ENCORE: EXACT NON-EQUILIBRIUM CONTROL WITH REPLICA EXCHANGE FOR DIFFUSION GENERATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Inference-time control steers a pretrained generative model towards a target distribution without retraining. We study tilted targets $\pi_0 \propto G_0\, p_0$, where p_0 is the sampler output distribution and G_0 is an evaluable reweighting function. Existing approaches rely on sequential annealing with sequential Monte Carlo (SMC) or parallel annealing with replica exchange (RE). Sequential control is exact but needs large particle populations, whereas no exact parallel control method exists: existing RE corrections approximate an intractable time reversal and are biased. We propose ***E**xact **N**on-equilibrium **CO**ntrol with **R**eplica **E**xchange* (ENCORE), the first exact parallel control method. Each replica stores its generation trajectory, so the upward move is a truncation and the intractable time reversal is never simulated. We prove target invariance and show that the resulting dynamics are those of non-equilibrium replica exchange with the exact time reversal as forward proposal. Under regularity conditions, our diffusion analysis shows that both sequential and parallel control become unstable under refinement of the time discretisation without guidance, whereas guided proposals remain stable and yield diagnostics for tuning the schedule and the computational budget. Across synthetic targets, Boltzmann sampling of biomolecules, and image generation, ENCORE achieves competitive accuracy and diversity, remains robust to sampler perturbations, and applies to distilled samplers where existing RE corrections are unavailable.

1 INTRODUCTION

Diffusion models (Ho et al., 2020; Song et al., 2021) have achieved strong performance in image (Dhariwal & Nichol, 2021), video (Ho et al., 2022), and molecular generation (Hoogeboom et al., 2022). However, downstream tasks often require samples to align with preferences, observations, or constraints. *Inference-time control* steers generation without retraining the model. Many such tasks, including reward tilting (Uehara et al., 2025) and posterior sampling (Wu et al., 2023a; Chung et al., 2023), can be formulated as endpoint reweighting,

$$\pi_0(\mathrm{d}x_0) \propto G_0(x_0)\, p_0(\mathrm{d}x_0), \tag{1.1}$$

where p_0 is the sampler's output distribution and G_0 is a non-negative, evaluable reweighting function. For instance, it steers image generation towards human preferences (Xu et al., 2023; Singhal et al., 2025), while in biomolecular applications it supports sampling at different temperatures (Skreta et al., 2025; He et al., 2026b) and exploring rare conformations for free-energy estimation (Xie et al., 2026).

Guidance (Dhariwal & Nichol, 2021; Ho & Salimans, 2022) steers generation by modifying the reverse-time score or drift. The exact correction is the Doob transform (Del Moral, 2004), but it depends on an intractable expectation over the remaining generation. In practice, guidance methods rely on approximations and therefore need not preserve the intended endpoint distribution (Bradley & Nakkiran, 2024; Dandapanthula & Boffi, 2026).

Monte Carlo methods provide a complementary approach by targeting π_0 directly. Sequential Monte Carlo (SMC). (Del Moral, 2004; Del Moral et al., 2006) propagates, reweights, and resamples a particle population through a sequence of intermediate distributions (sequential annealing) Nevertheless, SMC can suffer from weight and ancestral degeneracy, particularly with limited particle populations.

Replica exchange (RE), or parallel tempering (PT) (Swendsen & Wang, 1986; Geyer, 1991; Hukushima & Nemoto, 1996), instead maintains replicas across a ladder of intermediate distribu-

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

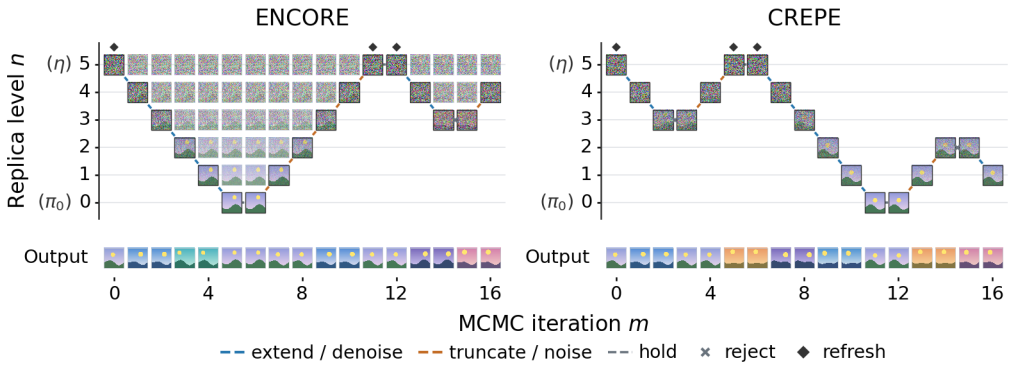

Figure 1: Replica movement (top) and output chains (bottom) for ENCORE and CREPE.

tions and exchanges neighbouring replicas through Metropolis–Hastings moves (parallel annealing). Whereas SMC propagates particles through the ladder once, RE evolves a persistent set of replicas over successive MCMC iterations (Syed et al., 2026). This gives an *online* advantage over SMC: sampling continues without restarting.

However, continued sampling recovers the intended target only if the chain preserves it. TDS (Wu et al., 2023a) and FKD (Singhal et al., 2025) compute exact SMC weights for their chosen potentials without evaluating sampler marginals. In contrast, CREPE (He et al., 2026b) approximates intermediate marginal-density ratios using denoising and noising transition densities. This substitution relies on an exact time-reversal relation that may be violated by model or discretisation error. Any resulting bias persists with longer runs. CREPE also requires a tractable paired noising kernel, which may be unavailable for, e.g., distilled few-step samplers. No exact parallel control method exists.

We address these limitations with ***E**xact **N**on-equilibrium **CO**ntrol with **R**eplica **E**xchange* (ENCORE). The key idea is to lift each replica from its current state to a partial generation trajectory. An upward exchange then truncates the stored trajectory rather than simulating a noising transition, allowing common sampler-path terms to cancel exactly in the Metropolis–Hastings ratio. The resulting replica movement across levels and target-level chains are illustrated in Fig. 1.

We show that ENCORE preserves the endpoint target without evaluating intermediate marginals or requiring a paired noising kernel. We show that ENCORE has the rejection rate of non-equilibrium replica exchange with the exact time reversal as forward proposal, which vanishes precisely under the Doob relations. For diffusion samplers, we show under regularity conditions that refining the time discretisation can cause accumulated rejection to diverge without guidance, whereas guided proposals admit a finite *global communication barrier* Λ, which measures the deviation of the guidance from the Doob transform and sets the number of levels needed. Experiments on synthetic reward tilting, molecular sampling, and few-step image generation demonstrate competitive accuracy and diversity, robustness to sampler perturbations, and applicability where existing RE corrections are unavailable.

2 Background

2.1 Diffusion Samplers as Markov Chains

We consider samplers defined by a tractable reference distribution η on a state space $\mathbb{X}$ and a sequence of Markov kernels $P_{n-1|n}$ for $n = N, \dots, 1$. The sampler draws $x_N \sim \eta$ and then $x_{n-1} \sim P_{n-1|n}(\mathrm{d}x_{n-1} \mid x_n)$ for $n = N, \dots, 1$, inducing the marginals p_n of x_n,

$$p_N = \eta, \qquad p_{n-1} = p_n P_{n-1|n}, \tag{2.1}$$

with p_0 the output distribution of the sampler.[1] Diffusion models are the motivating example (Ho et al., 2020; Song et al., 2021): $P_{n-1|n}$ is one or more steps of a learned denoiser, and p_0 approximates the data distribution. We assume only that $P_{n-1|n}$ can be simulated and that its transition density can be

[1]For a law p, a kernel K and a function f, $pK(\mathrm{d}y) := \int p(\mathrm{d}x)K(\mathrm{d}y \mid x)$, $Kf(x) := \int K(\mathrm{d}y \mid x)f(y)$, and $p \otimes K(\mathrm{d}x, \mathrm{d}y) := p(\mathrm{d}x)K(\mathrm{d}y \mid x)$ is the joint law of a draw from p propagated by K.

evaluated. In particular, the marginals p_n for $n < N$ are intractable, since they require integrating over $P_{n-1|n}$, and the reverse kernels $P_{n|n-1}$, for which $p_{n-1}P_{n|n-1} = p_n$, are unavailable.

2.2 Inference-Time Control

Given an evaluable, non-negative function $G_0 : \mathbb{X} \to [0, \infty)$ inference-time control asks for samples from the tilted target without retraining the model,

$$\pi_0(\mathrm{d}x_0) \propto G_0(x_0)\, p_0(\mathrm{d}x_0). \tag{2.2}$$

For example, reward tilting uses $G_0(x_0) = \exp\{\kappa r(x_0)\}$ for a reward r and strength $\kappa \geq 0$ (Singhal et al., 2025), while posterior sampling uses the likelihood $G_0(x_0) = p(y \mid x_0)$ of an observation y, with p_0 as the prior (Wu et al., 2023a; Chung et al., 2023).

The natural approach keeps the reference η, replaces each sampler kernel with a proposal $Q_{n-1|n}(\mathrm{d}x_{n-1} \mid x_n)$ that steers the chain towards the target, and returns x_0 as an approximate sample from π_0. General proposals introduce bias, and an exact construction follows from the Doob h-transform of the sampler by G_0 (Del Moral, 2004). Define the Doob potentials recursively and the corresponding proposals by

$$G_n^* := P_{n-1|n} G_{n-1}^*, \qquad Q_{n-1|n}^*(\mathrm{d}x_{n-1} \mid x_n) := \frac{G_{n-1}^*(x_{n-1})}{G_n^*(x_n)} P_{n-1|n}(\mathrm{d}x_{n-1} \mid x_n), \tag{2.3}$$

with $G_0^* := G_0$, so that $G_n^*(x_n) = \mathbb{E}[G_0(X_0) \mid X_n = x_n]$ under the sampler law. Running Q^* from the tilted reference $\eta^* \propto G_N^*\, \eta$ yields $x_0 \sim \pi_0$.[2] The potentials and the twisted kernels are intractable, and the tractable approximations used in practice need not preserve π_0 (Sec. 3.4). The annealed Monte Carlo methods below correct them.

2.3 Annealed Monte Carlo

Consider a ladder $\pi_N, \ldots, \pi_0$ on $\mathbb{X}$ from a tractable reference π_N to a target π_0. Annealed Monte Carlo generates, for every level n, particles x_n^m, $m = 1, \ldots, M$, whose empirical law approximates π_n. The two constructions below differ in the order in which the array (x_n^m) is filled. Sequential annealing propagates M particles level by level, so m indexes particles, and parallel annealing evolves one replica per level over M iterations, so m indexes iterations. Both move a particle from level n to $n-1$ with a backward proposal $Q_{n-1|n}(\mathrm{d}x_{n-1} \mid x_n)$, describe the reverse move with a forward proposal $Q'_{n|n-1}(\mathrm{d}x_n \mid x_{n-1})$, and correct it with the incremental weight, the Radon–Nikodym derivative of the two joint laws of (x_{n-1}, x_n) the proposals induce,

$$w_{n-1,n} := \frac{\mathrm{d}(\pi_{n-1} \otimes Q'_{n|n-1})}{\mathrm{d}(\pi_n \otimes Q_{n-1|n})} \propto \frac{\pi_{n-1}(x_{n-1})\, Q'_{n|n-1}(x_n \mid x_{n-1})}{\pi_n(x_n)\, Q_{n-1|n}(x_{n-1} \mid x_n)}. \tag{2.4}$$

Both produce consistent estimators as $M \to \infty$ and comparable accuracy for comparable N and M, so they are *computationally dual* (Syed et al., 2026). App. A gives the details.

Parallel annealing evolves N replicas, one per level, in parallel over M iterations, in $O(Nd)$ memory, and records the level-0 replica at every iteration. The general framework is non-equilibrium replica exchange (Ballard & Jarzynski, 2009), also known as accelerated parallel tempering (Zhang et al., 2026), which targets $\pi_N \otimes \cdots \otimes \pi_0$ by MCMC and, in place of the state swaps of parallel tempering (Swendsen & Wang, 1986; Geyer, 1991; Hukushima & Nemoto, 1996), proposes transitions between levels with non-reversible communication (Syed et al., 2022). Here x_n^m is the state at level n after iteration m. At iteration m the reference is redrawn, $x_N^m \sim \eta$, and for every $n \equiv m \pmod 2$ the pair $(x_{n-1}^{m-1}, x_n^{m-1})$ proposes $y_{n-1}^m \sim Q_{n-1|n}(\cdot \mid x_n^{m-1})$, $y_n^m \sim Q'_{n|n-1}(\cdot \mid x_{n-1}^{m-1})$ and sets (x_{n-1}^m, x_n^m) to (y_{n-1}^m, y_n^m) if accepted and to $(x_{n-1}^{m-1}, x_n^{m-1})$ if rejected, with acceptance probability $\alpha_{n-1,n}^m := \alpha_{n-1,n}(x_{n-1}^{m-1}, x_n^{m-1}; y_{n-1}^m, y_n^m)$ for

$$\alpha_{n-1,n}(x_{n-1}, x_n; y_{n-1}, y_n) = \min\left\{1,\ \frac{w_{n-1,n}(y_{n-1}, x_n)}{w_{n-1,n}(x_{n-1}, y_n)}\right\}. \tag{2.5}$$

The remaining levels keep $x_n^m = x_n^{m-1}$. Optionally, a π_n-invariant MCMC move is then applied to each x_n^m, in parallel over n.

[2] In diffusion models, x_N is typically nearly independent of x_0, so G_N^* is approximately constant and $\eta^* \approx \eta$.

Table 1: Monte Carlo corrections for inference-time control with M samples and N levels. *RNE is exact when taking the same construction as TDS (He et al., 2026a, App. C.5).

Method	Annealing	Exact	Parallel over	Moves/iter.	# Iter.	Memory
FKD (Singhal et al., 2025)	sequential	yes	M particles	M	N	$O(Md)$
TDS (Wu et al., 2023a)	sequential	yes	M particles	M	N	$O(Md)$
RNE (He et al., 2026a)	sequential	no / yes*	M particles	M	N	$O(Md)$
CREPE (He et al., 2026b)	parallel	no	N replicas	N	M	$O(Nd)$
ENCORE (ours)	parallel	yes	N replicas	$N/2$	M	$O(N^2d)$

Particle diversity is measured by the restart rate τ, the rate at which reference draws reach the target (Katzgraber et al., 2006). Under the efficient local exploration assumption of Syed et al. (2022), that the log-weights of consecutive exchange attempts at a level are independent,

$$\tau = \left(2 + 2\sum_{n=1}^{N} \frac{r_{n-1,n}}{1 - r_{n-1,n}}\right)^{-1} \leq \frac{1}{2}, \tag{2.6}$$

where $r_{n-1,n} := 1 - \mathbb{E}[\alpha_{n-1,n}]$ is the stationary rejection rate of the exchange at $(n-1, n)$. By Zhang et al. (2026, Thm. 1),

$$r_{n-1,n} = r\big(\pi_{n-1} \otimes Q'_{n|n-1},\ \pi_n \otimes Q_{n-1|n}\big), \qquad r(p, q) := \mathrm{TV}(p \otimes q,\ q \otimes p). \tag{2.7}$$

When the π_n discretise a continuum of annealing distributions and $N \to \infty$ with the rejection rates vanishing, $\sum_n r_{n-1,n} \to \Lambda$, the global communication barrier, and $\tau \to (2 + 2\Lambda)^{-1}$. An equi-rejection schedule attains stable communication with $N \approx 2\Lambda$ levels, beyond which additional levels do not increase τ and M should be as large as affordable (Syed et al., 2022).

Sequential annealing evolves M particles in parallel over the N levels, in $O(Md)$ memory, and terminates with a weighted batch. The general framework is sequential Monte Carlo samplers (Del Moral et al., 2006). Particle m at level n carries a weight w_n^m, the product of the incremental weights along its path. Initialise $x_N^m \sim \pi_N$ with $w_N^m = 1$. At each level move each particle down, $x_{n-1}^m \sim Q_{n-1|n}(\mathrm{d}x_{n-1} \mid x_n^m)$, update its weight, $w_{n-1}^m = w_n^m\, w_{n-1,n}(x_{n-1}^m, x_n^m)$, and resample the population when the weights degenerate, optionally with a π_{n-1}-invariant mutation, so that at level 0 the weighted population targets π_0. Particle diversity is measured by the effective sample size of the weights, and the analogous barrier governs the number of levels needed for stability (Syed et al., 2026). SMC never simulates Q' and only evaluates Eq. (2.4). Sequential annealing is preferable when wall-clock time is the binding constraint and parallel annealing when memory is.

3 Methods

3.1 Sequential versus Parallel Control

Our goal is consistent estimation under the target π_0, given the denoising kernels $P_{n-1|n}$ and potentials $G_0, \ldots, G_N$ with $G_N \equiv 1$, where $G_1, \ldots, G_{N-1}$ are supplied by a guidance method. We run annealed Monte Carlo on the ladder $\pi_n(\mathrm{d}x_n) \propto G_n(x_n)\, p_n(\mathrm{d}x_n)$, $n = N, \ldots, 0$, which interpolates from the reference $\pi_N = \eta$ to the target π_0. Both samplers of Sec. 2.3 require backward and forward kernels $Q_{n-1|n}$, $Q'_{n|n-1}$ for which the incremental weight Eq. (2.4) is evaluable. The ratio $\mathrm{d}\pi_{n-1}/\mathrm{d}\pi_n$ contains the marginal ratio $\mathrm{d}p_{n-1}/\mathrm{d}p_n$, which is unavailable. We write it through the time reversal of $P_{n-1|n}$, the kernel $P_{n|n-1}$ defined by $p_{n-1} \otimes P_{n|n-1} = p_n \otimes P_{n-1|n}$, formally

$$P_{n|n-1}(\mathrm{d}x_n \mid x_{n-1}) := \frac{p_n(\mathrm{d}x_n)}{p_{n-1}(\mathrm{d}x_{n-1})}\, P_{n-1|n}(\mathrm{d}x_{n-1} \mid x_n). \tag{3.1}$$

The weight then splits into a backward and a forward factor, $w_{n-1,n} = w_{n-1|n}\, w_{n|n-1}$ with

$$w_{n-1|n} := \frac{G_{n-1}(x_{n-1})}{G_n(x_n)}\, \frac{P_{n-1|n}(x_{n-1} \mid x_n)}{Q_{n-1|n}(x_{n-1} \mid x_n)}, \qquad w_{n|n-1} := \frac{Q'_{n|n-1}(x_n \mid x_{n-1})}{P_{n|n-1}(x_n \mid x_{n-1})}, \tag{3.2}$$

both functions of (x_{n-1}, x_n). We assume throughout that $P_{n-1|n}(\cdot \mid x_n) \ll Q_{n-1|n}(\cdot \mid x_n)$, so that the backward weight $w_{n-1|n}$ exists and is evaluable. The forward weight $w_{n|n-1}$ involves $P_{n|n-1}$,

which depends on p_n and can be neither evaluated nor simulated. We call a correction *exact* if it evaluates $w_{n-1,n}$ without asymptotic bias, and the proposal *unguided* if $Q_{n-1|n} = P_{n-1|n}$, in which case $w_{n-1|n} = G_{n-1}/G_n$ involves no model density, and *guided* otherwise.

Exact control. In practice, we want to preserve the intended target induced by the sampler. This is straightforward for *sequential control*: SMC never simulates $Q'_{n|n-1}$ or evaluates the corresponding forward weight, so we may take $Q'_{n|n-1} = P_{n|n-1}$, giving $w_{n|n-1} = 1$ and $w_{n-1,n} = w_{n-1|n}$. This gives the exact weights used by TDS (Wu et al., 2023a) with a guided proposal and FKD (Singhal et al., 2025) with an unguided one.

However, for *parallel control*, exactness remains challenging. Replica exchange must simulate $Q'_{n|n-1}$, while $P_{n|n-1}$ is generally unavailable for simulation. CREPE (He et al., 2026b) substitutes the model's noising kernel, $Q'_{n|n-1} = \widehat{P}_{n|n-1}$, and sets $w_{n|n-1}$ to one, as in the Radon–Nikodym estimator of He et al. (2026a). This is exact only if the sampler reverses the noising process exactly. Otherwise, the error accumulates over the model steps spanned by an exchange (Prop. B.1). It is also undefined for samplers without a noising kernel. The bias can be reduced by increasing N, at the expense of the computational efficiency of parallel annealing. No exact parallel control method currently exists. We obtain one by lifting the replicas to path space: rather than approximating $P_{n|n-1}$, each replica stores its denoising trajectory, so the forward move need not be simulated.

3.2 Replica Exchange on the Trajectory Ladder

The marginals p_n are intractable, but the law of the denoising path is not: the partial trajectory $\mathbf{x}_n = (x_N, \dots, x_n) \in \mathbb{X}_n := \mathbb{X}^{N-n+1}$ generated down to level n has the product law

$$\mathbf{p}_N = p_N, \qquad \mathbf{p}_{n-1} = \mathbf{p}_n \otimes P_{n-1|n}, \tag{3.3}$$

whose x_n-marginal is p_n. If the replica at level n stores the trajectory $\mathbf{x}_n$ that produced its state, the forward move is a truncation rather than a simulation: under the level law the truncated state is a draw of $P_{n|n-1}$ given the lower state, so $w_{n|n-1} = 1$ with nothing simulated. Bold symbols denote trajectories and objects defined on them, and the subscript of a trajectory is the level of its last state. Let $\mathcal{T}_{n|n-1} : \mathbb{X}_{n-1} \to \mathbb{X}_n$ drop the last state, so that $\mathbf{x}_{n-1} = (\mathcal{T}_{n|n-1}(\mathbf{x}_{n-1}), x_{n-1})$. Level n holds $\mathbf{x}_n \in \mathbb{X}_n$ and the level distributions are

$$\boldsymbol{\pi}_n(\mathrm{d}\mathbf{x}_n) \propto G_n(x_n)\, \mathbf{p}_n(\mathrm{d}\mathbf{x}_n), \tag{3.4}$$

so the replicas jointly target $\boldsymbol{\pi}_N \otimes \cdots \otimes \boldsymbol{\pi}_0$. The x_n-marginal of $\boldsymbol{\pi}_n$ is π_n: it is the ladder of Sec. 3.1 lifted to trajectories, with $\boldsymbol{\pi}_N = \eta$ and the target π_0 as the x_0-marginal of level 0. ENCORE (Alg. 1) is the parallel annealing sampler of Sec. 2.3 on this ladder with the following path exchange.

Path exchange. We mirror Eq. (2.5). Let $\mathbf{x}_n$ and $\mathbf{x}_{n-1} = (\mathbf{y}_n, x_{n-1})$ be the replicas at levels n and $n-1$, where $\mathbf{y}_n := \mathcal{T}_{n|n-1}(\mathbf{x}_{n-1})$ is the lower replica truncated to level n. Propose $y_{n-1} \sim Q_{n-1|n}(\mathrm{d}y_{n-1} \mid x_n)$, set $\mathbf{y}_{n-1} := (\mathbf{x}_n, y_{n-1})$, and exchange $(\mathbf{x}_{n-1}, \mathbf{x}_n) \leftarrow (\mathbf{y}_{n-1}, \mathbf{y}_n)$: level $n-1$ receives the upper replica extended by one proposed state, and level n receives the truncation $\mathbf{y}_n$ of the lower replica, with nothing simulated. On the common space $\mathbb{X}_{n-1}$ the proposal $\mathbf{y}_{n-1}$ has law $\boldsymbol{\pi}_n \otimes Q_{n-1|n}$ and the lower replica has law $\boldsymbol{\pi}_{n-1}$, and by Eq. (3.3) the ratio between $\boldsymbol{\pi}_{n-1}$ and $\boldsymbol{\pi}_n \otimes Q_{n-1|n}$ is $w_{n-1|n}(y_{n-1}, x_n)$ up to a constant: every factor of $\mathbf{p}_n$ cancels, no marginal appears, and the weight depends on the trajectory only through its last state. The exchange is therefore Eq. (2.5) with the forward proposal replaced by memory, and it has the same acceptance probability with the evaluable factor $w_{n-1|n}$ alone,

$$\boldsymbol{\alpha}_{n-1,n}(\mathbf{x}_{n-1}, \mathbf{x}_n; \mathbf{y}_{n-1}, \mathbf{y}_n) = \min\left\{1,\ \frac{w_{n-1|n}(y_{n-1}, x_n)}{w_{n-1|n}(x_{n-1}, y_n)}\right\}. \tag{3.5}$$

This is a reversible-jump exchange (Green, 1995) whose reverse re-extends $\mathbf{y}_n$ by $x_{n-1} \sim Q_{n-1|n}(\mathrm{d}x_{n-1} \mid y_n)$. Eq. (3.5) is the acceptance probability Eq. (2.5) with the exact time reversal Eq. (3.1) as forward proposal, $Q'_{n|n-1} = P_{n|n-1}$. Thm. 3.1 shows that ENCORE is exact for any potentials and proposals, that the Doob potentials and proposals are the optimal choice, and that departures from them cost rejections rather than bias.

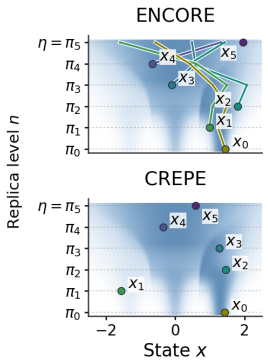

Figure 2: Stored paths (ENCORE) vs. states (CREPE). Blue shading: level densities.

Algorithm 1: ENCORE sampler

Input: reference η; sampler $(P_{n-1|n})$; potentials (G_n); proposals $(Q_{n-1|n})$; initial replicas $\mathbf{x}_n^0 \in \mathbb{X}_n$; iterations M.

for $m = 1, \dots, M$ **do**
- Draw $\mathbf{x}_N^m = x_N^m \sim \eta$ *// refresh reference replica*
- **for all** $n \in \{1, \dots, N\}$ *with* $n \equiv m \pmod 2$, *in parallel* **do**
  - Let $\mathbf{x}_n = \mathbf{x}_n^{m-1}$, $\mathbf{x}_{n-1} = \mathbf{x}_{n-1}^{m-1}$ and $\mathbf{y}_n = \mathcal{T}_{n|n-1}(\mathbf{x}_{n-1})$ *// truncate lower replica*
  - Draw $y_{n-1} \sim Q_{n-1|n}(\mathrm{d}y_{n-1} \mid x_n)$ and set $\mathbf{y}_{n-1} \leftarrow (\mathbf{x}_n, y_{n-1})$ *// extend upper replica*
  - With probability $\boldsymbol{\alpha}_{n-1,n}^m := \boldsymbol{\alpha}_{n-1,n}(\mathbf{x}_{n-1}, \mathbf{x}_n; \mathbf{y}_{n-1}, \mathbf{y}_n)$ of Eq. (3.5), set $(\mathbf{x}_{n-1}^m, \mathbf{x}_n^m) \leftarrow (\mathbf{y}_{n-1}, \mathbf{y}_n)$ *// exchange*
- Set $\mathbf{x}_n^m \leftarrow \mathbf{x}_n^{m-1}$ for every level not exchanged
- Optionally apply a π_n-invariant MCMC move to each x_n^m, in parallel over n *// local move*
- **Record** the last state x_0^m of $\mathbf{x}_0^m$ *// target sample*

State-level replica exchange simulates both $Q_{n-1|n}$ and $Q'_{n|n-1}$ at every iteration, whereas ENCORE simulates only $Q_{n-1|n}$, since moving up is a truncation of the stored trajectory. Each iteration therefore requires only $\lfloor N/2 \rfloor$ simulations from $Q_{n-1|n}$, performed in parallel, one for each active pair (Alg. 1). The trade-off is that we must store the entire trajectory of $N - n + 1$ states for the replica at level n, which requires $O(N^2 d)$ memory.

Theorem 3.1. *(a) The ENCORE transition leaves $\boldsymbol{\pi}_N \otimes \cdots \otimes \boldsymbol{\pi}_0$ invariant and the recorded state has stationary law π_0. The stationary rejection rate for the n-th exchange is*

$$r_{n-1,n} := 1 - \mathbb{E}[\boldsymbol{\alpha}_{n-1,n}] = r\big(\pi_{n-1} \otimes P_{n|n-1},\ \pi_n \otimes Q_{n-1|n}\big). \tag{3.6}$$

(b) If $w_{n-1|n}$ is uniformly bounded above and below for each n, the chain is uniformly ergodic.

