Understanding Solver-Induced Variance Distortion in Conditional Diffusion Regression

Jaemin Song
cantor@jnu.ac.kr
Jaegi Jeon
jaegijeon@jnu.ac.kr
Graduate School of Data Science, Chonnam National University
Gwangju, Republic of Korea

Abstract

Conditional diffusion regression quantifies predictive uncertainty through samples $y \sim \hat{p}_\theta(\cdot \mid x)$ generated by finite-step reverse-time sampling. For a fixed learned score field s_θ, reverse-time grid, and noise schedule, we study how the finite-step solver affects input-dependent conditional variance. For stochastic reverse-time SDE solvers, we prove an early-stopped total variation (TV) bound. For scalar responses, we introduce a solver-matched oracle reference and derive a learned–oracle variance-gap recursion separating transport, injection, and nonlinear-remainder effects, with the probability-flow ODE (PF-ODE) entering the recursion as a deterministic zero-injection case. Controlled heteroskedastic experiments with known data variance and oracle scores show that changing only the solver materially changes variance fidelity and interval calibration.

1. Introduction

In conditional prediction, uncertainty must vary correctly with the input, not only the predictive mean. This is central in heteroskedastic regression, where $x \mapsto \mathrm{Var}(Y \mid X = x)$ is part of the target structure (White, 1980). Classical and modern methods model this input-dependent scale through heteroskedastic models or variance networks (Le et al., 2005; Cao et al., 2021; Skafte et al., 2019). Bayesian and ensemble methods instead use posterior or ensemble variability (Gal & Ghahramani, 2016; Lakshminarayanan et al., 2017). In both cases, uncertainty is usually reported only at the endpoint, through predictive variances, intervals, or calibration scores. These summaries do not show how the inference procedure creates, transports, or distorts conditional variance.

Conditional diffusion regression represents the predictive law through samples. For a fixed input x, one draws

$$y^{(1)}, \ldots, y^{(M)} \sim \hat{p}_\theta(\cdot \mid x),$$

and estimates moments, intervals, and calibration statistics from the empirical distribution. This view underlies conditional diffusion methods such as CARD (Han et al., 2022) and diffusion-ensemble approaches to zero-shot uncertainty quantification (Shu & Farimani, 2024). It also makes the reverse-time solver part of the uncertainty-estimation pipeline. Even with the learned score field, reverse-time grid, and noise schedule fixed, changing the finite-step solver can change the sampled conditional law.

Existing theory mainly explains global distributional accuracy. Non-asymptotic analyses control score approximation, initialization mismatch, and time discretization in KL, total variation (TV), or Wasserstein distance (Lee et al., 2022; Chen et al., 2023a;b; Benton et al., 2023; Li et al., 2023). Recent work extends these guarantees to broader target assumptions and faster finite-step solvers (Liang et al., 2024). Conditional diffusion regression theory is also emerging, including CARD-type non-asymptotic bounds (Li, 2025) and statistical guarantees for uncertainty quantification (Jiao et al., 2025). Complementary supervised diffusion methods, including Treeffuser, diffusion–boosting hybrids, supervised score-based boosting, and diffusion-based probabilistic regression, focus on improving endpoint uncertainty estimates (Beltran Velez et al., 2024; Han & Zhou, 2024; Zhao et al., 2025; Kneissl et al., 2025). A separate line of

Accepted to FoGen 2026: Foundations of Deep Generative Models: Understanding Memorization, Generalization, and Reasoning, an ICML 2026 workshop (non-archival).

work highlights reverse-transition and sampling choices, including learned reverse-process variance, richer reverse-time posteriors, and empirical control of ensemble dispersion through the reverse-step budget (Nichol & Dhariwal, 2021; De Bortoli et al., 2025; Merizzi et al., 2026). Together, these results suggest that finite-step inference choices can affect uncertainty. However, they do not isolate which reverse transitions amplify, attenuate, inject, or distort input-dependent conditional variance when only the solver is changed.

We therefore study a fixed-score, fixed-grid setting. The learned score field, reverse-time grid, and noise schedule are held fixed, and only the finite-step reverse solver is varied. This design isolates solver-induced variance distortion. Our contributions are threefold. First, we adapt non-asymptotic stochastic-solver TV analysis to conditional diffusion regression, obtaining an early-stopped distribution-level baseline for Euler–Maruyama (EM) and an exponential integrator (ExpInt). Second, for scalar responses, we derive a solver-matched learned–oracle variance-gap recursion for EM, ExpInt, and a locally linearized probability-flow ODE (PF-ODE). The recursion separates transport, injection, and nonlinear-remainder effects, and treats PF-ODE as the deterministic zero-injection case. Third, we use controlled heteroskedastic tasks with known data variance and oracle scores to connect endpoint variance errors with learned–oracle gap paths and unrolled mechanism summaries.

2. Problem Setup and Variance Propagation

2.1. Conditional Diffusion Regression Setup

Let $(X, Y_0) \sim P_{\mathrm{data}}$, where X takes values in $\mathcal{X} \subseteq \mathbb{R}^{d_X}$ and $Y_0 \in \mathbb{R}^{d_Y}$. For each fixed $x \in \mathcal{X}$, let $p_0(\cdot \mid x)$ denote the conditional law of Y_0 given $X = x$.

Forward process. We adopt an Ornstein–Uhlenbeck (OU) noising process in the response space,

$$\mathrm{d}Y_t = -\tfrac{1}{2} Y_t \,\mathrm{d}t + \mathrm{d}W_t, \qquad Y_0 \sim p_0(\cdot \mid x), \qquad t \in [0, T]. \tag{1}$$

This choice is convenient because the OU transition density is available in closed form, so forward perturbations used for score learning can be simulated exactly (Chen et al., 2023b; Benton et al., 2023). In particular,

$$p_{t|0}(y \mid y_0) = \mathcal{N}\big(y;\, e^{-t/2} y_0,\, (1 - e^{-t}) I_{d_Y}\big),$$

and the OU-smoothed conditional marginal is the Gaussian convolution

$$p_t(y \mid x) = \int_{\mathbb{R}^{d_Y}} p_{t|0}(y \mid y_0)\, p_0(y_0 \mid x)\, \mathrm{d}y_0.$$

For clarity of exposition, we assume the response has been standardized so that the Gaussian reference is $\mathcal{N}(0, I_{d_Y})$. The OU semigroup Gaussianizes the conditional law and contracts the KL divergence to the reference,

$$\mathrm{KL}\Big(p_t(\cdot \mid x) \,\Big\|\, \mathcal{N}(0, I_{d_Y})\Big) \le e^{-t}\, \mathrm{KL}\Big(p_0(\cdot \mid x) \,\Big\|\, \mathcal{N}(0, I_{d_Y})\Big),$$

under mild regularity conditions (Bakry et al., 2013; Chen et al., 2023a). We take T sufficiently large and initialize reverse-time sampling from $Y_T \sim \mathcal{N}(0, I_{d_Y})$.

Reverse process and early stopping. Let $p_t(\cdot \mid x)$ denote the OU-smoothed conditional law (1). Diffusion time reversal characterizes the reversed trajectory $\bar{Y}_{t'} := Y_{T-t'}$, whose drift depends on the conditional score (Anderson, 1982). For sampling, we work directly in the original time variable t and integrate from $t = T$ down to $t = \delta$. The reverse-time dynamics can be written as

$$\mathrm{d}Y_t = \Big(-\tfrac{1}{2} Y_t - s^\star(t, Y_t, x) \Big) \mathrm{d}t + \mathrm{d}W_t, \qquad t : T \to \delta, \tag{2}$$

where W_t is a standard Brownian motion in the reverse-time parametrization and

$$s^\star(t, y, x) := \nabla_y \log p_t(y \mid x) \tag{3}$$

is the conditional score of the OU-smoothed law. We target the early-stopped law $p_\delta(\cdot \mid x)$ for a small $\delta > 0$ to avoid potential instability of the score field near $t \downarrow 0$ (Song et al., 2020; Benton et al., 2023).

Score matching. Simulating (2) requires access to the conditional score $s^\star(t, y, x)$ in the reverse drift. Since $p_t(\cdot \mid x)$ is unknown, we learn a neural score model $s_\theta(t, y, x)$ via score matching (Hyvärinen, 2005; Vincent, 2011). A canonical population objective on $[\delta, T]$ is the time-averaged L^2 score error

$$\mathcal{L}(s_\theta) := \frac{1}{T-\delta}\int_\delta^T \mathbb{E}\Big[\big\|s_\theta(t, Y_t, X) - s^\star(t, Y_t, X)\big\|_2^2\Big]\mathrm{d}t, \qquad Y_t \sim p_t(\cdot \mid X).$$

Although $\mathcal{L}$ is not directly estimable, denoising score matching provides a tractable denoising objective with the same population minimizer, by sampling from analytically simulable OU perturbations. Using $Y_t = e^{-t/2}Y_0 + \sigma_t Z$ with $\sigma_t^2 = 1 - e^{-t}$ and $Z \sim \mathcal{N}(0, I_{d_Y})$, we minimize

$$\min_\theta \ \mathbb{E}_{t\sim\mathrm{Unif}[\delta,T]}\,\mathbb{E}_{(X,Y_0)\sim P_{\mathrm{data}}}\,\mathbb{E}_{Z\sim\mathcal{N}(0,I_{d_Y})}\Big\|s_\theta\big(t,\, e^{-t/2}Y_0 + \sigma_t Z,\, X\big) + \tfrac{1}{\sigma_t}Z\Big\|_2^2.$$

Finite-step reverse-time solvers. Given s_θ, conditional generation discretizes the reverse-time dynamics on the grid (4) and integrates backward from $t_K = T$ to $t_0 = \delta$.

$$\delta = t_0 < t_1 < \cdots < t_K = T, \qquad h_k := t_k - t_{k-1} > 0, \tag{4}$$

where each update maps $Y_{t_k} \mapsto Y_{t_{k-1}}$.

We consider two stochastic reverse-SDE solvers and one deterministic PF-ODE solver. The stochastic solvers are Euler–Maruyama (EM) and an exponential integrator (ExpInt), both of which can be written in the affine form

$$Y_{t_{k-1}} = A_k Y_{t_k} + B_k s_\theta(t_k, Y_{t_k}, x) + \sigma_k \xi_k, \qquad \xi_k \sim \mathcal{N}(0, I_{d_Y}). \tag{5}$$

For EM, the coefficients come from the direct discretization of (2). For ExpInt, they exploit the linear OU structure (Zhang & Chen, 2022; De Bortoli, 2022; Chen et al., 2023a; Benton et al., 2023):

$$(A_k, B_k, \sigma_k^2) = \begin{cases} (1 + h_k/2,\ h_k,\ h_k), & \text{EM}, \\ (e^{h_k/2},\ 2(e^{h_k/2} - 1),\ e^{h_k} - 1), & \text{ExpInt}. \end{cases}$$

As a deterministic counterpart, we also consider the probability-flow ordinary differential equation (PF-ODE) associated with the same score field (Song et al., 2020; 2021),

$$\frac{\mathrm{d}}{\mathrm{d}t}Y_t = -\tfrac{1}{2}Y_t - \tfrac{1}{2}s_\theta(t, Y_t, x), \qquad t : T \to \delta,$$

and integrate it by Heun's deterministic update. With $f(t, y, x) := -\frac{1}{2}y - \frac{1}{2}s_\theta(t, y, x)$,

$$\widetilde{Y} = Y_{t_k} - h_k\, f(t_k, Y_{t_k}, x), \qquad Y_{t_{k-1}} = Y_{t_k} - \frac{h_k}{2}\Big(f(t_k, Y_{t_k}, x) + f(t_{k-1}, \widetilde{Y}, x)\Big).$$

Unlike the reverse-SDE solvers, PF-ODE introduces no explicit per-step noise injection.

2.2. Variance Targets and Oracle Reference

We study conditional uncertainty under finite-step sampling by tracking intermediate laws and moments along the reverse chain. On the time grid (4), let $q_{t_k}(\cdot \mid x)$ denote the conditional law of Y_{t_k} produced by the K-step solver driven by s_θ. For each $x \in \mathcal{X}$ and grid time t_k, define

$$m_{t_k}(x) := \mathbb{E}\big[Y_{t_k} \mid X = x\big], \qquad V_{t_k}(x) := \mathrm{Var}\big(Y_{t_k} \mid X = x\big).$$

For scalar targets ($d_Y = 1$), we write $V_{t_k}(x) := v_{t_k}(x)$.

To separate solver effects from score-learning effects, we define, for each solver S, a solver-matched oracle chain: the same discretization as the learned chain, but driven by the true score $s^\star$. In the scalar case, let $v_{S,t_k}^{\mathrm{learn}}(x)$ and $v_{S,t_k}^{\star}(x)$

denote the learned and solver-matched oracle conditional variances, and let $v^{\mathrm{data}}(x) := \mathrm{Var}(Y_0 \mid X = x)$ be the conditional data variance.

$$v^{\mathrm{learn}}_{S,t_0}(x) - v^{\mathrm{data}}(x) = \big\{v^{\mathrm{learn}}_{S,t_0}(x) - v^{\star}_{S,t_0}(x)\big\} + \big\{v^{\star}_{S,t_0}(x) - v^{\mathrm{data}}(x)\big\}. \tag{6}$$

The first term is the solver-matched learned–oracle variance gap analyzed below. The second term is the solver-matched oracle endpoint bias, which is measured separately in the experiments. We write $G_S(t_k, x) := v^{\mathrm{learn}}_{S,t_k}(x) - v^{\star}_{S,t_k}(x)$ for the learned–oracle variance-gap path. In fixed-solver derivations, we omit the solver index S.

Proposition 2.1 (Stepwise mean–variance propagation)**.** *Assume $d_Y = 1$ and fix $x \in \mathcal{X}$. For EM and ExpInt, use the affine stochastic update (5). This unified form covers EM and ExpInt through solver-specific coefficients (A_k, B_k, σ_k^2). Define the local Taylor remainder*

$$\rho_k := s_\theta(t_k, Y_{t_k}, x) - s_\theta\big(t_k, m_{t_k}(x), x\big) - s_y\big(t_k, m_{t_k}(x), x\big)\big(Y_{t_k} - m_{t_k}(x)\big),$$

where $s_y(t, y, x) := \partial_y s_\theta(t, y, x)$. Define the local transport factor

$$J_k(x) := A_k + B_k\, s_y\big(t_k, m_{t_k}(x), x\big).$$

Then the conditional mean and variance satisfy

$$m_{t_{k-1}}(x) = A_k\, m_{t_k}(x) + B_k\, s_\theta\big(t_k, m_{t_k}(x), x\big) + B_k\, \mathbb{E}[\rho_k \mid X = x], \tag{7}$$

and

$$v_{t_{k-1}}(x) = J_k(x)^2\, v_{t_k}(x) + \sigma_k^2 + r_k(x), \tag{8}$$

where

$$r_k(x) := B_k^2\, \mathrm{Var}(\rho_k \mid X = x) + 2J_k(x)B_k\, \mathrm{Cov}\big(Y_{t_k} - m_{t_k}(x),\, \rho_k \mid X = x\big). \tag{9}$$

For PF-ODE with Heun, write the deterministic one-step map as

$$Y_{t_{k-1}} = \Phi^{\theta}_{k,\mathrm{PF}}(Y_{t_k}; x).$$

Define

$$J_{k,\mathrm{PF}}(x) := \partial_y \Phi^{\theta}_{k,\mathrm{PF}}(m_{t_k}(x); x)$$

and

$$\zeta_k := \Phi^{\theta}_{k,\mathrm{PF}}(Y_{t_k}; x) - \Phi^{\theta}_{k,\mathrm{PF}}(m_{t_k}(x); x) - J_{k,\mathrm{PF}}(x)(Y_{t_k} - m_{t_k}(x)).$$

Then PF-ODE satisfies the same template with $\sigma_k^2 = 0$, $J_k = J_{k,\mathrm{PF}}$, and

$$r_{k,\mathrm{PF}}(x) = \mathrm{Var}(\zeta_k \mid X = x) + 2J_{k,\mathrm{PF}}(x)\, \mathrm{Cov}(Y_{t_k} - m_{t_k}(x), \zeta_k \mid X = x).$$

The oracle recursions are obtained by replacing s_θ and $\Phi^{\theta}_{k,\mathrm{PF}}$ with $s^\star$ and $\Phi^{\star}_{k,\mathrm{PF}}$. Full proofs are given in Appendices B.1 and B.4.

3. Theoretical Analysis

This section presents two finite-step results with different roles. Proposition 3.1 provides a distribution-level baseline for stochastic reverse-SDE solvers at the early-stopping time $t = \delta$. It controls the learned law $q_\delta(\cdot \mid x)$ relative to the positive-time OU target $p_\delta(\cdot \mid x)$, rather than directly to the original data law $p_0(\cdot \mid x)$. This result applies to EM and ExpInt because their learned and oracle one-step kernels are nondegenerate Gaussians with common covariance. It does not apply to deterministic PF-ODE.

Proposition 3.2 gives a scalar conditional-moment decomposition. It subtracts and unrolls the learned and solver-matched oracle one-step variance recursions from Proposition 2.1. This variance-gap analysis covers EM, ExpInt, and locally linearized PF-ODE, with PF-ODE treated as the deterministic zero-injection case.

3.1. Early-Stopped Total Variation Control

Throughout this subsection, fix a stochastic reverse-SDE solver family $S \in \{\mathrm{EM}, \mathrm{ExpInt}\}$. The learned and oracle laws $q_t(\cdot \mid x)$ and $q_t^\star(\cdot \mid x)$, and the coefficients A_k, B_k, σ_k, are those of this fixed solver. We suppress the solver index in the notation.

We first record a baseline distribution-level guarantee at the early-stopping time $t = \delta$. For covariates X drawn from the data distribution, we measure convergence through the mean conditional TV distance

$$\mathcal{E}_{\mathrm{TV}}(\delta) = \mathbb{E}_X\Big[\,\mathrm{TV}\left(q_\delta(\cdot \mid X), p_\delta(\cdot \mid X)\right)\Big].$$

Proposition 3.1 (Early-stopped TV control for stochastic reverse-SDE solvers)**.** *Under the assumptions stated in Appendix A, including the stochastic Gaussian solver scope, the admissible reverse-grid condition, and $R_{p\to q} < \infty$, there exists a constant $C > 0$ such that*

$$\mathcal{E}_{\mathrm{TV}}(\delta) \le C\,\varepsilon_s + R_{p\to q} + C\sqrt{\Sigma_K} + C\sqrt{d_Y + M_2^{\max}}\,e^{-T/2}. \tag{10}$$

The four terms in (10) are, respectively, the solver-weighted score error, the forward-to-reverse transfer residual, the oracle discretization error, and Gaussian initialization mismatch. We now define these quantities for the fixed stochastic solver.

Define

$$\Delta_k(y, x) = \|s_\theta(t_k, y, x) - s^\star(t_k, y, x)\|_2^2.$$

The solver-weighted forward-marginal score error is

$$\varepsilon_s^2 = \mathbb{E}_X\left[\sum_{k=1}^{K} \|B_k\|_F^2 \sigma_k^{-2}\, \mathbb{E}_{Y\sim p_{t_k}(\cdot\mid X)}\big[\Delta_k(Y, X)\big]\right].$$

For the fixed stochastic reverse-SDE solver $S \in \{\mathrm{EM}, \mathrm{ExpInt}\}$, let

$$\vartheta_t^2 = 1 - e^{-t}, \qquad \Gamma_K = \sum_{k=1}^{K} h_k^2 \vartheta_{t_{k-1}}^{-4}.$$

The solver-specific discretization quantity is

$$\Sigma_K = \begin{cases} d_Y^2 \Gamma_K + M_2^{\max} \displaystyle\sum_{k=1}^{K} h_k^3, & S = \mathrm{EM}, \\ d_Y^2 \Gamma_K, & S = \mathrm{ExpInt}. \end{cases}$$

Finally,

$$M_2^{\max} = \sup_{x\in\mathcal{X}} \mathbb{E}[\|Y_0\|_2^2 \mid X = x],$$

and the forward-to-reverse transfer residual is

$$R_{p\to q}^2 = \mathbb{E}_X\left[\left(\frac{1}{2}\sum_{k=1}^{K} \|B_k\|_F^2 \sigma_k^{-2}\Big(\mathbb{E}_{q_{t_k}(\cdot\mid X)}[\Delta_k(Y, X)] - \mathbb{E}_{p_{t_k}(\cdot\mid X)}[\Delta_k(Y, X)]\Big)\right)_+\right].$$

3.2. Solver-Matched Variance-Gap Decomposition

We analyze the learned–oracle variance-gap path $G_S(t, x)$ defined in Section 2.2. Fix a solver S and suppress its index:

$$v_{t_k}(x) := v_{S,t_k}^{\mathrm{learn}}(x), \qquad v_{t_k}^\star(x) := v_{S,t_k}^\star(x).$$

The quantities $J_k(x)$, σ_k^2, and $r_k(x)$ below also refer to this fixed solver. The typical-set construction in Appendix B.2 is used only to control local transport quantities. It does not define event-restricted conditional variances. The proposition below subtracts and unrolls the learned and solver-matched oracle one-step variance recursions.

Proposition 3.2 (Solver-matched variance-gap decomposition). *Assume $d_Y = 1$. Fix $x \in \mathcal{X}$ and a grid $\delta = t_0 < t_1 < \cdots < t_K = T$. Let $v_{t_k}(x)$ and $v^\star_{t_k}(x)$ be the learned and solver-matched oracle conditional variances at grid time t_k. Under the local regularity and finite-moment conditions stated in Appendix B, suppose that the learned and solver-matched oracle conditional variances satisfy, for each $k = 1, \ldots, K$,*

$$\begin{aligned} v_{t_{k-1}}(x) &= J_k(x)^2\, v_{t_k}(x) + \sigma_k^2 + r_k(x), \\ v^\star_{t_{k-1}}(x) &= \big(J^\star_k(x)\big)^2\, v^\star_{t_k}(x) + (\sigma^\star_k)^2 + r^\star_k(x), \end{aligned} \tag{11}$$

where $J_k(x), J^\star_k(x) \in \mathbb{R}$ are local transport factors, $\sigma_k^2, (\sigma^\star_k)^2 \geq 0$ are per-step injected variances, and $r_k(x), r^\star_k(x)$ collect higher-order effects from the local linearization, including remainder-variance and covariance contributions.

Define the learned-chain transport products

$$\Pi_{1:0}(x) := 1, \qquad \Pi_{1:m}(x) := \prod_{j=1}^{m} J_j(x)^2, \qquad m \geq 1.$$

Then, for this fixed x, the following deterministic conditional-moment bound holds:

$$\begin{aligned} \big|v_{t_0}(x) - v^\star_{t_0}(x)\big| \leq{}& \Pi_{1:K}(x)\,\big|v_{t_K}(x) - v^\star_{t_K}(x)\big| + \sum_{k=1}^{K} \Pi_{1:k-1}(x)\,\big|J_k(x)^2 - \big(J^\star_k(x)\big)^2\big|\, v^\star_{t_k}(x) \\ &+ \sum_{k=1}^{K} \Pi_{1:k-1}(x)\,\big|\sigma_k^2 - (\sigma^\star_k)^2\big| + \sum_{k=1}^{K} \Pi_{1:k-1}(x)\,\big|r_k(x) - r^\star_k(x)\big|. \end{aligned} \tag{12}$$

In particular, if the variance initialization at $t = T$ matches (i.e., $v_{t_K}(x) = v^\star_{t_K}(x)$), the first term vanishes.

