WHEN STABILITY BINDS: HIGH-ACCURACY STOCHASTIC-GRADIENT LANGEVIN SAMPLING UNDER FINITE BUDGETS

Anonymous authors
Paper under double-blind review

ABSTRACT

Superlinear drifts can make explicit Langevin samplers unstable. But stabilization is not free: once it becomes necessary, the stabilizer itself can become a major source of sampling error. This creates a two-stage problem. A sampler must first remain stable, and its stabilization must then disturb the intended dynamics as little as possible. We study Relative-Growth Localized Taming (RELTA), an overdamped stochastic-gradient Langevin sampler with localized quadratic stabilization. It acts strongly in the tails while remaining inactive throughout the localized typical region. Despite using overdamped dynamics, stochastic gradients, and no Metropolis correction, RELTA achieves high finite-budget sampling accuracy across controlled, correlated, and real-data targets, and can outperform carefully tuned kinetic and exact-gradient methods under finite gradient-access budgets. We quantify how strongly stabilization perturbs the sampling dynamics as the step size decreases, and show how this perturbation contributes to stationary bias. Our results suggest a simple principle: when stability binds, stabilize first, but intervene as little as possible.

1 INTRODUCTION

Langevin algorithms are widely used for modern Monte Carlo sampling, particularly when gradients of the log target density are available. Under regular drift conditions, their finite-time and finite-step behavior is by now well understood (Dalalyan, 2017; Durmus and Moulines, 2017; Teh et al., 2016; Vollmer et al., 2016; Raginsky et al., 2017; Brosse et al., 2018). With superlinearly growing drifts or stochastic gradients, however, explicit Langevin updates can lose moment control or diverge (Hutzenthaler et al., 2011). Taming and related stabilization mechanisms are then used to recover stable explicit schemes (Hutzenthaler et al., 2012; Sabanis, 2013; Brosse et al., 2019; Lovas et al., 2023; Neufeld et al., 2025; Lytras et al., 2025; Wang and Zhang, 2026).

A common way to stabilize an explicit Langevin update is to tame the drift,

$$X_{n+1} = X_n - \lambda \frac{H(X_n, Z_{n+1})}{D_\lambda(X_n)} + \sqrt{2\lambda/\beta}\,\xi_{n+1}, \qquad D_\lambda(x) \geq 1.$$

The denominator D_λ controls excessive drift growth, but also changes the dynamics being sampled: replacing h by $h_\lambda = h/D_\lambda$ can introduce stationary bias. The central design question is therefore how to obtain sufficient tail stabilization while disturbing the typical-region dynamics as little as possible.

To quantify this tradeoff, we separate the far-tail strength of a stabilizer from its local perturbation order. Since

$$h(x) - h_\lambda(x) = h(x)\big(1 - D_\lambda(x)^{-1}\big),$$

the relative drift perturbation is $1 - D_\lambda(x)^{-1}$. If, on bounded-score states,

$$1 - D_\lambda(x)^{-1} = O(\lambda^\rho),$$

then stabilization changes one drift update only at order $O(\lambda^{\rho+1})$. We call ρ the *sampling order* of the stabilizer. Larger ρ means that stabilization disappears more rapidly from the finite-step dynamics as $\lambda \to 0$. For the denominator family introduced below, $\rho = p$.

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.

Relative-Growth Localized Taming (RELTA), introduced in prior work for nonconvex stochastic-gradient Langevin learning (Zhou and Chen, 2026), is designed to separate these two roles. Here we study its sampling-order theory, stationary bias, and finite-budget sampling behavior. It uses

$$X_{n+1} = X_n - \lambda \frac{H(X_n, Z_{n+1})}{\left(1 + (K\lambda S_{s_0}(X_n))^2\right)^{1/2}} + \sqrt{2\lambda/\beta}\,\xi_{n+1}.$$

Here S_{s_0} is a localized drift score constructed from the relative-growth condition that drives the Lyapunov stability analysis. Localization determines *where* stabilization acts, while the quadratic denominator controls *how strongly* it perturbs the typical-region dynamics. The method therefore remains strongly stabilizing in the tail while retaining sampling order two on bounded-score states, within the overdamped stochastic-gradient Langevin framework.

We compare RELTA against fourteen baselines spanning explicit stabilization, adaptive preconditioning, momentum, kinetic splitting, and corrected MCMC. Our experiments examine both the stable regime and the regime where tail stability becomes a limiting factor. Among overdamped stochastic-gradient methods, RELTA consistently achieves lower finite-budget sampling error. Momentum and high-accuracy kinetic methods become competitive at larger step sizes or budgets, where stability can begin to bind. In this regime, RELTA remains strictly stable, whereas several competing methods either become unstable or lose accuracy under stronger stabilization. This separation distinguishes ordinary finite-step error from the additional error induced by stabilization. The same stability–accuracy pattern persists after breaking coordinate separability and on a real-data Bayesian neural regression posterior.

Together, the theory and experiments support a simple design principle: stabilize the tail as strongly as necessary, but perturb the typical-region dynamics as little as possible.

2 Related Work

Finite-step Langevin sampling. Langevin Monte Carlo has been extensively analyzed under smoothness and convexity-type assumptions, with nonasymptotic bounds that separate finite-time and discretization errors (Dalalyan, 2017; Durmus and Moulines, 2017). Stochastic-gradient Langevin dynamics was introduced by Welling and Teh (2011), and later work studied its convergence and finite-step bias, including the additional error caused by stochastic gradients (Teh et al., 2016; Vollmer et al., 2016; Raginsky et al., 2017; Brosse et al., 2018).

Stabilization under superlinear and noisy gradients. Explicit Euler schemes can fail for non-globally Lipschitz SDEs (Hutzenthaler et al., 2011), while tamed schemes restore convergence under superlinear growth (Hutzenthaler et al., 2012; Sabanis, 2013). Langevin variants include TULA and stochastic-gradient or modified taming schemes (Brosse et al., 2019; Lovas et al., 2023; Neufeld et al., 2025; Lytras et al., 2025), with later work weakening regularity and dissipativity conditions (Lytras and Mertikopoulos, 2025). Projection gives another explicit stabilization route for nonconvex superlinear targets (Pang et al., 2025). For stability under noisy minibatch gradients, Mensch et al. (2026) propose the bounded-update SGLRW and analyze Clipped-SGLD, which stabilizes stochastic-gradient drift by clipping. Prior work by Zhou and Chen (2026) studies RELTA for nonconvex stochastic-gradient Langevin learning; here we study its sampling behavior and stabilization-induced stationary error.

Preconditioning, momentum, splitting, and corrected MCMC. pSGLD introduces RMSProp-style diagonal preconditioning into stochastic-gradient Langevin dynamics (Li et al., 2016), providing an adaptive baseline complementary to explicit tail stabilization. SGHMC introduces stochastic-gradient momentum (Chen et al., 2014), with broader stochastic-gradient MCMC constructions developed by Ma et al. (2015); tSGHMC extends this direction to superlinear stochastic gradients (Wang and Zhang, 2026). BAOAB provides a high-configurational-accuracy kinetic splitting baseline (Leimkuhler and Matthews, 2013). We also compare with MALA, HMC, and NUTS (Roberts and Tweedie, 1996; Neal, 2011; Hoffman and Gelman, 2014). Together these baselines test adaptive geometry, momentum, refined kinetic integration, and target correction across the controlled comparisons below. Their diagnostic roles are summarized in Appendix B.

3 STABILIZATION AS A SAMPLING PERTURBATION

Let

$$\pi_\beta(dx) \propto e^{-\beta U(x)}dx, \qquad h = \nabla U,$$

and let $H(x, Z)$ be an unbiased stochastic-gradient oracle, $\mathbb{E}[H(x, Z)] = h(x)$. We consider

$$X_{n+1} = X_n - \lambda \frac{H(X_n, Z_{n+1})}{D_\lambda(X_n)} + \sqrt{2\lambda/\beta}\,\xi_{n+1}, \qquad D_\lambda \geq 1, \tag{1}$$

with $\xi_n \sim \mathcal{N}(0, I_d)$. Write $R_\lambda f(x) = \mathbb{E}[f(X_{n+1}) \mid X_n = x]$ and

$$I(x) = \langle x, h(x)\rangle, \qquad Q(x) = \mathbb{E}\|H(x, Z)\|^2, \qquad V(x) = 1 + \|x\|^2.$$

3.1 TAIL STABILITY REQUIREMENT

A one-step calculation gives

$$R_\lambda V(x) - V(x) = -2\lambda \frac{I(x)}{D_\lambda(x)} + \lambda^2 \frac{Q(x)}{D_\lambda(x)^2} + 2\lambda\beta^{-1}d. \tag{2}$$

For the polynomial-growth class considered here, assume outside a compact set

$$Q(x) \leq C_Q I(x) S_{s_0}(x), \qquad S_{s_0}(x) = \big[\|x\|^{2r} - s_0\big]_+. \tag{3}$$

The exponent $r > 0$ specifies the tail growth scale of the score and is chosen so that Assumption A2 holds; in the quartic benchmark, $r = 1$. We call (3) the *relative-growth condition*. Because the diffusion term in (2) does not grow with $\|x\|$, the denominator must operate on the same far-tail scale as $K\lambda S_{s_0}(x)$. In particular, we require the lower bound

$$D_\lambda(x) \geq K\lambda S_{s_0}(x), \tag{4}$$

with K sufficiently large. The family introduced below also satisfies $D_{\lambda,p}(x) \leq 1 + K\lambda S_{s_0}(x)$. The lower bound controls the positive quadratic Q-term, while the matching upper scale preserves sufficient inward drift. Thus tail stability constrains the far-tail growth scale of D_λ, but not how much it should perturb ordinary states.

