REFLECTED PROXIMAL DYNAMICS FOR ACCELERATED HIGH-ACCURACY LOG-CONCAVE SAMPLING

Anonymous authors
Paper under double-blind review

ABSTRACT

We study sampling from $\pi(x) \propto \mathrm{e}^{-V(x)}$ on $\mathbb{R}^d$, where V is α-strongly convex and β-smooth with condition number $\kappa = \beta/\alpha$. For this problem, proximal sampler preserves the target exactly and recovers overdamped Langevin dynamics in the small-step limit, but its Gaussian resampling step discards directional information. Underdamped Langevin dynamics retains such information through momentum to accelerate exploration, motivating a proximal construction that combines target preservation with directional persistence. We introduce *reflected proximal dynamics*, which combines proximal sampling with an exact reflection of the auxiliary variable, conditional kinetic motion, and controlled refreshment. We develop two realizations: a flux-resampled underdamped half-turn and a persistent-momentum quarter-turn. From a warm start with bounded order-two Rényi divergence, both methods achieve total variation error ε using $\widetilde{O}(\sqrt{\kappa}d^{1/4})$ expected oracle queries. We further remove the supplied warm-start assumption by combining either sampler with existing posterior samplers in a coarse-scale proximal initialization stage and optimizing the proximal step size. Accounting for both initialization and sampling yields $\widetilde{O}(\kappa^{6/7}d^{1/4})$ expected oracle queries. These results identify reflected proximal samplers as a promising framework for accelerated high-accuracy log-concave sampling.

1 INTRODUCTION

We study sampling from a distribution $\pi(x) \propto \mathrm{e}^{-V(x)}$ on $\mathbb{R}^d$, where V is α-strongly convex and β-smooth, with condition number $\kappa = \beta/\alpha$. This problem appears in fields including Bayesian inference, machine learning, and statistical physics (Gelman et al., 1995; Ghosal & van der Vaart, 2017; Newman & Barkema, 1999; Chewi, 2023), and a classical approach is Markov chain Monte Carlo (MCMC). An MCMC method constructs a Markov chain with stationary distribution π and, under suitable conditions on the transition kernel and target distribution, converges to π. Such samples allow expectations under π to be approximated even when direct integration and the normalizing constant of π are unavailable. Two workhorse continuous-time processes are (overdamped) Langevin dynamics (LD), whose Euler discretization is Langevin Monte Carlo (LMC, Dalalyan (2017); Durmus et al. (2019); Vempala & Wibisono (2019); Chewi et al. (2022)), and underdamped Langevin dynamics (ULD), which retains momentum to suppress diffusive motion (Cheng et al., 2018; Dalalyan & Riou-Durand, 2020; Ma et al., 2021; Betancourt & Girolami, 2019). LD is simple to discretize, whereas ULD couples position and momentum and requires more careful finite-step schemes (Shen & Lee, 2019; Zhang et al., 2023).

The Gaussian proximal sampler gives a target-preserving alternative for the overdamped case (Lee et al., 2021; Chen et al., 2022). It alternates a Gaussian update with a localized conditional draw, preserves π for every proximal scale, and reduces to LD in the small-step limit (Mitra & Wibisono, 2025). This comparison exposes a gap. The proximal sampler has an elegant target-preserving two-conditional structure, but it redraws the Gaussian displacement at every iteration and therefore discards directional memory. ULD retains precisely this memory, but its standard discretizations do not inherit the proximal sampler's exact two-conditional structure.

We bridge this gap with **reflected proximal dynamics**. Instead of redrawing $y \sim \mathcal{N}(x, \eta I_d)$, we set $y^+ = 2x - y$. This reflection preserves the augmented target while retaining the current Gaussian displacement. The subsequent conditional transition converts this retained displacement into persistent

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.

motion while preserving the appropriate conditional law. Occasional full conditional refreshes remove correlations that the persistent update alone may not eliminate. The resulting framework brings the key acceleration mechanism of ULD, directional persistence, into a target-preserving construction.

We give two realizations. The *flux-resampled half-turn* uses momentum only within one conditional update and resamples the incoming velocity at each collision. The *persistent quarter-turn* instead carries momentum across outer iterations and invokes independent conditional sampling only on rare refresh steps. From a bounded order-two Rényi warm start, both methods use $\widetilde{O}(\sqrt{\kappa}d^{1/4})$ expected oracle queries, up to logarithmic factors. The persistent method invokes an independent RGO draw only on rare refresh steps, resulting in a polylogarithmic expected number of RGO calls. When no warm start is supplied, we first use a tunable coarse-scale proximal stage to construct a partially warm law and then run either reflected sampler. Charging the warm-up and final sampling stages together gives $\widetilde{O}(\kappa^{6/7}d^{1/4})$ expected oracle queries.

Concurrent work on the proximal bouncy particle sampler (PBPS, Chen et al. (2026b)) uses the same Gaussian augmentation and outer reflection and attains the same $\widetilde{O}(\sqrt{\kappa}d^{1/4})$ expected oracle-query bound from a bounded order-two Rényi warm start. The methods differ in their conditional dynamics: PBPS uses an undamped harmonic half-turn with deterministic momentum reflections, whereas our realizations use a damped flux-resampled half-turn and a persistent position–momentum quarter-turn. PBPS and our flux method use an independent RGO reference in essentially every outer iteration, while the persistent method needs only a polylogarithmic expected number of RGO calls. App. A gives a detailed comparison.

2 Preliminaries

Notations. For probability densities ν and π, write the total variation distance as $\mathrm{TV}(\nu,\pi) := \sup_A|\nu(A)-\pi(A)|$. When $\nu \ll \pi$, the order-r Rényi divergence for $r>1$ is $D_r(\nu\|\pi) := \frac{1}{r-1}\log\int(\frac{\nu(x)}{\pi(x)})^r\pi(x)\,\mathrm{d}x$. In particular, the χ^2 divergence is $\chi^2(\nu\|\pi) := \int(\frac{\nu(x)}{\pi(x)}-1)^2\pi(x)\,\mathrm{d}x = \mathrm{e}^{D_2(\nu\|\pi)}-1$, and $\mathrm{TV}(\nu,\pi)\le\frac{1}{2}\sqrt{\chi^2(\nu\|\pi)}$. We call ν a constant-D_2 warm start if $D_2(\nu\|\pi)=O(1)$. The notation $\widetilde{O}(g)$ suppresses factors polylogarithmic in the displayed problem parameters, and $\widetilde{\Theta}(g)$ denotes matching upper and lower bounds up to such factors. $[a]_+ := \max\{a,0\}$. We write $\mathrm{Law}(X)$ for the law of a random variable X.

Assumption 1 (Target distribution)**.** *We assume that the target density is $\pi(x)\propto\mathrm{e}^{-V(x)}$, where the potential $V:\mathbb{R}^d\to\mathbb{R}$ is α-strongly convex and β-smooth, i.e., $\frac{\alpha}{2}\|x-x'\|^2\le V(x)-V(x')-\langle\nabla V(x'),x-x'\rangle\le\frac{\beta}{2}\|x-x'\|^2$ for all $x,x'\in\mathbb{R}^d$. We denote the condition number by $\kappa:=\beta/\alpha$.*

Query conventions. We assume oracle access to V and ∇V, and let N_{qry} denote the total number of oracle evaluations used by an algorithm. This count includes rejected proposals, event-rate evaluations, mode-finding steps, and all work internal to an approximate RGO.

2.1 Overdamped and Underdamped Langevin Dynamics

For standard Brownian motion $(B_t)_{t\ge0}$, (overdamped) Langevin dynamics (LD) follows $\mathrm{d}X_t = -\nabla V(X_t)\,\mathrm{d}t+\sqrt{2}\,\mathrm{d}B_t$, whereas underdamped Langevin dynamics (ULD) introduces momentum and follows $\mathrm{d}X_t=P_t\,\mathrm{d}t$ and $\mathrm{d}P_t=-\nabla V(X_t)\,\mathrm{d}t-\gamma P_t\,\mathrm{d}t+\sqrt{2\gamma}\,\mathrm{d}B_t$. Their stationary distributions are π and $\pi\otimes\gamma_d$, respectively, where $\gamma_d=\mathcal{N}(0,I_d)$. The distinction relevant to this work is how the two dynamics move through space. In LD, Brownian noise acts directly on position, so motion over short time scales is diffusive and carries no separate directional state from one interval to the next; ULD instead stores direction in the momentum: before friction substantially changes P_t, the position moves approximately as $X_{t+h}\approx X_t+hP_t$. This short-time transport can cross a region more coherently than a position-only diffusion. More precisely, let $\mu_t=\mathrm{Law}(X_t)$ for LD and $\widehat{\mu}_t=\mathrm{Law}(X_t,P_t)$ for ULD. For initial laws with finite χ^2 divergence from their respective stationary distributions, under Assump. 1, there are universal constants $c,C>0$ such that $\chi^2(\mu_t\|\pi)\le\mathrm{e}^{-2\alpha t}\chi^2(\mu_0\|\pi)$, $\chi^2(\widehat{\mu}_t\|\pi\otimes\gamma_d)\le C\mathrm{e}^{-c\sqrt{\alpha}t}\chi^2(\widehat{\mu}_0\|\pi\otimes\gamma_d)$, where the ULD bound uses the low-friction choice $\gamma\asymp\sqrt{\alpha}$ (Vempala & Wibisono, 2019; Chewi, 2023; Chen et al., 2026d). Thus, after normalizing $\beta=1$, the two continuous-time complexities are respectively $O(\kappa)$ and $O(\sqrt{\kappa})$.

2.2 THE PROXIMAL SAMPLER

The proximal sampler (Lee et al., 2021; Chen et al., 2022) augments $X \sim \pi$ with an auxiliary random variable Y under the joint distribution $\Pi_\eta(x, y) \propto \exp(-V(x) - \|x - y\|^2/(2\eta))$, where $\eta > 0$ is the proximal scale. Its X-marginal is π, and the conditional distributions are

$$Y \mid X = x \sim \mathcal{N}(x, \eta I_d), \qquad X \mid Y = y \sim q_y(x) \propto \exp(-V(x) - \|x - y\|^2/(2\eta)). \tag{1}$$

An exact independent draw from q_y is called an ideal **restricted Gaussian oracle (RGO)** draw. The standard proximal sampler is the Gibbs chain that draws $Y_{k+1} \sim \mathcal{N}(X_k, \eta I_d)$ and then $X_{k+1} \sim q_{Y_{k+1}}$. q_y has condition number $(\beta + \eta^{-1})/(\alpha + \eta^{-1})$, which approaches 1 as $\eta \to 0$. The small-η limit recovers LD, while the ideal Gibbs chain preserves π for every η (Mitra & Wibisono, 2025).

3 REFLECTED PROXIMAL DYNAMICS

The standard proximal sampler redraws the Gaussian auxiliary variable at every iteration. This makes each update simple, but it also discards the direction of the previous move. Reflected proximal dynamics modifies this outer update while preserving the augmented target. We first present the framework with state (X, Y); a momentum-augmented version is described at the end of the section.

3.1 THE FRAMEWORK

Recall the conditional distribution $q_y(x)$ in (1). For each fixed y, let M_y be a Markov kernel that preserves q_y. Unlike an RGO draw which depends only on Y_{k+1} when sampling X_{k+1}, the kernel $M_{Y_{k+1}}$ may also use the previous iterate X_k. The alternating updates of the ideal framework are given in Alg. 1.

The algorithm is *ideal* because the RGO branch draws exactly from q_y. The realizations in Secs. 4 and 5 replace the required exact conditional draws by controlled approximate subroutines, yielding implementable algorithms. We call the second branch *accelerated* because it is where directional persistence is introduced; invariance of an arbitrary M_y alone does not imply acceleration.

Algorithm 1: Ideal reflected proximal dynamics

Require: Proximal scale η, initial states (X_0, Y_0), total steps N, refresh probability ρ, and kernels M_y preserving q_y for any $y \in \mathbb{R}^d$

for $k = 0, \dots, N-1$ **do**
  $Y_{k+1} \leftarrow 2X_k - Y_k$ ▷ reflect
  Draw $U_k \sim \mathrm{Unif}[0, 1]$
  if $U_k \le \rho$ **then** ▷ RGO branch
    Draw $X_{k+1} \sim q_{Y_{k+1}}$ independently
  else ▷ accelerated branch
    Draw $X_{k+1} \sim M_{Y_{k+1}}(X_k, \cdot)$
  end if
end for
return X_N

3.2 WHY REFLECTION AND REFRESHMENT?

Suppose the current state is (X, Y), and write $D = Y - X$ for the Gaussian displacement. Standard proximal sampling replaces D by fresh Gaussian noise. We instead apply

$$A(x, y) = (x, 2x - y), \qquad Y^+ = 2X - Y.$$

The position X does not move, while the displacement changes from D to $-D$. Reflection therefore keeps the previous directional information available for the accelerated branch. It generally changes the current law before convergence, but it preserves Π_η: the map fixes x, sends $x - y$ to $-(x - y)$, and has unit absolute Jacobian.

The two branches play different roles. At stationarity, both return a position with conditional marginal q_y. The RGO branch produces this position independently of the previous X_k, whereas the accelerated branch starts from X_k and may keep its output correlated with the input. This correlation allows the displacement retained by reflection to affect later iterations. Standard proximal sampling takes the RGO branch at every iteration; reflected proximal dynamics uses it only occasionally.

We include the RGO branch because preservation of q_y does not by itself guarantee rapid mixing. For example, a nearly deterministic accelerated branch may retain weakly mixed components. A small *refresh probability* ρ preserves direction on most iterations while still removing these residual correlations.

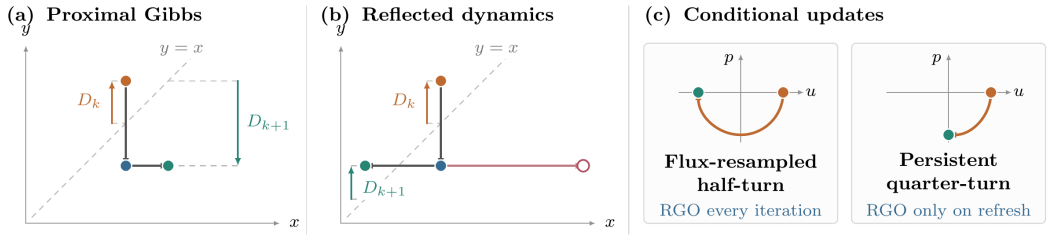

Figure 1: Panels (a)–(b) show the chain in the (x, y)-plane, with $D_k = Y_k - X_k$. Black segments show the ordinary auxiliary and conditional updates, whereas the red segment denotes the independent RGO branch. Orange, blue, and green mark the incoming, post-Y-update, and outgoing states; red marks an RGO branch. In panel (c), $u = (x - y)/\sqrt{\eta}$ is the normalized conditional position and p is momentum. **(a)** Standard proximal Gibbs discards the previous Gaussian displacement. **(b)** The reflected framework first sets $Y^+ = 2X - Y$, then chooses between the accelerated conditional update M_{Y^+} and an independent RGO branch. **(c)** Two realizations of the accelerated branch.

Proposition 1 (Exactness of the framework)**.** *If M_y preserves q_y for every y, then Alg. 1 preserves Π_η. Its position marginal therefore has invariant law π.*

Proof. The reflection preserves Π_η by the argument above. At fixed $Y = y$, both M_y and the independent RGO draw preserve q_y, so their mixture preserves the conditional law. Composing the two updates therefore preserves Π_η, whose position marginal is π. □

3.3 A One-Dimensional Picture of Acceleration

Consider the quadratic potential $V(x) = \alpha x^2/2$. The target has standard deviation $1/\sqrt{\alpha}$, whereas one local proximal update moves on the scale $\sqrt{\eta}$. In the standard proximal sampler, each Gaussian redraw forgets the previous direction, so the signed increments are centered and approximately uncorrelated. As a result, the typical net displacement after n steps is only $\sqrt{n\eta}$. Reaching the target scale requires roughly $1/(\alpha\eta)$ iterations. Reflection keeps the previous displacement available to the accelerated branch. If this branch makes several successive moves point in a correlated direction, their signed increments add approximately linearly rather than cancelling as a random walk. Achieving net displacement $1/\sqrt{\alpha}$ then takes roughly $1/\sqrt{\alpha\eta}$ iterations. The occasional RGO branch prevents the chain from remaining on a periodic trajectory without erasing its direction at every step. In the next two sections, we prove distributional contraction for the two resulting chains. For both of our proposed realizations, one iteration uses $\widetilde{O}(1 + \beta\eta\sqrt{d})$ oracle queries. Multiplying this cost by $\widetilde{O}((\alpha\eta)^{-1/2})$ iterations and choosing an optimal $\eta = \widetilde{\Theta}((\beta\sqrt{d})^{-1})$ gives $\widetilde{O}(\sqrt{\kappa}d^{1/4})$ oracle queries.

Extension with persistent momentum. Alg. 1 deliberately keeps only (X, Y) in the Markov-chain state for simplicity. The flux half-turn (Sec. 4) is a proper instance: it may use temporary momentum inside M_y, but discards that momentum after the conditional update. One can instead augment the state with momentum P, require the accelerated branch to preserve $q_y \otimes \gamma_d$, and redraw both X and P on the RGO branch. This is used by the persistent quarter-turn (Sec. 5).

4 Flux-Resampled Underdamped Half-Turn

We first instantiate the accelerated branch of Alg. 1 without retaining momentum between outer iterations. At the reflected value $Y_{k+1} = y$, the update starts from $X_k = x$, draws a Gaussian momentum, runs an event-driven underdamped half-turn, and returns only its terminal position. At a high level, it freezes the nonlinear force at an independent conditional reference, follows the resulting exactly solvable affine motion, and uses random events to correct the residual force.

4.1 The Flux-Resampled Update

We present the flux-resampled accelerated branch, $X_{k+1} \sim F_{Y_{k+1}}(X_k, \cdot)$, in Alg. 2. We choose a fixed numerical constant $\gamma_0 \in (0, 1]$, and set $\gamma = \gamma_0\sqrt{\frac{\alpha\beta\eta^2}{1+\alpha\eta}}$, $\omega = \sqrt{1 - \frac{\gamma^2}{4}}$, and $T_\gamma = \pi/\omega$. The

condition $\beta\eta \le 1/2$ used below (Prop. 2) ensures $\gamma < 2$. (Z_t, P_t) denotes the within-row position–momentum trajectory. An *event* is an instantaneous momentum update: when it happens at time t, Z_t remains continuous, while the momentum changes immediately from P_{t-} to P_t.[1] All driving random sources used by one call are fresh, i.e., conditional on (x, y), the reference draw, initial momentum, Brownian motion, driving Poisson clock, and collision marks are mutually independent.

We remark that the rate λ_t defines the ideal event clock; the implementable algorithm uses capped Poisson thinning and evaluates $\nabla V(Z_t)$ only at finitely many proposed event times (App. B.4.1). Its cap error and query cost are included in Thm. 2. Notice that the RGO branch uses its conditional draw as the new position, whereas the accelerated branch uses an independent RGO draw only as the reference $\widetilde{X}$. Thus, this realization makes one RGO call in every outer iteration.

Algorithm 2: Flux-resampled accelerated branch F_y

Require: $X_k = x$, $Y_{k+1} = y$, and proximal scale η

1: Draw $\widetilde{X} \sim q_y$ independently of x, draw $P_0 \sim \gamma_d$, and set $Z_0 \leftarrow x$
2: Until $t = T_\gamma$, evolve between events according to $\mathrm{d}Z_t = \sqrt{\eta} P_t \,\mathrm{d}t$ and $\mathrm{d}P_t = -\{(Z_t - y)/\sqrt{\eta} + \sqrt{\eta}\nabla V(\widetilde{X})\}\,\mathrm{d}t - \gamma P_t\,\mathrm{d}t + \sqrt{2\gamma}\,\mathrm{d}B_t$
3: Set residual gradient $h_t \leftarrow \nabla V(Z_t) - \nabla V(\widetilde{X})$ and generate events at rate $\lambda_t \leftarrow \sqrt{\eta}[P_{t-}^\top h_t]_+$
4: At each event, set $n_t \leftarrow h_t/\|h_t\|$, draw $R_t > 0$ with density $r \mapsto r\mathrm{e}^{-r^2/2}$, and update $P_t \leftarrow P_{t-} - (P_{t-}^\top n_t + R_t)n_t$
5: **return** $X_{k+1} \leftarrow Z_{T_\gamma} \sim F_y(x, \cdot)$

4.2 Why the Update is a Half-Turn

We now provide intuition for the design of Alg. 2. At fixed y, the gradient of the conditional potential splits as

$$-\nabla_x \log q_y(x) = \tfrac{x-y}{\eta} + \nabla V(x) = \big\{\tfrac{x-y}{\eta} + \nabla V(\widetilde{X})\big\} + \big\{\nabla V(x) - \nabla V(\widetilde{X})\big\}.$$

The first term is a harmonic force centered at $x_{\mathrm{ctr}} = y - \eta\nabla V(\widetilde{X})$; the second is the residual gradient. If the frozen gradient $\nabla V(\widetilde{X})$ were replaced by $\nabla V(Z_t)$, the diffusion in Alg. 2 would be ULD for the conditional law q_y. Freezing it makes the between-event SDE affine and exactly solvable, while the events correct the omitted residual gradient. Drawing $\widetilde{X} \sim q_y$ keeps this residual small on average because both the stationary position and reference lie in the same localized distribution.

The event intensity λ_t is the positive inner product of the velocity and the residual gradient. Hence events occur only while the motion increases the corresponding residual potential. Let $V_{\mathrm{res}} = V(x) - V(\widetilde{X}) - \nabla V(\widetilde{X})^\top(x - \widetilde{X})$ be the Taylor remainder omitted by the frozen affine force. We have $\frac{\mathrm{d}}{\mathrm{d}t} V_{\mathrm{res}}(Z_t) = \sqrt{\eta} P_t^\top h_t$, so λ_t is precisely the positive rate at which Z_t climbs the omitted potential. The algorithm samples this clock by proposing times from a capped dominating Poisson process. Under stationarity, observing the process at event times rate-weights a Gaussian normal velocity into the Rayleigh density $r\mathrm{e}^{-r^2/2}$; replacing the incoming normal component by an independent draw from this flux law in the opposite direction balances forward and reversed paths (Vanetti et al., 2017; Michel et al., 2020). This makes the conditional update exact.

Without residual events, the centered position $Z_t - x_{\mathrm{ctr}}$ follows a damped harmonic oscillator, and $T_\gamma = \pi/\omega$ is its half-period. At this time the coefficient of P_0 is zero, while $Z_0 - x_{\mathrm{ctr}}$ is multiplied by $-\mathrm{e}^{-\gamma T_\gamma/2}$. The choice of γ balances the resulting damping error with the residual-gradient error; the weighted estimate below makes this statement precise.

Proposition 2 (Conditional invariance and reversibility)**.** *Suppose $\beta\eta \le 1/2$. Conditional on the reflected auxiliary value $Y_{k+1} = y$, if the input position satisfies $X_k \mid Y_{k+1} = y \sim q_y$, then the output satisfies $X_{k+1} \mid Y_{k+1} = y \sim q_y$ as well. Moreover, F_y satisfies detailed balance with respect to q_y, i.e., $q_y(x)F_y(x, x') = q_y(x')F_y(x', x)$. Consequently, Alg. 1 with accelerated branch F_y (Alg. 2) preserves Π_η, whose position marginal is π.*

For the fixed-reference update, let $z^- = P_{t-}^\top n_t > 0$ denote the momentum component along n_t immediately before a collision. Under stationary Gaussian momentum this component is standard normal, but the event rate observes its positive values in proportion to their size. Its event-time

[1] By default, the process is always assumed to be càdlàg (right-continuous with left limits).

density is therefore $r\mathrm{e}^{-r^2/2}$. The update draws R_t from this Rayleigh density and sets the outgoing component to $z^+ = P_t^\top n_t = -R_t < 0$, leaving the orthogonal momentum unchanged. Between collisions, the affine Langevin motion handles the frozen part of the gradient, while this collision rule corrects the residual part. Together they balance a stationary path with its momentum-reversed path run backward. Averaging over the independent reference and discarding the temporary momentum gives the stated position transition. The proof that only finitely many collisions occur in finite time and the complete path-reversal argument are in App. B.1.

4.3 Convergence and Oracle Complexity

The local estimate that drives the outer-chain proof is most transparent in the normalized coordinate $W = (X - y)/\sqrt{\eta}$ for $X \sim q_y$. Let $\bar{w}_y$ and Σ_y be its mean and covariance, and write q_y^w for its law. If W^+ is the normalized output of F_y, define $(F_y^w f)(w) = \mathbb{E}[f(W^+) \mid W = w]$. For $u \in \mathbb{R}^d$, define the centered linear function $\ell_u(w) := u^\top(w - \bar{w}_y)$. An exact half-turn would satisfy $(F_y^w \ell_u)(w) = -\ell_u(w)$; the next result bounds the mean-square error in this identity. In other words, its left-hand side measures, under a stationary incoming position, how far the expected outgoing coordinate in direction u is from the exact sign reversal of the incoming coordinate.

Proposition 3 (Weighted half-turn estimate). *There is a universal $C > 0$ such that, if $\beta\eta \le 1/2$, then for every $y, u \in \mathbb{R}^d$,*

$$\|\ell_u + F_y^w \ell_u\|^2_{L^2(q_y^w)} \le C\beta\eta\, u^\top(I_d - \Sigma_y)u.$$

The matrix $I_d - \Sigma_y$ is the conditional covariance deficit relative to the Gaussian reference and, simultaneously, measures how strongly each direction is transferred by the outer reflection. Strong convexity ensures that every eigenvalue of Σ_y is at most one, so this weight is nonnegative. This matching weight is the key new estimate: the half-turn needs to be accurate only in directions that affect the complete chain. Its proof solves the affine motion exactly and bounds the event forcing by conditional covariance identities; see App. B.2.

Let $K_{\eta,\rho}$ be one complete reflected iteration with accelerated branch F_y. For an initial law ν, $\nu K_{\eta,\rho}^m$ denotes the law after m successive iterations. Let $\rho_\star > 0$ be the fixed universal numerical constant furnished by the theorem below, and set

$$\rho = \rho_\star \sqrt{\alpha\eta \log\left(\tfrac{\mathrm{e}}{\alpha\eta}\right)}. \tag{2}$$

Theorem 1 (Ideal flux contraction). *There exist universal constants $c_0, c, C > 0$ and a fixed $\rho_\star \in (0,1)$, independent of all problem parameters, such that, if $\beta\eta \le c_0$, then every integer $n_{\mathrm{blk}} \ge C(\alpha\eta)^{-1/2}\sqrt{\log(\mathrm{e}/(\alpha\eta))}$ satisfies, for every $\nu \ll \Pi_\eta$ and every nonnegative integer J,*

$$\chi^2(\nu K_{\eta,\rho}^{Jn_{\mathrm{blk}}} \| \Pi_\eta) \le \mathrm{e}^{-J}\chi^2(\nu\|\Pi_\eta). \tag{3}$$

The proof separates variation across Y from the remaining conditional variation in X. Reflection transfers the first component into the second, the weighted half-turn reverses the transferred component, and the RGO branch damps what remains. A modified L^2 energy then contracts over blocks of $n_{\mathrm{blk}} = \widetilde{O}((\alpha\eta)^{-1/2})$ iterations. The complete operator argument is in App. B.3.