(c) The restart rate Eq. (2.6) attains $\tau = 1/2$ if and only if G_n and $Q_{n-1|n}$ satisfy the Doob relations Eq. (2.3) for every n.

See App. C for proofs. Where the potential is constant and $Q_{n-1|n} = P_{n-1|n}$, every exchange is accepted whatever the spacing of the levels, so a coarse ladder is tolerated in the high-noise region.

3.3 Diffusion Models

We now specialise the sampler and proposal kernels to a diffusion and study the limit of increasing parallelism, $N \to \infty$. Let $1 = t_N > \cdots > t_0 = 0$ with steps $\Delta t_n := t_n - t_{n-1}$, and let the potentials discretise a smooth path $t \mapsto G_t$ with $G_1 \equiv 1$ and G_0 the reweighting function, so that $\pi_t \propto G_t p_t$, $G_n = G_{t_n}$ and $\pi_n = \pi_{t_n}$. The proposal $Q_{n-1|n}$ takes one Euler–Maruyama step from X_t to $t' = t - \Delta t$,

$$Y_{t'} = X_t + \big(b_t(X_t) + \mathbf{1}_{\text{guided}}\, \sigma_t \sigma_t^\top \nabla \log G_t(X_t)\big)\, \Delta t + \sigma_t \sqrt{\Delta t}\, \xi, \qquad \xi \sim N(0, I), \tag{3.7}$$

which for $\mathbf{1}_{\text{guided}} = 0$ is the denoiser step $P_{n-1|n}$, the Euler–Maruyama step of the generative SDE $\mathrm{d}X_t = b_t(X_t)\,\mathrm{d}t + \sigma_t\,\mathrm{d}B_t$ run from $t = 1$ to $t = 0$, and for $\mathbf{1}_{\text{guided}} = 1$ the guided denoiser. Composing several model steps per level changes only the remainders below.

Consider the exchange between $t = t_n$ and $t' = t_{n-1}$. Given X_t and ξ, the step Eq. (3.7) determines $Y_{t'}$, and so the weight $w_{n-1|n}(Y_{t'}, X_t)$ is a function $w_{t'|t}(X_t, \xi)$ of the state and the noise,

$$\log w_{t'|t}(X_t, \xi) := \underbrace{\log \frac{G_{t'}(Y_{t'})}{G_t(X_t)}}_{\text{potential}} + \mathbf{1}_{\text{guided}} \underbrace{\log \frac{P_{t'|t}(Y_{t'} \mid X_t)}{Q_{t'|t}(Y_{t'} \mid X_t)}}_{\text{proposal correction}}, \tag{3.8}$$

and the rejection rate Eq. (3.6) is approximately $\frac{1}{2}\mathbb{E}|\ell_n - \ell'_n|$ for two independent copies ℓ_n, ℓ'_n of this log-weight with $X_t \sim \pi_t$. Expanding along the step,

$$\log w_{t'|t}(X_t, \xi) = \begin{cases} \sqrt{\Delta t}\, g_t^\top \xi + O(\Delta t) & \text{unguided}, \\ \Delta t\, \Psi_t(X_t, \xi) + O(\Delta t^{3/2}) & \text{guided}, \end{cases} \tag{3.9}$$

where $g_t := \sigma_t^\top \nabla \log G_t$ and $\Psi_t(X_t, \xi) = -\partial_{t'} \log w_{t'|t}(X_t, \xi)\big|_{t'=t}$ is the pathwise derivative of the guided log-weight with ξ held fixed. Explicitly, Ψ_t is the rate of change of $\log G_t$ along the sampler plus a mean-zero curvature fluctuation,

$$\Psi_t(x, \xi) = \frac{\mathcal{A}_t G_t(x)}{G_t(x)} + \tfrac{1}{2}\big(\xi^\top M_t(x)\xi - \operatorname{tr} M_t(x)\big), \tag{3.10}$$

where $\mathcal{A}_t := -\partial_t + b_t \cdot \nabla + \frac{1}{2}\operatorname{tr}(\sigma_t \sigma_t^\top \nabla^2)$ is the generator of the generative SDE and $M_t := \sigma_t^\top \nabla^2 \log G_t \, \sigma_t$. In continuous time the Doob relations Eq. (2.3) read $\mathcal{A}_t G_t^* = 0$ with $G_0^* = G_0$, solved by $G_t^*(x) = \mathbb{E}[G_0(X_0) \mid X_t = x]$, together with the guided step for G_t^*. For X_t, X_t' i.i.d. from π_t and $\xi, \xi' \sim N(0, I)$ independent, define the local barriers

$$\lambda^{\mathrm{u}}(t) := \tfrac{1}{\sqrt{2\pi}}\, \mathbb{E}\sqrt{\|g_t(X_t)\|^2 + \|g_t(X_t')\|^2}, \qquad \lambda(t) := \tfrac{1}{2}\, \mathbb{E}\big|\Psi_t(X_t, \xi) - \Psi_t(X_t', \xi')\big|. \tag{3.11}$$

Theorem 3.2. *Assume σ_t is invertible and continuous, and b_t and $\log G_t$ are smooth with bounded derivatives. Then λ^{u} and λ are continuous on $[0, 1]$, and if $N \to \infty$ with $\max_n \Delta t_n \to 0$, the following hold uniformly in n.*

(a) If $\mathbf{1}_{\mathrm{guided}} = 0$ in Eq. (3.7), $r_{n-1,n} = \lambda^{\mathrm{u}}(t_n)\sqrt{\Delta t_n} + O(\Delta t_n)$, so $\sum_n r_{n-1,n} \to \infty$.

(b) If $\mathbf{1}_{\mathrm{guided}} = 1$ in Eq. (3.7), $r_{n-1,n} = \int_{t_{n-1}}^{t_n} \lambda(t)\,\mathrm{d}t + o(\Delta t_n)$, so $\sum_n r_{n-1,n} \to \Lambda := \int_0^1 \lambda(t)\,\mathrm{d}t < \infty$, with $\Lambda = 0$ if and only if $G_t \propto G_t^*$ and $\log G_t$ is affine in x for every t.

See App. D for the proof. Unguided, the log-weight is a mean-zero fluctuation of order $\sqrt{\Delta t}$, which a mean absolute difference does not average out. Guided, the proposal correction is the log-ratio of two Gaussians with equal covariance and cancels the $\sqrt{\Delta t}$ term exactly. $\Lambda = 0$ exactly when the guided proposal is the Doob proposal, so Λ measures the deviation of the guided ladder from the Doob h-transform, and by Eq. (2.6) the restart rate tends to zero for the unguided ladder and to $(2 + 2\Lambda)^{-1}$ for the guided one.

Guided and unguided ladders therefore differ qualitatively, and the gap widens with parallelism: the unguided rejection is set by the diffusivity of the sampler, not by the potentials, so additional levels only add rejections, whereas the guided rejection is set by the Doob residual and its total is fixed. Λ also sets the schedule. A guided ladder with $N \approx 2\Lambda$ equi-rejection levels attains stable communication, after which a fresh reference draw reaches the target every $2+2\Lambda$ iterations whatever N, so a longer run rather than a finer ladder increases the sample size. An unguided ladder should have levels about 1.2 apart in log-weight standard deviation, merged wherever $\nabla \log G_t$ is small, and where $G_t \equiv 1$ a single level suffices, the coarse-to-fine ladder of Sec. 4.

3.4 Related Work

Parallel control. CREPE (He et al., 2026b) is the only other parallel control method (Tab. 1). Exactness costs memory, $O(N^2 d)$ against $O(Nd)$, but the gap is smaller in practice: CREPE's reversal error accumulates over the model steps an exchange spans and forces a fine ladder, whereas ENCORE tolerates a coarse one (Sec. 4). ENCORE also returns the law of the full trajectory rather than its marginals, and applies without a noising kernel and to deterministic samplers.

Sequential control. FKD (Singhal et al., 2025) and TDS (Wu et al., 2023a) are the sequential counterparts of unguided and guided ENCORE. Against their $O(Md)$, ENCORE uses less memory once $M > N^2$, i.e. once $M \geq 4\Lambda^2$ at the $N \approx 2\Lambda$ levels of Sec. 3.3. The expansions behind Thm. 3.2 govern their weights as well: the unguided log-weight variance does not vanish under refinement while the guided one does, recovering the annealed-SMC barrier of Syed et al. (2026) (App. D.7). RNE (He et al., 2026a) estimates the marginal ratio by the transition ratio $P/\widehat{P}$ and is not exact, and FKC (Skreta et al., 2025) derives weights in continuous time for marginal-based targets and reduces to TDS on the ladder of Sec. 3.1.

Guidance and search. Guidance (Dhariwal & Nichol, 2021; Ho & Salimans, 2022; Chung et al., 2023; Song et al., 2023; Bansal et al., 2024) constructs $Q_{n-1|n}$ from an estimate of G_n^*, and search methods such as SVDD (Li et al., 2025) and LiDAR (Kim et al., 2026) select among candidate continuations. Neither targets π_0 (Bradley & Nakkiran, 2024; Dandapanthula & Boffi, 2026), and both supply the proposals and potentials that ENCORE corrects.

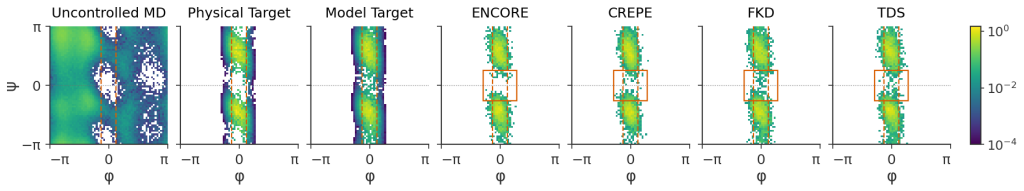

Figure 3: Alanine transition-region sampling at 800 K; each method panel shows 5,000 samples.

4 Experiments

In this section, we evaluate the empirical performance of ENCORE, with experimental details and additional results provided in App. E. We first test robustness to perturbed transition kernels in synthetic settings, then evaluate generalisation to molecular and image-generation tasks.

4.1 Synthetic Reward Tilting with Sampler Perturbations

We evaluate reward tilting on 6-mode 1D and 40-mode 2D Gaussian mixtures with exact scores under several reward settings. To assess robustness to sampler misspecification, we compare nominal transitions with perturbed kernels obtained by altering the transition variance while keeping the reward fixed. Because each perturbation changes the sampler output law p_0, every method is evaluated against the corresponding target $\pi_0 \propto G_0 p_0$. ENCORE remains competitive across all evaluated settings and is notably robust to transition perturbations. In the 1D example in Fig. 4, ENCORE and CREPE both closely match the nominal target; after perturbation, ENCORE remains accurate ($W_2 = 0.015$), whereas CREPE substantially distorts the target mode weights ($W_2 = 0.650$). See App. E.1 for detailed 1D and 2D results.

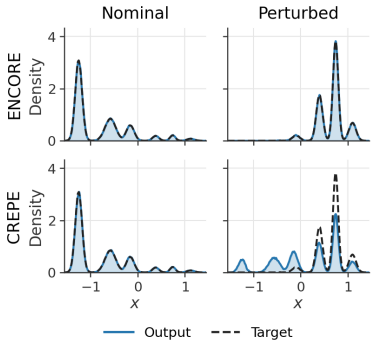

Figure 4: 1D densities and sampler-specific targets.

4.2 Sampling Molecular Transition Regions

In molecular dynamics, transitions between metastable states pass through low-probability reactive regions that are rarely sampled by unbiased dynamics (Du et al., 2026; Omidi et al., 2026). Efficiently targeting these regions is therefore important but challenging. We study alanine dipeptide, whose slow conformational changes are well described by the backbone torsion angles ϕ and ψ. To bias sampling towards a transition region, we apply a harmonic restraint of width 0.2 rad to ϕ in an 800 K sampler while leaving ψ unconstrained. We construct numerical references for the tilted model distribution and the corresponding physical ensemble using importance sampling over model outputs and molecular-dynamics (MD) frames, respectively. In Fig. 3, the red boxes shows residual density connecting the two populated regions for CREPE, FKD, and TDS, whereas ENCORE separates them more clearly. Across three seeds, ENCORE achieves the lowest mean regional TVD against the model reference (0.012), compared with 0.027 for CREPE and 0.037 for both FKD and TDS.

4.3 Boltzmann Tempering of Molecular Ensembles

Following Skreta et al. (2025); He et al. (2026b), we temper 800 K samplers to 300 K for alanine dipeptide (ALDP), 500 K for tetrapeptide (A4), and 600 K for hexapeptide (A6). For ENCORE and FKD, we formulate the temperature change as endpoint Boltzmann reweighting, $G_0(x) = \exp\{-(\beta_{T_t} - \beta_{T_s})U(x)\}$, with potential energy U evaluated using OpenMM (Eastman et al., 2023); CREPE, RNE, and FKC use the model-density tempering constructions. Across all evaluations, ENCORE achieves the lowest mean error in most cases and the second lowest in all others; see Tabs. 2 and 10. The A6 TICA densities in Fig. 5 show broader conformational coverage for ENCORE and CREPE than for FKC-Noise and RNE, illustrating the diversity retained by RE in this setting.

Table 2: A6 Boltzmann-tempering error using 5,000 samples per evaluation (mean ± standard deviation over three subsampling evaluations from one sampler run).

System	Method	Energy TVD↓	Distance TVD↓	Aligned W_2 ↓	Torsion MMD^2↓
Alanine Hexapeptide (A6) 800 → 600 K 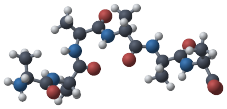	ENCORE	**0.098** ± 0.0017	**0.005** ± 0.0003	1.253 ± 0.0030	0.002 ± 0.0002
	CREPE	0.337 ± 0.0070	0.014 ± 0.0007	1.288 ± 0.0019	0.004 ± 0.0002
	RNE	0.160 ± 0.0042	0.023 ± 0.0008	1.591 ± 0.0029	0.009 ± 0.0006
	FKC-Score	0.169 ± 0.0047	0.012 ± 0.0005	1.596 ± 0.0054	0.011 ± 0.0006
	FKC-Noise	0.135 ± 0.0047	0.010 ± 0.0009	1.584 ± 0.0037	0.008 ± 0.0003
	FKD	0.196 ± 0.0061	0.006 ± 0.0003	**1.193** ± 0.0030	**0.001** ± 0.0000
	Uncontrolled	0.890 ± 0.0017	0.020 ± 0.0013	1.262 ± 0.0020	0.009 ± 0.0002

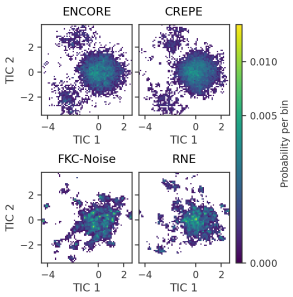

Figure 5: A6 target-fitted TICA densities; 49,000 samples per method.

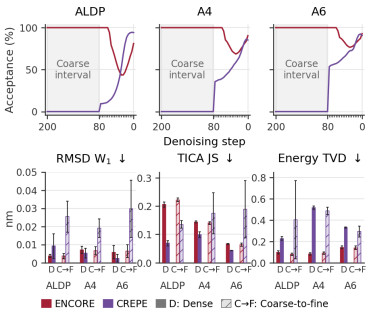

Figure 6: Coarse-to-fine acceptance (top) and sampling errors under both ladders (bottom).



Figure 7: SDXL-Turbo samples for *a photograph of an astronaut riding a horse*.

Ladder spacing and exchange acceptance. We further compare dense and coarse-to-fine PT ladders, where the latter use a large initial jump followed by finer spacing; see Fig. 6. Across the coarse high-noise interval, ENCORE accepts every exchange in our experiments, whereas CREPE rejects every proposal. CREPE's RMSD error increases substantially under the coarse-to-fine ladder, while ENCORE remains robust to the wider spacing. ENCORE's unit acceptance follows from $Q = P$ and $G_n = G_{n-1} \equiv 1$. For all three peptides, ENCORE achieves lower mean energy TVD under both ladders and lower RMSD W_1 on the coarse-to-fine ladder. See App. E.3.1 for details.

4.4 Text-to-Image Reward Tilting

We apply ImageReward tilting (Xu et al., 2023) to four-step SDXL-Turbo (Sauer et al., 2024) and DMD2-UNet (Yin et al., 2024), using 32 prompts and 32 outputs per prompt. ENCORE applies to both without a paired noising kernel. CREPE uses an approximate pairing on SDXL-Turbo and is not evaluated on DMD2, whose native transitions lack the required pairing. At matched chain counts, ENCORE and CREPE achieve similar mean ImageReward on SDXL-Turbo. Compared with FKD, 16-chain ENCORE achieves comparable reward and prompt alignment with higher diversity scores on both samplers; on DMD2, their mean DINO-Vendi scores are 3.276 and 1.473, respectively (Tabs. 12 and 13). Fig. 7 provides a qualitative comparison. Since exact target samples are unavailable, these comparisons assess the reward–diversity trade-off rather than distributional accuracy.

5 Limitations and Future Work

ENCORE achieves exact replica exchange by augmenting each replica with a partial generation trajectory instead of retaining only its current state. This incurs additional memory cost, and its practical efficiency depends on the choice of replica ladder, intermediate potentials, and proposal kernels. Moreover, the controlled target is defined relative to the sampler output distribution p_0, so ENCORE does not correct misspecification of the underlying generative model. Future work could develop more memory-efficient trajectory representations and adaptive designs for the ladder, potentials, and proposals, potentially drawing on ideas from SMC-based diffusion control.

AI USE STATEMENT

In this work, we used generative AI tools for assist in the writing of proofs, design or provide feedback on research methodology or experiments, implement methods, and interpret results. We have not used generative AI tools for help develop theoretical models or conceptual frameworks, formulate mathematical claims, provide critical ingredients for proving mathematical claims, propose or refine hypotheses, and clean and reformat dataset, and generate synthetic data sets, assist with translation, and support qualitative and thematic data analysis are not applicable to this work. Additionally, we used generative AI tools for sourcing/searching for information. We have reviewed all AI-assisted work. We checked all LLM-generated proofs, verified and tested LLM-generated code for correctness. We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

REPRODUCIBILITY STATEMENT

We provide all theoretical results, assumptions, and proofs in the appendix or refer to established results where appropriate. We also report the experimental settings and hyperparameters in App. E. To support reproducibility, we plan to release the code upon acceptance.

REFERENCES

Andrew J. Ballard and Christopher Jarzynski. Replica exchange with nonequilibrium switches. *Proceedings of the National Academy of Sciences*, 106(30):12224–12229, 2009.

Arpit Bansal, Hong-Min Chu, Avi Schwarzschild, Roni Sengupta, Micah Goldblum, Jonas Geiping, and Tom Goldstein. Universal Guidance for Diffusion Models. In *International Conference on Learning Representations*, pp. 51304–51323, 2024.

Arwen Bradley and Preetum Nakkiran. Classifier-Free Guidance is a Predictor-Corrector, 2024. arXiv preprint arXiv:2408.09000 [cs.LG].

Hyungjin Chung, Jeongsol Kim, Michael Thompson Mccann, Marc Louis Klasky, and Jong Chul Ye. Diffusion Posterior Sampling for General Noisy Inverse Problems. In *The Eleventh International Conference on Learning Representations*, 2023.

Sanjit Dandapanthula and Nicholas M. Boffi. Are we really tilting? The mechanics of reward guidance in flow and diffusion models, 2026. arXiv preprint arXiv:2606.02884 [cs.LG].

Pierre Del Moral. *Feynman-Kac Formulae: Genealogical and Interacting Particle Systems with Applications*. Springer Science & Business Media, 2004. ISBN 978-1-4684-9393-1.

Pierre Del Moral, Arnaud Doucet, and Ajay Jasra. Sequential Monte Carlo Samplers. *Journal of the Royal Statistical Society Series B (Statistical Methodology)*, 68(3):411–436, 2006. ISSN 1369-7412.

Prafulla Dhariwal and Alexander Nichol. Diffusion Models Beat GANs on Image Synthesis. In *Advances in Neural Information Processing Systems*, volume 34, pp. 8780–8794, 2021.

Yuanqi Du, Jiajun He, Francisco Vargas, Yuanqing Wang, Carla Gomes, José Miguel Hernández-Lobato, and Eric Vanden-Eijnden. FEAT: Free energy Estimators with Adaptive Transport. In *Advances in Neural Information Processing Systems*, volume 38, pp. 142130–142167, 2025.

Yuanqi Du, Jiajun He, Dinghuai Zhang, Eric Vanden-Eijnden, and Carles Domingo-Enrich. Rare Event Analysis via Stochastic Optimal Control, 2026. arXiv preprint arXiv:2604.13213 [stat.ML].

Peter Eastman, Raimondas Galvelis, Raúl P. Peláez, Charlles R. A. Abreu, Stephen E. Farr, Emilio Gallicchio, Anton Gorenko, Michael M. Henry, Frank Hu, Jing Huang, Andreas Krämer, Julien Michel, Joshua A. Mitchell, Vijay S. Pande, João PGLM Rodrigues, Jaime Rodriguez-Guerra, Andrew C. Simmonett, Sukrit Singh, Jason Swails, Philip Turner, Yuanqing Wang, Ivy Zhang, John D. Chodera, Gianni De Fabritiis, and Thomas E. Markland. OpenMM 8: Molecular Dynamics Simulation with Machine Learning Potentials. *Journal of Physical Chemistry B*, 128(1):109–116, 2023. ISSN 1520-6106.

Dan Friedman and Adji Bousso Dieng. The Vendi Score: A Diversity Evaluation Metric for Machine Learning. *Transactions on Machine Learning Research*, 2023. ISSN 2835-8856.

Charles J. Geyer. Markov Chain Monte Carlo Maximum Likelihood. In *Proceedings of the 23rd Symposium on the Interface*, pp. 156–163, 1991.

Peter J. Green. Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika*, 82(4):711–732, 1995. ISSN 0006-3444.

Arthur Gretton, Karsten M. Borgwardt, Malte J. Rasch, Bernhard Schölkopf, and Alexander Smola. A Kernel Two-Sample Test. *Journal of Machine Learning Research*, 13(25):723–773, 2012. ISSN 1533-7928.

Jiajun He, José Miguel Hernández Lobato, Yuanqi Du, and Francisco Vargas. RNE: Plug-and-play diffusion inference-time control and energy-based training. In *International Conference on Learning Representations*, pp. 75732–75786, 2026a.

Jiajun He, Paul Jeha, Peter Potaptchik, Leo Zhang, José Miguel Hernández Lobato, Yuanqi Du, Saifuddin Syed, and Francisco Vargas. CREPE: Controlling diffusion with REPlica Exchange. In *International Conference on Learning Representations*, pp. 99464–99494, 2026b.

Jonathan Ho and Tim Salimans. Classifier-Free Diffusion Guidance, 2022. arXiv preprint arXiv:2207.12598 [cs].

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising Diffusion Probabilistic Models. In *Advances in Neural Information Processing Systems*, volume 33, pp. 6840–6851, 2020.

Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David Fleet. Video Diffusion Models. In *Advances in Neural Information Processing Systems*, volume 35, pp. 8633–8646, 2022.

Emiel Hoogeboom, Víctor Garcia Satorras, Clément Vignac, and Max Welling. Equivariant Diffusion for Molecule Generation in 3D. In *Proceedings of the 39th International Conference on Machine Learning*, pp. 8867–8887, 2022.

Koji Hukushima and Koji Nemoto. Exchange Monte Carlo Method and Application to Spin Glass Simulations. *Journal of the Physical Society of Japan*, 65(6):1604–1608, 1996. ISSN 0031-9015.

W. Kabsch. A solution for the best rotation to relate two sets of vectors. *Acta Crystallographica Section A: Crystal Physics, Diffraction, Theoretical and General Crystallography*, 32(5):922–923, 1976. ISSN 0567-7394.

Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the Design Space of Diffusion-Based Generative Models. In *Advances in Neural Information Processing Systems*, volume 35, pp. 26565–26577, 2022.

Helmut G Katzgraber, Simon Trebst, David A Huse, and Matthias Troyer. Feedback-optimized parallel tempering Monte Carlo. *Journal of Statistical Mechanics: Theory and Experiment*, 2006 (03):P03018, 2006. ISSN 1742-5468.

Yeongmin Kim, Donghyeok Shin, Byeonghu Na, Minsang Park, Richard Lee Kim, and Il-chul Moon. Lookahead Sample Reward Guidance for Test-Time Scaling of Diffusion Models. In *Forty-Third International Conference on Machine Learning*, 2026.

Xiner Li, Yulai Zhao, Chenyu Wang, Gabriele Scalia, Gokcen Eraslan, Surag Nair, Tommaso Biancalani, Shuiwang Ji, Aviv Regev, Sergey Levine, and Masatoshi Uehara. Derivative-Free Guidance in Continuous and Discrete Diffusion Models with Soft Value-based Decoding. In *Advances in Neural Information Processing Systems*, volume 38, pp. 95507–95545, 2025.

Sean Meyn and Richard L. Tweedie. *Markov Chains and Stochastic Stability*. Cambridge Mathematical Library. Cambridge University Press, Cambridge, 2 edition, 2009. ISBN 978-0-521-73182-9. doi: 10.1017/CBO9780511626630.

Laurence Illing Midgley, Vincent Stimper, Gregor N. C. Simm, Bernhard Schölkopf, and José Miguel Hernández-Lobato. Flow Annealed Importance Sampling Bootstrap. In *The Eleventh International Conference on Learning Representations*, 2023.

Alireza Omidi, Jiajun He, Jörg Gsponer, and Saifuddin Syed. METALICA: METAdynamics and repLICA exchange for enhanced diffusion sampling, 2026. arXiv preprint arXiv:2609.17823 [stat.ML].

Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy V. Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, Mido Assran, Nicolas Ballas, Wojciech Galuba, Russell Howes, Po-Yao Huang, Shang-Wen Li, Ishan Misra, Michael Rabbat, Vasu Sharma, Gabriel Synnaeve, Hu Xu, Herve Jegou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. DINOv2: Learning Robust Visual Features without Supervision. *Transactions on Machine Learning Research*, 2024. ISSN 2835-8856.

Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning Transferable Visual Models From Natural Language Supervision. In *Proceedings of the 38th International Conference on Machine Learning*, pp. 8748–8763, 2021.

Axel Sauer, Dominik Lorenz, Andreas Blattmann, and Robin Rombach. Adversarial Diffusion Distillation. In Aleš Leonardis, Elisa Ricci, Stefan Roth, Olga Russakovsky, Torsten Sattler, and Gül Varol (eds.), *Computer Vision – ECCV 2024*, pp. 87–103, Cham, 2024. ISBN 978-3-031-73016-0.

Raghav Singhal, Zachary Horvitz, Ryan Teehan, Mengye Ren, Zhou Yu, Kathleen Mckeown, and Rajesh Ranganath. A General Framework for Inference-time Scaling and Steering of Diffusion Models. In *Proceedings of the 42nd International Conference on Machine Learning*, pp. 55810–55827, 2025.

Marta Skreta, Tara Akhound-Sadegh, Viktor Ohanesian, Roberto Bondesan, Alan Aspuru-Guzik, Arnaud Doucet, Rob Brekelmans, Alexander Tong, and Kirill Neklyudov. Feynman-Kac Correctors in Diffusion: Annealing, Guidance, and Product of Experts. In *Proceedings of the 42nd International Conference on Machine Learning*, pp. 55906–55949, 2025.

Jiaming Song, Qinsheng Zhang, Hongxu Yin, Morteza Mardani, Ming-Yu Liu, Jan Kautz, Yongxin Chen, and Arash Vahdat. Loss-Guided Diffusion Models for Plug-and-Play Controllable Generation. In *Proceedings of the 40th International Conference on Machine Learning*, pp. 32483–32498, 2023.

Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-Based Generative Modeling through Stochastic Differential Equations. In *International Conference on Learning Representations*, 2021.

Robert H. Swendsen and Jian-Sheng Wang. Replica Monte Carlo Simulation of Spin-Glasses. *Physical Review Letters*, 57(21):2607–2609, 1986.

Saifuddin Syed, Alexandre Bouchard-Côté, George Deligiannidis, and Arnaud Doucet. Non-Reversible Parallel Tempering: A Scalable Highly Parallel MCMC Scheme. *Journal of the Royal Statistical Society Series B (Statistical Methodology)*, 84(2):321–350, 2022. ISSN 1369-7412.