3.2 SAMPLING ORDER

We therefore consider

$$D_{\lambda,p}(x) = \big(1 + (K\lambda S_{s_0}(x))^p\big)^{1/p}, \qquad p \geq 1. \tag{5}$$

Every member has the same required far-tail scaling, while p controls the small-λ intervention order.

Proposition 1 (Tail strength and sampling order)**.** *For every $p \geq 1$,*

$$D_{\lambda,p}(x) \geq K\lambda S_{s_0}(x). \tag{6}$$

Moreover, on any set $\mathcal{B}$ with $\sup_{x\in\mathcal{B}} S_{s_0}(x) < \infty$,

$$1 - D_{\lambda,p}(x)^{-1} = \frac{K^p}{p} S_{s_0}(x)^p \lambda^p + O(\lambda^{2p}) \tag{7}$$

uniformly on $\mathcal{B}$. Hence stabilization changes one Euler drift update by $O(\lambda^{p+1})$.

We call p the *sampling order* of the stabilizer on bounded-score sets. It is the order of the stabilization perturbation, not the convergence order of the underlying Euler sampler. Proposition 1 is proved in Appendix A.2.

3.3 RELTA AND LOCALIZATION

RELTA takes the quadratic member $p = 2$:

$$D_\lambda^{\mathrm{RELTA}}(x) = \sqrt{1 + \big(K\lambda S_{s_0}(x)\big)^2}. \tag{8}$$

Below the localization threshold, $S_{s_0} = 0$ and the stabilizer is exactly inactive. On bounded-score states its relative intervention is $O(\lambda^2)$, while in the far tail $D_\lambda^{\mathrm{RELTA}} \sim K\lambda S_{s_0}$. Thus the sub-threshold core is exactly unperturbed; the $O(\lambda^2)$ statement concerns active states whose score remains bounded as $\lambda \to 0$. Localization therefore determines *where* stabilization acts, while the quadratic denominator determines how strongly an active stabilizer perturbs the typical-region dynamics.

The underlying Euler discretization has leading weak error $O(\lambda)$, whereas the stabilizer perturbs the dynamics at order $O(\lambda^p)$. Thus $p = 1$ can modify the leading bias coefficient, while any $p > 1$ moves the direct stabilization error to higher order. We use $p = 2$ as the simplest integer above this threshold; larger p can further reduce finite-step intervention but cannot remove the first-order Euler error.

4 THEORY

We next establish the stability, stationary consistency, and weak-bias properties of RELTA. Full assumptions and proofs are given in Appendix A.

Theorem 1 (Uniform Lyapunov stability)**.** *Suppose Assumptions A1–A2 and the chain-regularity condition A3(a) of Appendix A hold. Let $p \geq 1$ and choose $K > C_Q$. Then there exist $\lambda_0, c, C > 0$ such that, for every $0 < \lambda \leq \lambda_0$,*

$$R_{\lambda,p}V(x) \leq (1 - c\lambda)V(x) + C\lambda, \qquad x \in \mathbb{R}^d. \tag{9}$$

Consequently,

$$\sup_{n \geq 0} \mathbb{E}V(X_n) \leq V(X_0) + \frac{C}{c},$$

and the chain admits a unique invariant law $\pi_{\lambda,p}$ with $\pi_{\lambda,p}(V) \leq C/c$. In particular, the quadratic RELTA *choice $p = 2$ retains the same far-tail stability mechanism as $p = 1$.*

This theorem provides a Lyapunov basis for tail recovery: from far-tail initializations, the expected Lyapunov value contracts geometrically toward the typical scale rather than remaining trapped in the tail.

Proposition 2 (Stationary consistency)**.** *Under Assumptions A1–A3 of Appendix A, for every fixed $p \geq 1$,*

$$\pi_{\lambda,p} \Rightarrow \pi_\beta \qquad \text{as } \lambda \downarrow 0.$$

Together with the uniform second-moment bound in Theorem 1, this weak convergence further implies

$$W_1(\pi_{\lambda,p}, \pi_\beta) \longrightarrow 0. \tag{10}$$

Stationary consistency establishes correctness. To compare consistent stabilizers, we next examine the leading stationary weak bias.

Theorem 2 (Leading stationary weak bias)**.** *Suppose Assumptions A1–A4 of Appendix A hold, and let f belong to the observable class specified there. Then*

$$\pi_{\lambda,p}(f) - \pi_\beta(f) = \begin{cases} \lambda(C_{\mathrm{E}}(f) + C_{\mathrm{S}}(f)) + o(\lambda), & p = 1, \\ \lambda C_{\mathrm{E}}(f) + o(\lambda), & p > 1. \end{cases} \tag{11}$$

Here C_{E} is the leading weak-bias coefficient of the underlying stochastic-gradient Euler step, while C_{S} is the additional coefficient induced by first-order stabilization.

Theorem 1 supplies the uniform second-moment control used for stationary consistency; the higher polynomial moments and weighted remainder bounds required for the refined weak-bias expansion are additional conditions in Assumption A4.

For $p = 2$, the stabilizer first enters the modified generator at order λ^2 and therefore leaves the leading $O(\lambda)$ Euler coefficient in (11) unchanged. The underlying Euler discretization remains first order. The $p = 1, 2, 4$ ablation tests the empirical consequences of this sampling-order distinction.

5 EXPERIMENTS

We compare RELTA across stabilized, momentum, kinetic, and corrected samplers, then test generalization and stabilizer order. Additional clipping, lattice-update, and projection comparisons appear in Appendix B.7.

Benchmark. We use a $d = 50$ quartic finite-sum target

$$\pi_\beta(x) \propto \exp\{-\beta U(x)\},$$

where

$$U(x) = \sum_{j=1}^{d} \left[\frac{1}{4m} \sum_{k=1}^{m} (x_j - b_{jk})^4 + \frac{\gamma}{2} x_j^2 \right],$$

with $m = 64$, $\beta = 10$, and $\gamma = 0.05$, where $b_{jk} = 1 + 0.35\,\varepsilon_{jk}$ and $\varepsilon_{jk} \sim \mathcal{N}(0, 1)$. The quartic potential induces cubic-gradient growth and hence a controlled superlinear stability stress. Meanwhile, the product structure allows each marginal reference distribution to be computed accurately by one-dimensional numerical quadrature.

Evaluation. We evaluate each sampler along two axes: accuracy and strict stability. Accuracy is measured by exact-reference marginal W_1. Strict stability requires both that the chain remain numerically well-behaved over the full run and that it recover efficiently from far-tail excursions. To test the latter, we initialize all coordinates at 5 and ask whether the chain returns to the prescribed typical-radius region within a fixed horizon. For methods with different per-transition costs, we match gradient-evaluation costs and report full-gradient equivalents (FGE).

5.1 OVERDAMPED STABILIZATION: NECESSITY AND DESIGN

We compare RELTA with vanilla SGLD and three existing overdamped stabilizers. SGLD provides the unstabilized control, while TUSLA provides a stochastic-gradient Langevin baseline with globally active radial taming, allowing us to compare stabilization mechanisms under the same stochastic-gradient setting. mTULA and kTULA are stronger exact-gradient baselines: mTULA improves the known Wasserstein discretization rates over TUSLA, while kTULA provides further improved KL and Wasserstein guarantees (Lovas et al., 2023; Neufeld et al., 2025; Lytras et al., 2025). Thus, the comparison tests not only whether stabilization is necessary, but also whether stronger existing taming designs and exact gradients translate into lower finite-step sampling error under active stability constraints.

All hyperparameters are selected on separate tuning instances and then frozen for ten held-out instances. We evaluate the main stress band

$$\lambda \in \{0.08, 0.085, 0.09, 0.095\}.$$

Vanilla SGLD fails numerical survival from the typical-region initialization on all ten held-out instances at every step size in the stress band. Thus, stabilization is genuinely necessary in this regime.

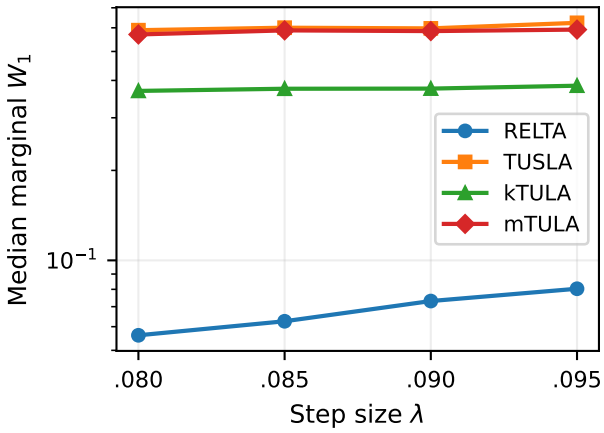


Figure 1: Overdamped exact-reference marginal W_1; SGLD fails numerically.