Theorem 2 (Implementable flux half-turn). *Under Assump. 1, let $0 < \varepsilon < 1/4$, $\Delta \ge 1$, and let ν_0 be an initial position law satisfying $D_2(\nu_0\|\pi) \le \Delta$. There exist universal numerical constants $c, C > 0$, independent of $d, \kappa, \Delta, \varepsilon$, such that the following choice is valid:*

$$L = \Delta + \log\tfrac{Cd\kappa}{\varepsilon}, \qquad \eta = \tfrac{c}{\beta(\sqrt{dL}+L)}. \tag{4}$$

Initialize $X_0 \sim \nu_0$ and $Y_0 \mid X_0 \sim \mathcal{N}(X_0, \eta I_d)$. Then run Alg. 1 with the accelerated branch in Alg. 2, using the approximate RGO and capped Poisson clock specified in App. B.4. Its output distribution $\widehat{\nu}$ satisfies $\mathrm{TV}(\widehat{\nu}, \pi) \le \varepsilon$. Writing N_{qry} for its total number of oracle queries,

$$\mathbb{E}N_{\mathrm{qry}} \le C\sqrt{\kappa}\,(d+L)^{1/4}L^{3/4}\left(\Delta + \log\tfrac{4}{\varepsilon}\right)\log(\mathrm{e}\kappa L). \tag{5}$$

In particular, a constant-D_2 warm start requires $\widetilde{O}(\sqrt{\kappa}d^{1/4})$ expected oracle queries.

The implementable algorithm replaces every exact q_y draw by an approximate RGO, uses capped Poisson thinning for collisions, and caps the optimization loop used to prepare the RGO. Tail bounds control the probability that any cap is reached during the run. Coupling this algorithm to the ideal chain until their first disagreement gives only polylogarithmic dependence on $1/\varepsilon$; the complete error and query ledger is in App. B.4.

5 Persistent-Momentum Quarter-Turn

The flux method obtains the accelerated rate with temporary momentum, but still makes an independent RGO call in every iteration. We now retain momentum across iterations so that the accelerated branch passes both its terminal position and momentum to the next iteration. The analysis first uses an ideal conditional reference; the implementable algorithm replaces that per-iteration reference by a Gaussian construction and therefore calls an RGO only on rare refresh steps.

Algorithm 3: Persistent-momentum accelerated branch G_y^ζ

Require: Incoming pair (x, p), fixed auxiliary value y, scale η, and reference law ζ_y

1: Draw $E \sim \zeta_y$ independently of (x, p), and set $(Z_0, Q_0) \leftarrow (x, p)$
2: Between collisions, evolve $\dot{Z}_t = \sqrt{\eta} Q_t$ and $\dot{Q}_t = -(Z_t - y)/\sqrt{\eta} - \sqrt{\eta}\nabla V(E)$
3: Set gradient residual $h_t \leftarrow \nabla V(Z_t) - \nabla V(E)$ and collision intensity $\lambda_t \leftarrow \sqrt{\eta}[Q_{t-}^\top h_t]_+$
4: At a collision, set $n_t \leftarrow h_t/\|h_t\|$ and update $Q_t \leftarrow Q_{t-} - 2(Q_{t-}^\top n_t) n_t$
5: At time $\pi/2$, set $(X^+, P^+) \leftarrow (Z_{\pi/2}, Q_{\pi/2})$
6: **return** $(X^+, P^+) \sim G_y^\zeta((x, p), \cdot)$

5.1 The Persistent Update

At fixed $Y = y$, let ζ_y be a reference law that depends only on y. The ideal comparison chain takes $\zeta_y = q_y$; the implementable algorithm uses the Gaussian reference described below. For an incoming position–momentum pair (x, p), Alg. 3 defines this branch as the phase-space kernel $G_y^\zeta((x, p), \cdot)$. The complete iteration first reflects $Y_{k+1} = 2X_k - Y_k$, draws $C_k \sim \mathrm{Bernoulli}(\rho)$, and then samples

$$(X_{k+1}, P_{k+1}) \sim \begin{cases} q_{Y_{k+1}} \otimes \gamma_d, & C_k = 1, \\ G_{Y_{k+1}}^\zeta((X_k, P_k), \cdot), & C_k = 0. \end{cases}$$

In Alg. 3, (Z_t, Q_t) is the local position–momentum path. As in Alg. 2, we call an instantaneous momentum update $Q_{t-} \mapsto Q_t$ a *collision*, and assume all driving random sources are fresh and conditionally independent given (x, p, y). The intensity λ_t defines the ideal clock, and the implementable algorithm uses capped Poisson thinning, so that apart from one evaluation of $\nabla V(E)$, it evaluates $\nabla V(Z_t)$ only at proposed collision times (Lem. 20).

5.2 Why a Quarter-Turn?

For a fixed reference E, the affine motion is centered at $x_{\mathrm{ctr}} = y - \eta\nabla V(E)$. If no collisions occur, the normalized position $U_t = (Z_t - x_{\mathrm{ctr}})/\sqrt{\eta}$ satisfies $\dot{U} = Q$ and $\dot{Q} = -U$. One quarter-period therefore gives $(U, Q) \mapsto (Q, -U)$. A half-turn would map (U, Q) to $(-U, -Q)$, keeping position and momentum in separate coordinates. The quarter-turn instead exchanges them, so centered position information is stored in momentum and remains available to later iterations. Collisions correct the residual gradient, while occasional RGO steps damp correlations retained by this nearly Hamiltonian motion. Unlike the flux method, this update uses specular reflections and carries its terminal momentum into the next iteration.

Proposition 4 (Conditional invariance)**.** *For every fixed y and reference E, the accelerated-branch motion preserves $q_y \otimes \gamma_d$. Consequently, G_y^ζ preserves $q_y \otimes \gamma_d$ whenever $E \sim \zeta_y$ is independent of the incoming position and momentum. The complete persistent iteration above therefore preserves $\widehat{\Pi}_\eta := \Pi_\eta \otimes \gamma_d$.*

The affine motion preserves its harmonic energy, while the collision intensity and momentum reflection correct the residual gradient. Reversing momentum pairs every stationary path with a path of the same probability run backward. Energy conservation also prevents infinitely many collisions in finite time; the complete argument is in App. C.1.

The following estimate quantifies the approximate rotation. For $X \sim q_y$, let $W = (X - y)/\sqrt{\eta}$, with mean $\bar{w}_y$, covariance Σ_y, and law q_y^w. Define $\ell_{a,y}(w) = a^\top(w - \bar{w}_y)$ and $m_a(p) = a^\top p$, and let (W^+, Q^+) be the output of the accelerated branch at fixed y using an independent reference $E \sim q_y$. Its observable operator is $(S_y f)(w, p) = \mathbb{E}[f(W^+, Q^+) \mid W = w, Q = p]$. All norms in the next proposition are L^2 norms under $q_y^w \otimes \gamma_d$. For an exact quarter-turn, $S_y \ell_{a,y} = m_a$ and $S_y m_a = -\ell_{a,y}$, so the left side below vanishes; it measures the mean-square failure of these two identities.

Proposition 5 (Approximate quarter-turn). *If $\beta\eta \le 1/2$, then for every $y, a \in \mathbb{R}^d$,*

$$\|S_y \ell_{a,y} - m_a\|_2^2 + \|S_y m_a + \ell_{a,y}\|_2^2 \le \tfrac{22\beta\eta}{1-\beta\eta} a^\top (I_d - \Sigma_y) a.$$

Thus the nonlinear update retains the exact quarter-rotation structure up to an error weighted by the same conditional covariance deficit that couples the position to the outer reflection. The proof is in Lem. 13.

5.3 Convergence and Implementation

Let K_{per} be the observable operator of the ideal chain with $\zeta_y = q_y$. For a distribution μ, let $L_0^2(\mu)$ denote its square-integrable mean-zero functions. Set $g = \sqrt{\alpha\eta}$, $\ell_g = 1 + \log(1/g)$, and choose the RGO probability $\rho = c_\rho g \sqrt{\ell_g}$, where $c_\rho > 0$ is the fixed universal constant supplied by the contraction proof.

Theorem 3 (Persistent-quarter contraction). *For the choice of ρ above, there are universal $c_0, c, C > 0$ such that, if $0 < \eta \le c_0/\beta$ and $\rho \le 1/4$, then*

$$\|K_{\mathrm{per}}^n\|_{L_0^2(\widehat{\Pi}_\eta) \to L_0^2(\widehat{\Pi}_\eta)} \le 2\mathrm{e}^{-cng/\sqrt{\ell_g}}. \tag{6}$$

Let ν_0 be an initial position law with $D_2(\nu_0 \| \pi) \le \Delta$, initialize $Y_0 \mid X_0 \sim \mathcal{N}(X_0, \eta I_d)$ and $P_0 \sim \gamma_d$ independently, and let ν_n be the position law after n iterations. For $0 < \varepsilon \le 1/2$, $\mathrm{TV}(\nu_n, \pi) \le \varepsilon$ whenever

$$n \ge \tfrac{C{\ell_g}^{1/2}}{g} \left(\Delta + \log \tfrac{4}{\varepsilon}\right). \tag{7}$$

Since $g = \sqrt{\alpha\eta}$, the theorem gives a constant-factor reduction every $\widetilde{O}((\alpha\eta)^{-1/2})$ iterations; iterating it from the stated warm start yields (7). The proof follows how an observable can be stored in Y, conditional position, and momentum. Reflection creates position dependence, the quarter-turn rotates it into momentum and back, and the RGO branch removes the remaining conditional dependence. The estimate above controls the nonlinear error; a three-coordinate resolvent argument converts this transfer mechanism into (6). Details are in Apps. C.2 and C.3.

Implementable algorithm with rare RGO calls. Conditional invariance holds for every fixed reference, so the accelerated branch does not require a draw from q_y. A short gradient-descent solve of the conditional potential produces a certified point $\widehat{x}_y$, and the implementable algorithm takes $E = y - \eta \nabla V(\widehat{x}_y) + \sqrt{\eta} G$, where $G \sim \mathcal{N}(0, I_d)$. This reference depends only on y and fresh Gaussian randomness, so Prop. 4 gives exact conditional invariance; its closeness to the ideal reference is needed only for contraction. At the chosen scale, the certified solve and an approximate RGO both cost $\widetilde{O}(1)$ oracle queries; the gain is fewer RGO calls, used only on the RGO branch, rather than a lower per-step query order.

Theorem 4 (Implementable persistent sampler). *Under Assump. 1, let ν_0 satisfy $D_2(\nu_0 \| \pi) \le \Delta$, where $\Delta \ge 0$, and let $0 < \varepsilon \le 1/2$. There exist fixed universal constants $c, C > 0$ so that the following holds. Define $\Lambda_\varepsilon = \Delta + \log(Cd\kappa/\varepsilon)$ and $\eta = c/\{\beta(\sqrt{d\Lambda_\varepsilon} + \Lambda_\varepsilon)\}$. Initialize $X_0 \sim \nu_0$, $Y_0 \mid X_0 \sim \mathcal{N}(X_0, \eta I_d)$, and $P_0 \sim \gamma_d$, independently conditional on X_0. Run the implementable algorithm specified in App. C.4; its output law $\widehat{\nu}$ satisfies $\mathrm{TV}(\widehat{\nu}, \pi) \le \varepsilon$. If N_{qry} counts all oracle queries and N_{RGO} counts approximate-RGO calls, then*

$$\begin{aligned} \mathbb{E}N_{\mathrm{qry}} &\le \widetilde{O}(\sqrt{\kappa}\,(\Delta + \log(8/\varepsilon))(\sqrt{d\Lambda_\varepsilon} + \Lambda_\varepsilon)^{1/2}), \\ \mathbb{E}N_{\mathrm{RGO}} &\le \widetilde{O}(\Delta + \log(8/\varepsilon)). \end{aligned} \tag{8}$$

In particular, a constant-D_2 warm start gives $\widetilde{O}(\sqrt{\kappa} d^{1/4})$ expected oracle queries and only $\mathrm{polylog}(d, \kappa, 1/\varepsilon)$ expected RGO calls.

The ideal horizon is $\widetilde{O}(g^{-1}\{\Delta + \log(1/\varepsilon)\})$, whereas $\rho = \widetilde{O}(g)$; their product explains the polylogarithmic RGO count. Shifted-reference comparison, capped Poisson thinning, and maximal coupling of the approximate RGO give the finite error and query ledger in App. C.4. An optional quadratic-baseline refinement is stated and proved in Cor. 3.

6 Removing the Supplied Warm Start

The rates in Secs. 4 and 5 assume a constant-D_2 warm start. Our cold-start pipeline has three stages: localization constructs a law with finite D_2 divergence, a coarse proximal chain reduces this divergence to a constant, and either reflected sampler performs the final high-accuracy sampling. When a finite bound $D_2(\nu_0\|\pi) \le B$ is available, we skip localization and choose between direct reflected sampling and the coarse warm-up route. All stages used are included in the query count.

At coarse scale $u = r/\beta$, the exact comparison chain becomes constant-warm after $\widetilde{O}((\kappa/r)\log(\mathrm{e}B))$ proximal Gibbs iterations. At this scale, each conditional has condition number at most $1 + r$ (Chen et al., 2022; Kook & Vempala, 2025). Thus increasing r shortens the coarse chain but makes each conditional draw harder. We optimize this tradeoff over two implementable conditional samplers: smoothed Picard HMC followed by reflected finalization, or exact diffusion (Chen et al., 2026c;d;a). The former is cheaper in the intermediate- and high-dimensional regimes, whereas the latter gives the low-dimensional endpoint. Optimizing the scale and selecting the cheaper implementable conditional sampler produces three regimes; their boundaries occur when the unconstrained optimizer reaches an endpoint of $1 \le r \le \kappa$. The resulting cost is

$$\Psi(\kappa, d) := \begin{cases} \kappa^{2/3} d^{1/3}, & d \le \kappa^{16/7}, \\ \kappa d^{3/16}, & \kappa^{16/7} < d < \kappa^{8}, \\ \sqrt{\kappa} d^{1/4}, & d \ge \kappa^{8}. \end{cases}$$

Theorem 5 (Warm-start removal)**.** Finite-divergence input. *Suppose a sample from an initial law ν_0 is available and the known number $B \ge 1$ satisfies $D_2(\nu_0\|\pi) \le B$. For every $0 < \varepsilon < 1/4$, the wrapper using either reflected sampler as its final stage returns an output law $\widehat{\nu}$ with $\mathrm{TV}(\widehat{\nu}, \pi) \le \varepsilon$. If N_{qry} is the total number of oracle queries, then*

$$\mathbb{E}N_{\mathrm{qry}} \le \widetilde{O}(\min\{\sqrt{\kappa}B(\sqrt{dB} + B)^{1/2}, \Psi(\kappa, d)\log(\mathrm{e}B)\}). \tag{9}$$

If B is polynomially bounded in $d, \kappa, 1/\varepsilon$, this is at most $\widetilde{O}(\kappa^{6/7} d^{1/4})$.

Supplied point. *Given $x_0 \in \mathbb{R}^d$ and $0 < \varepsilon < 1/4$, let x_* be the minimizer of V, define $\log_+(t) = \max\{1, \log t\}$, and set $Q_{\mathrm{loc}} = \sqrt{\kappa}\log_+[\kappa\{2 + \sqrt{\beta}\|x_0 - x_*\|\}]$. An implementable algorithm returns an output law $\widehat{\nu}$ with $\mathrm{TV}(\widehat{\nu}, \pi) \le \varepsilon$ and*

$$\mathbb{E}N_{\mathrm{qry}} \le \widetilde{O}(Q_{\mathrm{loc}} + \Psi(\kappa, d)) \le \widetilde{O}(Q_{\mathrm{loc}} + \kappa^{6/7} d^{1/4}). \tag{10}$$

The first term in (9) applies a reflected sampler directly to ν_0; the second first pays for the coarse warm-up. The proof runs the exact coarse chain only as a comparison process: its endpoint is constant-D_2-warm, while the algorithmic endpoint is close to it in total variation. Applying the same capped reflected finalizer to both endpoints transfers the convergence guarantee by data processing.

For a supplied point, accelerated optimization stops at the observable certificate $\|\nabla V(x_{\mathrm{ref}})\| \le \alpha/\sqrt{\beta}$. One draw from $\mathcal{N}(x_{\mathrm{ref}}, (2\beta)^{-1} I_d)$ then has $D_2 \le 1 + \frac{d}{2}\log(2\kappa)$, so the finite-divergence result applies. The algorithm does not use x_*; it appears only in the localization bound. The regime optimization, error handoff, and localization proof are given in App. D.

7 Conclusion

Reflected proximal dynamics accelerate the outer chain by combining exact Gaussian reflection, conditional kinetic transport, and controlled refreshment. From a constant-D_2 warm start, both realizations use $\widetilde{O}(\sqrt{\kappa} d^{1/4})$ expected oracle queries, while the persistent construction makes only a polylogarithmic expected number of approximate-RGO calls. Including localization and warm-up yields the fully charged $\widetilde{O}(\kappa^{6/7} d^{1/4})$ bound. Extending the macroscopic estimates from strong convexity to functional inequalities is a natural direction beyond the smooth strongly log-concave setting.

AI USE STATEMENT

The idea of accelerating the proximal sampler using underdamped dynamics originated with the authors. We used generative AI tools to brainstorm implementation details and algorithmic variants, assist in drafting and revising mathematical proofs. AI-generated suggestions were treated as unverified drafts rather than as authoritative sources. The authors reviewed and revised all AI-assisted material, checked the final theorem statements and proofs, and verified the citations used in the paper against the referenced works. We did not use generative AI to produce data or empirical results; no such results are reported in this paper. We take responsibility for the final content of this work, including all text, claims, and artifacts produced with the aid of generative AI.

REFERENCES

Jason M. Altschuler and Sinho Chewi. Faster high-accuracy log-concave sampling via algorithmic warm starts. *J. ACM*, 71(3), June 2024. ISSN 0004-5411. doi: 10.1145/3653446. URL https://doi.org/10.1145/3653446.

Michael Betancourt and Mark Girolami. Hamiltonian Monte Carlo for hierarchical models. In *Current trends in Bayesian methodology with applications*, pp. 119–142. Chapman and Hall/CRC, 2019.

Joris Bierkens, Sebastiano Grazzi, Kengo Kamatani, and Gareth Roberts. The boomerang sampler. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 908–918. PMLR, 2020.

Alexandre Bouchard-Côté, Sebastian J. Vollmer, and Arnaud Doucet. The bouncy particle sampler: A nonreversible rejection-free markov chain monte carlo method. *Journal of the American Statistical Association*, 113(522):855–867, 2018. doi: 10.1080/01621459.2017.1294075.

Fan Chen, Sinho Chewi, Constantinos Daskalakis, and Alexander Rakhlin. High-accuracy sampling for diffusion models and log-concave distributions. *arXiv preprint arXiv:2602.01338*, 2026a.

Fan Chen, Sinho Chewi, Jianfeng Lu, and Matthew S Zhang. Accelerated high-accuracy sampling from a warm start via the proximal bouncy particle sampler. *arXiv preprint arXiv:2609.06905*, 2026b.

Fan Chen, Sinho Chewi, Jianfeng Lu, and Matthew S Zhang. Smoothed Picard Hamiltonian Monte Carlo. *arXiv preprint arXiv:2609.06906*, 2026c.

Fan Chen, Sinho Chewi, Alexander Rakhlin, and Matthew S Zhang. Exact simulation of diffusions and improved algorithms for log-concave sampling. *arXiv preprint arXiv:2608.05022*, 2026d.

Yongxin Chen, Sinho Chewi, Adil Salim, and Andre Wibisono. Improved analysis for a proximal algorithm for sampling. In Po-Ling Loh and Maxim Raginsky (eds.), *Proceedings of Thirty Fifth Conference on Learning Theory*, volume 178 of *Proceedings of Machine Learning Research*, pp. 2984–3014. PMLR, 02–05 Jul 2022. URL https://proceedings.mlr.press/v178/chen22c.html.

Xiang Cheng, Niladri S. Chatterji, Peter L. Bartlett, and Michael I. Jordan. Underdamped Langevin MCMC: A non-asymptotic analysis. In Sébastien Bubeck, Vianney Perchet, and Philippe Rigollet (eds.), *Proceedings of the 31st Conference On Learning Theory*, volume 75 of *Proceedings of Machine Learning Research*, pp. 300–323. PMLR, 06–09 Jul 2018. URL https://proceedings.mlr.press/v75/cheng18a.html.

Sinho Chewi. Log-concave sampling. *Book draft available at https://chewisinho. github. io*, 9:17–18, 2023.

Sinho Chewi, Murat A Erdogdu, Mufan Li, Ruoqi Shen, and Shunshi Zhang. Analysis of Langevin Monte Carlo from Poincaré to log-Sobolev. In Po-Ling Loh and Maxim Raginsky (eds.), *Proceedings of Thirty Fifth Conference on Learning Theory*, volume 178 of *Proceedings of Machine Learning Research*, pp. 1–2. PMLR, 02–05 Jul 2022. URL https://proceedings.mlr.press/v178/chewi22a.html.

Arnak S. Dalalyan. Theoretical guarantees for approximate sampling from smooth and log-concave densities. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 79(3):651–676, 06 2017. ISSN 1369-7412. doi: 10.1111/rssb.12183. URL https://doi.org/10.1111/rssb.12183.

Arnak S. Dalalyan and Lionel Riou-Durand. On sampling from a log-concave density using kinetic Langevin diffusions. *Bernoulli*, 26(3):1956 – 1988, 2020. doi: 10.3150/19-BEJ1178. URL https://doi.org/10.3150/19-BEJ1178.

Alain Durmus, Éric Moulines, and Marcelo Pereyra. Efficient Bayesian computation by proximal Markov chain Monte Carlo: When Langevin meets Moreau. *SIAM Journal on Imaging Sciences*, 11(1):473–506, 2018. doi: 10.1137/16M1108340. URL https://doi.org/10.1137/16M1108340.

Alain Durmus, Szymon Majewski, and Błażej Miasojedow. Analysis of Langevin Monte Carlo via convex optimization. *Journal of Machine Learning Research*, 20(73):1–46, 2019. URL http://jmlr.org/papers/v20/18-173.html.

Jiaojiao Fan, Bo Yuan, and Yongxin Chen. Improved dimension dependence of a proximal algorithm for sampling. 195:1473–1521, 12–15 Jul 2023. URL https://proceedings.mlr.press/v195/fan23a.html.

Andrew Gelman, John B Carlin, Hal S Stern, and Donald B Rubin. *Bayesian data analysis*. chapman and hall/cRc, 1995.

Subhashis Ghosal and Aad van der Vaart. *Fundamentals of Nonparametric Bayesian Inference*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2017.

Sivakanth Gopi, Yin Tat Lee, Daogao Liu, Ruoqi Shen, and Kevin Tian. Algorithmic aspects of the log-Laplace transform and a non-Euclidean proximal sampler. In Gergely Neu and Lorenzo Rosasco (eds.), *Proceedings of Thirty Sixth Conference on Learning Theory*, volume 195 of *Proceedings of Machine Learning Research*, pp. 2399–2439. PMLR, 12–15 Jul 2023. URL https://proceedings.mlr.press/v195/gopi23a.html.

Nicolaï Gouraud, Lucas Journel, and Pierre Monmarché. The velocity jump langevin process and its splitting scheme: Long time convergence and numerical accuracy. *The SMAI Journal of Computational Mathematics*, 11:533–584, 2025. doi: 10.5802/smai-jcm.134.

Ye He, Alireza Mousavi-Hosseini, Krishnakumar Balasubramanian, and Murat A. Erdogdu. A separation in heavy-tailed sampling: Gaussian vs. stable oracles for proximal samplers. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.), *Advances in Neural Information Processing Systems*, volume 37, pp. 65260–65296. Curran Associates, Inc., 2024. doi: 10.52202/079017-2084. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/77f33e8bd80345de4aea8554bbe5a4da-Paper-Conference.pdf.

Yunbum Kook and Santosh S Vempala. Zeroth-order logconcave sampling. *arXiv preprint arXiv:2507.18021*, 2025.

Yunbum Kook, Santosh S. Vempala, and Matthew S. Zhang. In-and-out: Algorithmic diffusion for sampling convex bodies. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.), *Advances in Neural Information Processing Systems*, volume 37, pp. 108354–108388. Curran Associates, Inc., 2024. doi: 10.52202/079017-3440. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/c4006ff54a7bbda74c09bad6f7586f5b-Paper-Conference.pdf.

Yin Tat Lee, Ruoqi Shen, and Kevin Tian. Structured logconcave sampling with a restricted Gaussian oracle. In Mikhail Belkin and Samory Kpotufe (eds.), *Proceedings of Thirty Fourth Conference on Learning Theory*, volume 134 of *Proceedings of Machine Learning Research*, pp. 2993–3050. PMLR, 15–19 Aug 2021. URL https://proceedings.mlr.press/v134/lee21a.html.

Jiaming Liang and Yongxin Chen. A proximal algorithm for sampling. *Transactions on Machine Learning Research*, 2023. ISSN 2835-8856. URL https://openreview.net/forum?id=CkXOwlhf27.

Jiaming Liang and Yongxin Chen. Proximal oracles for optimization and sampling. *Journal of Optimization Theory and Applications*, 208(110), 2026. doi: 10.1007/s10957-026-02937-z.

Yi-An Ma, Niladri S. Chatterji, Xiang Cheng, Nicolas Flammarion, Peter L. Bartlett, and Michael I. Jordan. Is there an analog of Nesterov acceleration for gradient-based MCMC? *Bernoulli*, 27 (3):1942 – 1992, 2021. doi: 10.3150/20-BEJ1297. URL https://doi.org/10.3150/20-BEJ1297.

Manon Michel, Alain Durmus, and Stéphane Sénécal. Forward event-chain monte carlo: Fast sampling by randomness control in irreversible markov chains. *Journal of Computational and Graphical Statistics*, 29(4):689–702, 2020. doi: 10.1080/10618600.2020.1750417.

Siddharth Mitra and Andre Wibisono. Fast convergence of ϕ-divergence along the unadjusted Langevin algorithm and proximal sampler. In Gautam Kamath and Po-Ling Loh (eds.), *Proceedings of The 36th International Conference on Algorithmic Learning Theory*, volume 272 of *Proceedings of Machine Learning Research*, pp. 846–869. PMLR, 24–27 Feb 2025. URL https://proceedings.mlr.press/v272/mitra25a.html.

