

Counterfactual Generation via Flow Matching: Coupling-Sensitive End-to-End Rates

Yunrui Guan^{*1}, Krishnakumar Balasubramanian^{2,3}, and Shiva Prasad Kasiviswanathan³

¹Department of Applied Mathematics and Statistics, Johns Hopkins University.

²Department of Statistics, University of California, Davis.

³Amazon

Abstract

Counterfactual generation seeks to sample outcomes under a hypothetical intervention or decision using observational data collected under the factual assignment mechanism. We develop a flow-matching approach that combines a sample-split, doubly robust training objective with a learned coupling between observed source outcomes and target outcomes drawn from a fitted conditional outcome model. To enable finite-step generation, we leverage a score-corrected stochastic sampler based on a Gaussian-smoothed interpolation. Our main theoretical contribution is a coupling-sensitive KL bound for constant-step Euler discretization: the error is controlled by moments of the source–target displacement under the chosen coupling, rather than by global uniform regularity of the velocity field, and has near-linear dependence on the ambient dimension. We also establish finite-sample non-parametric guarantees for the learned velocity and score fields when both the conditional outcome model and the source–target coupling are estimated from data. These bounds separate approximation, coupling-replacement, nuisance-estimation, generalization, and Monte Carlo errors and, combined with the sampler analysis, yield an end-to-end guarantee for counterfactual generation. Experiments on synthetic and semi-synthetic image benchmarks support the coupling-dependent theory and show that, at finite discretization budgets, the stochastic sampler can outperform the corresponding deterministic ODE sampler.

Contents

1	Introduction	3
2	Problem Setup and Method	4
2.1	Gaussian-Smoothed Flow Matching	4
2.2	Sample Splitting and Learned Coupling	5
2.3	Doubly Robust Field Training	5
2.4	Score-Corrected SDE Sampler	6
3	Coupling-Dependent Sampling and End-to-End Guarantees	7
3.1	Coupling-Dependent Finite-Step Sampling	7
3.2	Finite-Sample Field Estimation and End-to-End Guarantee	8

^{*}This work was done during an internship at Amazon. Email: yguan34@jh.edu

4	Experiments	10
4.1	Angular Mode Reallocation	11
4.2	Image Benchmarks	11
5	Conclusion	11
A	Additional Experiment Details	17
A.1	Angular-Mode Experiment Details	17
A.2	Image-Benchmark Details and Additional Results	17
B	Additional Related-Work Details	18
B.1	Comparison With Dance et al. (2026)	20
C	Reader’s Guide and Proof Roadmap	22
C.1	Summary of Notation	24
D	Algorithms and Auxiliary-Model Details	25
D.1	Prediction-Oracle Accuracy	25
D.2	Score-Field Training Objective	26
D.3	Algorithms	26
E	Proof of Example 1	29
F	Proof of Example 2	31
G	Proof of Theorem 1	34
G.1	Path-Space Comparison via Girsanov	36
G.2	Posterior Derivative Identities	41
G.3	Integrated Derivative Bounds	49
G.4	Euler Increment Bounds	60
G.5	Completion of the KL Sampling Bound	66
G.6	Wasserstein Conversion to the Unsmoothed Target	68
H	Coupling Construction and Replacement	71
H.1	The Learned EOT Kernel	72
H.2	Finite-Sample Target Approximation	77
I	Proof of Proposition 1	79
I.1	Velocity-Field Bound	83
I.1.1	Velocity Network Class and Approximation	83
I.1.2	Completion of the Velocity Bound	89
I.2	Score-Field Bound	91
I.2.1	Score Loss and Exact Risk Identity	92
I.2.2	Score Network Class and Approximation	92
I.2.3	Completion of the Score Bound	100
J	Proof of Theorem 2	102

1 Introduction

Counterfactual generation seeks to model the distribution of outcomes that would arise if a treatment were assigned throughout a target population. The full distribution reveals heterogeneity, variability, and tail behavior that an average alone cannot capture. With observational data, however, treatment assignment may depend on background covariates. Individuals who receive treatment a can therefore differ systematically from the target population, so a generator trained directly on their outcomes may reproduce this selection bias.

Concretely, we observe data (C, A, Y) , where C denotes background covariates, A is the treatment, and $Y \in \mathbb{R}^d$ is the outcome. Fix a treatment value a . The observed and target outcome distributions average the same conditional outcome law over different covariate distributions:

$$\underbrace{\mathbb{P}_{Y|A=a} = \int \mathbb{P}_{Y|C=c, A=a} d\mathbb{P}_{C|A=a}(c)}_{\text{Observed outcome distribution}} \quad \text{and} \quad \underbrace{\int \mathbb{P}_{Y|C=c, A=a} d\mathbb{P}_C(c)}_{\text{Target outcome distribution}} .$$

The observed law weights covariates as they occur among individuals who received treatment a ; the target instead retains the covariate distribution of the full population. Under overlap, this target is identified from the observed-data law. If consistency and conditional exchangeability also hold, it coincides with the interventional outcome distribution under treatment a (Hernán and Robins, 2020; Dance et al., 2026).

This structure suggests transporting observed outcomes to the target rather than generating the target from an unrelated reference distribution. Given a coupled source–target pair (Y_0, Y_1) , flow matching learns a time-dependent velocity field by regressing the displacement $Y_1 - Y_0$ along paths between the paired endpoints. The coupling determines which source and target outcomes are connected. Structured choices, such as optimal-transport couplings, can shorten these displacements and facilitate flow learning (Pooladian et al., 2023; Tong et al., 2024; Liu et al., 2023).

Counterfactual flow matching nevertheless poses three linked problems: the target-dependent training loss must be estimated from observational data, the source–target coupling must itself be learned, and the learned dynamics must be simulated accurately with finitely many steps.

Recent work primarily addresses the first problem through doubly robust training. DoubleGen (Luedtke and Fukumizu, 2025) combines a prediction oracle $\hat{\mathbb{P}}_{Y|C=c, A=a}$ for $\mathbb{P}_{Y|C=c, A=a}$ with an estimate of the inverse propensity $\alpha_a(c) = 1/\mathbb{P}(A = a | C = c)$. Oracle samples provide conditional training losses, while observed treatment- a outcomes provide inverse-propensity-weighted residual corrections. This construction trains a generator directly for the counterfactual target. In a complementary approach, Dance et al. (2026) use $\mathbb{P}_{Y|A=a}$ as the flow source and derive a debiased objective that accounts for its dependence on the observed-data law. They also extend the construction to learned optimal-transport couplings. Appendix B.1 gives a detailed comparison with their method.

Our focus is the joint effect of coupling estimation, field estimation, and finite-step simulation. Existing theory relates flow straightness to finite-step Wasserstein error (Bansal et al., 2026; Liu et al., 2023), but straightness is a property of the induced flow and cannot be read directly from the selected coupling. Other analyses rely on global Lipschitz bounds (Zhou and Liu, 2025; Liu et al., 2026; Guan et al., 2026b), which reflect worst-case regularity over the state space. Neither approach yields a guarantee controlled directly by the source–target displacement under the chosen coupling. Appendix B provides a broader comparison with related finite-step sampling results.

Examples 1 and 2 in Section 2.4 isolate the two resulting obstacles. Even with fixed Gaussian smoothing, Example 1 constructs couplings whose displacement approaches zero while the velocity derivative diverges and the Euler ODE sampler has infinite KL error. Without smoothing, Example 2

shows that arbitrarily small displacement does not control the integrated squared derivative of the score. Together, the examples motivate stochastic sampling, positive smoothing, and regularity control averaged along the interpolation rather than uniformly over the state space.

We address these obstacles with a Gaussian-smoothed interpolation and a score-corrected SDE. With exact velocity and score fields, the SDE preserves the interpolation marginals and reaches the smoothed target. We estimate the propensity model, prediction oracle, and entropic optimal-transport coupling on an auxiliary data split, then train the velocity and score fields on an independent split using doubly robust objectives. Our analysis combines displacement-dependent discretization control with finite-sample field bounds, separating smoothing, coupling, estimation, and numerical errors in an end-to-end guarantee. Our main contributions are as follows.

- **Score-Corrected Stochastic Sampling.** We analyze a score-corrected SDE for a Gaussian-smoothed interpolation with constant smoothing. The two counterexamples delineate the respective roles of stochasticity and positive smoothing in obtaining coupling-dependent finite-step control.
- **Coupling-Sensitive Discretization Theory.** For constant-step Euler discretization, we prove a KL error bound controlled by

$$D_5 = (\mathbb{E}\|Y_1 - Y_0\|^5)^{1/5}, \quad (1)$$

rather than by global uniform regularity. The fifth moment is the highest displacement moment required by our derivative bounds; its empirical analogue is a coupling diagnostic. With exact fields and the prescribed noise tuning, Euler-step complexity is near-linear in dimension for fixed smoothing, D_5 , and target accuracy.

- **Finite-Sample Learning and End-to-End Guarantees.** We derive finite-sample bounds for sample-split, doubly robust velocity and score estimators when the conditional outcome model, propensity model, and source–target coupling are learned from data. The bounds separate neural approximation, coupling replacement, nuisance estimation, generalization, and Monte Carlo errors. Combined with the sampling analysis, they yield an end-to-end W_1 guarantee relative to the unsmoothed counterfactual target.
- **Empirical Validation.** Experiments on synthetic and semi-synthetic image benchmarks support the predicted dependence on coupling quality and show that, at finite discretization budgets, the score-corrected SDE can outperform the corresponding deterministic ODE sampler.

2 Problem Setup and Method

This section develops our counterfactual flow-matching procedure. We first define the Gaussian-smoothed interpolation and its population fields, then construct the learned source–target coupling and doubly robust field objectives, and finally introduce the score-corrected SDE sampler.

