

MARGINAL GIRSANOV REWEIGHTING: STABLE VARIANCE REDUCTION VIA NEURAL RATIO ESTIMATION

Anonymous authors

Paper under double-blind review

ABSTRACT

Recovering unbiased properties from biased or perturbed simulations is a central challenge in rare-event sampling. Classical Girsanov Reweighting (GR) offers a principled solution by yielding exact pathwise probability ratios between perturbed and reference processes. However, the variance of GR weights grows rapidly with time, rendering it impractical for long-horizon reweighting. We introduce Marginal Girsanov Reweighting (MGR), which mitigates variance explosion by marginalizing over intermediate paths, producing stable and scalable weights for long-timescale dynamics. Experiments demonstrate that MGR (i) accurately recovers kinetic properties from umbrella-sampling trajectories in molecular dynamics, and (ii) enables efficient Bayesian parameter inference for stochastic differential equations with temporally sparse observations.

1 INTRODUCTION

Estimating expectations under complex stochastic systems is central to modern machine learning in the sciences and engineering application. Direct sampling from the target law is often infeasible, so reweighting techniques express these expectations as averages under a more tractable reference law. Indeed, these reweighting strategies are fundamental tools in computational chemistry (Mey et al., 2014), mathematics (Beskos & Roberts, 2005), finance (Pascucci, 2011), and machine learning (Domke & Sheldon, 2018). Unfortunately, their practical use is often limited by estimator variance.

This reweighting principle is important in the analysis of stochastic dynamics, where probability laws are defined over path space (Sørensen, 2004; Donati et al., 2017). For example, in molecular dynamics (MD), perturbed potentials are used to overcome free energy barriers [due to slow mixing or interest in rare events](#), and reweighting is then required to recover unbiased quantities such as free energies and transition rates (Torrie & Valleau, 1977; Kästner, 2011; Mey et al., 2014). Similarly, in parameter inference for stochastic differential equations (SDEs), simulations are often carried out under a reference parameter, with likelihood ratios reweighting alternative candidates to form posterior distributions (Li et al., 2020; Ghosh et al., 2022). Across both settings, a common way to compute the importance weights is provided by Girsanov’s theorem (Girsanov, 1960).

Girsanov reweighting (GR) (Girsanov, 1960) provides a pathwise ratio, which enables a change of measure between SDEs that share the diffusion but differ in drift (Sørensen, 2004; Donati et al., 2017; 2022). Despite this elegance, two practical limitations arise. First, weight variance grows rapidly with time horizon, causing the exponentiated weights to explode or vanish in long timescales. Second, the weights are inherently tied to specific trajectories. Yet many quantities of interest—such as MD transition probabilities or SDE likelihoods—depend only on endpoints and require marginal ratios that integrate over all intermediate paths. Together, these issues severely restrict the use of Girsanov reweighting in long timescales and large systems.

To address these limitations, we propose a machine learning-based approach: **Marginal Girsanov Reweighting (MGR)**. Instead of relying on full-path Girsanov weights, which become numerically unstable for long trajectories, MGR learns marginal density ratios between end-points of trajectories. The key idea is to leverage accurate short-lag Girsanov weights and iteratively compose them into longer-lag ratios using neural classifiers. We formulate ratio estimation as a binary classification problem, where a neural network distinguishes between samples from reference and target distributions (Menon & Ong, 2016; Choi et al., 2021). By combining the mathematical foundation of

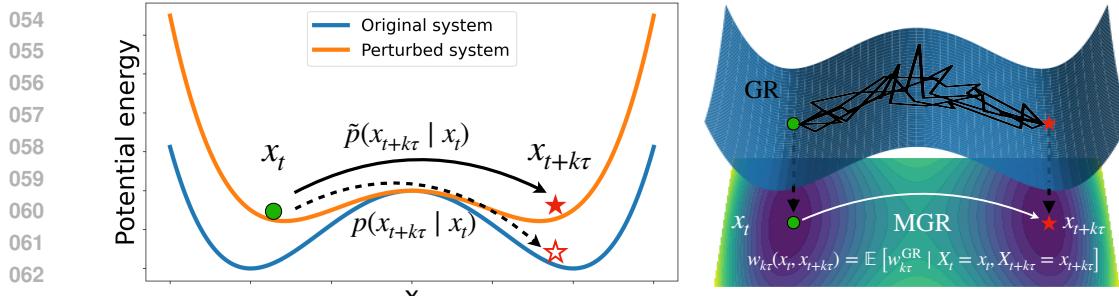


Figure 1: Marginal Girsanov Reweighting (MGR). For the given pairs, MGR defines the marginal weight as expectation of pathwise Girsanov reweighting (GR) factors as introduced in Section 4.1.

Girsanov reweighting with the flexibility of neural ratio estimation, MGR enables reliable inference across domains where traditional estimators break down. To summarize:

- We propose Marginal Girsanov Reweighting (MGR), an iterative learning approach, which estimates transition-based density ratios by marginalizing over intermediate paths.
- We implement MGR as a binary classification task with a weighted cross-entropy objective.
- We demonstrate the effectiveness of MGR on reweighting problems in two distinct domains (i) recovering unbiased thermodynamic and kinetic properties from biased MD simulations, and (ii) parameter posterior inference for SDEs from noisy observations.

2 RELATED WORK

Accelerated sampling via perturbations In many physical and molecular systems, direct sampling from the target dynamics is computationally infeasible due to slow mixing or the rarity of the events of interest. This has motivated a broad class of perturbation-based accelerated sampling methods that modify the dynamics to facilitate exploration. In molecular dynamics, metadynamics (Huber et al., 1994; Barducci et al., 2008) and umbrella sampling (Torrie & Valleau, 1977; Kästner, 2011) add additional energy (Grumbüller, 1995) to reaction coordinates in order to accelerate exploration. (Swendsen & Wang, 1986; Sugita & Okamoto, 1999; Wang et al., 2011) change the temperature of the system to help cross barriers.

Reweighting of Dynamics and Path Sampling Reweighting is a necessary technique to recover the original dynamic from the accelerated data (Kamenik et al., 2022). For thermodynamic quantities, methods such as the weighted histogram analysis method (WHAM) (Gallicchio et al., 2005; Souaille & Roux, 2001) and the multistate Bennett acceptance ratio (MBAR) (Shirts & Chodera, 2008), provide efficient estimators of equilibrium energies and have been applied to multiple windows enhanced sampling. Recent work also (Dibak et al., 2022; Wang et al., 2022; Moqvist et al., 2025; Invernizzi et al., 2022) uses machine learning to estimate the energy differences between different temperatures.

For kinetics, finite-lag transition densities help characterize long time behavior, where intermediate states can be ignored (Wu et al., 2017; Schreiner et al., 2023; Klein et al., 2023; Diez et al., 2024; 2025). With enhanced-simulation data, unbiased kinetics can be recovered by: combining multiple equilibrium ensembles (Mey et al., 2014) within a reversible Markov-state-model framework (Prinz et al., 2011; Husic & Pande, 2018), or directly learning the eigenfunctions of transfer operators from the biased simulations (Devergne et al., 2024). Recently, Girsanov-based path reweighting, which explicitly calculates the Radon–Nikodym derivative in configuration space (Donati et al., 2017; 2022; Schafer & Keller, 2024), has drawn our attention.

Bayesian inference for SDEs In classical Bayesian inference for parametric SDEs, the drift and diffusion structures are given analytically, but contain unknown parameters. With only discrete and sparsely observed data, the likelihood of these observations is analytically intractable. A common strategy is to simulate latent paths under a reference SDE that is numerically stable or computationally convenient, and then express the likelihood ratio relative to the target SDE via Girsanov’s theorem (Girsanov, 1960; Sørensen, 2004). This approach underlies a wide range of inference

108 algorithms, including particle MCMC (Andrieu et al., 2010), and series-expansion-based inference (Lyons et al., 2014; Ghosh et al., 2022). It has wide applications in life sciences (Fuchs, 2013; 109 Bunin, 2017) and financial modeling (Jones, 1998; Eraker, 2001). Recent advances in simulation-based inference (SBI) provide neural posterior by amortizing inference over large collections of 110 simulated data pairs (Cranmer et al., 2020; Gloeckler et al., 2024; Cai et al., 2023). These methods 111 require extensive simulation across a broad parameter range, and aim to learn a surrogate and 112 black-box simulator. In our paper, we focus on the former—classical Bayesian inference—which 113 operates in the physics-informed regime where the underlying SDE mechanism is known. The main 114 challenge arises from the high variance of path-space likelihood ratios. 115

3 PRELIMINARIES

3.1 BROWNIAN DYNAMICS AND MIXING

121 A diffusion process described by a stochastic differential equation (SDE) satisfies

$$122 \quad dX_t = f(X_t, t)dt + g(t) dW_t, \quad X_0 = x_0, \quad (1)$$

124 where $X_t \in \mathbb{R}^d$ denotes the state of the stochastic process at time t , $f(\cdot, t) : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is a drift 125 vector, $g(t) \in \mathbb{R}$ is the diffusion coefficient, and $W_t \in \mathbb{R}^d$ is a standard Wiener process. In a fixed 126 time horizon $t \in [0, T]$, the associated path probability measure induced by Eq. 1 is denoted by μ .

127 However, direct simulation from the target law is often infeasible (Sørensen, 2004; Vanden-Eijnden 128 et al., 2010), such as slow mixing in molecular dynamics and unknown parameters in Bayesian 129 inference, as illustrated in Section 2. This motivates us to introduce a perturbation to the drift, 130 replacing $f(\cdot, t)$ with $\tilde{f}(\cdot, t)$, while keeping the diffusion coefficient $g(t)$ unchanged. The associated 131 path probability measure is denoted by $\tilde{\mu}$.

132 Although such perturbed dynamics are easier to simulate, they alter the underlying probability 133 measure and thus distort the statistics of the original system. By calculating the Radon-Nikodym derivative 134 $\frac{d\mu}{d\tilde{\mu}}(\mathbf{x})$ between two path measures, we can analyze the original dynamics in Eq. 1 by reweighting 135 the trajectories $\mathbf{x} = \{x_t\}_{t \in [0, T]}$ under the perturbed paths. 136

3.2 GIRSANOV REWEIGHTING THEORY

139 A common approach to compute the Radon–Nikodym derivative of μ respect to $\tilde{\mu}$ is provided by 140 Girsanov’s theorem (Girsanov, 1960; Donati et al., 2017). If two diffusion processes share the same 141 diffusion coefficient but have different drifts, their path measures can be transformed from one to 142 the other.

143 Here, we consider the trajectory segment under the perturbed dynamics, $\mathbf{x}_{t, \tau} = \{x_s\}_{s=t}^{t+\tau}$ from x_t 144 to $x_{t+\tau}$, $t \in [0, T - \tau]$. Using an Euler–Maruyama discretization, the corresponding discrete-time 145 trajectory $\{x_t = x^0, x^1, x^2, \dots, x_{t+\tau} = x^N\}$ is observed with discretization step $\Delta t = \tau/N$. Then 146 the likelihood ratio between the original path μ and the perturbed path $\tilde{\mu}$, conditional on the same 147 starting state x_t , can be calculated as 148

$$149 \quad \log w_\tau^{\text{GR}}(\mathbf{x}_{t, \tau}) = \log \frac{d\mu}{d\tilde{\mu}}(\mathbf{x}_{t, \tau} | x_t) \\ 150 \\ 151 \quad \approx \sum_{k=0}^{N-1} \left(\frac{(f(x^k, t^k) - \tilde{f}(x^k, t^k))^{\top}}{g(t^k)} \sqrt{\Delta t} \xi^k - \frac{\Delta t}{2} \left\| \frac{(f(x^k, t^k) - \tilde{f}(x^k, t^k))^{\top}}{g(t^k)} \right\|^2 \right), \quad (2)$$

154 where $\sqrt{\Delta t} \xi^k = \frac{x^{k+1} - x^k - \tilde{f}(x^k, t^k) \Delta t}{g(t^k)}$ $\stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, \Delta t \cdot I_d)$ represents the Wiener increment 155 associated with the simulation step $x^k \rightarrow x^{k+1}$ under the perturbed dynamics. 156

158 Eq. 2 provides a solution for recovering original properties from perturbed simulations. However, 159 two major difficulties arise in practice. First, in computing $\log w_\tau^{\text{GR}}$, long time horizons τ will 160 introduce an accumulation of noise terms, e.g. ξ^k in Eq. 2, which causes the variance to grow with 161 trajectory length. Upon exponentiation, the weights will explode or vanish. Detailed analysis can be 162 found in Appendix A. Second, when comparing the τ -lag transition probabilities from x_t to $x_{t+\tau}$

under the perturbed and original dynamics, Girsanov reweighting yields a ratio along a specific simulated path, which is a biased proxy for the true marginal quantity (see Section 4.1 for details). Motivated by these two problems, our work develops a model that can estimate a more stable and accurate marginal weights for long time intervals and large systems.

4 MARGINAL GIRSANOV REWEIGHTING

4.1 SETUP

In stochastic dynamics, probability laws are typically defined over path space, capturing either a single long trajectory or multiple trajectories $\mathbf{x} = \{x_t\}_{t \in [0, T]}$ under a given stochastic process. However, in many practical applications, our focus shifts from full paths to the paired data $\{(x_t, x_{t+\tau})\}_{t \in [0, T-\tau]}$ between two time points separated by a fixed lag time τ . **Under the standard ergodic assumption, the transition-based representation provides meaningful dynamical information.**

For example, in molecular dynamics (MD), the analysis of kinetics and transport properties (Prinz et al., 2011; Tiwary et al., 2015; Mardt et al., 2018) is often carried out. The central object is then the transition probability between metastable states over a specified lag time. Similarly, in SDE parameter inference, discrete-time observations are available at sparse intervals of time horizons, and inference relies on computing the likelihood of these observation pairs (Sørensen, 2004; Golightly & Wilkinson, 2008). In these contexts, the object of interest is not the full path law $\mu(\mathbf{x})$, but rather the finite-time transition probability $p(x_{t+\tau} | x_t)$ induced at two time points.

