1018.

730 45. Eastman, P.; Swails, J.; Chodera, J.D.; McGibbon, R.T.; Zhao, Y.; Beauchamp, K.A.; Wang, L.P.; Simmonett, A.C.; Harrigan, M.P.; Stern, C.D.; Wiewiora, R.P.; Brooks, B.R.; Pande, V.S. OpenMM 7: Rapid development of high performance algorithms for molecular dynamics. *PLOS Computational Biology* **2017**, *13*, 1–17.

731 46. Lam, S.K.; Pitrou, A.; Seibert, S. Numba: A LLVM-Based Python JIT Compiler. Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC; ACM: New York, NY, USA, 2015; LLVM '15, pp. 7:1–7:6.

732 47. Chodera, J.; Rizzi, A.; Naden, L.; Beauchamp, K.; Grinaway, P.; Fass, J.; Rustenburg, B.; Ross, G.A.; Simmonett, A.; Swenson, D.W. Openmmtools: 0.14.0 - Exact Treatment of Alchemical PME Electrostatics, Water Cluster Test System, Optimizations.

733 48. Eastman, P.; Swails, J.; Chodera, J.D.; McGibbon, R.T.; Zhao, Y.; Beauchamp, K.A.; Wang, L.P.; Simmonett, A.C.; Harrigan, M.P.; Stern, C.D.; Wiewiora, R.P.; Brooks, B.R.; Pande, V.S. OpenMM 7: Rapid Development of High Performance Algorithms for Molecular Dynamics. *PLOS Computational Biology* **2017**, *13*, e1005659.

734 49. Akhmatkaya, E.; Bou-Rabee, N.; Reich, S. Erratum to “A Comparison of Generalized Hybrid Monte Carlo Methods with and without Momentum Flip” [J. Comput. Phys. 228 (2009) 2256–2265]. *Journal of Computational Physics* **2009**, *228*, 7492–7496.

735 50. Adjancor, G.; Athènes, M. Gibbs Free-Energy Estimates from Direct Path-Sampling Computations. *The Journal of Chemical Physics* **2005**, *123*, 234104.

736 51. Adib, A.B. Comment on “On the Crooks fluctuation theorem and the Jarzynski equality” [J. Chem. Phys. 129, 091101 (2008)]. *The Journal of Chemical Physics* **2009**, *130*, 247101.

737 52. Shirts, M.R.; Chodera, J.D. Statistically Optimal Analysis of Samples from Multiple Equilibrium States. *J. Chem. Phys.* **2008**, *129*, 124105.