Table 1: Overdamped comparison at $\lambda = 0.09$; "–" denotes numerical failure.

Method	Gradient	Strict stability	median W_1
RELTA	mini-batch	10/10	**0.0731**
SGLD	mini-batch	0/10	–
TUSLA	mini-batch	1/10	0.5973
kTULA	full	10/10	0.3756
mTULA	full	1/10	0.5850

RELTA achieves the lowest W_1 throughout the stress band while remaining strictly stable on all held-out instances. kTULA is also $10/10$ strictly stable at every tested step size, so the large accuracy gap persists even among strictly stable methods. TUSLA and mTULA both survive from the typical-region initialization but recover poorly from the tail, reaching only $1/10$ strict stability at $\lambda = 0.09$. Thus, exact gradients alone do not guarantee strict stability. The large gap to strictly stable kTULA further shows that exact gradients alone do not explain the accuracy difference.

5.2 Momentum: accuracy and the stability boundary

We next ask whether momentum alone removes the accuracy advantage of RELTA. We compare against carefully tuned SGHMC, selecting the friction parameter on separate tuning instances and freezing it before evaluation on fresh targets. This yields two distinct regimes.

Stage I: comparison before stability becomes binding. We first consider

$$\lambda \in \{0.070, 0.075, 0.080, 0.085, 0.090, 0.095\},$$

where both RELTA and tuned SGHMC satisfy strict stability on all $20/20$ fresh instances. Thus, differences in W_1 in this regime cannot be attributed to failures of the momentum baseline.

Table 2: RELTA versus tuned SGHMC in the strictly stable regime.

λ	RELTA W_1	SGHMC W_1	RELTA wins
0.070	**0.0469**	0.0804	20/20
0.075	**0.0527**	0.0841	20/20
0.080	**0.0607**	0.0822	19/20
0.085	**0.0675**	0.0825	19/20
0.090	**0.0780**	0.0827	12/20
0.095	0.0862	**0.0817**	9/20

Across the strictly stable regime, RELTA is substantially more accurate at smaller step sizes, with the advantage narrowing as λ increases and disappearing around $\lambda = 0.09$–0.095. Thus, momentum does not automatically improve sampling accuracy even in the absence of active stability pressure.

Stage II: the large-step advantage reaches the tail-stability boundary. We next increase the step size into the regime favorable to SGHMC. At $\lambda = 0.12$, SGHMC attains the lowest median W_1, but one far-tail trajectory diverges despite all typical-region chains surviving, reducing strict stability to $9/10$.

Table 3: Momentum comparison at $\lambda = 0.12$.

Method	Strict stability	median W_1
SGHMC	9/10	**0.0935**
RELTA	10/10	0.1473
tSGHMC	10/10	1.0288

We therefore compare against tSGHMC, a tamed SGHMC variant designed for superlinear stochastic gradients.

The standard tSGHMC stabilization eliminates the far-tail divergence observed above and restores $10/10$ strict stability, but its median W_1 increases from 0.0935 for untamed SGHMC to 1.0288. In contrast, RELTA remains $10/10$ strictly stable with median $W_1 = 0.1473$. Thus, once stabilization becomes binding, the choice of stabilization can matter more than momentum alone.

Overall, RELTA is substantially more accurate than tuned SGHMC at smaller step sizes, while SGHMC becomes competitive only at larger step sizes and eventually encounters a tail-stability boundary. This reinforces the need for stabilization that minimally perturbs the typical region.

5.3 Adaptive preconditioning and kinetic splitting

We next compare RELTA with two strong practical baselines that improve finite-budget sampling through complementary mechanisms: pSGLD (Li et al., 2016) uses adaptive diagonal preconditioning, while BAOAB uses high-accuracy kinetic splitting. Hyperparameters are tuned on separate targets and frozen before evaluation.

Table 4: Strong practical baselines under matched data-access budgets.

Budget	Method	Configuration	Typical	Tail	median W_1
2000 FGE	RELTA	$\lambda = .014$	10/10	10/10	**0.01071**
	pSGLD	$h = .05$	10/10	10/10	0.01153
	BAOAB (accuracy)	$(\lambda, \gamma) = (.82, .35)$	9/10	0/10	0.01420
	BAOAB (stability-constr.)	$(.34, .30)$	10/10	10/10	0.02224
5000 FGE	RELTA	$\lambda = .009$	10/10	10/10	**0.00885**
	pSGLD	$h = .03$	10/10	10/10	0.00958
	BAOAB (accuracy)	$(.78, .175)$	10/10	0/10	0.00953
	BAOAB (stability-constr.)	$(.30, .25)$	10/10	9/10	0.01414
10000 FGE	RELTA	$\lambda = .006$	10/10	10/10	0.00823
	pSGLD	$h = .02$	10/10	10/10	0.00918
	BAOAB (accuracy)	$(.75, .20)$	10/10	0/10	**0.00742**
	BAOAB (stability-constr.)	$(.36, .50)$	10/10	10/10	0.01004

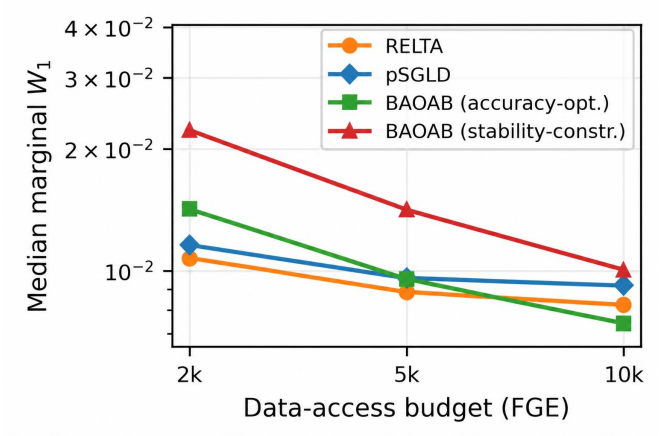


Figure 2: Finite-budget accuracy against pSGLD and BAOAB.

pSGLD remains stable across all budgets, but RELTA achieves lower median W_1 throughout. BAOAB reaches the lowest median W_1 at 10000 FGE, but the parameter setting that achieves this accuracy fails tail recovery on all ten targets. Enforcing tail stability removes this advantage and reveals the accuracy–stability tradeoff.

5.4 Corrected MCMC under finite FGE budgets

We finally ask whether target correction improves sampling accuracy under a finite gradient-access budget. MALA, HMC, and NUTS use correction mechanisms to preserve the target distribution but require full-data information. We therefore compare them with RELTA under matched FGE budgets.

We report typical-region sampling accuracy in the main comparison. Tail-start stability checks and the stability-oriented configurations used for fixed-step MALA and HMC are described in Appendix B.

Table 5: Finite-budget comparison with corrected MCMC methods.

Method	5000 FGE	10000 FGE
RELTA	**0.00890**	**0.00752**
HMC	0.01764	0.01260
NUTS	0.02250	0.01440
MALA	0.03327	0.02273

RELTA has the lowest median W_1 at both budgets and wins every paired held-out comparison. NUTS improves substantially as the budget doubles, but remains above RELTA at 10000 FGE.

Thus, asymptotic exactness does not automatically translate into higher accuracy under a substantial but finite FGE budget.

5.5 Generalization beyond the controlled benchmark

The product-form quartic target gives exact marginal ground truth, but its separability is deliberately controlled. We therefore test two departures from this setting at 2000 FGE; full constructions and tuning protocols are in Appendix B.

Dense correlated target. A fixed dense, non-orthogonal transform introduces non-separability and moderate anisotropy. On previously unused targets, RELTA and pSGLD are both $10/10$ strictly stable, with median sliced W_1 0.02016 and 0.02089, respectively; RELTA wins $7/10$ paired comparisons. Accuracy-tuned BAOAB has median sliced W_1 0.02654 and fails all tail-recovery tests, while stability-constrained BAOAB is $10/10$ strictly stable but has median sliced W_1 0.14587.

Real-data Bayesian neural regression. On the Diabetes data (Efron et al., 2004), we use a 385-parameter shallow Bayesian regressor and select hyperparameters on validation before test evaluation. At the selected $h = .020$, RELTA never binds on typical test trajectories, so matched-h SGLD gives identical predictive results. Under stress, RELTA succeeds $10/10$ while SGLD succeeds $0/10$; strict SGHMC is also $10/10$ but has slightly higher test NLL (Appendix B.10).

Table 6: Generalization beyond the controlled benchmark at 2000 FGE.