Let $p(x_{t+\tau} | x_t)$ denote the transition probability induced by the original process in Eq. 1, and let $\tilde{p}(x_{t+\tau} | x_t)$ denote the transition probability under the perturbed dynamics. The ratio between these two transition quantities defines a reweighting factor $w_\tau(x_t, x_{t+\tau}) = \frac{p(x_{t+\tau} | x_t)}{\tilde{p}(x_{t+\tau} | x_t)}$.

Instead of computing the intractable conditional density directly, we consider the joint distribution over such pairs. Letting $\rho_\tau(x_t, x_{t+\tau})$ and $\tilde{\rho}_\tau(x_t, x_{t+\tau})$ denote the joint densities under the original and perturbed processes respectively, the reweighting ratio can be expressed as

$$w_\tau(x_t, x_{t+\tau}) = \frac{p(x_{t+\tau} | x_t) \rho(x_t)}{\tilde{p}(x_{t+\tau} | x_t) \rho(x_t)} = \frac{\rho_\tau(x_t, x_{t+\tau})}{\tilde{\rho}_\tau(x_t, x_{t+\tau})}.$$

This representation highlights that the ratio of transition density ratio can be seen as the ratio of joint distributions. Furthermore, it admits a natural transformation in terms of the Radon–Nikodym derivative over path space

$$w_\tau(x_t, x_{t+\tau}) = \mathbb{E}_{\tilde{\mu}} \left[\frac{d\mu}{d\tilde{\mu}}(\mathbf{x}_{t,\tau}) \mid X_t = x_t, X_{t+\tau} = x_{t+\tau} \right]. \quad (3)$$

Detailed proof can be found in Appendix B. Based on the pathwise Girsanov weights $\frac{d\mu}{d\tilde{\mu}}(\mathbf{x}_{t,\tau})$, we propose the Marginal Girsanov Reweighting (MGR) to learn the reweighting factor for given paired data with long time interval (see Figure 1).

4.2 TRAINING ALGORITHM

We denote by $w_{k\tau}(x_t, x_{t+k\tau})$ the likelihood ratio from configuration x_t at time t to $x_{t+k\tau}$ at time $t + k\tau$. For the case $k = 1$, i.e., over a short lag time, Girsanov reweighting w_τ^{GR} in Eq. 2 provides a relatively stable and pathwise estimate of the ratio. Our goal is to develop a Marginal Girsanov Reweighting (MGR) approach that can reliably estimate the ratio for longer lag times with $k \gg 1$.

Our method, MGR, adopts an iterative training strategy based on either a long discretized simulation trajectory or multiple discretized trajec-

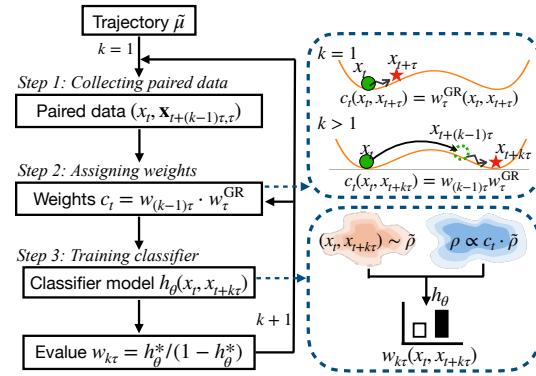


Figure 2: The training algorithm of MGR as illustrated in Sections 4.2 and 4.3

216 tories under perturbed path $\tilde{\mu}$. Suppose the ratio $w_{(k-1)\tau}$ has already been obtained. Then the ratio
 217 $w_{k\tau}$ at lag time $k\tau$ can be constructed as follows:
 218

219 **Step 1.** We collect pairs of the form $(x_t, \mathbf{x}_{t+(k-1)\tau, \tau})$ for $t \in [0, T - k\tau]$ from the perturbed
 220 simulation trajectory, where each pair consists of a state and a trajectory segment.
 221

222 **Step 2.** In this step, we construct approximations of perturbed and original distributions of
 223 $(x_t, x_{t+k\tau})$ based on the following identity:
 224

$$\begin{aligned} \mathbb{E}_{\rho_{k\tau}(x_t, x_{t+k\tau})}[O(x_t, x_{t+k\tau})] &= \mathbb{E}_{\rho_{(k-1)\tau}(x_t, x_{t+(k-1)\tau}) \cdot \mu(\mathbf{x}_{t+(k-1)\tau, \tau} | x_{t+(k-1)\tau})}[O(x_t, x_{t+k\tau})] \\ &= \mathbb{E}_{\tilde{\rho}_{(k-1)\tau}(x_t, x_{t+(k-1)\tau}) \cdot \tilde{\mu}(\mathbf{x}_{t+(k-1)\tau, \tau} | x_{t+(k-1)\tau})}[c_t O(x_t, x_{t+k\tau})], \end{aligned}$$

225 for any bounded measurable test function $O(\cdot, \cdot) : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$. The first equality holds because
 226 $\rho_{k\tau}$ can be regarded as the marginal distribution of $(x_t, x_{t+k\tau})$ defined by $\rho_{(k-1)\tau}(x_t, x_{t+(k-1)\tau}) \cdot$
 227 $\mu(\mathbf{x}_{t+(k-1)\tau, \tau} | x_{t+(k-1)\tau})$, and the second equality follows from the principle of importance sam-
 228 pling, with weight
 229

$$c_t = w_{(k-1)\tau}(x_t, x_{t+(k-1)\tau}) \cdot w_{\tau}^{\text{GR}}(\mathbf{x}_{t+(k-1)\tau, \tau}). \quad (4)$$

230 Since $w_{(k-1)\tau}$ is inherited from the previous iteration and w_{τ}^{GR} can be computed as described in
 231 Section 3.2, the pairs collected in *Step 1* can be used to approximate the perturbed and original joint
 232 distributions $\tilde{\rho}_{k\tau}$ and $\rho_{k\tau}$ as
 233

$$\begin{aligned} \tilde{\rho}_{k\tau}(x, y) &\approx \frac{1}{T - k\tau} \sum_t \delta(x - x_t) \delta(y - x_{t+k\tau}), \\ \rho_{k\tau}(x, y) &\approx \sum_t \frac{c_t \delta(x - x_t) \delta(y - x_{t+k\tau})}{\sum_{t'} c_{t'}}, \end{aligned}$$

234 where δ denotes the Dirac delta function. When accurate estimates of $w_{(k-1)\tau}$ are available and the
 235 dataset size is sufficiently large, the above approximations can be shown to be consistent. A detailed
 236 proof is provided in Appendix C.
 237

238 **Step 3.** Finally, using the joint distributions $\tilde{\rho}_{k\tau}(x, y)$, $\rho_{k\tau}(x, y)$ in *Step 2*, we approximate the
 239 marginal weight $w_{k\tau}(x, y)$ by a classifier-based density ratio estimator (see Section 4.3 for details).
 240 The inferred ratio from the optimal binary classifier is then used for the next iteration.
 241

242 This three-step procedure effectively extends short-time pathwise Girsanov weights to long-time
 243 marginal ratio. First, we train the model w_{τ} using the short τ -lag Girsanov weights. Then, we
 244 iterate three steps above to progressively learn $w_{k\tau}$ until $k \gg 1$. The workflow is summarized in
 245 Figure 2. In the next subsection, we introduce the ratio estimation method used in MGR *Step 3*.
 246

247 4.3 CLASSIFIER-BASED DENSITY RATIO ESTIMATION

248 Many machine learning methods have been proposed for density ratio estimation (Menon & Ong,
 249 2016; Choi et al., 2021; 2022; Yu et al., 2025). A widely used approach is probabilistic classification
 250 (Menon & Ong, 2016), which reformulates likelihood ratio estimation as a binary classification task.
 251 In this setting, a binary classifier with a sigmoid output $h(\cdot, \cdot) : \mathbb{R}^d \times \mathbb{R}^d \rightarrow [0, 1]$ is trained to
 252 discriminate between paired samples $(x_t, x_{t+k\tau})$ drawn from the perturbed distribution $\tilde{\rho}_{k\tau}$ and the
 253 original distribution $\rho_{k\tau}$.
 254

255 We define the optimal classifier h_{θ}^* as the probability that a given pair $(x_t, x_{t+k\tau})$ comes from $\rho_{k\tau}$,
 256 i.e., $h_{\theta}^*(x_t, x_{t+k\tau}) = \rho_{k\tau}(x_t, x_{t+k\tau}) / (\rho_{k\tau}(x_t, x_{t+k\tau}) + \tilde{\rho}_{k\tau}(x_t, x_{t+k\tau}))$. Then, the density ratio
 257 can be estimated as
 258

$$w_{k\tau}(x_t, x_{t+k\tau}) = \frac{\rho_{k\tau}(x_t, x_{t+k\tau})}{\tilde{\rho}_{k\tau}(x_t, x_{t+k\tau})} = \frac{h_{\theta}^*(x_t, x_{t+k\tau})}{1 - h_{\theta}^*(x_t, x_{t+k\tau})}. \quad (5)$$

259 Unlike standard density ratio estimation, where samples from both distributions are available and
 260 the ratio can be learned via cross-entropy loss, our setting in MGR is different. We only have
 261 samples from the perturbed distribution $\tilde{\rho}_{k\tau}$ together with the corresponding weight c_t assigned to
 262 each sample. To address this, we employ a weighted cross-entropy loss:
 263

$$\mathcal{L}(\theta) = -\mathbb{E}_t [c_t \log h_{\theta}(x_t, x_{t+k\tau}) + \log(1 - h_{\theta}(x_t, x_{t+k\tau}))], \quad (6)$$

270 where $(x_t, x_{t+\tau})$ denotes the paired data collected from the perturbed trajectory with lagtime $k\tau$.
 271 The weights c_t are constructed from the output of the previous model $w_{(k-1)\tau}$ and short-lag Gir-
 272 sanov reweighting w_τ^{GR} according to Eq. 4, and are normalized over the entire dataset.
 273

274 After sufficient training on either a single long trajectory or multiple trajectories, the estimator $w_{k\tau}$ in
 275 Eq. 5 converges to the marginal ratio, and then serves as a component for the next iteration. Through
 276 the iterative training scheme, MGR can attain stable marginal ratios under long lag times and in
 277 complex systems. The complete training procedure is summarized in Algorithm 1 and illustrated
 278 in Figure 2. For completeness, we compare a range of ablations and discuss other potential model
 279 choices (see Appendix F.1), but defer a more systematic investigation to future work.
 280

281 5 EXPERIMENTS

282 5.1 MOLECULAR DYNAMICS

284 MGR has particularly useful applications in the analysis of kinetic properties in molecular dynamics
 285 (MD) (Donati et al., 2017; 2022). In particular, it can be employed in the construction of Markov
 286 state models (MSMs) (Prinz et al., 2011) at a given lag time.

287 In an MSM framework, the dynamics are characterized by a transition probability matrix P_τ ,
 288 where each entry is obtained by normalizing the corresponding cross-correlation $C_{ij}(\tau)$. This
 289 quantity measures the probability of observing a transition from state i at time t to state j at
 290 time $t + \tau$. Under biased or enhanced sampling trajectories $\tilde{\mu}$, we want to recover the unbi-
 291 ased properties using reweighting. For each observed transition pair $(x_t, x_{t+\tau})$, a marginal weight
 292 $w_\tau(x_t, x_{t+\tau})$ is assigned to the cross-correlation $C_{ij}(\tau) = \mathbb{E}_{\tilde{\mu}} [w_\tau(x_t, x_{t+\tau}) \mathbf{1}_{B_i}(x_t) \mathbf{1}_{B_j}(x_{t+\tau})]$,
 293 where $w_\tau(x_t, x_{t+\tau})$ is estimated by MGR. In our evaluation, we focus on the following several key
 294 indicators. Detailed definitions and explanations are provided in Appendix E.1.
 295

- 296 • Effective sample size (ESS): Since the true transition density ratio is analytically intractable,
 297 relative ESS is then used as a reference indicator for weight stability (Freeman, 1966). A
 298 larger ESS implies more reliable statistical estimates, and the results are shown in Ap-
 299 pendix E.1.
- 300 • Implied timescales (ITS): Each eigenvalue $\lambda_i(\tau)$ of the transition probability matrix defines
 301 an implied timescale $t_i(\tau) = -\frac{\tau}{\log \lambda_i(\tau)}$. We examine several dominant implied timescales
 302 (ITS) as well as their sum. For each individual ITS, it shows the slow physical timescale of
 303 the system. For their sum, an accurate tracking of the exponential decay with increasing lag
 304 time is expected, which reflects the intrinsic relaxation behavior of the system dynamics.
- 305 • Dominant eigenfunctions and stationary distribution: The leading eigenfunctions characterize
 306 the slow dynamical modes, while the stationary distribution represents the long-time equilib-
 307 rium of the system. Accurate recovery of these quantities is essential for validating kinetic
 308 properties.

309 We primarily compare the results obtained from Marginal Girsanov Reweighting (MGR) with those
 310 produced by pathwise Girsanov reweighting (GR) (Donati et al., 2017). Guidelines for selecting a
 311 suitable short lag τ and network for training in MGR are provided in the Appendices F.2 and F.4.
 312

313 5.1.1 ONE DIMENSIONAL FOUR WELL

315 We first consider a one-dimensional four-well potential system (Prinz et al., 2011), which serves as a
 316 prototypical example for testing reweighting methods. The unbiased energy landscape contains four
 317 metastable states separated by barriers, with the two intermediate wells located at higher energies. To
 318 accelerate sampling, we introduce a biased potential that lowers the energy of the two intermediate
 319 wells. Detailed energy function and simulation information can be found in Appendix E.1.1.