Pierre Monmarché, Mathias Rousset, and Pierre-André Zitt. Exact targeting of Gibbs distributions using velocity-jump processes. *Stochastics and Partial Differential Equations: Analysis and Computations*, 11(3):908–947, 2023. doi: 10.1007/s40072-022-00247-9.

Mark E. J. Newman and Gerard T. Barkema. *Monte Carlo Methods in Statistical Physics*. Oxford University Press, 02 1999. ISBN 9780198517962. doi: 10.1093/oso/9780198517962.001.0001. URL https://doi.org/10.1093/oso/9780198517962.001.0001.

Adil Salim and Peter Richtarik. Primal dual interpretation of the proximal stochastic gradient Langevin algorithm. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 3786–3796. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/2779fda014fbadb761f67dd708c1325e-Paper.pdf.

Ruoqi Shen and Yin Tat Lee. The randomized midpoint method for log-concave sampling. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL https://proceedings.neurips.cc/paper_files/paper/2019/file/eb86d510361fc23b59f18c1bc9802cc6-Paper.pdf.

Paul Vanetti, Alexandre Bouchard-Côté, George Deligiannidis, and Arnaud Doucet. Piecewise-deterministic markov chain monte carlo. *arXiv preprint arXiv:1707.05296*, 2017.

Santosh Vempala and Andre Wibisono. Rapid convergence of the unadjusted Langevin algorithm: Isoperimetry suffices. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL https://proceedings.neurips.cc/paper_files/paper/2019/file/65a99bb7a3115fdede20da98b08a370f-Paper.pdf.

Andre Wibisono. Mixing time of the proximal sampler in relative Fisher information via strong data processing inequality. *IEEE Transactions on Information Theory*, 72(5):2667–2688, 2026. doi: 10.1109/TIT.2026.3670723.

Bo Yuan, Jiaojiao Fan, Jiaming Liang, Andre Wibisono, and Yongxin Chen. On a class of gibbs sampling over networks. In *Proceedings of the 36th Conference on Learning Theory*, volume 195 of *Proceedings of Machine Learning Research*, pp. 5754–5780. PMLR, 2023. URL https://proceedings.mlr.press/v195/yuan23a.html.

Shunshi Zhang, Sinho Chewi, Mufan Li, Krishna Balasubramanian, and Murat A. Erdogdu. Improved discretization analysis for underdamped Langevin Monte Carlo. In Gergely Neu and Lorenzo Rosasco (eds.), *Proceedings of Thirty Sixth Conference on Learning Theory*, volume 195 of

Proceedings of Machine Learning Research, pp. 36–71. PMLR, 12–15 Jul 2023. URL https://proceedings.mlr.press/v195/zhang23a.html.

A Related Work and Detailed Comparison

Proximal and kinetic samplers. Earlier proximal Langevin methods use Moreau regularization or proximal mappings to discretize overdamped Langevin dynamics (Durmus et al., 2018; Salim & Richtarik, 2020). The Gaussian proximal sampler instead separates global contraction from sampling a localized conditional (Lee et al., 2021; Chen et al., 2022). Approximate-RGO algorithms improve the local implementation (Fan et al., 2023; Chen et al., 2026a); our methods instead modify the outer dynamics to retain information across conditional updates. Kinetic Langevin samplers obtain acceleration by retaining velocity, with specialized discretizations needed to control the resulting finite-step error (Cheng et al., 2018; Dalalyan & Riou-Durand, 2020; Ma et al., 2021; Shen & Lee, 2019; Zhang et al., 2023). We use the same principle inside each localized conditional problem, while an exact Gaussian reflection transports memory between successive conditionals.

Event-driven kinetic updates. The bouncy particle sampler corrects deterministic motion through velocity events triggered by the positive directional derivative of the potential (Bouchard-Côté et al., 2018). Event-driven Hamiltonian samplers and the Boomerang sampler instead integrate an exactly solvable reference motion and use such events to correct the residual force (Vanetti et al., 2017; Bierkens et al., 2020). Gaussian flux-resampling of the force-aligned velocity component appears in the framework of Vanetti et al. (2017) and Forward Event-Chain Monte Carlo (Michel et al., 2020); related samplers combine Langevin motion with randomized velocity events (Monmarché et al., 2023; Gouraud et al., 2025). Our flux half-turn places these ingredients inside a reflected proximal iteration and selects the duration and damping to obtain the weighted reversal estimate used by the outer-chain analysis.

End-to-end complexity. Tab. 1 compares representative high-accuracy results under smooth strong log-concavity. We report the error metric in which each paper states its guarantee and the initialization that its bound requires. The exact-ULD/FORS row is not a proximal method; we include it as the closest global benchmark with a different tradeoff between d and κ.

Extensions of proximal sampling. The two-variable Gaussian augmentation has also been extended to bipartite networks, yielding distributed Gibbs samplers with non-asymptotic guarantees for structured log-concave targets (Yuan et al., 2023). The proximal viewpoint has also been developed beyond smooth, strongly log-concave Euclidean targets. Alternating-sampling and proximal-oracle implementations address nonsmooth, weakly smooth, and composite potentials (Liang & Chen, 2023; 2026). Non-Euclidean variants replace the Gaussian geometry by log-Laplace regularization (Gopi et al., 2023), while related algorithmic-diffusion constructions treat uniform sampling from convex bodies (Kook et al., 2024). For heavy-tailed targets, Gaussian proximal oracles can face an intrinsic high-accuracy barrier, motivating stable-oracle analogues (He et al., 2024). Recent strong-data-processing analyses also establish convergence of the proximal sampler in relative Fisher information (Wibisono, 2026). These works broaden the target, geometry, oracle, or convergence criterion; our methods instead retain the smooth strongly log-concave setting and modify the outer dynamics to carry directional memory.

The closest comparison is concurrent PBPS (Chen et al., 2026b). PBPS uses the same augmentation and outer reflection, but its conditional transition is a harmonic BPS half-turn with specular bounces and a fresh independent reference from q_y at every iteration. PBPS and both of our methods have $\widetilde{O}(\sqrt{\kappa}d^{1/4})$ warm-start oracle-query complexity. The flux method replaces specular reflection by Gaussian-flux resampling and adds underdamped forcing. The persistent method instead uses a quarter-turn, carries momentum between ordinary iterations, and calls a conditional-resampling subroutine only on rare refresh iterations.

B Proofs for the Flux-Resampled Half-Turn

This appendix proves Props. 2 and 3 and Thms. 1 and 2. The proof follows the algorithm:

Table 1: Representative high-accuracy bounds for smooth, strongly log-concave sampling. The table reports oracle-query counts in expectation unless marked otherwise. The notation $\widetilde{O}$ suppresses polylogarithmic factors in $d, \kappa, 1/\varepsilon$.

Paper	Method	Initialization	Metric	Complexity
Chen et al. (2022)	Gaussian proximal sampler with rejection RGO	polynomial KL/Rényi divergence	KL, χ^2, Rényi	$\widetilde{O}(\kappa d)$
Fan et al. (2023)	approximate-rejection proximal sampler	Gaussian cold start	TV, W_2, χ^2	$\widetilde{O}(\kappa d^{1/2})$
Altschuler & Chewi (2024)	ULMC warm-up with MALA/proximal reduction	known or located minimizer; no warm law	TV, KL, χ^2, W_2	$\widetilde{O}(\kappa d^{1/2})$ w.h.p.
Mitra & Wibisono (2025)	proximal sampler with a basic rejection compiler	polynomial Φ-divergence	Φ-divergence	$\widetilde{O}(\kappa d)$
Chen et al. (2026a)	proximal sampler with a FORS conditional sampler	finite-Rényi Gaussian start	fixed-order Rényi	$\widetilde{O}(\kappa d^{1/2})$
Chen et al. (2026d)	exact ULD/FORS	Gaussian with $D_{2q} = O(d \log \kappa)$, fixed q	fixed-order Rényi	$\widetilde{O}(\kappa^{2/3} d^{1/3})$
Chen et al. (2026c;b)	SPHMC warm-up with a PBPS finalizer	no supplied warm distribution	TV	$\widetilde{O}(\kappa^{7/6} d^{1/6} + \sqrt{\kappa} d^{1/4})$
This work	reflected proximal dynamics with optimized initialization	no supplied warm distribution	TV	$\widetilde{O}(\kappa^{6/7} d^{1/4})$

Table 2: Comparison of reflected proximal samplers from a constant-D_2 warm start.

	PBPS	Flux half-turn	Persistent quarter-turn
Conditional motion	harmonic BPS	underdamped, flux-resampled	fixed-reference bouncy
Momentum across outer steps	discarded	discarded	retained
Independent q_y reference	every step	every step	refresh steps only
Oracle queries	$\widetilde{O}(\sqrt{\kappa} d^{1/4})$	$\widetilde{O}(\sqrt{\kappa} d^{1/4})$	$\widetilde{O}(\sqrt{\kappa} d^{1/4})$
Conditional-resampling calls	$\widetilde{O}(\sqrt{\kappa} d^{1/4})$	$\widetilde{O}(\sqrt{\kappa} d^{1/4})$	$\mathrm{polylog}(d, \kappa, 1/\varepsilon)$

1. establish the geometry of the localized conditional law and show that the event-driven half-turn preserves it;
2. prove that one conditional half-turn approximately reverses every affine position coordinate that is relevant to the outer chain;
3. combine this estimate with the outer reflection and occasional conditional refreshment to obtain accelerated L^2 contraction;

4. replace exact conditional draws and unbounded event clocks by finite implementations, then account for total-variation error and oracle queries.

Standard ingredients are cited rather than reproved: Brascamp–Lieb and Poincaré inequalities for strongly log-concave measures, skew balance for velocity-jump processes, and the existing approximate-RGO implementation. The new parts are the covariance-weighted half-turn estimate and the modified-energy argument that turns it into contraction of the complete reflected proximal chain.

Throughout this appendix, set

$$a = \alpha\eta, \qquad s = \beta\eta, \qquad g_y(w) = \sqrt{\eta}\,\nabla V(y + \sqrt{\eta}w).$$

Thus $0 < a \le s$. Under the change of variables $w = (x-y)/\sqrt{\eta}$, the conditional density q_y becomes $q_y^w(w) \propto \exp\{-\|w\|^2/2 - V(y+\sqrt{\eta}w)\}$. We write I for the identity on the space at hand, use I_d when its matrix dimension is worth displaying, and set $\gamma_d = \mathcal{N}(0, I_d)$. Unless a measure is displayed explicitly, inner products and norms below are taken in the stationary L^2 space of the objects under discussion; gradients of functions of y are with respect to y. Operators act on observables; probability densities therefore evolve by their adjoints. The few operators used in the contraction proof are:

symbol	meaning
P	conditional expectation given Y
$\mathsf{Q} = I - \mathsf{P}$	remaining conditional variation
U	outer reflection $(x, y) \mapsto (x, 2x - y)$
F	conditional flux half-turn
K	one complete ideal iteration
$\mathsf{T} = \mathsf{PUP}$	reflection seen on Y-only observables
$\Gamma = (I - \mathsf{T}^2)^{1/2}$	amount transferred out of the Y-only part

B.1 Conditional Geometry and Validity of the Half-Turn

We first collect the conditional identities used throughout the proof and then verify that the event-driven motion is an exact conditional update. Neither step requires a classical Hessian: the identities hold in the weak sense for the strongly convex, smooth potential in Assump. 1.

Let $W = (X - y)/\sqrt{\eta}$ under q_y, and denote its mean and covariance by $\bar{w}_y$ and Σ_y. The following lemma says that the departure of this conditional law from a standard Gaussian is measured by the same matrix $I - \Sigma_y$ that controls the curvature of the Y-marginal.

Lemma 1 (Localized conditional geometry)**.** *For every y,*

$$\frac{1}{1+s} I \preceq \Sigma_y \preceq \frac{1}{1+a} I, \qquad \frac{a}{1+a} I \preceq I - \Sigma_y \preceq \frac{s}{1+s} I.$$

If $\Pi_\eta^Y(y) \propto \mathrm{e}^{-U_\eta(y)}$, then

$$\nabla U_\eta(y) = \eta^{-1}\{y - \mathbb{E}[X \mid Y = y]\}, \qquad \nabla^2 U_\eta(y) = \eta^{-1}(I - \Sigma_y),$$

and hence $\alpha(1+a)^{-1} I \preceq \nabla^2 U_\eta \preceq \beta(1+s)^{-1} I$. Moreover, when $s < 1$,

$$\mathbb{E} g_y(W) = -\bar{w}_y, \qquad \mathrm{Cov}(W, g_y(W)) = I - \Sigma_y, \qquad \mathrm{Cov}(g_y(W)) \preceq \frac{s}{1-s}(I - \Sigma_y). \tag{11}$$

Proof. The normalized conditional potential has curvature between $(1+a)I$ and $(1+s)I$. Brascamp–Lieb gives the upper covariance bound. The information inequality gives $\mathrm{Cov}(W) \succeq \{\mathbb{E}\nabla^2[-\log q_y^w(W)]\}^{-1} \succeq (1+s)^{-1} I$, proving the remaining covariance bounds. Differentiating the conditional normalizer gives the two identities for U_η.

For (11), use $\nabla \log q_y^w(w) = -w - g_y(w)$. Integration by parts first gives $\mathbb{E}(W + g_y(W)) = 0$ and $\Sigma_y + \mathrm{Cov}(W, g_y(W)) = I$. Applying it once more to g_y gives

$$\mathbb{E}\nabla g_y(W) = I - \Sigma_y + \mathrm{Cov}(g_y(W)).$$

The conditional law has Poincaré constant at most one, while $0 \preceq \nabla g_y \preceq sI$. Therefore $\mathrm{Cov}(g_y(W)) \preceq s\mathbb{E}\nabla g_y(W)$; rearrangement proves the last claim. These standard Brascamp–Lieb and integration-by-parts facts are reviewed, for example, by Chewi (2023). The extension from smooth potentials and test functions to the regularity in Assump. 1 is collected in App. B.4.2. □

We next verify the conditional process. Fix y and a reference $\widetilde{w}$. To lighten notation write $g = g_y$, $b_{\mathrm{ref}} = g(\widetilde{w})$, and $h(w) = g(w) - b_{\mathrm{ref}}$. In normalized coordinates, the generator of the diffusion between events is

$$\mathcal{L}_0 = p^\top \nabla_w - (w + b_{\mathrm{ref}})^\top \nabla_p + \gamma(\Delta_p - p^\top \nabla_p),$$

and events occur at rate $[p^\top h(w)]_+$. At an event, write $n = h/\|h\|$, decompose $p = v + zn$, draw $R > 0$ with density $r\mathrm{e}^{-r^2/2}$, and replace p by $v - Rn$. When $h = 0$, the event rate is zero and the direction n is never evaluated. More explicitly, if $\mathcal{J}$ denotes the event part of the generator, then

$$(\mathcal{J}f)(w,p) = [p^\top h(w)]_+ \int_0^\infty \left[f\big(w, p - \{p^\top n(w) + r\}n(w)\big) - f(w,p)\right] r\mathrm{e}^{-r^2/2}\,\mathrm{d}r,$$

with $\mathcal{J}f = 0$ when $h(w) = 0$. The full generator used below is $\mathcal{L} = \mathcal{L}_0 + \mathcal{J}$.

Proposition 6 (Fixed-reference validity)**.** *For every fixed $(y, \widetilde{w})$, the preceding process is non-explosive and preserves $q_y^w \otimes \gamma_d$. If $\mathsf{S}_t^{y,\widetilde{w}}$ is its transition operator and $(Jf)(w,p) = f(w,-p)$, then*

$$(\mathsf{S}_t^{y,\widetilde{w}})^* = J\mathsf{S}_t^{y,\widetilde{w}}J.$$

It is a strongly continuous contraction semigroup on $L^2(q_y^w \otimes \gamma_d)$. Consequently, inserting an independent Gaussian momentum, running for a fixed time, and discarding momentum gives a reversible position kernel for q_y^w. The conclusion remains true after averaging over any reference law independent of the incoming (W, P) conditional on $Y = y$.

Proof. The calculation follows the usual skew-balance argument for event-driven kinetic samplers (Vanetti et al., 2017; Bierkens et al., 2020; Michel et al., 2020; Monmarché et al., 2023); we give the semigroup step because our reference motion also contains Langevin noise.

The proof has three steps. First, a trajectory with no event can be reversed after flipping momentum; its no-event probability supplies exactly the likelihood correction caused by the residual force. Second, the Rayleigh update balances the stationary probability flow through a single collision. Finally, we sum these identities over the number and times of the collisions. This proves reversal of the complete process as well as invariance.

Non-explosion. With $\mathcal{V} = 1 + \{\|w + b_{\mathrm{ref}}\|^2 + \|p\|^2\}/2$, the full generator obeys $\mathcal{L}\mathcal{V} \le C_{\widetilde{w}}(1 + \mathcal{V})$, and the event rate is at most $C_{\widetilde{w}}(1 + \mathcal{V})$, where $C_{\widetilde{w}}$ is a finite constant that may depend on the fixed reference. Indeed, a jump changes $\mathcal{V}$ by $(R^2 - z^2)/2$, $\mathbb{E}R^2 = 2$, and $\|h(w)\| \le s\|w - \widetilde{w}\|$. A stopped Dynkin formula and Gronwall bound both $\mathbb{E}\mathcal{V}(W_t, P_t)$ and the expected event count on every finite interval. The recursive integrated-hazard construction is therefore well defined and non-explosive.

Reversal between and at events. Let $Z_t^0 = (W_t^0, P_t^0)$ be the affine Langevin diffusion without events, let $\lambda(w,p) = [p^\top h(w)]_+$, and define its killed semigroup—that is, its transition weighted by the probability of seeing no event—by

$$\mathsf{D}_t f(z) = \mathbb{E}_z\left[\mathrm{e}^{-\int_0^t \lambda(Z_u^0)\,\mathrm{d}u} f(Z_t^0)\right].$$

The affine diffusion is invariant and momentum-reversible for the Gaussian density $\mu_{\mathrm{ref}} \propto \exp\{-\|w + b_{\mathrm{ref}}\|^2/2 - \|p\|^2/2\}$. For the desired density $\mu = q_y^w \otimes \gamma_d$, the ratio $\varrho = \mathrm{d}\mu/\,\mathrm{d}\mu_{\mathrm{ref}}$ satisfies $\nabla \log \varrho = -h$. Since the position path is absolutely continuous,

$$\frac{\varrho(W_t^0)}{\varrho(W_0^0)} = \exp\left\{-\int_0^t P_u^{0\top} h(W_u^0)\,\mathrm{d}u\right\} = \exp\left\{-\int_0^t (\lambda - \lambda \circ J)(Z_u^0)\,\mathrm{d}u\right\}.$$

Stationary path reversal for the affine diffusion therefore gives $\mathsf{D}_t^* = J\mathsf{D}_t J$ first for bounded nonnegative functions. Taking one test function equal to one shows $\mu\mathsf{D}_t \le \mu$, and Jensen's inequality extends the identity to $L^2(\mu)$.

Let $\mathcal{G}$ be the gain kernel: it applies one collision, weighted by the rate at which that collision occurs. At fixed w, write the incoming momentum as $p = v + zn$, where $v \perp n$ and $z > 0$. Apart from a common constant, the stationary marked flux is

$$q_y^w(w)\|h(w)\|\varphi_{d-1}(v) z\mathrm{e}^{-z^2/2} r\mathrm{e}^{-r^2/2}\,\mathrm{d}w\,\mathrm{d}v\,\mathrm{d}z\,\mathrm{d}r, \qquad z, r > 0.$$

Here φ_{d-1} is the standard Gaussian density on the subspace orthogonal to n. It is unchanged by $(v, z, r) \mapsto (-v, r, z)$, which turns a marked transition into the momentum reversal of its time-reversed transition. Thus $\mathcal{G}^* = J\mathcal{G}J$ in the corresponding kernel sense.

The contribution of paths with exactly m events is the time-simplex integral of $\mathsf{D}_{t_0}\mathcal{G}\mathsf{D}_{t_1}\cdots\mathcal{G}\mathsf{D}_{t_m}$. Reversing the simplex coordinates and using the preceding two identities shows that its adjoint is obtained by conjugating with J. Tonelli justifies this calculation for nonnegative tests, and non-explosion makes the sum over m conservative, meaning that it retains total probability one. Summing proves $(\mathsf{S}_t^{y,\widetilde{w}})^* = J\mathsf{S}_t^{y,\widetilde{w}}J$; taking one test function equal to one proves invariance. Jensen gives L^2 contractivity, while stochastic continuity and density of bounded continuous compactly supported functions give strong continuity.

Finally, if $\iota f(w,p) = f(w)$ and $\mathsf{M}u(w) = \int u(w,p)\gamma_d(\,\mathrm{d}p)$, then $\mathsf{M} = \iota^*$, $J\iota = \iota$, and $\mathsf{M}J = \mathsf{M}$. Thus inserting a Gaussian momentum and then discarding it are adjoint operations. Suppressing the fixed superscripts $(y, \widetilde{w})$, $(\mathsf{M}\mathsf{S}_t\iota)^* = \mathsf{M}\mathsf{S}_t\iota$. Conditioning on an independent reference and applying Fubini proves the last assertion. □

Proof of Prop. 2. Use Prop. 6 with $\widetilde{W} \sim q_y^w$, independently of the incoming position. It shows that the conditional position kernel in Alg. 2 is reversible and invariant for q_y. The outer map $(x, y) \mapsto (x, 2x - y)$ preserves Π_η, because it only changes the sign of $y - x$ and has unit absolute Jacobian. Prop. 1 now proves invariance of the complete row and of its position marginal π. □

B.2 The Covariance-Weighted Half-Turn Estimate

The next proposition is the local estimate used by the outer-chain proof. It does not ask the conditional update to approximate a reflection uniformly over all functions. It asks only that centered affine position coordinates be nearly negated, with a smaller error in directions where the localized conditional is closer to the Gaussian reference.

Proposition 7 (Conditional affine half-turn)**.** *There is a universal C such that, whenever $s = \beta\eta \le 1/2$, the normalized-coordinate position kernel F_y^w of Alg. 2 satisfies, uniformly in $y, u \in \mathbb{R}^d$,*

$$\|(I + F_y^w)\ell_{u,y}\|^2_{L^2(q_y^w)} \le Cs\, u^\top(I - \Sigma_y)u, \qquad \ell_{u,y}(w) = u^\top(w - \bar{w}_y). \tag{12}$$

The proof has two ingredients. First, the damped harmonic flight can be solved exactly on affine observables. At the chosen half-period, dependence on the fresh initial momentum vanishes and the incoming position receives a negative coefficient. Second, the event correction appears as a forcing term whose second moment is controlled by $s(I - \Sigma_y)$ through (11).

Fix y and a reference $\widetilde{w}$, and retain the notation g, b_{ref}, h from App. B.1. Let $\bar{r} = \mathbb{E}R = \sqrt{\pi/2}$, $\psi(z) = z_+(z + \bar{r})$, and $n(w) = h(w)/\|h(w)\|$ when $h(w) \neq 0$. For affine position and momentum observables $\ell_{u,y}(w) = u^\top(w - \bar{w}_y)$ and $m_u(p) = u^\top p$, define

$$k_{u,\widetilde{w}}(w,p) = -u^\top\{\bar{w}_y + b_{\mathrm{ref}}\} - u^\top h(w)\psi(p^\top n(w)),$$

where the second term is zero when $h(w) = 0$.

Lemma 2 (Affine evolution)**.** *The two observables belong to the $L^2(q_y^w \otimes \gamma_d)$ generator domain and satisfy*

$$\mathcal{L}\ell_{u,y} = m_u, \qquad \mathcal{L}m_u = -\ell_{u,y} - \gamma m_u + k_{u,\widetilde{w}}.$$

Writing $\omega = \sqrt{1 - \gamma^2/4}$, set

$$A_\gamma(t) = \mathrm{e}^{-\gamma t/2}\left\{\cos(\omega t) + \frac{\gamma}{2\omega}\sin(\omega t)\right\}, \qquad B_\gamma(t) = \mathrm{e}^{-\gamma t/2}\frac{\sin(\omega t)}{\omega}.$$

Then, for the fixed-reference semigroup,

$$\mathsf{S}_t^{y,\widetilde{w}}\ell_{u,y} = A_\gamma(t)\ell_{u,y} + B_\gamma(t)m_u + \int_0^t B_\gamma(t - r)\mathsf{S}_r^{y,\widetilde{w}}k_{u,\widetilde{w}}\,\mathrm{d}r.$$

At $T_\gamma = \pi/\omega$, this becomes

$$(I + \mathsf{S}_{T_\gamma}^{y,\widetilde{w}})\ell_{u,y} = (1 - \mathrm{e}^{-\gamma T_\gamma/2})\ell_{u,y} + \int_0^{T_\gamma} B_\gamma(T_\gamma - r)\mathsf{S}_r^{y,\widetilde{w}}k_{u,\widetilde{w}}\,\mathrm{d}r. \tag{13}$$

Proof. The deterministic part of the generator gives $\mathcal{L}\ell_{u,y} = m_u$ and $-\ell_{u,y} - \gamma m_u - u^\top(\bar{w}_y + b_{\mathrm{ref}})$ in the momentum coordinate. At an event, the momentum change is $-(Z+R)u^\top n$, where $Z = p^\top n > 0$; multiplying its conditional mean by the rate $\|h\|Z$ gives the remaining term $-u^\top h\,\psi(Z)$.

These quantities are square integrable because q_y^w has Gaussian tails, h grows at most linearly, and $p^\top n$ is standard normal conditional on the position. Applying the stopped Dynkin formula and then removing the stopping time using the Lyapunov bound in Prop. 6 gives, after suppressing the fixed superscripts $(y, \widetilde{w})$,

$$\mathsf{S}_t\ell_{u,y} - \ell_{u,y} = \int_0^t \mathsf{S}_r m_u \,\mathrm{d}r, \qquad \mathsf{S}_t m_u - m_u = \int_0^t \mathsf{S}_r(-\ell_{u,y} - \gamma m_u + k_{u,\widetilde{w}})\,\mathrm{d}r.$$

Strong continuity and division by t prove the generator-domain statements. Solving the resulting two-dimensional linear equation by variation of constants gives the displayed formula. At T_γ, $B_\gamma(T_\gamma) = 0$ and $A_\gamma(T_\gamma) = -\mathrm{e}^{-\gamma T_\gamma/2}$, which proves (13). □

Proof of Props. 3 and 7. Let $W, \widetilde{W}$ be independent with law q_y^w, and average first over the Gaussian momentum; we denote this last expectation by $\mathbb{E}_P$. Direct one-dimensional Gaussian integration gives $\mathbb{E}\psi(Z) = 1$ and $\mathrm{Var}\psi(Z) = 5/2 + \pi/4$ for $Z \sim \mathcal{N}(0,1)$. Using $\mathbb{E}g_y(W) = -\bar{w}_y$, we obtain

$$\mathbb{E}_P k_{u,\widetilde{W}} = -u^\top\{g_y(W) - \mathbb{E}g_y(W)\},$$

while the conditional momentum variance is $(5/2+\pi/4)[u^\top\{g_y(W) - g_y(\widetilde{W})\}]^2$. Independence of W and $\widetilde{W}$ therefore gives

$$\mathbb{E}k^2_{u,\widetilde{W}} = \left(6 + \frac{\pi}{2}\right)\mathrm{Var}(u^\top g_y(W)) \le Cs\,u^\top(I - \Sigma_y)u, \tag{14}$$

where the last step is (11).