2.1 Gaussian-Smoothed Flow Matching

Given a coupled source–target pair (Y_0, Y_1) , let I_d denote the $d \times d$ identity matrix, and let $G \sim \mathcal{N}(0, I_d)$ be independent of (Y_0, Y_1) . For $t \in [0, 1]$, define

$$S_t = Y_0 + t(Y_1 - Y_0), \quad X_t = S_t + \sigma G,$$

where $\sigma > 0$ is constant in time. This Gaussian-smoothed interpolation is a special case of the stochastic interpolant framework (Albergo et al., 2025; Liu et al., 2026). Write q_t for the density of

X_t . Its endpoint distributions are those of $Y_0 + \sigma G$ and $Y_1 + \sigma G$; thus, the terminal population law is a Gaussian-smoothed approximation to the target.

For $t \sim \text{Unif}[0, 1]$, independently of (Y_0, Y_1, G) , the population flow-matching loss and its minimizer are

$$\mathcal{L}_{\text{FM}}(v) = \mathbb{E} \left[\|v(X_t, t) - (Y_1 - Y_0)\|^2 \right], \quad v_t(x) := v^*(x, t) = \mathbb{E}[Y_1 - Y_0 \mid X_t = x].$$

Gaussian smoothing gives the score as the population minimizer of the score-matching loss:

$$\mathcal{L}_{\text{score}}(s) = \mathbb{E} \left[\left\| s(X_t, t) + \frac{G}{\sigma} \right\|^2 \right], \quad s_t(x) := s^*(x, t) = \nabla_x \log q_t(x) = \mathbb{E} \left[-\frac{G}{\sigma} \mid X_t = x \right].$$

Appendix D.2 verifies the score identity. Under suitable regularity conditions, the ODE with velocity v_t , initialized from q_0 , has marginal density q_t . Section 2.4 introduces stochastic dynamics that preserve the same marginal path.

2.2 Sample Splitting and Learned Coupling

For a fixed treatment value a , we observe $2n$ independent samples from the observed-data law and split them into $\mathcal{D}_{\text{tr}} = \{(C_i, A_i, Y_i)\}_{i=1}^n$, and $\mathcal{D}_{\text{aux}} = \{(C_i, A_i, Y_i)\}_{i=n+1}^{2n}$. The auxiliary split is used to estimate the propensity $\hat{\mathbb{P}}(A = a \mid C = c)$, prediction oracle $\hat{\mathbb{P}}_{Y|C=c, A=a}$, and coupling kernel $\hat{Q}_n(dy_0 \mid y_1)$. The velocity and score fields are then trained on $\mathcal{D}_{\text{tr}}$. Conditioning on the auxiliary split fixes these estimated objects during the training analysis.

The empirical source distribution is

$$\hat{\mathbb{P}}_{Y|A=a}^{\text{aux}} = \frac{\sum_{i=n+1}^{2n} \mathbf{1}\{A_i = a\} \delta_{Y_i}}{\sum_{i=n+1}^{2n} \mathbf{1}\{A_i = a\}}.$$

This distribution is defined when the auxiliary split contains at least one treatment- a observation. The sampler is initialized from its Gaussian convolution: $\hat{\mathbb{P}}_{Y|A=a}^{\text{aux}} * \mathcal{N}(0, \sigma^2 I_d)$.

We approximate the target distribution using $\hat{\mathbb{P}}_{Y|C=c, A=a}$, implemented as a conditional flow trained on $\mathcal{D}_{\text{aux}}$; see Appendix D.1. For each auxiliary covariate C_i , we draw M_Q independent oracle samples $\hat{Y}_{i,1}^Q, \dots, \hat{Y}_{i,M_Q}^Q$, yielding

$$\hat{\mathbb{P}}_{Y_1}^Q = \frac{1}{nM_Q} \sum_{i=n+1}^{2n} \sum_{r=1}^{M_Q} \delta_{\hat{Y}_{i,r}^Q}.$$

We construct an entropic optimal-transport (EOT) coupling with regularization parameter $\varepsilon_{\text{OT}} > 0$ between $\hat{\mathbb{P}}_{Y|A=a}^{\text{aux}}$ and $\hat{\mathbb{P}}_{Y_1}^Q$. Disintegration with respect to the target coordinate gives the kernel $\hat{Q}_n(dy_0 \mid y_1)$; its extension to target queries outside the finite Monte Carlo sample is described in Appendix H. The prediction-based target is used only to construct this coupling kernel. The field objectives use the doubly robust correction described next.

2.3 Doubly Robust Field Training

With the auxiliary objects fixed, a training loss is constructed by pairing each target endpoint y_1 with $y_0 \sim \hat{Q}_n(\cdot \mid y_1)$, and then sampling the interpolation time and Gaussian noise. For each training covariate, let $M \geq 1$ be the number of independent prediction-oracle draws used to approximate the plug-in loss. Let $\hat{\ell}_i^{\text{aux}}(v)$ denote the average flow-matching loss using oracle-generated outcomes at

C_i , and let $\widehat{\ell}_i^{\text{obs}}(v)$ denote the corresponding loss using the observed outcome Y_i when $A_i = a$. For a minibatch $\mathcal{B} \subseteq \{1, \dots, n\}$, the loss is

$$\widehat{L}_{\mathcal{B}, \text{DR}}(v) = \frac{1}{|\mathcal{B}|} \sum_{i \in \mathcal{B}} \left[\widehat{\ell}_i^{\text{aux}}(v) + \frac{\mathbf{1}\{A_i = a\}}{\widehat{\mathbb{P}}(A = a \mid C = C_i)} \left\{ \widehat{\ell}_i^{\text{obs}}(v) - \widehat{\ell}_i^{\text{aux}}(v) \right\} \right].$$

The first term is the prediction-oracle plug-in loss. Observed treatment- a outcomes provide an inverse-propensity-weighted correction to its residual error. We train the score using the same construction with $\mathcal{L}_{\text{FM}}$ replaced by $\mathcal{L}_{\text{score}}$. Algorithm 2 gives the complete minibatch procedure.

2.4 Score-Corrected SDE Sampler

Let $N \geq 1$, $h = 1/N$, and recall the coupling-displacement moment Equation (1). We seek a finite-step error bound controlled by D_5 , rather than by global uniform regularity, which need not follow from small coupling displacement (Zhou and Liu, 2025; Liu et al., 2026; Guan et al., 2026b,a; Stéphanovitch, 2026). Examples 1 and 2 isolate the two resulting obstacles: stochasticity addresses the deterministic Euler obstruction, whereas positive smoothing provides the score regularity needed for displacement-dependent control.

Let $\widehat{X}_N^{\text{ODE}}$ denote the terminal constant-step Euler iterate using the exact velocity v_t , initialized from q_0 . The following example shows that arbitrarily small D_5 guarantees neither small uniform velocity regularity nor finite KL error for deterministic Euler sampling.

Example 1. Let $p \in (0, 1)$, and let $B \sim \text{Bernoulli}(p)$ and $U \sim \mathcal{N}(0, 1)$ be independent. Define

$$(Y_0, Y_1) = \begin{cases} (U, U), & B = 0, \\ (2U, 2(1 - \lambda)U), & B = 1, \end{cases} \quad \lambda = \frac{4 + \sigma^2}{4h}.$$

As $p \rightarrow 0$,

$$D_5 \rightarrow 0, \quad \sup_{t \in [0, 1], x \in \mathbb{R}} |\partial_x v_t(x)| \rightarrow \infty, \quad \text{but} \quad \int_0^1 \mathbb{E} |\partial_x v_t(X_t)|^2 dt \rightarrow 0.$$

Moreover, for every $p \in (0, 1)$, we have $\text{KL}(\mathcal{L}(X_1) \parallel \mathcal{L}(\widehat{X}_N^{\text{ODE}})) = +\infty$.

Appendix E proves the claim. It motivates controlling derivatives averaged along the interpolation rather than uniformly over the state space. Following the Itô-based approach of Xu et al. (2026); Benton et al. (2024), we therefore use a stochastic sampler for which this integrated regularity is sufficient.

For $\tau > 0$, the exact-field sampling process is

$$dZ_t = \left(v_t(Z_t) + \frac{\tau^2}{2} s_t(Z_t) \right) dt + \tau dB_t, \quad Z_0 \sim q_0,$$

where B_t is a standard Brownian motion. Under the conditions of Lemma 1, Z_t has density q_t for every $t \in [0, 1]$. Thus the score correction preserves the interpolation's marginal path while introducing stochasticity. Applying Itô's formula to v_t and s_t along Z_t reduces the analysis to integrated derivative bounds, which Gaussian smoothing controls through moments of $Y_1 - Y_0$.

Initialize the learned sampler according to $\widehat{Z}_0 \sim \widehat{\mathbb{P}}_{Y|A=a}^{\text{aux}} * \mathcal{N}(0, \sigma^2 I_d)$. For $t_k = kh$, the learned Euler sampler is

$$\widehat{Z}_{k+1} = \widehat{Z}_k + h \left(\widehat{v}(\widehat{Z}_k, t_k) + \frac{\tau^2}{2} \widehat{s}(\widehat{Z}_k, t_k) \right) + \tau \sqrt{h} G_k, \quad G_k \sim \mathcal{N}(0, I_d),$$

where the G_k are mutually independent and independent of the initialization. Since $t_N = Nh = 1$, the terminal Euler output $\widehat{Z}_N$ is compared with $Z_{t_N} = Z_1$. Algorithm 1 gives the complete procedure. Setting $\tau = 0$ recovers deterministic Euler sampling, although our KL analysis requires $\tau > 0$. Appendix B discusses related stochastic samplers and finite-step comparisons.

Gaussian smoothing is essential for controlling the required score regularity through D_5 . The following example shows that this control fails in general when $\sigma = 0$.