Subtracting the two recursions in (11) gives the one-step gap identity, and repeated substitution gives (12). Details are given in Appendices B.4 and B.5. The bound in (12) is a mechanism decomposition for the learned–oracle variance gap, not a convergence bound to v^{data}.

For the pathwise mechanism summaries used in Section 4, define the signed local defects

$$\begin{aligned} D_k^{\mathrm{trans}}(x) &:= \{J_k(x)^2 - (J^\star_k(x))^2\}\, v^\star_{t_k}(x), \\ D_k^{\mathrm{inj}}(x) &:= \sigma_k^2 - (\sigma^\star_k)^2, \\ D_k^{\mathrm{rem}}(x) &:= r_k(x) - r^\star_k(x). \end{aligned}$$

These three terms correspond to the three local sources in the decomposition. The transport defect measures mismatch in local variance multipliers and is scaled by the oracle variance $v^\star_{t_k}(x)$. The injection defect measures differences in explicit per-step noise; it is zero in our matched EM and ExpInt comparisons, and also zero for PF-ODE because PF-ODE has no additive noise. The remainder defect collects nonlinear local-map effects not captured by the first-order transport factor.

In empirical summaries, we also record the residual defect

$$\begin{aligned} D_k^{\mathrm{res}}(x) := \big(v_{t_{k-1}}(x) - v^\star_{t_{k-1}}(x)\big) - J_k(x)^2\big(v_{t_k}(x) - v^\star_{t_k}(x)\big) \\ - D_k^{\mathrm{trans}}(x) - D_k^{\mathrm{inj}}(x) - D_k^{\mathrm{rem}}(x). \end{aligned}$$

It is zero under the exact population recursion and records Monte Carlo, recording, or diagnostic approximation error. For each component $a \in \{\mathrm{trans}, \mathrm{inj}, \mathrm{rem}, \mathrm{res}\}$, define the transported contribution path

$$C_K^a(x) = 0, \qquad C_{k-1}^a(x) = J_k(x)^2 C_k^a(x) + D_k^a(x).$$

These contribution paths are oracle-aided mechanism summaries for the learned–oracle gap path $G_S(t, x)$.

4. Results

This section first reports endpoint predictive and variance summaries, then uses the oracle-aided path quantities defined above to explain the resulting solver-dependent variance distortion.

Table 1. **Main predictive summary on the synthetic heteroskedastic tasks, grouped by model.** RMSE_y is computed from predictive means, and QICE measures multi-level interval calibration. Boldface indicates the best value within each dataset and metric.

	Xsinx–Hetero		Sine–Hetero		Curve–Hetero		Cubic–Step–Hetero	
Model	**RMSE**	**QICE**	**RMSE**	**QICE**	**RMSE**	**QICE**	**RMSE**	**QICE**
EM	1.1420	**0.00790**	1.4561	0.00865	0.5246	**0.00664**	6.0018	0.00644
ExpInt	1.1420	0.00795	1.4570	0.00939	0.5248	0.00687	6.0078	0.00716
PF-ODE	1.3122	0.00962	1.5816	0.01403	0.5981	0.01625	6.4252	0.01823
MC Dropout[†]	1.1424	0.01248	1.4534	0.01099	0.5240	0.00966	5.9948	0.00738
DeepHetero[‡]	**1.1371**	0.01031	**1.4516**	**0.00813**	**0.5229**	0.00895	**5.9911**	**0.00577**

[†] MC Dropout follows the approximate Bayesian dropout baseline of Gal & Ghahramani (2016). [‡] DeepHetero denotes our Gaussian mean–variance MLP baseline, following classical mean–variance estimation (Nix & Weigend, 1994) and variance-network regression (Skafte et al., 2019).

4.1. Experiment Setup and Predictive Summary

Experimental setting. We evaluate four controlled scalar heteroskedastic tasks $Y = \mu(X) + \sigma(X)\epsilon$, $\epsilon \sim \mathcal{N}(0,1)$, which provide both $v^{\mathrm{data}}(x) = \sigma^2(x)$ and a closed-form OU-smoothed oracle score for endpoint variance evaluation and solver-matched oracle chains. Within each dataset and run, the trained score network, noise schedule, reverse grid, and Monte Carlo budget are fixed across solvers. Predictive moments are reported on the original target scale; full implementation and mechanism-summary details are given in Appendix D.1–D.6.

Main predictive summary. Table 1 reports predictive-mean RMSE and interval calibration (QICE). DeepHetero is a well-matched baseline because the tasks are conditionally Gaussian and it is trained with an explicit Gaussian mean–variance likelihood. EM and ExpInt remain close to the best RMSE values and achieve the lowest QICE on two of the four tasks, despite using a sample-based score model with no explicit variance head. PF-ODE is less competitive on both RMSE and QICE. These endpoint summaries motivate the variance-profile and pathwise analyses below.

4.2. Solver-Dependent Variance Fidelity

Figure 1 shows that solver choice changes the endpoint conditional-variance profile even when the learned score is fixed. EM and ExpInt track $v^{\mathrm{data}}(x)$ closely across most regions, while PF-ODE can substantially overstate variance in high-variance regions.

Table 2 assesses how strongly the oracle-aided path summaries based on $G_S(t,x)$ are associated with endpoint variance MSE. Across the 60 dataset–solver–run points, the oracle-gap mean is the most strongly associated summary. This pattern is not merely an artifact of using a solver-matched oracle rather than v^{data} as the target. Appendix D.9.2, Table 10, decomposes the final endpoint error into the learned–oracle gap, oracle endpoint bias, and their cross term. Across solvers and datasets, the learned–oracle endpoint gap is the dominant component, while the solver-matched oracle bias is typically much smaller.

4.3. Oracle-Gap Path Summaries and Unrolled Mechanism Analysis

Pathwise-unrolled mechanism summary. Figure 2 and Table 3 summarize the learned–oracle variance-gap path and transported contribution paths on Xsinx–Hetero. The quantities $G_S(t,x)$, D_k^{trans}, D_k^{inj}, D_k^{rem}, and C_k^a are defined in Sections 2.2 and 3.2. Thus panel (a) reports the mean absolute oracle-gap path, while panel (b) reports the AUCs of transported contribution paths. These are oracle-aided mechanism summaries, not endpoint error bounds.

For Xsinx–Hetero, PF-ODE has a smaller learned–oracle gap than EM and ExpInt in the low- and middle-variance regimes. The high-variance regime is different. There, the PF-ODE gap increases to 0.45, while EM and ExpInt remain near 0.11 and 0.12. The displayed pathwise-unrolled attribution also becomes strongly transport-weighted for PF-ODE, with a 74% transport share. Thus the main Xsinx evidence is not that PF-ODE is always larger, but that its high-variance regime shows a much larger pathwise learned–oracle gap and a transport-heavy attribution.

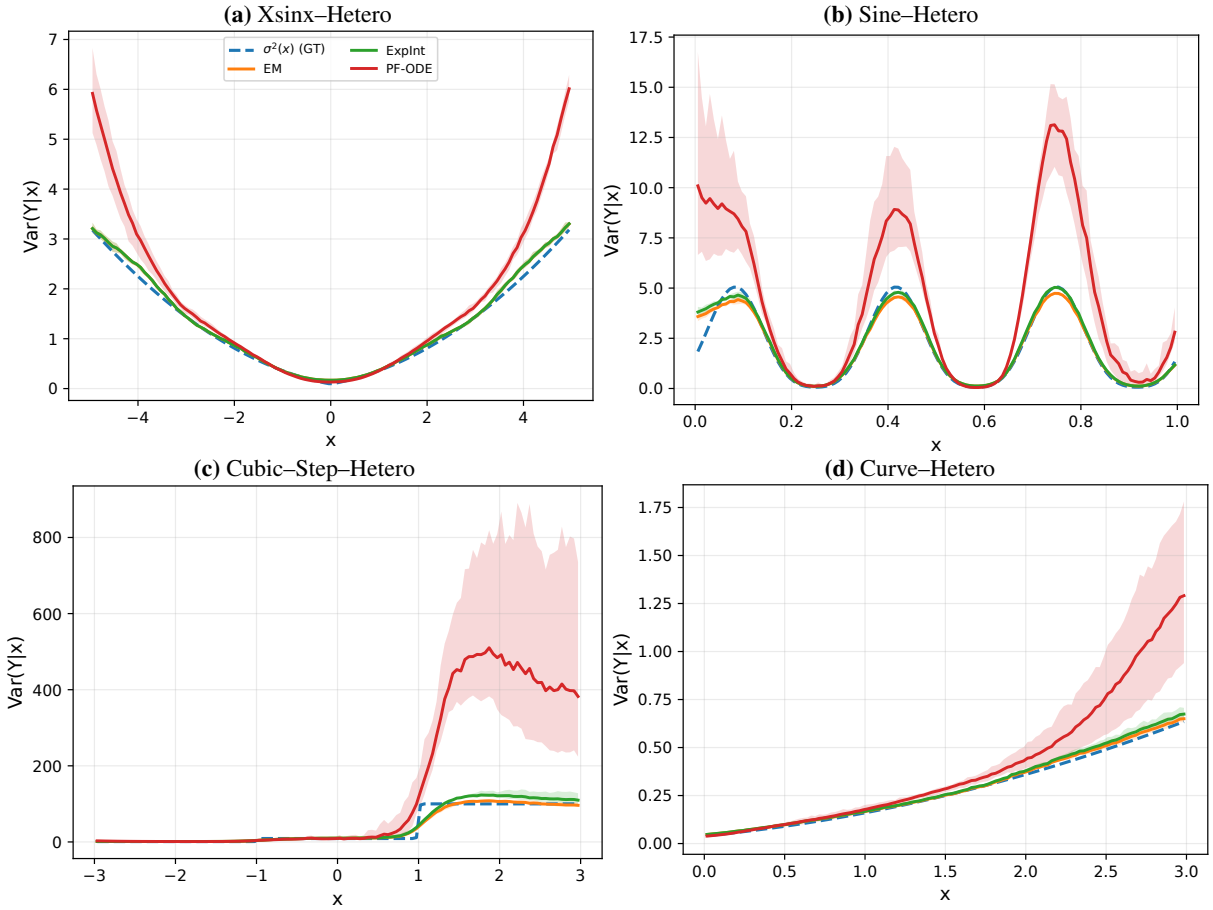

Figure 1. **Solver-dependent conditional variance profiles across the four synthetic tasks.** Panels (a)–(d) show Xsinx–Hetero, Sine–Hetero, Cubic–Step–Hetero, and Curve–Hetero, respectively. Solid curves are across-run means of binned endpoint variance estimates mapped to the original target scale for each solver. Shaded bands are pointwise empirical 95% bands across five independent runs. The dashed curve is the data variance target $v^{\mathrm{data}}(x) = \sigma^2(x)$.

Table 2. **Path summaries for endpoint variance distortion.** Across 60 dataset–solver–run points, oracle-gap path summaries are compared with a marginal KL-type path proxy as explanatory summaries of endpoint variance MSE. These are empirical correlations, not theoretical predictions of Proposition 3.2.

Explanatory summary	**Pearson** r	**Spearman** ρ	R^2
Oracle-gap AUC	0.875	0.753	0.766
Oracle-gap mean	0.915	0.894	0.837
KL path mean	0.851	0.500	0.724

Representative variance-regime mechanism summaries. The remaining variance-regime summaries, reported in Appendix D.9.2, Table 9, show the same qualitative lesson. High-variance regimes often enlarge the learned–oracle gap, but the dominant mechanism is not universal. In transport-heavy regimes, the local transport factorization in Appendix D.9.2, Eqs. (68)–(69), shows that the step-k contribution contains $\Pi_{1:k-1}(x)\big|J_k(x)^2 - (J_k^\star(x))^2\big|v_{t_k}^\star(x)$. Thus a comparable local transport mismatch can have a larger downstream effect when the transported oracle variance is large. Other regimes are remainder-heavy, indicating that nonlinear local-map effects can dominate even when this high-variance amplification channel is present.

5. Limitations and Future Work

The analysis is limited to finite-step second-moment fidelity. It does not characterize tail behavior, multimodality, higher-order shape, or full distributional calibration. Extending the variance-gap analysis to vector-valued responses and covariance matrices, broader forward processes beyond OU dynamics, and distributional guarantees for deterministic

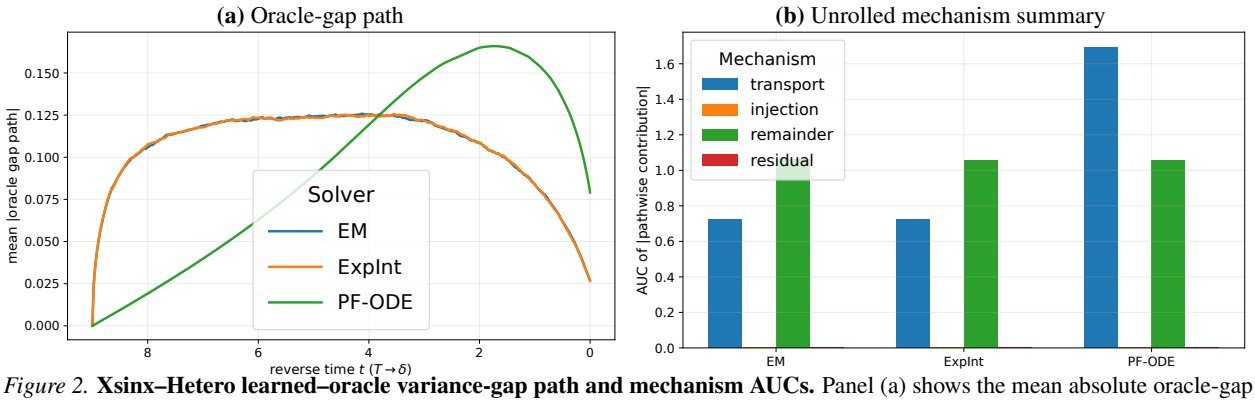

Figure 2. **Xsinx–Hetero learned–oracle variance-gap path and mechanism AUCs.** Panel (a) shows the mean absolute oracle-gap path over inputs. Panel (b) reports recorded-grid AUCs of the transported contribution paths for transport, injection, nonlinear remainder, and residual defects. These oracle-aided summaries are not endpoint error bounds.

Table 3. **Xsinx–Hetero mechanism summaries.** Gap is the average of $|G_S(t,x)|$ over recorded reverse times and inputs in each variance regime. Trans. and Rem. are the normalized AUC shares of the transported contribution paths C^{trans} and C^{rem}.

Regime	Solver	Gap	Trans.	Rem.
Low	EM	0.0665	34	66
	ExpInt	0.0665	34	66
	PF-ODE	0.0493	55	45
Middle	EM	0.0782	29	71
	ExpInt	0.0783	29	71
	PF-ODE	0.0620	54	46
High	EM	0.1100	57	43
	ExpInt	0.1200	57	43
	PF-ODE	0.4500	74	26

Note. Trans. and Rem. denote transport and remainder shares in percent. Injection and residual shares are omitted. Full summaries are in Table 9.

PF-ODE remain open problems. The present experiments are designed for controlled mechanism identification. Future work should test whether the same solver-dependent variance mechanisms persist in non-Gaussian, multimodal, real tabular, and scientific regression benchmarks. Another important direction is to develop oracle-free diagnostics that approximate or stress-test the learned–oracle variance gap when the true score and solver-matched oracle paths are unavailable.

6. Conclusion

We examined finite-step conditional diffusion regression with the learned score field, reverse-time grid, and noise schedule fixed, thereby isolating the reverse solver as a source of variation in conditional variance. The TV result serves as a distributional baseline for stochastic solvers, while the solver-matched variance-gap decomposition explains how variance mismatch is propagated by transport, injection, and nonlinear remainders. In controlled heteroskedastic experiments, solver changes affected both endpoint calibration and the path of variance mismatch. Endpoint accounting shows that the learned–oracle gap dominates solver-matched oracle bias across solvers and datasets.

The oracle-aided analysis is not a deployment-time diagnostic. Instead, it shows that diffusion-regression uncertainty estimates depend on both the learned score and the finite-step solver used at inference time. Solver sensitivity should therefore be evaluated under the chosen reverse-time grid in uncertainty-sensitive applications.

Impact Statement

This paper presents work whose goal is to advance the understanding of conditional diffusion regression. There are many potential societal consequences of our work, none of which we feel must be specifically highlighted here.

References

Anderson, B. D. Reverse-time diffusion equation models. *Stochastic Processes and their Applications*, 12(3):313–326, 1982.

Bakry, D., Gentil, I., and Ledoux, M. *Analysis and geometry of Markov diffusion operators*, volume 348. Springer Science & Business Media, 2013.

Beltran Velez, N., Grande, A. A., Nazaret, A., Kucukelbir, A., and Blei, D. Treeffuser: probabilistic prediction via conditional diffusions with gradient-boosted trees. *Advances in Neural Information Processing Systems*, 37: 118296–118325, 2024.

Benton, J., De Bortoli, V., Doucet, A., and Deligiannidis, G. Nearly d-linear convergence bounds for diffusion models via stochastic localization. *arXiv preprint arXiv:2308.03686*, 2023.

Cao, K., Chen, Y., Lu, J., Aréchiga, N., Gaidon, A., and Ma, T. Heteroskedastic and imbalanced deep learning with adaptive regularization. In *ICLR*, 2021.

Chen, H., Lee, H., and Lu, J. Improved analysis of score-based generative modeling: User-friendly bounds under minimal smoothness assumptions. In *International Conference on Machine Learning*, pp. 4735–4763. PMLR, 2023a.

Chen, S., Daras, G., and Dimakis, A. Restoration-degradation beyond linear diffusions: A non-asymptotic analysis for ddim-type samplers. In *International Conference on Machine Learning*, pp. 4462–4484. PMLR, 2023b.

De Bortoli, V. Convergence of denoising diffusion models under the manifold hypothesis. *arXiv preprint arXiv:2208.05314*, 2022.

De Bortoli, V., Galashov, A., Guntupalli, J. S., Zhou, G., Murphy, K. P., Gretton, A., and Doucet, A. Distributional diffusion models with scoring rules. In *International Conference on Machine Learning*, pp. 12632–12676. PMLR, 2025.

Gal, Y. and Ghahramani, Z. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pp. 1050–1059. PMLR, 2016.

Han, X. and Zhou, M. Diffusion boosted trees. *arXiv preprint arXiv:2406.01813*, 2024.

Han, X., Zheng, H., and Zhou, M. Card: Classification and regression diffusion models. *Advances in Neural Information Processing Systems*, 35:18100–18115, 2022.

Hyvärinen, A. Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research*, 6:695–709, 2005. URL https://www.jmlr.org/papers/volume6/hyvarinen05a/hyvarinen05a.pdf.

Jiao, Y., Kang, L., Liu, J., Peng, H., and Zuo, H. Model free prediction with uncertainty assessment. *IEEE Transactions on Information Theory*, 2025.

Kneissl, C., Bülte, C., Scholl, P., and Kutyniok, G. Improved probabilistic regression using diffusion models. In *Northern Lights Deep Learning Conference Abstracts 2026*, 2025.

Lakshminarayanan, B., Pritzel, A., and Blundell, C. Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in neural information processing systems*, 30, 2017.

Le, Q. V., Smola, A. J., and Canu, S. Heteroscedastic gaussian process regression. In *Proceedings of the 22nd International Conference on Machine Learning (ICML)*, 2005. URL https://www.machinelearning.org/proceedings/icml2005/papers/033_Heteroskedastic_LeSmolaCanu.pdf.

Lee, H., Lu, J., and Tan, Y. Convergence for score-based generative modeling with polynomial complexity. *Advances in Neural Information Processing Systems*, 35:22870–22882, 2022.

Li, G., Wei, Y., Chen, Y., and Chi, Y. Towards faster non-asymptotic convergence for diffusion-based generative models. *arXiv preprint arXiv:2306.09251*, 2023.

Li, M. Non-asymptotic convergence bound of conditional diffusion models. *arXiv preprint arXiv:2508.10944*, 2025.

Liang, Y., Ju, P., Liang, Y., and Shroff, N. Broadening target distributions for accelerated diffusion models via a novel analysis approach. *arXiv preprint arXiv:2402.13901*, 2024.

Merizzi, F., Evangelista, D., and Loukos, H. Controlling ensemble variance in diffusion models: an application for reanalyses downscaling. *Neural Computing and Applications*, 38(4):60, 2026.

Nichol, A. Q. and Dhariwal, P. Improved denoising diffusion probabilistic models. In *International conference on machine learning*, pp. 8162–8171. PMLR, 2021.

Nix, D. A. and Weigend, A. S. Estimating the mean and variance of the target probability distribution. In *Proceedings of 1994 ieee international conference on neural networks (ICNN'94)*, volume 1, pp. 55–60. IEEE, 1994.

Shu, D. and Farimani, A. B. Zero-shot uncertainty quantification using diffusion probabilistic models. *arXiv preprint arXiv:2408.04718*, 2024.

Skafte, N., Jørgensen, M., and Hauberg, S. Reliable training and estimation of variance networks. *Advances in Neural Information Processing Systems*, 32, 2019.

Song, Y., Sohl-Dickstein, J., Kingma, D. P., Kumar, A., Ermon, S., and Poole, B. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020.

Song, Y., Durkan, C., Murray, I., and Ermon, S. Maximum likelihood training of score-based diffusion models. *Advances in neural information processing systems*, 34:1415–1428, 2021.

Vincent, P. A connection between score matching and denoising autoencoders. *Neural Computation*, 23(7):1661–1674, 2011. URL https://www.iro.umontreal.ca/~vincentp/Publications/smdae_techreport.pdf.

White, H. A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity. *Econometrica: journal of the Econometric Society*, pp. 817–838, 1980.

Wong-Toi, E., Boyd, A., Fortuin, V., and Mandt, S. Understanding pathologies of deep heteroskedastic regression. In *Uncertainty in Artificial Intelligence*, pp. 3765–3790. PMLR, 2024.

Zhang, Q. and Chen, Y. Fast sampling of diffusion models with exponential integrator. *arXiv preprint arXiv:2204.13902*, 2022.

Zhao, C., Du, H., Liu, G., and Niyato, D. Supervised score-based modeling by gradient boosting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 22768–22776, 2025.

A. Distribution-Level TV Control

This appendix proves Proposition 3.1. We first compare the learned reverse chain with the solver-matched oracle chain. We then compare the oracle chain with the early-stopped target p_δ. The final bound follows by combining these two estimates and averaging over X. The argument applies to stochastic reverse-SDE discretizations with nondegenerate Gaussian one-step noise.