Saifuddin Syed, Alexandre Bouchard-Côté, Kevin Chern, and Arnaud Doucet. Optimized annealed sequential Monte Carlo samplers. *Journal of the Royal Statistical Society Series B (Statistical Methodology)*, pp. qkag082, 2026. ISSN 1369-7412.

Masatoshi Uehara, Xingyu Su, Yulai Zhao, Xiner Li, Aviv Regev, Shuiwang Ji, Sergey Levine, and Tommaso Biancalani. Reward-Guided Iterative Refinement in Diffusion Models at Test-Time with Applications to Protein and DNA Design. In *Proceedings of the 42nd International Conference on Machine Learning*, pp. 60515–60529, 2025.

Luhuan Wu, Brian Trippe, Christian Naesseth, David Blei, and John Cunningham. Practical and Asymptotically Exact Conditional Sampling in Diffusion Models. In *Advances in Neural Information Processing Systems*, volume 36, pp. 31372–31403, 2023a.

Xiaoshi Wu, Yiming Hao, Keqiang Sun, Yixiong Chen, Feng Zhu, Rui Zhao, and Hongsheng Li. Human Preference Score v2: A Solid Benchmark for Evaluating Human Preferences of Text-to-Image Synthesis, 2023b. arXiv preprint arXiv:2306.09341 [cs.CV].

Wangang Xie, Paul O. Lewis, Yu Fan, Lynn Kuo, and Ming-Hui Chen. Improving Marginal Likelihood Estimation for Bayesian Phylogenetic Model Selection. *Systematic Biology*, 60(2): 150–160, 2011. ISSN 1063-5157.

Yu Xie, Ludwig Winkler, Lixin Sun, Sarah Lewis, Adam E. Foster, José Jiménez Luna, Tim Hempel, Michael Gastegger, Yaoyi Chen, Iryna Zaporozhets, Cecilia Clementi, Christopher M. Bishop, and Frank Noé. Enhanced Diffusion Sampling: Efficient Rare Event Sampling and Free Energy Calculation with Diffusion Models, 2026. arXiv preprint arXiv:2602.16634 [stat.ML].

Jiazheng Xu, Xiao Liu, Yuchen Wu, Yuxuan Tong, Qinkai Li, Ming Ding, Jie Tang, and Yuxiao Dong. ImageReward: Learning and Evaluating Human Preferences for Text-to-Image Generation. In *Advances in Neural Information Processing Systems*, volume 36, pp. 15903–15935, 2023.

Tianwei Yin, Michaël Gharbi, Taesung Park, Richard Zhang, Eli Shechtman, Frédo Durand, and William T. Freeman. Improved Distribution Matching Distillation for Fast Image Synthesis. In *Advances in Neural Information Processing Systems*, volume 37, pp. 47455–47487, 2024.

Leo Zhang, Peter Potaptchik, Jiajun He, Yuanqi Du, Arnaud Doucet, Francisco Vargas, Hai-Dang Dau, and Saifuddin Syed. Accelerated Parallel Tempering via Neural Transports. In *International Conference on Learning Representations*, pp. 90874–90907, 2026.

Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The Unreasonable Effectiveness of Deep Features as a Perceptual Metric. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 586–595, 2018.

A Background

This appendix expands Sec. 2: the diffusion sampler and its time reversal, the Doob h-transform in discrete and continuous time, and the two annealed Monte Carlo samplers of Sec. 2.3 in full. Results quoted from the literature are stated without proof.

A.1 Diffusion Samplers

The sampler of Sec. 2.1 draws $x_N \sim \eta$ and then, for $n = N, \dots, 1$, $x_{n-1} \sim P_{n-1|n}(\mathrm{d}x_{n-1} \mid x_n)$. Its marginals and its path law are

$$p_N = \eta, \qquad p_{n-1}(\mathrm{d}x_{n-1}) = \int p_n(\mathrm{d}x_n)\, P_{n-1|n}(\mathrm{d}x_{n-1} \mid x_n), \tag{A.1}$$

$$\mathbf{p}_N = \eta, \qquad \mathbf{p}_{n-1} = \mathbf{p}_n \otimes P_{n-1|n}, \tag{A.2}$$

the second line being Eq. (3.3), where $\mathbf{x}_n = (x_N, \dots, x_n)$ is the partial trajectory generated down to level n. The kernels can be simulated, and for a stochastic sampler they have densities. The marginals p_n with $n < N$ are integrals over the learned transitions and are not available in closed form. This is the obstacle that every correction in this paper works around.

Time reversal. The kernel $P_{n|n-1}$ of Eq. (3.1) is the conditional law of x_n given x_{n-1} under the sampler, defined by the identity of laws of the pair (x_{n-1}, x_n),

$$p_{n-1} \otimes P_{n|n-1} = p_n \otimes P_{n-1|n}, \quad \text{i.e.}$$
$$p_{n-1}(\mathrm{d}x_{n-1})\, P_{n|n-1}(\mathrm{d}x_n \mid x_{n-1}) = p_n(\mathrm{d}x_n)\, P_{n-1|n}(\mathrm{d}x_{n-1} \mid x_n). \tag{A.3}$$

It exists whenever the joint law on the right admits a regular conditional distribution, which is the case on the Euclidean spaces used here. When densities exist,

$$P_{n|n-1}(x_n \mid x_{n-1}) = \frac{p_n(x_n)}{p_{n-1}(x_{n-1})}\, P_{n-1|n}(x_{n-1} \mid x_n), \tag{A.4}$$

which involves the ratio of two marginals. It can therefore be neither evaluated nor simulated.

The noising kernel. A diffusion model is trained on a noising process, and the transition $\widehat{P}_{n|n-1}$ of that process from t_{n-1} to t_n is tractable, typically Gaussian. If the learned sampler were the exact time reversal of the noising process at the chosen discretisation, then $\widehat{P}_{n|n-1}$ would equal $P_{n|n-1}$. Model error and discretisation error make the two differ, and the local reversal error

$$\delta_n(x_n, x_{n-1}) := \log \frac{P_{n|n-1}(x_n \mid x_{n-1})}{\widehat{P}_{n|n-1}(x_n \mid x_{n-1})} \tag{A.5}$$

is what App. B quantifies. Distilled few-step samplers are not trained as reversals of a noising process and have no $\widehat{P}_{n|n-1}$ at all. A deterministic sampler has kernels without densities. The unguided ENCORE weight $w_{n-1|n} = G_{n-1}/G_n$ involves no kernel density and applies to it (App. C.1).

A.2 The Doob h-Transform

Discrete time. For the target $\pi_0 \propto G_0\, p_0$ of Eq. (2.2), the Doob potentials and proposals of Eq. (2.3) are

$$G_0^* := G_0, \qquad G_n^* := P_{n-1|n} G_{n-1}^*,$$
$$Q_{n-1|n}^*(\mathrm{d}x_{n-1} \mid x_n) := \frac{G_{n-1}^*(x_{n-1})}{G_n^*(x_n)}\, P_{n-1|n}(\mathrm{d}x_{n-1} \mid x_n). \tag{A.6}$$

Unrolling the recursion gives

$$G_n^*(x_n) = \int P_{n-1|n}(\mathrm{d}x_{n-1} \mid x_n) \cdots P_{0|1}(\mathrm{d}x_0 \mid x_1)\, G_0(x_0) = \mathbb{E}\big[G_0(X_0) \mid X_n = x_n\big] \tag{A.7}$$

under the sampler, so $G_n^*(X_n)$ is a martingale in the generation direction. $Q^*_{n-1|n}$ is a probability kernel because $\int G^*_{n-1}(x_{n-1})P_{n-1|n}(\mathrm{d}x_{n-1} \mid x_n) = G^*_n(x_n)$. The chain that starts from $\eta^* \propto G^*_N\eta$ and follows Q^* has path law

$$\eta^*(\mathrm{d}x_N)\prod_{n=N}^{1} Q^*_{n-1|n}(\mathrm{d}x_{n-1} \mid x_n) = \frac{G^*_N(x_N)\,\eta(\mathrm{d}x_N)}{\mathbb{E}_\eta[G^*_N]}\prod_{n=N}^{1}\frac{G^*_{n-1}(x_{n-1})}{G^*_n(x_n)}\,P_{n-1|n}(\mathrm{d}x_{n-1} \mid x_n)$$
$$= \frac{G^*_0(x_0)}{\mathbb{E}_\eta[G^*_N]}\,\eta(\mathrm{d}x_N)\prod_{n=N}^{1} P_{n-1|n}(\mathrm{d}x_{n-1} \mid x_n), \tag{A.8}$$

since the potentials telescope, and $\mathbb{E}_\eta[G^*_N] = \mathbb{E}_{p_0}[G_0]$ by the martingale property. The right side is the sampler's path law reweighted by $G_0(x_0)$, so its x_0-marginal is $G_0\,p_0/\mathbb{E}_{p_0}[G_0] = \pi_0$ (Del Moral, 2004). This is the exact guidance of Sec. 2.2. It is intractable because each G^*_n is a conditional expectation over the remaining generation.

The Doob relations on the ladder. Thm. 3.1(c) refers to the two relations of Eq. (2.3) for a general ladder $\pi_n \propto G_n p_n$ with proposals $Q_{n-1|n}$:

$$\begin{aligned} &\text{potential relation:} && G_n \propto P_{n-1|n}G_{n-1}, \\ &\text{proposal relation:} && Q_{n-1|n}(\mathrm{d}x_{n-1} \mid x_n) \propto G_{n-1}(x_{n-1})\,P_{n-1|n}(\mathrm{d}x_{n-1} \mid x_n). \end{aligned} \tag{A.9}$$

Iterating the potential relation from G_0 gives $G_n \propto G^*_n$ for every n, the constants compounding along the way and cancelling in π_n. The proposal relation then reads $Q_{n-1|n} = Q^*_{n-1|n}$. At the top level the potential relation reads $1 = G_N \propto G^*_N$, which holds only if G^*_N is constant, i.e. if x_N is independent of x_0 under the sampler. For a diffusion this holds up to the small dependence of x_0 on the initial noise, so $\eta^* \approx \eta$ and the top exchange carries a correspondingly small residual rejection.

Continuous time. Let the sampler discretise the generative SDE

$$\mathrm{d}X_t = b_t(X_t)\,\mathrm{d}t + \sigma_t\,\mathrm{d}B_t, \qquad \text{run from } t = 1 \text{ to } t = 0, \tag{A.10}$$

whose generator, with the sign of ∂_t reflecting that t decreases along generation, is

$$\mathcal{A}_t := -\partial_t + b_t \cdot \nabla + \tfrac{1}{2}\,\mathrm{tr}\left(\sigma_t\sigma_t^\top \nabla^2\right). \tag{A.11}$$

The continuous-time Doob potential

$$G^*_t(x) := \mathbb{E}\big[G_0(X_0) \mid X_t = x\big] \tag{A.12}$$

solves the backward Kolmogorov equation $\mathcal{A}_t G^*_t = 0$ on $(0, 1]$ with terminal condition $G^*_0 = G_0$, and is its unique solution with bounded logarithm under the assumptions of App. D (Del Moral, 2004, Ch. 1). The h-transformed process

$$\mathrm{d}X_t = \big(b_t(X_t) + \sigma_t\sigma_t^\top \nabla \log G^*_t(X_t)\big)\,\mathrm{d}t + \sigma_t\,\mathrm{d}B_t, \qquad X_1 \sim \eta^* \propto G^*_1\eta, \tag{A.13}$$

has terminal law π_0: this is the continuous-time proposal relation, and the guided step of Eq. (3.7) is its Euler–Maruyama discretisation with a general G_t in place of G^*_t.

Two facts about this discretisation are used in App. D. The first is that for a potential whose logarithm is affine in x, the guided step is the Doob proposal of that potential exactly. Indeed, the sampler step is Gaussian, $P_{t'|t}(\cdot \mid x) = N(\mu(x), \Sigma)$ with $\mu(x) = x + b_t(x)\Delta t$ and $\Sigma = \sigma_t\sigma_t^\top \Delta t$, and if $\log G_{t'}(y) = c^\top y + d$ then completing the square gives

$$G_{t'}(y)\,P_{t'|t}(\mathrm{d}y \mid x) \propto \exp\Big(c^\top y - \tfrac{1}{2}(y - \mu(x))^\top \Sigma^{-1}(y - \mu(x))\Big)\,\mathrm{d}y$$
$$\propto \exp\Big(-\tfrac{1}{2}\big(y - \mu(x) - \Sigma c\big)^\top \Sigma^{-1}\big(y - \mu(x) - \Sigma c\big)\Big)\,\mathrm{d}y, \tag{A.14}$$

so the twisted kernel is $N(\mu(x) + \Sigma c, \Sigma)$: the sampler step with its mean shifted by $\Sigma\nabla \log G_{t'} = \sigma_t\sigma_t^\top \nabla \log G_{t'}\,\Delta t$, which is the guided step. The second fact is that the rate of change of $\log G_t$ along the sampler is $\mathcal{A}_t G_t / G_t$, and that it is constant in x if and only if $G_t \propto G^*_t$ (Lem. D.8). Equivalently, G_t is an eigenfunction of $\mathcal{A}_t$, $\mathcal{A}_t G_t = c(t)G_t$. The Doob potential is the eigenfunction with eigenvalue zero, the space-time harmonic function of the diffusion.

A.3 Parallel Annealing: Non-Equilibrium Replica Exchange

Replicas. A *replica* is a state of $\mathbb{X}$ carried by a level of the ladder. The sampler holds one replica per level, $x_0, \dots, x_N$, and targets $\pi_N \otimes \cdots \otimes \pi_0$ by a Markov chain on $\mathbb{X}^{N+1}$. A replica keeps its identity when it is exchanged between levels, so one can follow it up and down the ladder over the iterations. A fresh replica enters at the reference level N when that level is redrawn from π_N, and a replica leaves the ladder when the level it occupies is redrawn.

The sampler. Non-equilibrium replica exchange (Ballard & Jarzynski, 2009; Zhang et al., 2026) is Alg. 2. It proposes, for a pair of adjacent replicas (x_{n-1}, x_n), a downward transition $y_{n-1} \sim Q_{n-1|n}(\cdot \mid x_n)$ from the upper replica and an upward transition $y_n \sim Q'_{n|n-1}(\cdot \mid x_{n-1})$ from the lower one, and replaces the pair by (y_{n-1}, y_n) with the probability Eq. (2.5). Parallel tempering (Swendsen & Wang, 1986; Geyer, 1991; Hukushima & Nemoto, 1996) is the special case $Q_{n-1|n} = Q'_{n|n-1} = \mathrm{Id}$, for which the incremental weight is $w_{n-1,n}(x_{n-1}, x_n) = \pi_{n-1}(x_{n-1})/\pi_n(x_n)$ and Eq. (2.5) reduces to the swap acceptance

$$\alpha_{n-1,n} = \min\left\{1,\ \frac{\pi_{n-1}(x_n)\,\pi_n(x_{n-1})}{\pi_{n-1}(x_{n-1})\,\pi_n(x_n)}\right\}. \tag{A.15}$$

The even–odd schedule attempts the pairs $(n-1, n)$ with $n \equiv m \pmod 2$ at iteration m. The pairs attempted in one iteration are disjoint, so the exchanges run in parallel, and consecutive iterations attempt complementary sets of pairs, which is what makes the communication non-reversible and lets a replica move several levels in one direction without turning back (Syed et al., 2022).

Algorithm 2: Non-equilibrium replica exchange

Input: iterations M; ladder $(\pi_n)_{n=0}^N$; kernels $(Q_{n-1|n})$, $(Q'_{n|n-1})$.

Initialise one replica x_n^0 per level $n = 0, \dots, N$
for $m = 1, \dots, M$ **do**
  Draw $x_N^m \sim \pi_N$ *// refresh reference replica*
  foreach $n \in \{1, \dots, N\}$ *with* $n \equiv m \pmod 2$, *in parallel* **do**
    Draw $y_{n-1}^m \sim Q_{n-1|n}(\mathrm{d}y_{n-1} \mid x_n^{m-1})$ and $y_n^m \sim Q'_{n|n-1}(\mathrm{d}y_n \mid x_{n-1}^{m-1})$
    With probability $\alpha_{n-1,n}^m := \alpha_{n-1,n}(x_{n-1}^{m-1}, x_n^{m-1}; y_{n-1}^m, y_n^m)$ of Eq. (2.5), set $(x_{n-1}^m, x_n^m) \leftarrow (y_{n-1}^m, y_n^m)$ *// exchange*
  Set $x_n^m \leftarrow x_n^{m-1}$ for every level not exchanged
  Optionally apply a π_n-invariant MCMC move to each x_n^m, in parallel over n
  Record x_0^m

Restarts and round trips. Following a replica through the iterations, a *restart* occurs when a replica that entered at the reference level N reaches level 0. The recorded sample is then a fresh draw from the reference, transported to the target through accepted exchanges, and it is independent of the samples recorded before the previous restart. A *round trip* is a restart followed by the return of the same replica to level N. The *restart rate* τ is the long-run number of restarts per iteration (Katzgraber et al., 2006; Syed et al., 2022). It is the natural measure of particle diversity for parallel annealing: between restarts the recorded samples are correlated through the replica that produced them, and the number of effectively independent samples after M iterations is of order τM.

Estimators. In the notation of Alg. 2, x_0^m is the level-0 state recorded at iteration m and (y_{n-1}^m, x_n^{m-1}) the downward pair proposed at iteration m at level n. The recorded states estimate expectations under the target by the Monte Carlo average, and the incremental weights of the downward proposals estimate the ratio of normalising constants $Z := \mathcal{Z}_0/\mathcal{Z}_N$ level by level:

$$\widehat{\pi}_0^{\mathrm{RE}}[f] := \frac{1}{M}\sum_{m=1}^{M} f(x_0^m), \qquad \widehat{Z}^{\mathrm{RE}} := \prod_{n=1}^{N}\Big(\frac{1}{M}\sum_{m=1}^{M} w_{n-1,n}(y_{n-1}^m, x_n^{m-1})\Big). \tag{A.16}$$

The second is the stepping-stone estimator (Xie et al., 2011): at stationarity the downward pair has law $\pi_n \otimes Q_{n-1|n}$, and the incremental weight Eq. (2.4) has expectation $\mathcal{Z}_{n-1}/\mathcal{Z}_n$ under that law, so each factor is consistent for one ratio of adjacent normalising constants and the product is consistent for Z as $M \to \infty$.

The rejection rate. At stationarity the pair (x_{n-1}^{m-1}, y_n^m) proposed upwards has law $\pi_{n-1} \otimes Q'_{n|n-1}$ and the pair (y_{n-1}^m, x_n^{m-1}) proposed downwards has law $\pi_n \otimes Q_{n-1|n}$, and Zhang et al. (2026, Theorem 1) show that

$$r_{n-1,n} = r\big(\pi_{n-1} \otimes Q'_{n|n-1},\ \pi_n \otimes Q_{n-1|n}\big), \qquad r(p,q) := \mathrm{TV}(p \otimes q,\ q \otimes p), \tag{A.17}$$

which is Eq. (2.7). Here $r(p,q)$ is the rejection rate of a swap between two independent replicas with laws p and q on a common space. It is symmetric in (p,q), vanishes if and only if $p = q$, and satisfies $\mathrm{TV}(p,q) \le r(p,q) \le 2\ \mathrm{TV}(p,q)$. Lem. C.3 below records these facts, since they are used in the proof of Thm. 3.1.

Schedule and number of levels. How τ depends on the schedule and on N, the global communication barrier Λ that the summed rejection rates converge to as the schedule is refined, the equi-rejection schedule, and the recommendation $N \approx 2\Lambda$ quoted in Sec. 2.3 are developed in Syed et al. (2022). App. D computes Λ for the ladder of Sec. 3.1 on a diffusion.

A.4 Sequential Annealing: Sequential Monte Carlo Samplers

Sequential Monte Carlo samplers (Del Moral et al., 2006) are Alg. 3. Particle m at level n is a state x_n^m with weight w_n^m, and the weighted population approximates π_n in the sense that

$$\frac{\sum_{m=1}^M w_n^m\, f(x_n^m)}{\sum_{m=1}^M w_n^m} \xrightarrow[M\to\infty]{} \pi_n(f) \qquad \text{for bounded measurable } f. \tag{A.18}$$

Without resampling, w_0^m is the product of the incremental weights Eq. (2.4) along the particle's path, and $\frac{1}{M}\sum_m w_0^m$ is an unbiased estimator of the ratio of normalising constants of π_0 and π_N. Resampling replaces the population by a multinomial draw from the normalised weights when they degenerate, and a π_{n-1}-invariant mutation may follow to restore diversity. Forward proposals appear only inside Eq. (2.4), which is evaluated and never simulated.

Algorithm 3: Sequential Monte Carlo sampler

Input: number of particles M; ladder $(\pi_n)_{n=0}^N$; kernels $(Q_{n-1|n})$, $(Q'_{n|n-1})$; resampling rule.

Draw $x_N^m \sim \pi_N$ and set $w_N^m \leftarrow 1$ for $m = 1, \dots, M$
for $n = N, \dots, 1$ **do**
  foreach $m = 1, \dots, M$ *in parallel* **do**
    Draw $x_{n-1}^m \sim Q_{n-1|n}(\mathrm{d}x_{n-1} \mid x_n^m)$ *// propagate*
    $w_{n-1}^m \leftarrow w_n^m\, w_{n-1,n}(x_{n-1}^m, x_n^m)$ *// reweight*
  if *the resampling rule fires* **then**
    Resample $x_{n-1}^{1:M}$ from the normalised weights and set $w_{n-1}^{1:M} \leftarrow 1$, optionally followed by a π_{n-1}-invariant mutation
return the weighted population $(x_0^{1:M}, w_0^{1:M})$, which targets π_0

Estimators. The weighted population at level 0 estimates expectations under the target, and the average weights accumulated between resampling events estimate $Z = \mathcal{Z}_0/\mathcal{Z}_N$:

$$\widehat{\pi}_0^{\mathrm{SMC}}[f] := \frac{\sum_{m=1}^M w_0^m\, f(x_0^m)}{\sum_{m=1}^M w_0^m}, \qquad \widehat{Z}^{\mathrm{SMC}} := \prod_{k=1}^R \Big(\frac{1}{M}\sum_{m=1}^M w_{(k)}^m\Big), \tag{A.19}$$

where R is the number of resampling events and $w_{(k)}^m$ is the weight of particle m accumulated since the previous resampling event, so that with resampling at every level $\widehat{Z}^{\mathrm{SMC}} = \prod_{n=1}^N \big(\frac{1}{M}\sum_m w_{n-1,n}(x_{n-1}^m, x_n^m)\big)$. The expectation estimator is consistent, and $\widehat{Z}^{\mathrm{SMC}}$ is unbiased for Z for any M and any adaptive resampling rule, and consistent as $M \to \infty$ (Del Moral, 2004; Syed et al., 2026).

Relative variance and the discrepancy. Since $\widehat{Z}^{\mathrm{SMC}}$ is unbiased, its relative variance $\mathrm{Var}[\widehat{Z}^{\mathrm{SMC}}/Z]$ is the natural measure of particle diversity for sequential annealing, and it is what an

annealing schedule should minimise (Syed et al., 2026). Under the performance model of that paper (independent incremental weights across levels and particles, and efficient local exploration in the sense that the weight at each level is distributed as under $\pi_n \otimes Q_{n-1|n}$), the relative variance depends on the schedule only through the incremental discrepancies

$$\mathcal{D}_{n-1,n} := \log\Big(1 + \chi^2\big(\pi_{n-1} \otimes Q'_{n|n-1} \,\big\|\, \pi_n \otimes Q_{n-1|n}\big)\Big) = \log \frac{\mathbb{E}[w_{n-1,n}^2]}{\mathbb{E}[w_{n-1,n}]^2}, \tag{A.20}$$

the Rényi 2-divergence between the backward target and the forward proposal at level n, the expectations under $\pi_n \otimes Q_{n-1|n}$. Note that these are the same two joint laws whose swap rejection rate is Eq. (2.7). With R resampling events accumulating discrepancies $\mathcal{D}_{(1)}, \dots, \mathcal{D}_{(R)}$ between them (Syed et al., 2026, Theorem 1),

$$\mathrm{Var}\left[\frac{\widehat{Z}^{\mathrm{SMC}}}{Z}\right] = \prod_{k=1}^{R}\left(1 + \frac{e^{\mathcal{D}_{(k)}} - 1}{M}\right) - 1, \qquad \mathcal{D}_{0,N} := \sum_{n=1}^{N} \mathcal{D}_{n-1,n} = \sum_{k=1}^{R} \mathcal{D}_{(k)}. \tag{A.21}$$

For a fixed number of resampling events the right side is an increasing function of the total discrepancy $\mathcal{D}_{0,N}$, minimised over its distribution when the $\mathcal{D}_{(k)}$ are equal, and for $\mathcal{D}_{0,N}$ small it is $\approx \mathcal{D}_{0,N}/M$. Optimising the annealing schedule for sequential annealing therefore means minimising $\mathcal{D}_{0,N}$, the sum of the Rényi 2-divergences between adjacent levels. In the small-weight regime $\mathcal{D}_{n-1,n} = \mathrm{Var}(\ell_n) + O(\mathbb{E}|\ell_n|^3)$ with $\ell_n := \log w_{n-1,n}$, so $\mathcal{D}_{0,N}$ is the summed log-weight variance to leading order. App. D.7 evaluates it for the sequential control weights of FKD and TDS.

Computational duality. Both samplers fill the same array (x_n^m) of Sec. 2.3, sequential annealing one level at a time in parallel over m and parallel annealing one iteration at a time in parallel over n. The estimators Eq. (A.16) and Eq. (A.19) are the same functions of this array with the roles of m and n interchanged, and Syed et al. (2026, Section 6.3) show that under the two performance models $\widehat{Z}^{\mathrm{RE}}$ and $\widehat{Z}^{\mathrm{SMC}}$ with resampling at every level have the same law. Consequently the schedule that minimises the variance of $\widehat{Z}^{\mathrm{SMC}}$ also minimises the variance of $\widehat{Z}^{\mathrm{RE}}$, and the same divergences between $\pi_{n-1} \otimes Q'_{n|n-1}$ and $\pi_n \otimes Q_{n-1|n}$ govern both: the rejection rate Eq. (2.7) through $r(p,q) \le \sqrt{\chi^2(q\|p)/2}$ (Lem. C.3), the discrepancy through $\log(1+\chi^2)$. The restart-rate barrier of parallel annealing and the variance barrier of sequential annealing scale with each other, and the schedules that optimise the two criteria are comparable (Syed et al., 2026, Section E.1).

B CREPE AND THE REVERSAL ERROR

CREPE (He et al., 2026b) is Alg. 2 on the ladder of Sec. 3.1 with the noising kernel as forward proposal, $Q'_{n|n-1} = \widehat{P}_{n|n-1}$, and with the backward weight $w_{n-1|n}$ of Eq. (3.2) used in place of the full incremental weight in the acceptance probability Eq. (2.5). In the notation of Eq. (3.2), the forward weight it drops is

$$w_{n|n-1}(x_{n-1}, x_n) = \frac{\widehat{P}_{n|n-1}(x_n \mid x_{n-1})}{P_{n|n-1}(x_n \mid x_{n-1})} = e^{-\delta_n(x_n, x_{n-1})}, \tag{B.1}$$

with δ_n the local reversal error Eq. (A.5), and CREPE sets it to one.

Proposition B.1. *The CREPE exchange at $(n-1, n)$ is $\pi_n \otimes \pi_{n-1}$-reversible if and only if $\delta_n \equiv 0$, i.e. $\widehat{P}_{n|n-1} = P_{n|n-1}$. Otherwise, with x_{n-1}, x_n the current replicas and $y_{n-1} \sim Q_{n-1|n}(\cdot \mid x_n)$, $y_n \sim \widehat{P}_{n|n-1}(\cdot \mid x_{n-1})$ the proposals, its log acceptance ratio differs from the correct one by*

$$\epsilon_n = \delta_n(x_n, y_{n-1}) - \delta_n(y_n, x_{n-1}), \tag{B.2}$$

one reversal error for each of the two pairs it exchanges. When the exchange spans several model steps, the errors of those steps add. The ENCORE ratio has $\epsilon_n \equiv 0$ for every sampler.

Proof. Write $w_{n-1,n}^{\widehat{P}}$ for the incremental weight Eq. (2.4) with forward proposal $\widehat{P}_{n|n-1}$, and $A_{n-1,n}^{\star}$ for the Metropolis–Hastings ratio of the move for the target $\pi_n \otimes \pi_{n-1}$.