Setting	Method	Typical	Stress/tail	Primary metric ↓
Correlated quartic (sliced W_1)	RELTA	10/10	10/10	**0.02016**
	pSGLD	10/10	10/10	0.02089
	BAOAB, accuracy	10/10	0/10	0.02654
	BAOAB, stability-constr.	10/10	10/10	0.14587
Bayesian regression (test NLL)	RELTA ($h = .020$)	10/10	10/10	**0.97812**
	SGLD, matched-h	10/10	0/10	**0.97812**
	SGHMC, strict	10/10	10/10	0.98197
	SGHMC, accuracy	10/10	0/10	0.98125

5.6 Mechanism check: sampling order of the stabilizer

We isolate the stabilizer order using

$$D_{\lambda,p} = (1 + (K\lambda S)^p)^{1/p}, \qquad p \in \{1, 2, 4\},$$

with all other settings fixed. The choices share the same far-tail scaling $D_{\lambda,p} \sim K\lambda S$, while their bounded-score perturbations are $O(\lambda^p)$. As predicted, $p = 2$ improves W_1 over $p = 1$ in all $60/60$ paired comparisons with strict stability, while $p = 4$ improves a further $59/60$ by a smaller margin. We therefore use $p = 2$ as the simplest higher-order choice.

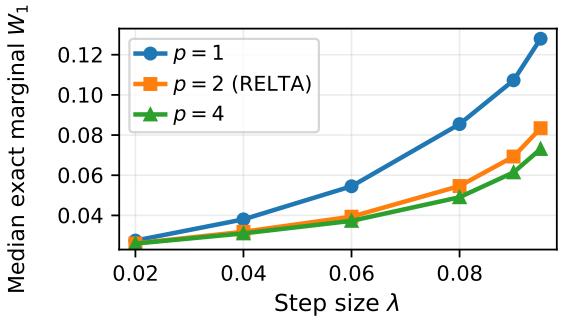


Figure 3: Sampling-order ablation. Higher-order denominators reduce sampling error while preserving strict stability.

6 Discussion and Conclusion

RELTA separates the tail scale required for stability from intervention on ordinary states. For $p > 1$, the required far-tail growth is retained while the leading $O(\lambda)$ stationary weak-bias coefficient remains that of Euler. Finite-budget experiments and the $p = 1, 2, 4$ ablation show that reducing local intervention improves accuracy without sacrificing strict stability.

Natural extensions include higher-order discretizations and geometry-aware localization. Moving from this conceptual design to a practical sampler will require automatic stabilization and step-size calibration while preserving local accuracy.

Reproducibility Statement

Appendix A provides complete assumptions and proofs, and Appendix B specifies the target construction, tuning/held-out splits, stability criteria, computational accounting, and frozen hyperparameters for the experiments. The anonymous supplementary material includes the experimental code, raw per-seed outputs, summary tables, and audit scripts needed to reproduce the reported results.

AI Use Statement

Generative AI tools were used to assist with experimental design; coding and debugging of experimental scripts; interpretation of experimental results; editing, adaptation, and checking of selected mathematical derivations and proof text; literature search; and assistance in drafting selected passages, manuscript editing, and language refinement. All theoretical claims and AI-assisted mathematical work were verified by the author, and experimental code and reported results were rerun and/or audited by the author as documented in the reproducibility materials. References identified or suggested with AI assistance were checked against the original sources. The author reviewed all AI-assisted work and takes responsibility for the final content of this paper.

References

N. Brosse, A. Durmus, and E. Moulines. The promises and pitfalls of stochastic gradient Langevin dynamics. In *Advances in Neural Information Processing Systems*, volume 31, pages 8278–8288, 2018.

N. Brosse, A. Durmus, E. Moulines, and S. Sabanis. The tamed unadjusted Langevin algorithm. *Stochastic Processes and their Applications*, 129(10):3638–3663, 2019.

T. Chen, E. Fox, and C. Guestrin. Stochastic gradient Hamiltonian Monte Carlo. In *Proceedings of the 31st International Conference on Machine Learning*, volume 32 of *Proceedings of Machine Learning Research*, pages 1683–1691, 2014.

A. S. Dalalyan. Theoretical guarantees for approximate sampling from smooth and log-concave densities. *Journal of the Royal Statistical Society: Series B*, 79(3):651–676, 2017.

A. Durmus and E. Moulines. Nonasymptotic convergence analysis for the unadjusted Langevin algorithm. *The Annals of Applied Probability*, 27(3):1551–1587, 2017.

B. Efron, T. Hastie, I. Johnstone, and R. Tibshirani. Least angle regression. *The Annals of Statistics*, 32(2):407–451, 2004.

M. D. Hoffman and A. Gelman. The No-U-Turn sampler: Adaptively setting path lengths in Hamiltonian Monte Carlo. *JMLR*, 15:1593–1623, 2014.

M. Hutzenthaler, A. Jentzen, and P. E. Kloeden. Strong and weak divergence in finite time of Euler's method for stochastic differential equations with non-globally Lipschitz coefficients. *Proceedings of the Royal Society A*, 467:1563–1576, 2011.

M. Hutzenthaler, A. Jentzen, and P. E. Kloeden. Strong convergence of an explicit numerical method for SDEs with nonglobally Lipschitz continuous coefficients. *The Annals of Applied Probability*, 22(4):1611–1641, 2012.

B. Leimkuhler and C. Matthews. Rational construction of stochastic numerical methods for molecular sampling. *Applied Mathematics Research eXpress*, 2013(1):34–56, 2013.

C. Li, C. Chen, D. Carlson, and L. Carin. Preconditioned stochastic gradient Langevin dynamics for deep neural networks. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, pages 1788–1794, 2016.

A. Lovas, I. Lytras, M. Rasonyi, and S. Sabanis. Taming neural networks with TUSLA: Nonconvex learning via adaptive stochastic gradient Langevin algorithms. *SIAM Journal on Mathematics of Data Science*, 5(2):323–345, 2023.

I. Lytras and P. Mertikopoulos. Tamed Langevin sampling under weaker conditions. In *Proceedings of the 28th International Conference on Artificial Intelligence and Statistics*, volume 258 of *Proceedings of Machine Learning Research*, pages 847–855, 2025.

I. Lytras, S. Sabanis, and Y. Zhang. kTULA: A Langevin sampling algorithm with improved KL bounds under super-linear log-gradients. *arXiv:2506.04878*, 2025.

Y.-A. Ma, T. Chen, and E. B. Fox. A complete recipe for stochastic gradient MCMC. In *Advances in Neural Information Processing Systems*, volume 28, pages 2917–2925, 2015.

Z. Mensch, L. Holdijk, S. Duffield, M. Aifer, P. J. Coles, M. Welling, and M. C. N. Cheng. Robust stochastic gradient posterior sampling with lattice based discretisation. In *Proceedings of the 43rd International Conference on Machine Learning*, 2026.

R. M. Neal. MCMC using Hamiltonian dynamics. In S. Brooks, A. Gelman, G. L. Jones, and X.-L. Meng, editors, *Handbook of Markov Chain Monte Carlo*, pages 113–162. Chapman & Hall/CRC, 2011.

A. Neufeld, M. C. E. Ng, and Y. Zhang. Non-asymptotic convergence bounds for modified tamed unadjusted Langevin algorithm in non-convex setting. *Journal of Mathematical Analysis and Applications*, 543(1):128892, 2025.

C. Pang, X. Wang, and Y. Wu. Projected Langevin Monte Carlo algorithms in non-convex and super-linear setting. *Journal of Computational Physics*, 526:113754, 2025.

M. Raginsky, A. Rakhlin, and M. Telgarsky. Non-convex learning via stochastic gradient Langevin dynamics: A nonasymptotic analysis. In *Proceedings of the 2017 Conference on Learning Theory*, volume 65 of *Proceedings of Machine Learning Research*, pages 1674–1703, 2017.

G. O. Roberts and R. L. Tweedie. Exponential convergence of Langevin distributions and their discrete approximations. *Bernoulli*, 2(4):341–363, 1996.

S. Sabanis. A note on tamed Euler approximations. *Electronic Communications in Probability*, 18(47):1–10, 2013.

Y. W. Teh, A. H. Thiery, and S. J. Vollmer. Consistency and fluctuations for stochastic gradient Langevin dynamics. *Journal of Machine Learning Research*, 17(7):1–33, 2016.

S. J. Vollmer, K. C. Zygalakis, and Y. W. Teh. Exploration of the (non-)asymptotic bias and variance of stochastic gradient Langevin dynamics. *Journal of Machine Learning Research*, 17(159):1–48, 2016.

Z. Wang and Y. Zhang. Tamed stochastic gradient Hamiltonian Monte Carlo. *arXiv:2607.14862*, 2026.

M. Welling and Y. W. Teh. Bayesian learning via stochastic gradient Langevin dynamics. In *Proceedings of the 28th International Conference on Machine Learning*, pages 681–688, 2011.

Y. Zhou and Z. Chen. RELTA-SGLD: Relative-Growth Localized Taming for Nonconvex Stochastic-Gradient Langevin Learning. *arXiv:2607.19544*, 2026.

A Theory Details and Proofs

This appendix gives a self-contained assumption interface and proofs for the results in Sections 3–4.