A.1. Proof architecture and standing quantities

Let $x \in \mathbb{R}^{d_X}$ and $y \in \mathbb{R}^{d_Y}$, and denote the conditional data law by $p_0(\cdot \mid x)$. For probability measures P, Q on $\mathbb{R}^{d_Y}$, define

$$\mathrm{KL}(P\|Q) := \int \log\Big(\frac{\mathrm{d}P}{\mathrm{d}Q}\Big)\,\mathrm{d}P, \qquad \mathrm{TV}(P,Q) := \sup_{A\in\mathcal{B}(\mathbb{R}^{d_Y})} |P(A) - Q(A)|.$$

We write $\|\cdot\|_2$ for the Euclidean norm and $\|\cdot\|_F$ for the Frobenius norm.

For each $t \in [\delta, T]$, let $p_t(\cdot \mid x)$ denote the OU-smoothed conditional law. Let $q_t(\cdot \mid x)$ denote the law induced by the learned finite-step reverse solver, and let $q_t^\star(\cdot \mid x)$ denote the law induced by the same solver family driven by the true score $s^\star$. Fix solver coefficients A_k, B_k, σ_k and write $h_k := t_k - t_{k-1}$.

Define

$$\Delta_k(y,x) := \|s_\theta(t_k, y, x) - s^\star(t_k, y, x)\|_2^2.$$

The first quantity measures solver-weighted score error under the forward conditional marginals, which is where score learning is naturally controlled.

$$\varepsilon_s^2 := \mathbb{E}_X\left[\sum_{k=1}^{K} \|B_k\|_F^2 \sigma_k^{-2}\, \mathbb{E}_{Y\sim p_{t_k}(\cdot|X)}\big[\Delta_k(Y,X)\big]\right]. \tag{13}$$

The second quantity records the price of transferring this score error from forward marginals to the learned reverse marginals that appear in the KL chain rule.

$$R_{p\to q}^2 := \mathbb{E}_X\left[\left(\frac{1}{2}\sum_{k=1}^{K} \|B_k\|_F^2 \sigma_k^{-2}\Big(\mathbb{E}_{q_{t_k}(\cdot|X)}[\Delta_k] - \mathbb{E}_{p_{t_k}(\cdot|X)}[\Delta_k]\Big)\right)_+\right]. \tag{14}$$

We use the following two regularity conditions.

Assumption A.1 (Conditional averaged score error)**.** The quantity ε_s^2 in (13) is finite.

Assumption A.2 (Uniform conditional second moment)**.**

$$M_2^{\max} := \sup_{x\in\mathcal{X}} \mathbb{E}[\|Y_0\|_2^2 \mid X = x] < \infty.$$

Stochastic solver scope. Throughout Appendix A, the distribution-level TV argument is restricted to stochastic reverse-SDE discretizations whose learned and oracle one-step kernels have the common nondegenerate Gaussian covariance $\sigma_k^2 I_{d_Y}$, with $\sigma_k^2 > 0$. Concretely, for each step $k = 1, \dots, K$,

$$Q_k^\theta(\cdot \mid y, x) = \mathcal{N}\big(A_k y + B_k s_\theta(t_k, y, x), \sigma_k^2 I_{d_Y}\big), \qquad Q_k^\star(\cdot \mid y, x) = \mathcal{N}\big(A_k y + B_k s^\star(t_k, y, x), \sigma_k^2 I_{d_Y}\big).$$

It is the condition that makes the one-step Gaussian KL comparison in Section A.2 available. Deterministic PF-ODE solvers are discussed separately in Section A.5.

A.2. Learned-to-oracle control on the reverse grid

This subsection controls the learned finite-step reverse chain relative to the solver-matched oracle chain. The argument follows the stochastic-solver learned-to-oracle decomposition of Li et al. (2023, Section 5.3): one first compares the learned and oracle Gaussian one-step kernels, and then accumulates the one-step KL errors by a reverse-grid KL chain rule. Our adaptation is conditional in x, averages the final bound over X, and separates the forward-to-reverse transfer residual $R_{p\to q}$, because the score error is naturally defined under the forward marginals $p_{t_k}(\cdot \mid X)$, whereas the KL chain rule produces expectations under the learned reverse marginals $q_{t_k}(\cdot \mid X)$.

One-step Gaussian kernel comparison. The first lemma is purely local. At a fixed step, common Gaussian covariance turns score mismatch into a one-step KL cost.

Lemma A.3 (One-step Gaussian kernel KL with equal covariance)**.** *Fix $k \in \{1, \dots, K\}$ and $x \in \mathcal{X}$. Consider*

$$Q_k(\cdot \mid y, x) := \mathcal{N}\Big(A_k y + B_k s_\theta(t_k, y, x),\, \sigma_k^2 I_{d_Y}\Big),$$
$$P_k(\cdot \mid y, x) := \mathcal{N}\Big(A_k y + B_k s^\star(t_k, y, x),\, \sigma_k^2 I_{d_Y}\Big).$$

Then

$$\mathrm{KL}\Big(Q_k(\cdot \mid y, x) \,\Big\|\, P_k(\cdot \mid y, x)\Big) = \frac{1}{2\sigma_k^2}\|B_k(s_\theta(t_k, y, x) - s^\star(t_k, y, x))\|_2^2.$$

Proof. Write

$$Q_k(\cdot \mid y, x) = \mathcal{N}(m_\theta, \sigma_k^2 I_{d_Y}), \qquad P_k(\cdot \mid y, x) = \mathcal{N}(m_\star, \sigma_k^2 I_{d_Y}),$$

where

$$m_\theta := A_k y + B_k s_\theta(t_k, y, x), \qquad m_\star := A_k y + B_k s^\star(t_k, y, x).$$

Since the two Gaussian laws have the same covariance,

$$\mathrm{KL}(\mathcal{N}(m_\theta, \Sigma)\|\mathcal{N}(m_\star, \Sigma)) = \frac{1}{2}(m_\theta - m_\star)^\top \Sigma^{-1}(m_\theta - m_\star), \qquad \Sigma = \sigma_k^2 I_{d_Y}.$$

Substituting $\Sigma^{-1} = \sigma_k^{-2} I_{d_Y}$ and $m_\theta - m_\star = B_k(s_\theta - s^\star)$ proves the claim. □

Reverse-grid KL chain rule. The previous lemma controls one transition. The next lemma accumulates those local transition costs along the reverse Markov path.

Lemma A.4 (KL chain rule along the reverse grid)**.** *Fix $x \in \mathcal{X}$ and let $t_0 = \delta < t_1 < \cdots < t_K = T$. Let $(q_{t_k})_{k=0}^K$ and $(p_{t_k})_{k=0}^K$ be the marginals of two Markov chains with one-step kernels $(Q_k)_{k=1}^K$ and $(P_k)_{k=1}^K$, respectively. Assume $q_T(\cdot \mid x) \ll p_T(\cdot \mid x)$ and $Q_k(\cdot \mid Y_{t_k}, x) \ll P_k(\cdot \mid Y_{t_k}, x)\, q_{t_k}(\cdot \mid x)$-a.s. Then*

$$\begin{aligned}\mathrm{KL}\big(\mathsf{Q}_{0:K}(\cdot \mid x) \,\big\|\, \mathsf{P}_{0:K}(\cdot \mid x)\big) &= \mathrm{KL}\big(q_T(\cdot \mid x) \,\big\|\, p_T(\cdot \mid x)\big) \\ &\quad + \sum_{k=1}^K \mathbb{E}_{q_{t_k}(\cdot|x)}\big[\mathrm{KL}\big(Q_k(\cdot \mid Y_{t_k}, x) \,\big\|\, P_k(\cdot \mid Y_{t_k}, x)\big)\big]\,.\end{aligned} \tag{15}$$

Moreover,

$$\mathrm{KL}\big(q_\delta(\cdot \mid x) \,\big\|\, p_\delta(\cdot \mid x)\big) \le \mathrm{KL}\big(\mathsf{Q}_{0:K}(\cdot \mid x) \,\big\|\, \mathsf{P}_{0:K}(\cdot \mid x)\big). \tag{16}$$

Proof. Let $\mathsf{Q}_{0:K}$ and $\mathsf{P}_{0:K}$ denote the path laws on $(Y_{t_0}, \dots, Y_{t_K})$. By the Markov property,

$$\mathsf{Q}_{0:K}(\mathrm{d}y_{0:K} \mid x) = q_T(\mathrm{d}y_K \mid x) \prod_{k=1}^K Q_k(\mathrm{d}y_{k-1} \mid y_k, x),$$

and similarly

$$\mathsf{P}_{0:K}(\mathrm{d}y_{0:K} \mid x) = p_T(\mathrm{d}y_K \mid x) \prod_{k=1}^K P_k(\mathrm{d}y_{k-1} \mid y_k, x).$$

Hence

$$\frac{\mathrm{d}\mathsf{Q}_{0:K}}{\mathrm{d}\mathsf{P}_{0:K}}(y_{0:K}) = \frac{\mathrm{d}q_T}{\mathrm{d}p_T}(y_K) \prod_{k=1}^K \frac{\mathrm{d}Q_k(\cdot \mid y_k, x)}{\mathrm{d}P_k(\cdot \mid y_k, x)}(y_{k-1}).$$

Taking expectation of the log-likelihood ratio under $\mathsf{Q}_{0:K}$ yields

$$\mathrm{KL}(\mathsf{Q}_{0:K}\|\mathsf{P}_{0:K}) = \mathbb{E}_{\mathsf{Q}_{0:K}}\left[\log \frac{\mathrm{d}q_T}{\mathrm{d}p_T}(Y_{t_K})\right] + \sum_{k=1}^K \mathbb{E}_{\mathsf{Q}_{0:K}}\left[\log \frac{\mathrm{d}Q_k(\cdot \mid Y_{t_k}, x)}{\mathrm{d}P_k(\cdot \mid Y_{t_k}, x)}(Y_{t_{k-1}})\right].$$

The first term is $\mathrm{KL}(q_T \| p_T)$. Conditioning the k-th term on Y_{t_k} gives

$$\mathbb{E}_{\mathbf{Q}_{0:K}}\left[\log \frac{\mathrm{d}Q_k(\cdot \mid Y_{t_k}, x)}{\mathrm{d}P_k(\cdot \mid Y_{t_k}, x)}(Y_{t_{k-1}})\right] = \mathbb{E}_{q_{t_k}(\cdot|x)}\left[\mathrm{KL}\big(Q_k(\cdot \mid Y_{t_k}, x) \,\|\, P_k(\cdot \mid Y_{t_k}, x)\big)\right],$$

which proves (15). Finally, (16) follows by data processing under the projection $\pi(y_{0:K}) = y_0$. □

Conditional learned-to-oracle control. The following proposition is the Li-style stochastic path comparison specialized to the conditional setting. We now combine the local Gaussian KL formula and the pathwise chain rule. The only additional bookkeeping step is to replace reverse-marginal score-error expectations by forward-marginal ones plus $R_{p\to q}$.

Proposition A.5 (Conditional learned-to-oracle KL control)**.** *Assume Assumption A.1 and the stochastic Gaussian solver scope stated in Section A.1. Then*

$$\mathbb{E}_X\Big[\mathrm{KL}\big(q_\delta(\cdot \mid X) \,\|\, q^\star_\delta(\cdot \mid X)\big)\Big] \le \frac{1}{2}\,\varepsilon_s^2 + \mathcal{R}^2_{p\to q}. \tag{17}$$

Proof. Fix $x \in \mathcal{X}$. Let $Q^\theta_k(\cdot \mid y, x)$ and $Q^\star_k(\cdot \mid y, x)$ denote the one-step kernels from t_k to t_{k-1} under s_θ and $s^\star$, respectively. By Lemma A.3,

$$\mathrm{KL}\big(Q^\theta_k(\cdot \mid y, x) \,\|\, Q^\star_k(\cdot \mid y, x)\big) = \frac{1}{2\sigma_k^2}\|B_k(s_\theta(t_k, y, x) - s^\star(t_k, y, x))\|_2^2.$$

Using

$$\|B_k u\|_2^2 \le \|B_k\|_{\mathrm{op}}^2 \|u\|_2^2 \le \|B_k\|_F^2 \|u\|_2^2,$$

we obtain

$$\mathrm{KL}\big(Q^\theta_k(\cdot \mid y, x) \,\|\, Q^\star_k(\cdot \mid y, x)\big) \le \frac{1}{2}\|B_k\|_F^2 \sigma_k^{-2} \Delta_k(y, x).$$

Applying Lemma A.4 to the two discrete chains driven by s_θ and $s^\star$, which share the same Gaussian initialization at time T, gives

$$\mathrm{KL}(q_\delta(\cdot \mid x) \| q^\star_\delta(\cdot \mid x)) \le \frac{1}{2}\sum_{k=1}^K w_k\, \mathbb{E}_{q_{t_k}(\cdot|x)}[\Delta_k], \qquad w_k := \|B_k\|_F^2 \sigma_k^{-2}.$$

Now decompose

$$\mathbb{E}_{q_{t_k}(\cdot|x)}[\Delta_k] = \mathbb{E}_{p_{t_k}(\cdot|x)}[\Delta_k] + \Big(\mathbb{E}_{q_{t_k}(\cdot|x)}[\Delta_k] - \mathbb{E}_{p_{t_k}(\cdot|x)}[\Delta_k]\Big).$$

Multiplying by $w_k/2$, summing over k, and averaging over X, we obtain

$$\mathbb{E}_X\Big[\mathrm{KL}(q_\delta(\cdot \mid X) \| q^\star_\delta(\cdot \mid X))\Big] \le \frac{1}{2}\,\mathbb{E}_X\left[\sum_{k=1}^K w_k\, \mathbb{E}_{p_{t_k}(\cdot|X)}[\Delta_k]\right] + \mathbb{E}_X\left[\left(\frac{1}{2}\sum_{k=1}^K w_k \big(\mathbb{E}_{q_{t_k}(\cdot|X)}[\Delta_k] - \mathbb{E}_{p_{t_k}(\cdot|X)}[\Delta_k]\big)\right)_+\right].$$

The first term is bounded by $\frac{1}{2}\,\varepsilon_s^2$ by (13), and the second is exactly $\mathcal{R}^2_{p\to q}$ by (14). This proves (17). □

A.3. Oracle stochastic-solver discretization control

This subsection controls the solver-matched oracle chain relative to the early-stopped target $p_\delta(\cdot \mid x)$. This is the Chen-style oracle discretization block, adapted to conditional OU marginals. The learned score does not appear in this subsection: the oracle chain is driven by $s^\star(t, y, x) = \nabla_y \log p_t(y \mid x)$. The role of this block is to prove a conditional bound of the form

$$\mathrm{KL}\big(p_\delta(\cdot \mid x) \,\|\, q^\star_\delta(\cdot \mid x)\big) \le C_{\mathrm{init}}(d_Y + M_2^{\max})e^{-T} + C_{\mathrm{disc},S}\Sigma_K.$$

Chen-style discretization quantities. The following quantity is the solver-specific discretization budget used by the oracle bound. This is the only place in this block where EM and ExpInt differ. Fix a stochastic solver family $S \in \{\mathrm{EM}, \mathrm{ExpInt}\}$. Define

$$\vartheta_t^2 := 1 - e^{-t}, \qquad \Gamma_K := \sum_{k=1}^{K} \frac{h_k^2}{\vartheta_{t_{k-1}}^4},$$

The solver-specific discretization quantity is

$$\Sigma_K := \begin{cases} d_Y^2\, \Gamma_K + M_2^{\max} \sum_{k=1}^{K} h_k^3, & S = \mathrm{EM}, \\ d_Y^2\, \Gamma_K, & S = \mathrm{ExpInt}. \end{cases} \tag{18}$$

The Chen-style oracle discretization argument uses the reverse-grid condition

$$\frac{h_k}{\vartheta_{t_{k-1}}^2} \le \frac{1}{K_0 d_Y}, \qquad k = 1, \dots, K, \tag{19}$$

for a sufficiently large universal constant $K_0 > 0$. This grid condition is invoked only in the oracle-to-target discretization block and in the final TV assembly.

The first oracle error is the mismatch between the finite terminal initialization $\mathcal{N}(0, I_{d_Y})$ and the true OU marginal $p_T(\cdot \mid x)$.

Lemma A.6 (Conditional Gaussian initialization bound (cf. Lemma C.4 of Chen et al. (2023a))). *For every $T > 0$ and for μ_X-a.e. x,*

$$\mathrm{KL}\big(p_T(\cdot \mid x) \,\|\, \mathcal{N}(0, I_{d_Y})\big) \le \frac{d_Y}{2} \log\!\Big(\frac{1}{1 - e^{-T}}\Big) + \frac{1}{2} e^{-T} \mathbb{E}[\|Y_0\|_2^2 \mid X = x]. \tag{20}$$

Consequently, under Assumption A.2, if $T \ge 1$, then

$$\mathrm{KL}\big(p_T(\cdot \mid x) \,\|\, \mathcal{N}(0, I_{d_Y})\big) \le C(d_Y + M_2^{\max}) e^{-T}. \tag{21}$$

Proof. Fix $x \in \mathcal{X}$. Let $p_{T|0}(\cdot \mid y, x)$ denote the conditional OU transition density given $Y_0 = y$. Since $u \mapsto u \log u$ is convex,

$$\int p_T(z \mid x) \log p_T(z \mid x)\, \mathrm{d}z = \int \Big(\int p_{T|0}(z \mid y, x)\, p_0(\mathrm{d}y \mid x) \Big) \log \Big(\int p_{T|0}(z \mid y, x)\, p_0(\mathrm{d}y \mid x) \Big)\, \mathrm{d}z$$

$$\le \int \Big(\int p_{T|0}(z \mid y, x) \log p_{T|0}(z \mid y, x)\, p_0(\mathrm{d}y \mid x) \Big)\, \mathrm{d}z = \int \Big(\int p_{T|0}(z \mid y, x) \log p_{T|0}(z \mid y, x)\, \mathrm{d}z \Big)\, p_0(\mathrm{d}y \mid x).$$

Since

$$Y_T \mid (Y_0 = y, X = x) \sim \mathcal{N}(e^{-T/2} y, (1 - e^{-T}) I_{d_Y}),$$

we have

$$\int p_{T|0}(z \mid y, x) \log p_{T|0}(z \mid y, x)\, \mathrm{d}z = -\frac{d_Y}{2} \log\big(2\pi(1 - e^{-T})\big) - \frac{d_Y}{2}.$$

Therefore,

$$\int p_T(z \mid x) \log p_T(z \mid x)\, \mathrm{d}z \le -\frac{d_Y}{2} \log\big(2\pi(1 - e^{-T})\big) - \frac{d_Y}{2}.$$

Using

$$\mathrm{KL}\big(p_T(\cdot \mid x) \,\|\, \mathcal{N}(0, I_{d_Y})\big) = \int p_T(z \mid x) \log p_T(z \mid x)\, \mathrm{d}z + \mathbb{E}[\|Y_T\|_2^2 \mid X = x]/2 + \frac{d_Y}{2} \log(2\pi),$$

we obtain

$$\mathrm{KL}\big(p_T(\cdot \mid x) \,\|\, \mathcal{N}(0, I_{d_Y})\big) \le \frac{d_Y}{2} \log\!\Big(\frac{1}{1 - e^{-T}}\Big) + \frac{1}{2} e^{-T} \mathbb{E}[\|Y_0\|_2^2 \mid X = x].$$

This proves (20). If $T \geq 1$, then

$$\log\left(\frac{1}{1-e^{-T}}\right) \leq Ce^{-T},$$

and Assumption A.2 gives

$$\mathbb{E}[\|Y_0\|_2^2 \mid X = x] \leq M_2^{\max}.$$

The simplified bound (21) follows. □

After initialization, the oracle discretization error is reduced to integrated local errors along the exact reverse path.

Proposition A.7 (Conditional oracle KL reduction (adapted from Proposition C.3 of Chen et al. (2023a))). *Fix $x \in \mathcal{X}$. Let $t'_k := T - t_{K-k}$ for $k = 0, \ldots, K$, so that*

$$0 = t'_0 < t'_1 < \cdots < t'_K = T - \delta.$$

Let $q^\star_\delta(\cdot \mid x)$ be the time-δ marginal of the oracle stochastic solver driven by $s^\star$ and initialized at time T from $\mathcal{N}(0, I_{d_Y})$. Then, for the exponential integrator,

$$\mathrm{KL}\big(p_\delta(\cdot \mid x) \,\|\, q^\star_\delta(\cdot \mid x)\big) \leq \mathrm{KL}\big(p_T(\cdot \mid x) \,\|\, \mathcal{N}(0, I_{d_Y})\big) + \sum_{k=1}^{K} \int_{t_{k-1}}^{t_k} \mathbb{E}\big[\|\nabla_y \log p_t(Y_t \mid x) - \nabla_y \log p_{t_k}(Y_{t_k} \mid x)\|_2^2\big]\, \mathrm{d}t. \tag{22}$$

For Euler–Maruyama, the same bound holds with the additional term

$$\sum_{k=1}^{K} \int_{t_{k-1}}^{t_k} \mathbb{E}\big[\|Y_t - Y_{t_k}\|_2^2\big]\, \mathrm{d}t. \tag{23}$$

Proof. Fix $k \in \{0, \ldots, K-1\}$ and $a \in \mathbb{R}^{d_Y}$. Consider the exact reverse SDE on $[t'_k, t'_{k+1}]$.

$$\mathrm{d}\bar{Y}^{k,a}_{t'} = \Big(\frac{1}{2}\bar{Y}^{k,a}_{t'} + \nabla_y \log p_{T-t'}(\bar{Y}^{k,a}_{t'} \mid x)\Big)\mathrm{d}t' + \mathrm{d}W_{t'}, \qquad \bar{Y}^{k,a}_{t'_k} = a. \tag{24}$$

Let $\bar{p}_{t'|t'_k}(\cdot \mid a, x)$ denote its transition density.

For the oracle exponential integrator, define the local interpolating process by

$$\mathrm{d}\widehat{Y}_{t'} = \Big(\frac{1}{2}\widehat{Y}_{t'} + \nabla_y \log p_{t_{K-k}}(a \mid x)\Big)\mathrm{d}t' + \mathrm{d}W_{t'}, \qquad \widehat{Y}_{t'_k} = a. \tag{25}$$

Let $\widehat{q}_{t'|t'_k}(\cdot \mid a, x)$ denote its transition density.