The correct ratio. By Eq. (2.5),

$$A^{\star}_{n-1,n} = \frac{w^{\widehat{P}}_{n-1,n}(y_{n-1}, x_n)}{w^{\widehat{P}}_{n-1,n}(x_{n-1}, y_n)}. \tag{B.3}$$

The ratio CREPE uses. By Eq. (3.2) with $Q' = \widehat{P}$,

$$w^{\widehat{P}}_{n-1,n}(x_{n-1}, x_n) = w_{n-1|n}(x_{n-1}, x_n)\, \frac{\widehat{P}_{n|n-1}(x_n \mid x_{n-1})}{P_{n|n-1}(x_n \mid x_{n-1})} \tag{B.4}$$

$$= w_{n-1|n}(x_{n-1}, x_n)\, e^{-\delta_n(x_n, x_{n-1})}. \tag{B.5}$$

CREPE replaces each $w^{\widehat{P}}_{n-1,n}$ by $w_{n-1|n}$, so

$$A^{\mathrm{CREPE}}_{n-1,n} = \frac{w_{n-1|n}(y_{n-1}, x_n)}{w_{n-1|n}(x_{n-1}, y_n)} \tag{B.6}$$

$$= \frac{w^{\widehat{P}}_{n-1,n}(y_{n-1}, x_n)\, e^{\delta_n(x_n, y_{n-1})}}{w^{\widehat{P}}_{n-1,n}(x_{n-1}, y_n)\, e^{\delta_n(y_n, x_{n-1})}} \tag{B.7}$$

$$= A^{\star}_{n-1,n}\, \exp\left\{\delta_n(x_n, y_{n-1}) - \delta_n(y_n, x_{n-1})\right\} \tag{B.8}$$

$$= A^{\star}_{n-1,n}\, e^{\epsilon_n}. \tag{B.9}$$

Reversibility. If $\delta_n \equiv 0$ then $A^{\mathrm{CREPE}} = A^{\star}$ and the exchange is a Metropolis–Hastings kernel for $\pi_n \otimes \pi_{n-1}$, hence reversible. Conversely, if $\epsilon_n = 0$ for almost every configuration then δ_n is constant, $P_{n|n-1} = e^c \widehat{P}_{n|n-1}$, and $c = 0$ because both are probability kernels. If $\epsilon_n \not\equiv 0$, the accepted moves are those of a Metropolis–Hastings kernel with a tilted ratio, which does not preserve $\pi_n \otimes \pi_{n-1}$.

Compounding. When the exchange spans K model steps and the substitution is made step by step along the simulated intermediate path, as CREPE does, the reversal error of the composite move is the sum of the per-step errors along that path,

$$\delta_n = \sum_{k=1}^{K} \delta^{(k)}. \tag{B.10}$$

These do not cancel in general, so $|\epsilon_n|$ grows with K, and $\min\{1, e^{\epsilon_n} A^{\star}\}$ is driven towards zero across a coarse interval.

ENCORE. There is no forward proposal: by Prop. C.2 the ENCORE ratio is exactly the ratio of two values of $w_{n-1|n}$, and it is the correct Metropolis–Hastings ratio for $\boldsymbol{\pi}_n \otimes \boldsymbol{\pi}_{n-1}$. Hence $\epsilon_n \equiv 0$. □

C ENCORE: Derivations and Proofs

Throughout, in an exchange between levels n and $n-1$, $\mathbf{x}_n$ is the upper replica, $\mathbf{x}_{n-1} = (\mathbf{y}_n, x_{n-1})$ the lower one with $\mathbf{y}_n = \mathcal{T}_{n|n-1}(\mathbf{x}_{n-1})$ its truncation, $y_{n-1} \sim Q_{n-1|n}(\cdot \mid x_n)$ the proposed extension and $\mathbf{y}_{n-1} = (\mathbf{x}_n, y_{n-1})$, as in Sec. 3.2. $\mathcal{T}_{n|n-1} : \mathbb{X}_{n-1} \to \mathbb{X}_n$ drops the last state. The backward proposal $Q_{n-1|n}$ is fixed, and $w_{n-1|n}$ is the backward weight of Eq. (3.2),

$$w_{n-1|n}(x_{n-1}, x_n) = \frac{G_{n-1}(x_{n-1})}{G_n(x_n)}\, \frac{P_{n-1|n}(x_{n-1} \mid x_n)}{Q_{n-1|n}(x_{n-1} \mid x_n)}. \tag{C.1}$$

For a law $\boldsymbol{\pi}_n$ on $\mathbb{X}_n$ and a kernel $Q_{n-1|n}$, the product $\boldsymbol{\pi}_n \otimes Q_{n-1|n}$ is the law on $\mathbb{X}_{n-1}$ of "draw $\mathbf{x}_n \sim \boldsymbol{\pi}_n$, then $x_{n-1} \sim Q_{n-1|n}(\mathrm{d}x_{n-1} \mid x_n)$". Primes denote independent copies. Where densities are assumed to exist they are denoted by the same symbol as the measure. App. C.1 removes the assumption.

C.1 The Path Exchange

The exchange compares two laws on $\mathbb{X}_{n-1}$: the extended upper replica $\mathbf{y}_{n-1}$, which has law $\boldsymbol{\pi}_n \otimes Q_{n-1|n}$, and the lower replica $\mathbf{x}_{n-1}$, which has law $\boldsymbol{\pi}_{n-1}$. The first lemma is the statement in Sec. 3.2 that their ratio is the backward weight.

Lemma C.1 (Size-biasing identity)**.** *Let $\mathcal{Z}_n := \mathbb{E}_{\mathbf{p}_n}[G_n(X_n)]$ be the normalising constant of $\boldsymbol{\pi}_n$. Then*

$$\boldsymbol{\pi}_{n-1}(\mathrm{d}\mathbf{x}_{n-1}) = \frac{\mathcal{Z}_n}{\mathcal{Z}_{n-1}}\, w_{n-1|n}(x_{n-1}, x_n)\; (\boldsymbol{\pi}_n \otimes Q_{n-1|n})(\mathrm{d}\mathbf{x}_{n-1}). \tag{C.2}$$

Consequently, writing $W_{n-1|n} := w_{n-1|n}(x_{n-1}, x_n)$ for the weight of a draw from $\boldsymbol{\pi}_n \otimes Q_{n-1|n}$,

$$\mathbb{E}[W_{n-1|n}] = \frac{\mathcal{Z}_{n-1}}{\mathcal{Z}_n} \qquad \textit{and} \qquad \frac{\mathrm{d}\boldsymbol{\pi}_{n-1}}{\mathrm{d}(\boldsymbol{\pi}_n \otimes Q_{n-1|n})}(\mathbf{x}_{n-1}) = \frac{w_{n-1|n}(x_{n-1}, x_n)}{\mathbb{E}[W_{n-1|n}]}. \tag{C.3}$$

Proof. By the definition Eq. (3.4) of $\boldsymbol{\pi}_{n-1}$ and the path law Eq. (3.3),

$$\begin{aligned}\boldsymbol{\pi}_{n-1}(\mathrm{d}\mathbf{x}_{n-1}) &= \frac{G_{n-1}(x_{n-1})}{\mathcal{Z}_{n-1}}\, \mathbf{p}_{n-1}(\mathrm{d}\mathbf{x}_{n-1}) \\ &= \frac{G_{n-1}(x_{n-1})}{\mathcal{Z}_{n-1}}\, \mathbf{p}_n(\mathrm{d}\mathbf{x}_n)\, P_{n-1|n}(\mathrm{d}x_{n-1} \mid x_n).\end{aligned}$$

Substitute $\mathbf{p}_n = \mathcal{Z}_n \boldsymbol{\pi}_n / G_n$, which is Eq. (3.4) at level n, and write $P_{n-1|n}(\mathrm{d}x_{n-1} \mid x_n)$ as its density with respect to $Q_{n-1|n}(\mathrm{d}x_{n-1} \mid x_n)$, which exists because $P_{n-1|n}(\cdot \mid x_n) \ll Q_{n-1|n}(\cdot \mid x_n)$:

$$\begin{aligned}\boldsymbol{\pi}_{n-1}(\mathrm{d}\mathbf{x}_{n-1}) &= \frac{G_{n-1}(x_{n-1})}{\mathcal{Z}_{n-1}}\, \frac{\mathcal{Z}_n\, \boldsymbol{\pi}_n(\mathrm{d}\mathbf{x}_n)}{G_n(x_n)}\, \frac{P_{n-1|n}(x_{n-1} \mid x_n)}{Q_{n-1|n}(x_{n-1} \mid x_n)}\, Q_{n-1|n}(\mathrm{d}x_{n-1} \mid x_n) \\ &= \frac{\mathcal{Z}_n}{\mathcal{Z}_{n-1}}\, w_{n-1|n}(x_{n-1}, x_n)\; \boldsymbol{\pi}_n(\mathrm{d}\mathbf{x}_n)\, Q_{n-1|n}(\mathrm{d}x_{n-1} \mid x_n).\end{aligned}$$

Every factor of $\mathbf{p}_n$ has cancelled, which is why no marginal appears in the weight. Integrating both sides over $\mathbb{X}_{n-1}$ gives $1 = (\mathcal{Z}_n/\mathcal{Z}_{n-1})\mathbb{E}[W_{n-1|n}]$, which is the first statement of Eq. (C.3), and dividing Eq. (C.2) by $\boldsymbol{\pi}_n \otimes Q_{n-1|n}$ gives the second. □

Proposition C.2 (Acceptance ratio)**.** *The Metropolis–Hastings ratio of the path exchange $(\mathbf{x}_{n-1}, \mathbf{x}_n) \mapsto (\mathbf{y}_{n-1}, \mathbf{y}_n)$ of Sec. 3.2 for the target $\boldsymbol{\pi}_n \otimes \boldsymbol{\pi}_{n-1}$ is*

$$A_{n-1,n} = \frac{w_{n-1|n}(y_{n-1}, x_n)}{w_{n-1|n}(x_{n-1}, y_n)}, \tag{C.4}$$

so that $\boldsymbol{\alpha}_{n-1,n} = \min\{1, A_{n-1,n}\}$ is Eq. (3.5).

Proof. The move is a reversible-jump exchange (Green, 1995). The forward proposal draws $y_{n-1} \sim Q_{n-1|n}(\mathrm{d}y_{n-1} \mid x_n)$ and maps

$$\big(\mathbf{x}_n,\ \mathbf{x}_{n-1}\big) = \big(\mathbf{x}_n,\ (\mathbf{y}_n, x_{n-1})\big) \longmapsto \big(\mathbf{y}_n,\ (\mathbf{x}_n, y_{n-1})\big) = \big(\mathbf{y}_n,\ \mathbf{y}_{n-1}\big);$$

the reverse proposal draws $x_{n-1} \sim Q_{n-1|n}(\mathrm{d}x_{n-1} \mid y_n)$ and maps back. On the augmented space $\mathbb{X}_n \times \mathbb{X}_{n-1} \times \mathbb{X}$ the dimension-matching map is a permutation of coordinates, so its Jacobian is one, and the ratio is the target density at the proposed pair times the reverse proposal density, over the target density at the current pair times the forward proposal density:

$$A_{n-1,n} = \frac{\boldsymbol{\pi}_n(\mathbf{y}_n)\; \boldsymbol{\pi}_{n-1}(\mathbf{x}_n, y_{n-1})\; Q_{n-1|n}(x_{n-1} \mid y_n)}{\boldsymbol{\pi}_n(\mathbf{x}_n)\; \boldsymbol{\pi}_{n-1}(\mathbf{y}_n, x_{n-1})\; Q_{n-1|n}(y_{n-1} \mid x_n)} \tag{C.5}$$

$$= \frac{\boldsymbol{\pi}_{n-1}(\mathbf{x}_n, y_{n-1}) \;/\; \big[\boldsymbol{\pi}_n(\mathbf{x}_n)\, Q_{n-1|n}(y_{n-1} \mid x_n)\big]}{\boldsymbol{\pi}_{n-1}(\mathbf{y}_n, x_{n-1}) \;/\; \big[\boldsymbol{\pi}_n(\mathbf{y}_n)\, Q_{n-1|n}(x_{n-1} \mid y_n)\big]} \tag{C.6}$$

$$= \frac{(\mathcal{Z}_n/\mathcal{Z}_{n-1})\; w_{n-1|n}(y_{n-1}, x_n)}{(\mathcal{Z}_n/\mathcal{Z}_{n-1})\; w_{n-1|n}(x_{n-1}, y_n)} \tag{C.7}$$

$$= \frac{w_{n-1|n}(y_{n-1}, x_n)}{w_{n-1|n}(x_{n-1}, y_n)}. \tag{C.8}$$

The second line regroups the six factors into two densities of $\boldsymbol{\pi}_{n-1}$ with respect to $\boldsymbol{\pi}_n \otimes Q_{n-1|n}$, one at the proposed pair $(\mathbf{x}_n, y_{n-1})$ and one at the current pair $(\mathbf{y}_n, x_{n-1})$. The third applies Lem. C.1 to each, and the constant cancels. □

Unguided exchange. For $Q_{n-1|n} = P_{n-1|n}$ the proposal correction in $w_{n-1|n}$ is one and the ratio reduces to

$$A_{n-1,n} = \frac{G_n(y_n)\, G_{n-1}(y_{n-1})}{G_n(x_n)\, G_{n-1}(x_{n-1})}, \tag{C.9}$$

the potentials of the two proposed states over the potentials of the two current ones, with no model density.

Kernels without densities. The proof used the densities of P and Q only through the Radon–Nikodym derivative $\mathrm{d}P_{n-1|n}(\cdot \mid x_n)/\mathrm{d}Q_{n-1|n}(\cdot \mid x_n)$, which exists under the assumption $P_{n-1|n}(\cdot \mid x_n) \ll Q_{n-1|n}(\cdot \mid x_n)$ of Sec. 3.1, so $w_{n-1|n}$, Lem. C.1 and Prop. C.2 are defined for kernels without densities. For $Q_{n-1|n} = P_{n-1|n}$ the derivative equals one almost surely, so the unguided ratio holds even when $P_{n-1|n}$ is deterministic, as for a distilled or ODE sampler.

Path-dependent potentials and proposals. Nothing above requires the potentials or proposals to depend on the last state only. With $G_n : \mathbb{X}_n \to [0, \infty)$ and $Q_{n-1|n}(\mathrm{d}z \mid \mathbf{x}_n)$, the weight becomes

$$w_{n-1|n}(\mathbf{x}_n, y_{n-1}) = \frac{G_{n-1}(\mathbf{x}_n, y_{n-1})}{G_n(\mathbf{x}_n)} \frac{P_{n-1|n}(y_{n-1} \mid x_n)}{Q_{n-1|n}(y_{n-1} \mid \mathbf{x}_n)}, \tag{C.10}$$

and Lem. C.1, Prop. C.2 and Thm. 3.1(a),(b) hold verbatim with $G_n(\mathbf{x}_n)$ in place of $G_n(x_n)$. This covers rewards on whole trajectories, at the cost that the exchange then reads the stored history rather than its last state. The same weight is the incremental weight of an SMC sampler on the trajectory ladder: extending a particle $\mathbf{x}_n$ with $Q_{n-1|n}$ and weighting by $\mathrm{d}\boldsymbol{\pi}_{n-1}/\mathrm{d}(\boldsymbol{\pi}_n \otimes Q_{n-1|n})$ gives $w_{n-1|n}(y_{n-1}, x_n)$ up to a constant by Lem. C.1. SMC uses $w_{n-1|n}$ to reweight a population. ENCORE uses the ratio of two evaluations of it.

C.2 Proof of Thm. 3.1

The three parts are proved in turn. Part (a) is invariance and the rejection-rate identity, part (b) is uniform ergodicity under bounded weights, and part (c) is the characterisation of the maximal restart rate.

C.2.1 Part (a): invariance and the rejection rate

Proof of invariance. The exchange at $(n-1, n)$ is a Metropolis–Hastings kernel with the ratio of Prop. C.2, hence reversible with respect to $\boldsymbol{\pi}_n \otimes \boldsymbol{\pi}_{n-1}$. Extended by the identity on the other levels, it leaves $\boldsymbol{\pi}_N \otimes \cdots \otimes \boldsymbol{\pi}_0$ invariant. Since $G_N \equiv 1$, $\boldsymbol{\pi}_N = \eta$, so redrawing $x_N \sim \eta$ leaves the level-N marginal invariant, and any $\boldsymbol{\pi}_n$-invariant local move leaves $\boldsymbol{\pi}_N \otimes \cdots \otimes \boldsymbol{\pi}_0$ invariant. Invariance is preserved under composition, so one iteration of Alg. 1 leaves $\boldsymbol{\pi}_N \otimes \cdots \otimes \boldsymbol{\pi}_0$ invariant. The $\mathbf{x}_0$-marginal of $\boldsymbol{\pi}_N \otimes \cdots \otimes \boldsymbol{\pi}_0$ is $\boldsymbol{\pi}_0 \propto G_0(x_0)\, \mathbf{p}_0$, whose x_0-marginal is

$$\int_{\mathbf{x}_0:\ \text{last state } x_0} \boldsymbol{\pi}_0(\mathrm{d}\mathbf{x}_0) \ \propto\ G_0(x_0) \int_{\mathbf{x}_0:\ \text{last state } x_0} \mathbf{p}_0(\mathrm{d}\mathbf{x}_0) = G_0(x_0)\, p_0(\mathrm{d}x_0),$$

which is π_0. Hence the recorded state has stationary law π_0. □

The rejection rate is computed in three steps. The path exchange is a swap between two laws on $\mathbb{X}_{n-1}$. The rejection rate of a swap is a total variation. And the total variation on trajectory space equals the total variation of the projections onto the last two states, which are the two laws in Eq. (3.6). Fix n and abbreviate

$$p := \boldsymbol{\pi}_n \otimes Q_{n-1|n} \quad \text{(law of the extended upper replica)}, \qquad q := \boldsymbol{\pi}_{n-1} \quad \text{(law of the lower replica)}, \tag{C.11}$$

both on $\mathbb{X}_{n-1}$. Let $W_{n-1|n} := w_{n-1|n}(x_{n-1}, x_n)$ under p, let $W'_{n-1|n}$ be an independent copy, and let $\bar{W}_{n-1|n} := W_{n-1|n}/\mathbb{E}[W_{n-1|n}]$, so that $\mathrm{d}q/\mathrm{d}p = \bar{W}_{n-1|n}$ by Lem. C.1.

Lemma C.3 (Rejection rate of a swap)**.** *Let p, q be laws on a common space, $u \sim p$ and $v \sim q$ independent, and consider the swap $(u, v) \mapsto (v, u)$ accepted with the Metropolis–Hastings probability $\min\{1, (\mathrm{d}q/\mathrm{d}p)(u)/(\mathrm{d}q/\mathrm{d}p)(v)\}$ for the target $p \otimes q$. Its rejection rate is*

$$r(p,q) = \mathrm{TV}(p \otimes q,\ q \otimes p) = \tfrac{1}{2}\,\mathbb{E}\big|\bar{W}(u) - \bar{W}(u')\big|, \qquad \bar{W} := \frac{\mathrm{d}q}{\mathrm{d}p}, \quad u, u' \sim p \text{ i.i.d.}, \tag{C.12}$$

and it satisfies

$$\mathrm{TV}(p,q) \;\le\; r(p,q) \;\le\; \min\Big\{2\ \mathrm{TV}(p,q),\ \sqrt{\mathrm{Var}(\bar{W})/2}\Big\}. \tag{C.13}$$

Proof. The rejection rate is a total variation. The acceptance probability is an expectation of a minimum, which becomes an integral of a minimum of two products of densities:

$$\begin{aligned} 1 - r(p,q) &= \mathbb{E}\Big[\min\Big\{1, \frac{(\mathrm{d}q/\mathrm{d}p)(u)}{(\mathrm{d}q/\mathrm{d}p)(v)}\Big\}\Big] \\ &= \iint p(u)\,q(v)\ \min\Big\{1, \frac{q(u)\,p(v)}{p(u)\,q(v)}\Big\}\,\mathrm{d}u\,\mathrm{d}v \\ &= \iint \min\big\{p(u)q(v),\ q(u)p(v)\big\}\,\mathrm{d}u\,\mathrm{d}v. \end{aligned}$$

For probability densities f, g one has $\int \min\{f, g\} = 1 - \mathrm{TV}(f, g)$. Taking $f = p \otimes q$ and $g = q \otimes p$ gives $r(p,q) = \mathrm{TV}(p \otimes q, q \otimes p)$.

The weight form. Substitute $q = \bar{W}p$ in the last integral and rewrite it as an integral over two independent draws from p:

$$\begin{aligned} \mathrm{TV}(p \otimes q,\ q \otimes p) &= \tfrac{1}{2}\iint \big|p(u)q(u') - q(u)p(u')\big|\,\mathrm{d}u\,\mathrm{d}u' \\ &= \tfrac{1}{2}\iint \big|\bar{W}(u') - \bar{W}(u)\big|\,p(u)\,p(u')\,\mathrm{d}u\,\mathrm{d}u' = \tfrac{1}{2}\,\mathbb{E}\big|\bar{W}(u) - \bar{W}(u')\big|. \end{aligned}$$

Lower bound. Since $\mathbb{E}\bar{W} = 1$, $\mathrm{TV}(p,q) = \frac{1}{2}\mathbb{E}|\bar{W}(u) - 1|$. Conditioning on $\bar{W}(u)$ and applying Jensen's inequality to the convex map $s \mapsto |\bar{W}(u) - s|$,

$$\mathbb{E}\big[\,|\bar{W}(u) - \bar{W}(u')|\ \big|\ \bar{W}(u)\big] \;\ge\; \big|\bar{W}(u) - \mathbb{E}\bar{W}(u')\big| = |\bar{W}(u) - 1|,$$

and taking expectations gives $r(p,q) \ge \mathrm{TV}(p,q)$.

Upper bounds. The triangle inequality $|\bar{W}(u) - \bar{W}(u')| \le |\bar{W}(u) - 1| + |1 - \bar{W}(u')|$ gives $r(p,q) \le 2\ \mathrm{TV}(p,q)$. By Cauchy–Schwarz and independence,

$$\mathbb{E}|\bar{W}(u) - \bar{W}(u')| \;\le\; \sqrt{\mathbb{E}\big[(\bar{W}(u) - \bar{W}(u'))^2\big]} = \sqrt{2\,\mathrm{Var}(\bar{W})},$$

hence $r(p,q) \le \sqrt{\mathrm{Var}(\bar{W})/2}$. □

Proof of the rejection rate Eq. (3.6). The path exchange is a swap between p and q. Under $\boldsymbol{\pi}_N \otimes \cdots \otimes \boldsymbol{\pi}_0$ the replicas at levels n and $n-1$ are independent, so the proposal $u := \mathbf{y}_{n-1} = (\mathbf{x}_n, y_{n-1})$ has law p and the lower replica $v := \mathbf{x}_{n-1}$ has law q, independently. By Prop. C.2 and Lem. C.1 the acceptance ratio is

$$A_{n-1,n} = \frac{w_{n-1|n}(u)}{w_{n-1|n}(v)} = \frac{\bar{W}_{n-1|n}(u)}{\bar{W}_{n-1|n}(v)} = \frac{(\mathrm{d}q/\mathrm{d}p)(u)}{(\mathrm{d}q/\mathrm{d}p)(v)},$$

which is the acceptance ratio of the swap in Lem. C.3. Hence

$$r_{n-1,n} = r(p,q) = \tfrac{1}{2}\,\mathbb{E}|\bar{W}_{n-1|n} - \bar{W}'_{n-1|n}|, \tag{C.14}$$

with $\bar{W}_{n-1|n}$ the normalised backward weight, and the bounds Eq. (C.13) apply. $\mathrm{Var}(\bar{W}_{n-1|n})$ is the squared coefficient of variation of the sequential control weight $w_{n-1|n}$ under p, so a level at which SMC would lose particle diversity is a level at which ENCORE rejects.

Projection to the state space. Let pr map a trajectory to its last two states, $\mathrm{pr}(\mathbf{x}_{n-1}) = (x_n, x_{n-1})$. By Lem. C.1, $\mathrm{d}q/\mathrm{d}p = \bar{W}_{n-1|n}$ depends on the trajectory only through pr, say $\bar{W}_{n-1|n} = h \circ \mathrm{pr}$. Then the projected laws $p_{\mathrm{pr}} := p \circ \mathrm{pr}^{-1}$ and $q_{\mathrm{pr}} := q \circ \mathrm{pr}^{-1}$ have the same density: for every measurable B,

$$q_{\mathrm{pr}}(B) = q(\mathrm{pr}^{-1}B) = \int_{\mathrm{pr}^{-1}B} h \circ \mathrm{pr}\, \mathrm{d}p = \int_B h\, \mathrm{d}p_{\mathrm{pr}},$$

so $\mathrm{d}q_{\mathrm{pr}}/\mathrm{d}p_{\mathrm{pr}} = h$. Since $r(p,q) = \frac{1}{2}\mathbb{E}|\bar{W}_{n-1|n}(u) - \bar{W}_{n-1|n}(u')|$ depends only on the law of $\bar{W}_{n-1|n}$ under p, which is the law of h under p_{pr}, we have $r(p,q) = r(p_{\mathrm{pr}}, q_{\mathrm{pr}})$.

The projections. Under p, the last state x_n of $\mathbf{x}_n \sim \boldsymbol{\pi}_n$ has law π_n and $x_{n-1} \sim Q_{n-1|n}(\cdot \mid x_n)$, so $p_{\mathrm{pr}} = \pi_n \otimes Q_{n-1|n}$. Under q, by Eq. (3.4), Eq. (3.3) and the time reversal Eq. (A.3),

$$\begin{aligned} q_{\mathrm{pr}}(\mathrm{d}x_n, \mathrm{d}x_{n-1}) &\propto G_{n-1}(x_{n-1}) \int_{\mathbf{x}_n:\ \text{last state } x_n} \mathbf{p}_n(\mathrm{d}\mathbf{x}_n)\, P_{n-1|n}(\mathrm{d}x_{n-1} \mid x_n) \\ &= G_{n-1}(x_{n-1})\, p_n(\mathrm{d}x_n)\, P_{n-1|n}(\mathrm{d}x_{n-1} \mid x_n) \\ &= G_{n-1}(x_{n-1})\, p_{n-1}(\mathrm{d}x_{n-1})\, P_{n|n-1}(\mathrm{d}x_n \mid x_{n-1}), \end{aligned}$$

so $q_{\mathrm{pr}} = \pi_{n-1} \otimes P_{n|n-1}$. Therefore

$$r_{n-1,n} = r\big(\pi_{n-1} \otimes P_{n|n-1},\ \pi_n \otimes Q_{n-1|n}\big),$$

which is Eq. (3.6), and by Eq. (2.7) this is the rejection rate of Alg. 2 with $Q'_{n|n-1} = P_{n|n-1}$. □

C.2.2 Part (b): uniform ergodicity

Assume $0 < \underline{w}_{n-1|n} \le w_{n-1|n} \le \overline{w}_{n-1|n}$ for every n and set

$$c_{n-1,n} := \frac{\underline{w}_{n-1|n}}{\overline{w}_{n-1|n}} \in (0,1]. \tag{C.15}$$

Since $A_{n-1,n}$ is a ratio of two values of $w_{n-1|n}$, every exchange is accepted with probability at least $c_{n-1,n}$, whatever the state of the ladder. The proof is a Doeblin minorisation of the whole ladder over a window of iterations long enough for every replica to be discarded and replaced.

Proof. Realise every acceptance as $\mathbf{1}\{U \le \boldsymbol{\alpha}_{n-1,n}\}$ with independent uniforms U, and every draw from η, from a proposal Q and from a local move as a measurable function of the current state and of independent noise. Write $\mathbf{X} = (\mathbf{x}_N, \dots, \mathbf{x}_0)$ for the ladder and $\mathcal{K}$ for the one-iteration kernel of Alg. 1. A minorisation of a single level does not suffice, because the future of one replica depends on the others, so we minorise the whole chain.