Example 2. Fix $\varepsilon > 0$, set $\delta = 2^{1/5}\varepsilon$, and let $Y_0 \sim \mathcal{N}(0, 1)$. For each integer $m \geq 0$, define

$$K_m = \frac{(2m+1)\pi}{\delta}, \quad p_m(x) = \frac{1}{4} + \frac{1}{8} \sin(K_m x).$$

Conditional on $Y_0 = x$, let

$$Y_1 = \begin{cases} x + \delta, & \text{with probability } p_m(x), \\ x - \delta, & \text{with probability } p_m(x), \\ x, & \text{with probability } 1 - 2p_m(x). \end{cases}$$

For the unsmoothed interpolation $Y_t = (1-t)Y_0 + tY_1$, let $q_t^{(m)}$ be its density and $s_t^{(m)} = \partial_x \log q_t^{(m)}$; the dependence of Y_t on m is suppressed in the notation.. Every coupling in this sequence satisfies $D_5 = \varepsilon$, but

$$\int_0^1 \mathbb{E} \left[|\partial_x s_t^{(m)}(Y_t)|^2 \right] dt \geq \frac{35}{32768} K_m^4 (1 - e^{-2K_m^2})^4 \rightarrow +\infty.$$

Thus, even arbitrarily small D_5 does not uniformly control the integrated score regularity without smoothing.

3 Coupling-Dependent Sampling and End-to-End Guarantees

This section proves a coupling-dependent finite-step sampling bound, derives finite-sample guarantees for the doubly robust velocity and score estimators, and combines them into an end-to-end sample-complexity result. Unless otherwise stated, $\lesssim$ suppresses universal numerical constants and may also hide logarithmic factors in the displayed parameters. A complete notation list is provided in Appendix C.1.

3.1 Coupling-Dependent Finite-Step Sampling

For $t_k = kh$, where $h = 1/N$ (so $t_N = 1$), write $\widehat{v}_{t_k}(x) = \widehat{v}(x, t_k)$ and $\widehat{s}_{t_k}(x) = \widehat{s}(x, t_k)$. We use the [Standard Sampling Conditions](#) in Appendix G. For $\sigma, \tau > 0$ and $D_5 < \infty$, Lemma 1 provides a jointly continuous exact drift and a weak solution Z with $\mathcal{L}(Z_t) = q_t$; throughout the sampling analysis, we use this solution. The following Assumption 1 defines $\varepsilon_{\text{vel}}^2$ and $\varepsilon_{\text{score}}^2$ as the grid-averaged mean-squared differences between the learned and population fields evaluated at Z_{t_k} , and assumes both are finite. This is standard in related sampling analyses ([Liu et al., 2025, 2026](#); [Li et al., 2025](#); [Benton et al., 2024](#)). Section 3.2 later provides finite-sample conditions that verify this assumption.

Assumption 1 (Integrated Estimation Error). Assume

$$\varepsilon_{\text{vel}}^2 := h \sum_{k=0}^{N-1} \mathbb{E} \left[\|v_{t_k}(Z_{t_k}) - \widehat{v}_{t_k}(Z_{t_k})\|^2 \right] < \infty, \quad \varepsilon_{\text{score}}^2 := h \sum_{k=0}^{N-1} \mathbb{E} \left[\|s_{t_k}(Z_{t_k}) - \widehat{s}_{t_k}(Z_{t_k})\|^2 \right] < \infty.$$

Theorem 1 (KL Sampling Bound). *Let $N, d \geq 1$, $h = 1/N$, $\sigma, \tau > 0$, and $D_5 \in (0, \infty)$. Under [Standard Sampling Conditions](#) and [Assumption 1](#), define*

$$A_h = \frac{4(9d + 51)}{3} \frac{h^2 D_5^4}{\sigma^2} + 2\varepsilon_{\text{vel}}^2, \quad B_h = \frac{124}{3} \frac{h^2 (d + 6)^3 D_5^2}{\sigma^4} + \frac{1}{2} \varepsilon_{\text{score}}^2.$$

Then

$$\text{KL}(\mathcal{L}(Z_{t_N}) \parallel \mathcal{L}(\hat{Z}_N)) \leq \frac{A_h}{\tau^2} + B_h \tau^2 + \frac{hd(d+3)}{\sigma^4} \tau^4 + \frac{68h^2(d+6)^3}{3\sigma^6} \tau^6 + \frac{2h(d+1)D_5^2}{\sigma^2}. \quad (2)$$

The first two terms are balanced by $\tau = (A_h/B_h)^{1/4}$.

Remark 1 (Coupling-Dependent Scaling). *In the exact-field case $\varepsilon_{\text{vel}} = \varepsilon_{\text{score}} = 0$, choosing $\tau^2 = \sigma D_5/d$ gives*

$$\text{KL}(\mathcal{L}(Z_{t_N}) \parallel \mathcal{L}(\hat{Z}_N)) \lesssim \frac{hd}{\sigma^2} D_5^2 + \frac{h^2 d^2}{\sigma^3} D_5^3. \quad (3)$$

Thus, for any $\varepsilon_{\text{target}} \in (0, 1)$,

$$h \lesssim \frac{1}{d D_5^2 / (\varepsilon_{\text{target}}^2 \sigma^2) + d D_5 / \sigma}$$

gives KL error of order $\varepsilon_{\text{target}}^2$, corresponding to $\tilde{O}(d[\frac{D_5^2}{\varepsilon_{\text{target}}^2 \sigma^2} + \frac{D_5}{\sigma}])$ Euler steps.

Recall that [Example 2](#) motivates the importance of smoothing. Compared with bounds for constant-step Euler discretizations of stochastic-interpolant SDE and ODE samplers in ([Liu et al., 2025, 2026](#)), our result improves the dimension dependence from quadratic to near-linear; [Appendix B](#) gives a detailed comparison. This dimension comparison holds with smoothing, displacement moment, and target accuracy fixed.

3.2 Finite-Sample Field Estimation and End-to-End Guarantee

We now give sufficient conditions for [Assumption 1](#). Below, d is fixed, so implicit constants may depend on d in addition to the logarithmic factors suppressed by $\lesssim$.

Assumption 2 (Clipping, Overlap, and Bounded Support). *We work with outcomes clipped to $[-1, 1]^d$ and assume that, for every $c \in \text{supp}(\mathbb{P}_C)$,*

$$\begin{aligned} \mathbb{P}_{Y|C=c, A=a}([-1, 1]^d) &= \hat{\mathbb{P}}_{Y|C=c, A=a}([-1, 1]^d) = 1, \\ \mathbb{P}(A = a \mid C = c) &\geq \pi_{\min} > 0, \quad 0 \leq \hat{\alpha}(c) \leq 2\pi_{\min}^{-1}. \end{aligned}$$

Here $\hat{\alpha}(c) = 1/\hat{\mathbb{P}}(A = a \mid C = c)$. Thus every endpoint pair satisfies $Y_0, Y_1 \in [-1, 1]^d$ almost surely, and $\hat{Q}_n([-1, 1]^d \mid y) = 1$ for every $y \in \mathbb{R}^d$.

Assumption 3 (Propensity Accuracy). *Let $\alpha_a(c) = 1/\mathbb{P}(A = a \mid C = c)$. Conditional on the auxiliary sample used to fit $\hat{\alpha}$, assume*

$$\|\hat{\alpha}/\alpha_a - 1\|_{L^2(\mathbb{P}_C)} \leq \varepsilon_\alpha.$$

[Assumption 2](#) bounds outcomes, oracle-generated samples, and inverse-propensity weights while enforcing overlap. [Assumption 3](#) requires relative $L^2(\mathbb{P}_C)$ accuracy of the estimated inverse propensity.

Assumption 4 (Prediction-Oracle Accuracy). Define $\varepsilon_{\text{pred}} = [(\varepsilon_{\text{pred,train}})^2 + \varepsilon_{\text{pred,gen}}^2]^{1/2}$, and assume

$$\left[\mathbb{E}_C W_1 \left(\mathbb{P}_{Y|C=c, A=a}, \hat{\mathbb{P}}_{Y|C=c, A=a} \right)^2 \right]^{1/2} \leq \varepsilon_{\text{pred}}.$$

Assumption 4 controls the oracle's average conditional W_1 error over population covariates. Appendix D.1 derives this requirement from empirical-fit and generalization conditions and describes conditional-flow training. Let $\mu_1 := \int \mathbb{P}_{Y|C=c, A=a} d\mathbb{P}_C(c)$. For any coupling kernel $Q(dy_0 | y_1)$, define the joint endpoint law

$$\Gamma_{\mu_1, Q}(dy_0, dy_1) := Q(dy_0 | y_1) \mu_1(dy_1).$$

To emphasize the role of coupling explicitly, under $(Y_0, Y_1) \sim \Gamma_{\mu_1, Q}$, we let $q_{t, Q}$ denote the density of $X_t = Y_0 + t(Y_1 - Y_0) + \sigma G$, and define

$$v_Q^*(x, t) := \mathbb{E}[Y_1 - Y_0 | X_t = x], \quad s_Q^*(x, t) := \nabla_x \log q_{t, Q}(x).$$

Population Coupling. Let $\mathcal{E}_a = \{\sum_{i=n+1}^{2n} \mathbf{1}\{A_i = a\} \geq 1\}$. On $\mathcal{E}_a$, optimally couple $Y_1 \sim \mu_1$ and $\hat{Y}_1 \sim \hat{\mathbb{P}}_{Y_1}^Q$ in W_1 , and then draw $Y_0 \sim \hat{Q}_n(\cdot | \hat{Y}_1)$. We define $Q(\cdot | Y_1)$ as the resulting conditional law of Y_0 given Y_1 . Thus $\Gamma_{\mu_1, Q}$ has target marginal μ_1 and source marginal $\hat{\mathbb{P}}_{Y|A=a}^{\text{aux}}$. Lemma 10 formalizes this construction and proves the coupling-replacement bound used below.