For the Euler–Maruyama local interpolation, use

$$\mathrm{d}\widehat{Y}_{t'} = \Big(\frac{1}{2}a + \nabla_y \log p_{t_{K-k}}(a \mid x)\Big)\mathrm{d}t' + \mathrm{d}W_{t'}, \qquad \widehat{Y}_{t'_k} = a. \tag{26}$$

Both local processes have the same unit diffusion coefficient. For $t' \in (t'_k, t'_{k+1}]$, the differential KL identity gives

$$\begin{aligned}\frac{\mathrm{d}}{\mathrm{d}t'} \mathrm{KL}\big(\bar{p}_{t'|t'_k}(\cdot \mid a, x) \,\|\, \widehat{q}_{t'|t'_k}(\cdot \mid a, x)\big) &= -J\big(\bar{p}_{t'|t'_k}(\cdot \mid a, x) \,\|\, \widehat{q}_{t'|t'_k}(\cdot \mid a, x)\big) \\ &\quad + \mathbb{E}\left[\Big\langle F_1(\bar{Y}^{k,a}_{t'}, t') - F_2(\bar{Y}^{k,a}_{t'}, t'), \nabla \log \frac{\bar{p}_{t'|t'_k}(\bar{Y}^{k,a}_{t'} \mid a, x)}{\widehat{q}_{t'|t'_k}(\bar{Y}^{k,a}_{t'} \mid a, x)} \Big\rangle\right],\end{aligned}$$

where F_1 and F_2 are the drifts of the exact and local approximate processes. Since

$$\langle u, v \rangle \leq \frac{1}{2}\|u\|_2^2 + \frac{1}{2}\|v\|_2^2,$$

the Fisher-information term cancels and yields

$$\frac{\mathrm{d}}{\mathrm{d}t'}\,\mathrm{KL}\big(\bar{p}_{t'|t'_k}(\cdot \mid a,x)\,\|\,\widehat{q}_{t'|t'_k}(\cdot \mid a,x)\big) \le \frac{1}{2}\mathbb{E}\big[\|F_1(\bar{Y}^{k,a}_{t'},t') - F_2(\bar{Y}^{k,a}_{t'},t')\|_2^2\big].$$

For the exponential integrator,

$$F_1(\bar{Y}^{k,a}_{t'},t') - F_2(\bar{Y}^{k,a}_{t'},t') = \nabla_y \log p_{T-t'}(\bar{Y}^{k,a}_{t'} \mid x) - \nabla_y \log p_{t_{K-k}}(a \mid x).$$

Hence

$$\frac{\mathrm{d}}{\mathrm{d}t'}\,\mathrm{KL}\big(\bar{p}_{t'|t'_k}(\cdot \mid a,x)\,\|\,\widehat{q}_{t'|t'_k}(\cdot \mid a,x)\big) \le \frac{1}{2}\,\mathbb{E}\big[\|\nabla_y \log p_{T-t'}(\bar{Y}^{k,a}_{t'} \mid x) - \nabla_y \log p_{t_{K-k}}(a \mid x)\|_2^2\big].$$

For Euler–Maruyama,

$$F_1(\bar{Y}^{k,a}_{t'},t') - F_2(\bar{Y}^{k,a}_{t'},t') = \frac{1}{2}(\bar{Y}^{k,a}_{t'} - a) + \nabla_y \log p_{T-t'}(\bar{Y}^{k,a}_{t'} \mid x) - \nabla_y \log p_{t_{K-k}}(a \mid x).$$

Hence

$$\frac{\mathrm{d}}{\mathrm{d}t'}\,\mathrm{KL}\big(\bar{p}_{t'|t'_k}(\cdot \mid a,x)\,\|\,\widehat{q}_{t'|t'_k}(\cdot \mid a,x)\big) \le \mathbb{E}\big[\|\nabla_y \log p_{T-t'}(\bar{Y}^{k,a}_{t'} \mid x) - \nabla_y \log p_{t_{K-k}}(a \mid x)\|_2^2\big] + \frac{1}{4}\mathbb{E}\big[\|\bar{Y}^{k,a}_{t'} - a\|_2^2\big].$$

Since the two processes start from the same point a at time t'_k, the local KL tends to zero as $t' \downarrow t'_k$. Integrating over $[t'_k, t'_{k+1}]$ gives

$$\mathrm{KL}\big(\bar{p}_{t'_{k+1}|t'_k}(\cdot \mid a,x)\,\|\,\widehat{q}_{t'_{k+1}|t'_k}(\cdot \mid a,x)\big) \le \frac{1}{2}\int_{t'_k}^{t'_{k+1}} \mathbb{E}\big[\|\nabla_y \log p_{T-t'}(\bar{Y}^{k,a}_{t'} \mid x) - \nabla_y \log p_{t_{K-k}}(a \mid x)\|_2^2\big]\,\mathrm{d}t'$$

for the exponential integrator, and the same bound plus

$$\frac{1}{4}\int_{t'_k}^{t'_{k+1}} \mathbb{E}\big[\|\bar{Y}^{k,a}_{t'} - a\|_2^2\big]\,\mathrm{d}t'$$

for Euler–Maruyama.

Now average in a against the exact reverse marginal at time t'_k. Apply the KL chain rule over the reverse grid and sum over $k = 0, \dots, K-1$. Finally set $t = T - t'$, so that $t \in [t_{K-k-1}, t_{K-k}]$. Re-indexing by $j = K - k$ yields (22) and (23). □

The remaining lemmas close the local error terms produced by the oracle KL reduction. The first one controls the additional state-increment term that appears for Euler–Maruyama.

Lemma A.8 (Conditional reverse increment bound (cf. Lemma C.5 of Chen et al. (2023a))). *Suppose $h_k \le 1$ for all k. Then for $t_{k-1} \le t \le t_k$,*

$$\mathbb{E}[\|Y_t - Y_{t_k}\|_2^2 \mid X = x] \le C d_Y (t_k - t) + C(M_2^{\max} + d_Y)(t_k - t)^2.$$

Consequently,

$$\sum_{k=1}^{K}\int_{t_{k-1}}^{t_k} \mathbb{E}[\|Y_t - Y_{t_k}\|_2^2 \mid X = x]\,\mathrm{d}t \le C d_Y \sum_{k=1}^{K} h_k^2 + C M_2^{\max} \sum_{k=1}^{K} h_k^3. \tag{27}$$

Proof. For $t \le t_k$,

$$Y_{t_k} - Y_t = -\frac{1}{2}\int_t^{t_k} Y_u\,\mathrm{d}u + (W_{t_k} - W_t).$$

Hence

$$\mathbb{E}[\|Y_{t_k} - Y_t\|_2^2 \mid X = x] \le 2\,\mathbb{E}\Big[\Big\|\int_t^{t_k} Y_u\,\mathrm{d}u\Big\|_2^2 \mid X = x\Big] + 2\,\mathbb{E}[\|W_{t_k} - W_t\|_2^2].$$

By Cauchy–Schwarz,

$$\mathbb{E}\Big[\Big\| \int_t^{t_k} Y_u \, \mathrm{d}u\Big\|_2^2 \mid X = x\Big] \le (t_k - t) \int_t^{t_k} \mathbb{E}[\|Y_u\|_2^2 \mid X = x] \, \mathrm{d}u.$$

For the OU forward process,

$$Y_u = e^{-u/2} Y_0 + \sqrt{1 - e^{-u}}\, Z_u, \qquad Z_u \sim \mathcal{N}(0, I_{d_Y}),$$

so

$$\mathbb{E}[\|Y_u\|_2^2 \mid X = x] \le \mathbb{E}[\|Y_0\|_2^2 \mid X = x] + d_Y \le M_2^{\max} + d_Y.$$

Also

$$\mathbb{E}[\|W_{t_k} - W_t\|_2^2] = d_Y (t_k - t).$$

Combining the above proves the pointwise bound. Integrating over $t \in [t_{k-1}, t_k]$ and summing over k yields (27). □

Next, the score mismatch across two times is reduced to a spatial displacement at a fixed time.

Lemma A.9 (Conditional time-to-space reduction (cf. Lemma C.6 of Chen et al. (2023a))). *For any $0 \le t \le s \le T$,*

$$\begin{aligned}\mathbb{E}\big[\|\nabla_y \log p_t(Y_t \mid x) - \nabla_y \log p_s(Y_s \mid x)\|_2^2\big] \le 4\,\mathbb{E}\big[\|\nabla_y \log p_t(Y_t \mid x) - \nabla_y \log p_t(\alpha_{t,s}^{-1} Y_s \mid x)\|_2^2\big] \\ +2(1 - \alpha_{t,s}^{-1})^2 \mathbb{E}[\|\nabla_y \log p_t(Y_t \mid x)\|_2^2],\end{aligned} \tag{28}$$

where $\alpha_{t,s} := e^{-(s-t)/2}$.

Proof. Since

$$Y_s \mid (Y_t = y, X = x) \sim \mathcal{N}(\alpha_{t,s} y, (1 - \alpha_{t,s}^2) I_{d_Y}),$$

the conditional density $p_{t|s}(\cdot \mid \cdot, x)$ gives the posterior representation

$$\nabla_y \log p_s(z \mid x) = \alpha_{t,s}^{-1}\, \mathbb{E}_{U \sim p_{t|s}(\cdot \mid z, x)}[\nabla_y \log p_t(U \mid x)].$$

Therefore

$$\begin{aligned}&\mathbb{E}\big[\|\nabla_y \log p_t(\alpha_{t,s}^{-1} Y_s \mid x) - \nabla_y \log p_s(Y_s \mid x)\|_2^2\big] \\ &= \mathbb{E}\Big[\Big\|\nabla_y \log p_t(\alpha_{t,s}^{-1} Y_s \mid x) - \alpha_{t,s}^{-1} \mathbb{E}[\nabla_y \log p_t(Y_t \mid x) \mid Y_s, X = x]\Big\|_2^2\Big] \\ &\le \mathbb{E}\big[\|\alpha_{t,s}^{-1} \nabla_y \log p_t(Y_t \mid x) - \nabla_y \log p_t(\alpha_{t,s}^{-1} Y_s \mid x)\|_2^2\big] \\ &\le 2(1 - \alpha_{t,s}^{-1})^2 \mathbb{E}[\|\nabla_y \log p_t(Y_t \mid x)\|_2^2] + 2\,\mathbb{E}\big[\|\nabla_y \log p_t(Y_t \mid x) - \nabla_y \log p_t(\alpha_{t,s}^{-1} Y_s \mid x)\|_2^2\big].\end{aligned}$$

Finally,

$$\begin{aligned}\|\nabla_y \log p_t(Y_t \mid x) - \nabla_y \log p_s(Y_s \mid x)\|_2^2 \le 2\|\nabla_y \log p_t(Y_t \mid x) - \nabla_y \log p_t(\alpha_{t,s}^{-1} Y_s \mid x)\|_2^2 \\ +2\|\nabla_y \log p_t(\alpha_{t,s}^{-1} Y_s \mid x) - \nabla_y \log p_s(Y_s \mid x)\|_2^2.\end{aligned}$$

Combining the two displays gives (28). □

The spatial displacement is controlled through a conditional Hessian bound for the OU-smoothed density.

Lemma A.10 (Conditional sub-exponential Hessian bound (conditional version of Lemma C.7 of Chen et al. (2023a))). *For each fixed $x \in \mathcal{X}$ and $t > 0$,*

$$\|\nabla_y^2 \log p_t(Y_t \mid x)\|_{F,\psi_1} \le C\, \frac{d_Y}{\vartheta_t^2}. \tag{29}$$

Proof. Fix $x \in \mathcal{X}$. Conditional on $X = x$, the OU perturbation has the fixed-base-law representation

$$Y_t = e^{-t/2} Y_0 + \vartheta_t Z, \qquad Z \sim \mathcal{N}(0, I_{d_Y}),$$

where $Y_0 \sim p_0(\cdot \mid x)$ and $\vartheta_t^2 = 1 - e^{-t}$. Let p_t^x denote the density of this Gaussian perturbation. This is exactly the fixed-base-law setting of Lemma C.7 of Chen et al. (2023a), applied pointwise to the conditional base law $p_0(\cdot \mid x)$. Therefore,

$$\|\nabla_y^2 \log p_t^x(Y_t)\|_{F,\psi_1} \le C \, \frac{d_Y}{\vartheta_t^2}.$$

Since $p_t^x(\cdot) = p_t(\cdot \mid x)$, this is precisely

$$\|\nabla_y^2 \log p_t(Y_t \mid x)\|_{F,\psi_1} \le C \, \frac{d_Y}{\vartheta_t^2}.$$

The conditioning on $X = x$ only fixes the base distribution; the Gaussian perturbation argument and the ψ_1-norm bound are otherwise unchanged. The constant C is universal and does not depend on x. This proves (29). □

Combining the Hessian control with a fixed-base-law change-of-measure estimate gives the local space-discretization estimate below. The density-ratio step is adapted from Appendix C.1 of Chen et al. (2023a) and is not reproved here.

Lemma A.11 (Conditional space-discretization bound (conditional fixed-x adaptation of Lemma C.8 and Appendix C.1 of Chen et al. (2023a))). *There exists a universal constant $K_{\mathrm{disc}} > 0$ such that if*

$$0 \le t \le s \le T, \qquad \frac{s-t}{\vartheta_t^2} \le \frac{1}{K_{\mathrm{disc}} d_Y},$$

then

$$\mathbb{E}\big[\|\nabla_y \log p_t(Y_t \mid x) - \nabla_y \log p_t(\alpha_{t,s}^{-1} Y_s \mid x)\|_2^2\big] \le C \, \frac{d_Y^2 (s-t)}{\vartheta_t^4}. \tag{30}$$

Proof. Write

$$Z_{t,s} := \alpha_{t,s}^{-1} Y_s - Y_t.$$

Then $Z_{t,s} \sim \mathcal{N}(0, (e^{s-t} - 1) I_{d_Y})$ and $Z_{t,s}$ is independent of Y_t. By the fundamental theorem of calculus,

$$\nabla_y \log p_t(Y_t \mid x) - \nabla_y \log p_t(\alpha_{t,s}^{-1} Y_s \mid x) = -\int_0^1 \nabla_y^2 \log p_t(Y_t + a Z_{t,s} \mid x) \, Z_{t,s} \, \mathrm{d}a.$$

Therefore,

$$\mathbb{E}\big[\|\nabla_y \log p_t(Y_t \mid x) - \nabla_y \log p_t(\alpha_{t,s}^{-1} Y_s \mid x)\|_2^2\big] \le \int_0^1 \mathbb{E}\big[\|\nabla_y^2 \log p_t(Y_t + a Z_{t,s} \mid x) \, Z_{t,s}\|_2^2\big] \, \mathrm{d}a.$$

At this point we invoke the fixed-base-law change-of-measure estimate from Appendix C.1 of Chen et al. (2023a). Conditional on $X = x$, the base law is $p_0(\cdot \mid x)$, and Lemma A.10 gives the same Hessian sub-exponential control with constants independent of x. Therefore, under

$$\frac{s-t}{\vartheta_t^2} \le \frac{1}{K_{\mathrm{disc}} d_Y},$$

the shifted point $Y_t + a Z_{t,s}$ is absolutely continuous with respect to $p_t(\cdot \mid x)$ with the corresponding controlled density-ratio bound, uniformly over $a \in [0,1]$. Combining this imported density-ratio estimate with the Hessian sub-exponential bound (29) and the Gaussian moment

$$\mathbb{E}[\|Z_{t,s}\|_2^2] \asymp d_Y (s-t),$$

gives

$$\mathbb{E}\big[\|\nabla_y^2 \log p_t(Y_t + a Z_{t,s} \mid x) \, Z_{t,s}\|_2^2\big] \le C \, \frac{d_Y^2 (s-t)}{\vartheta_t^4}, \qquad a \in [0,1].$$

Integrating over $a \in [0,1]$ proves (30). □

The final Chen-style proposition bundles initialization mismatch and local discretization into the compact oracle bound involving Σ_K.

Proposition A.12 (Conditional early-stopped oracle KL bound (adapted from Theorem 2.2 and Appendix C of Chen et al. (2023a)))**.** *Assume Assumption A.2, $T \geq 1$, and the reverse-grid condition (19). Then*

$$\mathrm{KL}\big(p_\delta(\cdot \mid x) \,\|\, q^\star_\delta(\cdot \mid x)\big) \leq C_{\mathrm{init}}(d_Y + M_2^{\max})e^{-T} + C_{\mathrm{disc},S}\,\Sigma_K. \tag{31}$$

Proof. By Proposition A.7,

$$\mathrm{KL}\big(p_\delta(\cdot \mid x) \,\|\, q^\star_\delta(\cdot \mid x)\big) \leq \mathrm{KL}\big(p_T(\cdot \mid x) \,\|\, \mathcal{N}(0, I_{d_Y})\big) + \mathcal{D}_{\mathrm{score}} + \mathbf{1}_{\{S=\mathrm{EM}\}}\mathcal{D}_{\mathrm{EM}},$$

where

$$\mathcal{D}_{\mathrm{score}} := \sum_{k=1}^{K} \int_{t_{k-1}}^{t_k} \mathbb{E}\big[\|\nabla_y \log p_t(Y_t \mid x) - \nabla_y \log p_{t_k}(Y_{t_k} \mid x)\|_2^2\big]\,\mathrm{d}t$$

and

$$\mathcal{D}_{\mathrm{EM}} := \sum_{k=1}^{K} \int_{t_{k-1}}^{t_k} \mathbb{E}[\|Y_t - Y_{t_k}\|_2^2]\,\mathrm{d}t.$$

By the $T \geq 1$ consequence of Lemma A.6,

$$\mathrm{KL}\big(p_T(\cdot \mid x) \,\|\, \mathcal{N}(0, I_{d_Y})\big) \leq C(d_Y + M_2^{\max})e^{-T}.$$

Apply Lemma A.9 with $s = t_k$ to $\mathcal{D}_{\mathrm{score}}$,

$$\mathbb{E}\big[\|\nabla_y \log p_t(Y_t \mid x) - \nabla_y \log p_{t_k}(Y_{t_k} \mid x)\|_2^2\big]$$

$$\leq 4\,\mathbb{E}\big[\|\nabla_y \log p_t(Y_t \mid x) - \nabla_y \log p_t(\alpha_{t,t_k}^{-1} Y_{t_k} \mid x)\|_2^2\big] + 2(1-\alpha_{t,t_k}^{-1})^2 \mathbb{E}[\|\nabla_y \log p_t(Y_t \mid x)\|_2^2].$$

Under (19), Lemma A.11 gives

$$\mathbb{E}\big[\|\nabla_y \log p_t(Y_t \mid x) - \nabla_y \log p_t(\alpha_{t,t_k}^{-1} Y_{t_k} \mid x)\|_2^2\big] \leq C\,\frac{d_Y^2(t_k - t)}{\vartheta_t^4}.$$

Since

$$(1 - \alpha_{t,t_k}^{-1})^2 \lesssim (t_k - t)^2,$$

and $\mathbb{E}[\|\nabla_y \log p_t(Y_t \mid x)\|_2^2] \lesssim d_Y/\vartheta_t^2$ for the OU perturbation, the time-discretization contribution is of lower order and is absorbed by the same bound. Therefore

$$\mathcal{D}_{\mathrm{score}} \leq C d_Y^2 \sum_{k=1}^{K} \int_{t_{k-1}}^{t_k} \frac{t_k - t}{\vartheta_t^4}\,\mathrm{d}t \leq C d_Y^2 \sum_{k=1}^{K} \frac{h_k^2}{\vartheta_{t_{k-1}}^4} = C d_Y^2 \Gamma_K.$$

For Euler–Maruyama, Lemma A.8 yields

$$\mathcal{D}_{\mathrm{EM}} \leq C d_Y \sum_{k=1}^{K} h_k^2 + C M_2^{\max} \sum_{k=1}^{K} h_k^3.$$

Since (19) holds, the term $d_Y \sum h_k^2$ is absorbed into $d_Y^2 \Gamma_K$.. Thus

$$\mathcal{D}_{\mathrm{score}} + \mathbf{1}_{\{S=\mathrm{EM}\}}\mathcal{D}_{\mathrm{EM}} \leq C_{\mathrm{disc},S}\,\Sigma_K.$$

Combining the previous displays proves (31). □

A.4. Conditional transfer-and-assembly step

We now combine the learned-to-oracle and oracle-to-target bounds. The residual $R_{p\to q}$ accounts for the mismatch between the measure used to define the score error and the measure produced by the learned-to-oracle KL chain rule. The score error in ε_s is measured under $p_{t_k}(\cdot \mid X)$, whereas the KL chain rule gives expectations under $q_{t_k}(\cdot \mid X)$. The final TV estimate follows from the triangle inequality, Pinsker's inequality, and Jensen's inequality.

The next proposition proves Proposition 3.1.

Proposition A.13 (Conditional TV bound at $t = \delta$)**.** *Assume Assumptions A.1–A.2 and $T \geq 1$. Suppose further that the stochastic Gaussian solver scope stated in Section A.1 holds and that the reverse grid satisfies (19). Then there exists a constant $C > 0$ such that*

$$\mathcal{E}_{\mathrm{TV}}(\delta) := \mathbb{E}_X\Big[\,\mathrm{TV}\big(q_\delta(\cdot \mid X), p_\delta(\cdot \mid X)\big)\Big] \leq C\,\varepsilon_s + \mathcal{R}_{p\to q} + C\sqrt{\Sigma_K} + C\sqrt{d_Y + M_2^{\max}}\,e^{-T/2}. \tag{32}$$

Proof. Fix $x \in \mathcal{X}$. After proving the conditional bound, we average over X.

Step 1 (Triangle decomposition). Apply the triangle inequality.

$$\mathrm{TV}(q_\delta, p_\delta) \leq \mathrm{TV}(q_\delta, q_\delta^\star) + \mathrm{TV}(q_\delta^\star, p_\delta).$$

Step 2 (Learned vs. oracle chain). Apply Pinsker's inequality and Jensen's inequality to Proposition A.5.

$$\mathbb{E}_X\big[\,\mathrm{TV}(q_\delta(\cdot \mid X), q_\delta^\star(\cdot \mid X))\big] \leq \mathbb{E}_X\Big[\sqrt{\tfrac{1}{2}\,\mathrm{KL}(q_\delta(\cdot \mid X)\|q_\delta^\star(\cdot \mid X))}\Big] \leq \frac{1}{2}\,\varepsilon_s + \mathcal{R}_{p\to q}.$$

Step 3 (Oracle chain vs. early-stopped target). Apply Pinsker's inequality and Jensen's inequality to Proposition A.12.

$$\begin{aligned}\mathbb{E}_X\big[\,\mathrm{TV}(q_\delta^\star(\cdot \mid X), p_\delta(\cdot \mid X))\big] &\leq \mathbb{E}_X\Big[\sqrt{\tfrac{1}{2}\,\mathrm{KL}(p_\delta(\cdot \mid X)\|q_\delta^\star(\cdot \mid X))}\Big] \\ &\leq \sqrt{\tfrac{1}{2}\,C_{\mathrm{disc},S}\,\Sigma_K} + \sqrt{\tfrac{1}{2}\,C_{\mathrm{init}}(d_Y + M_2^{\max})e^{-T}} \\ &\leq C\sqrt{\Sigma_K} + C\sqrt{d_Y + M_2^{\max}}\,e^{-T/2}.\end{aligned}$$

Step 4 (Combine). Combining Steps 1–3 yields (32). □

A.5. Scope: deterministic solvers and PF-ODE

The TV argument above is tailored to stochastic reverse-SDE discretizations with nondegenerate Gaussian one-step transitions. Deterministic probability-flow ODE solvers, including the PF-ODE solver used in our experiments, have $\sigma_k^2 \equiv 0$. They therefore fall outside the stochastic Gaussian solver scope stated in Section A.1, and the learned-to-oracle Gaussian-kernel KL comparison in Section A.2 does not apply. A corresponding distribution-level bound for such deterministic solvers would require separate control of the induced transport maps, for example through Jacobian, stability, or density-ratio bounds. Related non-asymptotic distributional analyses for deterministic solvers are available in the literature. See, e.g., Li et al. (2023). Thus PF-ODE is used in this paper as a contrasting finite-step solver in the experiments, but it is not covered by Proposition 3.1.