320 Figure 3 presents the results for the four-well system. The left panel shows the dominant left eigen-
 321 functions at lagtime $300\Delta t$ ($\Delta t = 0.001$). GR provides partial correction but still exhibits clear
 322 deviations, particularly in the central wells. By contrast, MGR yields eigenfunctions that are in
 323 close agreement with the unbiased results. It demonstrates that MGR enables accurate recovery of
 324 both equilibrium and kinetic properties. The results of dominant right eigenfunctions and the ESS

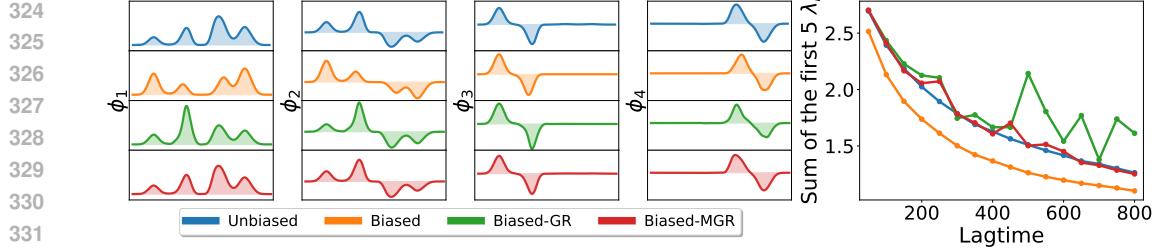


Figure 3: Results for the one-dimensional four-well potential system. **Left:** Dominant left ($\{\phi_i\}$) eigenfunctions of the transition matrix at lagtime $300\Delta t$. (ϕ_1 denotes the stationary distribution.) Unbiased results serve as reference, while Biased trajectories, Biased reweighted by GR (Biased-GR), and Biased reweighted by MGR (Biased-MGR) are compared. **Right:** sum of the first 5 eigenvalues, reflecting the intrinsic relaxation behavior of the dynamics, as a function of lag time.

curve are shown in Figure S.2, where the ESS for GR decays rapidly at long lag times, reflecting the instability. In contrast, our model MGR maintains substantially higher ESS across lag times, stabilizing reweighting .

We also compare the sum of the first five eigenvalues of the transition matrix in the right panel. The results of MGR closely follow the unbiased reference and exhibit smoother exponential decay, whereas GR suffers from strong fluctuations. This indicates that MGR more accurately captures the intrinsic relaxation dynamics of the system. For completeness, the several individual dominant implied timescales (ITS) are reported in Figure S.3.

5.1.2 ALANINE DIPEPTIDE

We next consider alanine dipeptide (Donati et al., 2017; Wu et al., 2020), a widely used benchmark system in molecular dynamics for studying conformational transitions. The backbone torsion angles ϕ and ψ are two important reaction coordinates, containing several important metastable basins. Sampling transitions between these basins is challenging due to the presence of high free-energy barriers. To accelerate exploration, we perform biased simulations by introducing umbrella potentials along ϕ and ψ , which distort the original equilibrium distribution and transition probabilities. Detailed energy function and simulation information can be found in Appendix E.1.2.

Figure 4 shows the results for alanine dipeptide. a) reports the sum of the first five eigenvalues as a function of lag time. MGR consistently follows the unbiased reference, whereas GR shows clear

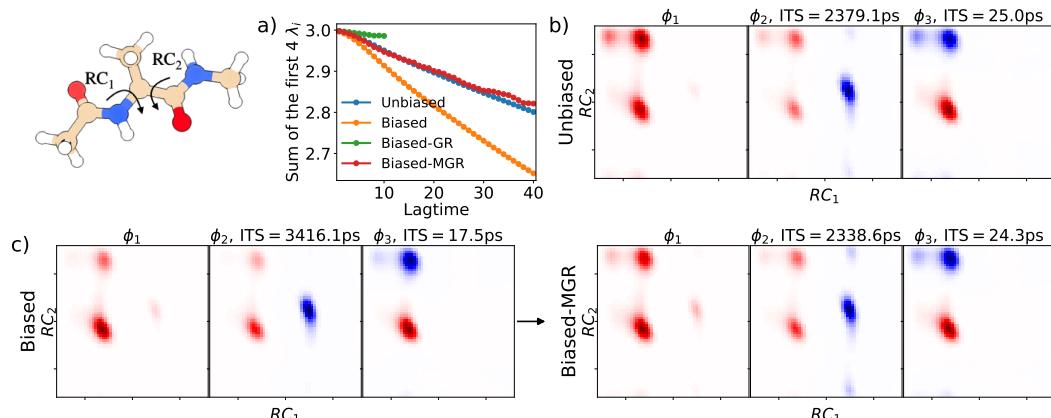


Figure 4: Results for the 22 atoms Alanine Dipeptide system. **a)** sum of the first five eigenvalues, reflecting the intrinsic relaxation behavior of the dynamics, as a function of lag time (lagtime unit is 40fs). **b)** Dominant left eigenfunctions of the transition matrix with corresponding ITS values are shown in the title at lag time 1.6ps. Unbiased results serve as reference. **c)** Dominant left eigenfunctions from the biased trajectory and the MGR-recovered results. GR fails to construct a convergent Markov state model due to excessive variance.

378 deviations. Due to the large variance and low ESS values, GR fails to provide a reliable MSM for
 379 lag times exceeding $10\tau = 400$ fs. b) shows the dominant left eigenfunctions and the corresponding
 380 implied timescales (ITS) at lag time $40\tau = 1.6$ ps under the unbiased trajectory. These serve as
 381 our evaluation reference. Using MGR, we reweight the biased trajectory to recover the unbiased
 382 properties. Since GR has already failed at this timescale, only MGR results are reported in c).
 383 Both the eigenfunctions and implied timescales demonstrate that MGR successfully reproduces the
 384 unbiased behavior. Detailed comparisons of ESS and individual ITS are provided in Figure S.6.

385 These two examples demonstrate that MGR achieves consistently strong performance in molecular
 386 dynamics. It maintains stable effective sample sizes, produces reliable implied timescales, and ac-
 387 curately recovers equilibrium and kinetic properties, even under large system and at long lag times
 388 where traditional GR fails. These results highlight the effectiveness of MGR, especially in mitigating
 389 variance growth and removing path dependence. As a result, MGR provides a practical framework
 390 for reweighting biased simulations in MD.

392 5.2 BAYESIAN INFERENCE OF SDE PARAMETERS

393 Another important application of MGR lies in the Bayesian posterior inference of SDEs (Sørensen,
 394 2004; Golightly & Wilkinson, 2008; Li et al., 2020; Ghosh et al., 2022). MGR is used to estimate
 395 the parameters of a stochastic differential equation from discrete observations. Detailed illustration
 396 can be found in Appendix E.2.

397 Consider a d -dimensional diffusion process, whose drift term contains unknown parameters θ . In-
 398 stead of simulating the SDE for every candidate θ , we set a reference parameter θ_0 , and generate
 399 reference trajectories under θ_0 .

400 Through the MGR model introduced in Section 4, we utilize short-lag Girsanov weights
 401 $w_{\tau}^{\theta, \theta_0}(\mathbf{x}_{t, \tau}) = \frac{d\mu_{\theta}}{d\mu_{\theta_0}}(\mathbf{x}_{t, \tau})$ in Eq. 2 iteratively, where the drift difference is $f(\cdot, t; \theta) - f(\cdot, t; \theta_0)$.

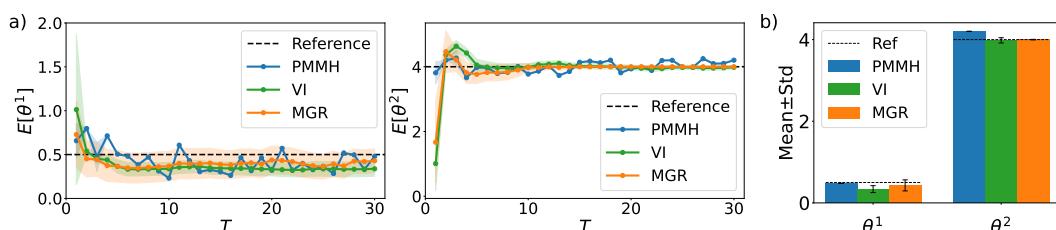
402 After training, stable marginal weights $w_{k\tau}^{\theta}(\cdot, \cdot)$ are used to estimate the likelihood of sparse obser-
 403 vations $\mathbf{y} = \{y(t=0) = y^0, y^1, \dots, y(t=T) = y^K\}$, where $k\tau = T/K$:

$$404 w_{k\tau}^{\theta}(\mathbf{y}) = \frac{p(\mathbf{y} \mid \theta)}{p(\mathbf{y} \mid \theta_0)} = \prod_i \frac{p(y^{i+1} \mid y^i; \theta)}{p(y^{i+1} \mid y^i; \theta_0)} = \prod_i w_{k\tau}^{\theta}(y^i, y^{i+1}).$$

405 This allows us to construct a stable likelihood ratio based on the given paired observations
 406 $\{(y^i, y^{i+1})\}$. Consequently, posterior expectations can be evaluated by importance sampling with
 407 observations \mathbf{y} :

$$408 \mathbb{E}[\theta \mid \mathbf{y}] = \int \theta p(\theta \mid \mathbf{y}) d\theta \approx \frac{\sum_{\theta} \theta w_{k\tau}^{\theta}(\mathbf{y}) p(\theta)}{\sum_{\theta} w_{k\tau}^{\theta}(\mathbf{y}) p(\theta)}.$$

410 Leveraging short-lag, pathwise Girsanov factors, MGR constructs reliable likelihood ratios over long
 411 observation intervals and time horizons, which enables accurate posterior estimation of the param-
 412 eters. In this example, we primarily compare MGR with the particle marginal Metropolis–Hastings
 413 (PMMH) algorithm (Golightly & Wilkinson, 2008; Hoffman et al., 2014) and a variational inference
 414 (VI) approach (Ghosh et al., 2022).



416 Figure 5: Results for Bayesian inference in the Graph Ornstein–Uhlenbeck process. **a)** Posterior es-
 417 timates of θ^1 and θ^2 under different observation horizons. Estimation of θ^1 is challenging, and MGR
 418 yields more stable posteriors. For θ^2 , MGR achieves highly concentrated and precise estimates. **b)**
 419 Posterior estimates with mean and standard deviation at the longest observation horizon $T = 30$.
 420

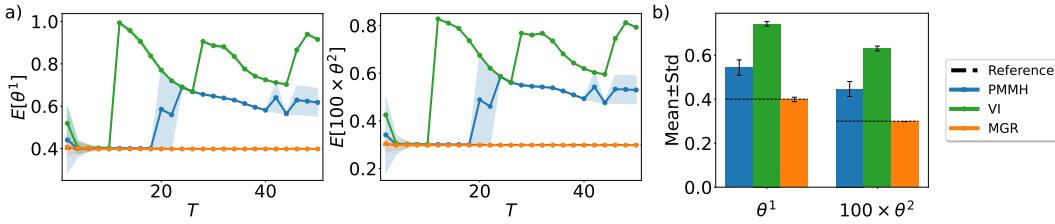


Figure 6: Results for Bayesian inference in the Stochastic Lotka-Volterra system. **a)** Posterior estimates of θ^1 and $100 \times \theta^2$ under different observation horizons. **b)** Posterior distributions by the mean of all time horizons.

5.2.1 GRAPH ORNSTEIN-UHLENBECK

we first consider the Ornstein–Uhlenbeck (OU) process (Courgeau & Veraart, 2022) defined on a graph structure,

$$dX_t = -A(\theta^1)(X_t - B(\theta^2)) dt + \sigma dW_t, X_t \in \mathbb{R}^2, A(\theta^1) = \begin{pmatrix} 1 & -\theta^1 \\ -\theta^1 & 1 \end{pmatrix}, B(\theta^2) = \begin{pmatrix} \theta^2 \\ 1 \end{pmatrix},$$

where $\sigma = 0.5$, and $t \in [0, 30]$. We set the reference parameter as $\theta_0^1 = 0.5, \theta_0^2 = 4$, and a latent trajectory are generated under this reference parameter. We consider a grid of 225 candidate parameters, where the detailed simulation settings and parameters candidates are provided in Appendix E.2.1. For the observations, we sample every unit time step $t = 1$ along the simulation path. At each sampling point, Gaussian noise is introduced to obtain observations $\mathbf{y} = \{y(t=0) = y^0, y^1, \dots, y(t=30) = y^{30}\}, y_i \sim \mathcal{N}(x_i, 0.25I)$.

Figure 5 reports the posterior estimates under different observation trajectory lengths. We observe that θ^1 is more difficult to identify: VI exhibits noticeable bias, while PMMH shows substantial fluctuations and unstable convergence. In contrast, MGR consistently yields better estimates across different horizons. For θ^2 , which is easier to infer due to its stronger signal in the dynamics, MGR again provides highly accurate and concentrated posteriors, clearly outperforming the alternatives.

5.2.2 STOCHASTIC LOTKA-VOLTERRA

Finally, we consider the stochastic Lotka–Volterra model (Wilkinson, 2018), which describes a population comprising of two competing species: prey grows intrinsically at rate θ^1 and is depleted due to predators at strength θ^2 ; predator dies intrinsically at rate θ^1 but increases through predation with the same interaction strength θ^2 . This system can be defined by

$$dX_t = A(X_t, \theta) dt + \sigma dW_t, X_t \in \mathbb{R}^2, A(X_t, \theta) = \begin{pmatrix} \theta^1 X_t^1 - \theta^2 X_t^1 X_t^2 \\ \theta^2 X_t^1 X_t^2 - \theta^1 X_t^2 \end{pmatrix},$$

where $\sigma = 0.1$ and $t \in [0, 50]$. We set the reference parameter as $\theta_0^1 = 0.4, \theta_0^2 = 0.003$, and simulate a latent trajectory under this reference parameter. We also consider a grid of 225 candidate parameters, where the detailed simulation settings and parameters candidates are provided in Appendix E.2.2. For the observations, we sample every unit time step $t = 2$ along the simulation path. At each sampling point, Gaussian noise is introduced to obtain observations $\mathbf{y} = \{y(t=0) = y^0, y^1, \dots, y(t=50) = y^{25}\}, y_i \sim \mathcal{N}(x_i, 25I)$.

Figure 6 reports the posterior estimates under different observation trajectory lengths. As the time horizon T (and thus the number of observations) increases, both PMMH and VI exhibit growing bias and variability. In contrast, MGR remains essentially flat across T and closely tracks accurate posterior estimates of the parameters.

By estimating the marginal likelihood ratio between paired observations, MGR can effectively avoid weight degeneracy that afflicts the baseline methods. Across these two examples, MGR shows reliable and efficient Bayesian posterior estimates under both long lag times and extended time horizons, which makes it a practical tool for posterior inference in complex stochastic systems.