Let $\mathcal{V}_n$ denote the clipped coordinatewise sparse-ReLU velocity-field class with input-clipping radius R and approximation tolerance ε , as defined in Appendix I.1, and let $\mathcal{S}_n$ denote the score-field class induced by clipped sparse-ReLU denoisers through $s_u(x, t) = \{u(x, t) - x\}/\sigma^2$, as defined in Appendix I.2. The next proposition gives explicit finite-sample bounds for the velocity- and score-field components of Assumption 1. If Assumption 7, which ensures that the auxiliary split contains at least one treatment- a observation, is added to the hypotheses of Proposition 1, then $\mathcal{E}_a$ holds by assumption and the term $e^{-n\pi_{\min}}$ can be omitted from the failure probability.

Proposition 1 (Doubly Robust Velocity and Score Field Error Bound). Let $n \geq 2$, $M, M_Q \geq 1$, and $0 < \sigma \leq 1$. Under Assumptions 2, 3, 4, let Q be the *Population Coupling* defined above. Let $\varepsilon_{\text{pred,train}} \leq \varepsilon_{\text{pred}}$ denote the empirical generator accuracy for prediction oracle, formally defined in Assumption 5. Define

$$\mathcal{A} := \frac{\varepsilon_{\text{OT}}^{-1}}{\sigma^2} \left[n^{-1/(d+2)} + \varepsilon_{\text{pred,train}} \right] + \frac{1 + \varepsilon_{\text{OT}}^{-1}}{\sigma^2} \varepsilon_\alpha \varepsilon_{\text{pred}}.$$

Then, with probability at least $1 - \mathcal{O}(n^{-1}) - e^{-n\pi_{\min}}$ over the data and Monte Carlo draws, $\mathcal{E}_a$ occurs and the following bound holds:

1. For velocity field, Let $\hat{v}_M$ be the exact empirical minimizer over $\mathcal{V}_n$, with $\varepsilon = n^{-1/(d+5)}$ and $R = 1 + \sigma[2 \log(32d^2/\varepsilon^2)]^{1/2}$.

$$\frac{1}{N} \sum_{k=0}^{N-1} \mathbb{E} \left[\|\hat{v}_M(X_{t_k}, t_k) - v_Q^*(X_{t_k}, t_k)\|^2 \right] \lesssim \mathcal{A} + \frac{1 + 2\pi_{\min}^{-1}}{\sigma^{d+1}} n^{-2/(d+5)}. \quad (4)$$

2. For score field, let $\hat{s}_M$ be the exact empirical minimizer over $\mathcal{S}_n$, with $\varepsilon = n^{-1/(d+5)}$ and $R = 1 + \sqrt{2}\sigma[\log(8d^2/(\sigma^4\varepsilon^2))]^{1/2}$.

$$\frac{1}{N} \sum_{k=0}^{N-1} \mathbb{E} \left[\|\hat{s}_M(X_{t_k}, t_k) - s_Q^*(X_{t_k}, t_k)\|^2 \right] \lesssim \sigma^{-4} \mathcal{A} + \frac{1 + 2\pi_{\min}^{-1}}{\sigma^{2d+6}} n^{-2/(d+5)}. \quad (5)$$

In Proposition 1, the finite- M and finite- M_Q errors are dominated by the displayed n -dependent terms, up to logarithmic factors.

Verification of Assumption 1. Set $\hat{v}_{t_k} = \hat{v}_M(\cdot, t_k)$ and $\hat{s}_{t_k} = \hat{s}_M(\cdot, t_k)$. For the interpolation induced by Q , $v_t = v_Q^*(\cdot, t)$ and $s_t = s_Q^*(\cdot, t)$, while the interpolation variable X_{t_k} and the exact SDE state Z_{t_k} both have density $q_{t_k} = q_{t_k, Q}$. Hence the expectations in the two field bounds equal the corresponding expectations in Assumption 1. Since $h = 1/N$, the two displayed grid averages are exactly $\varepsilon_{\text{vel}}^2$ and $\varepsilon_{\text{score}}^2$, respectively.

In each field bound, the three terms respectively capture coupling replacement, doubly robust nuisance error, and final-field estimation. The first vanishes for a coupling fixed independently of the prediction-based target, while the second is small when either nuisance model is accurate. Substituting both bounds into Theorem 1 controls finite-step KL error relative to the smoothed target. Corollary 1 then adds the smoothing bias and converts this control into a W_1 guarantee for the original target, yielding the following theorem.

Theorem 2 (End-to-End Sample Complexity With Negligible Nuisance Error). *Fix $d \geq 1$, $\pi_{\min} > 0$, and $\varepsilon_{\text{OT}} > 0$. Assume the conditions of Proposition 1. Let $\mu_1 = \int \mathbb{P}_{Y|C=c, A=a} d\mathbb{P}_C(c)$ and let $\overline{D}_5 > 0$ be a deterministic coupling-displacement envelope such that, almost surely on the event where Q is defined,*

$$\left(\mathbb{E}_{\Gamma_{\mu_1, Q}} \|Y_1 - Y_0\|^5 \right)^{1/5} \leq \overline{D}_5.$$

Suppose $\varepsilon_{\text{pred, train}}$ and $\varepsilon_\alpha \varepsilon_{\text{pred}}$ are negligible relative to the explicit n -dependent terms. For $\varepsilon_{\text{target}} \in (0, 1)$, choose $0 < \sigma \leq 1$ with $\sigma = \Theta(\varepsilon_{\text{target}})$, set $\tau^2 = \sigma \overline{D}_5 / d$, and choose $N \geq 1$ so that

$$h = 1/N \leq \frac{c_h}{d \overline{D}_5^2 / (\varepsilon_{\text{target}}^4 \sigma^2) + d \overline{D}_5 / \sigma},$$

where $c_h \leq 2$ is a scaling constant. Then, up to logarithmic factors, it suffices to take

$$n \gtrsim \max \left\{ \left[\max \left\{ \frac{2 + \varepsilon_{\text{OT}}^{-1}}{\overline{D}_5} \varepsilon_{\text{target}}^{-7}, \frac{\overline{D}_5}{\varepsilon_{\text{OT}}} \varepsilon_{\text{target}}^{-9} \right\} \right]^{d+2}, \right. \\ \left. \left[(1 + 2\pi_{\min}^{-1}) \max \left\{ \overline{D}_5^{-1} \varepsilon_{\text{target}}^{-(d+6)}, \overline{D}_5 \varepsilon_{\text{target}}^{-(2d+9)} \right\} \right]^{\frac{(d+5)}{2}} \right\}.$$

This yields $W_1(\mu_1, \mathcal{L}(\hat{Z}_N)) \lesssim \varepsilon_{\text{target}}$ with probability at least $1 - \mathcal{O}(n^{-1}) - e^{-n\pi_{\min}}$.

This gives an end-to-end guarantee for counterfactual flow matching. Under Assumption 2, we may clip the final generated sample coordinatewise to $[-1, 1]$. This post-processing won't increase its W_1 distance to the target and gives a sharper bound in W_1 ; see Corollary 2. Although Theorem 1 has near-linear dimension dependence, the end-to-end sample complexity still suffers from the curse of dimensionality, as expected under ambient-dimensional nonparametric assumptions and consistent with related flow-matching and diffusion results (Zhou and Liu, 2025; Su et al., 2025; Wibisono et al., 2024; Cai and Li, 2025).

4 Experiments

We first isolate the effects of coupling and discretization in two dimensions, then evaluate finite-step generation on two semi-synthetic image benchmarks.

4.1 Angular Mode Reallocation

We use a two-dimensional mixture of eight Gaussians arranged on a circle to isolate coupling effects. The observed source places total mass $3/4$ on the axis-aligned modes and $1/4$ on the diagonal modes, while the counterfactual target reverses these masses. We compare IID and entropic optimal-transport couplings across regularization levels and use the resulting exact empirical velocity and score fields as references for the learned models. Figures 1(a) and 1(b) show that smaller D_5 generally corresponds to smaller field-estimation error. Appendix A gives the complete mixture construction and experimental protocol. The remaining panels isolate finite-step sampling error using exact fields. Figure 1(c) shows how the 8-step KL error varies with D_5 when $\tau = 0.5\sqrt{\sigma D_5/d}$. For fixed h , σ , and d , Theorem 1 predicts the coupling-dependent order $\text{KL}(\mathcal{L}(Z_{t_N})\|\mathcal{L}(\hat{Z}_N)) \lesssim D_5^2 + D_5^3$, consistent with the empirical trend. Figure 1(d) gives a complementary mode-mass view: with 4 Euler steps, both the ODE and SDE samplers substantially misallocate mass, whereas with 64 steps both closely match the target and assign little mass outside the mode regions.

4.2 Image Benchmarks

We evaluate finite-step counterfactual generation on two semi-synthetic colored-image benchmarks: MNIST digits 1 and 6, adapted from Dance et al. (2026), and Fashion-MNIST sneakers and ankle boots. In both tasks, $C \sim \text{Unif}(0, 1)$ controls image color, A determines the class, and treatment assignment follows $\mathbb{P}(A = 1 \mid C = c) = \exp\{5(c - \frac{1}{2})\} / (1 + \exp\{5(c - \frac{1}{2})\})$. This selection mechanism biases the observed $A = 1$ images toward larger values of C , whereas C remains uniform under $\text{do}(A = 1)$; the clean counterfactual target can therefore be sampled directly. We vary the SDE noise level τ , compare with the corresponding $\tau = 0$ ODE and a pure $\sigma = 0$ ODE baseline, and use 50 Euler steps throughout. We report sliced Wasserstein-2 (SW2) distance to the clean target. Appendix A.2 gives the complete data construction, training, and evaluation protocols.