Step 1: what happens when every exchange is accepted. Under the even–odd schedule the pair $(n-1, n)$ is attempted at the iterations $m \equiv n \pmod 2$, so a replica at a level $0 < n < N$ takes part in exactly one attempted exchange per iteration. Suppose every attempted exchange in a window of iterations is accepted. A replica that moves down from level n to $n-1$ at iteration m is in the pair $(n-2, n-1)$ at iteration $m+1$, since $n-1 \equiv m+1 \pmod 2$, and moves down again. A replica that moves up keeps moving up in the same way. So each replica descends monotonically to level 0, idles there for one iteration, ascends monotonically to level N, and is discarded at the start of the next iteration, when the reference level is redrawn. The longest such journey, for a replica that has just entered at level N, is

$$\underbrace{N}_{\text{descend}} + \underbrace{1}_{\text{idle at 0}} + \underbrace{N}_{\text{ascend}} + \underbrace{1}_{\text{discard}} = 2N + 2 \text{ iterations},$$

so a replica present at the start of the window is discarded within $W := 2N + 2$ iterations. After W iterations every level holds a replica that entered at the reference level inside the window, extended by proposal draws and truncated along the way.

Step 2: the all-accept event is independent of the states. Let E be the event that every uniform used in the window satisfies $U \le c_{n-1,n}$ for its pair. Since $\{U \le c_{n-1,n}\} \subseteq \{U \le \boldsymbol{\alpha}_{n-1,n}\}$ whatever the states, every attempted exchange is accepted on E. Since E is an event in the uniforms alone, it is independent of the initial ladder $\mathbf{X}$ and of the reference, proposal and local-move noise. Each pair is attempted every second iteration, hence at most $N+1$ times in the window, so

$$\Pr(E) \;\ge\; \epsilon := \prod_{n=1}^{N} c_{n-1,n}^{N+1} > 0. \tag{C.16}$$

Step 3: minorisation. On E the evolution over the window is the all-accept dynamics of Step 1. By Step 1 the ladder at the end of the window is a function $F(\text{noise})$ of the noise drawn inside the window alone: every state that depended on $\mathbf{X}$ has been discarded, and the local moves only alter states that are themselves functions of that noise. Let ν be the law of $F(\text{noise})$. It does not depend on $\mathbf{X}$. For every measurable B,

$$\mathcal{K}^W(\mathbf{X}, B) \;\ge\; \Pr\big(E \cap \{F(\text{noise}) \in B\}\big) = \Pr(E)\,\nu(B) \;\ge\; \epsilon\,\nu(B), \tag{C.17}$$

using the independence of E from the noise. This is a Doeblin condition for $\mathcal{K}^W$, uniform in $\mathbf{X}$.

Step 4: conclusion. By the coupling argument for Doeblin chains (Meyn & Tweedie, 2009, Thm. 16.2.4), for every $\mathbf{X}$ and k,

$$\big\|\mathcal{K}^k(\mathbf{X}, \cdot) - \boldsymbol{\pi}_N \otimes \cdots \otimes \boldsymbol{\pi}_0\big\|_{\mathrm{TV}} \;\le\; (1-\epsilon)^{\lfloor k/W \rfloor}, \tag{C.18}$$

so the chain is uniformly ergodic, and the law of the recorded state x_0 converges to π_0 at this rate. □

Remark C.4. The hypothesis holds for bounded rewards, $|r| \le C$ with $G_n = \exp\{\kappa_n r\}$, $0 \le \kappa_n \le \kappa$, and $Q = P$, since then $e^{-2\kappa C} \le w_{n-1|n} \le e^{2\kappa C}$. It fails for unbounded tilts such as the Boltzmann reweighting, for which a drift condition on $\log G_n$ would be required.

C.2.3 Part (c): the maximal restart rate and the Doob relations

Proof. By Eq. (2.6), $\tau \le \frac{1}{2}$ with equality if and only if $r_{n-1,n} = 0$ for every n, since $s \mapsto s/(1-s)$ is non-negative on $[0,1)$ and vanishes only at $s = 0$. By Lem. C.3, $r_{n-1,n} = 0$ if and only if the two laws in Eq. (3.6) coincide,

$$\pi_{n-1} \otimes P_{n|n-1} = \pi_n \otimes Q_{n-1|n} \qquad \text{as laws of } (x_n, x_{n-1}). \tag{C.19}$$

Two laws on a product agree if and only if their x_n-marginals agree and their conditionals of x_{n-1} given x_n agree, so we compute both sides.

The two laws. The right side has density proportional to $G_n(x_n)\,p_n(x_n)\,Q_{n-1|n}(x_{n-1} \mid x_n)$. The left side, by the time reversal Eq. (A.3), has density proportional to

$$G_{n-1}(x_{n-1})\,p_{n-1}(x_{n-1})\,P_{n|n-1}(x_n \mid x_{n-1}) = G_{n-1}(x_{n-1})\,p_n(x_n)\,P_{n-1|n}(x_{n-1} \mid x_n).$$

Marginals. Integrating out x_{n-1},

$$\text{right side:} \quad \pi_n(\mathrm{d}x_n) \propto G_n(x_n)\,p_n(\mathrm{d}x_n),$$

$$\text{left side:} \quad \propto p_n(\mathrm{d}x_n) \int G_{n-1}(y)\,P_{n-1|n}(\mathrm{d}y \mid x_n) = \big(P_{n-1|n}G_{n-1}\big)(x_n)\,p_n(\mathrm{d}x_n).$$

The marginals agree if and only if $G_n \propto P_{n-1|n}G_{n-1}$, the potential relation of Eq. (A.9).

Conditionals. Given x_n, the right side draws $x_{n-1} \sim Q_{n-1|n}(\cdot \mid x_n)$, and the left side draws x_{n-1} from

$$\frac{G_{n-1}(y)\,P_{n-1|n}(\mathrm{d}y \mid x_n)}{(P_{n-1|n}G_{n-1})(x_n)}.$$

They agree if and only if $Q_{n-1|n} \propto G_{n-1}P_{n-1|n}$, the proposal relation of Eq. (A.9).

Hence $\tau = \frac{1}{2}$ if and only if the two Doob relations hold at every level n. By App. A.2 this is $G_n \propto G_n^*$ and $Q_{n-1|n} = Q^*_{n-1|n}$ for every n, with the top-level caveat noted there. In that case every acceptance ratio equals one, each reference draw reaches level 0 through accepted exchanges, and its trajectory is a draw from $\eta \otimes Q^*_{N-1|N} \otimes \cdots \otimes Q^*_{0|1}$, the Doob-transformed chain of Eq. (A.8), whose terminal marginal is π_0. □

C.3 Cost and Memory

Kernel simulations. Each iteration of Alg. 1 simulates one proposal $Q_{n-1|n}$ for each of the $N/2$ pairs active under the even–odd schedule, in parallel, and two consecutive iterations extend every level once. When the levels coincide with the model steps, two iterations therefore cost the same N kernel simulations as one pass of the sampler, and per sample the cost matches sequential control (Tab. 1). When a level spans several model steps, $Q_{n-1|n}$ composes them. The intermediate states are used once in the acceptance ratio and discarded. The optional local move costs one further proposal per level.

Memory. The upward move replaces the lower replica by its truncation $\mathcal{T}_{n|n-1}(\mathbf{x}_{n-1})$, so the initial segment must be stored. The replica at level n holds its states at the levels $N, \dots, n$ it has passed through, $N - n + 1$ states, and the ladder holds

$$\sum_{n=0}^{N}(N - n + 1) = \frac{(N+1)(N+2)}{2} = O(N^2) \tag{C.20}$$

states in all, $O(N^2 d)$ against $O(Nd)$ for state-space replica exchange. Only states at the levels are stored, not at model steps in between, so a coarse ladder is cheap in memory as well as in communication. For a deterministic sampler with a tractable inverse the stored prefix can be recomputed from the last state by inverting the transitions, which trades the $O(N^2 d)$ memory for extra model evaluations.

D Diffusion Models: Communication Barriers

This appendix proves Thm. 3.2 and develops its consequences in the order of Sec. 3.3: the guided proposal against the Doob proposal, schedule tuning, the optimal spacing of an unguided ladder, the corresponding SMC barrier, and a closed-form example.

D.1 Setting and Assumptions

We look at the exchange between the levels at times $t = t_n$ and $t' = t_{n-1} = t - \Delta t_n$, and index the objects attached to it by n. The same computation applies to every level. Let $X_t \sim \pi_t$ be the last state of the upper replica and $Y_{t'}$ the proposed extension produced by the step Eq. (3.7) with noise $\xi \sim N(0, I)$ independent of X_t. The log-weight of Eq. (3.8) at level n is

$$\ell_n := \log w_{t'|t}(X_t, \xi) = \log \frac{G_{t'}(Y_{t'})}{G_t(X_t)} + \mathbf{1}_{\text{guided}} \log \frac{P_{t'|t}(Y_{t'} \mid X_t)}{Q_{t'|t}(Y_{t'} \mid X_t)}, \tag{D.1}$$

ℓ'_n is an independent copy, and by Eq. (C.14) the rejection rate of the exchange is

$$r_{n-1,n} = \frac{\mathbb{E}|e^{\ell_n} - e^{\ell'_n}|}{2\,\mathbb{E}e^{\ell_n}}. \tag{D.2}$$

We write $m_{1,n} := \mathbb{E}|\ell_n|$ and $m_{4,n} := \mathbb{E}\ell_n^4$ for the moments that control the expansion of Eq. (D.2). Recall $g_t := \sigma_t^\top \nabla \log G_t$ and $M_t := \sigma_t^\top \nabla^2 \log G_t\, \sigma_t$ from Sec. 3.3, and define the Doob residual

$$\psi_t := \frac{\mathcal{A}_t G_t}{G_t} = -\partial_t \log G_t + b_t \cdot \nabla \log G_t + \tfrac{1}{2} \operatorname{tr} M_t + \tfrac{1}{2}\|g_t\|^2, \tag{D.3}$$

the second equality by the identity $G^{-1}\mathcal{A}_t G = \mathcal{A}_t \log G + \frac{1}{2}\|\sigma_t^\top \nabla \log G\|^2$, which follows from $\nabla^2 G/G = \nabla^2 \log G + \nabla \log G\, \nabla \log G^\top$. With this notation Eq. (3.10) reads

$$\Psi_t(x, \xi) = \psi_t(x) + \tfrac{1}{2}\big(\xi^\top M_t(x)\xi - \operatorname{tr} M_t(x)\big), \qquad \mathbb{E}\big[\Psi_t(X_t, \xi) \mid X_t\big] = \psi_t(X_t), \tag{D.4}$$

the conditional mean because $\mathbb{E}[\xi^\top M \xi] = \operatorname{tr} M$ for $\xi \sim N(0, I)$.

Assumptions.

(A1) σ_t is invertible for every t, and $t \mapsto \sigma_t$ and $t \mapsto \sigma_t^{-1}$ are continuous on $[0,1]$.

(A2) $(t,x) \mapsto b_t(x)$ is $C^{0,3}$ and $(t,x) \mapsto \log G_t(x)$ is $C^{1,4}$, with all these derivatives bounded on $[0,1] \times \mathbb{X}$.

(A3) The sampler and guided steps are the two cases of Eq. (3.7), with $\xi \sim N(0,I)$ independent of X_t.

(A2) makes $\nabla \log G_t$, $\nabla^2 \log G_t$ and $\partial_t \log G_t$, hence g_t, M_t and ψ_t, bounded, and every remainder below uniform in x. Under (A1)–(A2) the diffusion is uniformly elliptic with bounded smooth drift, so its transition densities are everywhere positive, and hence so are p_t and $\pi_t \propto G_t p_t$ for every $t < 1$. Replacing the Euler–Maruyama steps by exact transitions of the two SDEs changes only the $O(\Delta t_n^{3/2})$ remainders, by the same expansions applied through the Itô formula.

Remark D.1 (On the assumptions). (A2) is a convenience, not a description of practice: it is what makes the Taylor expansions uniform, and it excludes settings of interest. A Boltzmann tilt $\log G_0 = -\beta U$ with U of unbounded gradient violates it, as does constrained generation with $G_0 = \mathbf{1}_{\mathcal{C}}$, for which $\log G_0$ is not finite. The expansions remain valid in the first case on any region where the stated derivatives are bounded, and localisation arguments extend them under polynomial growth with Gaussian tails for π_t, at the price of a longer proof. The second case is outside the scope of an expansion in Δt_n altogether, since the weight is then an indicator. We state Thm. 3.2 under (A2) and read it as a guide to the behaviour of the sampler rather than as a covering theorem for every application.

D.2 Three Lemmas

The first lemma turns the rejection rate Eq. (D.2) into a mean absolute difference of log-weights, the second evaluates that mean for Gaussian log-weights, and the third is the pair of Taylor expansions behind Eq. (3.9).

Lemma D.2 (Small-weight expansion). *Suppose* $\sup_n \mathbb{E} e^{2|\ell_n|} \le C_0$. *Then there is a constant* C *depending only on* C_0 *such that, for every* n,

$$\left| r_{n-1,n} - \tfrac{1}{2}\,\mathbb{E}|\ell_n - \ell'_n| \right| \;\le\; C\big(m_{4,n}^{1/2} + m_{1,n}^2\big). \tag{D.5}$$

Lem. D.5 below verifies the hypothesis and bounds $m_{1,n}$ *and* $m_{4,n}$ *by powers of* Δt_n *uniformly in* n, *so that the right side is* $O(\Delta t_n)$ *in the unguided case and* $O(\Delta t_n^2)$ *in the guided case.*

Proof. Fix n and drop the index. The proof compares the numerator and the denominator of Eq. (D.2) with their first-order expansions.

Step 1: expand the exponential. Write $e^u = 1+u+\rho(u)$ with $|\rho(u)| \le \frac{1}{2}u^2 e^{|u|}$. By Cauchy–Schwarz and the hypothesis,

$$\mathbb{E}|\rho(\ell)| \;\le\; \tfrac{1}{2}\big(\mathbb{E}\ell^4\big)^{1/2}\big(\mathbb{E}e^{2|\ell|}\big)^{1/2} \;\le\; \tfrac{1}{2}C_0^{1/2} m_4^{1/2}.$$

Step 2: the numerator. With $\mathrm{N} := \mathbb{E}|e^{\ell} - e^{\ell'}|$ and $e^{\ell} - e^{\ell'} = (\ell - \ell') + (\rho(\ell) - \rho(\ell'))$, the triangle inequality gives

$$\mathrm{N} = \mathbb{E}|\ell - \ell'| + \rho_{\mathrm{N}}, \qquad |\rho_{\mathrm{N}}| \le \mathbb{E}|\rho(\ell)| + \mathbb{E}|\rho(\ell')| \le C_0^{1/2} m_4^{1/2}.$$

Step 3: the denominator. With $\mathrm{D} := \mathbb{E}e^{\ell} = 1 + \mathbb{E}\ell + \mathbb{E}\rho(\ell)$,

$$\mathrm{D} = 1 + \rho_{\mathrm{D}}, \qquad |\rho_{\mathrm{D}}| \le |\mathbb{E}\ell| + \mathbb{E}|\rho(\ell)| \le m_1 + \tfrac{1}{2}C_0^{1/2} m_4^{1/2}.$$

Step 4: combine. Since $r = \mathrm{N}/(2\mathrm{D})$,

$$r - \tfrac{1}{2}\mathbb{E}|\ell - \ell'| = \frac{\mathrm{N} - \mathbb{E}|\ell - \ell'|}{2\mathrm{D}} + \tfrac{1}{2}\mathbb{E}|\ell - \ell'|\Big(\frac{1}{\mathrm{D}} - 1\Big) = \frac{\rho_{\mathrm{N}}}{2\mathrm{D}} - \frac{\mathbb{E}|\ell - \ell'|}{2} \cdot \frac{\rho_{\mathrm{D}}}{\mathrm{D}}.$$

Suppose first that $|\rho_{\mathrm{D}}| \le \frac{1}{2}$, so that $\mathrm{D} \ge \frac{1}{2}$. Using $\mathbb{E}|\ell - \ell'| \le 2m_1$ and the bounds of Steps 2 and 3,

$$\left| r - \tfrac{1}{2}\mathbb{E}|\ell - \ell'| \right| \;\le\; |\rho_{\mathrm{N}}| + 2m_1|\rho_{\mathrm{D}}| \;\le\; C_0^{1/2} m_4^{1/2} + 2m_1^2 + C_0^{1/2} m_1 m_4^{1/2}.$$

The last term is at most $\frac{1}{2}C_0^{1/2}(m_1^2 + m_4)$ by the inequality $ab \le \frac{1}{2}(a^2 + b^2)$, and $m_4 \le m_4^{1/2}$ when $m_4 \le 1$, which gives Eq. (D.5). If $m_4 > 1$ or $|\rho_{\mathrm{D}}| > \frac{1}{2}$, then $m_4^{1/2} + m_1^2$ is bounded below by a constant depending only on C_0, while the left side is at most $1 + m_1 \le 2 + m_1^2$, so Eq. (D.5) holds with a larger C. The constant depends on C_0 only, hence not on n. □

Lemma D.3 (Gaussian mean absolute difference)**.** *Let $g : \mathbb{X} \to \mathbb{R}^d$ be bounded, X, X' i.i.d., and $\xi, \xi' \sim N(0, I)$ independent of (X, X') and of each other. Then*

$$\mathbb{E}\big|g(X)^\top \xi - g(X')^\top \xi'\big| = \sqrt{\tfrac{2}{\pi}}\, \mathbb{E}\sqrt{\|g(X)\|^2 + \|g(X')\|^2}. \tag{D.6}$$

Proof. Conditionally on (X, X'), $g(X)^\top \xi - g(X')^\top \xi'$ is Gaussian with mean zero and variance $\|g(X)\|^2 + \|g(X')\|^2$, and $\mathbb{E}|N(0, s^2)| = s\sqrt{2/\pi}$. Take expectations over (X, X'). □

Lemma D.4 (One-step expansions)**.** *Under (A1)–(A3), with all functions evaluated at X_t, the following hold uniformly in X_t.*

(a) For the sampler step ($\mathbf{1}_{\mathrm{guided}} = 0$),

$$\log \frac{G_{t'}(Y_{t'})}{G_t(X_t)} = \sqrt{\Delta t_n}\, g_t^\top \xi + \Delta t_n \Big[\Psi_t(X_t, \xi) - \tfrac{1}{2}\|g_t\|^2\Big] + O\big(\Delta t_n^{3/2}(1 + \|\xi\|^3)\big). \tag{D.7}$$

(b) For the guided step ($\mathbf{1}_{\mathrm{guided}} = 1$) the same holds with the bracket $\Psi_t(X_t, \xi) + \frac{1}{2}\|g_t\|^2$.

(c) For any f that is $C^{1,4}$ with bounded derivatives, the sampler step satisfies

$$\mathbb{E}\big[f_{t'}(Y_{t'}) \mid X_t = x\big] = f_t(x) + \Delta t_n\, (\mathcal{A}_t f_t)(x) + O(\Delta t_n^2). \tag{D.8}$$

(d) At the point $Y_{t'}$ produced by the guided step, the proposal correction is exactly

$$\log \frac{P_{t'|t}(Y_{t'} \mid X_t)}{Q_{t'|t}(Y_{t'} \mid X_t)} = -\sqrt{\Delta t_n}\, g_t^\top \xi - \tfrac{\Delta t_n}{2}\|g_t\|^2. \tag{D.9}$$

Proof. (a) Let $D_n := Y_{t'} - X_t = b_t \Delta t_n + \sigma_t \sqrt{\Delta t_n}\, \xi$, so $\|D_n\| = O(\sqrt{\Delta t_n}(1 + \|\xi\|))$. Taylor expand $\log G_{t'}(X_t + D)$ around (t, X_t) to first order in t and second order in x, with all derivatives evaluated at (t, X_t):

$$\log \frac{G_{t'}(Y_{t'})}{G_t(X_t)} = -\Delta t_n\, \partial_t \log G_t + \nabla \log G_t \cdot D_n + \tfrac{1}{2} D_n^\top \nabla^2 \log G_t\, D_n + O\big(\Delta t_n \|D_n\| + \|D_n\|^3 + \Delta t_n^2\big).$$

Substitute D_n into the linear term,

$$\begin{aligned} \nabla \log G_t \cdot D_n &= \sqrt{\Delta t_n}\, \nabla \log G_t^\top \sigma_t \xi + \Delta t_n\, b_t \cdot \nabla \log G_t \\ &= \sqrt{\Delta t_n}\, g_t^\top \xi + \Delta t_n\, b_t \cdot \nabla \log G_t, \end{aligned}$$

and into the quadratic term,

$$\begin{aligned} \tfrac{1}{2} D_n^\top \nabla^2 \log G_t\, D_n &= \tfrac{1}{2} \Delta t_n\, \xi^\top \sigma_t^\top \nabla^2 \log G_t\, \sigma_t \xi + O\big(\Delta t_n^{3/2}(1 + \|\xi\|)\big) \\ &= \tfrac{1}{2} \Delta t_n\, \xi^\top M_t \xi + O\big(\Delta t_n^{3/2}(1 + \|\xi\|)\big). \end{aligned}$$

Collecting the terms of order Δt_n and using Eq. (D.3) to rewrite $-\partial_t \log G_t + b_t \cdot \nabla \log G_t$ as $\psi_t - \frac{1}{2}\operatorname{tr} M_t - \frac{1}{2}\|g_t\|^2$, the bracket is

$$\begin{aligned} -\partial_t \log G_t + b_t \cdot \nabla \log G_t + \tfrac{1}{2}\xi^\top M_t \xi &= \psi_t + \tfrac{1}{2}\big(\xi^\top M_t \xi - \operatorname{tr} M_t\big) - \tfrac{1}{2}\|g_t\|^2 \\ &= \Psi_t(X_t, \xi) - \tfrac{1}{2}\|g_t\|^2 \end{aligned}$$

by Eq. (D.4). The remainder is $O(\Delta t_n^{3/2}(1+\|\xi\|^3))$ by (A2), since $\|D_n\| = O(\sqrt{\Delta t_n}(1+\|\xi\|))$.

(b) Replace b_t by $b_t + \sigma_t\sigma_t^\top \nabla \log G_t$ in D_n. The extra drift contributes, through the linear term,

$$\Delta t_n \, \nabla \log G_t^\top \sigma_t \sigma_t^\top \nabla \log G_t = \Delta t_n \, \|g_t\|^2,$$

and only $O(\Delta t_n^{3/2})$ terms elsewhere, so the bracket becomes $\Psi_t(X_t,\xi) - \frac{1}{2}\|g_t\|^2 + \|g_t\|^2 = \Psi_t(X_t,\xi) + \frac{1}{2}\|g_t\|^2$.

(c) Apply the expansion of (a) to f in place of $\log G$ and take expectations, using $\mathbb{E}\xi = 0$, $\mathbb{E}[\xi^\top M \xi] = \operatorname{tr} M$ and $\mathbb{E}\|\xi\|^3 < \infty$. The Δt_n term is $-\partial_t f_t + b_t \cdot \nabla f_t + \frac{1}{2}\operatorname{tr}(\sigma_t\sigma_t^\top \nabla^2 f_t) = \mathcal{A}_t f_t$ by Eq. (A.11).

(d) Both kernels are Gaussian with covariance $\Sigma := \sigma_t\sigma_t^\top \Delta t_n$. Their means are $X_t + b_t\Delta t_n$ and $X_t + b_t\Delta t_n + s$ with $s := \Sigma \nabla \log G_t$. The guided step produces $Y_{t'} = X_t + b_t\Delta t_n + s + e$ with $e := \sigma_t\sqrt{\Delta t_n}\,\xi$, so the log-ratio of the two densities at $Y_{t'}$ is

$$\begin{aligned}\log \frac{P_{t'|t}(Y_{t'} \mid X_t)}{Q_{t'|t}(Y_{t'} \mid X_t)} &= -\tfrac{1}{2}(s+e)^\top \Sigma^{-1}(s+e) + \tfrac{1}{2} e^\top \Sigma^{-1} e \\ &= -\tfrac{1}{2} s^\top \Sigma^{-1} s - s^\top \Sigma^{-1} e,\end{aligned}$$

and the two quadratic forms are

$$\begin{aligned} s^\top \Sigma^{-1} s &= \Delta t_n \, \nabla \log G_t^\top \sigma_t \sigma_t^\top \nabla \log G_t = \Delta t_n \, \|g_t\|^2, \\ s^\top \Sigma^{-1} e &= \nabla \log G_t^\top \sigma_t \sqrt{\Delta t_n}\, \xi = \sqrt{\Delta t_n}\, g_t^\top \xi. \end{aligned}$$

There is no remainder. □

The constants in Lem. D.4 depend only on the bounds in (A1)–(A2), which hold on all of $[0,1]\times\mathbb{X}$, so the expansions are uniform in n. The next lemma turns them into the moment bounds that Lem. D.2 needs, again uniformly in n.

Lemma D.5 (Moment bounds)**.** *Under (A1)–(A3) there are constants C, C_0 depending only on the bounds in (A1)–(A2) such that, for every schedule with $\max_n \Delta t_n \le 1$ and every n,*

$$\begin{aligned} &\sup_n \mathbb{E} e^{2|\ell_n|} \le C_0, \\ &m_{1,n} \le C\sqrt{\Delta t_n}, \qquad m_{4,n} \le C\Delta t_n^2 \qquad \textit{(unguided)}, \\ &m_{1,n} \le C\Delta t_n, \qquad m_{4,n} \le C\Delta t_n^4 \qquad \textit{(guided)}. \end{aligned} \tag{D.10}$$

Proof. Exponential moment. $\log G_t$ is Lipschitz in x and in t by (A2), so

$$\left|\log \frac{G_{t'}(Y_{t'})}{G_t(X_t)}\right| \le C\big(\Delta t_n + \|Y_{t'} - X_t\|\big) \le C\sqrt{\Delta t_n}\,(1+\|\xi\|),$$

using $\|Y_{t'} - X_t\| = \|D_n\| \le C\sqrt{\Delta t_n}(1+\|\xi\|)$, and the proposal correction Eq. (D.9) is affine in ξ with coefficients of the same order. Hence $|\ell_n| \le C\sqrt{\Delta t_n}(1+\|\xi\|)$ and, for $\Delta t_n \le 1$,

$$\mathbb{E} e^{2|\ell_n|} \le \mathbb{E} e^{2C(1+\|\xi\|)} =: C_0 < \infty,$$

with C_0 independent of n and of the schedule.

Unguided moments. By Eq. (D.7) without the proposal correction, $\ell_n = \sqrt{\Delta t_n}\, g_t^\top \xi + O(\Delta t_n(1+\|\xi\|^2))$, so $|\ell_n| \le C\sqrt{\Delta t_n}(1+\|\xi\|^2)$ for $\Delta t_n \le 1$. Taking expectations, $m_{1,n} \le C\sqrt{\Delta t_n}\,\mathbb{E}(1+\|\xi\|^2)$ and $m_{4,n} \le C\Delta t_n^2\,\mathbb{E}(1+\|\xi\|^2)^4$, and the Gaussian moments are finite constants.

Guided moments. By the proof of Thm. 3.2(b) below, which uses only Lem. D.4(b) and (d), $\ell_n = \Delta t_n \Psi_t(X_t,\xi) + O(\Delta t_n^{3/2}(1+\|\xi\|^3))$ with Ψ_t bounded by a quadratic in $\|\xi\|$, so $|\ell_n| \le C\Delta t_n(1+\|\xi\|^3)$ for $\Delta t_n \le 1$, and $m_{1,n} \le C\Delta t_n$, $m_{4,n} \le C\Delta t_n^4$ follow as before.

All remainders in Lem. D.4 are of the form $O(\Delta t_n^a(1+\|\xi\|^b))$ with ξ Gaussian, so they may be integrated term by term and contribute $O(\Delta t_n^a)$ to $m_{1,n}$ and to $\mathbb{E}|\ell_n - \ell'_n|$, and $O(\Delta t_n^{4a})$ to $m_{4,n}$, with constants independent of n. □

Lemma D.6 (Continuity of the local barriers). *Under (A1)–(A2), $t \mapsto \lambda^{\mathrm{u}}(t)$ and $t \mapsto \lambda(t)$ are continuous on $[0,1]$. If moreover $t \mapsto p_t$ has a density that is continuously differentiable in t, as holds for the uniformly elliptic diffusion Eq. (A.10) with the smooth bounded coefficients of (A1)–(A2) by parabolic regularity, then both are Lipschitz.*

Proof. We treat λ. The argument for λ^{u} is the same with a simpler integrand.

Step 1: write λ as an expectation of a fixed function. Let $\nu_t := \pi_t \otimes \pi_t \otimes N(0,I) \otimes N(0,I)$ be the joint law of (X_t, X_t', ξ, ξ') and

$$F(t,x,x',\xi,\xi') := \tfrac{1}{2}\big|\Psi_t(x,\xi) - \Psi_t(x',\xi')\big|, \qquad \text{so that} \qquad \lambda(t) = \int F(t,\cdot)\,\mathrm{d}\nu_t.$$

By Eq. (D.4) and the bounds of (A1)–(A2) on ψ_t and M_t,

$$0 \le F(t,x,x',\xi,\xi') \le C\big(1 + \|\xi\|^2 + \|\xi'\|^2\big) \tag{D.11}$$

uniformly in (t,x,x'), and F is continuous in (t,x,x') for each (ξ,ξ'), because $\psi_t(x)$ and $M_t(x)$ are continuous in (t,x).