Insert fresh Gaussian momentum in (13), average the independent reference, and discard the terminal momentum. Jensen's inequality, semigroup contractivity, and $\int_0^{T_\gamma}|B_\gamma(r)|\,\mathrm{d}r \le C$ yield

$$\|(I + F_y^w)\ell_{u,y}\|_2^2 \le C\left\{\gamma^2 u^\top\Sigma_y u + \mathbb{E}k^2_{u,\widetilde{W}}\right\}.$$

Here the norm is in $L^2(q_y^w)$, and the remaining expectation includes the independent position, reference, and Gaussian momentum. The parameter choice in Sec. 4 gives $\gamma^2 \le as/(1+a)$. Since $I - \Sigma_y \succeq a(1+a)^{-1}I$ and $\Sigma_y \preceq I$, the damping term is at most $su^\top(I - \Sigma_y)u$. Combining this with (14) proves (12), which is exactly the main-text estimate in Prop. 3. □

B.3 From the Conditional Half-Turn to Accelerated Contraction

We now prove the ideal-chain theorem. The argument follows two components of an observable: its average at fixed Y, and its remaining conditional variation. Outer reflection transfers part of the first component into the second. The flux half-turn approximately reverses precisely this transferred part, while an occasional independent conditional refresh damps all conditional variation.

Let $\mathfrak{H} = L^2(\Pi_\eta)$, let $\mathsf{P}f = \mathbb{E}[f \mid Y]$, and set $\mathsf{Q} = I - \mathsf{P}$. Let U be the pullback of the outer reflection and F the joint conditional half-turn. By Prop. 6, $\mathsf{F} = \mathsf{P} + \mathsf{F}_Q$, where F_Q is a self-adjoint contraction on $\mathrm{ran}(\mathsf{Q})$. The first term records observables depending only on Y; the second acts on their remaining conditional variation. Thus one ideal iteration has observable operator

$$\mathsf{K} = \mathsf{U}\{\mathsf{P} + (1-\rho)\mathsf{F}_Q\}. \tag{15}$$

Proposition 8 (Ideal L^2 contraction)**.** *There are universal constants $c, C > 0$ such that, when $s \le c$ and $\rho \asymp \sqrt{a\log(\mathrm{e}/a)}$, the ideal chain satisfies*

$$\|\mathsf{K}^n\|_{L^2_0(\Pi_\eta)\to L^2_0(\Pi_\eta)} \le \sqrt{3}\exp\left\{-\frac{cn\sqrt{a}}{\sqrt{\log(\mathrm{e}/a)}}\right\}.$$

The rest of this subsection proves the proposition. We first quantify what reflection transfers, then lift the conditional estimate in Prop. 7 to the complete state space, and finally construct an energy that decreases in one step.

B.3.1 What Outer Reflection Transfers

On the centered Y-only subspace $\mathcal{H}_Y = \mathrm{ran}(\mathsf{P}) \cap L^2_0(\Pi_\eta)$, define

$$\mathsf{T} = \mathsf{PUP}, \qquad \Gamma = (I - \mathsf{T}^2)^{1/2}, \qquad \mathsf{C} = \mathsf{QUP}.$$

Under stationarity, let $X \sim \pi$ and $Z \sim \mathcal{N}(0, I_d)$ be independent, and write $Y_+ = X + \sqrt{\eta}Z$, $Y_- = X - \sqrt{\eta}Z$. These are not two independent chain iterates: they share (X, Z) and are the auxiliary values before and after the same reflection, so $Y_- = 2X - Y_+$. Then $(\mathsf{T}f)(y) = \mathbb{E}[f(Y_-) \mid Y_+ = y]$, and $\|\mathsf{C}f\|^2 = \mathbb{E}\mathrm{Var}(f(Y_-) \mid Y_+)$. Thus T is the part retained as a function of Y_+, while C is the part transferred into conditional variation. The operator Γ measures the size of this transferred part.

Lemma 3 (Outer-reflection geometry)**.** *If $s \le 1/3$, then T is self-adjoint and, on $\mathcal{H}_Y$,*

$$-\tfrac{1}{2}I \preceq \mathsf{T} \preceq \frac{1-a}{1+a}I, \qquad \Gamma \succeq \gamma_a I, \qquad \gamma_a = \frac{2\sqrt{a}}{1+a}. \tag{16}$$

Moreover, $4\eta\|\nabla \mathsf{T}f\|^2 \le \|\Gamma f\|^2$, and $\mathsf{C} = \mathsf{V}\Gamma$ for an isometry $\mathsf{V} : \mathcal{H}_Y \to \mathrm{ran}(\mathsf{Q})$.

Proof. Exchangeability of (Y_+, Y_-) makes T self-adjoint. The conditional density of Y_- given $Y_+ = y$ is proportional to

$$\exp\left\{-V\left(\frac{y+u}{2}\right) - \frac{\|y-u\|^2}{8\eta}\right\}.$$

Its curvature in u is at least $(\alpha + \eta^{-1})I/4$, while the derivative with respect to u of its y-score has norm at most $(\eta^{-1} - \alpha)/4$. Conditional Brascamp–Lieb therefore gives

$$\eta\|\nabla \mathsf{T}f\|^2 \le \frac{(1-a)^2}{4(1+a)}\|\mathsf{C}f\|^2 \le \frac{1}{4}\|\Gamma f\|^2.$$

The Y-marginal has Poincaré constant at most $\alpha^{-1} + \eta$ by Lem. 1. Applying this inequality to the centered function $\mathsf{T}f$ and using $\|f\|^2 = \|\mathsf{T}f\|^2 + \|\mathsf{C}f\|^2$ gives $\|\mathsf{T}\| \le (1-a)/(1+a)$.

For the lower spectral edge, exchangeability and Gaussian Poincaré give $\langle h, (I - \mathsf{T})h\rangle \le 2\eta\|\nabla h\|^2$. Apply this with $h = \mathsf{T}f$ and use the preceding derivative bound. Spectral calculus yields $(I-\mathsf{T})^2(I + 2\mathsf{T}) \succeq 0$, hence $\mathsf{T} \succeq -I/2$. Because $(1-a)/(1+a) \ge 1/2$, these two edges imply the stated lower bound on Γ. Finally, unitarity of U gives $\mathsf{C}^*\mathsf{C} = I - \mathsf{T}^2 = \Gamma^2$, so $\mathsf{V} = \mathsf{C}\Gamma^{-1}$ is an isometry.

The weak-gradient and spectral-domain justification for this calculation is given in App. B.4.2. □

Reflection creates a conditional-position signal that is almost affine. For weakly differentiable $f \in \mathcal{H}_Y$, set

$$b_y = \mathbb{E}[\nabla f(Y_-) \mid Y_+ = y], \qquad L_f = 2\sqrt{\eta}\, b_{Y_+}^\top (W - \bar{w}_{Y_+}), \qquad W = \frac{X - Y_+}{\sqrt{\eta}}.$$

Here b_y is the average slope of $f(Y_-)$ conditional on the reflected value $Y_+ = y$, and L_f is the corresponding centered linear signal in the conditional position W. The next lemma separates this leading signal from a nonlinear remainder. Its quantities $\mathcal{R}_f$ and $\mathcal{D}_f$ measure, respectively, the gradient variation not retained by conditional averaging and the curvature-weighted size of the affine signal.

Lemma 4 (Affine approximation after reflection)**.** *There is an error e_f, centered conditionally on Y_+, such that $\mathsf{C}f = L_f + e_f$. If*

$$\mathcal{R}_f = \|\nabla f\|^2 - \|\mathsf{T}\nabla f\|^2, \qquad \mathcal{D}_f = \mathbb{E}[b_{Y_+}^\top (I - \Sigma_{Y_+}) b_{Y_+}],$$

then, for sufficiently small universal s,

$$\|e_f\|^2 \le 4\eta\mathcal{R}_f, \qquad \mathcal{R}_f + \mathcal{D}_f \le C\|\nabla(I - \mathsf{T})f\|\|\nabla f\|. \tag{17}$$

Proof. Here $\mathsf{T}\nabla f$ means that T is applied to every coordinate of the vector-valued gradient; in particular, $(\mathsf{T}\nabla f)(y) = b_y$. Conditionally on $Y_+ = y$, $Y_- = y + 2\sqrt{\eta}W$. Subtracting its best first-order conditional approximation gives

$$e_f = f(Y_-) - \mathsf{T}f(Y_+) - 2\sqrt{\eta}\, b_{Y_+}^\top (W - \bar{w}_{Y_+}).$$

Conditional Poincaré proves $\|e_f\|^2 \le 4\eta\mathcal{R}_f$.

For the second estimate, differentiate the conditional expectation weakly. Integration by parts and (11) give

$$\nabla(I-\mathsf{T})f = \nabla f - b_y + 2(I-\Sigma_y)b_y + \eta^{-1/2}\mathrm{Cov}(e_f, g_y(W) \mid Y_+ = y).$$

In this identity, ∇f on the right is evaluated at y, and the conditional covariance is over W given $Y_+ = y$. Write $b(y) = b_y$. Take its inner product with $\nabla f + b_y$ and average over y. The first term gives $\mathcal{R}_f$, and the leading covariance term gives $4\mathcal{D}_f$. Also, the spectral lower bound $\mathsf{T} \succeq -I/2$ implies

$$\|\nabla f - b\|^2 \le 3\{\|\nabla f\|^2 - \|b\|^2\} = 3\mathcal{R}_f,$$

The two remaining cross terms are bounded using this inequality, (11), the first inequality in (17), and $I - \Sigma_y \preceq sI$. They are at most

$$C\sqrt{s\mathcal{R}_f\mathcal{D}_f} + C\sqrt{s\mathcal{R}_f(s\mathcal{R}_f + \mathcal{D}_f)}.$$

The leading inner product is at most $2\|\nabla(I-\mathsf{T})f\|\|\nabla f\|$. Young's inequality absorbs the displayed cross terms when s is sufficiently small. The graph-norm extension to the stated weak-gradient domain is recorded in App. B.4.2. □

B.3.2 Lifting the Local Half-Turn Estimate

The next result is the bridge between the conditional dynamics and the outer Markov chain.

We first record the operator comparison that accounts for the logarithm in that bridge. It converts an interpolation-type estimate into a quadratic form estimate without assuming that the two operators commute.

Lemma 5 (Dyadic operator comparison)**.** *Suppose $mI \preceq B \preceq MI$ and a bounded operator A satisfies $\|Av\|^2 \le \|v\|\|Bv\|$. Then*

$$A^*A \preceq C\{1 + \log(M/m)\}B.$$

Proof. Partition the spectrum of B into dyadic intervals and write $v = \sum_{j=1}^{J} v_j$, where $J \le C\{1 + \log(M/m)\}$. If λ_j is the lower endpoint of the j-th interval, then the hypothesis gives $\|Av_j\| \le \sqrt{2\lambda_j}\|v_j\|$. Consequently,

$$\|Av\|^2 \le J\sum_j 2\lambda_j\|v_j\|^2 \le 2J\langle v, Bv\rangle,$$

which proves the claim. □

Proposition 9 (Global half-turn estimate)**.** *For sufficiently small universal s and every $v \in \mathcal{H}_Y$,*

$$\|(I + \mathsf{F}_Q)\mathsf{V}v\|^2 \le C\log(\mathrm{e}/\gamma_a)\|\Gamma v\|^2. \tag{18}$$

Proof. First restrict f to the spectral subspace of T corresponding to $[1/2, 1]$. Apply Prop. 7 conditionally to the affine part L_f of $\mathsf{C}f$, use $\|I + \mathsf{F}_Q\| \le 2$ on e_f, and then apply Lem. 4. This gives

$$\|(I + \mathsf{F}_Q)\mathsf{C}f\|^2 \le C\eta\|\nabla(I-\mathsf{T})f\|\|\nabla f\|.$$

On this spectral subspace, applying the derivative estimate in Lem. 3 to $\mathsf{T}^{-1}f$ and $\mathsf{T}^{-1}(I-\mathsf{T})f$ gives

$$\sqrt{\eta}\|\nabla f\| \le C\|\Gamma f\|, \qquad \sqrt{\eta}\|\nabla(I-\mathsf{T})f\| \le C\|\Gamma^3 f\|.$$

Since $\mathsf{C} = \mathsf{V}\Gamma$, setting $f = \Gamma^{-1}v$, which is allowed by $\Gamma \succeq \gamma_a I$, shows that $\mathsf{A}_{\mathrm{ht}} = (I + \mathsf{F}_Q)\mathsf{V}$ obeys $\|\mathsf{A}_{\mathrm{ht}}v\|^2 \le C\|v\|\|\Gamma^2 v\|$. Applying Lem. 5 to $A = \mathsf{A}_{\mathrm{ht}}$, $B = \Gamma^2$, and $m = \gamma_a^2$ gives the required bound without treating the generally non-self-adjoint half-turn error A_{ht} as a positive operator.

On the complementary spectral subspace, $\Gamma^2 \succeq 3I/4$, while $\|\mathsf{A}_{\mathrm{ht}}\| \le 2$, so the result is immediate. The two output vectors need not be orthogonal; the triangle inequality and $\|r_1 + r_2\|^2 \le 2\|r_1\|^2 + 2\|r_2\|^2$ prove (18) for an arbitrary v. □

B.3.3 A DECREASING ENERGY

For centered $f \in \mathfrak{H}$, write $f_Y = \mathsf{P}f$, $f_Q = \mathsf{Q}f$, and $f_V = \mathsf{V}^* f_Q$. Since V is an isometry, the conditional component has the orthogonal decomposition

$$f_Q = \mathsf{V}f_V + f_\perp, qquadf_\perp = (I - \mathsf{V}\mathsf{V}^*)f_Q.$$

The three pieces have distinct roles. The term f_Y is variation across values of Y; $\mathsf{V}f_V$ is the conditional signal created from it by reflection; and $f_\perp$ is the remaining conditional variation. The cross term below tracks the exchange between the first two pieces, whereas the RGO branch controls the full f_Q, including $f_\perp$. Define

$$\mathcal{E}_\vartheta(f) = \|f\|^2 + \vartheta\mathcal{C}(f_Y, f_V), \qquad \mathcal{C}(f_Y, f_V) = \frac{1}{2}(\|f_Y\|^2 - \|f_V\|^2) - \langle \mathsf{T}\Gamma^{-1} f_Y, f_V\rangle.$$

The form $\mathcal{C}$ is chosen so that, under an exact conditional sign reversal, the pair (f_Y, f_V) undergoes the rotation displayed below and $\mathcal{C}$ decreases by $\|f_Y\|^2 - \|f_V\|^2$.

Lemma 6 (One-step energy decay)**.** *Let*

$$\Lambda_\rho = \log(\mathrm{e}/\gamma_a) + (\rho/\gamma_a)^2, \qquad \vartheta = c_\star \rho/\Lambda_\rho$$

for a sufficiently small universal $c_\star$. If $0 < \rho \le 1/2$, then

$$\tfrac{1}{2}\|f\|^2 \le \mathcal{E}_\vartheta(f) \le \tfrac{3}{2}\|f\|^2, \qquad \mathcal{E}_\vartheta(\mathsf{K}f) \le \left(1 - c\frac{\rho}{\Lambda_\rho}\right)\mathcal{E}_\vartheta(f).$$

Proof. Because T commutes with Γ, the block operator associated with $2\mathcal{C}$ has square $\mathrm{diag}(\Gamma^{-2}, \Gamma^{-2})$. Hence $|\mathcal{C}| \le (2\gamma_a)^{-1}(\|f_Y\|^2 + \|f_V\|^2)$. The displayed choice of ϑ satisfies $\vartheta \le \gamma_a$, proving energy equivalence.

From (15) and contractivity of F_Q,

$$\|\mathsf{K}f\|^2 - \|f\|^2 \le -\rho\|f_Q\|^2. \tag{19}$$

To see how the correction changes, first replace the conditional move by an exact sign reversal and set $g = \mathsf{U}(\mathsf{P} - \mathsf{Q})f$. Unitarity of U and $\mathsf{C} = \mathsf{V}\Gamma$ give

$$\begin{pmatrix} g_Y \\ g_V \end{pmatrix} = \begin{pmatrix} \mathsf{T} & -\Gamma \\ \Gamma & \mathsf{T} \end{pmatrix}\begin{pmatrix} f_Y \\ f_V \end{pmatrix}, \qquad \mathcal{C}(g_Y, g_V) - \mathcal{C}(f_Y, f_V) = -\|f_Y\|^2 + \|f_V\|^2.$$

The actual conditional move differs from this sign reversal by

$$e_{\mathrm{ht}} = \mathsf{V}^*(I + \mathsf{F}_Q)f_Q, \qquad e_\rho = \rho f_V + (1-\rho)e_{\mathrm{ht}}.$$

The two tracked components of the actual update are $(\mathsf{K}f)_Y = g_Y + \Gamma e_\rho$ and $(\mathsf{K}f)_V = g_V - \mathsf{T}e_\rho$. Direct substitution gives

$$\mathcal{C}((\mathsf{K}f)_Y, (\mathsf{K}f)_V) - \mathcal{C}(f_Y, f_V) = -\|f_Y\|^2 + \|f_V\|^2 + \langle g_Y, \Gamma^{-1}e_\rho\rangle + \tfrac{1}{2}\|e_\rho\|^2.$$

By Prop. 9, $\mathsf{A}_{\mathrm{ht}}^*\mathsf{A}_{\mathrm{ht}} \preceq C\log(\mathrm{e}/\gamma_a)\Gamma^2$ for $\mathsf{A}_{\mathrm{ht}} = (I + \mathsf{F}_Q)\mathsf{V}$. The corresponding operator factorization gives a bounded B, with $\|B\|^2 \le C\log(\mathrm{e}/\gamma_a)$, such that $\mathsf{A}_{\mathrm{ht}} = B\Gamma$. Since $e_{\mathrm{ht}} = \mathsf{A}_{\mathrm{ht}}^* f_Q$, this gives $\|\Gamma^{-1}e_{\mathrm{ht}}\|^2 \le C\log(\mathrm{e}/\gamma_a)\|f_Q\|^2$. Therefore $\|\Gamma^{-1}e_\rho\|^2 \le C\Lambda_\rho\|f_Q\|^2$, while $\|e_\rho\| \le 2\|f_Q\|$. Young's inequality now yields

$$\mathcal{C}((\mathsf{K}f)_Y, (\mathsf{K}f)_V) - \mathcal{C}(f_Y, f_V) \le -\tfrac{1}{2}\|f_Y\|^2 + C\Lambda_\rho\|f_Q\|^2.$$

Multiply by ϑ, add (19), and choose $c_\star$ small enough. Energy equivalence then proves the claimed one-step decay. □

Proof of Prop. 8. By (16), $\gamma_a \asymp \sqrt{a}$. With $\rho \asymp \sqrt{a\log(\mathrm{e}/a)}$, we have $\Lambda_\rho \lesssim \log(\mathrm{e}/a)$ and $\rho/\Lambda_\rho \gtrsim \sqrt{a}/\sqrt{\log(\mathrm{e}/a)}$. Iterating Lem. 6 and using energy equivalence proves the proposition. □

Proof of Thm. 1. Choose n_{blk} as in the theorem, with a sufficiently large universal constant. Prop. 8 gives $\|\mathsf{K}^{n_{\mathrm{blk}}}\|^2_{L_0^2 \to L_0^2} \le \mathrm{e}^{-1}$. If $\varrho_\nu = \mathrm{d}\nu/\,\mathrm{d}\Pi_\eta$, then the density evolves by K^* and $\chi^2(\nu\|\Pi_\eta) = \|\varrho_\nu - 1\|^2$. Since an operator and its adjoint have the same norm, applying the block estimate J times proves (3). □

B.4 Implementable Algorithm and the Query Ledger

The ideal chain uses an exact conditional reference and an uncapped event clock. Its finite counterpart makes the following three replacements:

source	ideal operation	finite operation
reference	exact draw from q_y	certified approximate RGO
event clock	full state-dependent rate	capped Poisson thinning
center preparation	unrestricted solve	capped solver with early return

The proposition is stated first; the rest of the subsection controls the three corresponding disagreement events and counts their oracle queries.

Proposition 10 (Finite compiler at a fixed proximal scale)**.** *Let $D_2(\nu_0\|\pi) \le \Delta$, where $\Delta \ge 1$, and let $0 < \varepsilon < 1/4$. For $0 < \eta \le 1/\beta$, define*

$$\ell = K\left\{\Delta + \log\left(\frac{Kd\kappa}{\varepsilon\alpha\eta}\right)\right\}.$$

There are universal constants $c, K > 0$ such that, if

$$s\{\sqrt{d\ell} + \ell\} \le c, \tag{20}$$

the finite flux algorithm returns $\widehat{X}$ with $\mathrm{TV}(\mathrm{Law}(\widehat{X}), \pi) \le \varepsilon$ and

$$\mathbb{E}N_{\mathrm{qry}} \le \frac{K}{\sqrt{a}}\sqrt{\log(\mathrm{e}/a)}\left(\Delta + \log\frac{4}{\varepsilon}\right)\left\{\log(\mathrm{e}\kappa\ell) + s(\sqrt{d\ell} + \ell)\right\}. \tag{21}$$

B.4.1 Capped Poisson thinning

For fixed $(y, \widetilde{x})$, put $x_{\mathrm{ctr}} = y - \eta\nabla V(\widetilde{x})$ and $U = (X - x_{\mathrm{ctr}})/\sqrt{\eta}$. Between events, (U, P) follows a linear Gaussian diffusion. With $\omega = \sqrt{1 - \gamma^2/4}$, its conditional mean after time t is

$$\mathrm{e}^{-\gamma t/2}\begin{pmatrix} \cos(\omega t) + \frac{\gamma}{2\omega}\sin(\omega t) & \frac{\sin(\omega t)}{\omega} \\ -\frac{\sin(\omega t)}{\omega} & \cos(\omega t) - \frac{\gamma}{2\omega}\sin(\omega t)\end{pmatrix}\begin{pmatrix} U \\ P\end{pmatrix},$$

coordinatewise. Its covariance is obtained by integrating the transported Gaussian noise over this same linear motion. Both mean and covariance are explicit finite-dimensional operations. Thus, once $\nabla V(\widetilde{x})$ has been stored, no oracle query is needed between candidate event times.

Given a cap $\bar{\lambda}$, candidate times form a homogeneous Poisson process of rate $\bar{\lambda}$. At a candidate time, one oracle query computes

$$\lambda_t = \sqrt{\eta}[P_t^\top\{\nabla V(X_t) - \nabla V(\widetilde{x})\}]_+,$$

and the candidate is accepted with probability $\min\{1, \lambda_t/\bar{\lambda}\}$. This is exact thinning for the truncated rate $\min\{\lambda_t, \bar{\lambda}\}$.

Lemma 7 (Stationary rate tail and clock coupling)**.** *Conditionally on $Y = y$, let $X, \widetilde{X}$ be independent draws from q_y, and let $P \sim \gamma_d$ be independent. Then, for universal $c, C > 0$ and every $u \ge 0$,*

$$\Pr\left\{\sqrt{\eta}[P^\top\{\nabla V(X) - \nabla V(\widetilde{X})\}]_+ > Cs(\sqrt{du} + u)\right\} \le C\mathrm{e}^{-cu}.$$

If, in addition, $s\sqrt{d}$ is below a universal constant, then an augmented initial law ν satisfying $D_2(\nu\|\Pi_\eta) \le \Delta$ admits the following consequence: for every $\xi \in (0, 1/2)$, the exact and capped half-turns can be coupled to agree with probability at least $1 - \xi$, using

$$\bar{\lambda}_{\xi,\Delta} = Cs\{\sqrt{d\ell_{\xi,\Delta}} + \ell_{\xi,\Delta}\}, \qquad \ell_{\xi,\Delta} = C\{\Delta + \log(1/\xi)\}.$$

The capped half-turn uses exactly $T_\gamma\bar{\lambda}_{\xi,\Delta}$ candidate events in expectation.

Proof. The conditional law is $(\alpha + \eta^{-1})$-strongly log-concave. Norm concentration and its second-moment bound show that, outside an event of probability $2\mathrm{e}^{-u}$, both $\|X - \mathbb{E}X\|$ and $\|\widetilde{X} - \mathbb{E}X\|$ are at most $C\sqrt{\eta(d+u)}$. Smoothness then bounds the conditional standard deviation of the scalar product with P by $Cs(\sqrt{d} + \sqrt{u})$. A Gaussian tail bound proves the displayed estimate. Integrating it above $Cs(\sqrt{d\ell} + \ell)$ gives

$$\mathbb{E}(\lambda - \bar{\lambda})_+ \le Cs(\sqrt{d} + \ell + 1)\mathrm{e}^{-c\ell}.$$

Under Π_η, invariance from Prop. 6 makes this bound valid at every time. Couple the exact and capped clocks with the same Poisson random measure; they separate only if an excess-rate event occurs. Taking $\ell = C\{\Delta + \log(1/\xi)\}$ makes the stationary separation probability at most $\mathrm{e}^{-\Delta}\xi^2$. Under a D_2-warm initial law, the path likelihood ratio is only its initial density ratio, since all subsequent references, Gaussian noise, clocks, and marks use common kernels. Cauchy–Schwarz therefore raises the stationary failure probability by at most $\mathrm{e}^{\Delta/2}$, giving ξ. The final expectation is that of a homogeneous Poisson clock on an interval of length $T_\gamma = O(1)$. □

Approximate conditional draws. We use the following existing RGO result.

Proposition 11 (Approximate RGO)**.** *There is a universal K_{RGO} such that the following holds. Given y, $0 < \delta < 1/2$, and a point $\widehat{x}$ satisfying*

$$\|y - \eta\nabla V(\widehat{x}) - \widehat{x}\| \le \sqrt{d\eta},$$

if

$$\eta^{-1} \ge K_{\mathrm{RGO}}\beta\{\sqrt{d\log(1/\delta)} + \log(1/\delta)\},$$

the rejection sampler of Chen et al. (2026a, Theorems 3.1 and D.1) returns a law $\widehat{q}_y$ satisfying

$$\chi^2(\widehat{q}_y \| q_y) \le \delta^2, \qquad \chi^2(q_y \| \widehat{q}_y) \le \delta^2,$$

using at most K_{RGO} expected oracle queries after construction of the displayed certificate.