Digit ColorMNIST. Here $A = 0$ denotes digit 1, $A = 1$ denotes digit 6, and the observational sample size is $n = 20,000$. We evaluate smoothing levels $\sigma \in \{0.01, 0.015\}$; smaller σ reduces smoothing bias and places the smoothed target closer to the clean target. Figure 2 shows that, for both values of σ , some $\tau > 0$ improves on the corresponding $\tau = 0$ ODE, and the best results also outperform the pure $\sigma = 0$ baseline. For $\sigma = 0.01$, performance is best near $\tau = 0.075$. Along the corresponding SDE trajectory, the color-distribution KS statistic decreases from 0.266 at $t = 0$ to 0.079 at $t = 1$.

Fashion-MNIST. Here $A = 0$ denotes sneakers and $A = 1$ denotes ankle boots. We evaluate $\sigma \in \{0.005, 0.01\}$. As shown in Figure 3, the benefit of stochastic sampling depends strongly on the smoothing level. For $\sigma = 0.005$, performance improves with τ , is best near $\tau = 0.02$, and clearly outperforms the pure ODE baseline before deteriorating for larger τ . For $\sigma = 0.01$, varying τ yields no clear improvement over the pure ODE baseline within the reported variation, possibly because of the larger smoothing error. Along the $\sigma = 0.005$ trajectory, the red–blue color-statistic KS value decreases from 0.260 at $t = 0$ to 0.063 at $t = 1$. Overall, $\sigma = 0.005$ performs substantially better, consistent with its smaller smoothing error. Qualitative samples are provided in Appendix A.2.

5 Conclusion

We developed a counterfactual flow-matching framework that combines doubly robust field estimation, learned source–target couplings, Gaussian smoothing, and score-corrected stochastic sampling. We prove a coupling-sensitive finite-step guarantee in which the Euler sampling error depends explicitly on the source–target displacement under the chosen coupling. Separately, our finite-sample analysis separates coupling replacement, nuisance estimation, and statistical errors. Together, these results




Figure 1: Angular-mode experiment. (a)–(b) Velocity and score errors across couplings ($\sigma = 0.12$). (c) Exact-field 8-step KL versus D_5 , using 2,000 samples and $\tau = 0.5\sqrt{\sigma D_5/d}$. (d) Mode-mass allocation for 4 and 64 steps (EOT $\epsilon_{\text{OT}} = 0.5$, $\sigma = 0.02$): 4-step ODE/SDE misallocate mass, while 64-step ODE/SDE match the target modes.

yield an end-to-end guarantee relative to the original counterfactual target. Synthetic and image experiments show that stochastic sampling can improve over the corresponding ODE at fixed step budgets, with performance depending on smoothing and SDE noise.

AI Use Statement

In this work, we used Generative AI for formulate mathematical claims (e.g., finding counter examples), provide critical ingredients for proving mathematical claims (e.g., carry out some calculation that were checked by the authors), assist in the writing of proofs, design or provide feedback on research methodology or experiments (e.g., asking AI whether a specific theoretical claim can be improved toward a specific direction), implement methods (e.g., write code), literature search (e.g., ask if we missed any existing literature), language editing, proof auditing, and L^AT_EX assistance. The authors independently checked the mathematical arguments, citations, code, and experimental claims and take responsibility for the final manuscript.

Acknowledgement


We thank the authors of [Dance et al. \(2026\)](#) for sharing their code with us before its public release.

References


- M. S. Albergo, N. M. Boffi, and E. Vanden-Eijnden. Stochastic interpolants: A unifying framework for flows and diffusions. *Journal of Machine Learning Research*, 26(209):1–80, 2025. (Cited on page 4.)
- V. Bansal, S. Roy, A. Rinaldo, and P. Sarkar. On the convergence and straightness of rectified flow. In *The 29th International Conference on Artificial Intelligence and Statistics*, 2026. URL <https://openreview.net/forum?id=qKKxzTucgy>. (Cited on page 3.)
- J. Benton, V. De Bortoli, A. Doucet, and G. Deligiannidis. Nearly d -linear convergence bounds for diffusion models via stochastic localization. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=r5njV3BsuD>. (Cited on pages 6, 7, 19, and 22.)
- C. Cai and G. Li. Minimax optimality of the probability flow ODE for diffusion models, 2025. URL <https://arxiv.org/abs/2503.09583>. (Cited on page 10.)
- M. Chen, W. Liao, H. Zha, and T. Zhao. Distribution approximation and statistical estimation guarantees of generative adversarial networks. *arXiv preprint arXiv:2002.03938v3*, 2020. URL <https://arxiv.org/abs/2002.03938v3>. Version 3, revised July 2022. (Cited on page 86.)
- M. Chen, H. Jiang, W. Liao, and T. Zhao. Nonparametric regression on low-dimensional manifolds using deep ReLU networks: Function approximation and statistical recovery. *Information and Inference: A Journal of the IMA*, 11(4):1203–1253, 2022. doi: 10.1093/imaiai/iaac001. URL <https://doi.org/10.1093/imaiai/iaac001>. (Cited on page 88.)
- S. Chen, S. Chewi, H. Lee, Y. Li, J. Lu, and A. Salim. The probability flow ODE is provably fast. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023a. URL <https://openreview.net/forum?id=KD6MFeWSAd>. (Cited on pages 19 and 36.)
- S. Chen, S. Chewi, J. Li, Y. Li, A. Salim, and A. Zhang. Sampling is as easy as learning the score: theory for diffusion models with minimal data assumptions. In *The Eleventh International Conference on Learning Representations*, 2023b. URL https://openreview.net/forum?id=zyLVMgsZOU_. (Cited on pages 36 and 40.)
- H. Dance, J. Xi, P. Orbanz, and B. Bloem-Reddy. Debiased counterfactual generation via flow matching from observations. *arXiv preprint arXiv:2605.07665*, 2026. URL <https://arxiv.org/abs/2605.07665>. (Cited on pages 3, 11, 12, 17, 20, and 21.)
- A. Genevay, M. Cuturi, G. Peyré, and F. Bach. Stochastic optimization for large-scale optimal transport. In *Advances in Neural Information Processing Systems 29*, pages 3440–3448, 2016. (Cited on page 72.)
- Y. Guan, K. Balasubramanian, and S. Ma. Mirror flow matching with heavy-tailed priors for generative modeling on convex domains. In *The Fourteenth International Conference on Learning Representations*, 2026a. URL <https://openreview.net/forum?id=dZKl7uc0XQ>. (Cited on page 6.)
- Y. Guan, K. Balasubramanian, and S. Ma. Total variation rates for Riemannian flow matching. *arXiv preprint arXiv:2602.05174*, 2026b. URL <https://arxiv.org/abs/2602.05174>. (Cited on pages 3 and 6.)

- M. A. Hernán and J. M. Robins. *Causal Inference: What If*. Chapman & Hall/CRC, Boca Raton, FL, 2020. (Cited on page 3.)
- Y. Jiao, Y. Zhou, and G. Li. Optimal convergence analysis of DDPM for general distributions. *arXiv preprint arXiv:2510.27562*, 2025. URL <https://arxiv.org/abs/2510.27562>. (Cited on pages 19 and 22.)
- J.-F. Le Gall. *Brownian Motion, Martingales, and Stochastic Calculus*, volume 274 of *Graduate Texts in Mathematics*. Springer International Publishing, 2016. doi: 10.1007/978-3-319-31089-3. (Cited on page 40.)
- G. Li and Y. Jiao. Improved convergence rate for diffusion probabilistic models. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=S0d07Qxkw4>. (Cited on page 19.)
- G. Li and Y. Yan. $O(d/T)$ convergence theory for diffusion probabilistic models under minimal assumptions. *Journal of Machine Learning Research*, 26(292):1–55, 2025. (Cited on pages 19 and 22.)
- R. Li, Q. Di, and Q. Gu. Unified convergence analysis for score-based diffusion models with deterministic samplers. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=HrdVqFSn1e>. (Cited on page 7.)
- X. Liu, C. Gong, and Q. Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=XVjTT1nw5z>. (Cited on page 3.)
- Y. Liu, Y. Chen, R. Hu, and L. Huang. Finite-time analysis of discrete-time stochastic interpolants. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pages 38119–38143. PMLR, 2025. URL <https://proceedings.mlr.press/v267/liu25d.html>. (Cited on pages 7, 8, 18, and 22.)
- Y. Liu, Y. Chen, R. Hu, and L. Huang. Finite-time convergence analysis of ODE-based generative models for stochastic interpolants. In *The Fourteenth International Conference on Learning Representations*, 2026. URL <https://openreview.net/forum?id=VCStVrq0BZ>. (Cited on pages 3, 4, 6, 7, 8, 18, 19, and 22.)
- A. Luedtke and K. Fukumizu. DoubleGen: Debaised generative modeling of counterfactuals. *arXiv preprint arXiv:2509.16842*, 2025. URL <https://arxiv.org/abs/2509.16842>. (Cited on page 3.)
- A.-A. Pooladian and J. Niles-Weed. Entropic estimation of optimal transport maps. *arXiv preprint arXiv:2109.12004*, 2021. URL <https://arxiv.org/abs/2109.12004>. (Cited on page 72.)
- A.-A. Pooladian, H. Ben-Hamu, C. Domingo-Enrich, B. Amos, Y. Lipman, and R. T. Q. Chen. Multisample flow matching: Straightening flows with minibatch couplings. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 28100–28127. PMLR, 2023. URL <https://proceedings.mlr.press/v202/pooladian23a.html>. (Cited on page 3.)
- A. Stéphanovitch. Lipschitz regularity in flow matching and diffusion models: sharp sampling rates and functional inequalities. *arXiv preprint arXiv:2604.06065*, 2026. URL <https://arxiv.org/abs/2604.06065>. (Cited on page 6.)