Step 2: continuity in the law. Fix t_0 and let $t \to t_0$. Split

$$\lambda(t) - \lambda(t_0) = \underbrace{\int \big(F(t,\cdot) - F(t_0,\cdot)\big)\,\mathrm{d}\nu_t}_{=:I_1(t)} + \underbrace{\int F(t_0,\cdot)\,\mathrm{d}\nu_t - \int F(t_0,\cdot)\,\mathrm{d}\nu_{t_0}}_{=:I_2(t)}.$$

For I_1, the integrand converges to zero pointwise as $t \to t_0$ by continuity of F in t, and it is dominated by $2C(1 + \|\xi\|^2 + \|\xi'\|^2)$ by Eq. (D.11), whose ν_t-integral is a constant independent of t. Hence $I_1(t) \to 0$ by dominated convergence.

For I_2, the function $F(t_0,\cdot)$ is fixed, continuous, and bounded by a function of (ξ,ξ') alone. Integrate out (ξ,ξ') first: $\bar F(x,x') := \int F(t_0,x,x',\xi,\xi')\,N(\mathrm{d}\xi)N(\mathrm{d}\xi')$ is bounded and continuous in (x,x') by dominated convergence, and $I_2(t) = \int \bar F\,\mathrm{d}(\pi_t \otimes \pi_t) - \int \bar F\,\mathrm{d}(\pi_{t_0} \otimes \pi_{t_0})$. So it suffices that $t \mapsto \pi_t$ is weakly continuous. The law p_t is the marginal at time t of a diffusion with continuous paths, so $X_t \to X_{t_0}$ almost surely as $t \to t_0$ and $p_t \to p_{t_0}$ weakly. Since $G_t(x)$ is bounded, positive and continuous in (t,x), for every bounded continuous h,

$$\int h\,\mathrm{d}\pi_t = \frac{\int h\,G_t\,\mathrm{d}p_t}{\int G_t\,\mathrm{d}p_t} \longrightarrow \frac{\int h\,G_{t_0}\,\mathrm{d}p_{t_0}}{\int G_{t_0}\,\mathrm{d}p_{t_0}} = \int h\,\mathrm{d}\pi_{t_0},$$

the numerator and denominator converging because $hG_t \to hG_{t_0}$ uniformly on compacts and $p_t \to p_{t_0}$ weakly, with tightness of (p_t) supplied by path continuity. Hence $\pi_t \to \pi_{t_0}$ weakly, $I_2(t) \to 0$, and λ is continuous.

Step 3: Lipschitz regularity. Suppose p_t has a density $p_t(x)$ that is C^1 in t, with $\partial_t p_t$ integrable uniformly in t. Then $\pi_t(x) = G_t(x)p_t(x)/\mathcal{Z}_t$ is C^1 in t, and for the bounded function $\bar F$ of Step 2 the map $t \mapsto \int \bar F\,\mathrm{d}(\pi_t \otimes \pi_t)$ is C^1 by differentiation under the integral. Moreover F is Lipschitz in t uniformly in (x,x') and with a constant proportional to $1 + \|\xi\|^2 + \|\xi'\|^2$, since $\partial_t\psi_t$ and $\partial_t M_t$ are bounded by (A2). Writing $\lambda(t) - \lambda(s) = I_1 + I_2$ as in Step 2 with t_0 replaced by s, both terms are bounded by $C|t-s|$, so λ is Lipschitz. □

D.3 Proof of Thm. 3.2

The continuity of λ^{u} and λ asserted in the theorem is Lem. D.6. The two expansions are proved next, with all constants uniform in n by Lem. D.5.

Proof of (a): the unguided proposal. With $\mathbf{1}_{\mathrm{guided}} = 0$ the proposal correction is absent, so by Eq. (D.7) at $t = t_n$,

$$\ell_n = \sqrt{\Delta t_n}\, g_{t_n}(X_{t_n})^\top \xi + O\big(\Delta t_n(1 + \|\xi\|^2)\big).$$

The leading term is a mean-zero Gaussian fluctuation of order $\sqrt{\Delta t_n}$, and it survives the mean absolute difference: by Lem. D.5, $m_{1,n} = O(\sqrt{\Delta t_n})$ and $m_{4,n}^{1/2} = O(\Delta t_n)$ uniformly in n, so

Lem. D.2 gives $r_{n-1,n} = \frac{1}{2}\mathbb{E}|\ell_n - \ell'_n| + O(\Delta t_n)$. The $O(\Delta t_n)$ terms of $\ell_n - \ell'_n$ contribute $O(\Delta t_n)$ to the mean absolute difference by the triangle inequality, and Lem. D.3 evaluates the leading term:

$$\begin{aligned} r_{n-1,n} &= \tfrac{\sqrt{\Delta t_n}}{2}\,\mathbb{E}\big|g_{t_n}(X_{t_n})^\top \xi - g_{t_n}(X'_{t_n})^\top \xi'\big| + O(\Delta t_n) \\ &= \tfrac{\sqrt{\Delta t_n}}{2}\sqrt{\tfrac{2}{\pi}}\,\mathbb{E}\sqrt{\|g_{t_n}(X_{t_n})\|^2 + \|g_{t_n}(X'_{t_n})\|^2} + O(\Delta t_n) = \lambda^{\mathrm{u}}(t_n)\sqrt{\Delta t_n} + O(\Delta t_n), \end{aligned}$$

with λ^{u} as in Eq. (3.11), and the $O(\Delta t_n)$ term is bounded by $C\Delta t_n$ with C independent of n.

For the sum, λ^{u} is continuous in t by Lem. D.6, and it is not identically zero unless $\nabla \log G_t \equiv 0$ for every t, in which case there is no control. If $\lambda^{\mathrm{u}} \geq c > 0$ on a set of times of length $L > 0$, then on any schedule with $\max_n \Delta t_n \to 0$,

$$\sum_n r_{n-1,n} \;\geq\; c \sum_{t_n \in \text{that set}} \sqrt{\Delta t_n} - O(1) \;\geq\; c \sum_{t_n \in \text{that set}} \frac{\Delta t_n}{\sqrt{\max_n \Delta t_n}} - O(1) \;\geq\; \frac{cL}{\sqrt{\max_n \Delta t_n}} - O(1) \;\to\; \infty.$$

On a uniform schedule, $\Delta t_n = 1/N$, the Riemann sum gives

$$\sum_{n=1}^{N} r_{n-1,n} = \sqrt{N} \sum_{n=1}^{N} \lambda^{\mathrm{u}}(t_n)\,\Delta t_n + O(1) = \sqrt{N} \int_0^1 \lambda^{\mathrm{u}}(t)\,\mathrm{d}t + O(1) = \Theta(\sqrt{N}).$$

Since $r/(1-r) \geq r$, the restart rate Eq. (2.6) tends to zero. □

Proof of (b): the guided proposal. By Eq. (D.7) in the guided form (b) and by Eq. (D.9), both at $t = t_n$, the two terms of Eq. (D.1) are

$$\begin{aligned} \log \frac{G_{t'}(Y_{t'})}{G_t(X_t)} &= \sqrt{\Delta t_n}\, g_t^\top \xi + \Delta t_n \Big[\Psi_t(X_t, \xi) + \tfrac{1}{2}\|g_t\|^2\Big] + O\big(\Delta t_n^{3/2}(1 + \|\xi\|^3)\big), \\ \log \frac{P_{t'|t}(Y_{t'} \mid X_t)}{Q_{t'|t}(Y_{t'} \mid X_t)} &= -\sqrt{\Delta t_n}\, g_t^\top \xi - \tfrac{\Delta t_n}{2}\|g_t\|^2. \end{aligned}$$

Adding them, the $\sqrt{\Delta t_n}$ terms and the $\|g_t\|^2$ terms cancel exactly, leaving

$$\ell_n = \Delta t_n\, \Psi_{t_n}(X_{t_n}, \xi) + O\big(\Delta t_n^{3/2}(1 + \|\xi\|^3)\big),$$

which is the guided case of Eq. (3.9): guidance moves the proposal to where the potential is steep, and the correction it pays for doing so removes exactly the fluctuation that the steepness would otherwise cause. Now $m_{1,n} = O(\Delta t_n)$ and $m_{4,n}^{1/2} = O(\Delta t_n^2)$ uniformly in n by Lem. D.5, so Lem. D.2 with the triangle inequality on the remainder gives

$$r_{n-1,n} = \tfrac{\Delta t_n}{2}\,\mathbb{E}\big|\Psi_{t_n}(X_{t_n}, \xi) - \Psi_{t_n}(X'_{t_n}, \xi')\big| + R_n = \lambda(t_n)\,\Delta t_n + R_n, \qquad |R_n| \leq C\,\Delta t_n^{3/2}, \tag{D.12}$$

with C independent of n, by Lem. D.5 and the uniformity of the remainders in Lem. D.4. The error is therefore uniform in n along any schedule.

Integral form. By Lem. D.6, λ is continuous, so $\lambda(t_n)\Delta t_n = \int_{t_{n-1}}^{t_n} \lambda(t)\,\mathrm{d}t + o(\Delta t_n)$ uniformly in n by uniform continuity on $[0,1]$, and $O(\Delta t_n^2)$ when λ is Lipschitz. Hence Eq. (D.12) can be written as

$$r_{n-1,n} = \int_{t_{n-1}}^{t_n} \lambda(t)\,\mathrm{d}t + o(\Delta t_n), \tag{D.13}$$

uniformly in n, which is the statement of Thm. 3.2(b) and the form used to construct the equi-rejection schedule in App. D.5. The $o(\Delta t_n)$ needs only continuity of λ, hence only (A1)–(A3). It improves to $O(\Delta t_n^{3/2})$ under the Lipschitz regularity of Lem. D.6, and to $O(\Delta t_n^2)$ under the additional non-degeneracy condition of Prop. D.7 below.

The sum. Summing Eq. (D.13) over the schedule,

$$\sum_{n=1}^{N} r_{n-1,n} = \int_0^1 \lambda(t)\,\mathrm{d}t + \sum_{n=1}^{N} o(\Delta t_n) \;\to\; \Lambda < \infty,$$

since the error is $o(1)$ by uniformity of the $o(\Delta t_n)$ terms in n and $\sum_n \Delta t_n = 1$. Since $r_{n-1,n} \to 0$ uniformly in n, the odds satisfy $r/(1-r) = r + O(r^2)$ and $\sum_n r_{n-1,n}^2 = O(\max_n \Delta t_n) \to 0$, so the summed odds have the same limit and the restart rate Eq. (2.6) converges to $(2 + 2\Lambda)^{-1}$.

It remains to characterise $\Lambda = 0$, which is Lem. D.9 below. □

Proposition D.7 (Second-order accuracy)**.** *Assume (A1)–(A3), that λ is Lipschitz (Lem. D.6), and that for every t the difference $Y_t := \Psi_t(X_t, \xi) - \Psi_t(X'_t, \xi')$ of two independent copies has a density that is bounded near zero, uniformly in t. Then, for the guided proposal,*

$$r_{n-1,n} = \int_{t_{n-1}}^{t_n} \lambda(t)\,\mathrm{d}t + O(\Delta t_n^2) \qquad \textit{uniformly in } n. \tag{D.14}$$

Proof. The Lipschitz regularity of λ gives $\lambda(t_n)\Delta t_n = \int_{t_{n-1}}^{t_n} \lambda + O(\Delta t_n^2)$, so it suffices to show $r_{n-1,n} = \lambda(t_n)\Delta t_n + O(\Delta t_n^2)$. By Lem. D.2 and Lem. D.5 the error in $r_{n-1,n} = \frac{1}{2}\mathbb{E}|\ell_n - \ell'_n|$ is already $O(\Delta t_n^2)$ in the guided case, so the question is the remainder in ℓ_n itself.

The remainder is odd in the noise. Carrying the expansion of Lem. D.4 one order further, with all derivatives at (t, X_t),

$$\ell_n = \Delta t_n\, \Psi_t(X_t, \xi) + \Delta t_n^{3/2}\, R_t(X_t, \xi) + O\big(\Delta t_n^2 (1 + \|\xi\|^4)\big),$$

where R_t collects the terms of order $\Delta t_n^{3/2}$: the cross term $b_t^\top \nabla^2 \log G_t\, \sigma_t \xi$ from the quadratic part, the term $-\partial_t \nabla \log G_t^\top \sigma_t \xi$ from the mixed derivative, and the cubic term $\frac{1}{6}\nabla^3 \log G_t[\sigma_t\xi, \sigma_t\xi, \sigma_t\xi]$, together with the corresponding terms from the guided drift. Every term in R_t has an odd power of ξ, so $R_t(x, -\xi) = -R_t(x, \xi)$, whereas $\Psi_t(x, -\xi) = \Psi_t(x, \xi)$ by Eq. (D.4) and the $O(\Delta t_n^2)$ terms are even in ξ as well.

Parity. Let $Y := \Psi_t(X_t, \xi) - \Psi_t(X'_t, \xi')$ and $Z := R_t(X_t, \xi) - R_t(X'_t, \xi')$, and consider $f(\varepsilon) := \mathbb{E}|Y + \varepsilon Z|$. The map $(\xi, \xi') \mapsto (-\xi, -\xi')$ preserves the law of (X_t, X'_t, ξ, ξ'), fixes Y and negates Z, so $f(\varepsilon) = f(-\varepsilon)$. Under the density assumption on Y, f is twice continuously differentiable near 0 with $f'(0) = \mathbb{E}[\mathrm{sign}(Y)Z]$ and $f''(0) = \mathbb{E}[Z^2 \mid Y = 0]\, p_Y(0)$, both finite since Z has all moments by (A2). Evenness forces $f'(0) = 0$, hence $f(\varepsilon) = f(0) + O(\varepsilon^2)$ with a constant uniform in t. With $\varepsilon = \sqrt{\Delta t_n}$,

$$\tfrac{1}{2}\mathbb{E}|\ell_n - \ell'_n| = \tfrac{\Delta t_n}{2}\, f\big(\sqrt{\Delta t_n}\big) + O(\Delta t_n^2) = \tfrac{\Delta t_n}{2}\, \mathbb{E}|Y| + O(\Delta t_n^2) = \lambda(t_n)\, \Delta t_n + O(\Delta t_n^2),$$

the first equality absorbing the even $O(\Delta t_n^2(1 + \|\xi\|^4))$ remainder by the triangle inequality and the Gaussian moments. The constants depend only on the bounds in (A1)–(A2) and on the density bound, hence not on n. □

Without the density assumption the parity argument still gives $f(\varepsilon) - f(0) \le \varepsilon \mathbb{E}|Z|$, which is the $O(\Delta t_n^{3/2})$ of Eq. (D.12). The assumption holds, for instance, whenever $v(t)$ is bounded away from zero and $\Psi_t(X_t, \xi)$ has a bounded density, as it does when $M_t \not\equiv 0$ or $\psi_t(X_t)$ is non-degenerate.

Lemma D.8 (The Doob residual)**.** *Under (A1)–(A2), the following are equivalent for every t: (i) $\psi_t = \mathcal{A}_t G_t / G_t$ is constant in x, (ii) G_t is an eigenfunction of $\mathcal{A}_t$, $\mathcal{A}_t G_t = c(t) G_t$ for some $c(t)$, and (iii) $G_t \propto G_t^*$, where $G_t^*(x) = \mathbb{E}[G_0(X_0) \mid X_t = x]$ is the eigenfunction with eigenvalue zero.*

Proof. The first two statements are the same statement, since $\psi_t = \mathcal{A}_t G_t / G_t$. Suppose $\mathcal{A}_t G_t = c(t) G_t$. Set $\phi(t) := \int_0^t c(s)\,\mathrm{d}s$ and $\widetilde{G}_t := e^{\phi(t)} G_t$. Since $\mathcal{A}_t$ acts on the time dependence through $-\partial_t$ only,

$$\mathcal{A}_t \widetilde{G}_t = -\partial_t\big(e^{\phi(t)}\big) G_t + e^{\phi(t)} \mathcal{A}_t G_t = -c(t) e^{\phi(t)} G_t + e^{\phi(t)} c(t) G_t = 0,$$

and $\widetilde{G}_0 = G_0$. So $\widetilde{G}$ solves the backward equation $\mathcal{A}_t \widetilde{G}_t = 0$ with terminal condition G_0, whose unique solution with bounded logarithm is G_t^*: existence is the Feynman–Kac representation, and uniqueness holds because the difference of two bounded solutions is a bounded space-time harmonic function of the diffusion, hence constant along it and zero at $t = 0$ (Del Moral, 2004, Ch. 1). Hence $G_t = e^{-\phi(t)} G_t^* \propto G_t^*$.

Conversely, if $G_t = e^{-\phi(t)} G_t^*$ for a differentiable ϕ, then

$$\mathcal{A}_t G_t = -\partial_t\big(e^{-\phi(t)}\big) G_t^* + e^{-\phi(t)} \mathcal{A}_t G_t^* = \phi'(t)\, e^{-\phi(t)} G_t^* + 0 = \phi'(t)\, G_t,$$

so $\psi_t = \phi'(t)$ is constant in x. □

Lemma D.9 (Zero barrier). *Under (A1)–(A2), $\lambda(t) = 0$ if and only if $\log G_t$ is affine in x and $G_t \propto G_t^*$. Hence $\Lambda = 0$ if and only if both hold for every t, i.e. if and only if the guided proposal is the Doob proposal.*

Proof. $\lambda(t) = \frac{1}{2}\mathbb{E}|\Psi_t(X_t,\xi) - \Psi_t(X_t',\xi')|$ vanishes if and only if two independent copies of $\Psi_t(X_t,\xi)$ agree almost surely, i.e. if and only if $\Psi_t(X_t,\xi)$ is almost surely constant.

The curvature term must vanish. Condition on $X_t = x$. By Eq. (D.4),

$$\Psi_t(x,\xi) = \psi_t(x) + \tfrac{1}{2}\big(\xi^\top M_t(x)\xi - \operatorname{tr} M_t(x)\big)$$

is almost surely constant in $\xi \sim N(0,I)$, so the quadratic form $\xi \mapsto \xi^\top M_t(x)\xi$ is almost surely constant. A polynomial that is constant on a set of positive Lebesgue measure is constant everywhere, and a quadratic form is constant only if its matrix is zero, so $M_t(x) = 0$ for π_t-almost every x. By continuity of M_t in x and positivity of π_t, $M_t \equiv 0$. By (A1) σ_t is invertible, so $\nabla^2 \log G_t \equiv 0$ and $\log G_t$ is affine in x.

The residual must be constant. With $M_t \equiv 0$, $\Psi_t(x,\xi) = \psi_t(x)$, which is almost surely constant under π_t, hence everywhere by continuity and positivity. By Lem. D.8, $G_t \propto G_t^*$.

Conversely, if $\log G_t$ is affine then $M_t \equiv 0$, and if $G_t \propto G_t^*$ then ψ_t is constant in x by Lem. D.8. So Ψ_t is constant and $\lambda(t) = 0$. Since $\lambda \geq 0$ is continuous, $\Lambda = 0$ if and only if $\lambda \equiv 0$. The final equivalence is the affine case of App. A.2: for an affine log-potential the guided step is the Doob proposal of G_t by Eq. (A.14), and $G_t \propto G_t^*$ makes that the Doob proposal of the target. □

D.4 The Guided Proposal Against the Doob Proposal

For the chosen potentials G_t, the Doob proposal is the twisted kernel

$$Q^{\mathrm{tw}}_{t'|t}(\mathrm{d}y \mid x) := \frac{G_{t'}(y)\, P_{t'|t}(\mathrm{d}y \mid x)}{(P_{t'|t}G_{t'})(x)}, \tag{D.15}$$

which satisfies the proposal relation of Eq. (A.9) exactly and, by Thm. 3.1(c), minimises the rejection over all proposals for those potentials. The guided step approximates it. The next proposition computes its barrier and compares it with the guided one.

Proposition D.10 (Twisted proposal). *Under the assumptions of Thm. 3.2, the twisted proposal has $r_{n-1,n} = \lambda^{\mathrm{tw}}(t_n)\Delta t_n + O(\Delta t_n^2)$ with*

$$\lambda^{\mathrm{tw}}(t) := \tfrac{1}{2}\,\mathbb{E}\big|\psi_t(X_t) - \psi_t(X_t')\big| \;\leq\; \lambda(t), \tag{D.16}$$

with equality if and only if $\log G_t$ is affine in x, and $\Lambda^{\mathrm{tw}} := \int_0^1 \lambda^{\mathrm{tw}}(t)\,\mathrm{d}t = 0$ if and only if $G_t \propto G_t^*$ for all t.

Proof. The proposal correction of the twisted kernel cancels the potential ratio in y, and the weight depends on X_t alone:

$$w_{t'|t}(X_t,\xi) = \frac{G_{t'}(Y_{t'})}{G_t(X_t)} \cdot \frac{P_{t'|t}(Y_{t'} \mid X_t)}{Q^{\mathrm{tw}}_{t'|t}(Y_{t'} \mid X_t)} = \frac{G_{t'}(Y_{t'})}{G_t(X_t)} \cdot \frac{(P_{t'|t}G_{t'})(X_t)}{G_{t'}(Y_{t'})} = \frac{(P_{t'|t}G_{t'})(X_t)}{G_t(X_t)}.$$

The numerator is the one-step expectation of $G_{t'}$, so by Eq. (D.8) applied to $f = G$,

$$\begin{aligned}\ell_n &= \log \frac{G_t(X_t) + \Delta t_n\,(\mathcal{A}_t G_t)(X_t) + O(\Delta t_n^2)}{G_t(X_t)} \\ &= \log\big(1 + \Delta t_n\,\psi_t(X_t) + O(\Delta t_n^2)\big) = \Delta t_n\,\psi_t(X_t) + O(\Delta t_n^2),\end{aligned}$$

uniformly in X_t and in n. Hence $m_{1,n} = O(\Delta t_n)$, $m_{4,n}^{1/2} = O(\Delta t_n^2)$, and Lem. D.2 gives $r_{n-1,n} = \lambda^{\mathrm{tw}}(t_n)\Delta t_n + O(\Delta t_n^2)$.

For the comparison, $\mathbb{E}[\Psi_t(X_t,\xi) \mid X_t] = \psi_t(X_t)$ by Eq. (D.4), and conditional Jensen applied to the convex function $|\cdot|$ gives

$$\lambda(t) = \tfrac{1}{2}\,\mathbb{E}\big|\Psi_t(X_t,\xi) - \Psi_t(X_t',\xi')\big|$$

$$\geq \tfrac{1}{2}\,\mathbb{E}\Big|\,\mathbb{E}\big[\Psi_t(X_t,\xi)-\Psi_t(X_t',\xi')\;\big|\;X_t,X_t'\big]\Big|$$
$$=\tfrac{1}{2}\,\mathbb{E}\big|\psi_t(X_t)-\psi_t(X_t')\big|=\lambda^{\mathrm{tw}}(t),$$

with equality if and only if the curvature term $\frac{1}{2}(\xi^\top M_t\xi-\operatorname{tr} M_t)$ vanishes almost surely, i.e. $M_t\equiv 0$, i.e. $\log G_t$ affine, as in Lem. D.9. Finally $\lambda^{\mathrm{tw}}(t)=0$ if and only if ψ_t is constant in x, which is $G_t\propto G_t^*$ by Lem. D.8. □

The curvature term. When $\log G_{t'}$ is affine in x the guided step is the twisted kernel exactly, by Eq. (A.14). When it is not, the two kernels differ at second order. Write the sampler step as $P_{t'|t}(\cdot\mid x)=N(\mu_n(x),\Sigma_n)$ with $\mu_n(x)=x+b_t(x)\Delta t_n$ and $\Sigma_n=\sigma_t\sigma_t^\top\Delta t_n$, and expand $\log G_{t'}(y)\approx c_n^\top y+\frac{1}{2}y^\top H_n y$ with $H_n=\nabla^2\log G_{t'}$. Completing the square gives

$$Q^{\mathrm{tw}}_{t'|t}(\cdot\mid x)\approx N\Big((\Sigma_n^{-1}-H_n)^{-1}\big(\Sigma_n^{-1}\mu_n(x)+c_n\big),\;(\Sigma_n^{-1}-H_n)^{-1}\Big),$$

so the twisted kernel has covariance $(\Sigma_n^{-1}-H_n)^{-1}$ rather than Σ_n: narrower where $\log G_{t'}$ is concave ($H_n\prec 0$), wider where it is convex, with the mean shifted accordingly. The guided step keeps the covariance Σ_n and shifts only the mean, and the difference between the two is what produces the curvature term $\frac{1}{2}(\xi^\top M_t\xi-\operatorname{tr} M_t)$ in Ψ_t of Eq. (3.10). That term has mean zero given X_t, so no choice of the potentials removes it. It is the cost of a Gaussian proposal with the sampler's covariance, and it is what separates Λ from Λ^{tw}.

D.5 Tuning the Schedule

Under the guided proposal, Eq. (D.13) gives $r_{n-1,n}=\int_{t_{n-1}}^{t_n}\lambda(t)\,\mathrm{d}t+o(\Delta t_n)$ uniformly in n, so the summed rejection rates converge to $\Lambda=\int_0^1\lambda$ and a schedule controls only how Λ is distributed over the levels. By Eq. (2.6) the restart rate is decreasing in $\sum_n r_{n-1,n}/(1-r_{n-1,n})$, and $s\mapsto s/(1-s)$ is convex on $[0,1)$, so by Jensen's inequality, among schedules with a fixed number of levels and a fixed sum $\sum_n r_{n-1,n}\approx\Lambda$, the sum of odds is minimised when the per-level rejection rates are equal:

$$r_{n-1,n}=\int_{t_{n-1}}^{t_n}\lambda(t)\,\mathrm{d}t=\frac{\Lambda}{N}\qquad\text{for every } n. \tag{D.17}$$

This is the equi-rejection schedule of Syed et al. (2022) applied to the local barrier of Thm. 3.2, and their analysis applies: $N\approx 2\Lambda$ levels give stable communication, after which a fresh reference draw reaches the target every $2+2\Lambda$ iterations whatever N, and a longer run rather than a finer ladder increases the sample size. In practice λ is estimated from the rejection rates observed along a pilot run and the schedule is refined by inverting the resulting empirical barrier, as there. Under the unguided proposal no such limit exists and the spacing is instead set by Prop. D.12.

D.6 Optimal Spacing for the Unguided Ladder

Lemma D.11 (Log-normal weights). *If* $\log W_{n-1|n}\sim N(\mu_n,s_n^2)$ *under* $\pi_n\otimes Q_{n-1|n}$, *then*

$$r_{n-1,n}=\operatorname{erf}(s_n/2)=\frac{s_n}{\sqrt{\pi}}+O(s_n^3). \tag{D.18}$$

Proof. Write $W_{n-1|n}=e^A$, $W'_{n-1|n}=e^B$ with A,B i.i.d. $N(\mu_n,s_n^2)$. By symmetry,

$$\mathbb{E}|W_{n-1|n}-W'_{n-1|n}|=2\,\mathbb{E}\big[(e^A-e^B)\mathbf{1}_{A>B}\big]=2\big(\mathbb{E}[e^A\mathbf{1}_{A>B}]-\mathbb{E}[e^B\mathbf{1}_{A>B}]\big).$$

Under the tilted measure $\mathrm{d}\mathbb{Q}=e^{A-\mu_n-s_n^2/2}\mathrm{d}\mathbb{P}$ one has $A\sim N(\mu_n+s_n^2,s_n^2)$ and B unchanged, so $A-B\sim N(s_n^2,2s_n^2)$ under $\mathbb{Q}$ and

$$\mathbb{E}\big[e^A\mathbf{1}_{A>B}\big]=e^{\mu_n+s_n^2/2}\,\mathbb{Q}(A>B)=e^{\mu_n+s_n^2/2}\,\Phi\big(s_n/\sqrt{2}\big),$$

with Φ the standard normal distribution function. The same tilt applied to B gives $\mathbb{E}[e^B\mathbf{1}_{A>B}]=e^{\mu_n+s_n^2/2}\,\Phi(-s_n/\sqrt{2})$. Hence

$$\mathbb{E}|W_{n-1|n}-W'_{n-1|n}|=2e^{\mu_n+s_n^2/2}\big(2\Phi(s_n/\sqrt{2})-1\big),\qquad \mathbb{E}[W_{n-1|n}]=e^{\mu_n+s_n^2/2},$$

and $2\Phi(x)-1=\operatorname{erf}(x/\sqrt{2})$ gives $r_{n-1,n}=\mathbb{E}|W_{n-1|n}-W'_{n-1|n}|/(2\mathbb{E}W_{n-1|n})=\operatorname{erf}(s_n/2)$ by Eq. (D.2). □

The proof of Thm. 3.2(a) shows that with the unguided proposal, refining a time stretch always increases the summed rejection rate: if the log-weight increments over disjoint intervals are approximately Gaussian and independent with variances v_i, then

$$\sum_i \operatorname{erf}\left(\sqrt{v_i}/2\right) \;\geq\; \operatorname{erf}\left(\sqrt{\textstyle\sum_i v_i}\,/2\right),$$

because $v \mapsto \operatorname{erf}(\sqrt{v}/2)$ is concave with value 0 at 0, hence subadditive. Merging levels is limited by the cost $\sum_i r_i/(1-r_i)$ in Eq. (2.6), which diverges as $r_i \to 1$. The balance is the following.