The certificate itself uses only ordinary oracle queries. The conditional potential $V_y(x) = V(x) + \|x - y\|^2/(2\eta)$ has condition number at most two. Gradient descent with step $2/(\alpha + \beta + 2\eta^{-1})$, initialized at the incoming position x, reaches and certifies the required residual in at most

$$C\left[1 + \log\left\{1 + \frac{\|y - \eta\nabla V(x) - x\|}{\sqrt{d\eta}}\right\}\right] \tag{22}$$

oracle queries, by standard linear convergence for strongly convex smooth optimization.

The implementable algorithm and error allocation. Use ρ from (2), and set

$$n_{\mathrm{blk}} = \left\lceil \frac{K}{\sqrt{a}}\sqrt{\log(\mathrm{e}/a)} \right\rceil, \quad M = \left\lceil K\left(\Delta + \log\frac{4}{\varepsilon}\right)\right\rceil n_{\mathrm{blk}}, \quad \delta_{\mathrm{row}} = \frac{\varepsilon}{6M}.$$

Use the cap $\bar{\lambda} = Ks(\sqrt{d\ell} + \ell)$ and allow at most $Q_{\mathrm{ctr}} = \lceil K\log(\mathrm{e}\kappa\ell)\rceil$ center-solver queries per row. Initialize $X_0 \sim \nu_0$, $Y_0 \mid X_0 \sim \mathcal{N}(X_0, \eta I_d)$. Each row reflects Y, runs the capped center solver, draws an approximate reference from Prop. 11 with tolerance δ_{row}, and then uses that draw either as the independent refresh or as the reference of the capped flux half-turn. If the center solver does not certify its residual within its cap, the algorithm stops and returns the current position. This failure rule, together with the almost-sure finiteness of the rejection sampler and Poisson clock, makes the procedure almost surely finite from every initial state.

Proof of Prop. 10. Augmentation preserves warmness:

$$\frac{\mathrm{d}\,\mathrm{Law}(X_0, Y_0)}{\mathrm{d}\Pi_\eta}(x, y) = \frac{\mathrm{d}\nu_0}{\mathrm{d}\pi}(x), \qquad D_2(\mathrm{Law}(X_0, Y_0)\|\Pi_\eta) \le \Delta.$$

Thm. 1 and the choice of M put the ideal output within $\varepsilon/2$ of π. Because every ideal row preserves Π_η, data processing keeps its intermediate D_2 divergence at most Δ.

It remains to couple every finite row to its ideal counterpart. There are three possible disagreements.

Center cap. Under stationarity, write $Y = X + \sqrt{\eta}Z$. After reflection, the normalized residual at the incoming position is

$$\frac{\|\sqrt{\eta}Z + \eta\nabla V(X)\|}{\sqrt{d\eta}}.$$

Strong log-concavity, Gaussian concentration, and $\|\nabla V(X)\| \le \beta\|X - x^\star\|$, where $x^\star$ minimizes V, show that it exceeds $2\sqrt{\kappa}(1 + \sqrt{2u/d})$ with probability at most $2\mathrm{e}^{-u}$. Taking $u = \Delta + 2\log(1/\delta_{\mathrm{row}}) + \log 2$ and using (22) shows that the stationary cap-failure probability is at most $\mathrm{e}^{-\Delta}\delta_{\mathrm{row}}^2$. The warm-path Cauchy–Schwarz argument used in Lem. 7 makes it at most δ_{row} under the current ideal law.

Approximate reference. The definitions imply $\log(1/\delta_{\mathrm{row}}) \le C\ell$. The gate (20) therefore implies the small-scale condition in Prop. 11. Either directed χ^2 bound implies total variation at most $\delta_{\mathrm{row}}/2$, so maximal coupling matches the approximate and exact references except with probability at most δ_{row}. Because all laws have Borel densities on $\mathbb{R}^d$, these maximal couplings may be selected jointly measurably in y. On agreement, the same draw is used in both chains.

Capped clock. The same parameter bound gives $\ell_{\delta_{\mathrm{row}},\Delta} \le C\ell$. After increasing K, Lem. 7, applied conditionally on the common pre-row history, couples the exact and capped half-turns with failure probability at most δ_{row}.

A union bound over three events and M rows is at most $\varepsilon/2$. Together with the ideal mixing error, this proves the total-variation claim.

Early termination only reduces work. Each row uses at most Q_{ctr} center queries, a universal expected number of RGO queries, one query to cache $\nabla V(\widetilde{X})$ on the half-turn branch, and $T_\gamma\bar{\lambda}$ expected event-test queries. Hence

$$\mathbb{E}N_{\mathrm{qry}} \le CM\{Q_{\mathrm{ctr}} + s(\sqrt{d\ell} + \ell)\},$$

which is (21). □

Proof of Thm. 2. Take η and L as in (4), with the universal constant c sufficiently small. Then

$$a = \frac{c}{\kappa(\sqrt{dL} + L)}, \qquad s = \frac{c}{\sqrt{dL} + L}, \qquad \ell \le CL,$$

so (20) holds. Moreover,

$$\frac{1}{\sqrt{a}}\sqrt{\log(\mathrm{e}/a)} \le C\sqrt{\kappa}\,(d + L)^{1/4}L^{3/4},$$

and the braces in (21) are at most $C\log(\mathrm{e}\kappa L)$. Substitution proves (5) and the total-variation conclusion. □

For the warm-start removal argument we also record the dependence on a nonconstant divergence budget.

Corollary 1 (General warmness budget)**.** *Let $B \ge 1$, $D_2(\nu_0\|\pi) \le B$, and set*

$$L_B = B + \log\frac{Cd\kappa}{\varepsilon}, \qquad A_B = \sqrt{dL_B} + L_B, \qquad \eta = \frac{c}{\beta A_B}.$$

Then the flux algorithm returns an ε-TV sample using

$$\mathbb{E}N_{\mathrm{qry}} \le C\sqrt{\kappa}A_B^{1/2}\sqrt{\log(\mathrm{e}\kappa A_B)}\left(B + \log\frac{4}{\varepsilon}\right)\log(\mathrm{e}\kappa L_B)$$

expected oracle queries. Ignoring logarithmic factors, this is $\widetilde{O}(\sqrt{\kappa}B\{\sqrt{dB} + B\}^{1/2})$, and hence $\widetilde{O}(\sqrt{\kappa}d^{1/4}B^{5/4})$ when $1 \le B \le d$.

Proof. Apply Prop. 10 with $\Delta = B$. The displayed choice gives $a = c/(\kappa A_B)$, $s = c/A_B$, and $\ell \le CL_B$, so the finite gate holds. Substitution into (21) gives the stated bound. □

B.4.2 Regularity and Domain Extensions

For completeness, we collect here the regularity steps used above. The weak derivatives and integrations by parts are first justified for smooth potentials and test functions. Mollifying V preserves strong convexity and smoothness; the normalizers and conditional moments have a common Gaussian envelope, so the identities pass to the limit.

The derivative estimate in Lem. 3 also begins on smooth functions. Closedness of the weak gradient extends $\nabla \mathsf{T}$ by continuity. On the spectral subspace $[1/2, 1]$, T^{-1} is bounded. Applying the extended estimate to $\mathsf{T}^{-1} f$ and $\mathsf{T}^{-1}(I - \mathsf{T})f$ therefore gives approximations converging both in L^2 and in weak-gradient norm; in particular, no nonsmooth spectral projection is applied directly to a classical derivative. The same graph-norm approximation extends Lem. 4 from smooth functions to its stated weak-gradient domain.

The event-driven process itself needs no smoothing: its residual force is globally Lipschitz, its integrated hazard is continuous, and the direction $h/\|h\|$ is evaluated only at positive event rate.

The proof is complete relative to the earlier RGO result in Prop. 11 and standard strongly log-concave concentration and optimization facts.

C Proofs for the Persistent Quarter-Turn

This appendix proves Props. 4 and 5, Thms. 3 and 4, and Cor. 3. The argument is organized in the dependency order of the proof:

1. show that every fixed-reference conditional motion is exact;
2. quantify how outer reflection creates a conditional-position signal;
3. prove that a quarter-turn transfers this signal to and from momentum;
4. combine this transfer with occasional refreshment to obtain L^2 contraction;
5. replace the ideal reference, event clock, and refresh by implementable implementations and account for their errors;
6. absorb the known strong-convexity quadratic into the conditional motion to obtain the near-Gaussian refinement.

Standard ingredients are cited rather than reproved. In particular, we use established skew-balance results for velocity-jump processes, standard Poincaré and log-Sobolev consequences of strong log-concavity, and the existing first-order RGO implementation. The new arguments are the position–momentum transfer estimate, its unit-circle contraction analysis, and the shifted-reference implementation.

Throughout, π denotes the target. Uppercase X, Y, P, E denote position, auxiliary position, momentum, and reference random variables; lowercase letters denote their values. Operators act on observables, so laws evolve through their adjoints. Concretely, for a Markov transition A, $(\mathsf{A}f)(z) = \mathbb{E}[f(Z_{\mathrm{next}}) \mid Z_{\mathrm{current}} = z]$; thus operator products act from right to left on functions, even though the corresponding state updates occur in chronological order. The small set of operators used throughout is:

symbol	meaning
P	conditional expectation given Y
$\mathsf{Q} = I - \mathsf{P}$	remaining conditional variation
U	outer reflection
R	momentum reversal, $(x, y, p) \mapsto (x, y, -p)$
S	conditional quarter-turn
K	one complete ideal iteration
$\mathsf{T} = \mathsf{PUP}$	reflection seen on Y-only observables
$\Gamma = (I - \mathsf{T}^2)^{1/2}$	amount transferred out of the Y-only part

C.1 Exactness of the Persistent Conditional Move

We first isolate the only facts about the event-driven dynamics that are used later: it is well defined, it preserves the desired conditional law, and reversing momentum gives its stationary time reversal. These are standard skew-balance properties of bouncy and Boomerang processes (Bouchard-Côté et al., 2018; Vanetti et al., 2017; Bierkens et al., 2020; Monmarché et al., 2023). We verify that our conditional process satisfies the required identities.

Fix $Y = y$, use the normalized position $w = (x - y)/\sqrt{\eta}$, and write

$$q_y^w(\,\mathrm{d}w) \propto \exp\{-V(y + \sqrt{\eta}w) - \|w\|^2/2\}\,\mathrm{d}w, \qquad b_y(w) = \sqrt{\eta}\,\nabla V(y + \sqrt{\eta}w).$$

For a nonzero vector h, let $\mathrm{Ref}_h\, p = p - 2(p^\top h)h/\|h\|^2$, and set $\mathrm{Ref}_0\, p = p$. A physical reference E corresponds to $e = (E - y)/\sqrt{\eta}$. Conditional on (y, e), the generator is

$$\begin{aligned}(\mathcal{L}_{y,e}F)(w,p) = {} & p^\top \nabla_w F - \{w + b_y(e)\}^\top \nabla_p F \\ & + [p^\top\{b_y(w) - b_y(e)\}]_+\{F(w, \mathrm{Ref}_{b_y(w)-b_y(e)}\, p) - F(w,p)\}. \end{aligned} \tag{23}$$

The first line is a deterministic harmonic flight centered at $-b_y(e)$. The last line says that a state-dependent Poisson clock rings at rate $[p^\top\{b_y(w) - b_y(e)\}]_+$; when it rings, only the momentum is reflected across the residual-gradient direction $b_y(w) - b_y(e)$. Fixing e is a proof device. The accelerated branch later draws a fresh reference independently given y, and averages these fixed-reference kernels.

Lemma 8 (Fixed-reference validity)**.** *For every fixed y, e,* (23) *defines a non-explosive Markov process. If $S_t^{y,e}$ denotes its transition operator and $\mathsf{R}(w,p) = (w,-p)$, then*

$$q_y^w \otimes \gamma_d \text{ is invariant}, \qquad (S_t^{y,e})^* = \mathsf{R} S_t^{y,e} \mathsf{R}.$$

Both conclusions remain valid after averaging $S_t^{y,e}$ over any measurable reference law depending on y, provided the reference is conditionally independent of the incoming (W, P).

Proof. Put $a = b_y(e)$, $r(w) = b_y(w) - a$, and

$$\mathsf{J}(w,p) = (w, \mathrm{Ref}_{r(w)}\, p), \qquad s(w,p) = p^\top r(w).$$

Between events the process follows the harmonic flow

$$\phi_u(w,p) = \big(-a + (w+a)\cos u + p \sin u, -(w+a)\sin u + p\cos u\big).$$

Both this flow and J preserve $H(w,p) = \{\|w+a\|^2 + \|p\|^2\}/2$. Thus every set $K_R = \{H \le R\}$ is invariant. Lipschitz continuity of b_y bounds the event rate $\lambda = [s]_+$ on K_R by

$$\|p\|\,\|b_y(w) - b_y(e)\| \le \beta\eta\sqrt{2R}\{\sqrt{2R} + \|e + a\|\} =: \Lambda_R.$$

Thinning a homogeneous Poisson process of rate Λ_R therefore constructs the process from every point of K_R. Since every initial state lies in some K_R, this proves well-posedness and non-explosion.

We prove skew balance directly on paths. The idea is to pair every path with its time reversal after flipping momentum: collision factors agree one by one, while the change in target density exactly cancels the change in the survival probability. Write $z = (w,p)$, set

$$\Psi(w) = V(y + \sqrt{\eta}w) - a^\top w, \qquad \nu = q_y^w \otimes \gamma_d,$$

so that the Lebesgue density of ν is, up to normalization, $\varrho(z) = \exp\{-H(z) - \Psi(w)\}$. Along a flight, $\frac{\mathrm{d}}{\mathrm{d}u}\Psi(w_u) = s(z_u)$, while a collision changes neither H nor Ψ. Moreover,

$$\mathsf{R}\phi_u\mathsf{R} = \phi_{-u}, \qquad \mathsf{J}\mathsf{R} = \mathsf{R}\mathsf{J}, \qquad s \circ \mathsf{J} = s \circ \mathsf{R} = -s. \tag{24}$$

Fix $t > 0$. For $\boldsymbol{t} = (t_1, \dots, t_n)$ in the ordered simplex $\Delta_n(t) = \{0 < t_1 < \cdots < t_n < t\}$, let $z_u = z_u(z_0, \boldsymbol{t})$ be the deterministic skeleton obtained by following ϕ and inserting a collision J at every t_i. On K_R, the iterated-hazard expansion of the transition operator is

$$S_t^{y,e}F(z_0) = \sum_{n\ge 0}\int_{\Delta_n(t)} \exp\Big\{-\int_0^t \lambda(z_u)\,\mathrm{d}u\Big\}\prod_{i=1}^n \lambda(z_{t_i-})F(z_t)\,\mathrm{d}\boldsymbol{t}. \tag{25}$$

The bound by Λ_R makes this series absolutely convergent.

Pair this skeleton with the one starting at $\mathsf{R}z_t$ and having ordered collision times $(t-t_n,\dots,t-t_1)$. By (24), it is the momentum-reversed time reversal of the original skeleton and ends at $\mathsf{R}z_0$. The map

$$(z_0,t_1,\dots,t_n)\longmapsto(\mathsf{R}z_t,t-t_n,\dots,t-t_1)$$

preserves Lebesgue measure. Indeed, ϕ_u and R preserve volume, while J does so by applying, for every fixed w, an orthogonal change of variables in p. This argument does not require continuity of the reflection direction at $r(w)=0$, where the event rate vanishes.

If z^- is a forward pre-collision state, the corresponding reverse pre-collision state is $\mathsf{RJ}z^-$; hence

$$\lambda(\mathsf{RJ}z^-)=\lambda(z^-).$$

Thus the collision-rate products in (25) agree. The reverse survival hazard is $\int_0^t[-s(z_u)]_+\,\mathrm{d}u$, and

$$\Psi(w_t)-\Psi(w_0)=\int_0^t s(z_u)\,\mathrm{d}u=\int_0^t\big\{[s(z_u)]_+-[-s(z_u)]_+\big\}\,\mathrm{d}u.$$

Since $H(z_t)=H(z_0)$ and $\varrho(\mathsf{R}z)=\varrho(z)$, this is exactly the survival-weight identity

$$\varrho(z_0)\exp\!\Big\{-\int_0^t[s(z_u)]_+\,\mathrm{d}u\Big\}=\varrho(\mathsf{R}z_t)\exp\!\Big\{-\int_0^t[-s(z_u)]_+\,\mathrm{d}u\Big\}. \tag{26}$$

Let ν_R be ν conditioned on K_R. Pairing paths term by term in (25) using (26) gives, for bounded measurable F,G,

$$\int FS_t^{y,e}G\,\mathrm{d}\nu_R=\int(\mathsf{R}G)S_t^{y,e}(\mathsf{R}F)\,\mathrm{d}\nu_R=\int(\mathsf{R}S_t^{y,e}\mathsf{R}F)G\,\mathrm{d}\nu_R.$$

Letting $R\to\infty$ gives the same identity under ν. Taking $F=1$ proves invariance. Jensen's inequality then makes $S_t^{y,e}$ a contraction on $L^2(\nu)$, and density of bounded functions extends the identity to $L^2(\nu)$, proving $(S_t^{y,e})^*=\mathsf{R}S_t^{y,e}\mathsf{R}$.

Finally, the recursive hazard construction is jointly Borel in (y,e,z_0): the flow, rate, integrated hazard, and reflection are Borel, and successive event times are obtained by taking level-hitting times of the integrated hazard. Non-explosion gives the measurable pointwise limit. Consequently, conditioning on an incoming-state-independent reference and applying Fubini passes both invariance and the adjoint identity to every measurable reference mixture. □

We now pass from the conditional move to one complete proximal iteration. All operators below act on $L^2(\widehat{\Pi}_\eta)$, where $\widehat{\Pi}_\eta=\Pi_\eta\otimes\gamma_d$. Let U be the pullback of the outer reflection $(x,y,p)\mapsto(x,2x-y,p)$, let $\mathsf{P}f=\mathbb{E}[f\mid Y]$, and set $\mathsf{Q}=I-\mathsf{P}$. Finally, for the ideal chain let S be the quarter-period conditional kernel after drawing $E\sim q_y$, independently of the incoming state. Then $\mathsf{SP}=\mathsf{PS}=\mathsf{P}$. The conditional mixture and the observable operator of one ideal iteration are

$$\mathsf{B}_\rho=\rho\mathsf{P}+(1-\rho)\mathsf{S}=\mathsf{P}+(1-\rho)\mathsf{SQ},\qquad \mathsf{K}=\mathsf{UB}_\rho. \tag{27}$$

The first term is the independent conditional refresh: as an operator on observables, an exact draw from $q_y\otimes\gamma_d$ is precisely P. The second term is the persistent quarter-turn.

Lemma 9 (Invariance and refresh dissipation)**.** *The operator K preserves $\widehat{\Pi}_\eta$ and is a contraction. Moreover, for every $f\in L^2(\widehat{\Pi}_\eta)$,*

$$\|f\|^2-\|\mathsf{K}f\|^2\ge\rho\|\mathsf{Q}f\|^2. \tag{28}$$

Proof. The outer reflection is a measure-preserving involution because it sends $Y-X$ to its negative. Lem. 8 shows that the conditional mixture preserves the conditional position–momentum law; hence their composition preserves $\widehat{\Pi}_\eta$. Orthogonality of $\mathsf{P}f$ and $\mathsf{Q}f$, together with contractivity of S, gives

$$\|f\|^2-\|\mathsf{K}f\|^2=\|\mathsf{Q}f\|^2-(1-\rho)^2\|\mathsf{SQ}f\|^2\ge\{1-(1-\rho)^2\}\|\mathsf{Q}f\|^2,$$

which is stronger than (28). □

C.2 How Reflection Feeds the Quarter-Turn

This subsection proves the sampling-specific estimate used by the contraction argument. Reflection first turns a function of Y into an almost affine function of the conditional position. The quarter-turn then sends that affine position component into momentum and sends the corresponding momentum component back with the opposite sign.

At a semantic level, the signal flow is

$$\begin{aligned} f(Y_+) \xrightarrow{\mathsf{U}} f(Y_-) = \mathsf{T}f + \mathsf{C}f, \qquad & \mathsf{C} := \mathsf{QUP}, \\ \mathsf{C}f \approx L_f, \qquad & \mathsf{SC}f \approx M_f, \\ \mathsf{S}M_f \approx -\mathsf{C}f, \qquad & \mathsf{U}M_f \approx M_f. \end{aligned}$$

The proposition below controls the last three errors simultaneously; the intermediate affine observable L_f makes that control possible. Refreshment then damps whatever is not represented by these signals.

Notation and the result to be proved. Let $\mathsf{P}f = \mathbb{E}[f \mid Y]$, $\mathsf{Q} = I - \mathsf{P}$, and let U denote outer reflection. On the centered Y-only subspace $\mathcal{H}_Y = \mathrm{ran}(\mathsf{P}) \cap L^2_0$, define

$$\mathsf{T} = \mathsf{PUP}, \qquad \Gamma = (I - \mathsf{T}^2)^{1/2}, \qquad g = \sqrt{\alpha\eta}.$$

Under the stationary reflection coupling, write $Y_+ = Y$, $Y_- = 2X - Y$, and $W = (X - Y_+)/\sqrt{\eta}$. Conditional on $Y_+ = y$, denote the mean and covariance of W by $\bar{w}_y$ and Σ_y, and set $D_y = I - \Sigma_y$. For a differentiable $f \in \mathcal{H}_Y$, define

$$a_f(y) = \mathbb{E}[\nabla f(Y_-) \mid Y_+ = y], \quad L_f = 2\sqrt{\eta}\, a_f(Y_+)^\top (W - \bar{w}_{Y_+}), \quad M_f = 2\sqrt{\eta}\, a_f(Y_+)^\top P.$$

The letter L_f denotes the affine position observable; M_f is its momentum counterpart. Componentwise, the definition also says $a_f = \mathsf{T}\nabla f$, where T acts coordinatewise on vector-valued functions. The roles of the main quantities are summarized as follows.

quantity	role in the transfer argument
$\mathsf{T}f$	the part of the reflected observable predictable from Y_+ alone
Γf	the scale of the conditional-position signal created by reflection
$\mathsf{C}f$	the conditional-position signal itself
L_f	the leading affine approximation to that position signal
M_f	the momentum copy of L_f produced by a quarter-turn

These formulas are initially classical for differentiable f. The proposition below is an all–L^2 statement: the closed weak-gradient lemma that follows proves that every $f \in \mathcal{H}_+$ has the required weak gradient and defines a_f, M_f by graph closure. The auxiliary observable L_f is needed only in the smooth-core calculation. In particular, the proposition does not assume that a spectral projection preserves smoothness.

Proposition 12 (Position–momentum transfer)**.** *Let $\mathcal{H}_+$ be the spectral subspace of T corresponding to $[1/2, 1]$. If $\beta\eta$ is smaller than a universal constant, then every $f \in \mathcal{H}_+$ satisfies*

$$\begin{aligned} \|\mathsf{S}(\mathsf{QUP}f) - M_f\|^2 + \|\mathsf{S}M_f + \mathsf{QUP}f\|^2 + \|\mathsf{U}M_f - M_f\|^2 \\ \le C\|\Gamma f\|\|\Gamma^3 f\|. \end{aligned} \tag{29}$$

The proof has a technical prelude followed by three substantive steps. The prelude extends estimates from smooth observables to the closed weak-gradient domain. Step 1 analyzes the signal created by reflection, Step 2 analyzes its quarter-rotation, and Step 3 combines the two estimates and performs the single closure passage.

Technical prelude: extending smooth estimates to all L^2. We first isolate the only closure argument needed below. A graph core here means that every function with a weak gradient can be approximated together with its gradient by smooth compactly supported functions. Let $\nu = \Pi^Y_\eta$, let $D : \mathrm{Dom}(D) \subset L^2_0(\nu) \to L^2(\nu; \mathbb{R}^d)$ be the closed weak gradient, and put

$$\mathcal{C}_0 = C^\infty_c(\mathbb{R}^d) \cap L^2_0(\nu).$$

Lemma 10 (Closed weak-gradient smoothing)**.** *The set $\mathcal{C}_0$ is a graph core for D. Suppose that T is a self-adjoint contraction, $\mathsf{T}\mathcal{C}_0 \subset \mathrm{Dom}(D)$, and, on $\mathcal{C}_0$,*

$$\sqrt{\eta}\|D\mathsf{T}u\| \le c_D\|\Gamma u\|. \tag{30}$$

Then $D\mathsf{T}$ has an everywhere-defined bounded extension $\mathsf{R}_D : L^2_0(\nu) \to L^2(\nu;\mathbb{R}^d)$, and in fact

$$\mathsf{T}L^2_0(\nu) \subset \mathrm{Dom}(D), \qquad D\mathsf{T} = \mathsf{R}_D, \qquad \sqrt{\eta}\|\mathsf{R}_D u\| \le c_D\|\Gamma u\|.$$

Consequently, if $\Pi_+ = \mathbf{1}_{[1/2,1]}(\mathsf{T})$, then $\mathrm{ran}(\Pi_+) \subset \mathrm{Dom}(D)$ and

$$\sqrt{\eta}\|Df\| \le 2c_D\|\Gamma f\|,$$
$$\sqrt{\eta}\|D(I-\mathsf{T})f\| \le 2c_D\|\Gamma^3 f\|, \qquad f \in \mathrm{ran}(\Pi_+).$$

Moreover, for every $f \in \mathrm{Dom}(D)$,

$$D(I-\mathsf{T})f = Df - \mathsf{R}_D f. \tag{31}$$

Thus convergence in the graph norm of D automatically implies convergence of $D(I-\mathsf{T})f$; no additional graph norm is needed.

Proof. The density of ν is positive and continuous, hence is bounded above and below on every compact set. Cutoff followed by ordinary convolution therefore approximates every element of $\mathrm{Dom}(D)$ in the norm $\|f\|^2 + \|Df\|^2$. To preserve centering, subtract the mean of each approximant times a fixed $\psi \in C_c^\infty$ satisfying $\int \psi \, \mathrm{d}\nu = 1$. This proves the graph-core assertion.

If $u_n \in \mathcal{C}_0$ converges in $L^2(\nu)$, (30) makes $D\mathsf{T}u_n$ Cauchy, while $\mathsf{T}u_n$ converges. Closedness of D shows that the limit belongs to $\mathrm{Dom}(D)$, identifies its gradient, and gives the bounded extension $\mathsf{R}_D = D\mathsf{T}$.