486 6 CONCLUSION
487488 Our work introduces Marginal Girsanov Reweighting (MGR), an innovative approach for estimating
489 unbiased properties from perturbed paired data. Based on biased pathwise Girsanov reweighting
490 (GR)—which computes the ratio along a specific trajectory—MGR learns a marginal estimator by
491 integrating over intermediate states. This approach mitigates the variance blow-up of GR and yields
492 an unbiased estimate of the transition probability under long lags.493 We show that MGR can be implemented using standard density ratio ratio estimation setups with
494 only minor adaptions, and outperforms several base-lines on benchmarks spanning multiple different
495 fields.496 Despite MGR’s advantages, challenges remain. When the perturbed drift deviates substantially from
497 the reference dynamics, GR weights become unstable, which in turn hampers MGR training. Care-
498 ful choice of reaction coordinates is therefore essential in practice. Moreover, while we implemented
499 MGR using standard classifier-based ratio estimation, richer neural estimators and architectures hold
500 promise for further gains. **In addition, given the rapid development of physical systems and molecu-**
501 **lar dynamics, an interesting direction is to extend MGR toward simulation-based Bayesian inference**
502 **and multi-ensemble estimators.** Exploring these directions will broaden the scope of MGR, paving
503 the way for robust, ML-driven reweighting methods applicable across scales, domains, and dynam-
504 ical systems.505
506 ACKNOWLEDGEMENTS
507508 This work was partially supported by the Wallenberg AI, Autonomous Systems and Software Pro-
509 gram (WASP) funded by the Knut and Alice Wallenberg Foundation and by the National Natural
510 Science Foundation of China (NSFC) under grant number 12171367. Preliminary results were en-
511 abled by resources provided by the National Academic Infrastructure for Supercomputing in Sweden
512 (NAIIS) at Alvis (project: NAIIS 2025/22-463), partially funded by the Swedish Research Council
513 through grant agreement no. 2022-06725.514
515 REFERENCES
516517 Christophe Andrieu, Arnaud Doucet, and Roman Holenstein. Particle markov chain monte carlo
518 methods. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 72(3):269–
519 342, 2010.
520 Alessandro Barducci, Giovanni Bussi, and Michele Parrinello. Well-tempered metadynamics: a
521 smoothly converging and tunable free-energy method. *Physical review letters*, 100(2):020603,
522 2008.
523 Alexandros Beskos and Gareth O Roberts. Exact simulation of diffusions. 2005.
524 Leo Breiman. The strong law of large numbers for a class of markov chains. *The Annals of Mathe-*
525 *matical Statistics*, 31(3):801–803, 1960.
526 Guy Bunin. Ecological communities with lotka-volterra dynamics. *Physical Review E*, 95(4):
527 042414, 2017.
528 Xin Cai, Jingyu Yang, Zhibao Li, Hongqiao Wang, and Miao Huang. Simulation-based transition
529 density approximation for the inference of sde models. *arXiv preprint arXiv:2401.02529*, 2023.
530 Kristy Choi, Madeline Liao, and Stefano Ermon. Featurized density ratio estimation. In *Uncertainty*
531 *in Artificial Intelligence*, pp. 172–182. PMLR, 2021.
532 Kristy Choi, Chenlin Meng, Yang Song, and Stefano Ermon. Density ratio estimation via infinitesi-
533 mal classification. In *International Conference on Artificial Intelligence and Statistics*, pp. 2552–
534 2573. PMLR, 2022.
535 Valentin Courgeau and Almut ED Veraart. Likelihood theory for the graph ornstein-uhlenbeck
536 process. *Statistical Inference for Stochastic Processes*, 25(2):227–260, 2022.

540 Kyle Cranmer, Johann Brehmer, and Gilles Louppe. The frontier of simulation-based inference.
 541 *Proceedings of the National Academy of Sciences*, 117(48):30055–30062, 2020.
 542

543 Timothée Devergne, Vladimir R. Kostic, Michele Parrinello, and Massimiliano Pontil. From bi-
 544 ased to unbiased dynamics: An infinitesimal generator approach. In A. Globerson, L. Mackey,
 545 D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.), *Advances in Neural Information
 546 Processing Systems*, volume 37, pp. 75495–75521. Curran Associates, Inc., 2024. doi: 10.52202/
 547 079017-2404. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/89edef87915d31de3437b6b2ac5f79e7-Paper-Conference.pdf.
 548

549 Manuel Dibak, Leon Klein, Andreas Krämer, and Frank Noé. Temperature steerable flows and
 550 boltzmann generators. *Physical Review Research*, 4(4):L042005, 2022.
 551

552 Juan Viguera Diez, Mathias Schreiner, Ola Engkvist, and Simon Olsson. Boltzmann priors for
 553 implicit transfer operators. *arXiv preprint arXiv:2410.10605*, 2024.
 554

555 Juan Viguera Diez, Mathias Jacob Schreiner, Ola Engkvist, and Simon Olsson. Boltzmann priors
 556 for implicit transfer operators. In *The Thirteenth International Conference on Learning Repre-
 557 sentations*, 2025. URL <https://openreview.net/forum?id=pRCOZ11ZdT>.
 558

559 Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real nvp. *arXiv
 560 preprint arXiv:1605.08803*, 2016.
 561

562 Justin Domke and Daniel R Sheldon. Importance weighting and variational inference. *Advances in
 563 neural information processing systems*, 31, 2018.
 564

565 Lorenzo Donati, Carsten Hartmann, and Bettina G Keller. Girsanov reweighting for path ensembles
 566 and markov state models. *The Journal of chemical physics*, 146(24), 2017.
 567

568 Luca Donati, Marcus Weber, and Bettina G Keller. A review of girsanov reweighting and of
 569 square root approximation for building molecular markov state models. *Journal of Mathemat-
 570 ical Physics*, 63(12), 2022.
 571

572 Peter Eastman, Raimondas Galvelis, Raúl P Peláez, Charles RA Abreu, Stephen E Farr, Emilio Gal-
 573 licchio, Anton Gorenko, Michael M Henry, Frank Hu, Jing Huang, et al. Openmm 8: molecular
 574 dynamics simulation with machine learning potentials. *The Journal of Physical Chemistry B*, 128
 575 (1):109–116, 2023.
 576

577 Bjørn Eraker. Mcmc analysis of diffusion models with application to finance. *Journal of Business
 578 & Economic Statistics*, 19(2):177–191, 2001.
 579

580 Linton C Freeman. Kish: Survey sampling (book review). *Social Forces*, 45(1):132, 1966.
 581

582 Christiane Fuchs. *Inference for diffusion processes: with applications in life sciences*. Springer
 583 Science & Business Media, 2013.
 584

585 Emilio Gallicchio, Michael Andrec, Anthony K Felts, and Ronald M Levy. Temperature weighted
 586 histogram analysis method, replica exchange, and transition paths. *The Journal of Physical Chem-
 587 istry B*, 109(14):6722–6731, 2005.
 588

589 Sanmitra Ghosh, Paul J Birrell, and Daniela De Angelis. Differentiable bayesian inference of sde
 590 parameters using a pathwise series expansion of brownian motion. In *International Conference
 591 on Artificial Intelligence and Statistics*, pp. 10982–10998. PMLR, 2022.
 592

593 Igor Vladimirovich Girsanov. On transforming a certain class of stochastic processes by absolutely
 594 continuous substitution of measures. *Theory of Probability & Its Applications*, 5(3):285–301,
 595 1960.
 596

597 Manuel Gloeckler, Shoji Toyota, Kenji Fukumizu, and Jakob H Macke. Compositional simulation-
 598 based inference for time series. *arXiv preprint arXiv:2411.02728*, 2024.
 599

600 Andrew Golightly and Darren J Wilkinson. Bayesian inference for nonlinear multivariate diffusion
 601 models observed with error. *Computational Statistics & Data Analysis*, 52(3):1674–1693, 2008.
 602

594 Arthur Gretton, Alex Smola, Jiayuan Huang, Marcel Schmittfull, Karsten Borgwardt, Bernhard
 595 Schölkopf, et al. Covariate shift by kernel mean matching. *Dataset shift in machine learning*, 3
 596 (4):5, 2009.

597 Helmut Grubmüller. Predicting slow structural transitions in macromolecular systems: Conforma-
 598 tional flooding. *Physical Review E*, 52(3):2893, 1995.

599

600 Jack Hachigian. Collapsed markov chains and the chapman-kolmogorov equation. *The Annals of*
 601 *Mathematical Statistics*, 34(1):233–237, 1963.

602 Matthew D Hoffman, Andrew Gelman, et al. The no-u-turn sampler: adaptively setting path lengths
 603 in hamiltonian monte carlo. *J. Mach. Learn. Res.*, 15(1):1593–1623, 2014.

604

605 Thomas Huber, Andrew E Torda, and Wilfred F Van Gunsteren. Local elevation: a method for
 606 improving the searching properties of molecular dynamics simulation. *Journal of computer-aided*
 607 *molecular design*, 8(6):695–708, 1994.

608 Brooke E Husic and Vijay S Pande. Markov state models: From an art to a science. *Journal of the*
 609 *American Chemical Society*, 140(7):2386–2396, 2018.

610

611 Michele Invernizzi, Andreas Krämer, Cecilia Clementi, and Frank Noé. Skipping the replica ex-
 612 change ladder with normalizing flows. *The Journal of Physical Chemistry Letters*, 13(50):
 613 11643–11649, December 2022. ISSN 1948-7185. doi: 10.1021/acs.jpclett.2c03327. URL
 614 <http://dx.doi.org/10.1021/acs.jpclett.2c03327>.

615 Christopher S Jones. Bayesian estimation of continuous-time finance models. *manuscript University*
 616 *of Rochester*, 1998.

617

618 Anna S Kamenik, Stephanie M Linker, and Sereina Riniker. Enhanced sampling without borders:
 619 on global biasing functions and how to reweight them. *Physical Chemistry Chemical Physics*, 24
 620 (3):1225–1236, 2022.

621

622 Johannes Kästner. Umbrella sampling. *Wiley Interdisciplinary Reviews: Computational Molecular*
 623 *Science*, 1(6):932–942, 2011.

624

625 Leon Klein, Andrew Foong, Tor Fjelde, Bruno Mlodzeniec, Marc Brockschmidt, Sebastian
 626 Nowozin, Frank Noé, and Ryota Tomioka. Timewarp: Transferable acceleration of molecular
 627 dynamics by learning time-coarsened dynamics. *Advances in Neural Information Processing*
 628 *Systems*, 36:52863–52883, 2023.

629

630 Soheil Kolouri, Gustavo K Rohde, and Heiko Hoffmann. Sliced wasserstein distance for learning
 631 gaussian mixture models. In *Proceedings of the IEEE Conference on Computer Vision and Pattern*
 632 *Recognition*, pp. 3427–3436, 2018.

633

634 Xuechen Li, Ting-Kam Leonard Wong, Ricky TQ Chen, and David Duvenaud. Scalable gradients
 635 for stochastic differential equations. In *International Conference on Artificial Intelligence and*
 636 *Statistics*, pp. 3870–3882. PMLR, 2020.

637

638 Simon MJ Lyons, Simo Särkkä, and Amos J Storkey. Series expansion approximations of brown-
 639 ian motion for non-linear kalman filtering of diffusion processes. *IEEE Transactions on Signal*
 640 *Processing*, 62(6):1514–1524, 2014.

641

642 Andreas Mardt, Luca Pasquali, Hao Wu, and Frank Noé. Vampnets for deep learning of molecular
 643 kinetics. *Nature communications*, 9(1):5, 2018.

644

645 Aditya Menon and Cheng Soon Ong. Linking losses for density ratio and class-probability estima-
 646 tion. In *International Conference on Machine Learning*, pp. 304–313. PMLR, 2016.

647

648 Antonia SJS Mey, Hao Wu, and Frank Noé. xtram: Estimating equilibrium expectations from time-
 649 correlated simulation data at multiple thermodynamic states. *Physical Review X*, 4(4):041018,
 650 2014.

651

652 Selma Moqvist, Weilong Chen, Mathias Schreiner, Feliks Nüske, and Simon Olsson. Thermody-
 653 namic interpolation: A generative approach to molecular thermodynamics and kinetics. *Journal*
 654 *of Chemical Theory and Computation*, 21(5):2535–2545, 2025.

648 Andrea Pascucci. *PDE and martingale methods in option pricing*. Springer Science & Business
 649 Media, 2011.

650

651 Jan-Hendrik Prinz, Hao Wu, Marco Sarich, Bettina Keller, Martin Senne, Martin Held, John D
 652 Chodera, Christof Schütte, and Frank Noé. Markov models of molecular kinetics: Generation
 653 and validation. *The Journal of chemical physics*, 134(17), 2011.

654 Joana-Lysiane Schafer and Bettina G Keller. Implementation of girsanov reweighting in openmm
 655 and deptime. *The Journal of Physical Chemistry B*, 128(25):6014–6027, 2024.

656 Mathias Schreiner, Ole Winther, and Simon Olsson. Implicit transfer operator learning: Multiple
 657 time-resolution models for molecular dynamics. *Advances in Neural Information Processing
 658 Systems*, 36:36449–36462, 2023.

659

660 Michael R Shirts and John D Chodera. Statistically optimal analysis of samples from multiple
 661 equilibrium states. *The Journal of chemical physics*, 129(12), 2008.

662 Helle Sørensen. Parametric inference for diffusion processes observed at discrete points in time: a
 663 survey. *International statistical review*, 72(3):337–354, 2004.

664 Marc Souaille and Benoit Roux. Extension to the weighted histogram analysis method: combining
 665 umbrella sampling with free energy calculations. *Computer physics communications*, 135(1):
 666 40–57, 2001.

667

668 Yuji Sugita and Yuko Okamoto. Replica-exchange molecular dynamics method for protein folding.
 669 *Chemical physics letters*, 314(1-2):141–151, 1999.

670 Masashi Sugiyama, Taiji Suzuki, and Takafumi Kanamori. Density ratio estimation: A comprehen-
 671 sive review (statistical experiment and its related topics). 2010.