- M. Su, J. Y.-C. Hu, S. Pi, and H. Liu. On flow matching KL divergence. *arXiv preprint arXiv:2511.05480*, 2025. URL <https://arxiv.org/abs/2511.05480>. (Cited on pages 10 and 19.)
- A. Tong, K. Fatras, N. Malkin, G. Huguet, Y. Zhang, J. Rector-Brooks, G. Wolf, and Y. Bengio. Improving and generalizing flow-based generative models with minibatch optimal transport. *Transactions on Machine Learning Research*, 2024. ISSN 2835-8856. URL <https://openreview.net/forum?id=CD9Snc73AW>. Expert Certification. (Cited on page 3.)
- D. Trevisan. Well-posedness of multidimensional diffusion processes with weakly differentiable coefficients. *Electronic Journal of Probability*, 21(22):1–41, 2016. doi: 10.1214/16-EJP4453. (Cited on page 36.)
- J. Weed and F. Bach. Sharp asymptotic and finite-sample rates of convergence of empirical measures in Wasserstein distance. *Bernoulli*, 25(4A):2620–2648, 2019. doi: 10.3150/18-BEJ1065. (Cited on page 77.)
- A. Wibisono, Y. Wu, and K. Y. Yang. Optimal score estimation via empirical Bayes smoothing. In *Proceedings of Thirty Seventh Conference on Learning Theory*, volume 247 of *Proceedings of Machine Learning Research*, pages 4958–4991. PMLR, 2024. URL <https://proceedings.mlr.press/v247/wibisono24a.html>. (Cited on page 10.)
- X. Xu, Z. Zhang, Y. Nakahira, G. Qu, and Y. Chi. Polynomial convergence of Riemannian diffusion models. In *The Fourteenth International Conference on Learning Representations*, 2026. URL <https://openreview.net/forum?id=1L0FR3UPhZ>. (Cited on page 6.)
- R. Yang, Z. Zhu, B. Jiang, C. Chen, and S. Li. Elucidating rectified flow with deterministic sampler: Polynomial discretization complexity for multi and one-step models. *arXiv preprint arXiv:2508.08735*, 2025. URL <https://arxiv.org/abs/2508.08735>. (Cited on page 19.)
- M. S. Zhang, S. Huan, J. Huang, N. M. Boffi, S. Chen, and S. Chewi. Sublinear iterations can suffice even for DDPMs. *arXiv preprint arXiv:2511.04844*, 2025. URL <https://arxiv.org/abs/2511.04844>. (Cited on page 19.)
- Z. Zhou and W. Liu. An error analysis of flow matching for deep generative modeling. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pages 78903–78932. PMLR, 2025. URL <https://proceedings.mlr.press/v267/zhou251.html>. (Cited on pages 3, 6, and 10.)




(a) SW2
 $t = 0.5$



(b) $t = 0$
 $t = 1$




(c) $t = 0.5$




(d) $t = 1$


Figure 2: ColorMNIST 1–6. (a) Clean-target SW2 for $\sigma \in \{0.01, 0.015\}$ (5,000 samples, five seeds; dashed: pure ODE). (b)–(d) Color evolution along the $\sigma = 0.01$ SDE at $t = 0, 0.5, 1$ (5,000 samples; dashed: clean target).




(a) SW2
 $t = 0.5$



(b) $t = 0$
 $t = 1$



(c) $t = 0.5$



(d) $t = 1$

Figure 3: Fashion-MNIST shoes. (a) Clean-target SW2 versus τ for $\sigma \in \{0.005, 0.01\}$ (5,000 samples, five seeds; dashed: pure ODE). (b)–(d) Color evolution along the $\sigma = 0.005$ SDE at $t = 0, 0.5, 1$ (5,000 samples; dashed: clean target).

Appendix

A Additional Experiment Details

A.1 Angular-Mode Experiment Details

Data-Generating Process. Let $C \in \{0, \dots, 7\}$ index eight Gaussian components equally spaced on a circle, with

$$\mu_c = 2(\cos(c\pi/4), \sin(c\pi/4))^\top, \quad P_c = \mathcal{N}(\mu_c, \sigma_Y^2 I_2), \quad \sigma_Y = 0.08.$$

The population assigns weight $1/16$ to each axis-aligned component $c \in \{0, 2, 4, 6\}$ and $3/16$ to each diagonal component $c \in \{1, 3, 5, 7\}$. Treatment assignment satisfies $\mathbb{P}(A = 1 \mid C = c) = 0.9$ for axis-aligned components and 0.1 for diagonal components. Bayes' rule therefore gives conditional weights $3/16$ and $1/16$, respectively, among observations with $A = 1$.

Taking $\mathbb{P}_{Y|C=c,A=1} = P_c$, the observed source and counterfactual target are

$$\begin{aligned} \mathbb{P}_{Y|A=1} &= \sum_{c \in \{0,2,4,6\}} \frac{3}{16} P_c + \sum_{c \in \{1,3,5,7\}} \frac{1}{16} P_c, \\ \int \mathbb{P}_{Y|C=c,A=1} d\mathbb{P}_C(c) &= \sum_{c \in \{0,2,4,6\}} \frac{1}{16} P_c + \sum_{c \in \{1,3,5,7\}} \frac{3}{16} P_c. \end{aligned}$$

Thus, the two distributions share the same components but assign opposite weights to the axis-aligned and diagonal groups.

Couplings and Evaluation. For fixed empirical source and target samples, we compare IID and entropic optimal-transport couplings at different regularization levels. Each coupling assigns mass to source–target pairs; after Gaussian smoothing, these weighted pairs determine the interpolation law and the exactly computable empirical fields

$$v_t(x) = \mathbb{E}[Y_1 - Y_0 \mid X_t = x], \quad s_t(x) = \nabla_x \log q_t(x).$$

These exact fields serve as references for the learned models. The velocity and score experiments use three-hidden-layer MLPs of width 128 and $\sigma = 0.12$.

To isolate discretization error, we generate 2,000 exact-field samples with 8 Euler steps and estimate the generated density by kernel density estimation to compute KL divergence. We set $\tau = 0.5\sqrt{\sigma D_5/d}$, following the scaling $\tau \asymp \sqrt{\sigma D_5/d}$ suggested by Theorem 1. The mode-mass experiment uses the closed-form EOT coupling with $\varepsilon_{\text{OT}} = 0.5$ and $\sigma = 0.02$, and compares 4- and 64-step ODE and SDE samplers.

A.2 Image-Benchmark Details and Additional Results

Digit ColorMNIST. Following Dance et al. (2026), we draw $C \sim \text{Unif}(0, 1)$ and color each MNIST image using RGB value $(C, 0, 1 - C)$, with the original image intensity controlling brightness. Treatment $A = 0$ selects digit 1, $A = 1$ selects digit 6, and

$$\mathbb{P}(A = 1 \mid C = c) = \frac{\exp\{5(c - \frac{1}{2})\}}{1 + \exp\{5(c - \frac{1}{2})\}}.$$

Conditional on A , an image is sampled with replacement from the corresponding MNIST class and colored according to C . We generate $n = 20,000$ observational samples and target $a = 1$. Because C remains uniform under $\text{do}(A = 1)$, the clean counterfactual target can be sampled directly.

We train smoothed models with $\sigma \in \{0.01, 0.015\}$ for 1,000 epochs and evaluate the SDE over a range of τ , with $\tau = 0$ corresponding to the ODE for the same smoothed model. The pure ODE baseline is the same doubly robust flow trained with $\sigma = 0$; its best checkpoint through epoch 1,000 occurs at epoch 700. All samplers use 50 Euler steps. Generated images are clipped to $[0, 1]$, and SW2 is computed against the clean target using 5,000 samples over five random seeds. The baseline and smoothed models use the same fitted propensity model, conditional-flow prediction oracle, and conditional-sample reservoir. For $\sigma = 0.01$, the epoch-1000 model performs best near $\tau = 0.075$.

Fashion-MNIST. We use $C \sim \text{Unif}(0, 1)$, with $A = 0$ for sneakers and $A = 1$ for ankle boots, and use the same treatment-assignment probability as in the digit experiment. Given $C = c$, the grayscale image is colored using

$$(0.15 + 0.75c, 0.10 + 0.55 \sin^2(\pi c), 0.85 - 0.65c).$$

Consequently, observed ankle boots tend to have larger C , whereas the clean $\text{do}(A = 1)$ target contains the full population color range.

We train models with $\sigma = 0.01$ and $\sigma = 0.005$, together with the pure ODE baseline ($\sigma = 0$), for 400 epochs. Checkpoints are saved every 20 epochs. For $\sigma = 0.01$, the selected velocity and score checkpoints are epochs 100 and 220; for $\sigma = 0.005$, they are epochs 100 and 120. The selected pure ODE checkpoint is epoch 140. We vary τ , and every evaluation uses 50 Euler steps, 5,000 generated samples, 256 random projection directions, and five evaluation seeds. The color-distribution plots use the red–blue statistic $\hat{x} = R/(R + B)$.

Figure 5 shows images before and after transport. The source contains observed ankle boots with color-selection bias; the generated samples remain realistic ankle boots while covering colors closer to the clean target population.

Figures 4 and 6 provide complementary results for the shirt experiment: the former reports sampling performance and color evolution, while the latter shows that the generated samples preserve shirt structure while changing their colors toward the target distribution.