Proposition D.12 (Optimal unguided spacing)**.** *Let $\rho(u) := \operatorname{erf}(u)/\operatorname{erfc}(u)$ and $f(v) := \rho(\sqrt{v}/2)$. For any partition of a time stretch into steps with Gaussian log-weight variances v_i summing to S,*

$$\sum_i \frac{r_i}{1-r_i} = \sum_i f(v_i) \;\geq\; S \inf_{v>0} \frac{f(v)}{v} \;\approx\; 1.06\, S, \tag{D.19}$$

with equality up to integrality when all $v_i = v^* \approx 1.5$, i.e. log-weight standard deviation ≈ 1.2 per step and per-exchange rejection $r^* = \operatorname{erf}(0.61) \approx 0.61$.

Proof. By Lem. D.11, $r_i = \operatorname{erf}(\sqrt{v_i}/2)$, so $r_i/(1-r_i) = f(v_i)$. Then

$$\sum_i f(v_i) = \sum_i v_i \frac{f(v_i)}{v_i} \;\geq\; \Big(\sum_i v_i\Big) \inf_{v>0} \frac{f(v)}{v} = S \inf_{v>0} \frac{f(v)}{v},$$

with equality if and only if every v_i attains the infimum. Substituting $v = 4u^2$, $f(v)/v = \rho(u)/(4u^2)$ tends to $+\infty$ as $u \to 0$ (since $\rho(u) \sim 2u/\sqrt{\pi}$) and as $u \to \infty$ (since $\operatorname{erfc}(u) \to 0$), and a numerical minimisation gives $u^* \approx 0.61$ with value ≈ 1.06. □

Under Eq. (2.6) this is the spacing that maximises the restart rate across an unguided time stretch, and the effective cost of the stretch is its total log-weight variance S: consecutive levels should differ by about 1.2 in log-weight standard deviation, and the ladder should be coarsened wherever $\nabla \log G_t$ is small. Where $G_t \equiv 1$ the log-weight is identically zero and a single level suffices.

D.7 The SMC Barrier: FKD and TDS

The backward weights of Eq. (3.2) also drive the sequential control samplers FKD and TDS, and the expansions above determine how those samplers behave under refinement. Write $\ell_n := \log w_{t_{n-1}|t_n}$ for the log-weight of a particle at step n of the schedule, with variance taken under the proposal law $\pi_n \otimes Q_{n-1|n}$, and recall from Eq. (A.21) that the relative variance of the normalising-constant estimator is governed by the total discrepancy $\mathcal{D}_{0,N} = \sum_n \mathcal{D}_{n-1,n}$ of Eq. (A.20), which in the small-weight regime is the summed log-weight variance $\sum_n \operatorname{Var}(\ell_n)$ up to $O(\sum_n \mathbb{E}|\ell_n|^3)$, a term that vanishes with refinement in both cases below.

Unguided proposal (FKD). By the proof of Thm. 3.2(a), $\ell_n = \sqrt{\Delta t_n}\, g_{t_n}(X_{t_n})^\top \xi + O(\Delta t_n)$ with ξ independent of X_{t_n}. Conditionally on X_{t_n} the leading term has variance $\Delta t_n \|g_{t_n}(X_{t_n})\|^2$ and mean zero, so

$$\operatorname{Var}(\ell_n) = \mathbb{E}_{\pi_{t_n}}\left[\|g_{t_n}(X_{t_n})\|^2\right] \Delta t_n + O(\Delta t_n^{3/2}),$$

and summing the Riemann sum,

$$\mathcal{D}_{0,N} \xrightarrow[N\to\infty]{} \int_0^1 \mathbb{E}_{\pi_t}\left[\|\sigma_t^\top \nabla \log G_t(X_t)\|^2\right] \mathrm{d}t. \tag{D.20}$$

The limit is positive whenever the potential path is non-trivial and does not depend on the schedule: refining the sampler does not reduce the weight variance of the unguided sampler, and a stronger tilt increases it quadratically. This is why FKD collapses at small batch sizes irrespective of the number of steps.

Guided proposal (TDS). By the proof of Thm. 3.2(b), $\ell_n = \Delta t_n\, \Psi_{t_n}(X_{t_n}, \xi) + O(\Delta t_n^{3/2})$, so $\mathrm{Var}(\ell_n) = \Delta t_n^2\, v(t_n) + O(\Delta t_n^{5/2})$ with

$$v(t) := \mathrm{Var}\left(\Psi_t(X_t, \xi)\right) = \mathrm{Var}_{\pi_t}\left(\psi_t(X_t)\right) + \tfrac{1}{2}\,\mathbb{E}_{\pi_t}\|M_t(X_t)\|_F^2, \tag{D.21}$$

by the law of total variance: the first term is the variance of the conditional mean ψ_t, the second is the mean of the conditional variance of the curvature term, which is $\frac{1}{2}\|M_t\|_F^2$ since $\mathrm{Var}(\xi^\top M \xi) = 2\|M\|_F^2$ for symmetric M and $\xi \sim N(0, I)$. The twisted proposal drops the second term. Hence

$$\mathcal{D}_{0,N} = \sum_{n=1}^{N} \Delta t_n^2\, v(t_n) + o\big(\max_n \Delta t_n\big), \tag{D.22}$$

which is of order $1/N$ for any schedule with $\max_n \Delta t_n = O(1/N)$. The schedule that minimises it for given N equalises the per-step standard deviations $\Delta t_n \sqrt{v(t_n)}$: by Cauchy–Schwarz, with $\sum_n \Delta t_n = 1$,

$$\Big(\sum_{n=1}^{N} \Delta t_n \sqrt{v(t_n)}\Big)^2 \;\le\; N \sum_{n=1}^{N} \Delta t_n^2\, v(t_n),$$

with equality if and only if all $\Delta t_n \sqrt{v(t_n)}$ are equal, and the left side is a Riemann sum for Λ_{SMC}^2. Therefore

$$\min_{\text{schedules}} \mathcal{D}_{0,N} = \frac{\Lambda_{\mathrm{SMC}}^2}{N} + o(1/N), \qquad \Lambda_{\mathrm{SMC}} := \int_0^1 \sqrt{v(t)}\,\mathrm{d}t, \tag{D.23}$$

attained when $\Delta t_n \sqrt{v(t_n)} = \Lambda_{\mathrm{SMC}}/N$ for every n. Λ_{SMC} is the global barrier of annealed SMC in the sense of Syed et al. (2026), with Ψ_t as the rate of change of the log incremental weight, and it is finite under the assumptions of Thm. 3.2. By Eq. (A.21) the relative variance of $\widehat{Z}^{\mathrm{SMC}}$ under the optimal schedule is $\approx \Lambda_{\mathrm{SMC}}^2/(NM)$ for N large, which vanishes with refinement. This is why TDS is accurate with few particles, and its optimal schedule places steps where Ψ_t fluctuates most.

Relation between the two barriers. The two barriers are equivalent up to constants. Let $v(t) = \mathrm{Var}(\Psi_t(X_t, \xi))$ as in Eq. (D.21) and, where $v(t) > 0$, let

$$\kappa_t := \frac{\mathbb{E}\big[(\Psi_t(X_t, \xi) - \mathbb{E}\Psi_t(X_t, \xi))^4\big]}{v(t)^2} \tag{D.24}$$

be the kurtosis of $\Psi_t(X_t, \xi)$. Under (A1)–(A2) all moments of Ψ_t are bounded uniformly in t, so κ_t is finite wherever $v(t) > 0$.

Proposition D.13 (Equivalence of the barriers)**.** *For every t with $v(t) > 0$,*

$$\frac{\sqrt{v(t)}}{\sqrt{\kappa_t + 3}} \;\le\; \lambda(t) \;\le\; \frac{\sqrt{v(t)}}{\sqrt{2}}, \tag{D.25}$$

and both sides vanish with $\lambda(t)$ when $v(t) = 0$. Consequently, if $\bar{\kappa} := \sup\{\kappa_t : v(t) > 0\} < \infty$,

$$\frac{\Lambda_{\mathrm{SMC}}}{\sqrt{\bar{\kappa} + 3}} \;\le\; \Lambda \;\le\; \frac{\Lambda_{\mathrm{SMC}}}{\sqrt{2}}. \tag{D.26}$$

Proof. Let $Y := \Psi_t(X_t, \xi) - \Psi_t(X_t', \xi')$, a difference of two i.i.d. copies, so that $\lambda(t) = \frac{1}{2}\mathbb{E}|Y|$, $\mathbb{E}Y = 0$, $\mathbb{E}Y^2 = 2v(t)$, and

$$\mathbb{E}Y^4 = 2\,\mathbb{E}\big[(\Psi_t - \mathbb{E}\Psi_t)^4\big] + 6\,v(t)^2 = (2\kappa_t + 6)\,v(t)^2.$$

Upper bound. By Cauchy–Schwarz, $\mathbb{E}|Y| \le \sqrt{\mathbb{E}Y^2} = \sqrt{2v(t)}$, so $\lambda(t) \le \sqrt{v(t)/2}$.

Lower bound. Write $Y^2 = |Y|^{2/3}\,|Y|^{4/3}$ and apply Hölder's inequality with exponents $3/2$ and 3:

$$\mathbb{E}Y^2 \;\le\; \big(\mathbb{E}|Y|\big)^{2/3}\big(\mathbb{E}Y^4\big)^{1/3}.$$

Rearranging and substituting the moments,

$$\mathbb{E}|Y| \;\ge\; \frac{(\mathbb{E}Y^2)^{3/2}}{(\mathbb{E}Y^4)^{1/2}} = \frac{(2v(t))^{3/2}}{\sqrt{2\kappa_t + 6}\;v(t)} = \frac{2\sqrt{v(t)}}{\sqrt{\kappa_t + 3}},$$

so $\lambda(t) \geq \sqrt{v(t)}/\sqrt{\kappa_t + 3}$. When $v(t) = 0$, $\Psi_t(X_t, \xi)$ is almost surely constant and $\lambda(t) = 0 = \sqrt{v(t)}$.

Integrating Eq. (D.25) over t gives Eq. (D.26), with $\Lambda_{\mathrm{SMC}} = \int_0^1 \sqrt{v(t)}\,\mathrm{d}t$. □

Both barriers vanish on the Doob path with the twisted proposal, where $\psi_t \equiv 0$, and away from it they are comparable: the two samplers pay for the same mismatch in different currencies, parallel control through rejected exchanges that slow restarts, sequential control through weight variance that must be absorbed by particles. The kurtosis bound $\bar{\kappa} < \infty$ excludes only the degenerate case in which $v(t)$ vanishes at some times but not identically while Ψ_t has heavy tails there. It holds, for instance, whenever v is bounded away from zero on $[0, 1]$, the analogue of the non-degeneracy assumption of Syed et al. (2026).

D.8 A Closed-Form Example

Take a variance-preserving diffusion with forward rate γ_t and signal coefficient $m_t^2 = e^{-\int_0^t \gamma}$, with $p_0 = N(0, 1)$ so that $p_t = N(0, 1)$ for all t, exact Gaussian reverse kernels with $\sigma_t^2 = \gamma_t$, and the linear reward $r(x) = \kappa x$. Take the Doob path scaled by a fixed $\beta \in (0, 1]$, $G_t = (G_t^*)^\beta$. Since $X_0 \mid X_t = x \sim N(m_t x,\, 1 - m_t^2)$, with $s_t^2 := 1 - m_t^2$,

$$G_t^*(x) = \mathbb{E}\big[e^{\kappa X_0} \mid X_t = x\big] = \exp\{\kappa m_t x + \tfrac{1}{2}\kappa^2 s_t^2\}, \qquad \log G_t(x) = \beta\kappa m_t x + \text{const}, \qquad \pi_t = N(\beta\kappa m_t,\, 1).$$

Unguided. For $t' < t$ the reverse kernel is

$$X_{t'} \mid X_t \;\sim\; N\Big(\frac{m_t}{m_{t'}} X_t,\; 1 - \frac{m_t^2}{m_{t'}^2}\Big),$$

so $m_{t'} X_{t'} - m_t X_t$ is $N(0,\, m_{t'}^2 - m_t^2)$ given X_t, whatever its tilt. The log-weight of the exchange from t to t' is $\beta\kappa(m_{t'} X_{t'} - m_t X_t)$ up to a constant, hence exactly $N\big(0, \beta^2\kappa^2(m_{t'}^2 - m_t^2)\big)$, and by Lem. D.11

$$r(t \to t') = \operatorname{erf}\Big(\frac{\beta\kappa\sqrt{m_{t'}^2 - m_t^2}}{2}\Big). \tag{D.27}$$

Variances add over sub-intervals, with total $\beta^2\kappa^2$ from $t = 1$ ($m_1 \approx 0$) to $t = 0$ ($m_0 = 1$), so Prop. D.12 applies with $S = \beta^2\kappa^2(m_{t'}^2 - m_t^2)$ on a stretch $[t', t]$. For small Δt, $m_{t-\Delta t}^2 - m_t^2 \approx \gamma_t m_t^2 \Delta t$, and the rejection rate is $\approx \beta\kappa m_t\sqrt{\gamma_t \Delta t}/\sqrt{\pi}$, which matches $\lambda^{\mathrm{u}}(t)\sqrt{\Delta t}$: here $\|g_t\| = \beta\kappa m_t\sqrt{\gamma_t}$ is constant in x, so Eq. (3.11) gives $\lambda^{\mathrm{u}}(t) = \sqrt{2}\,\|g_t\|/\sqrt{2\pi} = \|g_t\|/\sqrt{\pi}$.

Guided. $\log G_t$ is affine in x, so $M_t = 0$, the guided step coincides with the twisted proposal, and ψ_t is constant in x because $G_t \propto G_t^*$. Hence $\lambda = \lambda^{\mathrm{tw}} \equiv 0$ and every exchange is accepted, whatever the schedule. For $\beta = 1$ the guided step is the Euler–Maruyama discretisation of the exact Doob transform Eq. (A.13). The gap between zero and infinity here is a degenerate case of Thm. 3.2, but it isolates the mechanism: with the gradient of the potential constant in x, there is nothing for a guided exchange to pay for, and everything for an unguided one.

E Experimental Details

Experiments use NVIDIA GeForce RTX 2080 Ti and NVIDIA A100 GPUs, and Intel Core i9-10900X CPUs at 3.70 GHz. Synthetic sampling uses CPUs; molecular and image models use GPUs.

E.1 Synthetic Reward Tilting with Sampler Perturbations

Target. We use a six-component 1D Gaussian mixture and the 40-component 2D FAB mixture of Midgley et al. (2023). The 1D means, standard deviations and weights are

$$\begin{aligned}\mu &= (-1.25, -0.58, -0.16, 0.38, 0.74, 1.12),\\ s &= (0.070, 0.110, 0.080, 0.065, 0.055, 0.075),\\ \pi &= (0.04, 0.09, 0.13, 0.12, 0.28, 0.34).\end{aligned}$$

The 2D benchmark has uniform weights and isotropic component standard deviation 1.3132616. For mixture responsibilities $\gamma_k(x)$, the smooth reward is $r(x) = \log \sum_k \gamma_k(x) e^{u_k}$, where $u_k = a\, v^\top \mu_k / \max_j |v^\top \mu_j|$. We use $(v, a, \lambda) = (1, -1, 3)$ in 1D and $((1, -1), 1, 4)$ in 2D, with $g(x) = e^{\lambda r(x)}$. The discontinuous 2D reward is the indicator that the most responsible component belongs to $\{9, 11, 20, 24, 26, 35, 37, 38\}$, with $\lambda = 4$.

Samplers. Exact mixture scores drive 200 stochastic reverse transitions on a Karras schedule with $(\sigma_{\min}, \sigma_{\max}, \rho) = (0.002, 80, 7)$ (Karras et al., 2022):

$$x_{\mathrm{lo}} = x_{\mathrm{hi}} + \Delta \nabla \log p_{\sigma_{\mathrm{hi}}}(x_{\mathrm{hi}}) + \sqrt{c(\sigma_{\mathrm{hi}})\Delta}\, \epsilon, \quad \Delta = \sigma_{\mathrm{hi}}^2 - \sigma_{\mathrm{lo}}^2, \quad \epsilon \sim \mathcal{N}(0, I). \tag{E.1}$$

Initialization uses the exact noised mixture distribution, and the endpoint map is the mixture posterior mean. Nominal sampling uses $c = 1$. For the perturbed sampler, we set $c = 0.01$ above noise thresholds 0.3 in 1D and 5.4341856 in 2D, and $c = 1$ otherwise, while leaving the drift unchanged. Each sampler kernel K induces its own endpoint law p_K and hence its own target,

$$q_K(\mathrm{d}x) \propto g(x)\, p_K(\mathrm{d}x). \tag{E.2}$$

The nominal and perturbed samplers have different endpoint targets despite sharing the same reward.

Methods. ENCORE and CREPE (He et al., 2026b) use 51 uniformly spaced replica levels, four independent ladders per seed, alternating exchanges, reference refreshes, and no within-level moves. The potentials are $G_i = g(\widehat{x}_0(x_i))^{(1-i/200)^5}$. Each ladder discards 4,000 half-sweeps and retains the subsequent 20,480 states, yielding 81,920 outputs per seed; rejected and inactive states are retained. ENCORE uses the reverse transition kernels as proposals, $L_n = K_n$. CREPE resolves every transition between adjacent replicas and applies its acceptance test at each exchange. TDS (Wu et al., 2023a) and FKD-Diff (Singhal et al., 2025) use 4, 32, or 256 particles. TDS retains the terminal importance weights, while FKD-Diff performs adaptive resampling at steps 40, 80, 120, 160, and 200 whenever the ESS falls below half the particle population. SVDD-MC (Li et al., 2025) uses ten candidates, a width-64 value network, 4,096 training rollouts, 2,048 validation rollouts, and 1,000 Adam steps with learning rate 10^{-3}. Uncontrolled sampling and SNIS use the same base draws, with SNIS assigning weights proportional to g. All methods use the same endpoint map.

Evaluation. We use eight seeds and compare methods either at 81,920 outputs per method or at forward-evaluation budgets of 5.12, 6.4, 12.8, and 25.6 million, including initialization, warm-up, endpoint decoding, and value-network training rollouts. Accuracy is measured by weighted empirical W_2 in 1D and sliced W_2 over 32 fixed directions in 2D. Metrics are computed directly from the empirical samples, with weighted methods evaluated as weighted empirical measures. We report means and sample standard deviations across seeds, pooling ladders within each seed.

E.1.1 Additional Results

Fig. 8 shows the 2D densities pooled across eight seeds. Tab. 4 compares methods at a common output count, while Tabs. 5 to 7 report accuracy under fixed forward-evaluation budgets. Tab. 8 varies the global variance multiplier and the number of uniformly spaced replicas while retaining the 200-step generator and all intervening transitions. Changing the variance modifies the induced target, whereas changing only the number of replicas leaves it unchanged. The tables report W_2 in 1D and sliced W_2 in 2D; parentheses indicate particle counts. SNIS and TDS are evaluated using their retained importance weights. For the zero-reward control, unguided ENCORE has exactly zero rejection, as predicted by App. C.2.1.

Exchange rejection. For post-warm-up attempted and rejected swaps A_n, D_n between adjacent replicas $n-1, n$, we report the rejection percentage $100 \sum_n D_n / \sum_n A_n$ and summed rejection $S_R = \sum_{n=1}^{R-1} D_n / A_n$. The 51-replica ladder has 50 adjacent replica pairs, with 40,960 attempts per pair per seed pooled across four ladders; hence its rejection percentage is $100 S_R / 50$. Warm-up and unconditional reference refreshes are excluded. Tab. 3 reports both quantities over eight seeds. The sum is dimensionless, can exceed one, and does not establish target accuracy.

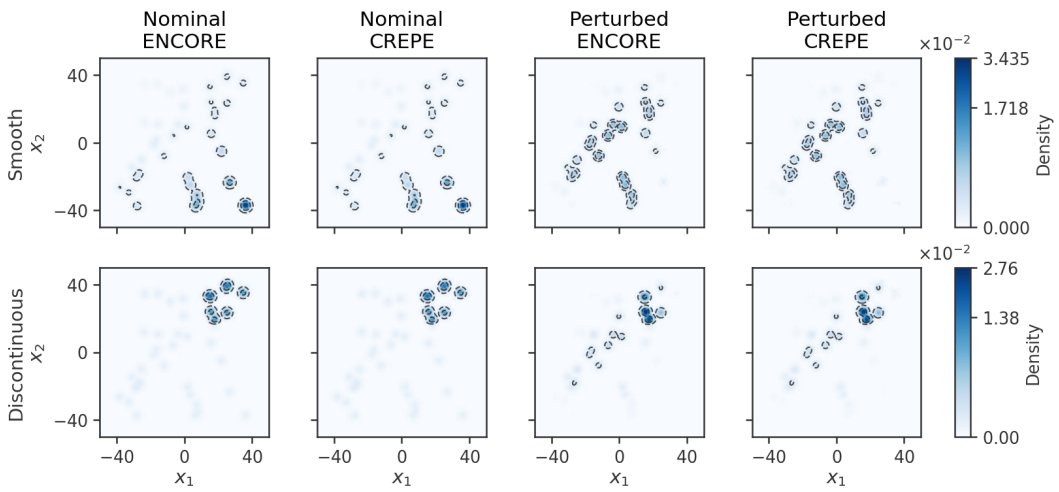

Figure 8: 2D densities for smooth (top) and discontinuous (bottom) rewards.

Table 3: Post-warm-up exchange rejection on the 51-replica ladder (mean ± standard deviation over eight seeds).

	ENCORE		CREPE	
Setting	Rejection (%)	Summed rejection	Rejection (%)	Summed rejection
1D smooth / nominal	2.55 ± 0.007	1.274 ± 0.0033	2.54 ± 0.008	1.272 ± 0.0039
1D smooth / perturbed	1.00 ± 0.009	0.500 ± 0.0043	1.81 ± 0.025	0.905 ± 0.0126
2D smooth / nominal	1.87 ± 0.009	0.936 ± 0.0043	1.87 ± 0.007	0.937 ± 0.0037
2D smooth / perturbed	0.86 ± 0.016	0.429 ± 0.0078	0.87 ± 0.014	0.433 ± 0.0070
2D discontinuous / nominal	1.51 ± 0.005	0.756 ± 0.0025	1.51 ± 0.005	0.756 ± 0.0025
2D discontinuous / perturbed	1.59 ± 0.008	0.793 ± 0.0041	1.59 ± 0.006	0.795 ± 0.0030

E.2 Sampling Alanine Transition Regions

Target. We apply a single harmonic umbrella to the 22-atom ALDP system at 800 K:

$$G_0(x) = \exp\left[-\frac{\widetilde{\phi}(x)^2}{2(0.2)^2}\right], \qquad q_{\text{model}}(x) \propto p_0(x)G_0(x), \tag{E.3}$$

where $\widetilde{\phi}$ is the torsion wrapped to $[-\pi, \pi)$; ψ is unrestrained. We report the band $|\widetilde{\phi}| \leq 0.4$ rad, split into upper ($\psi \geq 0$) and lower ($\psi < 0$) sectors, retaining outside-band samples. This is conformation sampling within a prescribed region, rather than a measurement of transition rates.

Samplers. All methods use the ALDP checkpoint from App. E.3, 200 Karras reverse transitions, terminal denoising and the same bond-constraint projection. The shared intermediate log potential is $\ell_i(z) = \gamma_i \log G_0(\mathcal{R}_i(z))$, with projected denoiser readout $\mathcal{R}_i$ and $\gamma_i = [((1 - i/200) - 5/7)/(2/7)]_{[0,1]}$. The model reference weights 262,144 independent model draws by G_0; the physical reference weights 150,000 equilibrium MD frames at 800 K. These estimate the tilted model and physical laws, respectively, and are kept distinct.

Methods. ENCORE and CREPE use eight ladders with levels $(200, 196, \ldots, 0)$ and run for 20,000 half-sweeps, discarding the first 5,000 and yielding 120,000 outputs per seed. ENCORE uses unguided native reverse proposals, while CREPE uses its fixed-CV noising kernel without temperature scaling. FKD and TDS use 256 particles per population. FKD checks for resampling at the replica levels using an ESS threshold of half the population, whereas TDS uses a guidance coefficient of 1, systematic resampling at every native step, and the Gaussian proposal correction of Wu et al. (2023a). Both methods resample at the endpoint. Their 315 and 337 populations yield 80,640 and 86,272 outputs,

Table 4: Synthetic reward-tilting accuracy at 81,920 outputs per seed; parentheses indicate particle counts (mean ± standard deviation over eight seeds).

	1D smooth $W_2 \downarrow$		2D smooth $SW_2 \downarrow$		2D discontinuous $SW_2 \downarrow$	
Method	Nominal	Perturbed	Nominal	Perturbed	Nominal	Perturbed
ENCORE	0.024 ± 0.0113	0.015 ± 0.0069	0.648 ± 0.1195	0.458 ± 0.1854	0.341 ± 0.1018	0.402 ± 0.1283
CREPE	0.029 ± 0.0090	0.650 ± 0.0331	0.700 ± 0.3051	0.609 ± 0.2473	0.345 ± 0.0924	0.513 ± 0.1155
TDS (4)	0.263 ± 0.0077	0.046 ± 0.0015	11.523 ± 0.1663	5.203 ± 0.0690	11.543 ± 0.0898	14.487 ± 0.0784
TDS (32)	0.097 ± 0.0062	0.026 ± 0.0040	6.991 ± 0.1277	1.744 ± 0.1064	7.981 ± 0.1788	5.695 ± 0.1984
TDS (256)	0.064 ± 0.0085	0.016 ± 0.0073	6.058 ± 0.2063	0.525 ± 0.1124	1.804 ± 0.2730	1.654 ± 0.2651
FKD (4)	0.974 ± 0.0045	0.178 ± 0.0012	21.461 ± 0.1046	7.925 ± 0.0247	18.273 ± 0.0862	15.172 ± 0.0816
FKD (32)	0.192 ± 0.0068	0.154 ± 0.0024	5.611 ± 0.1624	8.036 ± 0.0352	4.826 ± 0.0902	2.094 ± 0.0690
FKD (256)	0.057 ± 0.0060	0.179 ± 0.0026	1.750 ± 0.1502	8.291 ± 0.0288	3.236 ± 0.1102	3.028 ± 0.1010
SVDD-MC	0.337 ± 0.0606	0.154 ± 0.0149	10.961 ± 1.7239	7.554 ± 0.0439	15.062 ± 0.7323	19.925 ± 0.0916
SNIS	0.012 ± 0.0055	0.007 ± 0.0033	0.416 ± 0.2029	0.325 ± 0.1625	0.155 ± 0.0577	0.179 ± 0.0693
Uncontrolled	1.386 ± 0.0018	0.187 ± 0.0009	28.819 ± 0.0417	8.378 ± 0.0337	28.782 ± 0.0350	20.506 ± 0.0694

Table 5: 1D smooth reward-tilting accuracy at fixed forward-evaluation budgets (mean ± standard deviation over eight seeds).