672 Robert H Swendsen and Jian-Sheng Wang. Replica monte carlo simulation of spin-glasses. *Physical
 673 review letters*, 57(21):2607, 1986.

674

675 Matthew Tancik, Pratul Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh
 676 Singhal, Ravi Ramamoorthi, Jonathan Barron, and Ren Ng. Fourier features let networks learn
 677 high frequency functions in low dimensional domains. *Advances in neural information processing
 678 systems*, 33:7537–7547, 2020.

679 Pratyush Tiwary, Vittorio Limongelli, Matteo Salvalaglio, and Michele Parrinello. Kinetics of
 680 protein–ligand unbinding: Predicting pathways, rates, and rate-limiting steps. *Proceedings of
 681 the National Academy of Sciences*, 112(5):E386–E391, 2015.

682 Glenn M Torrie and John P Valleau. Nonphysical sampling distributions in monte carlo free-energy
 683 estimation: Umbrella sampling. *Journal of computational physics*, 23(2):187–199, 1977.

684

685 Eric Vanden-Eijnden et al. Transition-path theory and path-finding algorithms for the study of rare
 686 events. *Annual review of physical chemistry*, 61:391–420, 2010.

687 Lingle Wang, Richard A Friesner, and BJ Berne. Replica exchange with solute scaling: a more effi-
 688 cient version of replica exchange with solute tempering (rest2). *The Journal of Physical Chemistry
 689 B*, 115(30):9431–9438, 2011.

690

691 Yihang Wang, Lukas Herron, and Pratyush Tiwary. From data to noise to data for mixing physics
 692 across temperatures with generative artificial intelligence. *Proceedings of the National Academy
 693 of Sciences*, 119(32):e2203656119, 2022.

694

695 Darren J Wilkinson. *Stochastic modelling for systems biology*. Chapman and Hall/CRC, 2018.

696

697 Hao Wu, Feliks Nüske, Fabian Paul, Stefan Klus, Péter Kolai, and Frank Noé. Variational koopman
 698 models: Slow collective variables and molecular kinetics from short off-equilibrium simulations.
 699 *The Journal of chemical physics*, 146(15), 2017.

700

701 Hao Wu, Jonas Köhler, and Frank Noé. Stochastic normalizing flows. *Advances in neural informa-
 702 tion processing systems*, 33:5933–5944, 2020.

Hanlin Yu, Arto Klami, Aapo Hyvärinen, Anna Korba, and Omar Chehab. Density ratio estimation
 with conditional probability paths. *arXiv preprint arXiv:2502.02300*, 2025.

702 **A VARIANCE OF GIRSANOV REWEIGHTING**
 703

704 We analyze the variance behavior of Girsanov reweighting over a continuous time interval $[t, t + \tau]$.
 705 According to Girsanov theory Girsanov (1960), the log-weight under the Girsanov transformation
 706 can be expressed as

707
$$\log w_{\tau}^{\text{GR}}(\mathbf{x}_{t,\tau}) = \int_t^{t+\tau} u(x_s, s)^\top dW_s - \frac{1}{2} \int_t^{t+\tau} \|u(x_s, s)\|^2 ds,$$

 708

710 where $u(x_t, t) := \frac{f(x_t, t) - \tilde{f}(x_t, t)}{g(t)}$ denotes the rescaled drift difference.
 711

712 Given a fixed control path $\mathbf{u}_{t,\tau} := \{u(x_s, s)\}_{s=t}^{t+\tau}$, the expectation and variance of the log-weight are
 713

714
$$\mathbb{E} [\log w_{\tau}^{\text{GR}}(\mathbf{x}_{t,\tau}) | \mathbf{u}_{t,\tau}] = -\frac{1}{2} \int_t^{t+\tau} \|u(x_s, s)\|^2 ds,$$

 715
 716
$$\text{Var} (\log w_{\tau}^{\text{GR}}(\mathbf{x}_{t,\tau}) | \mathbf{u}_{t,\tau}) = \int_t^{t+\tau} \|u(x_s, s)\|^2 ds.$$

 717

718 We now consider a time-discretized version of the variance over an interval $[t, t + \tau]$ with N steps
 719 of size $\Delta t = \tau/N$. Denote the discretized control as $\mathbf{u}_{t,\tau}^{0:N-1} := \{u(x^k, t^k)\}_{k=0}^{N-1}$. Exponentiating
 720 this log-weight, the conditional variance of the log-weight becomes
 721

722
$$\text{Var} (w_{\tau}^{\text{GR}}(\mathbf{x}_{t,\tau}^{0:N}) | \mathbf{u}_{t,\tau}^{0:N-1}) = \exp \left(\sum_{k=0}^{N-1} \|u(x^k, t^k)\|^2 \Delta t \right).$$

 723

724 This reveals that the variance of the Girsanov weight grows exponentially with both trajectory duration τ and the dimension of control magnitude u , where $\|u(\cdot, t)\|^2$ scales with the dimension d .
 725

726 **B RATIO OF TRANSITION DENSITY**
 727

730 Let $\mathbf{x}_{t,\tau} = \{x_s\}_{s=t}^{t+\tau}$ denote the path space of continuous trajectories. We consider two probability
 731 measures, which are $\mu(\mathbf{x}_{t,\tau})$ (the original process defined in Eq. 1) and $\tilde{\mu}(\mathbf{x}_{t,\tau})$ (the perturbed
 732 process with drift term $\tilde{f}(\cdot, t)$). For any bounded measurable test function $O : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$,
 733

734
$$\begin{aligned} \mathbb{E}_{\rho_{\tau}(x_t, x_{t+\tau})} [O(x_t, x_{t+\tau})] &= \int O(x_t, x_{t+\tau}) \rho_{\tau}(x_t, x_{t+\tau}) dx_t dx_{t+\tau} \\ 735 &= \int O(x_t, x_{t+\tau}) \delta(X_t - x_t) \delta(X_{t+\tau} - x_{t+\tau}) d\mu(\mathbf{x}_{t,\tau}) \\ 736 &= \int O(x_t, x_{t+\tau}) \delta(X_t - x_t) \delta(X_{t+\tau} - x_{t+\tau}) \frac{d\mu}{d\tilde{\mu}}(\mathbf{x}_{t,\tau}) d\tilde{\mu}(\mathbf{x}_{t,\tau}) \\ 737 &= \int O(x_t, x_{t+\tau}) \tilde{\rho}_{\tau}(x_t, x_{t+\tau}) \frac{d\mu}{d\tilde{\mu}}(\mathbf{x}_{t,\tau}) d\tilde{\mu}(\mathbf{x}_{t,\tau}) \\ 738 &:= \mathbb{E}_{\tilde{\rho}_{\tau}(x_t, x_{t+\tau})} [w_{\tau}(x_t, x_{t+\tau}) O(x_t, x_{t+\tau})], \end{aligned}$$

 739

740 where $w_{\tau}(x_t, x_{t+\tau}) = \mathbb{E}_{\tilde{\mu}} \left[\frac{d\mu}{d\tilde{\mu}}(\mathbf{x}_{t,\tau}) \mid X_t = x_t, X_{t+\tau} = x_{t+\tau} \right]$. Here, $\rho_{\tau}(x_t, x_{t+\tau})$, $\tilde{\rho}_{\tau}(x_t, x_{t+\tau})$
 741 denote the joint marginal distributions under the original and perturbed processes respectively.
 742

743 This formulation suggests that the ratio of transition densities $w_{\tau}(x_t, x_{t+\tau})$ can be estimated by
 744 Girsanov reweighting. However, beyond the well-known issue of rapidly growing variance in Ap-
 745 pendix A, Girsanov reweighting computes weights tied to specific trajectories $\mathbf{x}_{t,\tau}$, whereas the
 746 desired transition ratio $w_{\tau}(x_t, x_{t+\tau})$ corresponds to an expectation over paths connecting the given
 747 endpoints.
 748

749 **C CONSISTENCY OF THE APPROXIMATE $\rho_{k\tau}(x, y)$**
 750

751 For the given iteration $k > 1$, let $Z := X_{t+(k-1)\tau}$ and denote the short lag- τ path $\mathbf{x}_{t+(k-1)\tau, \tau} =$
 752 $\{x_s\}_{s=t+(k-1)\tau}^{t+k\tau}$. For any bounded measurable $O : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$, an unbiased estimation of
 753

756 transition properties can be obtained by

$$758 \mathbb{E}_{\rho_{k\tau}(x_t, x_{t+k\tau})} [O(x_t, x_{t+k\tau})] = \mathbb{E}_{\tilde{\rho}_{(k-1)\tau}(x_t, z), \tilde{\mu}(\mathbf{x}_{t+(k-1)\tau, \tau} | z)} [c_t O(x_t, x_{t+k\tau})],$$

759 where $c_t = w_{(k-1)\tau}(x_t, z) w_{\tau}^{\text{GR}}(\mathbf{x}_{t+(k-1)\tau, \tau})$. Here, $w_{(k-1)\tau}(x_t, z)$ is the marginal weights
760 inherited from the previous iteration, and $w_{\tau}^{\text{GR}}(\mathbf{x}_{t+(k-1)\tau, \tau})$ is the pathwise Girsanov reweighting
761 introduced in Section 3.2.

762 In practice we approximate this expectation by Monte Carlo. Extract all paired data
763 $\{(x^i, \mathbf{x}_{i+(k-1)\tau, \tau}^{0:N})\}_{i=1}^M$ under $\tilde{\mu}$ from either a single long trajectory or multiple trajectories,
764 where each short segment $\mathbf{x}_{i+(k-1)\tau, \tau}^{0:N}$ connects the intermediate endpoints (z^i, y^i) with lag
765 τ , (thus the total lag between x^i and y^i is $k\tau$). With an assigned pathwise weight $c^i =$
766 $w_{(k-1)\tau}(x^i, z^i) w_{\tau}^{\text{GR}}(\mathbf{x}_{i+(k-1)\tau, \tau}^{0:N})$, the expectation can be estimated by

$$768 \mathbb{E}_{\rho_{k\tau}(x_t, x_{t+k\tau})} [O(x_t, x_{t+k\tau})] \approx \sum_{i=1}^M \frac{c^i O(x^i, y^i)}{\sum_{i=1}^M c^i}.$$

771 The estimation error approaches zero when $w_{(k-1)\tau}$ is accurate and $M \rightarrow \infty$.

772 *Proof:*

773 By Chapman–Kolmogorov (Hachigian, 1963), we have

$$774 \mathbb{E}_{\rho_{k\tau}(x_t, x_{t+k\tau})} [O(x_t, x_{t+k\tau})] = \int \rho_{k\tau}(x_t, x_{t+k\tau}) O(x_t, x_{t+k\tau}) dx_t dx_{t+k\tau}$$

$$775 = \int \rho_{(k-1)\tau}(x_t, z) p(x_{t+k\tau} | z) O(x_t, x_{t+k\tau}) dz dx_t dx_{t+k\tau}$$

$$776 = \int \frac{\rho_{(k-1)\tau}(x_t, z) p(x_{t+k\tau} | z)}{\tilde{\rho}_{(k-1)\tau}(x_t, z) \tilde{p}(x_{t+k\tau} | z)}.$$

$$777 \tilde{\rho}_{(k-1)\tau}(x_t, z) \tilde{p}(x_{t+k\tau} | z) O(x_t, x_{t+k\tau}) dz dx_t dx_{t+k\tau}.$$

778 According to Section 4.1, $\frac{p(x_{t+k\tau} | z)}{\tilde{p}(x_{t+k\tau} | z)} = \mathbb{E}_{\tilde{\mu}(\mathbf{x}_{t+(k-1)\tau, \tau})} [w_{\tau}^{\text{GR}}(\mathbf{x}_{t+(k-1)\tau, \tau}) | Z = z, X_{t+k\tau} = x_{t+k\tau}]$
779 and $w_{(k-1)\tau}(x_t, z) = \frac{\rho_{(k-1)\tau}(x_t, z)}{\tilde{\rho}_{(k-1)\tau}(x_t, z)}$. It yields

$$780 \mathbb{E}_{\rho_{k\tau}(x_t, x_{t+k\tau})} [O(x_t, x_{t+k\tau})] = \int \tilde{\rho}_{(k-1)\tau}(x_t, z) \tilde{p}(x_{t+k\tau} | z) w_{(k-1)\tau}(x_t, z) w_{\tau}^{\text{GR}}(\mathbf{x}_{t+(k-1)\tau, \tau})$$

$$781 O(x_t, x_{t+k\tau}) \tilde{\mu}(d\mathbf{x}_{t+(k-1)\tau, \tau} | z, x_{t+k\tau}) dz dx_t dx_{t+k\tau}$$

$$782 = \int \tilde{\rho}_{(k-1)\tau}(x_t, z) \tilde{\mu}(d\mathbf{x}_{t+(k-1)\tau, \tau} | z)$$

$$783 w_{(k-1)\tau}(x_t, z) w_{\tau}^{\text{GR}}(\mathbf{x}_{t+(k-1)\tau, \tau}) O(x_t, x_{t+k\tau}) dz dx_t$$

$$784 := \mathbb{E}_{\tilde{\rho}_{(k-1)\tau}(x_t, z), \tilde{\mu}(\mathbf{x}_{t+(k-1)\tau, \tau} | z)} [c_t O(x_t, x_{t+k\tau})].$$

785 Due to the fact $\mathbb{E}_{\rho_{k\tau}(x_t, x_{t+k\tau})} [O(x_t, x_{t+k\tau})] = \mathbb{E}_{\tilde{\rho}_{k\tau}(x_t, x_{t+k\tau})} [w_{k\tau}(x_t, x_{t+k\tau}) O(x_t, x_{t+k\tau})]$, we
786 can also prove $w_{k\tau}(x_t, x_{t+k\tau}) = \mathbb{E}_{\tilde{\rho}_{(k-1)\tau}(x_t, z), \tilde{\mu}(\mathbf{x}_{t+(k-1)\tau, \tau} | z)} [c_t]$, which provides the theoretical
787 consistency across all lag scales.