A.6. Finite-sample consequence for fixed stochastic solvers

This subsection records a finite-sample consequence of Proposition A.13 for a fixed stochastic solver. The solver-indexed quantities and all additional transfer, score-rate, and discretization inputs are stated inside the corollary where they are used.

Corollary A.14 (Finite-sample TV consequence for a fixed stochastic solver)**.** *Fix a stochastic reverse-SDE solver family $S \in \{\mathrm{EM}, \mathrm{ExpInt}\}$. For this solver, define*

$$\Delta_k(y, x) := \|s_\theta(t_k, y, x) - s^\star(t_k, y, x)\|_2^2,$$

$$\varepsilon_{s,S}^2 := \mathbb{E}_X\left[\sum_{k=1}^{K} \|B_k^{(S)}\|_F^2 (\sigma_k^{(S)})^{-2}\, \mathbb{E}_{Y\sim p_{t_k}(\cdot|X)}\big[\Delta_k\big]\right], \tag{33}$$

$$R_{p\to q,S}^2 := \mathbb{E}_X\left[\left(\frac{1}{2}\sum_{k=1}^{K} \|B_k^{(S)}\|_F^2 (\sigma_k^{(S)})^{-2}\Big(\mathbb{E}_{q_{t_k}^{(S)}(\cdot|X)}[\Delta_k] - \mathbb{E}_{p_{t_k}(\cdot|X)}[\Delta_k]\Big)\right)_+\right], \tag{34}$$

and

$$\Sigma_{K,S} := \begin{cases} d_Y^2 \displaystyle\sum_{k=1}^{K} h_k^2\, \vartheta_{t_{k-1}}^{-4} + M_2^{\max} \sum_{k=1}^{K} h_k^3, & S = \mathrm{EM}, \\ d_Y^2 \displaystyle\sum_{k=1}^{K} h_k^2\, \vartheta_{t_{k-1}}^{-4}, & S = \mathrm{ExpInt}. \end{cases} \tag{35}$$

Let

$$\mathcal{E}_{\mathrm{TV},S}(\delta) := \mathbb{E}_X\Big[\,\mathrm{TV}\big(q_\delta^{(S)}(\cdot \mid X), p_\delta(\cdot \mid X)\big)\Big].$$

Suppose $\varepsilon_{s,S}^2 < \infty$, Assumption A.2 holds, $T \geq 1$, the stochastic Gaussian solver scope stated in Section A.1 holds for S, and the reverse grid satisfies (19). Suppose further that there exists $\kappa_S \geq 1$ such that, for every $k = 1, \dots, K$ and $x \in \mathcal{X}$,

$$q_{t_k}^{(S)}(\cdot \mid x) \ll p_{t_k}(\cdot \mid x), \qquad \left\| \frac{\mathrm{d}q_{t_k}^{(S)}(\cdot \mid x)}{\mathrm{d}p_{t_k}(\cdot \mid x)} \right\|_{L^\infty(p_{t_k}(\cdot|x))} \leq \kappa_S. \tag{36}$$

Assume also that the solver discretization obeys

$$\sqrt{\Sigma_{K,S}} \leq C_{\mathrm{disc},S}\, d_Y\, \frac{\log(1/\delta) + T}{\sqrt{N}}\, r_{S,N}, \tag{37}$$

and that an external score-learning analysis gives, with probability at least $1 - 2e^{-u}$,

$$\varepsilon_{s,S}^2 \leq C_{\mathrm{score}}\left[P_s^{-\frac{2\beta}{1+d_X+d_Y}} (\log P_s)^{2c_1} \;+\; \sqrt{\frac{P_s \log n + u}{n}} \;+\; \mathrm{Tail}_R\right]. \tag{38}$$

Then, with probability at least $1 - 2e^{-u}$,

$$\begin{aligned} \mathcal{E}_{\mathrm{TV},S}(\delta) \leq C_\kappa\, C_{\mathrm{score}}^{1/2}\left[P_s^{-\frac{\beta}{1+d_X+d_Y}} (\log P_s)^{c_1} \;+\; \left(\frac{P_s \log n + u}{n}\right)^{1/4} \;+\; \mathrm{Tail}_R^{1/2}\right] \\ + C\, C_{\mathrm{disc},S}\, d_Y\, \frac{\log(1/\delta) + T}{\sqrt{N}}\, r_{S,N} + C_{\mathrm{init}} \sqrt{d_Y + M_2^{\max}}\, e^{-T/2}, \end{aligned} \tag{39}$$

where $C_\kappa := C + \sqrt{(\kappa_S - 1)_+/2}$.

Proof. Apply Proposition A.13 to the fixed solver family S. With the quantities defined in (33)–(35),

$$\mathcal{E}_{\mathrm{TV},S}(\delta) \leq C\varepsilon_{s,S} + R_{p\to q,S} + C\sqrt{\Sigma_{K,S}} + C\sqrt{d_Y + M_2^{\max}}\, e^{-T/2}. \tag{40}$$

The density-ratio condition (36) gives, for each k and x,

$$\mathbb{E}_{q_{t_k}^{(S)}(\cdot|x)}[\Delta_k] \leq \kappa_S\, \mathbb{E}_{p_{t_k}(\cdot|x)}[\Delta_k].$$

Substituting this inequality into (34) yields

$$R_{p\to q,S}^2 \leq \frac{(\kappa_S - 1)_+}{2}\, \varepsilon_{s,S}^2. \tag{41}$$

Combining (40) and (41) gives

$$\mathcal{E}_{\mathrm{TV},S}(\delta) \leq C_\kappa \varepsilon_{s,S} + C\sqrt{\Sigma_{K,S}} + C\sqrt{d_Y + M_2^{\max}}\, e^{-T/2}.$$

The score term is controlled by (38) and $\sqrt{a+b+c} \leq \sqrt{a} + \sqrt{b} + \sqrt{c}$. The discretization term is controlled by (37). The initialization term is unchanged and is absorbed into the constant C_{init}. Combining the three bounds proves (39). □

B. Conditional Variance Transport and Variance-Gap Analysis

This appendix gives the technical details for the one-step variance recursions used in Proposition 3.2. Appendix A is restricted to stochastic reverse-SDE discretizations. In contrast, this appendix focuses on the one-dimensional response setting $d_Y = 1$ and derives the local variance recursions used in the solver-matched variance-gap analysis. Sections B.1–B.4 give the detailed affine-recursion proof for the stochastic reverse-SDE solvers EM and ExpInt, whose updates have the form

$$A_k Y_{t_k} + B_k s_\theta(t_k, Y_{t_k}, x) + \sigma_k \xi_k.$$

We then record the corresponding zero-injection local-map recursion for PF-ODE with Heun. The final variance-gap proposition only requires that the learned-score and solver-matched oracle runs satisfy the stated one-step variance recursions

For fixed $x \in \mathcal{X}$, define

$$m_{t_k}(x) := \mathbb{E}[Y_{t_k} \mid X = x], \qquad v_{t_k}(x) := \mathrm{Var}(Y_{t_k} \mid X = x),$$

and, for the solver-matched true-score run,

$$m^\star_{t_k}(x) := \mathbb{E}[Y^\star_{t_k} \mid X = x], \qquad v^\star_{t_k}(x) := \mathrm{Var}(Y^\star_{t_k} \mid X = x).$$

B.1. Affine one-step reverse recursion for EM and ExpInt

Assumption B.1 (Local score regularity on bounded neighborhoods)**.** Fix $x \in \mathcal{X}$ and a reverse grid $\delta = t_0 < t_1 < \cdots < t_K = T$. For each k, assume that $y \mapsto s_\theta(t_k, y, x)$ is differentiable at $m_{t_k}(x)$. When a quadratic Taylor-remainder control is invoked, assume that for every $R < \infty$ there exists $L_s(R, x) < \infty$, uniformly over k, such that

$$\big|s_y(t_k, y, x) - s_y(t_k, z, x)\big| \le L_s(R, x)|y - z|$$

whenever

$$|y - m_{t_k}(x)| \vee |z - m_{t_k}(x)| \le R.$$

Assumption B.2 (Transport typical event)**.** There exist nonnegative numbers $\eta_1, \ldots, \eta_K$, a constant $c_\eta < \infty$, and a probability level $\eta_J \in (0, 1)$ such that

$$\mathcal{T}_\nabla(x) := \Big\{ |s_y(t_k, Y_{t_k}, x)| \le 1 + \eta_k \ \forall k,\ \sum_{k=1}^K \eta_k \le c_\eta \Big\}$$

satisfies

$$\mathbb{P}\big(\mathcal{T}_\nabla(x)\big) \ge 1 - \eta_J.$$

Assumption B.3 (Curvature moment control)**.** For each $k = 1, \ldots, K$, define

$$\rho_k := s_\theta(t_k, Y_{t_k}, x) - s_\theta\big(t_k, m_{t_k}(x), x\big) - s_y\big(t_k, m_{t_k}(x), x\big)\big(Y_{t_k} - m_{t_k}(x)\big).$$

Assume that the following conditional moments are finite:

$$\mathrm{Var}(\rho_k \mid X = x) < \infty, \qquad \big|\mathrm{Cov}\big(Y_{t_k} - m_{t_k}(x),\ \rho_k \mid X = x\big)\big| < \infty.$$

If the solver-matched oracle chain driven by $s^\star$ is also considered, define

$$\rho^\star_k := s^\star(t_k, Y^\star_{t_k}, x) - s^\star\big(t_k, m^\star_{t_k}(x), x\big) - s^\star_y\big(t_k, m^\star_{t_k}(x), x\big)\big(Y^\star_{t_k} - m^\star_{t_k}(x)\big),$$

and assume

$$\mathrm{Var}(\rho^\star_k \mid X = x) < \infty, \qquad \big|\mathrm{Cov}\big(Y^\star_{t_k} - m^\star_{t_k}(x),\ \rho^\star_k \mid X = x\big)\big| < \infty$$

for every $k = 1, \ldots, K$.

Proposition B.4 (One-step conditional mean and variance recursion)**.** *Fix $x \in \mathcal{X}$ and consider the scalar-target case $d_Y = 1$. For the affine stochastic solvers EM and ExpInt, let the one-step reverse update at grid point $t_k \to t_{k-1}$ be*

$$Y_{t_{k-1}} = A_k\, Y_{t_k} + B_k\, s_\theta(t_k, Y_{t_k}, x) + \sigma_k\, \xi_k, \qquad \xi_k \sim \mathcal{N}(0,1), \tag{42}$$

where ξ_k is independent of (Y_{t_k}, X). Then

$$m_{t_{k-1}}(x) = A_k\, m_{t_k}(x) + B_k\, s_\theta(t_k, m_{t_k}(x), x) + B_k\, \mathbb{E}[\rho_k \mid X = x], \tag{43}$$

and

$$v_{t_{k-1}}(x) = J_k(x)^2\, v_{t_k}(x) + \sigma_k^2 + r_k(x), \tag{44}$$

where

$$J_k(x) := A_k + B_k\, s_y(t_k, m_{t_k}(x), x),$$

and

$$r_k(x) := B_k^2\, \mathrm{Var}(\rho_k \mid X = x) + 2J_k(x)B_k\, \mathrm{Cov}\big(Y_{t_k} - m_{t_k}(x),\, \rho_k \mid X = x\big). \tag{45}$$

Proof. Suppress the explicit dependence on x to lighten notation. By the definition of ρ_k,

$$s_\theta(t_k, Y_{t_k}, x) = s_\theta(t_k, m_{t_k}, x) + s_y(t_k, m_{t_k}, x)\,(Y_{t_k} - m_{t_k}) + \rho_k.$$

Substituting this into (42) gives

$$Y_{t_{k-1}} = A_k Y_{t_k} + B_k s_\theta(t_k, m_{t_k}, x) + B_k s_y(t_k, m_{t_k}, x)(Y_{t_k} - m_{t_k}) + B_k \rho_k + \sigma_k\, \xi_k.$$

Taking conditional expectation given $X = x$ yields (43). Subtracting (43) from the update gives

$$Y_{t_{k-1}} - m_{t_{k-1}} = J_k(Y_{t_k} - m_{t_k}) + B_k\big(\rho_k - \mathbb{E}[\rho_k \mid X = x]\big) + \sigma_k\, \xi_k.$$

Now take conditional variance given $X = x$. Since ξ_k is independent of (Y_{t_k}, X) and has variance 1,

$$v_{t_{k-1}} = J_k^2\, v_{t_k} + B_k^2\, \mathrm{Var}(\rho_k \mid X = x) + 2J_k B_k\, \mathrm{Cov}\big(Y_{t_k} - m_{t_k},\, \rho_k \mid X = x\big) + \sigma_k^2,$$

which is exactly (44)–(45). □

Lemma B.5 (Local quadratic Taylor remainder)**.** *Under Assumption B.1, fix a finite $R > 0$ and define*

$$E_{k,R}(x) := \{|Y_{t_k} - m_{t_k}(x)| \le R\}.$$

On $E_{k,R}(x)$,

$$|\rho_k| \le \frac{L_s(R,x)}{2}\,(Y_{t_k} - m_{t_k}(x))^2.$$

Proof. Let $m := m_{t_k}(x)$, $Y := Y_{t_k}$, and $\Delta := Y - m$. On $E_{k,R}(x)$, $m + \tau\Delta \in I_{k,R}(x)$ for every $\tau \in [0,1]$. By the integral form of the Taylor remainder,

$$\rho_k = \int_0^1 \Big(s_y\big(t_k, m + \tau\Delta, x\big) - s_y(t_k, m, x)\Big)\Delta\, \mathrm{d}\tau.$$

Assumption B.1 gives, on this event,

$$\Big|s_y\big(t_k, m + \tau\Delta, x\big) - s_y(t_k, m, x)\Big| \le L_s(R,x)\,\tau|\Delta|.$$

Therefore

$$|\rho_k| \le \int_0^1 L_s(R,x)\,\tau|\Delta|^2\, \mathrm{d}\tau = \frac{L_s(R,x)}{2}\Delta^2.$$

□

Remark B.6 (Role of the curvature remainder r_k). The variance recursion in Proposition B.4 is an exact conditional-moment identity once ρ_k is defined and the finite-moment conditions in Assumption B.3 hold. Lemma B.5 is only a local control statement. On $E_{k,R}(x)$, with $\Delta_k := Y_{t_k} - m_{t_k}(x)$, it implies

$$\rho_k^2 \le \frac{L_s(R,x)^2}{4}\Delta_k^4, \qquad |\Delta_k \rho_k| \le \frac{L_s(R,x)}{2}|\Delta_k|^3.$$

Thus the curvature terms in r_k are locally controlled by third and fourth centered moments, whereas the leading transport term $J_k(x)^2 v_{t_k}(x)$ depends on the second centered moment. Unconditional bounds require the finite-moment condition in Assumption B.3, or additional tail control outside $E_{k,R}(x)$.

B.2. Typical set construction

The following typical events are used only as high-probability controls for local transport factors and solver trajectories. The conditional variance recursions in Appendix B.4 and the variance-gap identity in Appendix B.5 are deterministic conditional-moment identities once the corresponding moments are finite.

Transport typical event. Recall

$$\mathcal{T}_\nabla(x) := \Big\{ |s_y(t_k, Y_{t_k}, x)| \le 1 + \eta_k \ \forall k,\ \sum_{k=1}^K \eta_k \le c_\eta \Big\},$$

where the parameters $(\eta_k)_{k=1}^K$ and c_η are those appearing in Assumption B.2. Under Assumption B.2,

$$\mathbb{P}(\mathcal{T}_\nabla(x)) \ge 1 - \eta_J.$$

Noise typical event. Let $\xi_1, \dots, \xi_K$ be the i.i.d. Gaussian innovations used by the stochastic reverse solver. Define

$$\mathcal{T}_\xi := \Big\{ \max_{1\le k\le K} |\xi_k| \le r_\xi \Big\}.$$

Lemma B.7 (Gaussian noise typical set probability)**.** *Let $\xi_1, \dots, \xi_K$ be i.i.d. $\mathcal{N}(0,1)$. Then for any $\eta_\xi \in (0,1)$, choosing*

$$r_\xi := \sqrt{2\log\Big(\frac{2K}{\eta_\xi}\Big)}$$

yields $\mathbb{P}(\mathcal{T}_\xi) \ge 1 - \eta_\xi$.

Proof. By the union bound and the Gaussian tail $\mathbb{P}(|\xi| \ge r) \le 2e^{-r^2/2}$,

$$\mathbb{P}(\mathcal{T}_\xi^c) = \mathbb{P}\Big(\max_{1\le k\le K} |\xi_k| > r_\xi \Big) \le \sum_{k=1}^K \mathbb{P}(|\xi_k| > r_\xi) \le 2Ke^{-r_\xi^2/2} = \eta_\xi.$$

□

Reverse-solver typical event. Define

$$\mathcal{T}_{\mathrm{rev}}(x) := \mathcal{T}_\nabla(x) \cap \mathcal{T}_\xi.$$

Lemma B.8 (Reverse-solver typical set probability)**.** *Under Assumption B.2 and the setting of Lemma B.7, for any $\eta_\xi \in (0,1)$ we have*

$$\mathbb{P}\big(\mathcal{T}_{\mathrm{rev}}(x)\big) \ge 1 - (\eta_J + \eta_\xi).$$

Proof. By the union bound,

$$\mathbb{P}\big(\mathcal{T}_{\mathrm{rev}}(x)^c\big) = \mathbb{P}\big(\mathcal{T}_\nabla(x)^c \cup \mathcal{T}_\xi^c\big) \le \mathbb{P}\big(\mathcal{T}_\nabla(x)^c\big) + \mathbb{P}\big(\mathcal{T}_\xi^c\big) \le \eta_J + \eta_\xi.$$

□

B.3. Transport-product stability

For the affine stochastic solvers EM and ExpInt, define the local transport factors

$$J_k(x) := A_k + B_k\, s_y(t_k, m_{t_k}(x), x), \qquad J_k^\star(x) := A_k + B_k\, s_y^\star(t_k, m_{t_k}^\star(x), x),$$

and the learned-chain transport products

$$\Pi_{1:0}(x) := 1, \qquad \Pi_{1:m}(x) := \prod_{j=1}^{m} J_j(x)^2, \quad m \geq 1.$$

These products encode the cumulative transport of conditional variance along the reverse grid for EM and ExpInt. The PF-ODE Heun transport factor is defined separately below as the local derivative of the deterministic Heun update map.

B.4. One-step variance recursions for solver-specific updates

Lemma B.9 (Affine variance recursion for EM and ExpInt)**.** *Assume $d_Y = 1$, fix $x \in \mathcal{X}$, and work on the reverse grid*

$$\delta = t_0 < t_1 < \cdots < t_K = T.$$

Suppose Assumptions B.1 and B.3 hold. Then, for each $k = 1, \ldots, K$, the learned-score chain satisfies

$$v_{t_{k-1}}(x) = J_k(x)^2\, v_{t_k}(x) + \sigma_k^2 + r_k(x), \tag{46}$$

where $J_k(x)$ is as above and $r_k(x)$ is given by (45). If the analogous local differentiability and finite-moment conditions are imposed on $s^\star$ as well, then the solver-matched oracle chain satisfies

$$v_{t_{k-1}}^\star(x) = \left(J_k^\star(x)\right)^2 v_{t_k}^\star(x) + (\sigma_k^\star)^2 + r_k^\star(x), \tag{47}$$

where $r_k^\star(x)$ denotes the analogue of (45) obtained by replacing $(Y_{t_k}, m_{t_k}(x), \rho_k, J_k)$ with $(Y_{t_k}^\star, m_{t_k}^\star(x), \rho_k^\star, J_k^\star)$.

Proof. Apply Proposition B.4 to the learned-score chain. Assumption B.1, together with Lemma B.5, gives the local Taylor remainder representation. Assumption B.3 ensures that the resulting remainder-variance and covariance terms are finite for each step. For the solver-matched oracle chain, apply the same argument under the corresponding local differentiability and finite-moment conditions for $s^\star$. □

Deterministic PF-ODE with Heun. We now record the analogous local variance recursion for the deterministic PF-ODE solver used in the experiments. This result is separate from the affine recursion above: PF-ODE is not obtained by setting the Gaussian noise in EM or ExpInt to zero, but by applying a deterministic Heun update to the probability-flow ODE.

Define

$$f_\theta(t, y, x) := -\frac{1}{2} y - \frac{1}{2} s_\theta(t, y, x).$$

For a state y at time t_k, the Heun predictor is

$$\widetilde{y} = y - h_k f_\theta(t_k, y, x),$$

and the one-step PF-ODE Heun map is

$$\Phi_{k,\mathrm{PF}}^\theta(y; x) := y - \frac{h_k}{2}\left\{f_\theta(t_k, y, x) + f_\theta(t_{k-1}, \widetilde{y}, x)\right\}. \tag{48}$$

The solver-matched oracle map $\Phi_{k,\mathrm{PF}}^\star$ is defined by replacing s_θ with $s^\star$ in the same Heun update.

Assumption B.10 (Local regularity of the PF-ODE Heun map)**.** Fix $x \in \mathcal{X}$. For each $k = 1, \ldots, K$, assume that $y \mapsto \Phi_{k,\mathrm{PF}}^\theta(y; x)$ is differentiable on the line segment between $m_{t_k}(x)$ and Y_{t_k}, and that its derivative is locally Lipschitz on that segment with constant $L_\Phi < \infty$, uniformly over k. Assume the analogous condition for the solver-matched oracle map $\Phi_{k,\mathrm{PF}}^\star$ on the line segment between $m_{t_k}^\star(x)$ and $Y_{t_k}^\star$.

Define the PF-ODE transport factors

$$J_{k,\mathrm{PF}}(x) := \partial_y \Phi^{\theta}_{k,\mathrm{PF}}(m_{t_k}(x); x), \qquad J^{\star}_{k,\mathrm{PF}}(x) := \partial_y \Phi^{\star}_{k,\mathrm{PF}}(m^{\star}_{t_k}(x); x). \tag{49}$$

Writing $f_{\theta,y} := \partial_y f_\theta$, the learned Heun derivative is

$$\partial_y \Phi^{\theta}_{k,\mathrm{PF}}(y; x) = 1 - \frac{h_k}{2}\left[f_{\theta,y}(t_k, y, x) + f_{\theta,y}(t_{k-1}, \widetilde{y}, x)\left\{1 - h_k f_{\theta,y}(t_k, y, x)\right\}\right]. \tag{50}$$

Define the Heun-map linearization remainder

$$\zeta_k := \Phi^{\theta}_{k,\mathrm{PF}}(Y_{t_k}; x) - \Phi^{\theta}_{k,\mathrm{PF}}(m_{t_k}(x); x) - J_{k,\mathrm{PF}}(x)\big(Y_{t_k} - m_{t_k}(x)\big), \tag{51}$$

and define $\zeta^{\star}_k$ analogously for the oracle Heun map.