A.1 Assumptions and notation

For $p \geq 1$, let $R_{\lambda,p}$ be the Markov kernel generated by (1) with $D_\lambda = D_{\lambda,p}$ and $S = S_{s_0}$. Let

$$Lf = -\langle h, \nabla f \rangle + \beta^{-1}\Delta f$$

denote the generator of the target overdamped Langevin diffusion, and write

$$M_2(x) = \mathbb{E}\big[H(x,Z)H(x,Z)^\top\big].$$

Assumption A1 (basic oracle regularity). The oracle is unbiased,

$$\mathbb{E}[H(x,Z)] = h(x),$$

and

$$Q(x) := \mathbb{E}\|H(x,Z)\|^2$$

is locally bounded. The map $x \mapsto H(x,z)$ is continuous for almost every z, with a locally integrable envelope sufficient to make $R_{\lambda,p}$ a Feller kernel. The deterministic score S_{s_0} is locally bounded.

Assumption A2 (relative growth and radial coercivity). There exist $R, c_I, C_Q > 0$, with R taken large enough that $R^{2r} > s_0$, such that, for $\|x\| \geq R$,

$$I(x) \geq c_I V(x)\big(1 + S_{s_0}(x)\big), \qquad Q(x) \leq C_Q I(x) S_{s_0}(x). \tag{12}$$

On the ball $\{\|x\| \leq R\}$, I and Q are bounded. Thus the same score that controls the relative growth Q/I also ensures that $I/D_{\lambda,p}$ retains a lower bound proportional to V in the tail.

Assumption A3 (Markov regularity and stationary identification). *(a) Chain regularity.* For all sufficiently small λ, $R_{\lambda,p}$ is ψ-irreducible and aperiodic.

(b) Target identification. The target Langevin diffusion has unique invariant law π_β, and infinitesimal invariance on $C_c^\infty(\mathbb{R}^d)$ characterizes this law: if a probability measure μ with $\mu(V) < \infty$ satisfies

$$\mu(Lf) = 0, \qquad f \in C_c^\infty(\mathbb{R}^d),$$

then $\mu = \pi_\beta$.

(c) Weighted generator consistency. For every $f \in C_c^\infty(\mathbb{R}^d)$,

$$\left\| \frac{R_{\lambda,p} f - f}{\lambda} - Lf \right\|_V \longrightarrow 0, \qquad \|g\|_V := \sup_x \frac{|g(x)|}{V(x)}. \tag{13}$$

Clause (a) is used only for uniqueness of the invariant law in the stability result. Clauses (b)–(c) are used only to identify the small-step stationary limit.

Assumption A4 (additional regularity for the refined weak-bias expansion). For every observable f in a class $\mathcal{F}$, the centered Poisson equation

$$L\phi_f = f - \pi_\beta(f) \tag{14}$$

has a solution $\phi_f \in C^4$ whose derivatives up to order four have at most polynomial growth.

The oracle has finite fourth moments with polynomial growth: for some $m < \infty$,

$$\mathbb{E}\|H(x, Z)\|^4 \leq C(1 + \|x\|^m).$$

There is a coercive polynomial weight $W \geq V$ such that

$$\sup_{0<\lambda\leq\lambda_0} \pi_{\lambda,p}(W) < \infty$$

for the values of p under consideration. The coefficient functions $A_{\mathrm{E}}\phi_f$ and $A_{\mathrm{S}}\phi_f$ defined below are continuous and $o(W)$ at infinity. The Taylor and denominator remainders used below are $o(\lambda^2)$ in the weighted norm

$$\|g\|_W := \sup_x \frac{|g(x)|}{W(x)}.$$

A.2 Stability scale and tail stability

We first record the local sampling order of the denominator and then prove the common far-tail Lyapunov estimate.

Proof of Proposition 1. Set

$$a = K\lambda S(x).$$

Since

$$(1 + a^p)^{1/p} \geq a,$$

the stability-scale bound (6) follows.

On every set where S is bounded, $a \to 0$ uniformly as $\lambda \downarrow 0$, and

$$(1 + a^p)^{-1/p} = 1 - \frac{1}{p} a^p + O(a^{2p}).$$

Hence

$$1 - D_{\lambda,p}^{-1} = \frac{K^p}{p} S(x)^p \lambda^p + O(\lambda^{2p})$$

locally uniformly. Multiplying the relative drift change by the step length λ gives the $O(\lambda^{p+1})$ one-step stabilization perturbation. □

Proof of Theorem 1. Fix $p \geq 1$ and write

$$S = S_{s_0}(x), \qquad D = D_{\lambda,p}(x).$$

From (5),

$$D \geq K\lambda S, \qquad D \leq 1 + K\lambda S. \tag{15}$$

Using the first inequality and (12), outside the ball of radius R,

$$\lambda^2 \frac{Q}{D^2} \leq C_Q \lambda^2 \frac{IS}{D^2} = C_Q \lambda \frac{I}{D} \frac{\lambda S}{D} \leq \frac{C_Q}{K} \lambda \frac{I}{D}.$$

Substituting into (2) gives

$$R_{\lambda,p} V - V \leq -\left(2 - \frac{C_Q}{K}\right) \lambda \frac{I}{D} + 2\lambda\beta^{-1} d. \tag{16}$$

Since $K > C_Q$, the coefficient in parentheses is strictly positive.

Choose $\lambda_0 > 0$ sufficiently small and set

$$c_K := \max\{1, K\lambda_0\}.$$

For $0 < \lambda \leq \lambda_0$, the second inequality in (15) gives

$$D \leq c_K(1 + S).$$

By the coercivity condition in (12),

$$\frac{I(x)}{D_{\lambda,p}(x)} \geq \frac{c_I}{c_K} V(x), \qquad \|x\| \geq R.$$

Hence there exist constants $c_1, C_1 > 0$, independent of λ, such that

$$R_{\lambda,p} V(x) - V(x) \leq -c_1 \lambda V(x) + C_1 \lambda, \qquad \|x\| \geq R.$$

On the compact ball $\{\|x\| \leq R\}$, boundedness of I and Q, together with $D \geq 1$, gives the same inequality after increasing C_1. This proves (9).

Iterating (9) gives

$$\mathbb{E}V(X_n) \leq (1 - c\lambda)^n V(X_0) + C\lambda \sum_{j=0}^{n-1} (1 - c\lambda)^j \leq V(X_0) + \frac{C}{c}.$$

Let

$$\nu_N = \frac{1}{N} \sum_{n=0}^{N-1} \mathcal{L}(X_n)$$

be the Cesàro averages. The preceding bound implies tightness. Since $R_{\lambda,p}$ is Feller, the Krylov–Bogoliubov argument yields an invariant probability measure along a weakly convergent subsequence. Lower semicontinuity of V gives

$$\pi_{\lambda,p}(V) \leq \frac{C}{c}.$$

Assumption A3(a) gives uniqueness of the invariant law. □

The proof uses only the two denominator bounds in (15). Thus the far-tail stability mechanism is uniform over $p \geq 1$: changing p alters the local intervention order without changing the stability scale.

A.3 STATIONARY CONSISTENCY

Lemma 1 (Local generator consistency)**.** *Under Assumption A1, for every $f \in C_c^\infty(\mathbb{R}^d)$ and every compact set $\mathcal{K} \subset \mathbb{R}^d$,*

$$\sup_{x\in\mathcal{K}} \left| \frac{R_{\lambda,p} f(x) - f(x)}{\lambda} - Lf(x) \right| \longrightarrow 0 \qquad \text{as } \lambda \downarrow 0.$$

Proof. Write

$$\Delta X = -\lambda H(x, Z) D_{\lambda,p}(x)^{-1} + \sqrt{2\lambda/\beta}\, \xi.$$

On a fixed compact set $\mathcal{K}$, local boundedness of S_{s_0} gives

$$D_{\lambda,p}^{-1} \longrightarrow 1$$

uniformly. Assumption A1 also gives a uniform bound on $\mathbb{E}\|H(x, Z)\|^2$ over $\mathcal{K}$.

A second-order Taylor expansion gives

$$f(x + \Delta X) - f(x) = \nabla f(x)^\top \Delta X + \frac{1}{2} \Delta X^\top \nabla^2 f(x) \Delta X + \rho_f(x, \Delta X),$$

where the remainder is $o(\|\Delta X\|^2)$ locally uniformly as $\Delta X \to 0$. Taking conditional expectation,

$$\mathbb{E}[\Delta X] = -\lambda D_{\lambda,p}^{-1} h(x),$$

and

$$\mathbb{E}[\Delta X \Delta X^\top] = \frac{2\lambda}{\beta} I_d + \lambda^2 D_{\lambda,p}^{-2} M_2(x).$$

The second term above is $O(\lambda^2)$ uniformly on $\mathcal{K}$.

The Taylor remainder satisfies

$$\sup_{x\in\mathcal{K}} \frac{\mathbb{E}|\rho_f(x, \Delta X)|}{\lambda} \longrightarrow 0.$$

Indeed, one may split according to $\{\|\Delta X\| \le \eta\}$ and its complement. Uniform continuity of $\nabla^2 f$ controls the first part, while the Gaussian tail and the factor λ multiplying the oracle drift control the second using the local second-moment bound in Assumption A1.