788 Let $\{(x^i, y^i)\}_{i=1}^M$ be endpoint pairs sampled under $\tilde{\mu}$ and their pathwise weights drawn from $c^i =$
789 $w_{(k-1)\tau}(x^i, z^i) w_{\tau}^{\text{GR}}(\mathbf{x}_{i+(k-1)\tau, \tau}^{0:N})$. By the strong law of large numbers for Markov chains with
790 ergodic assumption (Breiman, 1960),

$$791 \sum_{i=1}^M \frac{c^i O(x^i, y^i)}{\sum_{i=1}^M c^i} \xrightarrow{\text{a.s.}} \mathbb{E}_{\tilde{\rho}_{(k-1)\tau}(x_t, z), \tilde{\mu}(\mathbf{x}_{t+(k-1)\tau, \tau} | z)} [c_t O(x_t, x_{t+k\tau})] = \mathbb{E}_{\rho_{k\tau}(x_t, x_{t+k\tau})} [O(x_t, x_{t+k\tau})].$$

806 D TRAINING ALGORITHM

807 To summarize the MGR procedure, we present the full training and evaluation process in Algo-
808 rithm 1. At each iteration indexed by lag time $k\tau$, the goal is to train a classifier to estimate the

810 marginal density ratio $w_{k\tau}(x_t, x_{t+k\tau})$ between the original and perturbed transition distributions.
 811 This is achieved by constructing pathwise weights c_t that combine the short-time Girsanov weight
 812 w_τ^{GR} with the model prediction from the previous iteration $w_{(k-1)\tau}$. The resulting $w_{k\tau}$ is then used
 813 in the next iteration, allowing the model to progressively extend from short to long lag times.

Algorithm 1 Marginal Girsanov Reweighting (MGR)

816 **Require:** Simulation trajectory $\{x_t\}_{t=0}^T$ from perturbed dynamics, lagtime τ , the maximum training
 817 iteration K , learning rate η .

818 1: **for** $k = 1$ to K **do**

819 **Step 1: Collect paired data**

820 2: Collect training pairs $\{(x_t, x_{t+k\tau})\}_{t=0}^{T-k\tau}$;

821 **Step 2: Compute pathwise weights**

822 3: Compute Girsanov weight $w_\tau^{\text{GR}}(\mathbf{x}_{t+(k-1)\tau, \tau}^{0:N})$ using Eq. 2;

823 4: Compute pathwise weight:

824 5: **if** $k = 1$ **then**

825 6: $c_t = w_\tau^{\text{GR}}(\mathbf{x}_{t+(k-1)\tau, \tau}^{0:N})$;

826 7: **else**

827 8: $c_t = w_{(k-1)\tau}(x_t, x_{t+(k-1)\tau}) w_\tau^{\text{GR}}(\mathbf{x}_{t+(k-1)\tau, \tau}^{0:N})$;

828 9: **end if**

829 **Step 3: Train density ratio estimator**

830 10: **for** each training epoch **do**

831 11: **for** each minibatch $\{(x^{(i)}, y^{(i)}, c^{(i)})\}_{i=1}^B$ drawn from $\{(x_t, x_{t+k\tau}, c_t)\}_{t=0}^{T-k\tau}$ **do**

832 12: Compute weighted binary classification loss with normalized weight $c^{(i)}$:

$$833 \quad 834 \quad 835 \quad \mathcal{L}(\theta) = -\frac{1}{B} \sum_{i=1}^B \left[c^{(i)} \log h_\theta(x^{(i)}, y^{(i)}) + \log(1 - h_\theta(x^{(i)}, y^{(i)})) \right];$$

836 13: Update parameters: $\theta \leftarrow \theta - \eta \cdot \nabla_\theta \mathcal{L}(\theta)$;

837 14: **end for**

838 15: **end for**

839 16: Update the marginal ratio $w_{k\tau}(x_t, x_{t+k\tau}) = \frac{h_\theta^*(x_t, x_{t+k\tau})}{1 - h_\theta^*(x_t, x_{t+k\tau})}$;

840 17: **end for**

841 18: **return** Marginal ratio $w_{k\tau}(x_t, x_{t+k\tau})$ by model

E EXPERIMENTAL DETAILS

E.1 MOLECULAR DYNAMICS

843 In this section, we provide further details on the experimental setup, evaluation metrics, and implementation of MGR in the molecular dynamics (MD) setting described in Section 5.1.

844 For each observed transition pair $(x_t, x_{t+\tau})$, a marginal weight $w_\tau(x_t, x_{t+\tau})$ is assigned. The cross-correlation then becomes

$$845 \quad 846 \quad 847 \quad \begin{aligned} C_{ij}(\tau) &= \mathbb{E}_\mu [\mathbf{1}_{B_i}(x_t) \mathbf{1}_{B_j}(x_{t+\tau})] \\ &= \mathbb{E}_{\tilde{\mu}} [w_\tau(x_t, x_{t+\tau}) \mathbf{1}_{B_i}(x_t) \mathbf{1}_{B_j}(x_{t+\tau})] \\ &\approx \frac{\sum_t w_\tau(x_t, x_{t+\tau}) \mathbf{1}_{B_i}(x_t) \mathbf{1}_{B_j}(x_{t+\tau})}{\sum_t w_\tau(x_t, x_{t+\tau})}. \end{aligned}$$

848 Here $w_\tau(x_t, x_{t+\tau})$ is estimated by MGR in our model. In the GR baseline, $w_\tau(x_t, x_{t+\tau})$ is computed via pathwise Girsanov reweighting. Once the corrected matrix is constructed, standard MSM analysis can be performed.

849 We adopt the following metrics, as described in the main text, to assess the performance of MGR:

850 **Effective Sample Size (ESS):** ESS is a popular metric to measure the variance of weights and is
 851 regarded as a metric for quantifying goodness of reweighting ratio on importance sampling. Denoting $\{w_{k\tau}(x^i, y^i)\}_{i=1}^M$ the weights assigned to data pairs, we compare the relative ESS, defined

864 as

$$\text{rESS}_{k\tau} = \frac{1}{M} \text{ESS}_{k\tau} = \frac{(\sum_t w_{k\tau}(x_t, x_{t+\tau}))^2}{M \sum_t w_{k\tau}(x_t, x_{t+\tau})^2}.$$

865 The relative ESS value lies in the interval $(0, 1]$, and a higher value indicates reduced variance and
 866 greater weight stability.
 867

870 **Implied Timescales (ITS):** ITS is obtained by eigendecomposition of the reweighted transition
 871 probability matrix. We evaluate both individual timescales and their cumulative behavior as a func-
 872 tion of lag time τ , comparing with the unbiased reference to assess how well dynamical relaxation
 873 is captured.
 874

875 For a lagtime τ transition matrix, let $1 = \lambda_1(\tau) > \lambda_2(\tau) \geq \lambda_3(\tau) \geq \dots$ be the leading eigenvalues.
 876 Each eigenvalue defines an implied timescale
 877

$$t_i(\tau) = -\frac{\tau}{\ln \lambda_i(\tau)}.$$

879 Each t_i measures the relaxation time of a distinct slow dynamical mode: larger t_i means slower
 880 decay (stronger metastability).
 881

882 Beyond inspecting individual t_i , it is important to track an aggregate relaxation measure over the
 883 first m slow modes:
 884

$$S_m(\tau) := \sum_{i=2}^{m+1} \lambda_i(\tau) = \sum_{i=2}^{m+1} \exp\left(-\frac{\tau}{t_i(\tau)}\right).$$

885 This quantity decays exponentially with lagtime τ , because $t_i(\tau)$ of a slow mode is usually approx-
 886 imately constant across τ . Accurate tracking of this decay indicates that the model captures the
 887 intrinsic relaxation behavior of the dynamics (i.e., the spectrum of slow processes). In practice, we
 888 report: (i) several dominant ITS $t_i(\tau)$ to assess mode-wise accuracy, and (ii) the aggregate $S_m(\tau)$
 889 over the first m ITS, to summarize overall slow relaxation captured by the estimated transition ma-
 890 trix.
 891

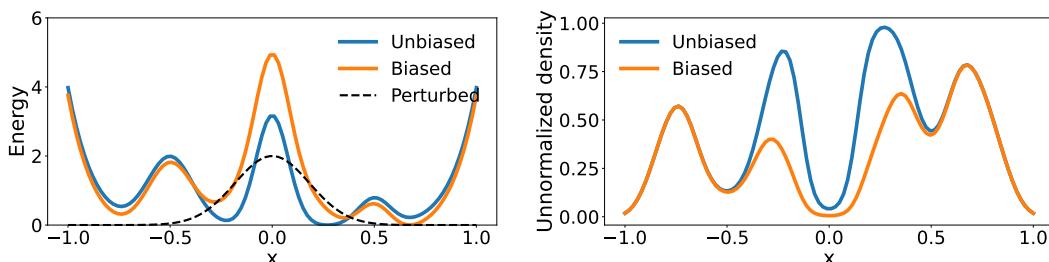
892 **Dominant eigenfunctions and stationary distribution:** For a lagtime τ transition matrix P , the
 893 stationary distribution π satisfies $\pi^\top P = \pi^\top$ and represents the long-time equilibrium. Equiv-
 894 alently, π is the first left eigenvector ϕ_1 of P with eigenvalue 1. The remaining leading eigen-
 895 functions (right eigenfunctions $\{\psi_2, \psi_3, \dots\}$, left eigenfunctions $\{\phi_2, \phi_3, \dots\}$) with eigenvalues
 896 $1 > \lambda_2 \geq \lambda_3 \geq \dots$ encode the slow dynamical modes: their sign pattern and level sets parti-
 897 tion state space into metastable regions and provide approximate coordinates. Accurate recovery of
 898 π validates thermodynamics, and accurate recovery of $\{\psi_i, \phi_i\}$ validates kinetics.
 899

900 E.1.1 1 DIMENSIONAL FOUR WELL

901 We consider an overdamped Langevin dynamics on a one-dimensional four-well landscape,
 902

$$dX_t = -\nabla V(X_t) dt + \sigma dW_t,$$

903 where $V(x) = 4 \left(x^8 + 0.8e^{-80x^2} + 0.2e^{-80(x-0.5)^2} + 0.5e^{-40(x+0.5)^2} \right)$, $\sigma = 1$.
 904



915 Figure S.1: Energy and density of 1D four-well system. **Left:** Energy profiles of unbiased poten-
 916 tial $V(x)$, biased potential $\tilde{V}(x) = V(x) + U(x)$, and the bias term $U(x)$. **Right:** Theoretical
 917 unnormalized probability of the system.

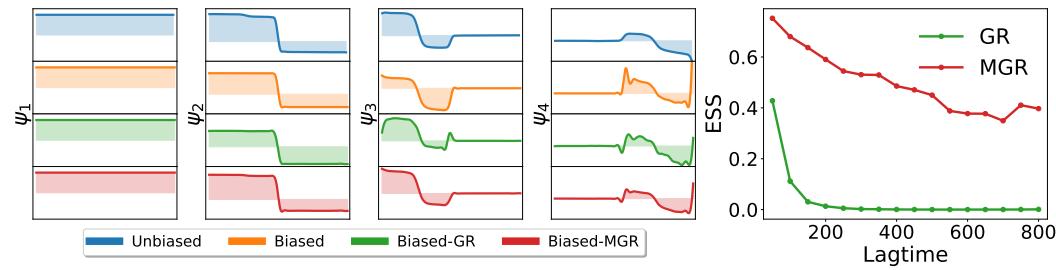


Figure S.2: **Left:** Dominant right ($\{\psi_i\}$) eigenfunctions of the transition matrix at lagtime $300\Delta t$. **Right:** Relative ESS as a function of lag time comparing GR and MGR.

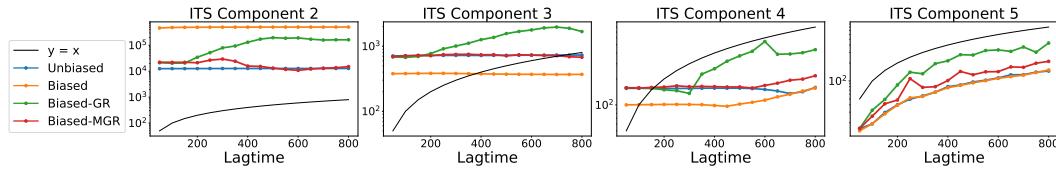


Figure S.3: The first four implied timescales (ITS) versus lag time. Each panel reports $t_i(\tau)$ for mode $i = 2, \dots, 5$ (log scale). MGR tracks the unbiased ITS with less fluctuation, while GR deviates at long lagtimes.

To accelerate barrier crossings between the two intermediate wells, we introduce a perturbed potential

$$dX_t = -\nabla \tilde{V}(X_t) dt + \sigma dW_t,$$

where $\tilde{V}(x) = V(x) + U(x)$, $U(x) = 2e^{-kx^2}$, $k = 15$. Illustrations of the potentials and stationary densities are provided in Figure S.1.

We use the Euler–Maruyama scheme with time step $\Delta t = 10^{-3}$. Each run is integrated up to $T = 10,000$, yielding $T/\Delta t = 10^7$ samples per trajectory. We initialize at $X_0 = 0$ and simulate a single trajectory. During the perturbed run, we record the discrete time Girsanov log-weights according to Eq. 2, i.e.,

$$\log w_{\Delta t}^{\text{GR}}(x) = -\nabla U(x)\sqrt{\Delta t} \xi - \frac{\Delta t}{2} (\nabla U(x))^2,$$

where ξ is the corresponding noise in the simulation.

We set $\tau = 50\Delta t$ as the reference short GR lagtime during training. At this short lagtime, the Girsanov weights are numerically stable, with $\text{ESS} \approx 0.43$. Guidelines for selecting a suitable short lag τ are provided in Appendix F.2. In each training iteration k , we normalize the pathwise weights

$$c_t = \frac{w_{(k-1)\tau}(x_t, x_{t+(k-1)\tau}) w_{\tau}^{\text{GR}}(\mathbf{x}_{t+(k-1)\tau, \tau}^{0:N})}{\frac{1}{M} \sum_t (w_{(k-1)\tau}(x_t, x_{t+(k-1)\tau}) w_{\tau}^{\text{GR}}(\mathbf{x}_{t+(k-1)\tau, \tau}^{0:N}))}, \quad (\text{S.1})$$

where M is the number of total paired data in our training dataset. After training, we obtain the estimated marginal ratio $w_{k\tau}(x_t, x_{t+k\tau}) = \frac{h_{\theta}^{\tau}(x_t, x_{t+k\tau})}{1 - h_{\theta}^{\tau}(x_t, x_{t+k\tau})}$ for reweighting.