Lemma B.11 (Quadratic Heun-map remainder bound)**.** *Under Assumption B.10,*

$$|\zeta_k| \le \frac{L_\Phi}{2}\big(Y_{t_k} - m_{t_k}(x)\big)^2.$$

The same bound holds for $\zeta^{\star}_k$ with the oracle map.

Proof. Let $m := m_{t_k}(x)$, $Y := Y_{t_k}$, and $\Phi(y) := \Phi^{\theta}_{k,\mathrm{PF}}(y; x)$. By the integral form of the Taylor remainder,

$$\zeta_k = \int_0^1 \left[\partial_y \Phi(m + \tau(Y - m)) - \partial_y \Phi(m)\right](Y - m)\, \mathrm{d}\tau.$$

Using the local Lipschitz condition in Assumption B.10,

$$|\partial_y \Phi(m + \tau(Y - m)) - \partial_y \Phi(m)| \le L_\Phi \tau |Y - m|.$$

Therefore,

$$|\zeta_k| \le \int_0^1 L_\Phi \tau |Y - m|^2\, \mathrm{d}\tau = \frac{L_\Phi}{2}(Y - m)^2.$$

The oracle statement follows by the same argument applied to $\Phi^{\star}_{k,\mathrm{PF}}$. □

Assumption B.12 (PF-ODE Heun remainder moment control)**.** For each $k = 1, \dots, K$, assume

$$\mathrm{Var}(\zeta_k \mid X = x) < \infty, \qquad \big|\mathrm{Cov}\big(Y_{t_k} - m_{t_k}(x), \zeta_k \mid X = x\big)\big| < \infty,$$

and impose the analogous finite-moment conditions on $\zeta^{\star}_k$.

Lemma B.13 (PF-ODE Heun variance recursion)**.** *Assume $d_Y = 1$, fix $x \in \mathcal{X}$, and suppose Assumptions B.10–B.12 hold. Then the learned-score PF-ODE Heun chain satisfies*

$$v_{t_{k-1}}(x) = J_{k,\mathrm{PF}}(x)^2 v_{t_k}(x) + r_{k,\mathrm{PF}}(x), \tag{52}$$

where

$$r_{k,\mathrm{PF}}(x) := \mathrm{Var}(\zeta_k \mid X = x) + 2J_{k,\mathrm{PF}}(x)\,\mathrm{Cov}\big(Y_{t_k} - m_{t_k}(x), \zeta_k \mid X = x\big). \tag{53}$$

The solver-matched oracle PF-ODE Heun chain satisfies

$$v^{\star}_{t_{k-1}}(x) = \big(J^{\star}_{k,\mathrm{PF}}(x)\big)^2 v^{\star}_{t_k}(x) + r^{\star}_{k,\mathrm{PF}}(x). \tag{54}$$

Equivalently, PF-ODE Heun fits the variance-gap template with

$$\sigma^2_{k,\mathrm{PF}} = (\sigma^{\star}_{k,\mathrm{PF}})^2 = 0.$$

Proof. Suppress the explicit dependence on x. By the definition of ζ_k,

$$\Phi^{\theta}_{k,\mathrm{PF}}(Y_{t_k}; x) = \Phi^{\theta}_{k,\mathrm{PF}}(m_{t_k}; x) + J_{k,\mathrm{PF}}(Y_{t_k} - m_{t_k}) + \zeta_k.$$

Since the PF-ODE Heun update is deterministic,

$$Y_{t_{k-1}} = \Phi^{\theta}_{k,\mathrm{PF}}(Y_{t_k}; x).$$

Taking conditional expectation gives

$$m_{t_{k-1}} = \Phi^{\theta}_{k,\mathrm{PF}}(m_{t_k}; x) + \mathbb{E}[\zeta_k \mid X = x].$$

Subtracting this mean recursion from the update yields

$$Y_{t_{k-1}} - m_{t_{k-1}} = J_{k,\mathrm{PF}}(Y_{t_k} - m_{t_k}) + \zeta_k - \mathbb{E}[\zeta_k \mid X = x].$$

Taking conditional variance gives

$$v_{t_{k-1}} = J^2_{k,\mathrm{PF}} v_{t_k} + \mathrm{Var}(\zeta_k \mid X = x) + 2J_{k,\mathrm{PF}} \,\mathrm{Cov}(Y_{t_k} - m_{t_k}, \zeta_k \mid X = x),$$

which proves (52)–(53). The oracle recursion follows by applying the same argument to $\Phi^{\star}_{k,\mathrm{PF}}$. Since the PF-ODE Heun update has no additive Gaussian innovation, the injected-variance term is zero. □

B.5. Variance-gap theory

This section provides the full proof of the main variance-fidelity result, Proposition 3.2, together with the recursion-unrolling lemmas that the proof calls.

B.5.1. Unrolling lemmas for transported recursions

Lemma B.14 (Unrolling an inhomogeneous transport recursion)**.** *Fix $x \in \mathcal{X}$ and $K \in \mathbb{N}$. Let $\{a_k(x)\}_{k=1}^K \subset \mathbb{R}$ and $\{b_k(x)\}_{k=1}^K \subset \mathbb{R}$ be arbitrary scalars, and let $\{u_k(x)\}_{k=0}^K$ be a sequence satisfying, for each $k = 1, \dots, K$,*

$$u_{k-1}(x) = a_k(x)\, u_k(x) + b_k(x). \tag{55}$$

Define

$$A_{1:0}(x) := 1, \qquad A_{1:m}(x) := \prod_{j=1}^{m} a_j(x), \quad m \geq 1.$$

Then, for every $m \in \{0, 1, \dots, K\}$,

$$u_0(x) = A_{1:m}(x)\, u_m(x) + \sum_{k=1}^{m} A_{1:k-1}(x)\, b_k(x). \tag{56}$$

Proof. We prove (56) by induction on m. For $m = 0$, the identity is tautological. Assume it holds for $m - 1 \geq 0$. Using (55) at step m,

$$u_{m-1} = a_m u_m + b_m.$$

Plugging this into the induction hypothesis gives

$$u_0 = A_{1:m-1}(a_m u_m + b_m) + \sum_{k=1}^{m-1} A_{1:k-1} b_k,$$

which equals

$$A_{1:m} u_m + \sum_{k=1}^{m} A_{1:k-1} b_k.$$

□

Proposition B.15 (Generic variance-gap bound)**.** *Fix $x \in \mathcal{X}$ and assume $d_Y = 1$. Let $v_{t_k}(x)$ be the conditional variance generated by the K-step solver driven by s_θ, and let $v^\star_{t_k}(x)$ be the conditional variance generated by the same solver family driven by $s^\star$. Suppose that, for each $k = 1, \dots, K$, the learned and solver-matched oracle conditional variance recursions are well-defined and satisfy*

$$\begin{aligned} v_{t_{k-1}}(x) &= J_k(x)^2\, v_{t_k}(x) + \sigma_k^2 + r_k(x), \\ v^\star_{t_{k-1}}(x) &= \big(J^\star_k(x)\big)^2\, v^\star_{t_k}(x) + \big(\sigma^\star_k\big)^2 + r^\star_k(x), \end{aligned} \tag{57}$$

where $J_k(x), J^\star_k(x) \in \mathbb{R}$ are local variance-transport factors, $\sigma_k^2, (\sigma^\star_k)^2 \geq 0$ are per-step injected variances, and $r_k(x), r^\star_k(x)$ collect higher-order effects from the local linearization used to obtain the one-step variance recursion. Define

$$\Pi_{1:0}(x) := 1, \qquad \Pi_{1:m}(x) := \prod_{j=1}^{m} J_j(x)^2, \quad m \geq 1. \tag{58}$$

Then,

$$\begin{aligned} \big|v_{t_0}(x) - v^\star_{t_0}(x)\big| &\leq \Pi_{1:K}(x)\, \big|v_{t_K}(x) - v^\star_{t_K}(x)\big| \\ &+ \sum_{k=1}^{K} \Pi_{1:k-1}(x)\, \big|J_k(x)^2 - \big(J^\star_k(x)\big)^2\big|\, v^\star_{t_k}(x) \\ &+ \sum_{k=1}^{K} \Pi_{1:k-1}(x)\, \big|\sigma_k^2 - (\sigma^\star_k)^2\big| \\ &+ \sum_{k=1}^{K} \Pi_{1:k-1}(x)\, \big|r_k(x) - r^\star_k(x)\big|. \end{aligned} \tag{59}$$

In particular, if the variance initialization at $t = T$ matches, i.e. $v_{t_K}(x) = v^\star_{t_K}(x)$, then the first term vanishes.

Proof. Starting from the conditional variance recursions in (57), suppress the explicit (x) dependence to lighten notation.

Step 1: subtract the one-step variance recursions. From (57),

$$v_{t_{k-1}} - v^\star_{t_{k-1}} = J_k^2 v_{t_k} - (J^\star_k)^2 v^\star_{t_k} + \big(\sigma_k^2 - (\sigma^\star_k)^2\big) + \big(r_k - r^\star_k\big).$$

Add and subtract $J_k^2 v^\star_{t_k}$ to isolate the transport mismatch:

$$J_k^2 v_{t_k} - (J^\star_k)^2 v^\star_{t_k} = J_k^2 (v_{t_k} - v^\star_{t_k}) + \big(J_k^2 - (J^\star_k)^2\big)\, v^\star_{t_k}.$$

Therefore,

$$\Delta_{k-1} = J_k^2\, \Delta_k + \big(J_k^2 - (J^\star_k)^2\big)\, v^\star_{t_k} + \big(\sigma_k^2 - (\sigma^\star_k)^2\big) + \big(r_k - r^\star_k\big), \tag{60}$$

where $\Delta_k := v_{t_k} - v^\star_{t_k}$.

Step 2: unroll the difference recursion. Equation (60) has the affine form in Lemma B.14 with transport coefficient $a_k := J_k^2$ and forcing term

$$b_k := \big(J_k^2 - (J^\star_k)^2\big)\, v^\star_{t_k} + \big(\sigma_k^2 - (\sigma^\star_k)^2\big) + \big(r_k - r^\star_k\big).$$

Applying Lemma B.14 with $u_k = \Delta_k$ yields

$$\Delta_0 = \Big(\prod_{j=1}^{K} J_j^2\Big)\Delta_K + \sum_{k=1}^{K} \Big(\prod_{j=1}^{k-1} J_j^2\Big)\, b_k.$$

Using the definition (58) of $\Pi_{1:m}$, this becomes

$$\Delta_0 = \Pi_{1:K}\, \Delta_K + \sum_{k=1}^{K} \Pi_{1:k-1} \Big(\big(J_k^2 - (J^\star_k)^2\big)\, v^\star_{t_k} + \big(\sigma_k^2 - (\sigma^\star_k)^2\big) + \big(r_k - r^\star_k\big)\Big).$$

Step 3: take absolute values and separate contributions. Taking absolute values and applying the triangle inequality term-by-term gives

$$|\Delta_0| \le \Pi_{1:K}\,|\Delta_K| + \sum_{k=1}^{K} \Pi_{1:k-1}\left|J_k^2 - (J_k^\star)^2\right| |v_{t_k}^\star| + \sum_{k=1}^{K} \Pi_{1:k-1}\left|\sigma_k^2 - (\sigma_k^\star)^2\right| + \sum_{k=1}^{K} \Pi_{1:k-1}\left|r_k - r_k^\star\right|.$$

Since $v_{t_k}^\star = \mathrm{Var}(\cdot \mid x) \ge 0$, we may replace $|v_{t_k}^\star|$ by $v_{t_k}^\star$. Reintroducing (x) yields exactly (59). □

Application to Proposition 3.2. The one-step recursions assumed in Proposition 3.2 are supplied by Lemma B.9 for EM and ExpInt and by Lemma B.13 for PF-ODE, with $\sigma_k^2 = (\sigma_k^\star)^2 = 0$ in the PF-ODE case. Applying Proposition B.15 gives the claimed solver-matched variance-gap bound.

C. Model and Algorithms

C.1. Score network definition

We use a time-conditioned ReLU network for the score interface. The network outputs a scaled-score head $\widehat{\epsilon}_\theta(t, y, x)$, which is passed to the reverse samplers as

$$s_\theta(t, y, x) = \frac{\widehat{\epsilon}_\theta(t, y, x)}{\vartheta_t}, \qquad \vartheta_t = \sqrt{1 - e^{-t}}.$$

The head is trained under the negative-noise DSM convention.

Embeddings. The ReLU-DNN score network conditions on diffusion time through the OU noise level. Given t, we encode $\log \vartheta_t$ using dyadic Fourier features with $F_\sigma = 8$ frequencies. The covariate x, response y, and noise-level embedding are passed through ReLU encoders to obtain h_x, h_y, and h_t.

Backbone and output. The network concatenates h_t, h_x, h_y, and $h_t \odot h_x$, then applies a ReLU trunk followed by a linear output layer. The hidden width, depth, dropout rate, and F_σ are reported in the hyperparameter table.

C.2. Finite-step reverse sampling pseudocode

D. Experimental Details

D.1. Experimental protocol and reproducibility

Data splits. We use fixed train/validation/test splits throughout. For synthetic datasets, we generate i.i.d. samples and use $N_{\text{train}} = 45{,}000$, $N_{\text{val}} = 2{,}500$, and $N_{\text{test}} = 2{,}500$. Hyperparameters are selected by minimizing the validation denoising score matching (DSM) loss, and all reported metrics are computed on the test split.

Standardization and unit conversion. Inputs and targets are standardized using training-set statistics, and metrics are reported in the original target units. Standardization improves numerical conditioning and keeps score evaluations in a stable scale regime, consistent with the regularity conditions assumed for the reverse dynamics in the main text. We compute $\mu_y = \frac{1}{N_{\text{train}}} \sum_{i=1}^{N_{\text{train}}} y_i$ and $\sigma_y = \sqrt{\frac{1}{N_{\text{train}}} \sum_{i=1}^{N_{\text{train}}} (y_i - \mu_y)^2}$, transform targets by $y^{\text{norm}} = (y - \mu_y)/\sigma_y$, run sampling in normalized space, and map samples back via $y^{\text{orig}} = \sigma_y\, y^{\text{norm}} + \mu_y$. Predictive variance rescales as

$$\widehat{v}^{\text{orig}}(x) = \widehat{v}^{\text{norm}}(x)\, \sigma_y^2.$$

D.2. Datasets

Synthetic heteroskedastic regression. We evaluate four one-dimensional heteroskedastic designs $y = \mu(x) + \sigma(x)\varepsilon$ with $\varepsilon \sim \mathcal{N}(0, 1)$, adapted from standard benchmarks (Skafte et al., 2019; Wong-Toi et al., 2024). These settings provide a known data variance target $v^{\text{data}}(x) = \sigma^2(x)$, which enables direct evaluation of endpoint variance fidelity.

Algorithm 1 Finite-step conditional sampling

Input: conditioning input x, reverse grid $t_K = T > \cdots > t_0 = \delta$, solver $S \in \{\mathrm{EM}, \mathrm{ExpInt}, \text{PF-ODE}\}$, number of samples M
Output: $\{y_{t_0}^{(m)}\}_{m=1}^{M}$, approximate samples from $q_{\delta,S}(\cdot \mid x)$
Initialize terminal samples $\{y_{t_K}^{(m)}\}_{m=1}^{M} \sim \mathcal{N}(0, I_{d_Y})$
for $k = K, K-1, \ldots, 1$ **do**
  for $m = 1, 2, \ldots, M$ **do**
    if $S = \mathrm{EM}$ **then**
      Draw $z_k^{(m)} \sim \mathcal{N}(0, I_{d_Y})$
      $g \leftarrow s_\theta(t_k, y_{t_k}^{(m)}, x)$
      $y_{t_{k-1}}^{(m)} \leftarrow y_{t_k}^{(m)} + h_k \left\{ \frac{1}{2} y_{t_k}^{(m)} + g \right\} + \sqrt{h_k}\, z_k^{(m)}$
    else if $S = \mathrm{ExpInt}$ **then**
      Draw $z_k^{(m)} \sim \mathcal{N}(0, I_{d_Y})$
      $g \leftarrow s_\theta(t_k, y_{t_k}^{(m)}, x)$
      $y_{t_{k-1}}^{(m)} \leftarrow e^{h_k/2} y_{t_k}^{(m)} + 2\big(e^{h_k/2} - 1\big) g + \sqrt{e^{h_k} - 1}\, z_k^{(m)}$
    else if $S = \text{PF-ODE}$ **then**
      $f_k \leftarrow -\frac{1}{2} y_{t_k}^{(m)} - \frac{1}{2} s_\theta(t_k, y_{t_k}^{(m)}, x)$
      $\widetilde{y} \leftarrow y_{t_k}^{(m)} - h_k f_k$
      $f_{k-1} \leftarrow -\frac{1}{2} \widetilde{y} - \frac{1}{2} s_\theta(t_{k-1}, \widetilde{y}, x)$
      $y_{t_{k-1}}^{(m)} \leftarrow y_{t_k}^{(m)} - \frac{h_k}{2} \{ f_k + f_{k-1} \}$
    end if
  end for
end for
return $\{y_{t_0}^{(m)}\}_{m=1}^{M}$

Algorithm 2 Stepwise variance and score-covariance summaries

Input: test inputs $\{x_n\}_{n=1}^{N}$, recorded times $\{t_r\}_{r=0}^{R-1}$, stored trajectories $\{y_{t_r}^{(m,n)}\}$, score network s_θ
Output: empirical variance paths $\widehat{v}_{r,n}$, score-variance paths $\widehat{u}_{r,n}$, and score–state covariance paths $\widehat{c}_{r,n}$
for $r = 0, 1, \ldots, R-1$ **do**
  for $n = 1, 2, \ldots, N$ **do**
    $\widehat{v}_{r,n} \leftarrow \mathrm{Var}_m\big(y_{t_r}^{(m,n)}\big)$
    Compute $g_{t_r}^{(m,n)} \leftarrow s_\theta(t_r, y_{t_r}^{(m,n)}, x_n)$ for all m
    $\widehat{u}_{r,n} \leftarrow \mathrm{Var}_m\big(g_{t_r}^{(m,n)}\big)$
    $\widehat{c}_{r,n} \leftarrow \mathrm{Cov}_m\big(y_{t_r}^{(m,n)}, g_{t_r}^{(m,n)}\big)$
  end for
end for
return $\{\widehat{v}_{r,n}, \widehat{u}_{r,n}, \widehat{c}_{r,n}\}_{r,n}$

Closed-form oracle score for synthetic tasks. Because each synthetic task has the form

$$Y_0 \mid X = x \sim \mathcal{N}(\mu(x), \sigma^2(x)),$$

the OU-smoothed conditional law is available in closed form. In normalized target units, let

$$\widetilde{\mu}(x) := \frac{\mu(x) - \mu_y}{\sigma_y}, \qquad \widetilde{\sigma}^2(x) := \frac{\sigma^2(x)}{\sigma_y^2},$$

where μ_y and σ_y are the training-set target mean and standard deviation used for standardization. With $a_t = e^{-t/2}$ and $\vartheta_t^2 = 1 - e^{-t}$,

$$Y_t \mid X = x \sim \mathcal{N}\big(a_t \widetilde{\mu}(x), a_t^2 \widetilde{\sigma}^2(x) + \vartheta_t^2\big)\,.$$

Table 4. **Synthetic one–dimensional regression datasets.** Each task follows $y = \mu(x) + \sigma(x)\varepsilon$ with $\varepsilon \sim \mathcal{N}(0,1)$.

Task	Data generating function	x	Noise $\sigma(x)$
Xsinx–Hetero	$y = x\sin x + \sigma(x)\varepsilon$	$x \sim \mathrm{Unif}(-5,5)$	$0.3 + 0.3\lvert x\rvert$
Sine–Hetero	$y = 2\sin(4\pi x) + \sigma(x)\varepsilon$	$x \sim \mathrm{Unif}(0,1)$	$1.25 + \sin(6\pi x)$
Curve–Hetero	$y = \log(1+x^2) + 0.5\sin(3x) + \sigma(x)\varepsilon$	$x \sim \mathrm{Unif}(0,3)$	$0.2 + 0.2x$
Cubic–Step–Hetero	$y = 0.3x^3 + \sigma(x)\varepsilon$	$x \sim \mathrm{Unif}(-3,3)$	$1\,\mathbb{I}_{\{x\le -1\}} + 3\,\mathbb{I}_{\{\lvert x\rvert<1\}} + 10\mathbb{I}_{\{1\le x\}}$

Therefore the oracle conditional score used in the solver-matched oracle chains is

$$s^\star(t,y,x) = \nabla_y \log p_t(y \mid x) = -\frac{y - a_t\widetilde{\mu}(x)}{a_t^2\widetilde{\sigma}^2(x) + \vartheta_t^2}.$$

All oracle-chain computations are performed in normalized target units and then mapped back to the original target scale for reporting variance metrics.

D.3. Models and solvers

Conditional diffusion regression and baselines. Within each dataset, we train a single score network and compare reverse-time solvers under a shared noise schedule and matched reverse-step budgets. The solver choices include EM, ExpInt, and PF-ODE, as reported in the tables and figures. We also compare against MC Dropout (Gal & Ghahramani, 2016) and a heteroscedastic deep regressor DeepHetero, implemented as a Gaussian MLP that predicts both mean and variance.

For MC Dropout, we use a depth-3, width-128 MLP with dropout probability 0.1 in hidden layers. Training uses Adam with learning rate 10^{-3}, batch size 256, and 500 epochs.

For DeepHetero, we use a single depth-3, width-128 MLP that outputs a Gaussian likelihood $\mathcal{N}(\mu_\phi(x), \sigma_\phi^2(x))$. The variance head is parameterized via $\log\sigma_\phi^2(x)$ to ensure positivity. Training minimizes the Gaussian negative log-likelihood:

$$\mathcal{L}(\phi) = \frac{1}{N}\sum_{n=1}^{N}\left[\frac{(y_n - \mu_\phi(x_n))^2}{2\,\sigma_\phi^2(x_n)} + \frac{1}{2}\log\sigma_\phi^2(x_n)\right].$$

At test time, predictive samples are drawn as $y^{(m)}(x) = \mu_\phi(x) + \sigma_\phi(x)\,\epsilon^{(m)}$ with $\epsilon^{(m)} \sim \mathcal{N}(0,1)$.

D.4. Evaluation metrics

For each test input x_n, the solver returns M predictive samples $\{y_n^{(m)}\}_{m=1}^M$.