On $\mathrm{ran}(\Pi_+)$, T^{-1} is bounded with norm at most two. For f in that subspace,

$$Df = \mathsf{R}_D\mathsf{T}^{-1}f, \qquad D(I-\mathsf{T})f = \mathsf{R}_D\mathsf{T}^{-1}(I-\mathsf{T})f.$$

The first estimate follows from $\gamma_t/t \le 2\gamma_t$ on $t \in [1/2,1]$, where $\gamma_t = (1-t^2)^{1/2}$. The second follows from

$$\frac{1-t}{t} = \frac{1-t^2}{t(1+t)} \le 2(1-t^2), \qquad t \in [1/2,1].$$

Finally, $\mathsf{T}f \in \mathrm{Dom}(D)$ for every $f \in L^2_0(\nu)$, so (31) follows by linearity and proves the last assertion. □

C.2.1 Step 1: What the Outer Reflection Creates

Lemma 11 (Outer-reflection geometry)**.** *Assume $\beta\eta \le 1/3$. On $\mathcal{H}_Y$,*

$$-\tfrac{1}{2}I \preceq \mathsf{T} \preceq \frac{1-g^2}{1+g^2}I, \qquad \Gamma \succeq \frac{2g}{1+g^2}I.$$

Moreover, $\mathsf{C} := \mathsf{QUP}$ can be written as $\mathsf{C} = \mathsf{V}\Gamma$, where V is an isometry, and

$$\mathsf{UP}f = \mathsf{PT}f + \mathsf{V}\Gamma f, \qquad \mathsf{UV}f = \mathsf{P}\Gamma f - \mathsf{VT}f.$$

On $\mathcal{H}_+$,

$$\sqrt{\eta}\|\nabla f\| \le C\|\Gamma f\|, \qquad \sqrt{\eta}\|\nabla(I-\mathsf{T})f\| \le C\|\Gamma^3 f\|.$$

Proof. The pair (Y_+, Y_-) is exchangeable and $(\mathsf{T}f)(y) = \mathbb{E}[f(Y_-) \mid Y_+ = y]$, so T is a self-adjoint contraction. The conditional density of Y_- given $Y_+ = y$ is proportional to

$$\exp\left\{-V\left(\frac{y+u}{2}\right) - \frac{\|y-u\|^2}{8\eta}\right\}.$$

Its score with respect to y is $s_y(u) = -\nabla V((y+u)/2)/2 - (y-u)/(4\eta)$. For $f \in \mathcal{C}_0$, weak differentiation under the integral gives

$$D\mathsf{T}f(y) = \mathrm{Cov}(f(Y_-), s_y(Y_-) \mid Y_+ = y).$$

Indeed, on compact sets of y, the score times the conditional density is dominated by an integrable Gaussian tail. Thus Fubini and weak integration by parts justify differentiation of both the numerator and normalizing constant. In particular, $\mathsf{T}\mathcal{C}_0 \subset \mathrm{Dom}(D)$. Applying conditional Poincaré to this covariance gives

$$\eta\|D\mathsf{T}f\|^2 \le \frac{(1-\alpha\eta)^2}{4(1+\alpha\eta)}\{\|f\|^2 - \|\mathsf{T}f\|^2\} \le \frac{1}{4}\|\Gamma f\|^2.$$

The Y-marginal has Poincaré constant at most $\alpha^{-1}+\eta$: this follows from the conditional covariance bound $\mathrm{Cov}(X \mid Y) \preceq \eta(1+\alpha\eta)^{-1}I$ and Brascamp–Lieb (Chewi, 2023). Equivalently, if W_η is the negative log-density of the Y-marginal, then

$$\nabla^2 W_\eta(y) = \eta^{-1}I - \eta^{-2}\mathrm{Cov}(X \mid Y=y) \succeq \frac{\alpha}{1+\alpha\eta}I.$$

Applying the resulting Poincaré inequality to the centered function $\mathsf{T}f$ yields $\|\mathsf{T}\|_{L^2_0} \le (1-g^2)/(1+g^2)$. This inequality is first obtained on $\mathcal{C}_0$ and extends to $L^2_0(\nu)$ by density.

For the lower edge, exchangeability and Gaussian Poincaré imply

$$\langle h, (I-\mathsf{T})h\rangle = \tfrac{1}{2}\mathbb{E}\{h(Y_+) - h(Y_-)\}^2 \le 2\eta\|\nabla h\|^2.$$

Take $h = \mathsf{T}f$ for $f \in \mathcal{C}_0$ and use the preceding derivative bound. The resulting inequality involves only bounded polynomials in T, so density extends it to all of $L^2_0(\nu)$. Spectral calculus gives $2\mathsf{T}^2(I-\mathsf{T}) \preceq I - \mathsf{T}^2$, equivalently $(I-\mathsf{T})^2(I+2\mathsf{T}) \succeq 0$, and hence $\mathsf{T} \succeq -I/2$. The two spectral edges imply the lower bound on Γ. Since $\mathsf{T}1 = 1$, the lower bound $\mathsf{T} \succeq -I/2$ also holds on the full scalar $L^2(\nu)$, and hence coordinatewise on vector fields such as Df, whose components need not be centered.

Finally, $\mathsf{C}^*\mathsf{C} = I - \mathsf{T}^2 = \Gamma^2$. Thus $\mathsf{V} = \mathsf{C}\Gamma^{-1}$ is an isometry, and $\mathsf{U}^2 = I$ gives the two block identities.

The core derivative estimate above is (30) with $c_D = 1/2$. Applying Lem. 10 proves that all of $\mathcal{H}_+$, not merely a smooth subset of it, belongs to the weak gradient domain and gives the two final inequalities. Notice that this argument never applies the discontinuous spectral projection Π_+ to a smooth function. □

Lemma 12 (Affine approximation after reflection)**.** *For every $f \in \mathcal{C}_0$, $\mathsf{QUP}f = L_f + e_f$, where $\mathbb{E}[e_f \mid Y_+] = 0$. With*

$$\mathcal{R}_f = \|\nabla f\|^2 - \|\mathsf{T}\nabla f\|^2, \qquad \mathcal{D}_f = \mathbb{E}[a_f(Y_+)^\top D_{Y_+} a_f(Y_+)],$$

one has

$$\|e_f\|^2 \le 4\eta\mathcal{R}_f, \qquad \mathcal{R}_f + \mathcal{D}_f \le C\|\nabla(I-\mathsf{T})f\|\|\nabla f\|. \tag{32}$$

Proof. Conditionally on $Y_+ = y$, we have $Y_- = y + 2\sqrt{\eta}W$. Set $e_f = \mathsf{QUP}f - L_f$. The choice of $a_f(y)$ makes both $\mathbb{E}[e_f \mid Y_+ = y]$ and $\mathbb{E}[\nabla_W e_f \mid Y_+ = y]$ zero. The conditional position law is one-strongly log-concave in W, so conditional Poincaré gives the first inequality.

For the second, let $b_y(w) = \sqrt{\eta}\,\nabla V(y + \sqrt{\eta}w)$. Weak integration by parts under the conditional density gives

$$\mathrm{Cov}(W, b_y(W) \mid y) = D_y, \qquad \mathrm{Cov}(b_y(W) \mid y) \preceq \frac{\beta\eta}{1-\beta\eta}D_y.$$

Conditional Poincaré applied to linear functions gives $\Sigma_y \preceq I$, while the Stein identity $\mathbb{E}[\nabla b_y] = D_y + \mathrm{Cov}(b_y)$ and $0 \preceq \nabla b_y \preceq \beta\eta I$ give $0 \preceq D_y \preceq \beta\eta I$. The second relation follows by applying conditional Poincaré to b_y and using $(\nabla b_y)^2 \preceq (\beta\eta)\nabla b_y$. These are standard Brascamp–Lieb/Stein calculations (Chewi, 2023); for a $C^{1,1}$ potential they follow by mollification and cutoff.

Differentiating $\mathsf{T}f$ weakly and inserting $\mathsf{QUP}f = L_f + e_f$ yields

$$\nabla(I-\mathsf{T})f = \nabla f - a_f + 2D_y a_f + \eta^{-1/2}\mathrm{Cov}(e_f, b_y(W) \mid Y_+ = y).$$

Take its inner product with $\nabla f + a_f$ and average over y. Using $a_f = \mathsf{T}\nabla f$, exact rearrangement gives

$$\mathcal{R}_f + 4\mathcal{D}_f = \langle \nabla(I-\mathsf{T})f, \nabla f + a_f\rangle - 2\mathbb{E}\big[a_f^\top D_{Y_+}(\nabla f - a_f)\big] \\ - \eta^{-1/2}\mathbb{E}\big[\mathrm{Cov}(e_f, b_{Y_+}(W)\mid Y_+)^\top(\nabla f + a_f)\big].$$

Cauchy–Schwarz, the two conditional covariance bounds above, and the first inequality control the right-hand side by

$$2\|\nabla(I-\mathsf{T})f\|\|\nabla f\| + C\sqrt{\beta\eta\,\mathcal{R}_f\mathcal{D}_f} + C\sqrt{\beta\eta\,\mathcal{R}_f(\beta\eta\,\mathcal{R}_f + \mathcal{D}_f)}.$$

For sufficiently small $\beta\eta$, Young's inequality absorbs the last two terms and proves (32). □

C.2.2 Step 2: What the Quarter-Turn does to Affine Information

At fixed y, write $\ell_{a,y}(w) = a^\top(w - \bar{w}_y)$ and $m_a(p) = a^\top p$. Let S denote the quarter-period kernel after drawing an independent exact reference from the conditional position law. All norms and expectations in the next lemma are under $q_y^w \otimes \gamma_d$, including the independent reference randomness.

Lemma 13 (Approximate quarter rotation)**.** *If $L_\eta := \beta\eta \le 1/2$, then, uniformly in y,*

$$\|\mathsf{S}\ell_{a,y} - m_a\|^2 + \|\mathsf{S}m_a + \ell_{a,y}\|^2 \le \frac{22L_\eta}{1-L_\eta}a^\top D_y a.$$

Proof. For fixed reference e, the generator in (23) maps $\ell_{a,y}$ to m_a and maps m_a to $-\ell_{a,y} + k_{a,e}$, where, with $r = b_y(w) - b_y(e)$ and $Z = p^\top r/\|r\|$,

$$k_{a,e} = \mathbb{E}[a^\top b_y(W)\mid y] - a^\top b_y(e) - 2a^\top r\, Z^2\mathbf{1}_{\{Z>0\}}.$$

When $r = 0$, only the final term is defined to be zero; the first two terms are unchanged. The localized Dynkin formula therefore gives, at $T = \pi/2$,

$$S_T^{y,e}\ell_{a,y} - m_a = \int_0^T S_s^{y,e}k_{a,e}\cos s\ \mathrm{d}s, \qquad S_T^{y,e}m_a + \ell_{a,y} = \int_0^T S_s^{y,e}k_{a,e}\sin s\ \mathrm{d}s.$$

This is the variation-of-constants identity for the two observable equations (equivalently, a localized Dynkin formula). Localization is legitimate because the harmonic energy in the proof of Lem. 8 is conserved; dominated convergence then removes the cutoff.

When e is an independent conditional draw, elementary Gaussian moments give $\mathbb{E}k_{a,e}^2 = 11\mathrm{Var}(a^\top b_y(W)\mid y)$. The covariance calculation in the proof of Lem. 12 bounds this variance by $L_\eta(1-L_\eta)^{-1}a^\top D_y a$. Contractivity of $S_s^{y,e}$ and Cauchy–Schwarz in time prove the claimed bound. □

C.2.3 Step 3: Assemble the Global Transfer Estimate

The single graph closure. For every centered $f \in \mathrm{Dom}(D)$, define

$$a_f = \mathsf{T}Df, \qquad M_f = 2\sqrt{\eta}\,a_f(Y)^\top P, \qquad \mathsf{C} = \mathsf{QUP}.$$

This agrees with the preceding definition on smooth functions. Since P is conditionally standard Gaussian and T is a contraction on vector-valued L^2, for $f, h \in \mathrm{Dom}(D)$,

$$\|M_f - M_h\|^2 = 4\eta\|\mathsf{T}D(f-h)\|^2 \le 4\eta\|D(f-h)\|^2. \tag{33}$$

Consider the stacked error map from $\mathrm{Dom}(D)$, equipped with the graph norm $\|f\|_D^2 = \|f\|^2 + \eta\|Df\|^2$, into three copies of the stationary L^2 space:

$$\mathbf{E}f = (\mathsf{SC}f - M_f, \mathsf{S}M_f + \mathsf{C}f, \mathsf{U}M_f - M_f).$$

The operators $\mathsf{C}, \mathsf{S}, \mathsf{U}$ are contractions, so (33) makes $\mathbf{E}$ graph-norm continuous. Moreover, (31) shows that $f \mapsto D(I-\mathsf{T})f$ is graph-norm continuous. These two facts will pass the final three-error inequality through the core in one step; none of its intermediate conditional error quantities needs a separate closure argument.

Proof of Prop. 12. First let $f \in \mathcal{C}_0$. Apply Lem. 13 conditionally with coefficient $2\sqrt{\eta}\, a_f(y)$, then add the remainder e_f from Lem. 12. The first two terms on the left of (29) are bounded by $C\{\beta\eta^2\mathcal{D}_f + \|e_f\|^2\}$.

For the third term, reflection exchanges Y_+ and Y_- without changing momentum. Since $\mathsf{T} \succeq -I/2$, $\|\mathsf{U}M_f - M_f\|^2 \le 4\eta\mathcal{R}_f$. (32) now bounds the sum of all three terms by

$$C\eta\|\nabla(I-\mathsf{T})f\|\|\nabla f\|.$$

Equivalently, on $\mathcal{C}_0$ we have the pre-spectral estimate

$$\|\mathbf{E}f\|^2 \le C\eta\|D(I-\mathsf{T})f\|\|Df\|. \tag{34}$$

Now take an arbitrary $f \in \mathcal{H}_+$. By Lem. 10, $f \in \mathrm{Dom}(D)$. Because $\mathcal{C}_0$ is a graph core, there are $f_n \in \mathcal{C}_0$ with $f_n \to f$ and $Df_n \to Df$. The approximants are not required to lie in $\mathcal{H}_+$. The graph continuity proved above gives

$$\mathbf{E}f_n \to \mathbf{E}f, \qquad D(I-\mathsf{T})f_n \to D(I-\mathsf{T})f.$$

Passing to the limit in (34) and only then applying the two sector derivative estimates in Lem. 11 proves (29). Thus no density assertion about $\mathcal{C}_0 \cap \mathcal{H}_+$, and no smoothness preservation by the spectral projection, is used. □

A bounded form used below. The transfer proposition is now complete. We next normalize its position and momentum signals to the same scale so that the contraction proof can treat them as two coordinates of one approximate rotation. The isometry V is the normalized position channel, A will be the normalized momentum channel, and F_1, F_2, F_3 record the three departures from an exact quarter-rotation. The later factorization $F_j = B_j\Gamma$ says that these errors shrink with the amount of signal being transferred.

On $\mathcal{H}_+$, the definition above and Lem. 11 make $f \mapsto M_f$ a bounded operator M: indeed $\|M_f\|^2 \le 4\eta\|\nabla f\|^2 \le C\|\Gamma f\|^2$. Put $\mathsf{A} = \mathsf{M}\Gamma^{-1}$, and recall $\mathsf{QUP} = \mathsf{V}\Gamma$. A is therefore bounded on $\mathcal{H}_+$. Define F_1, F_2, F_3 there by

$$\mathsf{SV} = \mathsf{A} + F_1, \qquad \mathsf{SA} = -\mathsf{V} + F_2, \qquad \mathsf{UA} = \mathsf{A} + F_3. \tag{35}$$

Position observables are even in momentum and momentum observables are odd, so the ranges of V and A are orthogonal. Momentum reversal in Lem. 8 gives the adjoint versions of these relations. When an adjoint is taken in the next subsection, $\mathsf{V}, \mathsf{A}, F_j$ mean their restrictions to $\mathcal{H}_+$, extended by zero on its orthogonal complement. Thus all subsequent adjoints are bounded, everywhere-defined Hilbert-space operators.

Corollary 2 (Controlled transfer errors)**.** *Let $\ell_g = 1 + \log(1/g)$. There are bounded operators B_j such that*

$$F_j = B_j\Gamma, \qquad F_j^* = \Gamma B_j^*, \qquad \|B_j\| \le C\sqrt{\ell_g}, \qquad j = 1,2,3. \tag{36}$$

Proof. (29), with f replaced by $\Gamma^{-1}f$, gives $\|F_jf\|^2 \le C\|f\|\|\Gamma^2 f\|$. Partition the spectrum of Γ^2 into dyadic intervals between cg^2 and 1. On each spectral component f_r, $\|F_jf_r\|^2 \le C\langle f_r, \Gamma^2 f_r\rangle$. There are $O(\ell_g)$ components, so the triangle inequality followed by Cauchy–Schwarz gives

$$F_j^*F_j \preceq C\ell_g\Gamma^2.$$

Since Γ is invertible on the centered subspace, setting $B_j = F_j\Gamma^{-1}$ proves (36). The factor order is important; no commutation between B_j and Γ is used. □

C.3 From Position–Momentum Transfer to Contraction

We now prove the ideal-chain bound. Reflection alone preserves norm, and the quarter-turn mainly rotates information between position and momentum, so a one-step Poincaré estimate cannot reveal the contraction. Instead, we rule out oscillations on the unit circle and then convert that bound into geometric decay.

The resolvent strategy asks whether $(z - \mathsf{K})f = b$ can have a large solution when $|z| = 1$. For each slowly varying Y-mode, the maps V and A expose its position and momentum companions; the

resulting three coordinates form an approximate finite-dimensional rotation. The refresh dissipation estimate controls the orthogonal conditional variation that these coordinates do not see. A uniform bound on the unit-circle resolvent then yields geometric decay through Lem. 15.

Proposition 13 (Ideal persistent contraction)**.** *Under the assumptions of Thm. 3, choosing $\rho \asymp g\sqrt{\ell_g}$, where $g = \sqrt{\alpha\eta}$ and $\ell_g = 1 + \log(1/g)$, gives*

$$\|\mathsf{K}^n\|_{L_0^2 \to L_0^2} \le 2\exp\left(-\frac{cng}{\sqrt{\ell_g}}\right).$$

The proof has three steps. We follow three slowly changing coordinates, invert their three-by-three update on the unit circle, and use refreshment to control all remaining conditional variation. The coordinates do not form a decomposition of the full state space; this is why the final refreshment step is essential.

We allow complex-valued observables below so that z may range over the complex unit circle; this does not change any operator norm. Let $k = 1 - \rho$, and let K be the one-step operator in (27). Recall $\mathsf{T}, \Gamma, \mathsf{V}, \mathsf{A}$ from Lem. 11 and (35).

Lemma 14 (Unit-circle resolvent bound)**.** *If $\rho \asymp g\sqrt{\ell_g}$, then*

$$\sup_{|z|=1} \|(z - \mathsf{K})^{-1}\|_{L_0^2 \to L_0^2} \le \frac{C\sqrt{\ell_g}}{g}.$$

Proof. Step 1: isolate the three slow coordinates. Fix $|z| = 1$ and solve $(z - \mathsf{K})f = b$. Write $f_\perp = \mathsf{Q}f$. Let $\Pi_{\mathrm{slow}} = \mathbf{1}_{[1/2,1]}(\mathsf{T})$, and split the Y-only component into $x_{\mathrm{slow}} = \Pi_{\mathrm{slow}}\mathsf{P}f$ and $x_{\mathrm{fast}} = (I - \Pi_{\mathrm{slow}})\mathsf{P}f$. In the three-coordinate calculation, every occurrence of $\mathsf{T}, \Gamma, \mathsf{V}, \mathsf{A}, F_j$ is restricted to $\mathcal{H}_+ = \mathrm{ran}(\Pi_{\mathrm{slow}})$; equivalently, V means $\mathsf{V}\Pi_{\mathrm{slow}}$, and similarly for the other maps. The other two tested coordinates are then $v = \mathsf{V}^* f_\perp$ and $m = \mathsf{A}^* f_\perp$.

coordinate	information it measures
x_{slow}	the slowly changing component stored in the auxiliary variable Y
$v = \mathsf{V}^* f_\perp$	conditional variation aligned with the position signal
$m = \mathsf{A}^* f_\perp$	conditional variation aligned with the momentum signal
x_{fast}	the remaining Y-modes, controlled directly by the spectral gap of $z - \mathsf{T}$

The next identities encode the approximate quarter-rotation in these coordinates; their error terms are precisely the factorized F_j's from the preceding subsection. The adjoint identities used below are

$$\mathsf{V}^*\mathsf{S}u = -\mathsf{A}^*u + F_1^*\mathsf{R}u, \quad \mathsf{A}^*\mathsf{S}u = \mathsf{V}^*u - F_2^*\mathsf{R}u, \quad \mathsf{V}^*\mathsf{U}\mathsf{Q} = -\mathsf{T}\mathsf{V}^*, \quad \mathsf{A}^*\mathsf{U} = \mathsf{A}^* + F_3^*.$$

They follow from (35), momentum reversal, and the even/odd momentum parity recorded there. No x_{fast} term is omitted below: spectral orthogonality gives $\mathsf{V}^*\mathsf{U}\mathsf{P}x_{\mathrm{fast}} = 0$, while momentum parity gives $\mathsf{A}^*\mathsf{U}\mathsf{P}x_{\mathrm{fast}} = 0$ and $F_3^*\mathsf{P}x_{\mathrm{fast}} = 0$.

Projecting the resolvent equation onto these coordinates gives

$$\begin{aligned} (z - \mathsf{T})x_{\mathrm{slow}} + k\Gamma m &= b_0 + k\Gamma r_1, \\ zv - \Gamma x_{\mathrm{slow}} - k\mathsf{T}m &= b_V - k\mathsf{T}r_1, \\ zm - kv &= b_A + kr_2, \end{aligned} \tag{37}$$

where $b_0 = \Pi_{\mathrm{slow}}\mathsf{P}b$, $b_V = \mathsf{V}^*b$, $b_A = \mathsf{A}^*b$, and

$$r_1 = F_1^*\mathsf{R}f_\perp, \qquad r_2 = -F_2^*\mathsf{R}f_\perp + F_3^*\mathsf{S}f_\perp.$$

Cor. 2 implies

$$\|\Gamma^{-1}r_1\| + \|\Gamma^{-1}r_2\| \le C\sqrt{\ell_g}\|f_\perp\|.$$

This is the only logarithmic loss. The factorization $F_j^* = \Gamma B_j^*$ is used in this order; no commutation is assumed.

Step 2: invert the three-coordinate system. All coefficients on the left of (37) are functions of the single self-adjoint operator T. At a scalar spectral value $t \in [1/2, 1)$, put $\gamma_t = \sqrt{1-t^2}$. The determinant of the coefficient matrix is

$$d_{t,k}(z) = z^3 - tz^2 - k^2 tz + k^2.$$

Taking the first row of the adjugate gives the exact operator identity

$$\begin{aligned} d_{\mathsf{T},k}(z) x_{\mathrm{slow}} = (z^2 - k^2\mathsf{T})b_0 - k^2\Gamma b_V - kz\Gamma b_A \\ + kz^2\Gamma r_1 - k^2 z\Gamma r_2. \end{aligned} \tag{38}$$

The cancellation that produces the final two terms in this form is what prevents an extra factor g^{-1}.

We use two elementary scalar estimates. If $|z - t| < \gamma_t/16$, then

$$|d_{t,k}(z)| \geq \frac{15}{32}\gamma_t^2, \qquad |z^2 - k^2 t| \leq 4(\rho + \gamma_t).$$

For the first, write $z = e^{2iu}$. The near condition gives $|u| \leq \gamma_t/16$, while the real part of $e^{-3iu} d_{t,k}(e^{2iu})$ is $(1+k^2)\{\cos(3u) - t\cos u\}$. The identity $\cos(3u) - t\cos u = (1-t)\cos u - 2\sin(2u)\sin u$, together with $1 - t = \gamma_t^2/(1+t)$, proves the lower bound. The second estimate follows from $|z^2 - k^2 t| \leq 2|z-1| + (1-t) + (1-k^2)$.

On the spectral region $|z-t| \geq \gamma_t/16$, the first row of (37) directly controls x_{slow}. On its complement, apply the two scalar estimates to (38). The terms involving b cost at most $C(\rho/g^2 + g^{-1})\|b\|$, while (36) bounds the two error terms as follows. If E denotes this near spectral region, then

$$\|\mathbf{1}_E d_{\mathsf{T},k}(z)^{-1}\Gamma F_j^* q\| = \|\mathbf{1}_E d_{\mathsf{T},k}(z)^{-1}\Gamma^2 B_j^* q\| \leq C\sqrt{\ell_g}\|q\|.$$

Here the lower bound on $d_{t,k}$ cancels the two powers of γ_t supplied by Γ^2. Since $\Gamma \succeq cgI$ and $\rho \asymp g\sqrt{\ell_g}$, both regions yield

$$\|\mathsf{P}f\| \leq C\sqrt{\ell_g}\left(\|f_\perp\| + \frac{\|b\|}{g}\right). \tag{39}$$

Here the fast component has not been dropped: on $[-1/2, 1/2)$, $z - \mathsf{T}$ has inverse norm at most two, so $\|x_{\mathrm{fast}}\| \leq 2(\|b\| + \|f_\perp\|)$.

Step 3: control everything not tested above. (28) and $\mathsf{K}f = zf - b$ imply $\rho\|f_\perp\|^2 \leq 2\|f\|\|b\|$. Combine this with (39), use $f = \mathsf{P}f + f_\perp$, and apply Young's inequality. The result is

$$\|f\| \leq \frac{C\sqrt{\ell_g}}{g}\|b\|.$$

Thus $z - \mathsf{K}$ is bounded below and has closed range.

It is also onto. Otherwise some nonzero u would satisfy $\mathsf{K}^* u = \bar{z}u$. Because K is a contraction and $|z| = 1$, equality of norms implies $\mathsf{K}\mathsf{K}^* u = u$, and hence $\mathsf{K}u = zu$. This contradicts the lower bound at z, proving the asserted resolvent estimate. □

The next elementary Hilbert-space lemma converts the resolvent estimate into a bound on powers. We retain its short proof because the precise linear dependence on M is needed for the final rate.