	$W_2 \downarrow$			
Forward evaluations	5,120,000	6,400,000	12,800,000	25,600,000
Nominal sampler				
ENCORE	0.047 ± 0.0263	0.041 ± 0.0224	0.026 ± 0.0181	0.024 ± 0.0085
CREPE	0.103 ± 0.0551	0.063 ± 0.0161	0.035 ± 0.0097	0.029 ± 0.0105
TDS (4)	0.272 ± 0.0116	0.272 ± 0.0110	0.264 ± 0.0135	0.264 ± 0.0101
TDS (32)	0.107 ± 0.0258	0.106 ± 0.0261	0.100 ± 0.0160	0.098 ± 0.0088
TDS (256)	0.062 ± 0.0112	0.065 ± 0.0101	0.060 ± 0.0094	0.063 ± 0.0111
FKD (4)	0.971 ± 0.0084	0.971 ± 0.0071	0.973 ± 0.0056	0.975 ± 0.0037
FKD (32)	0.196 ± 0.0084	0.194 ± 0.0079	0.193 ± 0.0072	0.193 ± 0.0063
FKD (256)	0.060 ± 0.0121	0.057 ± 0.0119	0.056 ± 0.0082	0.058 ± 0.0039
SVDD-MC	0.337 ± 0.0622	0.332 ± 0.0603	0.335 ± 0.0617	0.337 ± 0.0617
SNIS	0.015 ± 0.0085	0.021 ± 0.0066	0.014 ± 0.0102	0.009 ± 0.0030
Uncontrolled	1.386 ± 0.0046	1.388 ± 0.0028	1.387 ± 0.0022	1.387 ± 0.0017
Perturbed sampler				
ENCORE	0.031 ± 0.0192	0.027 ± 0.0130	0.018 ± 0.0080	0.013 ± 0.0061
CREPE	0.593 ± 0.1007	0.656 ± 0.0893	0.662 ± 0.0488	0.647 ± 0.0370
TDS (4)	0.048 ± 0.0049	0.047 ± 0.0043	0.046 ± 0.0027	0.046 ± 0.0023
TDS (32)	0.027 ± 0.0055	0.028 ± 0.0058	0.026 ± 0.0060	0.026 ± 0.0038
TDS (256)	0.024 ± 0.0104	0.021 ± 0.0096	0.017 ± 0.0059	0.015 ± 0.0076
FKD (4)	0.178 ± 0.0030	0.178 ± 0.0025	0.178 ± 0.0016	0.178 ± 0.0011
FKD (32)	0.155 ± 0.0047	0.155 ± 0.0038	0.155 ± 0.0021	0.154 ± 0.0020
FKD (256)	0.178 ± 0.0050	0.178 ± 0.0044	0.178 ± 0.0029	0.179 ± 0.0021
SVDD-MC	0.155 ± 0.0133	0.153 ± 0.0138	0.155 ± 0.0150	0.155 ± 0.0143
SNIS	0.012 ± 0.0052	0.011 ± 0.0059	0.008 ± 0.0041	0.007 ± 0.0020
Uncontrolled	0.187 ± 0.0018	0.187 ± 0.0015	0.187 ± 0.0011	0.187 ± 0.0008

respectively. ENCORE, FKD, and TDS each use approximately 17.34–17.36 million forward function evaluations (NFEs) per seed, including warm-up, while CREPE uses 20.76 million. TDS additionally requires 4.92 million backward evaluations.

Evaluation. We use three seeds. Regional TVD is the total variation distance between the probabilities assigned to the upper, lower, and out-of-band regions. Within the band, energy error is measured by the supremum distance between the empirical energy CDF and the model-reference CDF. Molecular potential energies are evaluated with OpenMM, excluding the umbrella potential and kinetic energy (Eastman et al., 2023). All quantitative metrics use the complete set of retained

Table 6: 2D smooth reward-tilting accuracy at fixed forward-evaluation budgets (mean ± standard deviation over eight seeds).

	$SW_2 \downarrow$			
Forward evaluations	5,120,000	6,400,000	12,800,000	25,600,000
Nominal sampler				
ENCORE	1.105 ± 0.3611	1.236 ± 0.3721	0.855 ± 0.1674	0.485 ± 0.2042
CREPE	3.319 ± 0.9357	1.982 ± 0.6748	1.217 ± 0.4132	0.814 ± 0.3448
TDS (4)	11.367 ± 0.3464	11.424 ± 0.2930	11.436 ± 0.2125	11.493 ± 0.1971
TDS (32)	7.001 ± 0.1957	7.077 ± 0.3123	7.034 ± 0.1980	6.967 ± 0.1900
TDS (256)	6.179 ± 0.6875	6.147 ± 0.7077	5.932 ± 0.3251	6.077 ± 0.2521
FKD (4)	21.556 ± 0.1374	21.520 ± 0.1308	21.463 ± 0.0986	21.425 ± 0.0873
FKD (32)	5.574 ± 0.3091	5.599 ± 0.2698	5.581 ± 0.1906	5.610 ± 0.1159
FKD (256)	1.967 ± 0.3034	1.966 ± 0.2272	1.788 ± 0.2128	1.707 ± 0.1150
SVDD-MC	10.973 ± 1.8865	10.988 ± 1.8396	10.867 ± 1.7939	10.937 ± 1.7837
SNIS	0.503 ± 0.1874	0.546 ± 0.2097	0.426 ± 0.1624	0.292 ± 0.1252
Uncontrolled	28.800 ± 0.0523	28.803 ± 0.0390	28.822 ± 0.0467	28.821 ± 0.0177
Perturbed sampler				
ENCORE	1.104 ± 0.4091	1.090 ± 0.3965	0.555 ± 0.2726	0.436 ± 0.2050
CREPE	2.546 ± 0.5044	1.284 ± 0.5484	0.790 ± 0.3732	0.630 ± 0.3012
TDS (4)	5.195 ± 0.1540	5.225 ± 0.1469	5.223 ± 0.0823	5.196 ± 0.0633
TDS (32)	1.860 ± 0.3219	1.921 ± 0.2425	1.825 ± 0.1826	1.709 ± 0.1214
TDS (256)	0.837 ± 0.2607	0.752 ± 0.1659	0.591 ± 0.2230	0.514 ± 0.1421
FKD (4)	7.947 ± 0.0672	7.930 ± 0.0568	7.914 ± 0.0329	7.932 ± 0.0287
FKD (32)	8.052 ± 0.0614	8.041 ± 0.0678	8.035 ± 0.0322	8.032 ± 0.0302
FKD (256)	8.294 ± 0.0757	8.283 ± 0.0719	8.293 ± 0.0319	8.294 ± 0.0250
SVDD-MC	7.634 ± 0.1963	7.623 ± 0.1996	7.560 ± 0.1374	7.554 ± 0.0959
SNIS	0.435 ± 0.1611	0.435 ± 0.1925	0.363 ± 0.1534	0.275 ± 0.0920
Uncontrolled	8.373 ± 0.0618	8.380 ± 0.0708	8.377 ± 0.0450	8.378 ± 0.0256

samples. Density plots show 5,000 samples from a single seed, selected before any conditioning, while reference samples retain their full importance weights.

E.2.1 Additional Results

Fig. 9 shows the conditional conformational densities. Tab. 9 reports accuracy against the model + IS reference under two forward-NFE budgets, spanning 8.08–10.36 and 17.34–20.76 million evaluations per seed, with exact budgets reported for each method. The lower- and higher-budget evaluations use nested portions of the same trajectories.

E.3 Alanine Tempering

Target. We temper alanine dipeptide (ALDP; 22 atoms), tetrapeptide (A4; 43 atoms), and hexapeptide (A6; 63 atoms) from 800 K to 300, 500, and 600 K, respectively. For molar potential energy U, the endpoint factor is $g(x) = \exp[-(\beta_{T_t} - \beta_{T_s})U(x)]$, with $\beta_T = 1/(RT)$. OpenMM evaluates U (Eastman et al., 2023); held-out target-temperature MD supplies the physical reference.

Samplers. We use the CREPE molecular diffusion implementation (He et al., 2026b): an EDM-preconditioned equivariant network with five message-passing layers and width 256, with 200 Karras steps and $(\sigma_{\min}, \sigma_{\max}, \rho) = (0.001, 10, 7)$. ALDP uses the 800 K RNE checkpoint (He et al., 2026a); A4 uses an EMA checkpoint at step 190,000, selected on source-temperature validation data after FEAT-style training (Du et al., 2025); A6 uses the width-256 checkpoint supplied by the CREPE authors.

Methods. ENCORE compresses the high-noise zero-potential region while retaining native transitions, then increases the reward to the full endpoint tilt. CREPE uses 51 levels, one local ULA update per level, reference refreshes and alternating exchanges. CREPE retains 49,000 outputs after 1,000 warm-up half-sweeps. RNE and FKC-Score/Noise use 50 populations of 1,000 particles with

Table 7: 2D discontinuous reward-tilting accuracy at fixed forward-evaluation budgets (mean ± standard deviation over eight seeds).

	$SW_2 \downarrow$			
Forward evaluations	5,120,000	6,400,000	12,800,000	25,600,000
Nominal sampler				
ENCORE	0.858 ± 0.4000	0.825 ± 0.2631	0.419 ± 0.1052	0.298 ± 0.0888
CREPE	2.247 ± 0.9686	1.461 ± 0.4610	0.640 ± 0.2679	0.370 ± 0.1292
TDS (4)	11.569 ± 0.3386	11.541 ± 0.2401	11.540 ± 0.1328	11.514 ± 0.0797
TDS (32)	7.888 ± 0.5637	7.908 ± 0.5490	7.904 ± 0.4265	7.967 ± 0.2195
TDS (256)	1.774 ± 0.6222	1.725 ± 0.3936	1.775 ± 0.4473	1.843 ± 0.2417
FKD (4)	18.237 ± 0.1115	18.257 ± 0.0885	18.283 ± 0.1135	18.272 ± 0.0253
FKD (32)	4.829 ± 0.1546	4.833 ± 0.1310	4.816 ± 0.1064	4.841 ± 0.0577
FKD (256)	3.283 ± 0.1340	3.268 ± 0.0950	3.242 ± 0.1193	3.169 ± 0.0967
SVDD-MC	14.915 ± 1.0940	14.970 ± 0.8346	15.216 ± 0.8642	15.103 ± 0.7567
SNIS	0.243 ± 0.0810	0.231 ± 0.0619	0.160 ± 0.0705	0.120 ± 0.0209
Uncontrolled	28.788 ± 0.0977	28.777 ± 0.0794	28.778 ± 0.0536	28.769 ± 0.0341
Perturbed sampler				
ENCORE	0.856 ± 0.2147	0.701 ± 0.2228	0.579 ± 0.1724	0.351 ± 0.1343
CREPE	1.753 ± 0.6221	1.230 ± 0.5334	0.702 ± 0.1470	0.530 ± 0.1242
TDS (4)	14.581 ± 0.2542	14.562 ± 0.2125	14.494 ± 0.1762	14.493 ± 0.1064
TDS (32)	5.826 ± 0.2727	5.882 ± 0.2538	5.786 ± 0.1990	5.737 ± 0.2345
TDS (256)	1.727 ± 0.3931	1.641 ± 0.4391	1.710 ± 0.3187	1.679 ± 0.2594
FKD (4)	15.183 ± 0.1652	15.192 ± 0.1134	15.175 ± 0.0904	15.152 ± 0.0558
FKD (32)	2.096 ± 0.0963	2.106 ± 0.0891	2.110 ± 0.0870	2.087 ± 0.0439
FKD (256)	3.028 ± 0.2118	3.004 ± 0.2208	3.050 ± 0.1115	3.051 ± 0.0903
SVDD-MC	19.785 ± 0.5226	19.787 ± 0.5443	19.904 ± 0.2690	19.902 ± 0.1780
SNIS	0.304 ± 0.1254	0.276 ± 0.1324	0.219 ± 0.0785	0.157 ± 0.0378
Uncontrolled	20.487 ± 0.1078	20.486 ± 0.0956	20.495 ± 0.0798	20.495 ± 0.0466

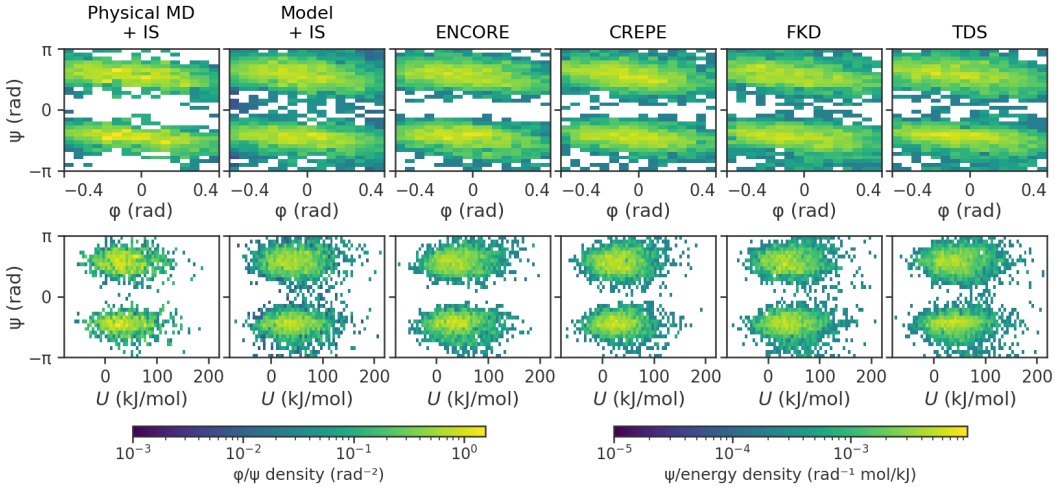

Figure 9: Conditional ϕ/ψ densities (top) and joint ψ/energy densities (bottom) within $|\phi| \leq 0.4$.

resampling every four transitions (He et al., 2026a; Skreta et al., 2025). FKC-Score multiplies the score by T_s/T_t; FKC-Noise uses score coefficient $(T_s/T_t + 1)/2$ and noise multiplier $(T_s/T_t)^{-1/2}$. FKD-Diff uses 32 particles, reward checks at steps $124, 132, \ldots, 164, 168$, and resampling below half-population ESS (Singhal et al., 2025). FKD and uncontrolled sampling each return 50,000 outputs. For the separate acceptance/diversity comparison, ENCORE and CREPE use dense $(200, 196, \ldots, 0)$ and coarse-to-fine $(200, 80, 76, \ldots, 0)$ ladders, one chain per seed and three seeds. ENCORE uses $L_n = K_n$ and an additional reward-only endpoint replica; both use 50,000 half-sweeps, including 1,000 warm-up.

Table 8: Synthetic reward tilting under variance and replica-schedule changes (mean ± standard deviation over eight seeds).

		W_2 (1D), SW_2 (2D) ↓		
Transition-variance multiplier (51 replicas)		0.25	1	4
1D smooth	ENCORE	0.050 ± 0.0175	0.024 ± 0.0113	0.007 ± 0.0020
	CREPE	0.344 ± 0.0172	0.029 ± 0.0090	0.056 ± 0.0059
2D smooth	ENCORE	0.985 ± 0.2812	0.648 ± 0.1195	0.473 ± 0.0922
	CREPE	0.865 ± 0.2375	0.700 ± 0.3051	0.528 ± 0.1812
2D discontinuous	ENCORE	0.452 ± 0.0993	0.341 ± 0.1018	0.277 ± 0.0672
	CREPE	0.630 ± 0.1465	0.345 ± 0.0924	0.406 ± 0.1526
Uniform replica count (nominal sampler)		6	11	51
1D smooth	ENCORE	0.031 ± 0.0128	0.029 ± 0.0150	0.024 ± 0.0113
	CREPE	0.025 ± 0.0091	0.028 ± 0.0147	0.029 ± 0.0090
2D smooth	ENCORE	1.069 ± 0.3577	0.813 ± 0.2756	0.648 ± 0.1195
	CREPE	0.815 ± 0.4486	0.826 ± 0.3959	0.700 ± 0.3051
2D discontinuous	ENCORE	0.401 ± 0.1643	0.434 ± 0.0934	0.341 ± 0.1018
	CREPE	0.514 ± 0.3543	0.445 ± 0.1279	0.345 ± 0.0924
Uniform replica count (perturbed sampler)		6	11	51
1D smooth	ENCORE	0.012 ± 0.0047	0.013 ± 0.0045	0.015 ± 0.0069
	CREPE	0.323 ± 0.0392	0.474 ± 0.0473	0.650 ± 0.0331
2D smooth	ENCORE	0.423 ± 0.1169	0.392 ± 0.1062	0.458 ± 0.1854
	CREPE	0.467 ± 0.2270	0.516 ± 0.1145	0.609 ± 0.2473
2D discontinuous	ENCORE	0.589 ± 0.2536	0.419 ± 0.1650	0.402 ± 0.1283
	CREPE	0.621 ± 0.2754	0.558 ± 0.1494	0.513 ± 0.1155

Table 9: Transition-region accuracy across forward-evaluation budgets (mean ± standard deviation over three seeds).

Method	Forward NFEs	Regional TVD↓	Energy CDF↓
ENCORE (8 chains)	8,681,760	0.024 ± 0.0235	0.031 ± 0.0188
	17,361,760	0.012 ± 0.0041	0.028 ± 0.0138
CREPE (8 chains)	10,361,600	0.033 ± 0.0209	0.148 ± 0.0124
	20,761,600	0.027 ± 0.0142	0.151 ± 0.0079
FKD (256 particles)	8,641,280	0.051 ± 0.0330	0.021 ± 0.0042
	17,337,600	0.037 ± 0.0120	0.016 ± 0.0071
TDS (256 particles)	8,078,592	0.030 ± 0.0016	0.025 ± 0.0107
	17,340,672	0.037 ± 0.0170	0.020 ± 0.0150

Evaluation. Boltzmann-tempering errors use 25,000 target-temperature MD frames and three evaluations drawing 5,000 configurations without replacement from both outputs and reference. Thus their standard deviations describe subsampling of one sampler run. Energy TVD uses 100-bin histograms of the negative regularised reduced potential; distance TVD pools unordered interatomic distances. Aligned W_2 uses centred, Kabsch-aligned coordinates in nm (Kabsch, 1976); torsion MMD^2 uses a Gaussian kernel on sine–cosine backbone torsions (Gretton et al., 2012). A4/A6 ENCORE and FKD outputs are projected onto OpenMM bond constraints. For acceptance/diversity, RMSD W_1 compares distributions from 100,000 backbone-heavy-atom pairs per sampler and MD reference, without energy filtering. TICA JS uses 72×72 target-fitted histograms and natural logarithms; Energy TVD averages three 100-bin evaluations per run. ENCORE is evaluated after constraint projection and before terminal denoising; CREPE uses its level-0 state directly. Histogram metrics retain out-of-range mass; acceptance excludes warm-up.

Table 10: ALDP and A4 Boltzmann-tempering accuracy with 5,000 samples per evaluation (mean ± standard deviation over three subsampling evaluations from one sampler run).

System	Method	Energy TVD↓	Distance TVD↓	Aligned W_2 ↓	Torsion MMD^2↓
Alanine Dipeptide (ALDP) 800 → 300 K 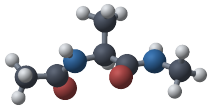	ENCORE	0.214 ± 0.0040	**0.015** ± 0.0008	**0.243** ± 0.0025	**0.004** ± 0.0004
	CREPE	**0.201** ± 0.0020	0.019 ± 0.0002	0.257 ± 0.0010	0.017 ± 0.0012
	RNE	0.363 ± 0.0070	0.017 ± 0.0006	0.277 ± 0.0014	0.005 ± 0.0004
	FKC-Score	0.369 ± 0.0092	0.019 ± 0.0007	0.283 ± 0.0005	0.006 ± 0.0005
	FKC-Noise	0.884 ± 0.0006	0.032 ± 0.0001	0.266 ± 0.0008	0.010 ± 0.0004
	FKD	0.564 ± 0.0055	0.037 ± 0.0009	0.367 ± 0.0012	0.084 ± 0.0019
	Uncontrolled	0.993 ± 0.0006	0.054 ± 0.0004	0.300 ± 0.0021	0.068 ± 0.0032
Alanine Tetrapeptide (A4) 800 → 500 K 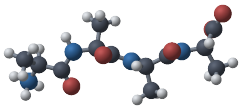	ENCORE	0.136 ± 0.0057	**0.008** ± 0.0003	**0.734** ± 0.0016	**0.002** ± 0.0003
	CREPE	0.502 ± 0.0076	0.014 ± 0.0004	0.748 ± 0.0029	0.007 ± 0.0003
	RNE	0.246 ± 0.0145	0.013 ± 0.0004	0.905 ± 0.0040	0.012 ± 0.0002
	FKC-Score	**0.117** ± 0.0143	0.025 ± 0.0001	0.995 ± 0.0010	0.024 ± 0.0002
	FKC-Noise	0.396 ± 0.0079	0.020 ± 0.0001	0.894 ± 0.0049	0.007 ± 0.0003
	FKD	0.517 ± 0.0081	0.024 ± 0.0003	0.879 ± 0.0015	0.031 ± 0.0006
	Uncontrolled	0.928 ± 0.0041	0.023 ± 0.0005	0.788 ± 0.0063	0.020 ± 0.0016

Table 11: Alanine conformational diversity and energy agreement under dense and coarse-to-fine replica ladders (mean ± standard deviation over three seeds).

		Dense			Coarse-to-fine		
System	Method	RMSD W_1 (nm)↓	TICA JS↓	Energy TVD↓	RMSD W_1 (nm)↓	TICA JS↓	Energy TVD↓
ALDP	ENCORE	**0.0037** ± 0.0008	0.205 ± 0.0088	**0.097** ± 0.0163	**0.0036** ± 0.0015	0.222 ± 0.0069	**0.079** ± 0.0120
	CREPE	0.0091 ± 0.0068	**0.068** ± 0.0094	0.225 ± 0.0174	0.0254 ± 0.0086	**0.135** ± 0.0136	0.403 ± 0.3661
A4	ENCORE	0.0070 ± 0.0020	0.144 ± 0.0034	**0.080** ± 0.0113	**0.0065** ± 0.0023	**0.140** ± 0.0050	**0.088** ± 0.0122
	CREPE	**0.0052** ± 0.0027	**0.099** ± 0.0098	0.516 ± 0.0151	0.0190 ± 0.0051	0.174 ± 0.0728	0.486 ± 0.0340
A6	ENCORE	0.0057 ± 0.0037	0.066 ± 0.0019	**0.144** ± 0.0122	**0.0063** ± 0.0035	**0.063** ± 0.0052	**0.139** ± 0.0161
	CREPE	**0.0024** ± 0.0021	**0.043** ± 0.0006	0.329 ± 0.0066	0.0296 ± 0.0157	0.188 ± 0.1009	0.291 ± 0.0504

E.3.1 Additional Results

Tab. 10 reports ALDP/A4 tempering errors; A6 appears in Tab. 2. Tab. 11 gives the full dense and coarse-to-fine comparison, with means and standard deviations across three seeds.

E.4 Text-to-Image Reward Tilting

Target. For prompt c, we target $q_0(\mathrm{d}x \mid c) \propto \exp\{100\, r_{\mathrm{IR}}(x, c)\}\, p_0(\mathrm{d}x \mid c)$, where r_{IR} is ImageReward (Xu et al., 2023), following He et al. (2026b).

Samplers. Four-step SDXL-Turbo uses Euler–Ancestral transitions at 512×512 (Sauer et al., 2024); four-step DMD2-UNet uses native LCM transitions at 1024×1024 (Yin et al., 2024). Both use timesteps $(999, 749, 499, 249)$, without classifier-free guidance or watermarking.

Methods. ENCORE uses five path levels with reward strengths $(0, 25, 50, 75, 100)$ and independent reference refreshes. With 16 or 8 chains it discards 64 or 128 iterations, then collects every 32 iterations to obtain 32 outputs. CREPE uses the same chain counts and warm-up, collecting every 50 iterations; its SDXL-Turbo adaptation excludes the deterministic final transition from exchange. We do not evaluate it on DMD2, whose native LCM sampler lacks the paired forward kernel required by its correction. FKD-Diff uses 2,048 particles with multinomial resampling after every transition and returns 32 particles (Singhal et al., 2025). LiDAR uses a 2,048-sample bank for its score correction and generates 32 trajectories (Kim et al., 2026). Both use reward strength 100. Best-of-64 returns the highest-reward output from each of 32 independent groups of 64 draws; uncontrolled sampling returns 32 independent base outputs.

Table 12: Reward, preference, prompt alignment, and diversity on SDXL-Turbo-4 (mean ± standard deviation over 32 prompts).

	Reward	*Preference / Alignment*		*Diversity*			
Method	IR ↑	HPS ↑	CLIP-T ↑	DINO-Vendi ↑	CLIP diversity ↑	LPIPS ↑	Unique (%) ↑
ENCORE (16 chains)	1.711 ± 0.2784	0.324 ± 0.0249	0.290 ± 0.0297	3.020 ± 0.9522	0.170 ± 0.0575	0.577 ± 0.0535	84.1 ± 11.12
ENCORE (8 chains)	**1.745** ± 0.2578	0.324 ± 0.0246	0.290 ± 0.0290	2.778 ± 0.7928	0.170 ± 0.0665	0.559 ± 0.0531	64.9 ± 22.59
CREPE (16 chains)	1.704 ± 0.2832	0.320 ± 0.0255	0.290 ± 0.0291	3.095 ± 0.9597	0.183 ± 0.0648	0.583 ± 0.0525	78.8 ± 10.49
CREPE (8 chains)	1.738 ± 0.2410	0.320 ± 0.0257	0.290 ± 0.0292	2.819 ± 0.8245	0.176 ± 0.0635	0.563 ± 0.0594	62.0 ± 19.12
FKD	1.703 ± 0.3255	**0.326** ± 0.0253	**0.291** ± 0.0316	1.720 ± 0.6493	0.081 ± 0.0634	0.299 ± 0.2132	44.9 ± 30.09
LiDAR	1.125 ± 0.5412	0.308 ± 0.0245	0.283 ± 0.0316	**4.390** ± 1.7833	**0.267** ± 0.0952	**0.614** ± 0.0473	**100.0** ± 0.00
Best-of-64	1.651 ± 0.3464	0.323 ± 0.0243	0.289 ± 0.0301	3.184 ± 1.0612	0.172 ± 0.0540	0.587 ± 0.0504	100.0 ± 0.00
Uncontrolled	1.048 ± 0.5492	0.311 ± 0.0255	0.285 ± 0.0300	3.905 ± 1.6762	0.218 ± 0.0854	0.609 ± 0.0472	100.0 ± 0.00

Table 13: Reward, preference, prompt alignment, and diversity on DMD2-UNet-4 (mean ± standard deviation over 32 prompts).

	Reward	*Preference / Alignment*		*Diversity*			
Method	IR ↑	HPS ↑	CLIP-T ↑	DINO-Vendi ↑	CLIP diversity ↑	LPIPS ↑	Unique (%) ↑
ENCORE (16 chains)	**1.661** ± 0.3086	**0.324** ± 0.0254	**0.299** ± 0.0335	**3.276** ± 1.1356	**0.180** ± 0.0811	**0.595** ± 0.0671	**78.3** ± 10.74
FKD	1.659 ± 0.3476	**0.324** ± 0.0238	0.298 ± 0.0375	1.473 ± 0.6032	0.054 ± 0.0557	0.194 ± 0.1744	38.5 ± 31.59
Best-of-64	1.609 ± 0.3661	0.323 ± 0.0250	0.298 ± 0.0336	3.633 ± 1.4605	0.180 ± 0.0879	0.606 ± 0.0623	100.0 ± 0.00
Uncontrolled	0.962 ± 0.6206	0.314 ± 0.0284	0.295 ± 0.0362	3.976 ± 1.4719	0.197 ± 0.0703	0.619 ± 0.0683	100.0 ± 0.00

Evaluation. We use 32 prompts from the ImageReward test split, with 32 outputs per prompt, decoded in a common evaluation environment. Metrics are ImageReward v1.0, HPSv2.1 (ViT-H/14), CLIP-T and CLIP diversity (ViT-L/14), DINO-Vendi (normalized DINOv2-Large embeddings), LPIPS (AlexNet), and the fraction of distinct decoded RGB images (Xu et al., 2023; Wu et al., 2023b; Radford et al., 2021; Friedman & Dieng, 2023; Oquab et al., 2024; Zhang et al., 2018). Reward, preference and alignment are averaged within each prompt; DINO-Vendi and uniqueness use the complete output set, and pairwise diversity averages all 496 unordered pairs. Tables report means and standard deviations across prompts.

E.4.1 Additional Results

Tabs. 12 and 13 report reward, preference, alignment and diversity for both samplers.