B Additional Related-Work Details

The closest comparisons are (Liu et al., 2025, 2026). For stochastic sampling, (Liu et al., 2025) considers a more general stochastic interpolant with time-varying smoothing γ_t . Under a uniform Euler discretization, their complexity has leading dimension dependence $\tilde{\mathcal{O}}(d^2)$ and polynomial dependence on the endpoint cutoffs t_0 and $1 - t_N$. These cutoff parameters are introduced to avoid the singular behavior near the two endpoints and play a role analogous to the fixed smoothing level σ in our analysis. In comparison, our uniform-step result achieves near-linear $\tilde{\mathcal{O}}(d)$ dimension dependence with leading smoothing dependence of order $1/\sigma^2$. This comparison holds with displacement moment and target accuracy fixed, using $\tau^2 = \sigma D_5/d$ and controlled field error; D_5 itself can grow with dimension. Liu et al. (2025) further considers a specially designed nonuniform time schedule, which improves the polynomial endpoint dependence (on $1/t_0$ and $1/(1 - t_N)$) to logarithmic dependence while retaining a leading $\tilde{\mathcal{O}}(d^2)$ dimension dependence. Thus, compared with their uniform discretization, our result improves the dimension dependence, while their nonuniform schedule achieves sharper endpoint dependence through a more specialized discretization. Although their bounds contain source–target displacement moments, they also retain coupling-independent score terms and do not develop a coupling-quality interpretation analogous to our D_5 control. The




Figure 4: Fashion-MNIST shirts experiment. (a) Clean-target SW2 as a function of τ for the two smoothing levels; shaded regions show variation across five evaluation seeds, and the dashed line denotes the pure ODE baseline. (b)–(d) Color distributions at $t = 0, 0.5, 1$ along the $\sigma = 0.005$ SDE trajectory; dashed distributions denote the clean target.

companion ODE analysis (Liu et al., 2026) likewise gives a $\tilde{O}(d^2)$ dimension dependence for Euler discretization.

Diffusion models admit sharper guarantees by exploiting the specific Gaussian noising structure of the diffusion probability path. Under weak distributional assumptions, (Benton et al., 2024) and (Li and Yan, 2025) obtain near-linear dependence on d for reverse-diffusion sampling. More recent results establish sublinear dependence under additional regularity. Jiao et al. (2025) introduce a non-uniform regularity parameter L and obtain dimension dependence $\tilde{O}(\sqrt{d} \min\{\sqrt{d}, L\})$ for the standard DDPM sampler. They showed L can be small for some specific class of targets, but in general the bound could still be linear in d .

We remark that related works such as (Zhang et al., 2025; Li and Jiao, 2025) obtain sublinear dimension dependence using randomized midpoint schemes, rather than the Euler discretization considered here. Specifically, (Zhang et al., 2025) obtain $\tilde{O}(\tilde{\beta}_0 \sqrt{d} + M_2^2)$, which reduces to $\tilde{O}(\sqrt{d})$ under bounded regularity constants, while (Li and Jiao, 2025) obtain $\tilde{O}(Ld^{1/3})$. These improved dimension rates depend explicitly on score regularity constants such as $\tilde{\beta}_0$, M_2 , and L .

Table 1 compares our sampling bound with the most relevant existing results, highlighting the dimension dependence, assumptions, and whether the bound reflects the quality of the source–target coupling.

Finally, we mention two additional related works. Yang et al. (2025) extends the analysis of (Chen et al., 2023a) to the flow-matching setting using a similar predictor–corrector framework, under assumptions including L_2 estimation error, bounded support, and uniform regularity of the learned vector field. Su et al. (2025) also establishes a KL-divergence error bound for flow matching with a deterministic ODE sampler. However, their analysis relies on strong global-in-space




Figure 5: Qualitative examples from the observed source distribution (top), samples generated by the $\sigma = 0.005$ SDE (bottom). The generated samples preserve the boots shape and structure while showing clear, coherent color changes.

regularity assumptions on the density path and velocity fields and does not analyze numerical ODE discretization error.

B.1 Comparison With Dance et al. (2026)

The main difference from Algorithm 3 of (Dance et al., 2026) is that our training loss is a direct minibatch implementation of the doubly robust loss. That is, for a fixed treatment a , the plug-in term is evaluated for every sample in the minibatch, while the observed correction is included only for samples with $A_i = a$:

$$\frac{1}{|\mathcal{B}|} \sum_{i \in \mathcal{B}} \left[\hat{\ell}_i^a(v) + \frac{\mathbf{1}\{A_i = a\}}{\hat{\pi}_a(C_i)} \left\{ \ell_i^{\text{obs},a}(v) - \hat{\ell}_i^a(v) \right\} \right].$$

Note that the plug-in term is averaged over the full covariate distribution and the inverse-propensity term corrects only the observed treatment- a samples.




Figure 6: Qualitative examples from the observed source distribution (top), samples generated by the $\sigma = 0.005$ SDE (bottom). The generated samples preserve the shirt shape and structure while showing clear, coherent color changes.

In contrast, Algorithm 3 of (Dance et al., 2026) forms, for each treatment a , a subset $\mathcal{B}_a = \{i \in \mathcal{B} : A_i = a\}$ and accumulates a loss only over $i \in \mathcal{B}_a$. This is not a direct implementation of the standard DR loss: the plug-in term is then averaged over the covariate distribution of the observed treatment- a subsample rather than over the full minibatch ¹.

A minor difference is how the coupling is handled. (Dance et al., 2026) recomputes the coupling within each minibatch, whereas our analysis estimates a single kernel $\hat{Q}_n$ and uses it for the full dataset of size n . Using one fixed global kernel gives a well-defined empirical loss associated with a specific coupling.

Moreover, we remark that (Dance et al., 2026) implements an ODE sampler, whereas we consider an SDE sampler.

¹This description refers to Algorithm 3 as stated in the first arXiv version of (Dance et al., 2026); the discrepancy may simply reflect a typo or pseudocode-level issue rather than the intended implementation.

Table 1: Comparison with related finite-step sampling results. Dimension dependence suppresses logarithmic factors and holds accuracy and the other problem parameters fixed; for our result this includes σ, D_5 , with $\tau^2 = \sigma D_5/d$ and controlled field error. “Partial” means that source–target displacement enters the bound, but coupling-independent score terms remain.

Work	Setting / Method	Metric	Dim.	Assumption	Coupling control
<i>Stochastic interpolants and flow matching</i>					
This work Main theorem	constant-smoothing SI SDE, Euler	KL	$\tilde{\mathcal{O}}(d)$	L_2 error	D_5
(Liu et al., 2025) Thm. 4.3; Sec. 5	SI SDE, Euler	KL	$\tilde{\mathcal{O}}(d^2)$	L_2 error	Partial
(Liu et al., 2026) Thm. 4.5	SI ODE, Euler	TV	$\tilde{\mathcal{O}}(d^2)$	drift/div. errors; unif. reg.	No
<i>Diffusion-model baselines</i>					
(Benton et al., 2024) Cor. 1	diffusion exp. integrator	KL	$\tilde{\mathcal{O}}(d)$	L_2 error	–
(Li and Yan, 2025) Thm. 1	DDPM	TV	$\tilde{\mathcal{O}}(d)$	L_2 error	–
(Jiao et al., 2025) Thm. 1	DDPM	KL	$\tilde{\mathcal{O}}(\sqrt{d})$ $\min\{\sqrt{d}, L\}$	weighted L_2 error; nonuniform reg.	–

C Reader’s Guide and Proof Roadmap

This section lists the subclaims needed for the end-to-end result and points to the proof of each one.

Appendix Organization. Appendix A provides experimental details and additional results. Appendix B gives additional related-work details first. We then collect the notation and algorithmic details, give the proofs of Examples 1 and 2, and present the proof of Theorem 1, the coupling and field-estimation analyses, and the proof of Theorem 2.

Motivating Examples. The proofs of Examples 1 and 2 are given in Appendices E and F, respectively.

Sampling Theorem. Theorem 1 follows from the following chain.

1. *Marginal Correctness of the Exact SDE.* Lemma 1 shows that q_t solves the Fokker–Planck equation of the score-corrected SDE and satisfies the integrability and continuity conditions required by the superposition principle. It therefore yields a weak solution Z with $Z_t \sim q_t$ for every $t \in [0, 1]$.
2. *Girsanov KL bound for the Euler process.* Lemma 2 localizes the exact-to-Euler drift perturbation, applies Girsanov before the stopping time, and passes to the unstopped process by weak convergence and lower semicontinuity of KL. Proposition 3 verifies the pathwise integrability condition when the drifts and sample paths are continuous.

3. *Posterior Derivative Formulas.* Lemmas 3 and 4 express the score, velocity Jacobian, velocity Hessian, score Jacobian, and score Laplacian as conditional moments of $D = Y_1 - Y_0$ and the Gaussian variable G .
4. *Integrated Derivative Bounds.* Lemmas 5 and 6 convert the posterior identities into $L^2(q_t)$ bounds using conditional variance contraction, Cauchy–Schwarz, and explicit Gaussian moments.
5. *Euler Increment Bounds.* Lemmas 7 and 8 apply Itô’s formula along the exact SDE. The score proof uses the continuity equation to cancel $D_x s_t v_t$, leaving only the Hessian, mixed velocity–score, and pure score terms controlled by the integrated derivative bounds.
6. *Collection of Constants.* The proof of Theorem 1 substitutes the Euler increment bounds into the Girsanov reduction and records the explicit coefficients A_h and B_h , together with the explicit τ^4 , τ^6 , and displacement terms. Corollary 1 then adds the Gaussian smoothing error and converts KL to W_1 using a maximal coupling, Pinsker’s inequality, and a second-moment bound for the Euler recursion.

Field-Estimation Propositions. The proof of Proposition 1 first derives Proposition 5 and then verifies its hypotheses for the velocity and score losses.