Lemma 15 (Unit-circle bound implies geometric decay)**.** *Let Q be a contraction on a Hilbert space. If $\sup_{|z|=1}\|(z-Q)^{-1}\| \leq M$ with $M \geq 1$, then $\|Q^n\| \leq 2\exp(-cn/M)$.*

Proof. For $a_j = \sin((j+1)\pi/(N+2))$, apply the resolvent bound to $\sum_{j=0}^{N} a_j z^{-j-1} Q^j f$ and integrate over the unit circle. Parseval's identity gives

$$\sum_{j=0}^{N} a_j^2 \|Q^j f\|^2 \leq M^2 \sum_{j=0}^{N+1} (a_j - a_{j-1})^2 \|Q^j f\|^2, \qquad a_{-1} = a_{N+1} = 0.$$

Contractivity and $\sum(a_j - a_{j-1})^2 = 2\{1 - \cos(\pi/(N+2))\}\sum a_j^2$ imply $\|Q^N\| \leq \pi M/(N+2)$. Taking $N_0 = \lceil 2\pi M \rceil$ gives $\|Q^{N_0}\| \leq 1/2$; writing $n = qN_0 + r$ proves the claim. □

Proof of Prop. 13. Lem. 9, (35), and Cor. 2 provide the invariance, dissipation, row identities, and controlled errors used in Lem. 14; these ultimately follow from Lem. 8 and Prop. 12. Apply Lem. 15 with $M = C\sqrt{\ell_g}/g$. This proves (6). □

Total-variation conclusion in Thm. 3. Initialize $Y_0 \mid X_0 \sim \mathcal{N}(X_0, \eta I_d)$ and $P_0 \sim \gamma_d$. The initial and stationary augmented laws have the same conditional law of (Y, P) given X, so augmentation preserves D_2. If r_0 is the resulting density ratio, then

$$\mathrm{TV}(\mathrm{Law}(X_n, Y_n, P_n), \widehat{\Pi}_\eta) \le \frac{1}{2}\|(\mathsf{K}^*)^n(r_0 - 1)\| \le \frac{1}{2}e^{\Delta/2}\|\mathsf{K}^n\|_{L_0^2 \to L_0^2}.$$

Substitution of the contraction estimate gives (7); marginalization proves the claim for X_n. □

C.4 Implementable Algorithm and Warm-Start Accounting

We begin with the result proved in this subsection and then verify its four implementation components: the shifted Gaussian reference, a deterministic iteration cap for the mode solver, a capped event intensity simulated by Poisson thinning, and the approximate RGO used on the RGO branch.

Theorem 6 (Implementation theorem)**.** *Let $D_2(\nu_0\|\pi) \le \Delta$, $0 < \varepsilon \le 1/2$, and define*

$$\Lambda_\varepsilon = \Delta + \log\frac{Cd\kappa}{\varepsilon}, \qquad \eta = \frac{c}{\beta(\sqrt{d\Lambda_\varepsilon} + \Lambda_\varepsilon)}.$$

Put $g = \sqrt{\alpha\eta}$, $\ell_g = 1 + \log(1/g)$, choose $\rho \asymp g\sqrt{\ell_g}$, and set

$$N = \left\lceil \frac{C\sqrt{\ell_g}}{g}\{\Delta + \log(8/\varepsilon)\} \right\rceil.$$

Initialize $X_0 \sim \nu_0$, $Y_0 \mid X_0 \sim \mathcal{N}(X_0, \eta I_d)$, and $P_0 \sim \gamma_d$, independently conditional on X_0. There is an implementable algorithm for the persistent iteration in Sec. 5, using Alg. 3 as its accelerated branch, that returns X_N with $\mathrm{TV}(\mathrm{Law}(X_N), \pi) \le \varepsilon$ and

$$\mathbb{E}N_{\mathrm{qry}} \le \widetilde{O}\Big[\sqrt{\kappa}\,(\Delta + \log(8/\varepsilon))(\sqrt{d\Lambda_\varepsilon} + \Lambda_\varepsilon)^{1/2}\Big], \qquad \mathbb{E}N_{\mathrm{RGO}} \le \widetilde{O}(\Delta + \log(8/\varepsilon)).$$

C.4.1 A Comparison Principle for the Complete Path

The chain induced by the implementable algorithm is not invariant because it contains truncation and approximation. We therefore keep the exact chain as a comparison process and couple the two chains until their first disagreement.

Lemma 16 (Warm path comparison)**.** *Let an exact Markov kernel preserve π, and let $\nu \ll \pi$ satisfy $D_2(\nu\|\pi) \le \Delta$. If A is an event determined by a Markov path together with all its internal randomness, then*

$$\Pr_\nu(A) \le e^{\Delta/2}\Pr_\pi(A)^{1/2}.$$

Consequently, if the stationary one-step probability that an exact and an approximate transition disagree is at most a, their N-step laws started from ν can be coupled with disagreement probability at most $Ne^{\Delta/2}\sqrt{a}$.

Proof. For $R = \mathrm{d}\nu/\mathrm{d}\pi$, conditional Cauchy–Schwarz gives $\mathbb{E}_\pi[R\Pr(A \mid X_0)] \le \|R\|_2 \Pr_\pi(A)^{1/2}$. An invariant exact kernel cannot increase D_2. Apply the same bound before each exact transition and take a union bound over the first disagreement. □

C.4.2 Replacing the Ideal Exact-Conditional Reference on the Accelerated Branch

The ideal contraction proof draws $E \sim q_y$ on the accelerated branch, but this draw is used only as an independent reference for the event-driven motion; it is distinct from the actual refresh on the RGO branch. The finite algorithm replaces this ideal reference by a Gaussian centered near the

conditional mode. The next comparison lemma controls the resulting normalized shift and is used only to re-establish the local quarter-turn error bound.

Fix y, let $x_*(y)$ minimize the conditional potential, and set $z = (x - x_*(y))/\sqrt{\eta}$. After removing a constant and a linear term, the normalized conditional density has the form

$$q(z) \propto e^{-\|z\|^2/2 - r(z)}, \qquad mI \preceq \nabla^2 r \preceq LI,$$

where $m = \alpha\eta$, $L = \beta\eta$, and $\nabla r(0) = 0$.

Lemma 17 (Shifted-Gaussian comparison)**.** *Let $\gamma_\delta = \mathcal{N}(\delta, I_d)$, and set $\tau = 2 + \log(L/m)$. If $\tau L \le c$ and $\tau(L^2 d + \|\delta\|^2) \le c$, then*

$$\left\| \frac{\mathrm{d}\gamma_\delta}{\mathrm{d}q} \right\|_{L^\tau(q)} \le e^{C\tau(L^2 d + \|\delta\|^2)}, \qquad \chi^2(\gamma_\delta \| q) \le C(L^2 d + \|\delta\|^2).$$

Proof. Normalize $r(0) = 0$. Convexity and L-smoothness give $0 \le r(z) \le L\|z\|^2/2$ and $\|\nabla r(z)\|^2 \le 2Lr(z)$. The Gaussian log-Sobolev inequality and the standard entropy argument (Chewi, 2023) imply, for $0 \le tL < 1$,

$$\log \mathbb{E}_\gamma e^{t(r - \mathbb{E}_\gamma r)} \le \frac{t^2 L^2 d}{2(1 - tL)}, \qquad \log \mathbb{E}_\gamma e^{-t(r - \mathbb{E}_\gamma r)} \le \frac{t^2 L^2 d}{2}.$$

For completeness, the second inequality is not obtained by simply changing the sign in the first. If $A_-(t) = \log \mathbb{E}_\gamma e^{-t(r - \mathbb{E}_\gamma r)}$, the same entropy calculation gives $tA'_-(t) - A_-(t) \le t^2 L\{\mathbb{E}_\gamma r - A'_-(t)\}$. Convexity and $A_-(0) = A'_-(0) = 0$ imply $A'_-(t) \ge 0$, so integrating $(A_-(t)/t)' \le L\mathbb{E}_\gamma r \le L^2 d/2$ proves the displayed negative-tilt bound. Applying these bounds to $\mathrm{d}\gamma/\,\mathrm{d}q$ gives $\log \mathbb{E}_q(\,\mathrm{d}\gamma/\,\mathrm{d}q)^\tau \le C\tau^2 L^2 d$. For the shift, $\mathrm{d}\gamma_\delta/\,\mathrm{d}\gamma = \exp\{\delta^\top z - \|\delta\|^2/2\}$. The law q is one-strongly log-concave, so its linear functionals are one-subgaussian after centering. Weak integration by parts gives $\mathbb{E}_q Z = -\mathbb{E}_q \nabla r(Z)$, while $\mathbb{E}_q\|Z\|^2 \le d$; hence $\|\mathbb{E}_q Z\| \le L\sqrt{d}$. Cauchy–Schwarz and Young's inequality therefore add at most $C\tau^2(L^2 d + \|\delta\|^2)$ to the logarithmic moment. Taking the τ-th root proves the first claim. Repeating the calculation at exponent two and using $e^u - 1 \le 2u$ in the stated smallness regime proves the χ^2 bound. Mollification and cutoff justify the calculation for $C^{1,1}$ potentials. □

Starting from y, run gradient descent on the conditional potential until it returns a measurable point $\widehat{x}_y$ satisfying

$$\|y - \eta\nabla V(\widehat{x}_y) - \widehat{x}_y\| \le \sqrt{d\eta}.$$

The accelerated branch then draws $G \sim \mathcal{N}(0, I_d)$ and uses

$$E_y = y - \eta\nabla V(\widehat{x}_y) + \sqrt{\eta}\, G. \tag{40}$$

Strong monotonicity of $x \mapsto x + \eta\nabla V(x) - y$ shows that the normalized shift of (40) from $x_*(y)$ is at most $L\sqrt{d}/(1 + \alpha\eta)$. Thus Lem. 17 applies whenever $\tau L \lesssim 1$ and $\tau L^2 d \lesssim 1$. Because the reference depends only on y and fresh Gaussian randomness, Lem. 8 gives exact conditional invariance before numerical caps are imposed.

Lemma 18 (Shifted-reference transfer)**.** *If $\tau L \lesssim 1$ and $\tau L^2 d \lesssim 1$, the conditional kernel using* (40) *satisfies Prop. 12 and Cor. 2, with possibly different universal constants.*

Proof. Let $h = \nabla r$, $\Sigma = \mathrm{Cov}_q(Z)$, and $D = I - \Sigma$. For the reference force $\mathcal{U} = h(G + \delta)$, direct expansion of the Gaussian momentum moments in the proof of Lem. 13 gives

$$\mathbb{E}k_{a,E}^2 = 6\mathrm{Var}_q(a^\top h) + 5\mathrm{Var}(a^\top \mathcal{U}) + 5\{\mathbb{E}_q(a^\top h) - \mathbb{E}(a^\top \mathcal{U})\}^2.$$

The weak integration-by-parts calculation used in Lem. 12 yields $\mathrm{Var}_q(a^\top h) \le CL\, a^\top D a$.

It remains to control the two Gaussian-reference terms. If $H = \nabla^2 r$ denotes the weak Hessian, then $\mathbb{E}_q H = D + \mathrm{Cov}_q(h) \preceq CD$. Hölder's inequality, Lem. 17, and $0 \preceq H \preceq LI$ give $\mathbb{E}_{\gamma_\delta}[a^\top H a] \le Ca^\top D a$. Indeed, $a^\top H a \le L\|a\|^2$, $\mathbb{E}_q[a^\top H a] \ge m\|a\|^2$, and the choice $\tau = 2 + \log(L/m)$ makes the resulting Hölder ratio universal. Gaussian Poincaré therefore implies $\mathrm{Var}_{\gamma_\delta}(a^\top h) \le CL\, a^\top D a$. The χ^2 part of Lem. 17, followed by Cauchy–Schwarz under q, bounds the squared difference of means by the same quantity. Hence the local quarter-turn estimate holds with the shifted reference. The affine lifting in Prop. 12, momentum parity, and the factorization proof are unchanged. □

C.4.3 Mode, Clock, and Refresh Subroutines

Lemma 19 (Mode certificate and deterministic cap)**.** *Assume $\beta\eta \le 1$. Gradient descent with step $(\beta + \eta^{-1})^{-1}$, initialized at y, finds the certificate used in* (40) *using at most*

$$C\left[1 + \log_+\left(\frac{\sqrt{\eta}\|\nabla V(y)\|}{\sqrt{d}}\right)\right]$$

oracle queries. Under the stationary augmented law, stopping after

$$J_B = C\left[1 + \log(2+\kappa) + \log\left(1 + \sqrt{B/d}\right)\right]$$

iterations fails with probability at most $2e^{-B}$. The uncapped output is a Borel function of y.

Proof. The conditional potential has Hessian between $\alpha + \eta^{-1}$ and $\beta + \eta^{-1}$, hence condition number at most two. One gradient step therefore reduces the gradient norm by a universal factor, which gives the pointwise logarithmic bound.

For the stationary tail, let $x^\star$ be the global minimizer of V and write $Y = X + \sqrt{\eta}Z$, where $Z \sim \mathcal{N}(0, I_d)$. Strong log-concavity gives $\mathbb{E}\|X - x^\star\|^2 \le d/\alpha$ and Gaussian concentration around that scale. Combining it with the Gaussian norm tail and $\|\nabla V(Y)\| \le \beta\|Y - x^\star\|$ gives, outside an event of probability $2e^{-B}$,

$$\frac{\sqrt{\eta}\|\nabla V(Y)\|}{\sqrt{d}} \le 2\sqrt{\kappa}\left(1 + \sqrt{2B/d}\right).$$

Substitution into the pointwise bound gives J_B. Every finite gradient-descent iterate and the first iterate meeting the certificate are Borel functions of y. Returning the complete pre-reflection state on cap failure therefore defines a total Markov kernel without violating conditional independence. □

Lemma 20 (Clock cap)**.** *Suppose the normalized residual potential r has L-Lipschitz gradient with $0 \le L \le 1$, and the normalized reference is $G + \delta_y$, where G is standard Gaussian, $y \mapsto \delta_y$ is measurable, and $\sup_y\|\delta_y\| \le L\sqrt{d}$. For $B \ge 1$, set $\bar{\lambda}_B = CL(\sqrt{dB} + B)$. Replacing the event rate by its minimum with $\bar{\lambda}_B$ changes one stationary transition with probability at most $C(1 + L\sqrt{d})e^{-cB}$. The capped flight is simulated by homogeneous Poisson thinning and uses at most $1 + (\pi/2)\bar{\lambda}_B$ expected oracle queries, including the query at the reference.*

Proof. Conditionally on $Y = y$, the normalized position is one-strongly log-concave, while G and the momentum are Gaussian. Their norm tails, the bound $\|\nabla r(w) - \nabla r(e)\| \le L\|w - e\|$, and a one-dimensional Gaussian tail give, uniformly in y,

$$\Pr\{\lambda_t > CL(\sqrt{du} + u)\} \le 6e^{-u}.$$

Invariance from Lem. 8 makes the bound valid at every time. Integrating the tail over the quarter-period bounds the expected excess hazard by $C(1 + L\sqrt{d})e^{-cB}$. Using the same Poisson random measure for the exact and capped clocks couples the paths until the first excess event. Conditioning on $Y = y$ and then integrating proves the claim for the measurable random center. Ordinary thinning gives the query bound. □

Proposition 14 (Approximate RGO)**.** *There is a universal K_{RGO} such that the following holds. Given y, $0 < \delta < 1/2$, and a point $\widehat{x}_y$ satisfying*

$$\|y - \eta\nabla V(\widehat{x}_y) - \widehat{x}_y\| \le \sqrt{d\eta},$$

if

$$\eta^{-1} \ge K_{\mathrm{RGO}}\beta\{\sqrt{d\log(1/\delta)} + \log(1/\delta)\},$$

the first-order rejection sampler of Chen et al. (2026a, Theorems 3.1 and D.1) returns $\widehat{q}_y$ with $\chi^2(q_y\|\widehat{q}_y) \le \delta^2$ and $\chi^2(\widehat{q}_y\|q_y) \le \delta^2$, using at most K_{RGO} expected oracle queries after construction of the mode certificate.

C.4.4 Proof of the Implementation Theorem

Proof of Thm. 6. Let $A_\varepsilon = \Delta + \log(8/\varepsilon)$ and take

$$\delta_{\mathrm{RGO}} = \frac{\varepsilon}{8(1+N\rho)}, \qquad B = C\left\{\Delta + \log\frac{8N}{\varepsilon} + \log(2+d+\kappa)\right\}.$$

The exact comparison chain uses the measurable shifted reference in (40), without mode or clock caps, and an exact refresh. Lem. 18 lets us apply Prop. 13; the choice of N makes its distance from stationarity at most $\varepsilon/2$.

The implementable algorithm first reflects y. It then runs the y-initialized mode solver for at most J_B steps. On failure it returns the complete pre-reflection state (x, y, p). On the refresh branch it invokes Prop. 14 with tolerance δ_{RGO}; otherwise it uses the shifted Gaussian reference and the clock cap in Lem. 20.

Couple both chains until their first disagreement. The complete error ledger is

source	total TV contribution
ideal mixing	$\varepsilon/2$
mode-solver cap	$\varepsilon/8$
event-clock cap	$\varepsilon/8$
approximate refresh	$N\rho\delta_{\mathrm{RGO}}/2 \le \varepsilon/16$.

For the mode and clock errors, Lem. 16 converts stationary exponential tails into path probabilities under the warm exact chain. The RGO error is uniform in y, so maximal coupling and kernel telescoping apply directly. The entries sum to less than ε, and projection to X cannot increase total variation.

We next check the local conditions. With $\tau = 2 + \log\kappa$ and $L = \beta\eta = c/(\sqrt{d\Lambda_\varepsilon} + \Lambda_\varepsilon)$, reducing c gives $\tau L \lesssim 1$ and $\tau L^2 d \lesssim 1$. Also $\log(1/\delta_{\mathrm{RGO}}) \lesssim \Lambda_\varepsilon$, so the same step size verifies the gate in Prop. 14.

The expected number of refresh calls is $N\rho \le C\ell_g A_\varepsilon$. Each iteration uses at most

$$C\{J_B + 1 + L(\sqrt{dB} + B) + \rho K_{\mathrm{RGO}}\}$$

expected oracle queries. Here J_B and ℓ_g are logarithmic, and the definition of η makes $L(\sqrt{dB} + B) = O(1)$. Substituting $g^2 = \alpha\eta$ proves the theorem and hence Thm. 4. □

C.4.5 Affine Form Used by the Quadratic Baseline

Lemma 21 (Affine local compiler)**.** *Suppose a conditional potential is*

$$\Phi_y(x) = R_y(x) + \frac{\omega}{2}\|x - y\|^2, \qquad m\omega I \preceq \nabla^2 R_y \preceq L\omega I, \qquad 0 < m \le L \le 1.$$

Let $x_*(y)$ denote the unique minimizer of Φ_y, let $G \sim \mathcal{N}(0, I_d)$, and fix $0 < \delta < 1/2$. Let $\tau = 2 + \log(L/m)$, and assume $\tau L \lesssim 1$ and $\tau L^2 d \lesssim 1$. If a measurable point $\widehat{x}_y$ satisfies

$$\|\omega(\widehat{x}_y - y) + \nabla R_y(\widehat{x}_y)\| \le \sqrt{d\omega},$$

then the reference

$$E_y = y - \omega^{-1}\nabla R_y(\widehat{x}_y) + \omega^{-1/2}G$$

satisfies the shifted-Gaussian comparison, the local quarter-turn forcing estimate, and the clock estimate with local smoothness L. The certificate is found from y using

$$C\left[1 + \log_+ \frac{\|\nabla R_y(y)\|}{\sqrt{d\omega}}\right]$$

oracle queries. The RGO gate becomes

$$1 \ge K_{\mathrm{RGO}} L\{\sqrt{d\log(1/\delta)} + \log(1/\delta)\}.$$

Proof. In the coordinate $z = \sqrt{\omega}(x - x_*(y))$, the residual has Hessian between mI and LI. Strong monotonicity of $x \mapsto \omega(x-y) + \nabla R_y(x)$ bounds the normalized Gaussian shift by $L\sqrt{d}/(1+m)$. Thus Lems. 17 and 20 apply after replacing $\beta\eta$ by L. The proof of Lem. 18 up to its local quarter-turn estimate applies in these normalized coordinates; lifting that estimate to the global reflected chain requires the separate scaling calculation in App. C.5. The condition number of Φ_y is at most $(1+L)/(1+m) \le 2$, giving the displayed logarithmic mode cost. Finally, applying Prop. 14 after the affine change of variables gives exactly the stated RGO gate. □

C.5 Quadratic-Baseline Refinement

Corollary 3 (Quadratic-baseline refinement)**.** *For $0 < \varepsilon \le 1/2$, if α is known and the initial position law ν_0 satisfies $D_2(\nu_0\|\pi) = O(1)$, a quadratic-baseline implementable algorithm under the same augmented initialization attains TV error ε with $\widetilde{O}(1 + \sqrt{\kappa - 1}\, d^{1/4})$ expected oracle queries and* $\mathrm{polylog}(d, \kappa, 1/\varepsilon)$ *expected RGO calls.*

We prove Cor. 3 by changing only the quadratic part integrated exactly by the conditional motion. The outer reflection and its gap are unchanged.

The parameter choice below is easiest to read from its design goals. Write $s = \kappa - 1$ for the curvature above the known α-quadratic baseline, and let θ be the desired normalized smoothness of the residual force. We choose $\lambda = 1 - \delta_0$ and η so that

(i) subtracting the $\lambda\alpha$-quadratic leaves residual strong convexity $\delta_0\alpha > 0$;

(ii) the residual smoothness divided by the new harmonic curvature is exactly θ;

(iii) the original small-step condition remains valid, so the outer reflection gap is still governed by $g = \sqrt{\alpha\eta}$.

The displayed formulas are chosen to meet these requirements; C_0 provides fixed slack in the first one.

Let $s = \kappa - 1$ and, for a sufficiently small universal c, set

$$\Lambda_\varepsilon = \Delta + \log\frac{Cd\kappa}{\varepsilon}, \qquad \theta = \frac{c}{\sqrt{d\Lambda_\varepsilon} + \Lambda_\varepsilon}.$$

Choose a sufficiently large universal C_0, put $\delta_0 = C_0\theta$ and $\lambda = 1 - \delta_0$, and select η by

$$\alpha\eta = \frac{\theta}{s + (C_0 - 1)\theta + C_0\theta^2}. \tag{41}$$

Choosing C_0 first and then reducing c ensures

$$0 < \lambda < 1, \qquad \beta\eta \le c_0, \qquad \theta \le c_0,$$

where c_0 is the universal small-step constant in Thm. 3. To see the middle inequality, split into $s \ge \theta$ and $s < \theta$ in $\beta\eta = (s+1)\theta/\{s + (C_0-1)\theta + C_0\theta^2\}$.

For the implementable algorithm, put

$$g = \sqrt{\alpha\eta}, \quad \ell_g = 1 + \log(1/g), \quad \rho \asymp g\sqrt{\ell_g}, \quad A_\varepsilon = \Delta + \log(8/\varepsilon),$$

and choose

$$N = \left\lceil \frac{C\sqrt{\ell_g}}{g} A_\varepsilon \right\rceil, \quad \delta_{\mathrm{RGO}} = \frac{\varepsilon}{8(1+N\rho)}, \quad B = C\left\{\Delta + \log\frac{8N}{\varepsilon} + \log(2 + d + \kappa)\right\}.$$

These are the same mixing horizon and error budgets used in Thm. 6, evaluated at the baseline step size in (41).

For fixed y, split the conditional potential as

$$V(x) + \frac{\|x-y\|^2}{2\eta} = R_y(x) + \frac{\omega_\eta}{2}\|x-y\|^2, \quad R_y(x) = V(x) - \frac{\lambda\alpha}{2}\|x-y\|^2, \quad \omega_\eta = \lambda\alpha + \eta^{-1}.$$

The residual curvature satisfies

$$\delta_0 \alpha I \preceq \nabla^2 R_y \preceq (\beta - \lambda\alpha) I, \qquad \frac{\beta - \lambda\alpha}{\omega_\eta} = \theta, \qquad \omega_\eta \eta \asymp 1.$$

Thus the harmonic motion integrates both the proximal quadratic and the known $\lambda\alpha$-quadratic baseline; collisions see only the remaining curvature.

Local validity and implementation. In the coordinate $z = \sqrt{\omega_\eta}(x - y)$, use the same fixed-reference process as in (23), with

$$b_y^{\mathrm{res}}(z) = \omega_\eta^{-1/2} \nabla R_y(y + \omega_\eta^{-1/2} z)$$

in place of b_y. Lem. 8 applies after this affine change of variables, so the process still preserves the original conditional law. Its normalized residual curvatures are

$$m_\star = \frac{\delta_0 \alpha}{\omega_\eta}, \qquad L_\star = \theta.$$

If $\tau_\star = 2 + \log(L_\star / m_\star)$, then

$$\frac{L_\star}{m_\star} = \frac{\omega_\eta}{C_0 \alpha} = 1 + \frac{s}{C_0 \theta}, \qquad \tau_\star \lesssim \Lambda_\varepsilon.$$

Consequently $\tau_\star L_\star \lesssim 1$ and $\tau_\star L_\star^2 d \lesssim 1$, after reducing c. The definition above also gives $\log(1/\delta_{\mathrm{RGO}}) \le C\Lambda_\varepsilon$. Hence

$$K_{\mathrm{RGO}} \theta \{ \sqrt{d \log(1/\delta_{\mathrm{RGO}})} + \log(1/\delta_{\mathrm{RGO}}) \} \le 1$$

after another universal reduction of c. This is exactly the affine RGO gate in Lem. 21.

The reference and its certificate are

$$\| \omega_\eta (\widehat{x}_y - y) + \nabla R_y(\widehat{x}_y) \| \le \sqrt{d \omega_\eta}, \qquad E_y = y - \omega_\eta^{-1} \nabla R_y(\widehat{x}_y) + \omega_\eta^{-1/2} G.$$

In the remainder of this subsection, S denotes the rescaled conditional quarter-turn averaged over this shifted-Gaussian reference. Lem. 21 therefore supplies the shifted Gaussian comparison, event-clock bound, mode solver, and RGO gate with local smoothness $L_\star = \theta$. The stationary mode cap remains logarithmic: since $\nabla R_Y(Y) = \nabla V(Y)$ and $\omega_\eta \ge \eta^{-1}$,

$$\frac{\eta}{d} \mathbb{E} \|\nabla V(Y)\|^2 \le \kappa(\beta\eta) + (\beta\eta)^2 \le \kappa + 1.$$

The concentration argument in Lem. 19 then applies without a polynomial query factor.

Transfer estimate after rescaling. It remains to verify that the ideal contraction theorem sees the same position–momentum geometry. Let

$$\Sigma_y = \eta^{-1} \mathrm{Cov}(X \mid Y = y), \qquad D_y = I - \Sigma_y, \qquad D_z(y) = I - \omega_\eta \eta \, \Sigma_y.$$

Thus Σ_y, D_y are exactly the conditional covariance quantities defined in App. C.2, while D_z is their rescaled counterpart. Because $\omega_\eta \eta \ge 1$,

$$D_z(y) = \omega_\eta \eta D_y - (\omega_\eta \eta - 1) I \preceq \omega_\eta \eta D_y.$$

The correctly scaled momentum observable is

$$M_f^{(\alpha)} = \frac{2}{\sqrt{\omega_\eta}} a_f(Y)^\top P.$$

Applying the local quarter-turn estimate in the z-coordinate yields

$$\| \mathsf{S} L_f - M_f^{(\alpha)} \|^2 + \| \mathsf{S} M_f^{(\alpha)} + L_f \|^2 \le C\theta\eta \mathcal{D}_f.$$

Its defect under the outer reflection is at most $C\eta\mathcal{R}_f$, and $\|M_f^{(\alpha)}\|^2 \le C\eta \|\nabla f\|^2$. Combining these bounds with (32) and Lem. 11 proves the same position–momentum transfer estimate (29). Momentum parity is unchanged, so Cor. 2 and Lem. 14 apply with the same universal constants. Notice that the outer gap is still $g = \sqrt{\alpha\eta}$; only the local collision scale changes from $\beta\eta$ to θ.