The reweighted transition counts yield an approximation of the unbiased transition matrix, enabling standard MSM analysis. The dominant right eigenfunctions and the relative Effective sample size (ESS) as a function of lag time are shown in Figure S.2 and Figure S.3 reports the first 4 implied timescales individually. Furthermore, Figure S.4 displays the density plot of the transfer operator, which can also be reweighted by MGR ratio

$$\begin{aligned} \rho_{k\tau}(x, y) &= \rho(x) p_{k\tau}(x, y) \\ &= e^{-V(x)} \tilde{\rho}(x) w_{k\tau}(x, y) \tilde{p}_{k\tau}(x, y) \\ &= e^{-V(x)} w_{k\tau}(x, y) \tilde{\rho}_{k\tau}(x, y). \end{aligned}$$

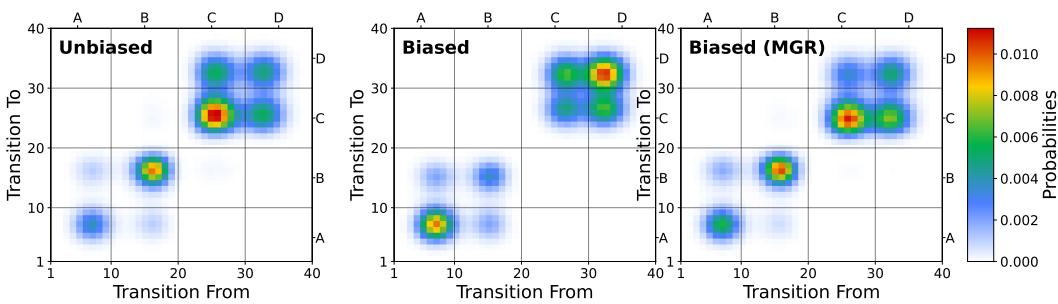


Figure S.4: Density plot of the transfer operator for Four Well, defined on the discrete spatial range $[0, 40]$ with a lag time of $300\Delta t$. Red indicates high transition probability, white zero transition probability.

In a four-well landscape, the first three slow kinetic modes are important. The slowest ITS (ITS component 2) corresponds to the global left-right rearrangement across the highest barriers. ITS component 3 and the third eigenfunctions in Figures 3, S.2 capture the exchanges between the two left wells (left-pair exchange), while ITS component 4 and associated eigenfunctions resolve the remaining transfer with the two right wells (right-pair exchange). The rest of ITS is small and unstable across lagtimes, reflecting fast relaxation rather than a meaningful slow process.

E.1.2 ALANINE DIPEPTIDE

We performed all-atom MD simulations of acetyl-alanine-methylamide (Ac-A-NHMe, alanine dipeptide) in implicit water. The simulation was carried out with the OPENMM 8.2 simulation package (Eastman et al., 2023) at 300K. The system employed the Amber14 force field with OBC2 implicit water (“`amber14-all.xml`”, “`implicit/obc2.xml`”). Dynamics were propagated with an Underdamped Langevin integrator with time step 2fs. The aggregated simulation time was $1\mu\text{s}$. Coordinates and Girsanov reweighting (GR) factors were saved every 20 steps (40fs) using the Girsanov-enabled OPENMM implementation (Schafer & Keller, 2024).

For enhanced sampling, we applied a dihedral bias to the backbone torsions ϕ and ψ . The unbiased dihedral potentials were

$$\begin{aligned} V(\phi) &= 0.27 \cos(2\phi) + 0.42 \cos(3\phi), \\ V(\psi) &= 0.45 \cos(\psi - \pi) + 1.58 \cos(2\psi - \pi) + 0.44 \cos(3\psi - \pi), \end{aligned}$$

and the perturbation was a quadratic restraint

$$U(\phi, \psi) = \frac{1}{2} \kappa_\phi \phi^2 + \frac{1}{2} \kappa_\psi \psi^2, \quad \kappa_\phi = \kappa_\psi = 1,$$

so that the biased potential is $\tilde{V}(\cdot) = V(\cdot) + U(\phi, \psi)$. The perturbation lowers the energy barrier in targeted regions of the (ϕ, ψ) free-energy surface, thereby facilitating transitions among metastable basins (Figure S.5), but it also distorts equilibrium and kinetics.

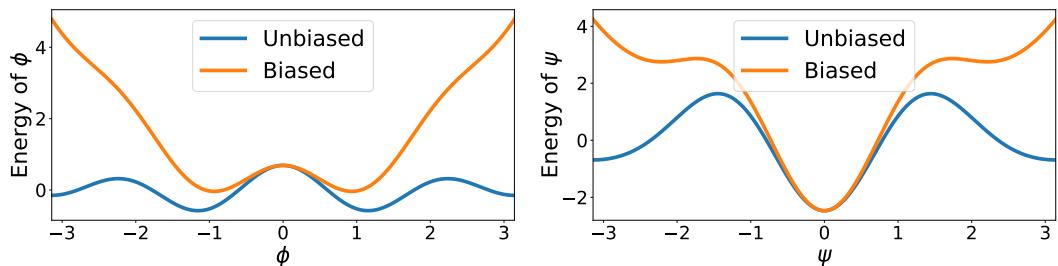


Figure S.5: Alanine dipeptide dihedral energy function. **Left:** ϕ torsion: unbiased energy $V(\phi)$ and biased function $V(\phi) + \frac{1}{2}\kappa_\phi\phi^2$. **Right:** ψ torsion: unbiased energy $V(\psi)$ and biased function $V(\psi) + \frac{1}{2}\kappa_\psi\psi^2$.

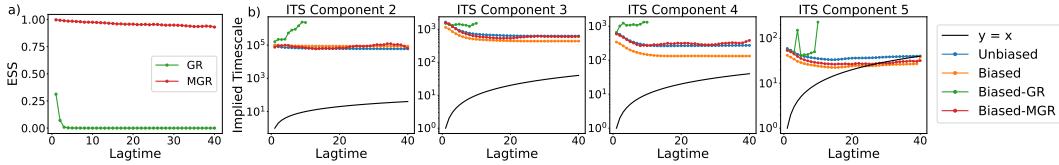


Figure S.6: Relative ESS and the first four implied timescales (ITS) versus lag time in Alanine dipeptide. a) Relative ESS as a function of lag time comparing GR and MGR. b) Each panel reports $t_i(\tau)$ for mode $i = 2, \dots, 5$ (log scale). MGR tracks the unbiased ITS with less fluctuation, while GR shows huge deviation. GR totally fails to construct a valid MSM beyond lag time longer than $10\tau = 400\text{fs}$.

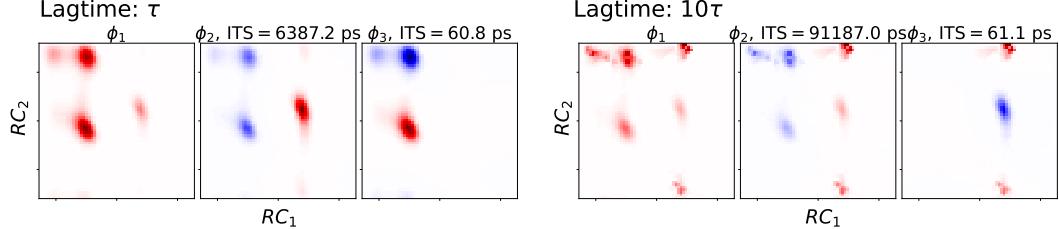


Figure S.7: Dominant eigenfunctions at lagtime τ and 10τ of biased trajectories, reweighting by GR.

We set $\tau = 40\text{fs}$ as the reference short GR lagtime during training. At this short lagtime, the Girvanov weights are numerically stable, with $\text{ESS} \approx 0.31$. In each training iteration k , we normalize the pathwise weights according to Eq. S.1.

The reweighted transition counts yield an approximation of the unbiased transition matrix, enabling standard MSM analysis. For alanine dipeptide, we construct the MSM in the key reaction coordinates (ϕ, ψ) space. In our model, we likewise use only two angles as inputs to the ratio estimator, which yields satisfactory performance.

The results are shown in Figure 4, and Figure S.6 further reports ESS and the first 4 ITS individually. As the lagtime increases, the variance of pathwise GR grows rapidly and the effective sample size collapses. Beyond 10τ , GR fails to produce a stable transition matrix and thus cannot construct a valid MSM. For illustration, we plot the dominant eigenfunctions at τ and 10τ reweighted by GR in Figure S.7. At τ , GR yields corrected modes, but at 10τ the modes become distorted and metastable partitions blur. In contrast, MGR retains stable eigenfunctions across lags as shown in Figure 4 that closely match the unbiased reference, consistent with its higher ESS and accurate ITS.

E.2 BAYESIAN INFERENCE OF SDE PARAMETERS

Consider a d -dimensional diffusion process that satisfies the following SDE:

$$dX_t = f(X_t, t; \theta)dt + g(t)dW_t, \quad X_0 \sim p_0,$$

where $X_t \in \mathbb{R}^d$ denotes the state at time t , $f(\cdot, t; \theta)$ is the drift depending on the unknown parameter vector θ , $g(\cdot)$ is the diffusion coefficient, and W_t is a standard Wiener process. Suppose we observe noisy measurements $\mathbf{y} = \{y(t=0) = y^0, y^1, \dots, y(t=T) = y^K\}$ at equally spaced time intervals with lag time T/K . The goal is to infer the unknown parameter vector θ from these observations. Within the Bayesian framework, we assign a prior distribution $p(\theta)$ and aim to recover the posterior

$$p(\theta | \mathbf{y}) \propto p(\mathbf{y} | \theta) p(\theta).$$

The main difficulty lies in computing the likelihood $p(\mathbf{y} | \theta)$. Since the process $\mathbf{X} = \{x(t=0) = x^0, x^1, \dots, x(t=T) = x^K\}$ is governed by the SDE with parameter θ , the likelihood requires integrating over all possible trajectories,

$$p(\mathbf{y} | \theta) = \int p(\mathbf{y} | \mathbf{X}, \theta) p(\mathbf{X} | \theta) d\mathbf{X},$$

which is generally intractable.

Instead of simulating the SDE for every candidate θ , we set a reference parameter θ_0 . We then generate reference trajectories under θ_0 and compute the ratio of likelihood ratio between candidates θ to θ_0 . According to Girsanov reweighting in Section 3.2, the ratio $\frac{dp(\mathbf{X}|\theta)}{dp(\mathbf{X}|\theta_0)}$ can be recorded during the simulation, where the drift difference in Eq. 2 is $f(X_t, t; \theta) - f(X_t, t; \theta_0)$.

Through the MGR model introduced in Section 4, we approximate the ratio with the given observation pair (y_i, y_{i+1})

$$w_{k\tau}^\theta(\mathbf{y}) = \frac{p(y^{0:K}|\theta)}{p(y^{0:K}|\theta_0)} = \Pi_i \frac{p(y^{i+1}|y^i; \theta)}{p(y^{i+1}|y^i; \theta_0)} = \Pi_i w_{k\tau}^\theta(y_i, y_{i+1}). \quad (\text{S.2})$$

Consequently, posterior expectations can be evaluated by importance sampling using observations under θ_0 :

$$\begin{aligned} \mathbb{E}[\theta | \mathbf{y}] &= \int \theta p(\theta | \mathbf{y}) d\theta \\ &= \int \theta \frac{p(\mathbf{y} | \theta)p(\theta)}{p(\mathbf{y})} d\theta \\ &= \int \theta \frac{p(\mathbf{y} | \theta)}{p(\mathbf{y} | \theta_0)} \frac{p(\mathbf{y} | \theta_0)p(\theta)}{p(\mathbf{y})} d\theta \\ &= \frac{\int \theta \frac{p(\mathbf{y} | \theta)}{p(\mathbf{y} | \theta_0)} p(\theta)p(\mathbf{y} | \theta_0) d\theta}{\int \frac{p(\mathbf{y} | \theta)}{p(\mathbf{y} | \theta_0)} p(\theta)p(\mathbf{y} | \theta_0) d\theta} \\ &\approx \frac{\sum_\theta \theta w_{k\tau}^\theta(\mathbf{y})p(\theta)}{\sum_\theta w_{k\tau}^\theta(\mathbf{y})p(\theta)}. \end{aligned} \quad (\text{S.3})$$

For the baseline particle marginal Metropolis–Hastings (PMMH) algorithm (Golightly & Wilkinson, 2008; Hoffman et al., 2014) and a variational inference (VI) approach (Ghosh et al., 2022). We use the default setting in their original paper.

E.2.1 GRAPH ORNSTEIN-UHLENBECK

We consider a two dimensional Graph Ornstein-Uhlenbeck (OU) process as introduced in Section 5.2.1. The Euler–Maruyama scheme with time step $\Delta t = 10^{-3}$ is used to generate a trajectory and the simulation is integrated up to $T = 30$, yielding $M = T/\Delta t = 3 \times 10^4$ samples. A set of 30 evenly spaced values constitute the observations \mathbf{y} , and each observation is corrupted with Gaussian noise of $\sigma = 0.5$, shown in Figure S.8. We consider 15 candidate values of θ^1 in the range $[0.1, 1.5]$ with increments of 0.1, and 15 values of θ^2 in the range $[0.5, 7.5]$ with increments of 0.5. The resulting grid consists of 225 parameter combinations, and for each of them we compute the corresponding Girsanov reweighting during simulation.

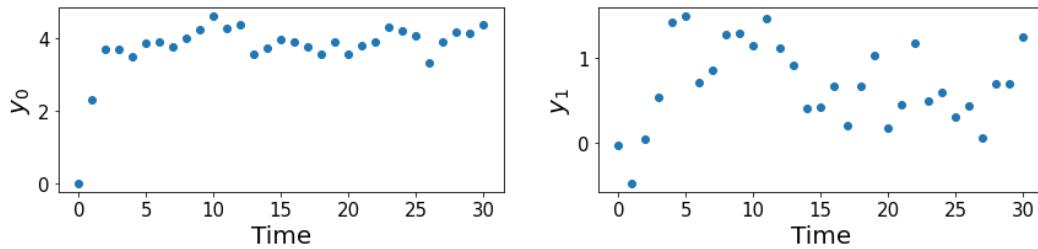


Figure S.8: Observations of Graph Ornstein-Uhlenbeck.