Therefore, uniformly for $x \in \mathcal{K}$,

$$\frac{R_{\lambda,p} f - f}{\lambda} = -D_{\lambda,p}^{-1} \langle h, \nabla f \rangle + \beta^{-1} \Delta f + o(1).$$

Since $D_{\lambda,p}^{-1} \to 1$ uniformly on $\mathcal{K}$, the right-hand side converges uniformly to

$$-\langle h, \nabla f \rangle + \beta^{-1} \Delta f = Lf.$$

□

The weighted condition (13) is used only to control the far-tail contribution when passing the invariant laws to the small-step limit.

Proof of Proposition 2. Let $\lambda_k \downarrow 0$. By Theorem 1,

$$\sup_k \pi_{\lambda_k,p}(V) < \infty,$$

so the family $\{\pi_{\lambda_k,p}\}$ is tight. Let

$$\pi_{\lambda_{k_j},p} \Rightarrow \pi_\star$$

along a convergent subsequence.

For $f \in C_c^\infty(\mathbb{R}^d)$, invariance gives

$$0 = \pi_{\lambda_{k_j},p} \left(\frac{R_{\lambda_{k_j},p} f - f}{\lambda_{k_j}} \right).$$

By Assumption A3(c),

$$\left|\pi_{\lambda_{k_j},p}\left(\frac{R_{\lambda_{k_j},p}f - f}{\lambda_{k_j}} - Lf\right)\right| \le \left\|\frac{R_{\lambda_{k_j},p}f - f}{\lambda_{k_j}} - Lf\right\|_V \pi_{\lambda_{k_j},p}(V) \longrightarrow 0.$$

Since $Lf \in C_c(\mathbb{R}^d)$, weak convergence then gives

$$\pi_\star(Lf) = 0, \qquad f \in C_c^\infty(\mathbb{R}^d).$$

The uniform V-moment bound and lower semicontinuity of V imply

$$\pi_\star(V) < \infty.$$

Assumption A3(b) therefore yields

$$\pi_\star = \pi_\beta.$$

Every convergent subsequence has the same limit, and hence

$$\pi_{\lambda,p} \Rightarrow \pi_\beta.$$

Finally,

$$V(x) = 1 + \|x\|^2$$

and the uniform V-moment bound imply uniform integrability of the first moments. Weak convergence therefore yields

$$W_1(\pi_{\lambda,p}, \pi_\beta) \longrightarrow 0,$$

which proves (10). □

A.4 Refined stationary weak-bias analysis

For the refined weak-bias analysis only, we additionally impose Assumption A4.

Proposition 3 (Local weak expansion)**.** *Under Assumptions A1 and A4, for every $f \in \mathcal{F}$ and its Poisson solution ϕ_f from* (14)*,*

$$R_{\lambda,p}\phi_f = \phi_f + \lambda L\phi_f + \lambda^2 A_{\mathrm{E}}\phi_f + \mathbf{1}_{\{p=1\}}\lambda^2 A_{\mathrm{S}}\phi_f + r_{\lambda,p}(\phi_f), \tag{17}$$

where

$$\|r_{\lambda,p}(\phi_f)\|_W = o(\lambda^2),$$

and

$$A_{\mathrm{E}}\varphi = \frac{1}{2}M_2 : \nabla^2\varphi - \beta^{-1}\langle h, \nabla\Delta\varphi\rangle + \frac{1}{2\beta^2}\Delta^2\varphi, \tag{18}$$

$$A_{\mathrm{S}}\varphi = KS_{s_0}\langle h, \nabla\varphi\rangle. \tag{19}$$

Thus first-order stabilization ($p = 1$) enters the $O(\lambda^2)$ one-step coefficient, whereas for every $p > 1$ it does not.

Proof. Write

$$\Delta X = -\lambda H(x, Z)D_{\lambda,p}(x)^{-1} + \sqrt{2\lambda/\beta}\,\xi.$$

Proposition 1 gives

$$D_{\lambda,p}^{-1} = 1 - K\lambda S_{s_0} + O(\lambda^2) \qquad \text{if } p = 1,$$

whereas

$$D_{\lambda,p}^{-1} = 1 + o(\lambda) \qquad \text{if } p > 1.$$

Hence

$$\mathbb{E}[\Delta X] = -\lambda h + \mathbf{1}_{\{p=1\}}K\lambda^2 S_{s_0}h + o_W(\lambda^2). \tag{20}$$

At second weak order, replacing $D_{\lambda,p}^{-1}$ by one in the second and higher increment moments affects only the remainder. Using independence of H and ξ and the Gaussian moments,

$$\mathbb{E}[\Delta X_i \Delta X_j] = \frac{2\lambda}{\beta}\delta_{ij} + \lambda^2 M_{2,ij} + o_W(\lambda^2),$$
$$\mathbb{E}[\Delta X_i \Delta X_j \Delta X_k] = -\frac{2\lambda^2}{\beta}\big(h_i\delta_{jk} + h_j\delta_{ik} + h_k\delta_{ij}\big) + o_W(\lambda^2),$$
$$\mathbb{E}[\Delta X_i \Delta X_j \Delta X_k \Delta X_\ell] = \frac{4\lambda^2}{\beta^2}\big(\delta_{ij}\delta_{k\ell} + \delta_{ik}\delta_{j\ell} + \delta_{i\ell}\delta_{jk}\big) + o_W(\lambda^2).$$

A fourth-order Taylor expansion of $\phi_f(x + \Delta X)$ followed by conditional expectation gives the order-λ term

$$-\langle h, \nabla\phi_f\rangle + \beta^{-1}\Delta\phi_f = L\phi_f.$$

The order-λ^2 contractions yield $A_{\mathrm{E}}\phi_f$, while the additional mean term in (20) yields $A_{\mathrm{S}}\phi_f$ when $p = 1$. Assumption A4 controls the weighted remainder. □

Proof of Theorem 2. Fix $f \in \mathcal{F}$ and let ϕ_f solve (14). Invariance gives

$$\pi_{\lambda,p}\big(R_{\lambda,p}\phi_f - \phi_f\big) = 0.$$

Divide (17) by λ and integrate:

$$0 = \pi_{\lambda,p}(L\phi_f) + \lambda\pi_{\lambda,p}(A_{\mathrm{E}}\phi_f) + \mathbf{1}_{\{p=1\}}\lambda\pi_{\lambda,p}(A_{\mathrm{S}}\phi_f) + o(\lambda).$$

Since

$$L\phi_f = f - \pi_\beta(f),$$

we have

$$\pi_{\lambda,p}(L\phi_f) = \pi_{\lambda,p}(f) - \pi_\beta(f).$$

By Proposition 2,

$$\pi_{\lambda,p} \Rightarrow \pi_\beta.$$

Assumption A4 gives uniform W-moment control and

$$A_{\mathrm{E}}\phi_f = o(W), \qquad A_{\mathrm{S}}\phi_f = o(W).$$

Hence the coefficient functions are uniformly integrable under $\pi_{\lambda,p}$, and continuity gives

$$\pi_{\lambda,p}(A_{\mathrm{E}}\phi_f) \longrightarrow \pi_\beta(A_{\mathrm{E}}\phi_f),$$

and

$$\pi_{\lambda,p}(A_{\mathrm{S}}\phi_f) \longrightarrow \pi_\beta(A_{\mathrm{S}}\phi_f).$$

Therefore

$$\pi_{\lambda,p}(f) - \pi_\beta(f) = -\lambda\pi_\beta(A_{\mathrm{E}}\phi_f) - \mathbf{1}_{\{p=1\}}\lambda\pi_\beta(A_{\mathrm{S}}\phi_f) + o(\lambda).$$

Defining

$$C_{\mathrm{E}}(f) = -\pi_\beta(A_{\mathrm{E}}\phi_f), \qquad C_{\mathrm{S}}(f) = -\pi_\beta(A_{\mathrm{S}}\phi_f),$$

gives (11). □

Thus, for $p = 2$, stabilization does not alter the leading stationary weak-bias coefficient.

B Experimental Protocols

B.1 Comparator roles

Table 7 summarizes the role of each baseline in the diagnostic benchmark.

Table 7: Comparator methods and their roles in the benchmark.

Method	Gradient	Mom.	Benchmark role
SGLD	mini	no	unstabilized stochastic-gradient control
TUSLA	mini	no	stabilized stochastic-gradient baseline
pSGLD	mini	no	strong adaptive SG-MCMC baseline
Clipped-SGLD	mini	no	componentwise drift-clipping baseline
SGLRW	mini	no	bounded lattice-update SG baseline
PLMC	full	no	projected exact-gradient baseline
mTULA	full	no	exact-gradient tamed baseline
kTULA	full	no	recent exact-gradient tamed baseline
SGHMC	mini	yes	untamed SG momentum baseline
tSGHMC	mini	yes	tamed SG momentum baseline
BAOAB	full	yes	high-accuracy kinetic baseline
MALA	full	no	Metropolis-corrected overdamped baseline
HMC	full	yes	Metropolis-corrected Hamiltonian baseline
NUTS	full	yes	adaptive Hamiltonian baseline

B.2 $d = 50$ EXACT-REFERENCE BENCHMARK

We use the quartic finite-sum target from Section 5, with $m = 64$, $\beta = 10$, and quadratic regularization 0.05. Its product structure allows each reference marginal to be computed accurately by one-dimensional numerical quadrature. Unless stated otherwise, hyperparameters are selected on disjoint tuning targets and frozen before held-out evaluation. Strict stability requires both numerical survival over the full run and recovery from the all-coordinates-equal-to-5 initialization to $\|x\|^2/d < 4$ within the prescribed tail horizon. The number of held-out targets, step-size range, and budget are reported separately for each experiment below.