Predictive moments. We define Monte Carlo predictive moments as

$$\widehat{m}(x_n) := \frac{1}{M}\sum_{m=1}^{M} y_n^{(m)}, \qquad \widehat{v}(x_n) := \mathrm{Var}_m\big(y_n^{(m)}\big).$$

Point metrics. Given targets $\{y_n\}_{n=1}^N$, we report RMSE computed from the predictive mean:

$$\mathrm{RMSE}_{\mathrm{mean}} := \Big(\frac{1}{N}\sum_{n=1}^{N}(\widehat{m}(x_n) - y_n)^2\Big)^{1/2},$$

Predictive intervals and calibration. For a nominal central level γ (we use $\gamma = 0.95$), we form a central prediction interval $[L_n, U_n]$ from empirical quantiles of $\{y_n^{(m)}\}_{m=1}^M$ and compute

$$\mathrm{PICP}_\gamma = \tfrac{1}{N}\sum_{n=1}^{N}\mathbf{1}\{y_n \in [L_n, U_n]\}.$$

Table 5. **Automatic reverse-step selection rule used in the synthetic experiments.**

Setting	Value
Schedule	hybrid
Minimum / maximum step cap	64 / 2000
Target max-ratio	1.01
Early-stopping time δ	10^{-3}
Rule	smallest K such that $\max_k h_k/\vartheta^2_{t_{k-1}} \le 1.01$

For multi-level calibration, we follow (Han et al., 2022) using $\tau \in \{0.50, 0.55, \dots, 0.95\}$, compute empirical coverages $\widehat{C}(\tau)$, and report

$$\mathrm{QICE} = \tfrac{1}{L}\sum_{\ell=1}^{L}\big|\widehat{C}(\tau_\ell) - \tau_\ell\big|.$$

Variance-profile and marginal-density summaries. To visualize input-dependent variance, we compute endpoint conditional-variance estimates $\widehat{v}_\delta(x_n)$ on the original target scale. For each run and solver, we aggregate pointwise estimates into binned curves along the input axis. In Figure 1, solid curves are across-run means of these binned curves, and shaded regions are pointwise empirical 95% bands across five independent runs. The reference curve is the data variance target $v^{\mathrm{data}}(x) = \sigma^2(x)$. To diagnose global spread distortions, we compare the pooled marginal density of generated samples $\{y_n^{(m)}\}_{n,m}$ against the empirical test-target marginal $\{y_n\}_{n=1}^N$.

D.5. Inference budgets and step rules

Predictive sampling budget. Unless stated otherwise, ordinary predictive metrics for diffusion samplers use $M_{\mathrm{eval}} = 200$ predictive samples per test input with $M_{\mathrm{per\ pass}} = 200$. The MC Dropout baseline uses 200 stochastic forward passes, and the DeepHetero baseline uses 200 predictive samples.

Auxiliary analysis budgets. We keep ordinary predictive sampling separate from auxiliary path and endpoint analyses. Stepwise variance-covariance, oracle-gap, score-error, and KL-proxy analyses use $M_{\mathrm{diag}} = 200$ samples per input unless stated otherwise. Endpoint variance fidelity and the solver-matched oracle endpoint accounting use a repeated-MC variance protocol with $M_{\mathrm{var}} = 200$ samples per repetition and $R_{\mathrm{var}} = 5$ repetitions in the reported artifacts.

Automatic reverse-step selection. Table 5 summarizes the automatic reverse-step rule used in the main synthetic experiments. Starting from the shared hybrid time-grid rule, we choose the smallest K_{auto}, subject to the step cap, satisfying the displayed max-ratio criterion, where $h_k := t_k - t_{k-1}$ and $\vartheta_t^2 := 1 - e^{-t}$ is the forward OU marginal variance coefficient. The selected K_{auto} is fixed across solvers within each dataset–run comparison. The DSM hyperparameter $K_{\mathrm{DSM}} = 8$ denotes the number of diffusion times sampled per training example in the DSM loss; it is not the number of reverse sampling steps.

D.6. Calculation details for the main figures and tables

Predictive moments and terminal variance estimates. All predictive moments use the Monte Carlo definitions in Appendix D.4 with the sampling budgets in Appendix D.5. Ordinary predictive summaries use $M_{\mathrm{eval}} = 200$ reverse-time samples per test input, while auxiliary path summaries use $M_{\mathrm{diag}} = 200$ samples per input unless stated otherwise. For memory efficiency, auxiliary path sampling is performed in chunks of at most 64 samples per pass while preserving the same total Monte Carlo budget. All reported variance summaries are computed after mapping samples back to the original target scale.

Figure 1 conditional-variance profiles. Figure 1 is constructed from endpoint variance estimates mapped to the original target scale and the analytic data variance target

$$v^{\mathrm{data}}(x_n) = \sigma^2(x_n)$$

on the synthetic tasks. For each run and solver, pointwise variance estimates are aggregated into an input-axis binned curve. The plotted solid curves are across-run means of these binned curves, and the shaded regions are pointwise empirical 95% bands across five independent runs.

Table 2 : explanatory summaries. For the explanatory-hierarchy analysis, we form one point for each dataset–solver–run combination. With four synthetic datasets, three reverse-time solvers, and five runs per dataset–solver pair, this yields $4 \times 3 \times 5 = 60$ points in total. The response variable is the endpoint variance distortion measured by VarMSE, and the explanatory variables are oracle-aided path summaries computed from the saved reverse-chain artifacts.

Using the oracle-gap path $G_S(t,x)$ defined in Section 2.2, the timewise mean absolute summary is

$$\bar{G}_S(t) = \mathbb{E}_x[\,|G_S(t,x)|\,],$$

and the oracle-gap AUC is

$$\mathrm{AUC}_{\text{oracle-gap},S} = \int_\delta^T \bar{G}_S(t)\,dt,$$

evaluated numerically by trapezoidal integration over the recorded reverse-time grid.

For comparison, we also compute a marginal Gaussian KL proxy between the learned reverse-chain marginal and the solver-matched oracle marginal. For each recorded time t_j and evaluation input x_i, let

$$\hat{m}^{\text{learn}}_{j,i} = \frac{1}{M}\sum_{r=1}^{M} Y^{\text{learn},(r)}_{t_j}(x_i), \qquad \hat{v}^{\text{learn}}_{j,i} = \frac{1}{M}\sum_{r=1}^{M}\left(Y^{\text{learn},(r)}_{t_j}(x_i) - \hat{m}^{\text{learn}}_{j,i}\right)^2,$$

and define $\hat{m}^{\text{oracle}}_{j,i}$ and $\hat{v}^{\text{oracle}}_{j,i}$ analogously from the solver-matched oracle chain. The pointwise KL proxy is

$$\kappa(t_j,x_i) = D_{\mathrm{KL}}\big(\mathcal{N}(\hat{m}^{\text{learn}}_{j,i}, \hat{v}^{\text{learn}}_{j,i}) \,\big\|\, \mathcal{N}(\hat{m}^{\text{oracle}}_{j,i}, \hat{v}^{\text{oracle}}_{j,i})\big)\,,$$

which in the scalar case is evaluated as

$$\kappa(t_j,x_i) = \frac{1}{2}\left[\log\frac{\hat{v}^{\text{oracle}}_{j,i}+\epsilon}{\hat{v}^{\text{learn}}_{j,i}+\epsilon} + \frac{\hat{v}^{\text{learn}}_{j,i} + \big(\hat{m}^{\text{learn}}_{j,i} - \hat{m}^{\text{oracle}}_{j,i}\big)^2}{\hat{v}^{\text{oracle}}_{j,i}+\epsilon} - 1\right],$$

with $\epsilon = 10^{-12}$ used only for numerical stability. We then average over the diagnostic inputs,

$$\bar{\kappa}(t_j) = \mathbb{E}_x[\kappa(t_j,x)] \approx \frac{1}{B}\sum_{i=1}^{B}\kappa(t_j,x_i),$$

and use the recorded-grid path mean

$$\mathrm{KL}_{\text{proxy-mean}} = \frac{1}{J}\sum_{j=1}^{J}\bar{\kappa}(t_j)$$

as the marginal KL-type explanatory summary. This quantity compares marginal Gaussian summaries at each reverse time and does not represent a path-space KL or a transition-level mechanism decomposition.

We use oracle-gap summaries as explanatory path summaries for terminal variance distortion rather than as alternative endpoint error measures. The KL proxy is included as a marginal distributional baseline because it compares learned and solver-matched oracle Gaussian marginals at each recorded time, but it does not encode the transport-weighted accumulation structure used by the variance-gap decomposition.

Pathwise-unrolled mechanism summary for Figure 2. We use the local defects and transported contribution paths defined in Section 3.2. This appendix specifies only the numerical aggregation used in the figure. On the recorded reverse-time grid, write $C^a(t_j,x) := C^a_j(x)$. The plotted mechanism bars report

$$\int_\delta^T \mathbb{E}_x|C^a(t,x)|\,dt, \tag{61}$$

computed by trapezoidal integration on the recorded grid. These bars summarize+pathwise-unrolled contribution paths, not endpoint-only attribution bounds. The same aggregation is reused in Appendix D.9.2 for variance-regime summaries.

Table 6. **Common hyperparameters used in most synthetic tasks.**

Category	Setting (value)
Architecture	hidden width $= 128$, depth $= 4$, dropout $= 0.1$, noise-level embedding frequencies $F_\sigma = 8$
Training	epochs $= 500$, batch size $= 1024$
Optimizer	learning rate $= 9.1685 \times 10^{-4}$, weight decay $= 10^{-5}$, warmup epochs $= 0$, scheduler $cosine$
EMA	ema decay=0.999
DSM loss	$\gamma = 0.7617$, $T_0 = 10^{-3}$, $K = 8$
Stability terms	step clip $= 1.0$, dt eps $= 4.891 \times 10^{-5}$, jitter $= 1.342 \times 10^{-8}$
Hybrid schedule	$\rho = 8.3587$, tail fraction $= 0.02$, $\kappa = 1.5$
Evaluation	Monte Carlo samples per test input $M = 200$

Table 7. **Dataset-specific exceptions (when different from Table 6).**

Dataset	Setting	Value
Synthetic Xsinx–Hetero	hidden width / epochs / batch size	256 / 300 / 512
	learning rate / weight decay / warmup epochs	2.0271×10^{-4} / 10^{-6} / 50
	DSM (γ, T_0)	$(0.8174,\ 10^{-3})$
	stability (step clip,dt eps)	$(2.0,\ 1.2193 \times 10^{-5})$

D.7. Sensitivity checks for Table 2

Blocked resampling check. Table 2 reports Pearson correlation, Spearman rank correlation, and linear R^2 between each explanatory summary and endpoint variance MSE. As an auxiliary robustness check, we also compute a blocked resampling analysis over observed dataset–run blocks in the released analysis artifacts. Each block contains the solver-specific points for one dataset and one run directory. This check is used only to assess whether the reported ordering is dominated by individual dataset–run blocks. We do not use it as an additional formal uncertainty statement in the main table.

D.8. Hyperparameters

Score model and training. Table 6 lists the default hyperparameters used across most datasets. Table 7 reports the small set of dataset-specific deviations.

Baselines. Table 8 reports baseline hyperparameters.

D.9. Additional synthetic-task figures

This appendix provides detailed synthetic-task evidence supporting the solver-dependent variance-fidelity and mechanism analyses in Section 4. We organize the supplementary results into three parts: predictive panels, grouped amplification summaries, and dataset-level unrolled mechanism overviews with variance-regime summaries. The predictive panels visualize how changing the reverse-time solver affects the sampled predictive law. The mechanism overviews in Appendix D.9.2 show the oracle-gap path together with pathwise-unrolled contribution paths. Table 9 summarizes the corresponding unrolled mechanism shares within representative variance regimes. These figures and tables are used as oracle-aided explanatory evidence in the synthetic setting, not as quantities available at ordinary test time.

D.9.1. Predictive panels

Figures 3–6 report solver-wise predictive means, 95% prediction intervals, final-sample panels, and pooled marginal densities. These panels provide a dataset-level visual check of how changing only the reverse-time solver affects the sampled predictive law. Across the four tasks, EM and ExpInt remain visually close in predictive means, interval shapes, final samples, and pooled marginal densities. PF-ODE shows the clearest departures from the stochastic solvers, especially in the spread and tail behavior of the sampled predictive distribution. These supplementary panels support the main-text observation that solver choice can change predictive uncertainty even when the learned score field and

Table 8. **Baseline hyperparameters.**

Method	Setting (value)
MC Dropout	hidden width $= 128$, dropout $= 0.1$, lr $= 10^{-3}$, epochs $= 500$, bs $= 256$
DeepHetero	depth $= 3$, hidden width $= 128$, Gaussian NLL, lr $= 10^{-3}$, epochs $= 500$, bs $= 256$

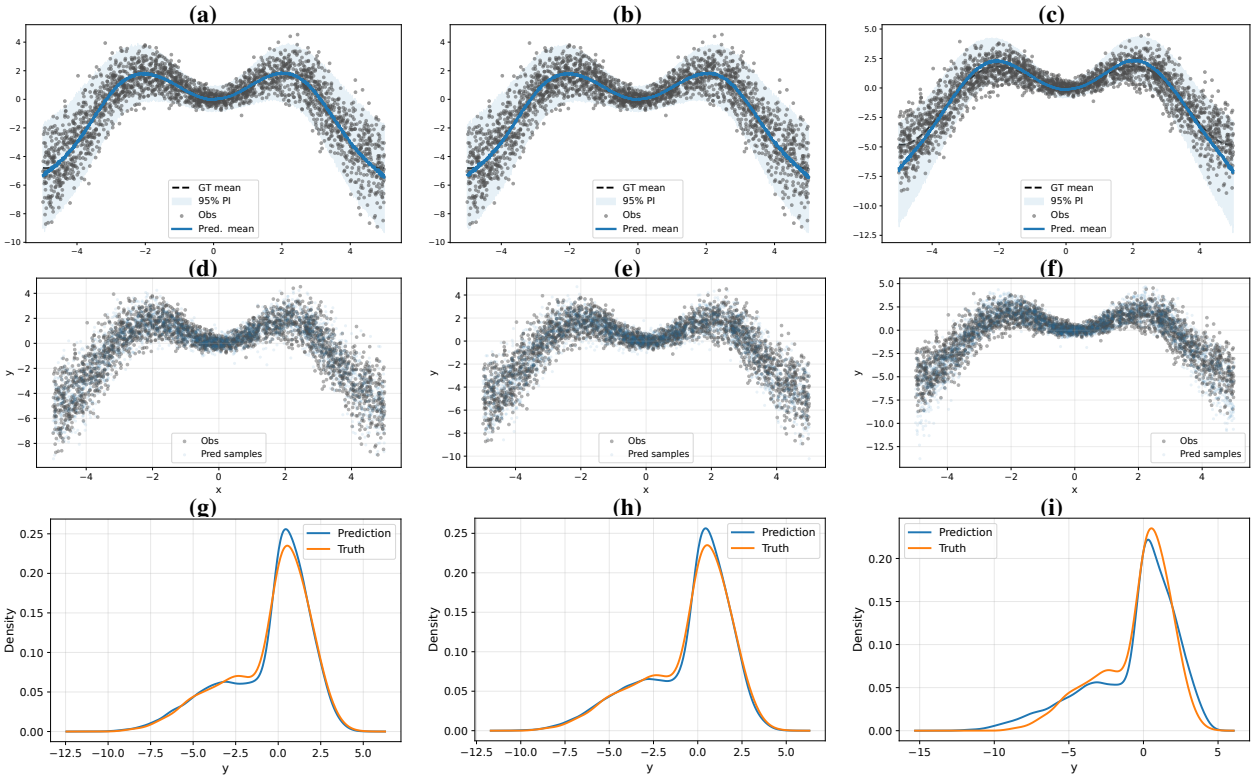

Figure 3. **Xsinx–Hetero: predictive mean/interval, final-sample, and pooled density panels.** Top row (a–c): predictive mean and 95% prediction interval for ExpInt, EM, and PF-ODE. Middle row (d–f): final predictive samples for ExpInt, EM, and PF-ODE. Bottom row (g–i): pooled marginal densities for ExpInt, EM, and PF-ODE.

noise schedule are fixed.

D.9.2. Dataset-level pathwise-unrolled mechanism overviews and variance-regime summaries

Figure 2 in the main text shows the Xsinx–Hetero pathwise-unrolled mechanism overview. The additional appendix overviews in Figures 7–9 provide the corresponding dataset-level summaries for Sine–Hetero, Curve–Hetero, and Cubic–Step–Hetero. Each overview shows the learned–oracle variance-gap path and pathwise-unrolled contribution summaries for transport, injection, remainder, and residual terms. Table 9 reports the full variance-regime summaries for all four synthetic tasks, including Xsinx–Hetero.

Local transport factorization. The transport component can be interpreted through the local variance multiplier of a one-step reverse update. For a local response map, the learned-chain variance recursion can be written as

$$v_{t_{k-1}}(x) = J_k(x)^2 v_{t_k}(x) + \sigma_k^2 + r_k(x), \tag{62}$$

where $J_k(x)$ is the local transport factor, σ_k^2 is the injected variance term, and $r_k(x)$ collects higher-order or non-affine remainder effects. For deterministic PF-ODE under the Heun convention, $\sigma_k^2 = 0$, and $J_k(x)$ is the local derivative of the Heun update.

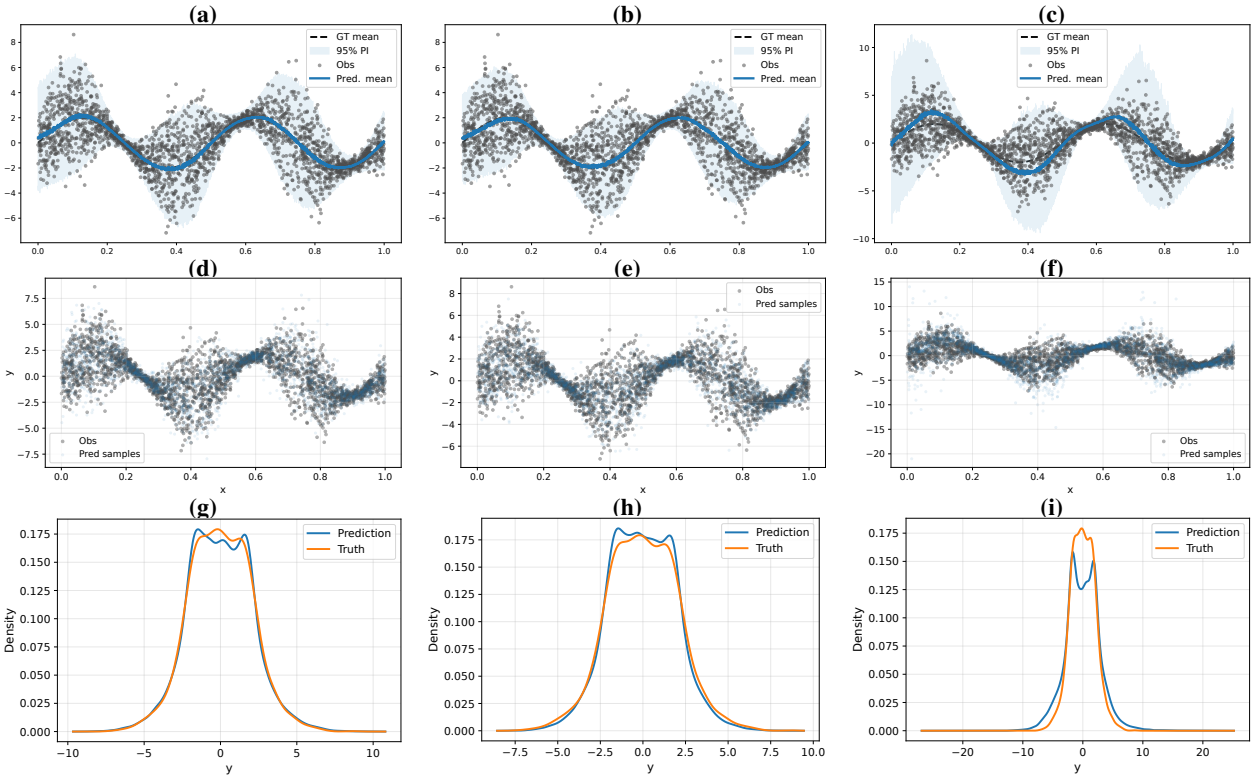

Figure 4. **Sine–Hetero: predictive mean/interval, final-sample, and pooled density panels.** Top row (a–c): predictive mean and 95% prediction interval for ExpInt, EM, and PF-ODE. Middle row (d–f): final predictive samples for ExpInt, EM, and PF-ODE. Bottom row (g–i): pooled marginal densities for ExpInt, EM, and PF-ODE.