E.2.2 STOCHASTIC LOTKA-VOLTERRA

We consider a two dimensional stochastic Lotka-Volterra system as introduced in Section 5.2.2. The Euler–Maruyama scheme with time step $\Delta t = 10^{-3}$ is used to generate a trajectory and the simulation is integrated up to $T = 50$, yielding $M = T/\Delta t = 5 \times 10^4$ samples. A set of 25 evenly spaced values from this path corrupted with Gaussian noise of $\sigma = 5$ constitute the observations \mathbf{y} ,

1134 shown in Figure S.9. We consider 15 candidate values of θ^1 in the range $[0.3, 0.5]$ with increments
 1135 of 0.014, and 15 values of $100 \times \theta^2$ in the range $[0.2, 0.4]$ with increments of 0.014. The resulting
 1136 grid consists of 225 parameter combinations, and for each of them we compute the corresponding
 1137 Girsanov reweighting during simulation.

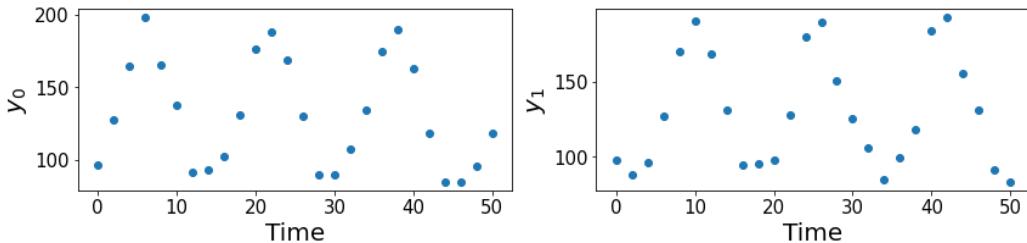


Figure S.9: Observations of Stochastic Lotka-Volterra.

1149 E.3 MODEL ARCHITECTURE

1151 Our classifier estimator $h_\theta(\cdot, \cdot)$ in MGR is modeled by 3 layer MLPs, augmented with Fourier feature
 1152 encodings (Tancik et al., 2020). We consider two types of encoding $\gamma(\mathbf{v})$, ($\mathbf{v} = [x_t, x_{t+k\tau}]$ is the
 1153 concatenation of the paired data):

1154 **Positional encoding:** $\gamma(\mathbf{v}) = [\sin(\mathbf{v}), \cos(\mathbf{v}), \sin(2\mathbf{v}), \cos(2\mathbf{v}), \dots, \sin(B\mathbf{v}), \cos(B\mathbf{v})]$, where
 1155 $B = 10$ denotes the scaling number.

1156 **Gaussian encoding:** $\gamma(\mathbf{v}) = [\cos(B\mathbf{v}), \sin(B\mathbf{v})]$, where $B \in \mathbb{R}^{m \times 2d}$ is sampled from $\mathcal{N}(0, 10)$.

1158 Both Fourier-feature variants converge faster during training and yield comparable estimation accuracy.
 1159 Detailed comparison is reported in Appendix F.4. In our experiments, we adopt a positional-
 1160 encoding MLP with ReLU activations.

1162 F ABLATION

1164 F.1 NEURAL RATIO ESTIMATION

1166 Ratio estimation (Sugiyama et al., 2010) is a fundamental technique for comparing two distributions.
 1167 Kernel moment matching, e.g. KMM (Gretton et al., 2009), matches all the moments with
 1168 reproducing kernels, which is effective and computationally efficient. Probabilistic classification
 1169 recasts ratio estimation as posteriors from a binary classifier (Menon & Ong, 2016), showing pow-
 1170 erful fitting capability. Featurized classification with normalizing flows (Choi et al., 2021) further
 1171 performs classification in a learned latent space, mitigating issues caused by large distributional dis-
 1172 crepancies. Path-based methods (Choi et al., 2022; Yu et al., 2025) connect the two distributions
 1173 via a continuous probability path and estimate the density ratio by integrating a learned time score.
 1174 By constructing consecutive path distributions, it alleviates the problems caused by poor overlap
 1175 between two densities.

1176 However, unlike the standard setting with samples from both distributions, here we only have sam-
 1177 ples from one distribution plus reference weights linking two distributions. We therefore conducted
 1178 minor adaptions to estimators below and compare their performance on Four well system. We also
 1179 adapted the path-based method (Yu et al., 2025), but it exhibited numerical instability in our setting.

1181 **Standard classifier (weighted BCE)** Following Section 4.3, we train a binary classifier on end-
 1182 point pairs $(x_t, x_{t+k\tau})$ with weighted cross-entropy in Eq. 6.

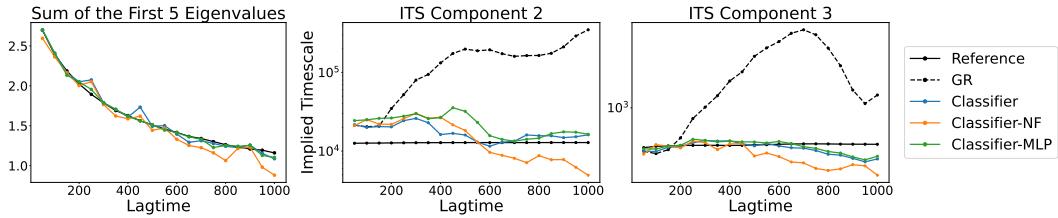
1184 **Featurized classifier (weighted BCE)** We first map each paired sample to a latent representation
 1185 $z_\phi = \Phi(x_t, x_{t+k\tau}; \phi)$, and then perform the classifier-based ratio estimation in this feature space. A
 1186 joint training objective is adopted (Choi et al., 2021):

$$1187 \mathcal{L}_{\text{joint}} = \alpha \mathcal{L}_{\text{BCE}}(\theta, \phi) + (1 - \alpha) \mathcal{L}_{\text{latent}}(\phi),$$

1188 where $\mathcal{L}_{\text{BCE}}(\theta, \phi) = -\mathbb{E}_t [c_t \log h_\theta(z_\phi) + \log(1 - h_\theta(z_\phi))]$ is the weighted binary cross entropy
 1189 in latent space, $\mathcal{L}_{\text{latent}}(\phi)$ denotes the objective for optimizing encoder network, and $\alpha = 0.5$ is a
 1190 hyperparameter.

1191 Here, we consider two encoders to map data into latent space: (i) an invertible normalizing-flow
 1192 encoder (Classifier-NF) (Choi et al., 2021), and (ii) a non-bijective MLP encoder (Classifier-MLP).
 1193 Classifier-NF guarantees ratios computed in feature space are equivalent to those in input space,
 1194 whereas Classifier-MLP uses more flexible, non-invertible networks at the cost of potential infor-
 1195 mation loss. For Classifier-NF, we train the encoder by maximum likelihood loss $\mathcal{L}_{\text{latent}}^{\text{KL}}(\phi)$ (Dinh
 1196 et al., 2016; Choi et al., 2021), and for Classifier-MLP, we minimize a sliced 2-wasserstein distance
 1197 $\mathcal{L}_{\text{latent}}^{\text{Wass}}(\phi)$ (Kolouri et al., 2018) between the latent variables and a standard Gaussian.

1198 Comparative results are reported in Figure S.10, where featurized classifiers did not show measur-
 1199 able improvement over a standard classifier. We adopt the standard classifier in Section 4.3, which
 1200 provides satisfying results. Alternative advanced ratio estimators, and their applications, require
 1201 further investigation.

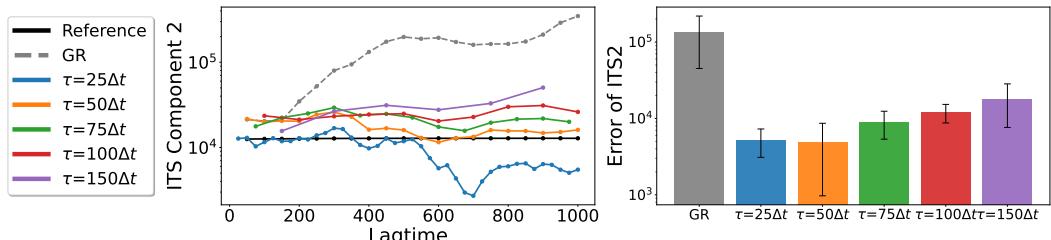


1202 Figure S.10: Neural ratio estimators ablation on the 1D Four well system. **Left:** Sum of the first
 1203 5 eigenvalues as a function of lagtime across different ratio estimators. **Right:** The first 2 ITS as a
 1204 function of lagtime across different ratio estimators.

F.2 LAGTIME CHOSEN

1205 Selecting the suitable short lag τ used in w_τ^{GR} is crucial for MGR. Large τ will inflate the variance
 1206 of Girsanov weights and can destabilize training, while small τ will require many iterations to reach
 1207 long timescales, accumulating approximation error and cost. As a practical rule, we recommend
 1208 choosing τ , which the relative ESS of the Girsanov weights falls in the range $0.3 \sim 0.5$. This strikes
 1209 a balance between weight degeneracy and excessive iteration depth.

1210 Take Four well system as an example. We train MGR with short-lag values $\tau \in \{25, 50, 75, 100,\$
 1211 $150\}\Delta t$. The corresponding relative ESS of w_τ^{GR} is $\{0.73, 0.43, 0.24, 0.11, 0.06\}$. All other settings
 1212 are kept identical. We then compare the dominant second implied timescale (ITS2) across evaluation
 1213 lags in Figure S.11. It shows that the model trained under $\tau = 50\Delta t$ produces the most stable ITS2
 1214 and matched the unbiased (reference) result most closely. As τ increases, the discrepancy between
 1215 the model results and the unbiased reference grows.



1216 Figure S.11: Lagtime ablation on the 1D Four well system. **Left:** ITS2 as a function of lagtime for
 1217 models trained with different short-lag values τ . **Right:** Mean and standard deviation of the ITS2
 1218 error, aggregated over all evaluation lags for models trained with different short-lag values τ .

F.3 THE RANGE OF CANDIDATE PARAMETERS

1219 A practical question in Bayesian inference is how to determine an appropriate range of candidate
 1220 parameters around the reference value. In principle, the Girsanov reweighting estimator is valid for
 1221 any parameter shift, but its variance grows rapidly when the reference and target dynamics exhibit

little overlap. Usually, the candidate parameters should typically be chosen as small perturbations around the reference value so that the path distributions remain close and the reweighting remains stable.

To investigate this effect, we performed an ablation study on the GraphOU and Lotka–Volterra (LV) SDEs using a broader candidate range. For each parameter value, we computed the effective sample size (ESS) of the raw Girsanov weights, which quantifies the degree of overlap between the reference and target dynamics. As shown in Figure S.12, the ESS values decay sharply as the parameter deviates from the reference, indicating that reliable reweighting is possible within a suitable perturbation region.

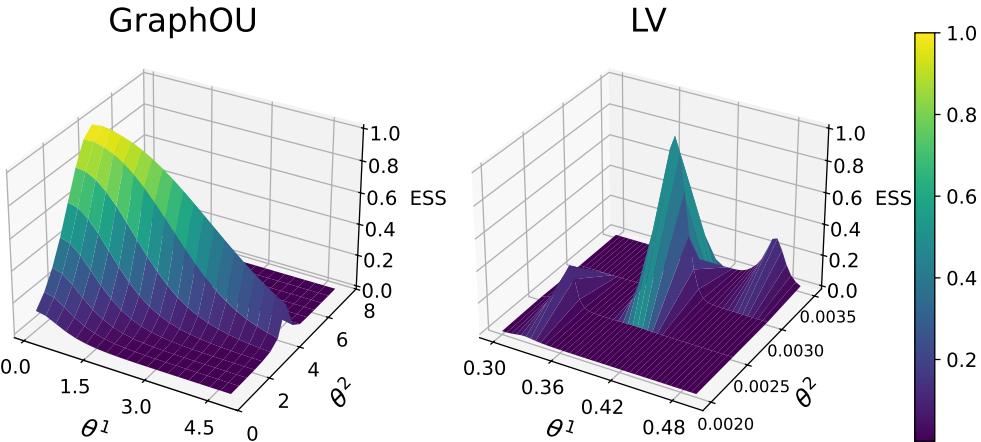


Figure S.12: ESS as a function of candidate parameters. ESS values computed from the raw Girsanov weights for the GraphOU (left) and LV (right) SDE examples.

F.4 MODEL ARCHITECTURE

To demonstrate network efficiency, we compare a plain MLP, a positional encoding MLP, and a Gaussian encoding MLP on Four well system in Figure S.13. Both Fourier-feature models (Positional and Gaussian) converge faster than the plain MLP (Base). Three models show comparable estimation accuracy during evaluation.

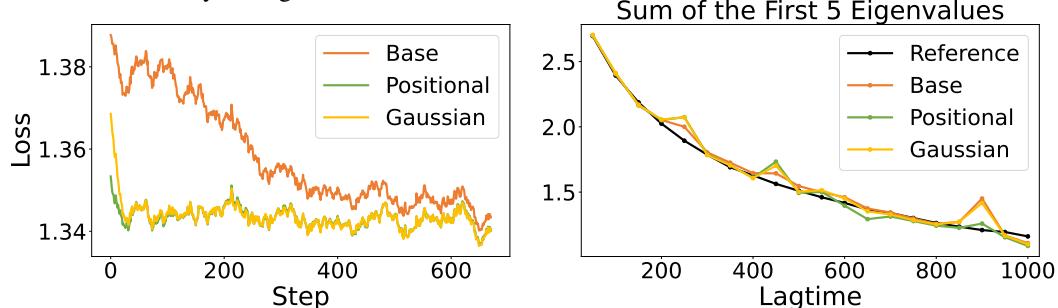


Figure S.13: Network ablation on the 1D Four well system. **Left:** Training curve as a function of training steps across different networks. **Right:** Sum of the first 5 eigenvalues as a function of lagtime across different networks.