B.3 UNIFIED OVERDAMPED CORE REPLICATION

The main overdamped comparison uses held-out targets 7203–7212 and $\lambda \in \{0.08, 0.085, 0.09, 0.095\}$ under one common runner. Typical-region chains start from $x = 0$, use physical burn-in 80 and sampling horizon 400, and retain at most 4001 states. Tail runs start from $x = 5\mathbf{1}$, use physical horizon 30, and count as recovered once $\|x\|^2/d < 4$.

The frozen methods are: RELTA with batch size 16, $K = 2$, $p = 2$, and score $[\|x\|^2 - 1.1d]_+$; vanilla SGLD with the same mini-batch oracle and no stabilization; TUSLA with batch size 16 and denominator $1 + \sqrt{\lambda}\|x\|^2$; kTULA with exact gradients and $(a, \epsilon_h) = (0.08, 0.5)$; and mTULA with exact gradients and the published modified taming $h_\lambda(x) = h(x)/\sqrt{1 + \lambda\|x\|^{2r}}$, with $r = 2$ for the quartic target. No held-out result is used to retune any configuration in this replication. The raw per-run outputs underlying the reported medians and stability counts are included in the reproducibility package.

B.4 kTULA TUNING

kTULA is given a broad theorem-compatible grid over its linear-floor parameter a and exponent ϵ_h. The broad sweep places the best configuration satisfying strict stability throughout the overdamped stress band near $\epsilon_h = 0.5$; a one-dimensional refinement freezes $a = 0.08$ before held-out evaluation. The frozen configuration is $10/10$ strictly stable at all four reported step sizes.

B.5 MOMENTUM PROTOCOLS

For the strictly stable SGHMC comparison, we evaluate

$$\lambda \in \{0.070, 0.075, 0.080, 0.085, 0.090, 0.095\}.$$

SGHMC friction is selected on tuning targets 7200–7202 from $\{0.1, 0.3, 1, 3, 10, 30, 100\}$; $\gamma = 10$ is selected at every reported step size. The frozen configurations are then evaluated on the same 20 fresh targets, 8100–8119, for both SGHMC and RELTA, using matched typical-region and far-tail stability protocols.

At the stability boundary $\lambda = 0.12$, SGHMC and tSGHMC friction parameters are selected on ten disjoint tuning targets, 7600–7609. Selection first maximizes strict stability, tail recovery, and typical-region survival, and then minimizes median W_1 among ties. This gives $\gamma = 12$ for SGHMC and $\gamma = 1$ for tSGHMC. These frozen configurations, together with the frozen RELTA configuration, are evaluated on fresh targets 7700–7709. All ten SGHMC typical-region chains survive, while one SGHMC far-tail trajectory diverges numerically; the tSGHMC and RELTA configurations are $10/10$ strictly stable.

B.6 pSGLD AND BAOAB FINITE-BUDGET PROTOCOLS

The pSGLD and BAOAB comparisons use the same held-out targets 8700–8709 as the corresponding RELTA runs. For pSGLD, the mini-batch size is 16 and the RMSProp preconditioner is

$$v_t = \alpha v_{t-1} + (1-\alpha) g_t^2, \qquad G_t = \operatorname{diag}\big((\sqrt{v_t} + \epsilon)^{-1}\big).$$

Disjoint tuning targets 7200–7202 select $\alpha = 0.999$ and $\epsilon = 1.0$, with budget-specific step sizes $h = .05$, .03, and .02 at 2000, 5000, and 10000 FGE, respectively. The main-text numbers use the divergence-corrected update, adding $h\beta^{-1}\Gamma_t$ with $(\Gamma_t)_i = \partial (G_t)_{ii}/\partial x_i$, computed analytically from the same sampled mini-batch. The frozen configurations are $10/10$ strictly stable at all three budgets and attain median marginal W_1 0.011532, 0.009576, and 0.009175. BAOAB uses force reuse, so after initialization one completed step requires one new exact full gradient. The 2000-, 5000-, and 10000-FGE entries use frozen configurations evaluated on the same fresh targets 8700–8709, with no retuning. The 2000-FGE row was rerun under this common protocol to replace an older summary-only result. The frozen BAOAB accuracy-optimal configurations are $(\lambda, \gamma) = (.82, .35)$, $(.78, .175)$, and $(.75, .20)$ at 2000, 5000, and 10000 FGE, respectively; the corresponding stability-constrained configurations are $(.34, .30)$, $(.30, .25)$, and $(.36, .50)$. The frozen RELTA step sizes are .014, .009, and .006.

For all three budget replications, RELTA and pSGLD sample 16 scalar finite-sum components (j, k) from the $dm = 50 \times 64$ component pool per stochastic-gradient step and rescale them to form an unbiased vector estimator. Hence each such step costs $16/(50 \cdot 64)$ FGE. BAOAB is charged one FGE per completed step after initialization. We use 25% burn-in and retain at most 4001 states for the marginal W_1 calculation. Tail runs start from $x = 5\mathbf{1}$ (and $v = 0$ for BAOAB), use physical-time horizon 30, and count as recovered when the final state satisfies $\|x\|^2/d < 4$.

B.7 ADDITIONAL EXPLICIT STABILIZATION BASELINES

We additionally compare RELTA with three explicit stabilization mechanisms on the same quartic targets and 2000/5000/10000-FGE budgets. SGLRW and Clipped-SGLD follow Mensch et al. (2026) and use the same batch-16 scalar-component oracle as RELTA, so each transition costs $16/(50 \cdot 64)$ FGE. In our temperature convention, SGLRW makes coordinatewise jumps $\pm\sqrt{2h/\beta}$ with

$$\Pr(\Delta X_i = +\sqrt{2h/\beta} \mid X) = \frac{1 - a_i}{2}, \qquad a_i = \operatorname{clip}\Big(\sqrt{\beta h/2}\, H_i(X, Z), -1, 1\Big),$$

where the clipping keeps the lattice probabilities valid outside their nominal range. Clipped-SGLD clips each component of the drift increment $hH(X, Z)$ at radius $\sqrt{2h/\beta}$ before adding Gaussian noise. Disjoint tuning targets 7200–7202 select SGLRW steps $h = (.004, .003, .002)$ and Clipped-SGLD steps $h = (.003, .0035, .00275)$ at the three budgets, respectively; frozen configurations are evaluated on targets 8700–8709.

PLMC follows Pang et al. (2025). With cubic drift growth, its projected state is

$$\mathcal{P}_\lambda(x) = \min\{1, R_\lambda/\|x\|\}\, x, \qquad R_\lambda = \vartheta(\beta d/\lambda)^{1/6}.$$

Each PLMC transition uses one exact full gradient and costs one FGE. A single globally frozen $\vartheta = 2.2$ is selected on the tuning targets, with $\lambda = (.34, .29, .30)$ at 2000/5000/10000 FGE. All methods use 25% burn-in, at most 4001 retained states, and the same far-tail recovery criterion as the main finite-budget benchmark.

Table 8: Additional explicit stabilization baselines under matched FGE budgets. Entries are median marginal W_1 over ten held-out targets; all configurations are $10/10$ strictly stable.

Method	2000 FGE	5000 FGE	10000 FGE
RELTA	**0.01071**	**0.00885**	**0.00823**
SGLRW	0.01924	0.01461	0.01264
Clipped-SGLD	0.02050	0.01414	0.01144
PLMC	0.02744	0.01924	0.01547

All three added baselines remain strictly stable throughout, so the comparison isolates finite-budget accuracy rather than numerical failure. RELTA wins all paired held-out comparisons against SGLRW, Clipped-SGLD, and PLMC at every reported budget.

B.8 Corrected MCMC and budget accounting

The corrected MCMC comparison uses held-out targets 7500–7509 and budgets of 5000 and 10000 FGE. RELTA uses batch size 16, $\lambda = 0.04$, $K = 2$, and the localized score $[\|x\|^2 - 1.1d]_+$. For typical-region accuracy, MALA uses $h = 0.08$ and HMC uses $(\epsilon, L) = (0.12, 5)$. Separate stability-oriented tail checks use $h = 0.01$ for MALA and $(\epsilon, L) = (0.04, 10)$ for HMC. NUTS uses NumPyro with diagonal mass-matrix adaptation, 200 warmup iterations, target acceptance 0.8, and maximum tree depth 10; the same adaptive method is run from both initializations.

RELTA is charged by component-gradient evaluations, MALA by one full gradient per proposal, and HMC/NUTS by the number of full-gradient evaluations along the trajectory. NUTS warmup is charged to the FGE budget. Full target-density evaluations are not included in the FGE accounting. We do not equalize the number of retained post-burn-in states in this comparison: the number of valid samples produced under the fixed FGE budget is part of the finite-budget efficiency being measured. The main-text table reports only typical-region W_1 so that accuracy is evaluated under a common initialization; tail-start results are used only for the stability diagnostic.