For affine reverse-SDE updates, the local transport factor takes the form

$$J_k(x) = A_k + B_k s_{\theta,y}(t_k, m_{t_k}(x), x), \tag{63}$$

with the corresponding solver-matched oracle factor

$$J_k^\star(x) = A_k + B_k s_y^\star(t_k, m_{t_k}^\star(x), x). \tag{64}$$

Thus the transport mismatch satisfies

$$\begin{aligned} J_k(x)^2 - \big(J_k^\star(x)\big)^2 &= \big(J_k(x) - J_k^\star(x)\big)\big(J_k(x) + J_k^\star(x)\big) \\ &= B_k\Big[s_{\theta,y}(t_k, m_{t_k}(x), x) - s_y^\star(t_k, m_{t_k}^\star(x), x)\Big] \\ &\quad \times \Big[2A_k + B_k\big\{s_{\theta,y}(t_k, m_{t_k}(x), x) + s_y^\star(t_k, m_{t_k}^\star(x), x)\big\}\Big]. \end{aligned} \tag{65}$$

The first bracket contains both score-slope error and center-drift effects:

$$\begin{aligned} &s_{\theta,y}(t_k, m_{t_k}(x), x) - s_y^\star(t_k, m_{t_k}^\star(x), x) \\ &= \underbrace{s_{\theta,y}(t_k, m_{t_k}(x), x) - s_y^\star(t_k, m_{t_k}(x), x)}_{\text{score-slope error at the learned center}} \\ &\quad + \underbrace{s_y^\star(t_k, m_{t_k}(x), x) - s_y^\star(t_k, m_{t_k}^\star(x), x)}_{\text{center-drift / curvature effect}}. \end{aligned} \tag{66}$$

If the true score is locally twice differentiable in y, the second term can be written by the mean-value theorem as

$$s_y^\star(t_k, m_{t_k}(x), x) - s_y^\star(t_k, m_{t_k}^\star(x), x) = s_{yy}^\star(t_k, \widetilde{m}_k(x), x)\big(m_{t_k}(x) - m_{t_k}^\star(x)\big), \tag{67}$$

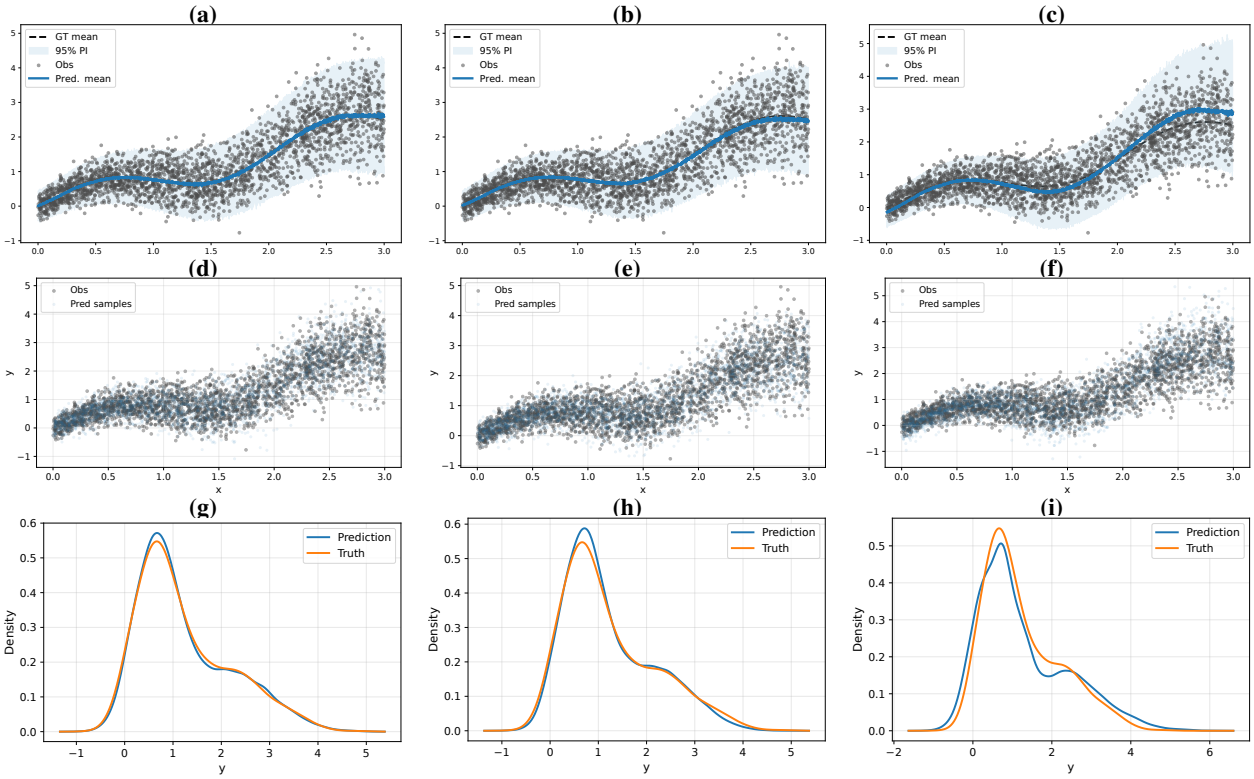

Figure 5. **Curve–Hetero: predictive mean/interval, final-sample, and pooled density panels.** Top row (a–c): predictive mean and 95% prediction interval for ExpInt, EM, and PF-ODE. Middle row (d–f): final predictive samples for ExpInt, EM, and PF-ODE. Bottom row (g–i): pooled marginal densities for ExpInt, EM, and PF-ODE.

for some $\widetilde{m}_k(x)$ between $m_{t_k}(x)$ and $m^\star_{t_k}(x)$.

In the variance-gap bound, the transport contribution is scaled by the oracle variance path:

$$\Pi_{1:k-1}(x)\big|J_k(x)^2 - \big(J^\star_k(x)\big)^2\big|v^\star_{t_k}(x). \tag{68}$$

Equivalently, in the affine solver-matched case this term is controlled by a product of a solver sensitivity factor, a local score-slope or center-drift mismatch, and the oracle variance being transported:

$$\begin{aligned}
&\Pi_{1:k-1}(x)\big|J_k(x)^2 - \big(J^\star_k(x)\big)^2\big|v^\star_{t_k}(x) \\
&\quad\lesssim \Pi_{1:k-1}(x)|B_k|\Big|2A_k + B_k\big\{s_{\theta,y}(t_k, m_{t_k}(x), x) + s^\star_y(t_k, m^\star_{t_k}(x), x)\big\}\Big| \\
&\qquad\times \Big(\big|s_{\theta,y}(t_k, m_{t_k}(x), x) - s^\star_y(t_k, m_{t_k}(x), x)\big| + \big|s^\star_{yy}(t_k, \widetilde{m}_k(x), x)\big|\big|m_{t_k}(x) - m^\star_{t_k}(x)\big|\Big)v^\star_{t_k}(x).
\end{aligned} \tag{69}$$

This factorization explains why a fixed local transport mismatch can have a larger effect when the transported oracle variance $v^\star_{t_k}(x)$ is larger. It does not require high-variance regimes to create a larger local mismatch; rather, larger transported variance can amplify the effect of the same score-slope or center-drift mismatch.

For a representative variance group $\mathcal{G}_g$, we average the recorded pathwise quantities over inputs in the group. Define the groupwise mean absolute oracle gap at step k by

$$\overline{|G|}_{k,g} := \frac{1}{|\mathcal{G}_g|}\sum_{x\in\mathcal{G}_g}\big|v_{t_k}(x) - v^\star_{t_k}(x)\big|.$$

The table reports

$$\mathrm{GapMean}_g := \frac{1}{K+1}\sum_{k=0}^{K}\overline{|G|}_{k,g}. \tag{70}$$

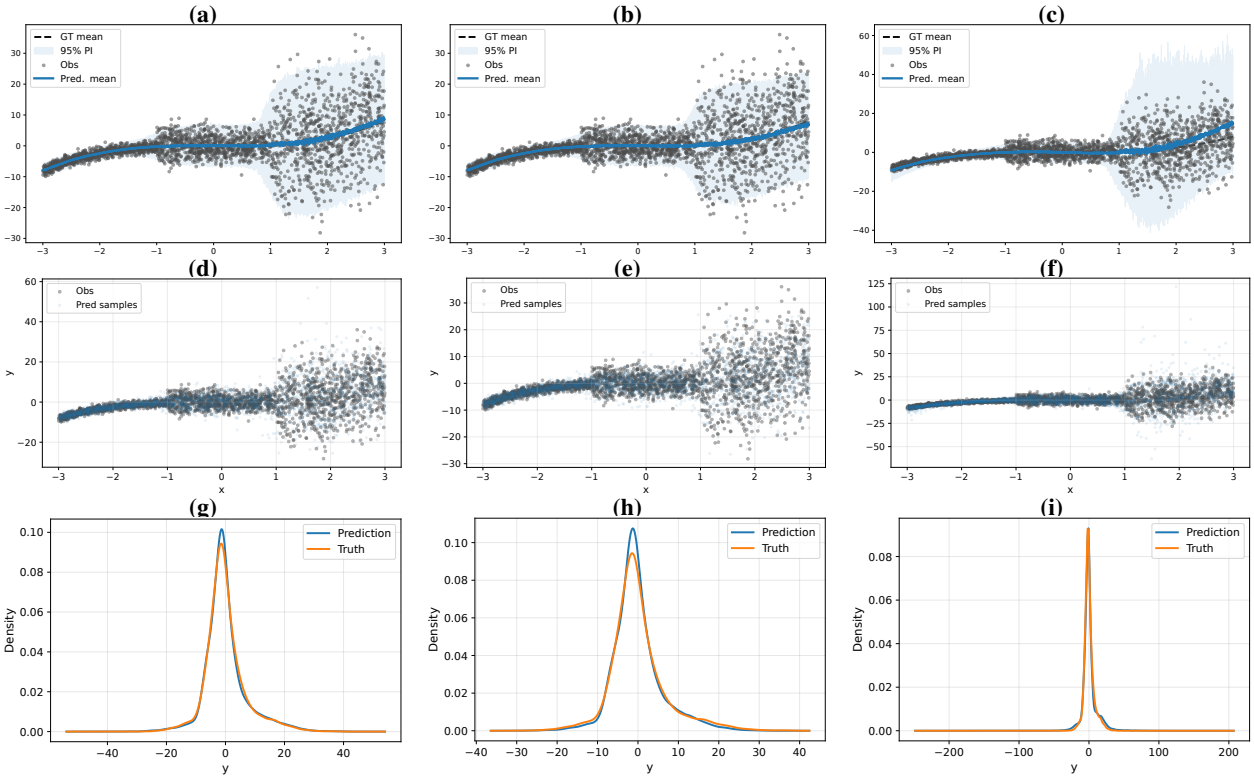

Figure 6. **Cubic–Step–Hetero: predictive mean/interval, final-sample, and pooled density panels.** Top row (a–c): predictive mean and 95% prediction interval for ExpInt, EM, and PF-ODE. Middle row (d–f): final predictive samples for ExpInt, EM, and PF-ODE. Bottom row (g–i): pooled marginal densities for ExpInt, EM, and PF-ODE.

For each mechanism component $a \in \{\mathrm{trans}, \mathrm{inj}, \mathrm{rem}, \mathrm{res}\}$, we use the nodewise contribution path $C_k^a(x)$ defined in Section 3.2. For a representative variance group $\mathcal{G}_g$, define

$$\overline{|C^a|}_{k,g} := \frac{1}{|\mathcal{G}_g|} \sum_{x \in \mathcal{G}_g} |C_k^a(x)|.$$

Let ω_k^{trap} denote the trapezoidal quadrature weights on the recorded grid. We compute the pathwise-unrolled contribution AUC

$$\mathrm{AUC}_g^a := \sum_{k=0}^{K} \omega_k^{\mathrm{trap}} \, \overline{|C^a|}_{k,g}. \tag{71}$$

The carry-over term $J_j(x)^2 C_j^a(x)$ propagates a previously created component contribution from one node to the next; it is not counted as a new local mechanism. In the displayed tables, we report only the non-negligible displayed components

$$\mathcal{A}_{\mathrm{disp}} := \{\mathrm{trans}, \mathrm{rem}\}.$$

The reported displayed shares are

$$\mathrm{Share}_g^{a,\mathrm{disp}} := \frac{\mathrm{AUC}_g^a}{\sum_{b \in \mathcal{A}_{\mathrm{disp}}} \mathrm{AUC}_g^b}, \qquad a \in \mathcal{A}_{\mathrm{disp}}. \tag{72}$$

Table 9 reports the remaining variance-regime summaries beyond the Xsinx–Hetero case in the main text. It includes Sine–Hetero, Curve–Hetero, and Cubic–Step–Hetero. Gap mean reports the average absolute learned–oracle variance-gap path within each regime. Transport and remainder report normalized shares of the displayed pathwise mechanism

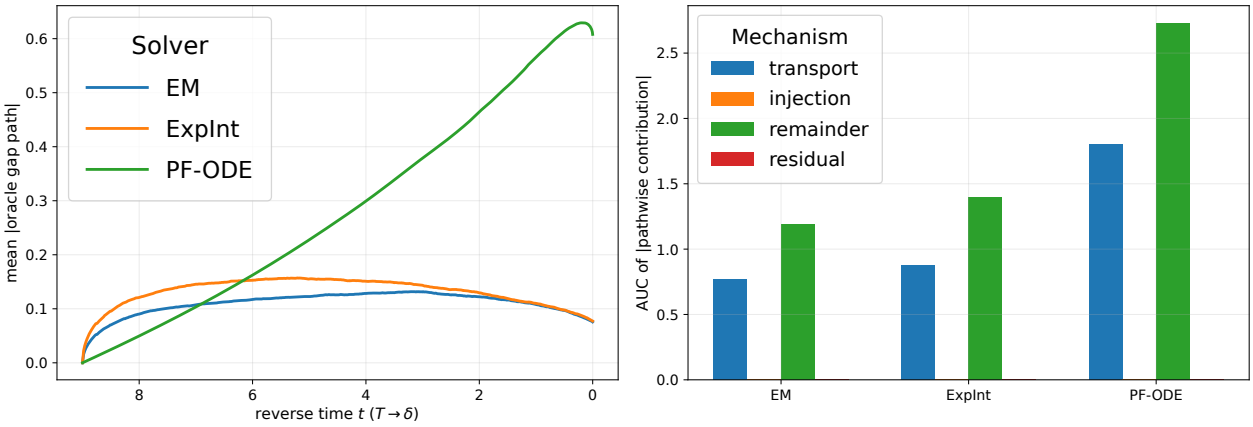

Figure 7. **Sine–Hetero learned–oracle variance-gap path and mechanism AUCs.** Left shows the mean absolute learned–oracle variance-gap path along the reverse chain. Right reports pathwise-unrolled mechanism contribution AUCs for transport, injection, remainder, and residual terms.

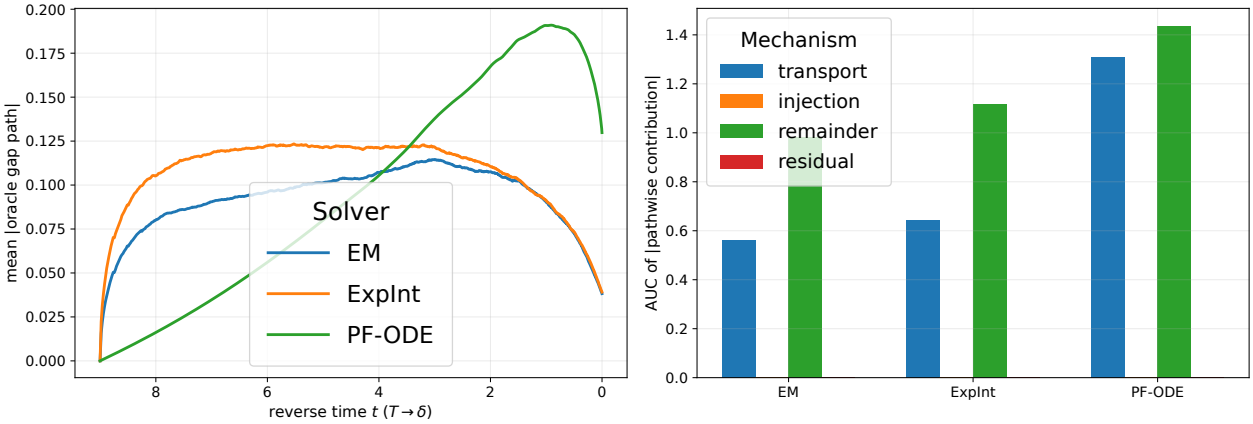

Figure 8. **Curve–Hetero learned–oracle variance-gap path and mechanism AUCs.** Left shows the mean absolute learned–oracle variance-gap path along the reverse chain. Right reports pathwise-unrolled mechanism contribution AUCs for transport, injection, remainder, and residual terms.

contributions. Across these remaining regimes, high-variance groups often increase the learned–oracle gap, but the dominant mechanism is not universally transport. Transport dominates in Curve–Hetero high, whereas Sine–Hetero high and Cubic–Step–Hetero at $\sigma^2 = 100$ are more remainder-heavy. This supports a pathwise interpretation in which transport amplification is one important route to variance distortion, but not the only one.

Solver-matched oracle endpoint accounting. The pathwise summaries above describe how the learned–oracle variance gap evolves along the reverse chain. To connect this pathwise gap to the final variance error against the data target, we use a separate endpoint accounting on the same evaluation inputs.

Let $v^{\mathrm{data}}(x) = \mathrm{Var}(Y \mid X = x)$. For a fixed solver S, define

$$A_S(x) = \widehat{v}^{\mathrm{learn}}_{S,t_0}(x) - \widehat{v}^{\star}_{S,t_0}(x), \qquad B_S(x) = \widehat{v}^{\star}_{S,t_0}(x) - v^{\mathrm{data}}(x).$$

Thus A_S is the learned–oracle endpoint gap, while B_S is the solver-matched oracle endpoint bias relative to the data variance target. For each input,

$$\widehat{v}^{\mathrm{learn}}_{S,t_0}(x) - v^{\mathrm{data}}(x) = A_S(x) + B_S(x).$$

Let $\mathcal{E}$ denote the common evaluation set and write

$$\langle f \rangle_{\mathcal{E}} := |\mathcal{E}|^{-1} \sum_{x \in \mathcal{E}} f(x).$$

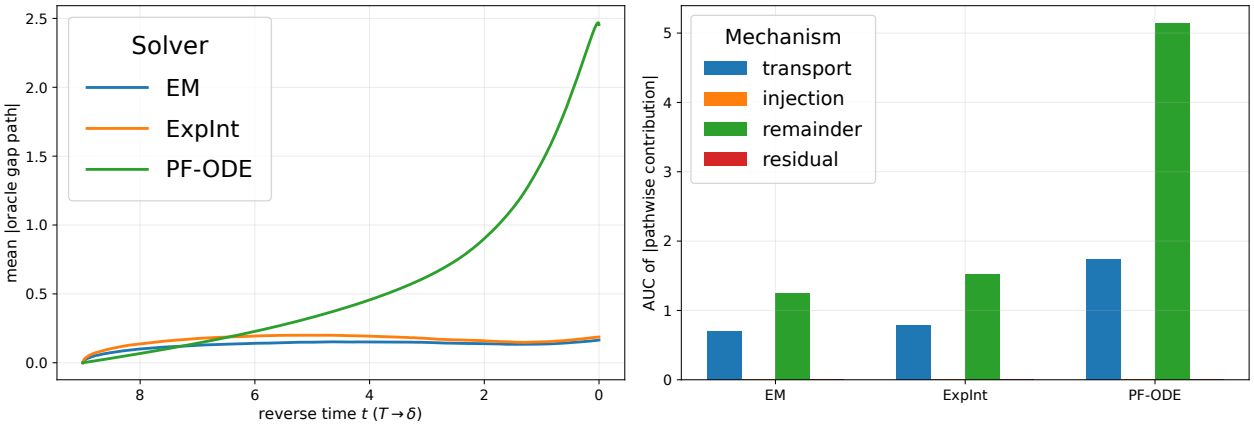

Figure 9. **Cubic–Step–Hetero learned–oracle variance-gap path and mechanism AUCs.** Left shows the mean absolute learned–oracle variance-gap path along the reverse chain. Right reports pathwise-unrolled mechanism contribution AUCs for transport, injection, remainder, and residual terms.

The reported plug-in quantities satisfy the empirical identity

$$\widehat{\mathrm{MSE}}_{\mathrm{learned},S} = \langle (A_S + B_S)^2 \rangle_{\mathcal{E}} = \langle A_S^2 \rangle_{\mathcal{E}} + \langle B_S^2 \rangle_{\mathcal{E}} + 2\langle A_S B_S \rangle_{\mathcal{E}}.$$

Table 10 reports this same-support endpoint accounting, which is separate from the pathwise mechanism attribution.

The endpoint accounting verifies that the learned–oracle variance gap studied pathwise is also the dominant component of the final variance error against v^{data}. In the largest PF-ODE failure rows, $\langle A_S^2 \rangle_{\mathcal{E}}$ is close to $\langle (A_S + B_S)^2 \rangle_{\mathcal{E}}$, whereas $\langle B_S^2 \rangle_{\mathcal{E}}$ is much smaller. The cross term is included through $2\langle A_S B_S \rangle_{\mathcal{E}}$ and does not change this dominance.

Table 9. **Additional variance-regime pathwise mechanism summaries.**

Dataset / regime	Solver	Gap mean	Transport %	Remainder %
Sine–Hetero / low	EM	0.0683	47	53
Sine–Hetero / low	ExpInt	0.0761	47	53
Sine–Hetero / low	PF-ODE	0.14	59	41
Sine–Hetero / middle	EM	0.12	33	67
Sine–Hetero / middle	ExpInt	0.14	32	68
Sine–Hetero / middle	PF-ODE	0.54	31	69
Sine–Hetero / high	EM	0.14	32	68
Sine–Hetero / high	ExpInt	0.16	31	69
Sine–Hetero / high	PF-ODE	0.94	27	73
Curve–Hetero / low	EM	0.0637	36	64
Curve–Hetero / low	ExpInt	0.0694	35	65
Curve–Hetero / low	PF-ODE	0.0486	45	55
Curve–Hetero / middle	EM	0.0767	37	63
Curve–Hetero / middle	ExpInt	0.0818	37	63
Curve–Hetero / middle	PF-ODE	0.0716	52	48
Curve–Hetero / high	EM	0.11	43	57
Curve–Hetero / high	ExpInt	0.13	46	54
Curve–Hetero / high	PF-ODE	0.57	68	32
Cubic–Step–Hetero / $\sigma^2 = 1$	EM	0.066	40	60
Cubic–Step–Hetero / $\sigma^2 = 1$	ExpInt	0.072	39	61
Cubic–Step–Hetero / $\sigma^2 = 1$	PF-ODE	0.0681	54	46
Cubic–Step–Hetero / $\sigma^2 = 9$	EM	0.0909	38	62
Cubic–Step–Hetero / $\sigma^2 = 9$	ExpInt	0.099	38	62
Cubic–Step–Hetero / $\sigma^2 = 9$	PF-ODE	0.24	49	51
Cubic–Step–Hetero / $\sigma^2 = 100$	EM	0.28	32	68
Cubic–Step–Hetero / $\sigma^2 = 100$	ExpInt	0.34	29	71
Cubic–Step–Hetero / $\sigma^2 = 100$	PF-ODE	3.77	16	84

Table 10. **Solver-matched oracle endpoint accounting.** Entries are plug-in empirical averages over the same evaluation inputs. A_S is the learned–oracle endpoint gap and B_S is the oracle endpoint bias relative to v^{data}. Cross reports $\langle A_S B_S \rangle_{\mathcal{E}}$ and enters the identity with factor 2.

Dataset	Solver	Learned $\langle (\boldsymbol{A}+\boldsymbol{B})^2 \rangle$	Oracle bias $\langle \boldsymbol{B}^2 \rangle$	Gap $\langle \boldsymbol{A}^2 \rangle$	Cross $\langle \boldsymbol{AB} \rangle$
Xsinx–Hetero	EM	0.0156	0.00515	0.0196	-0.00462
Xsinx–Hetero	ExpInt	0.0157	0.00516	0.0197	-0.00457
Xsinx–Hetero	PF-ODE	0.797	0.00534	0.778	0.00669
Sine–Hetero	EM	0.12	0.0223	0.115	-0.00869
Sine–Hetero	ExpInt	0.123	0.0152	0.138	-0.0151
Sine–Hetero	PF-ODE	15.49	0.0154	15.36	0.0614
Curve–Hetero	EM	0.000468	0.000242	0.00083	-0.000302
Curve–Hetero	ExpInt	0.000802	0.00022	0.000983	-0.0002
Curve–Hetero	PF-ODE	0.0569	0.000214	0.0562	0.000237
Cubic–Step–Hetero	EM	93.64	19.92	102.7	-14.49
Cubic–Step–Hetero	ExpInt	193.4	6.765	199.0	-6.195
Cubic–Step–Hetero	PF-ODE	51574.9	6.462	51411.8	78.33