Complexity. From (41),

$$\frac{1}{\alpha\eta} = \frac{s}{\theta} + C_0 - 1 + C_0\theta,$$

and hence

$$\frac{1}{\sqrt{\alpha\eta}} \le C\left(1 + \frac{\sqrt{\kappa-1}}{\sqrt{\theta}}\right) = \widetilde{O}\Big(1 + \sqrt{\kappa-1}\, d^{1/4}\Big).$$

Also $\theta\sqrt{d} = O(1)$, so the clock uses only $\widetilde{O}(1)$ expected oracle queries per iteration. The ideal contraction bound and our choice of N give TV error at most $\varepsilon/2$. With the explicit B above, Lems. 16, 19 and 20 bound the two cap errors by $\varepsilon/8$ each. Uniform maximal coupling for the approximate refresh contributes at most $N\rho\delta_{\mathrm{RGO}}/2 \le \varepsilon/16$. Thus the complete error is below ε. The per-iteration mode and clock costs are polylogarithmic, while $N\rho$ is polylogarithmic, so the same accounting gives

$$\mathrm{TV}(\mathrm{Law}(X_N), \pi) \le \varepsilon, \qquad \mathbb{E}N_{\mathrm{qry}} = \widetilde{O}\Big(1 + \sqrt{\kappa-1}\, d^{1/4}\Big), \qquad \mathbb{E}N_{\mathrm{RGO}} = \mathrm{polylog}(d, \kappa, 1/\varepsilon)$$

when $\Delta = O(1)$. This proves Cor. 3.

D Proof of the Warm-Start Wrapper

This appendix proves both parts of Thm. 5. For the finite-divergence part, we first state two implementable samplers for the localized conditionals and a capped interface for the final reflected sampler. We then show that exact coarse proximal steps reduce a finite D_2 budget to a constant, couple those steps to an algorithmic chain at adaptive observation centers, optimize the coarse scale, and transfer the final guarantee through the coupling. The last subsection proves the supplied-point part by constructing the required finite D_2 budget from one point.

The proof uses two external posterior-sampling results, stated precisely in Prop. 15, and the standard exact-proximal hypercontractive inequality stated in Lem. 23. All remaining steps—including the random observation centers, the distinction between comparison and algorithmic laws, the final error transfer, and the parameter optimization—are proved below.

At a high level, the wrapper compares two routes. The direct route applies a reflected sampler immediately to the initial law. The coarse route first constructs two endpoints in parallel for the analysis: an exact endpoint $\nu_m^\dagger$, which is proved to be constant-D_2-warm, and an algorithmic endpoint $\widehat{\nu}_m$, which is close to $\nu_m^\dagger$ in total variation. The same capped reflected kernel is then applied to both endpoints. Only $\widehat{\nu}_m$ is generated by the algorithm; $\nu_m^\dagger$ is a comparison law used to prove convergence.

D.1 Posterior Samplers Used by the Wrapper

At a coarse proximal scale $u > 0$, an observation y defines the conditional density

$$q_y(x) \propto \exp\left\{-V(x) - \frac{\|x-y\|^2}{2u}\right\}.$$

Its strong-convexity and smoothness parameters are $m_u = \alpha + u^{-1}$ and $L_u = \beta + u^{-1}$, so its condition number is $K_u = L_u/m_u$.

We use two ways to sample this conditional. The *regular interface* first constructs a constant-warm law and then applies one of our reflected samplers. The *harmonic-base interface* instead starts from a localized Gaussian and invokes exact-ULD/FORS. We abbreviate the latter by "HB" because its reference is the quadratic, or harmonic, part of the conditional potential; HB does not mean a heat-bath update. The regular interface is advantageous in the intermediate- and large-dimensional regimes, whereas the HB interface supplies the small-dimensional endpoint.

We will repeatedly turn expected-work routines into total kernels with a deterministic work limit. The following elementary observation records why this does not introduce a polynomial dependence on the requested accuracy.

Lemma 22 (Logarithmic work truncation). *Suppose a routine has a deterministic number of stages and every random-length operation is one of the following: a geometric trial count with uniformly positive success probability; a Poisson count with known mean Λ; or a linearly convergent optimization for which*

$$N(z) \le C\{1 + \log(1 + H(z)/a)\}, \qquad \mathbb{E}H(Z)^2 \le M^2$$

under the comparison law relevant to that call. For every $\xi > 0$, all such operations can be stopped at deterministic limits so that the chance of any stop is at most ξ. The limits add only logarithmic factors to the original work, apart from the already charged Poisson mean Λ. Returning a fixed fallback output after a stop changes the output law by at most ξ in total variation.

Proof. Geometric tails require $O(\log(1/\xi))$ trials. A Poisson variable of mean Λ is at most $\Lambda + C\{\sqrt{\Lambda \log(1/\xi)} + \log(1/\xi)\}$ except with probability ξ. Finally, Markov's inequality gives $H(Z) \le M/\sqrt{\xi}$ outside an event of probability ξ; because N depends logarithmically on H, this increases its limit by only $O(\log(1/\xi))$. Divide ξ by the deterministic number of operations and use a union bound. Coupling the capped and uncapped routines until the first stop proves the total-variation statement. □

Proposition 15 (Finite posterior interfaces). *Fix a law ν with $D_2(\nu\|\pi) \le B$, draw $X \sim \nu$ and, independently, a standard Gaussian vector $G \sim \mathcal{N}(0, I_d)$, where I_d is the d-dimensional identity; set $Y = X + \sqrt{u}G$, and denote the law of Y by λ. Suppose $u \le 1/\alpha$. For every $0 < \delta < 1/2$, each of the two interfaces above gives a jointly measurable, work-capped kernel $y \mapsto \widehat{P}_y$ such that*

$$\int \mathrm{TV}(\widehat{P}_y, q_y)\, \lambda(\,\mathrm{d}y) \le \delta. \tag{42}$$

For the regular interface, every call uses at most

$$C_{\mathrm{reg}}(K_u, d) := K_u^{7/6} d^{1/6} + K_u^{1/2} d^{1/4} \tag{43}$$

oracle queries up to factors logarithmic in $d, K_u, B, 1/\delta$. For the HB interface, the corresponding bound is

$$C_{\mathrm{HB}}(K_u, d) := 1 + K_u^{2/3} d^{1/3} \tag{44}$$

oracle queries, again up to the same type of logarithmic factors.

Proof. **Regular interface.** After rescaling a potential whose condition number is K, Chen et al. (2026c, Theorem 7.1(ii)) constructs the output law of an implementable algorithm and an analysis-only comparison law that are arbitrarily close in total variation, with the latter having constant order-two Rényi divergence from the target. At constant target divergence its query cost is $\widetilde{O}(K^{7/6}d^{1/6} + K)$; the second term is absorbed by the first for $K, d \ge 1$. Applying either reflected finalizer from this paper to the comparison law and the same capped Markov kernel to the algorithmic output law costs another $\widetilde{O}(K^{1/2}d^{1/4})$. Total- variation data processing and the triangle inequality give (43).

HB interface. First find a point z_y satisfying the observable test $\|\nabla\Phi_y(z_y)\| \le m_u/\sqrt{L_u}$, where $\Phi_y(x) = V(x) + \|x - y\|^2/(2u)$. Strong convexity then gives $\sqrt{L_u}\|z_y - x_y^*\| \le 1$, where x_y^* minimizes Φ_y. The Gaussian $G_y = \mathcal{N}(z_y, (2L_u)^{-1} I_d)$ then obeys $D_q(G_y\|q_y) = O_q(d\log(\mathrm{e}K_u))$ for every fixed $q > 1$, by the usual quadratic sandwich and completion of squares. Applying Chen et al. (2026d, Theorem 3.2(ii)) with output order two, initialize phase space from G_y tensored with stationary Gaussian momentum. The source theorem's initial budget is then $d + D_4(G_y\|q_y) = \widetilde{O}(d)$. Set its Rényi tolerance to $O(\delta^2)$, project the output to position, and use $\mathrm{TV}^2 \le D_2/2$. This gives (44). Its high-probability query bound also includes rejected trials.

Random centers and work caps. It remains to justify localization at the random center y. The conditional objective has gradient $\nabla V(y)$ at y. For the regular interface, the observable test $\|\nabla\Phi_y(z)\| \le \sqrt{2m_u d}$ implies $\Phi_y(z) - \min\Phi_y \le d$, as required by the imported warmifier; for the HB interface we use the test in the preceding paragraph. Accelerated optimization started at y reaches either certificate in $\widetilde{O}(\sqrt{K_u})$ queries, with only logarithmic dependence on $1 + \|y - x_*\|$. This cost is absorbed by both displayed bounds. Under the law λ, the tail needed to cap this search uniformly follows from Lem. 24 below. On cap failure the kernel returns a fixed point; the cap probability is

included in (42). The cap depends only on the known parameters $\alpha, \beta, d, B, \delta$, not on the unknown law λ.

We next verify the work cap rather than inferring it from an expected-query statement. The recursive warmifier has a deterministic number of stages. Its proof controls the second moments of every center passed to its gradient-descent subroutines, while its pointwise bound for each such solve is logarithmic in the center gradient; see Chen et al. (2026c, Theorems 6.1(ii) and 6.5, their proofs, and Lemmas D.3–D.4). Thus these solves satisfy the third case of Lem. 22. Its remaining rejection loops have geometric tails once their observable mode certificate holds; this is stated for the RGO component in Chen et al. (2026a, Theorems 3.1 and D.1). The appended reflected finalizer in the regular interface has only geometric rejection counts and Poisson event counts. The exact-ULD/FORS interface already gives a high-probability query bound. Applying Lem. 22 makes the displayed work bounds deterministic up to their logarithmic factors and charges every stop to (42). Finally, finite optimization iterates, Gaussian draws, acceptance tests, and capped stopping rules are Borel functions of y and the random seed. Thus $\widehat{P}_y$ is a genuine Markov kernel rather than a pointwise collection of samplers. □

We next isolate the final-stage property needed by the wrapper. Stating it as a capped Markov-kernel interface is important: after the algorithmic and comparison paths separate, the final sampler must still have a deterministic work limit.

Proposition 16 (Capped reflected finalizer)**.** *Fix $B \geq 1$ and $0 < \varepsilon, \xi < 1/2$. Either reflected sampler admits an implementable realization as a total Markov kernel $\widehat{K}_{B,\varepsilon,\xi}$ such that, for every ν satisfying $D_2(\nu\|\pi) \leq B$, $\mathrm{TV}(\nu\widehat{K}_{B,\varepsilon,\xi}, \pi) \leq \varepsilon + \xi$. Up to logarithmic factors, the kernel uses at most*

$$\widetilde{O}\Big(\sqrt{\kappa}B(\sqrt{dB} + B)^{1/2}\Big) \tag{45}$$

oracle queries. In particular, constant B costs $\widetilde{O}(\sqrt{\kappa}d^{1/4})$. The same kernel may be applied to a second, arbitrary input law: its deterministic work bound remains valid, and total-variation distance between the two inputs cannot increase.

Proof. The uncapped guarantees are Cor. 1 and (8). The implementable algorithms have a fixed outer horizon. Their mode solves have logarithmic pointwise work bounds and warm-path second-moment control, their rejection counts are geometric, and their event proposal counts are Poisson with the means already charged in their query ledgers. Hence Lem. 22 applies. For any additional error budget ξ, each implementation can therefore be replaced by a total Markov kernel with a deterministic work bound of the same polynomial order and with at most ξ additional total-variation error from every D_2-warm input. For the flux method the required estimates are in Prop. 10 and Lem. 7; for the persistent method it follows from Lems. 16, 19 and 20 together with the same elementary Poisson and geometric tail bounds. We use this capped version below so that the query ledger remains valid even if the algorithmic and comparison paths have already separated. The final assertion is the data-processing inequality for a common Markov kernel. □

D.2 Exact Coarse Proximal Contraction

Let S_u denote one exact proximal Gibbs update: observe $Y = X + \sqrt{u}G$, then draw $X' \sim q_Y$. The kernel is π-reversible. The following standard hypercontractive form makes its warm-start effect explicit.

Lemma 23 (Exact proximal warmification)**.** *For every density $f = \mathrm{d}\nu/\mathrm{d}\pi$, every $p > 1$, and every $u > 0$,*

$$\|S_u f\|_{1+(1+\alpha u)^2(p-1)} \leq \|f\|_p.$$

Consequently, if $G_m = \prod_{j=1}^m (1 + \alpha u_j)$, then

$$D_2(\nu S_{u_1} \cdots S_{u_m}\|\pi) \leq \frac{2D_2(\nu\|\pi)}{1 + G_m^2}. \tag{46}$$

Proof. Let π_u be the observation marginal and define $A_u^* f(y) = \mathbb{E}[f(X) \mid Y = y]$, which maps density ratios on π to density ratios on π_u. Its adjoint is $A_u g(x) = \mathbb{E}[g(Y) \mid X = x]$, and one Gibbs

row has density operator $S_u = A_u A_u^*$. Put $c_u = 1 + \alpha u$. The two-measure heat-flow inequality of Chen et al. (2022); Kook & Vempala (2025) states that $A_u^* : L^{p_0}(\pi) \to L^{p_1}(\pi_u)$ is a contraction whenever $p_1 - 1 = c_u(p_0 - 1)$. Duality gives the analogous gain for A_u: $A_u : L^{p_1}(\pi_u) \to L^{p_2}(\pi)$ is a contraction when $p_2 - 1 = c_u(p_1 - 1)$. To see the direction of this second relation, the conjugate-exponent form of the first one is $p_0' - 1 = c_u(p_1' - 1)$, where a prime denotes the Hölder conjugate.

Thus the exponent flow through one Gibbs row is

$$L^{p_0}(\pi) \xrightarrow{A_u^*} L^{p_1}(\pi_u) \xrightarrow{A_u} L^{p_2}(\pi), \qquad p_2 - 1 = c_u^2(p_0 - 1).$$

The two excess-exponent gains therefore multiply and prove the first claim. After iteration it gives $\|S_{u_1} \cdots S_{u_m} f\|_2 \le \|f\|_{1+G_m^{-2}}$.

For $p = 1 + G_m^{-2}$, interpolation between $L^1(\pi)$ and $L^2(\pi)$, using $\|f\|_1 = 1$, gives $\log\|f\|_p \le (p-1)D_2(\nu\|\pi)/p$. Since the output order-two divergence is twice the logarithm of its L^2 density norm, (46) follows. □

Set $r = \beta u$ and restrict $1 \le r \le \kappa$. Then $\alpha u = r/\kappa$, and $\log(1 + r/\kappa) \asymp r/\kappa$. Therefore (46) makes the exact endpoint constant-D_2 warm after

$$m = O\Big(\frac{\kappa}{r}\log(\mathrm{e}B)\Big) \tag{47}$$

equal-scale updates. Notice also that their conditional condition number is

$$K_u = \frac{\beta + u^{-1}}{\alpha + u^{-1}} = \frac{\kappa(1+r)}{\kappa + r} \le 1 + r \le 2r. \tag{48}$$

D.3 Random Centers and the Algorithmic Path

The posterior centers in the preceding chain are adaptive: the next center is sampled from the current position. We now verify that the conditional interfaces can be applied without assuming that these centers are fixed or uniformly bounded.

Lemma 24 (Random-center and path accounting)**.** *Let ν_0 satisfy $D_2(\nu_0\|\pi) \le B$, and run m exact proximal updates with scales at most $1/\alpha$. Then every position law ν_j satisfies*

$$\mathbb{E}_{\nu_j}\|X - x_*\|^2 \le \frac{C(d+B)}{\alpha}.$$

Moreover, all observation-dependent mode searches along the path can be capped so that the probability of any cap failure is at most δ, while their query cost changes only by a factor logarithmic in $d, \kappa, B, m, 1/\delta$.

Proof. Every exact proximal kernel preserves π, so data processing gives $D_2(\nu_j\|\pi) \le B$. Hence $D_{\mathrm{KL}}(\nu_j\|\pi) \le B$, and the T_2 inequality for an α-strongly log-concave law gives $W_2^2(\nu_j, \pi) \le 2B/\alpha$. Combining an optimal coupling with $\mathbb{E}_\pi\|X - x_*\|^2 \le d/\alpha$ proves the moment bound. Since $Y = X + \sqrt{u}G$ and $u \le 1/\alpha$, the same order of control holds for the observation.

For the tail statement, let $\mathbb{P}_{\nu_0}^{(m)}$ be the law of the entire exact path, including its observations, and let $\mathbb{P}_\pi^{(m)}$ be the law of the same path started at stationarity. Because their transition kernels agree,

$$\frac{\mathrm{d}\mathbb{P}_{\nu_0}^{(m)}}{\mathrm{d}\mathbb{P}_\pi^{(m)}} = \frac{\mathrm{d}\nu_0}{\mathrm{d}\pi}(X_0), \qquad \mathbb{P}_{\nu_0}^{(m)}(E) \le \mathrm{e}^{B/2}\mathbb{P}_\pi^{(m)}(E)^{1/2}$$

for every path event E, where the second inequality is Cauchy–Schwarz. Under the stationary path, strong log-concavity and the independent Gaussian observation give exponential tails for $\|Y - x_*\|$. Because the conditional mode solver converges geometrically and its stopping test is observable, a threshold with stationary failure probability at most $\mathrm{e}^{-B}\delta^2/(4m^2)$ increases its iteration limit only logarithmically. A union bound followed by the displayed path inequality makes the probability of any cap failure at most δ. This bound is for the whole adaptive path, so no independence between successive centers is used, and its threshold is determined entirely by the displayed problem parameters. □

Lemma 25 (Finite approximation of the coarse path)**.** *Let $\nu_m^\dagger = \nu_0 S_u^m$ be the exact endpoint with m as in* (47)*. Using either posterior interface in Prop. 15, there is an algorithmic endpoint $\widehat{\nu}_m$ such that $\mathrm{TV}(\widehat{\nu}_m, \nu_m^\dagger) \le \delta$. Its number of oracle queries is deterministically bounded, up to logarithmic factors, by $mC_{\mathrm{reg}}(K_u, d)$ or $mC_{\mathrm{HB}}(K_u, d)$, respectively.*

Proof. The row-wise coupling has the following form as long as the two paths have not yet disagreed:

$$X_{j-1}^\dagger \xrightarrow{Y_j^\dagger = X_{j-1}^\dagger + \sqrt{u}G_j} Y_j^\dagger \xrightarrow{q_{Y_j^\dagger}} X_j^\dagger,$$
$$\widehat{X}_{j-1} \xrightarrow{\widehat{Y}_j = \widehat{X}_{j-1} + \sqrt{u}G_j} \widehat{Y}_j \xrightarrow{\widehat{P}_{j,\widehat{Y}_j}} \widehat{X}_j.$$

The shared Gaussian G_j makes $Y_j^\dagger = \widehat{Y}_j$ whenever $X_{j-1}^\dagger = \widehat{X}_{j-1}$. At that common observation, we use a maximal coupling of the two conditional draws.

At row j, apply Prop. 15 with total error budget δ/m; this budget includes the posterior approximation and all of its mode and work caps. Couple the exact and approximate conditional draws maximally whenever the two paths have agreed so far. To avoid conditioning on the agreement event, first generate the entire exact path and let $e_j(y) = \mathrm{TV}(\widehat{P}_{j,y}, q_y)$. The probability that the first mismatch occurs in row j is at most $\mathbb{E}[\mathbf{1}_{\{\text{no earlier mismatch}\}} e_j(Y_j)]$, which is bounded by the unconditional exact-path average $\mathbb{E}e_j(Y_j) \le \delta/m$. Summing over the rows gives total mismatch at most δ. The simultaneous choice of all random-center caps is justified by Lem. 24.

This is ordinary total-variation telescoping; a Hellinger chain rule is not needed because the requested accuracy affects both posterior interfaces only logarithmically. Crucially, every call is work-capped. Thus the query bound remains valid after the first coupling disagreement, even though the later algorithmic centers no longer have the exact center law. The exact endpoint $\nu_m^\dagger$, rather than the algorithmic endpoint, is the law to which (46) is applied. □

D.4 Cost Optimization and Final Error Transfer

Proof of the finite-divergence part of Thm. 5. The optimization selects one of the following three coarse-stage implementations. The table states the resulting choices; the algebra below derives them.

Regime	Interface	r	Coarse-stage cost
$d \le \kappa^{16/7}$	HB	$r = \kappa$	$\widetilde{O}(\kappa^{2/3} d^{1/3} \log(\mathrm{e}B))$
$\kappa^{16/7} < d < \kappa^8$	Regular	$r = d^{1/8}$	$\widetilde{O}(\kappa d^{3/16} \log(\mathrm{e}B))$
$d \ge \kappa^8$	Regular	$r = \kappa$	$\widetilde{O}(\sqrt{\kappa} d^{1/4} \log(\mathrm{e}B))$

Fix $1 \le r \le \kappa$, set $u = r/\beta$, and use the regular posterior interface. Multiplying (43) by the row count in (47), and using (48), gives the coarse-stage ledger

$$Q_{\mathrm{reg}}(r) = \widetilde{O}\Big(\kappa\{r^{-1} + r^{1/6} d^{1/6} + r^{-1/2} d^{1/4}\} \log(\mathrm{e}B)\Big).$$

The r^{-1} term records the constant work of each row.

If $d \le \kappa^8$, the admissible choice $r = d^{1/8}$ balances the last two terms and gives $\widetilde{O}(\kappa d^{3/16} \log(\mathrm{e}B))$. If $d \ge \kappa^8$, choose $r = \kappa$. The two nonconstant terms are then $\kappa^{7/6} d^{1/6}$ and $\sqrt{\kappa} d^{1/4}$; the first is no larger than the second precisely when $d \ge \kappa^8$.

For the remaining endpoint, again choose $r = \kappa$, but use the HB interface in (44). There are only $O(\log(\mathrm{e}B))$ rows and $K_u = (\kappa + 1)/2$, so this route costs $\widetilde{O}(\kappa^{2/3} d^{1/3} \log(\mathrm{e}B))$. The harmonic and regular ledgers cross at $d = \kappa^{16/7}$, while the two regular choices cross at $d = \kappa^8$. Selecting the cheapest route gives exactly $\widetilde{O}(\Psi(\kappa, d) \log(\mathrm{e}B))$.

We now close the distributional argument. Choose the constants in (47) so that $D_2(\nu_m^\dagger \| \pi) \le 1$, and use Lem. 25 with $\delta = \varepsilon/4$. Lift both $\nu_m^\dagger$ and $\widehat{\nu}_m$ through the same Gaussian augmentation required by the selected finalizer, and append the same independent Gaussian momentum when needed. The lifted comparison law has the same density ratio relative to stationarity, while total-variation data processing preserves the $\varepsilon/4$ discrepancy.

Invoke Prop. 16 with final-stage convergence tolerance $\varepsilon/4$ and work-cap budget $\varepsilon/4$, and apply the resulting single kernel to both augmented laws. Its output from the comparison law is within $\varepsilon/2$ of stationarity. Because applying the same Markov kernel cannot increase total-variation distance, the two final outputs remain within $\varepsilon/4$ of one another. The triangle inequality therefore puts the algorithmic output within $3\varepsilon/4$ of stationarity. The finalizer's $\widetilde{O}(\sqrt{\kappa}d^{1/4})$ query cost is bounded by every branch of Ψ in its selected regime. This proves the second term in the minimum of (9). Alternatively, skip the coarse stage and apply the general-budget bound (45) directly to ν_0; this proves the first term.

Finally, every branch of Ψ is at most $\kappa^{6/7}d^{1/4}$. For the first branch the ratio is $\kappa^{-4/21}d^{1/12} \le 1$ because $d \le \kappa^{16/7}$. For the second it is $\kappa^{1/7}d^{-1/16} \le 1$ because $d \ge \kappa^{16/7}$, and the third follows from $\kappa \ge 1$. This proves the claimed envelope. □

D.5 Initialization from One Point

Let x_* be the minimizer of V, let $\log_+(t) = \max\{1, \log t\}$, and recall the localization budget

$$Q_{\mathrm{loc}} := \sqrt{\kappa}\,\log_+\Big[\kappa\{2 + \sqrt{\beta}\|x_0 - x_*\|\}\Big].$$

The quantity Q_{loc} appears only in the query bound: the algorithm does not know x_* and stops using the observable gradient test stated below.

Lemma 26 (Observable localization)**.** *Starting from x_0, accelerated optimization followed by one explicit Gaussian draw produces a law q_0 such that*

$$D_2(q_0\|\pi) \le 1 + \frac{d}{2}\log(2\kappa)$$

using $\widetilde{O}(Q_{\mathrm{loc}})$ oracle queries.

Proof. Run accelerated gradient descent until the observable certificate $\|\nabla V(x_{\mathrm{ref}})\| \le \alpha/\sqrt{\beta}$ holds. The standard strongly-convex accelerated rate and $V(x_0) - V(x_*) \le \beta\|x_0 - x_*\|^2/2$ give the stated Q_{loc} bound. Strong convexity then implies $\sqrt{\beta}\|x_{\mathrm{ref}} - x_*\| \le 1$.

Draw $q_0 = \mathcal{N}(x_{\mathrm{ref}}, (2\beta)^{-1}I_d)$. Strong convexity bounds the target normalizer by $Z \le \mathrm{e}^{-V(x_*)}(2\pi/\alpha)^{d/2}$, while smoothness gives $V(x) \le V(x_*) + \beta\|x - x_*\|^2/2$. Completing the square yields

$$\left\|\frac{\mathrm{d}q_0}{\mathrm{d}\pi}\right\|_\infty \le \mathrm{e}(2\kappa)^{d/2}.$$

An L^∞ density-ratio bound by M implies $D_2 \le \log M$, proving the result. □

Proof of the supplied-point part of Thm. 5. Apply Lem. 26, followed by Thm. 5 with $B = 1 + (d/2)\log(2\kappa)$. Since $\log(\mathrm{e}B)$ is polylogarithmic in d and κ, the coarse-wrapper branch costs $\widetilde{O}(\Psi(\kappa, d))$. Adding the one-time localization cost proves the first inequality in (10); the envelope proved above gives the second. □