B.9 Dense correlated-target generalization

The correlated-target experiment starts from the same latent $d = 50$ quartic family as the main benchmark and applies the fixed dense non-orthogonal map

$$T = Q_1 \operatorname{diag}(s_1, \ldots, s_{50}) Q_2^\top, \qquad x = Ty.$$

The matrices Q_1, Q_2 are obtained by QR factorization of independent standard-Gaussian matrices generated with seed 260909. Before normalization, the singular values are geometrically spaced from $1/\sqrt{6}$ to $\sqrt{6}$; they are then rescaled so that $50^{-1} \sum_i s_i^2 = 1$. Thus $\kappa(T) = 6$. The transform is used to define the observed target and to generate exact references; no sampler receives T or the latent coordinates as preconditioning information.

The budget is 2000 FGE, and targets 7200–7202 are used exclusively for tuning. To equalize the primary tuning budget, RELTA and pSGLD each receive 24 coarse configurations. For RELTA, $K \in \{8, 16, 32\}$ and $h \in \{.008, .010, .012, .014, .016, .018, .020, .022\}$; for pSGLD, $\alpha \in \{.99, .999\}$, $\epsilon \in \{.3, 1, 3\}$, and $h \in \{.014, .018, .022, .026\}$. Every configuration uses mini-batch size 16, the same 2000-FGE budget, and exactly 1500 retained states for the tuning metric. Eligibility requires strict stability on all three tuning targets; among eligible configurations, selection minimizes median sliced W_1, with mean sliced W_1 as a deterministic tie-breaker.

Each method then receives the same five-point local-h refinement with the other hyperparameters frozen: for RELTA, $K = 16$ and $h \in \{.012, .013, .014, .015, .016\}$; for pSGLD, $(\alpha, \epsilon) = (.99, 1)$ and $h \in \{.018, .020, .022, .024, .026\}$. Because the selected pSGLD decay parameter lies on the lower coarse-grid boundary, we additionally check $\alpha \in \{.90, .95, .98, .985\}$ with $\epsilon = 1$ over the same five step sizes, using only the tuning targets and the same selection rule; this check does not change the selected configuration. The resulting frozen settings are $(K, h) = (16, .014)$ for RELTA and $(\alpha, \epsilon, h) = (.99, 1, .022)$ for divergence-corrected pSGLD.

All reported results use previously unused targets 11000–11009. For each target, the reference is generated from 12001 independent latent samples obtained from the numerical one-dimensional

quantiles and then mapped by T. Sliced empirical W_1 is averaged over 42 fixed unit directions: 12 coordinate directions, the three largest and three smallest left-singular directions of T, and 24 Gaussian random directions generated with seed 20260909. All methods are compared using 1500 post-burn-in samples; longer RELTA chains are uniformly subsampled to this common size. Strict stability uses the same typical-survival and $x_0 = 5\mathbf{1}$ tail-recovery criteria as the controlled benchmark, with physical tail horizon 30.

On these targets, RELTA and pSGLD are both $10/10$ strictly stable; median sliced W_1 is 0.02016 for RELTA and 0.02089 for pSGLD, and RELTA wins $7/10$ paired comparisons. Accuracy-tuned BAOAB uses $(\gamma, h) = (.2, .38)$ and is $10/10$ typical but $0/10$ on tail recovery, with median sliced W_1 0.02654. Stability-constrained BAOAB uses $(\gamma, h) = (1, .025)$, is $10/10$ strictly stable, and has median sliced W_1 0.14587.

B.10 Real-data Bayesian neural regression

We use the sklearn Diabetes data with a fixed 309/66/67 train/validation/test split. The first split uses seed 260909 and the validation/test split uses seed 260910. Features and targets are standardized using training statistics only. The Bayesian model is

$$f_\theta(x) = a^\top \mathrm{ReLU}(Wx + b) + c,$$

with $W \in \mathbb{R}^{32\times 10}$, $b, a \in \mathbb{R}^{32}$, and $c \in \mathbb{R}$, for 385 parameters in total. The observation noise is Gaussian with standard deviation 0.55 on the standardized target and all parameters have independent $N(0, 1)$ priors. The sampler uses the per-datum potential

$$U(\theta) = \frac{1}{N}\sum_{i=1}^{N} \frac{(y_i - f_\theta(x_i))^2}{2(0.55)^2} + \frac{\|\theta\|^2}{2N}, \qquad \beta = N = 309,$$

so that $\exp(-\beta U)$ is the Bayesian posterior. SGLD, RELTA, and SGHMC all use mini-batches of size 32 and a 2000-FGE budget. Chains discard the first 25% as burn-in and retain at most 2500 samples; predictive NLL is computed from at most 1200 evenly spaced post-burn-in samples.

The MAP point is obtained by a fixed Adam procedure (6000 steps, learning rate .005, seed 11), giving $\|\theta_{\mathrm{MAP}}\|^2 = 79.66$. For RELTA we use the score $[\|\theta\|^2 - s_0]_+$ with $p = 2$. Three training-only pilot SGLD chains of 300 FGE each give $q_{0.99}^{\mathrm{pilot}} = 88.96$ and determine

$$s_0 = \max\left\{q_{0.99}^{\mathrm{pilot}},\ 1.25d,\ 2\|\theta_{\mathrm{MAP}}\|^2\right\} = 481.25, \qquad d = 385.$$

These pilot chains are a one-time anchor calibration and are separate from the 2000-FGE budget of each evaluated chain. The off-posterior stress diagnostic starts from $\theta_0 = 4\theta_{\mathrm{MAP}}$ and runs for fixed physical time $T = 20$, using $\lceil 20/h \rceil$ iterations. A stress success requires numerical boundedness and terminal contraction, $\|\theta_T\|^2 < \|\theta_0\|^2$.

A preliminary width-scaling development run showed that the initial RELTA step-size range was poorly scaled for the 385-parameter model. Before opening any fresh test-chain seed, we therefore reset only the validation-stage SGLD/RELTA grid to the common low-step range

$$h \in \{.006, .008, .010, .012, .016, .020, .024, .032, .040\},$$

with $K = 2$ fixed for RELTA. Selection uses tuning seeds 7200–7202 only. Stability-oriented candidates must survive all three typical tuning chains and all three stress runs, after which median validation NLL is minimized; accuracy-oriented candidates require only three-of-three typical survival and minimize the same validation metric. No final-seed result is used in this development-stage grid adjustment or selection. The broader SGHMC development sweep is reused unchanged:

$$\begin{aligned}
&\text{block 1: } h \in \{.04, .05, .06, .07, .08, .09, .10\}, \quad \gamma \in \{2, 5\},\\
&\text{block 2: } h \in \{.03, .04, .05, .06, .07, .08, .09, .10, .12, .14, .16\},\\
&\qquad\qquad \gamma \in \{0.5, 1, 1.5, 3, 7.5, 10\}.
\end{aligned}$$

The frozen configurations are RELTA $(h, K) = (.020, 2)$, matched-h SGLD $h = .020$, stability-constrained SGHMC $(h, \gamma) = (.03, 3)$, and accuracy-oriented SGHMC $(h, \gamma) = (.04, .5)$. No SGLD

point in the adjusted grid succeeds on the stress diagnostic, so there is no strictly stable SGLD configuration to report.

All reported test results use fresh chain seeds 9800–9809, opened only after the above configurations were frozen. RELTA is $10/10$ typical and $10/10$ stress-stable with median test NLL 0.97812. Matched-h SGLD is $10/10$ typical and $0/10$ on stress with the identical median test NLL 0.97812. Stability-constrained SGHMC is $10/10$ typical and $10/10$ stress-stable with median test NLL 0.98197, while accuracy-oriented SGHMC is $10/10$ typical and $0/10$ on stress with median test NLL 0.98125.

At the frozen RELTA point the stabilizer never binds on any of the ten typical test trajectories. Because the matched-h SGLD control uses the same seeds, mini-batch streams, and Gaussian-noise streams, the ordinary-chain outputs are identical seed by seed, including predictive NLL, RMSE, predictive spread, maximum parameter norm, and retained sample count. Under the common stress initialization, RELTA succeeds in all ten runs whereas SGLD succeeds in none; the median terminal squared norm is approximately 1.09×10^3 for RELTA versus 1.87×10^{15} for SGLD, from a common initial squared norm of approximately 1.275×10^3. Thus the real-data experiment isolates stabilization against off-posterior excursions without attributing a typical-region predictive gain to an inactive stabilizer.

B.11 Sampling-order ablation

The $p = 1, 2, 4$ sweep uses fresh targets 7303–7312 and $\lambda \in \{0.02, 0.04, 0.06, 0.08, 0.09, 0.095\}$. Within each target and step size, all three values of p share the same stochastic-gradient oracle, chain random seed, burn-in, sampling horizon, and tail-recovery horizon. The reported $60/60$ wins for $p = 2$ over $p = 1$ and $59/60$ wins for $p = 4$ over $p = 2$ are computed entirely within this matched three-way sweep. All 180 runs satisfy strict stability.