1. *Coupling Construction and Replacement Error.* Lemma 9 proves Lipschitz stability of the EOT conditional kernel. Lemma 10 constructs an ideal kernel Q with the required source and target marginals, and Corollary 3 controls the joint-law discrepancy between Q and $\hat{Q}_n$.
2. *Exact Excess-Risk Identities.* For both fields, conditional projection identifies excess population loss with the grid-averaged squared $L^2(q_{t_k})$ field error. The score proof uses the conditional-mean parametrization $s_u(x, t) = \{u(x, t) - x\}/\sigma^2$, for which the Gaussian term cancels from the loss.
3. *Approximation.* Propositions 6 and 7 construct network approximations uniformly on a bounded cube and control the Gaussian tail outside the cube.
4. *Loss Perturbations.* The coupling-replacement error is controlled using Corollary 3, together with the endpoint Lipschitz bound for the loss. For the doubly robust error, we first write the population bias as $\mathbb{E}_C[(1 - \hat{\alpha}/\alpha_a)(m_v - \hat{m}_v)]$. We then control $m_v - \hat{m}_v$ through the conditional-generator W_1 error and the stability of $\hat{Q}_n$, and apply Cauchy–Schwarz to obtain the final product bound.
5. *Statistical Generalization.* Lemma 13 controls the empirical-to-population difference for any bounded loss class from a proper finite cover. The two field analyses supply their own network covering numbers and loss-to-network transfer constants.
6. *Monte Carlo Approximation.* We first decompose the implementation error according to the structure of the doubly robust loss and define separate observed-correction and auxiliary Monte Carlo classes. For each class, the corresponding error can be written as the difference between a conditional expectation and its realized Monte Carlo average. Lemma 13 is then applied conditionally to each array. A key step is to condition on the training observations and fitted auxiliary objects, which gives the conditional independence required in Proposition 5.
7. *End-to-End Substitution.* Combining the preceding velocity- and score-field steps gives the integrated errors $\varepsilon_{\text{vel}}^2$ and $\varepsilon_{\text{score}}^2$. Substitution into the explicit sampling theorem, followed by the Wasserstein conversion, proves Theorem 2.

Technical Novelty. Our proof makes two main contributions. First, for fixed Gaussian smoothing, we obtain near-linear rather than quadratic dimension dependence in the number of uniform Euler steps, with D_5 , σ , and target accuracy held fixed. We bound velocity derivatives and mixed velocity–score terms using displacement moments, and keep the noise factors in the score terms explicit. With exact fields, choosing $\tau^2 = \sigma D_5/d$ then reduces the dimension dependence and makes every discretization term shrink with D_5 , including the pure-score terms. This gives both improved dimension dependence and full coupling-sensitive control of the discretization error.

Second, we give an end-to-end analysis of counterfactual generation with a learned source–target coupling, rather than only independent pairing. We analyze a doubly robust loss that combines prediction-based losses with weighted corrections from observed outcomes, rather than the usual flow-matching loss. Because the training loss itself is approximated using finitely many random draws from the prediction oracle and the coupling, we bound the error caused by this Monte Carlo approximation. We also separate this error from the finite-data error and the Monte Carlo error in the target samples used to learn the coupling. Combining these learning bounds with the Euler and smoothing errors gives a guarantee for the final output distribution.

C.1 Summary of Notation

Table 2 summarizes the notation used throughout the paper.

Table 2: Summary of notation used throughout the paper.

Notation	Meaning	Notation	Meaning
(C, A, Y)	Observed covariate, treatment, and outcome.	$\mathcal{V}_n, \mathcal{S}_n$	Clipped sparse-ReLU velocity class and score class induced by clipped sparse-ReLU denoisers.
a	Fixed target treatment value.	R, ε	Network input-clipping radius and approximation tolerance; field-specific choices appear in Proposition 1.
$\mathbb{P}_{Y C=c, A=a}$	Conditional outcome law at covariate value c under treatment a .	$\varepsilon_{\text{vel}}, \varepsilon_{\text{score}}$	Integrated estimation errors of the learned velocity and score fields.
μ_1	Target law $\int \mathbb{P}_{Y C=c, A=a} d\mathbb{P}_C(c)$.	τ	Noise level of the score-corrected SDE sampler.
$\mathbb{P}_{Y A=a}$	Observed outcome law among samples with treatment a .	N, h	Number of sampling steps and Euler step size, with $h = 1/N$.
$\mathcal{D}_{\text{tr}}, \mathcal{D}_{\text{aux}}$	Training and auxiliary splits, each containing n observations.	t_k	Sampling grid point $t_k = kh$.
$\hat{\mathbb{P}}_{Y C=c, A=a}$	Learned conditional outcome law, or prediction oracle.	M	Monte Carlo size in the final doubly robust training objective.
$\hat{\mathbb{P}}_{Y A=a}^{\text{aux}}$	Empirical source law formed from treatment- a outcomes in $\mathcal{D}_{\text{aux}}$.	M_Q	Prediction-oracle samples per auxiliary covariate used to construct $\hat{\mathbb{P}}_{Y_1}^Q$.
$\hat{\mathbb{P}}_{Y_1}^Q$	Finite-sample fitted-target approximation based on the prediction oracle.	$\alpha_a(c), \hat{\alpha}(c)$	True and estimated inverse propensities for treatment a .
$\hat{Q}_n(dy_0 y_1)$	Estimated coupling kernel between the empirical source and fitted target.	$\pi_{\min}$	Lower bound on $\mathbb{P}(A = a C = c)$.
$Q(dy_0 y_1), \Gamma_{\mu_1, Q}$	Population-valid kernel and joint end-point law $\Gamma_{\mu_1, Q} = Q\mu_1$.	ε_α	Estimation error of the inverse propensity model.
(Y_0, Y_1)	Source–target pair under the specified coupling.	$\varepsilon_{\text{pred, train}}$	Empirical prediction-oracle error on auxiliary covariates.
D_5	Coupling displacement moment $(\mathbb{E}\ Y_1 - Y_0\ ^5)^{1/5}$.	$\varepsilon_{\text{pred, gen}}$	Prediction-oracle generalization error.
S_t	Linear interpolation $Y_0 + t(Y_1 - Y_0)$.	ε_{OT}	Entropic regularization parameter for EOT.
X_t	Smoothed interpolation $S_t + \sigma G$, $G \sim \mathcal{N}(0, I_d)$.	$\mathcal{E}_a$	Event that $\mathcal{D}_{\text{aux}}$ contains at least one observation with $A_i = a$.

σ	Gaussian smoothing level.	$\text{supp}(P)$	Topological support: x such that $P(B(x, r)) > 0$ for every $r > 0$.
v_t, s_t	Population velocity and score fields for the smoothed interpolation.	d	Outcome dimension.
$q_{t,Q}, v_Q^*, s_Q^*$	Density, exact velocity, and exact score induced by $\Gamma_{\mu_1,Q}$.	$W_1, W_{1,\rho}$	Wasserstein distances on the outcome and coupled endpoint spaces.
$\hat{v}_M, \hat{s}_M$	Learned doubly robust velocity and score fields.	KL	Kullback–Leibler divergence.

Conditioning Convention. In the finite-sample analysis, we first condition on $\mathcal{D}_{\text{aux}}$ and the prediction-oracle target reservoir. The fitted nuisance models, $\hat{Q}_n, Q, \Gamma_{\mu_1,Q}, q_{t,Q}, v_Q^*, s_Q^*$, and the corresponding D_5 are then interpreted conditionally. Unconditional high-probability statements include the corresponding auxiliary-sample events in their failure probabilities.

D Algorithms and Auxiliary-Model Details

D.1 Prediction-Oracle Accuracy

We give sufficient conditions for Assumption 4. The first controls the prediction-oracle error on the auxiliary covariates used to construct the coupling.

Assumption 5 (Empirical Generator Accuracy on Auxiliary Covariates). *The learned conditional outcome model satisfies*

$$\left[\frac{1}{n} \sum_{i=n+1}^{2n} W_1 \left(\mathbb{P}_{Y|C=C^i, A=a}, \hat{\mathbb{P}}_{Y|C=C^i, A=a} \right)^2 \right]^{1/2} \leq \varepsilon_{\text{pred,train}}.$$

This assumption measures how well the learned conditional outcome model fits the conditional laws at the empirical auxiliary covariates. It is not a population- C quantity.

Assumption 6 (Generator Generalization Gap). *Let Ψ_n denote the class for prediction oracle, and let $\hat{\mathbb{P}}_{Y|C=c, A=a}^\psi$ be the conditional law generated by ψ . Assume*

$$\sup_{\psi \in \Psi_n} \left| \mathbb{E}_C \left[W_1 \left(\mathbb{P}_{Y|C=c, A=a}, \hat{\mathbb{P}}_{Y|C=c, A=a}^\psi \right)^2 \right] - \frac{1}{n} \sum_{i=n+1}^{2n} W_1 \left(\mathbb{P}_{Y|C=C^i, A=a}, \hat{\mathbb{P}}_{Y|C=C^i, A=a}^\psi \right)^2 \right| \leq \varepsilon_{\text{pred,gen}}^2.$$

Together, these two assumptions imply the population prediction-oracle accuracy used in the main text.

Proposition 2 (Population Generator Accuracy From Empirical Accuracy). *Under Assumptions 5 and 6,*

$$\left[\mathbb{E}_C W_1 \left(\mathbb{P}_{Y|C=c, A=a}, \hat{\mathbb{P}}_{Y|C=c, A=a} \right)^2 \right]^{1/2} \leq [(\varepsilon_{\text{pred,train}})^2 + \varepsilon_{\text{pred,gen}}^2]^{1/2}.$$

Hence Assumption 4 holds.

Proof. Since Assumption 6 holds uniformly over $\psi \in \Psi_n$, it also applies to the learned prediction oracle. Therefore,

$$\mathbb{E}_C W_1 \left(\mathbb{P}_{Y|C=c, A=a}, \hat{\mathbb{P}}_{Y|C=c, A=a} \right)^2 \leq \frac{1}{n} \sum_{i=n+1}^{2n} W_1 \left(\mathbb{P}_{Y|C=C^i, A=a}, \hat{\mathbb{P}}_{Y|C=C^i, A=a} \right)^2 + \varepsilon_{\text{pred,gen}}^2$$

$$\leq (\varepsilon_{\text{pred,train}})^2 + \varepsilon_{\text{pred,gen}}^2.$$

Taking square roots proves the claim. $\square$

The prediction oracle $\hat{\mathbb{P}}_{Y|C=c,A=a}$ can be obtained by training an auxiliary conditional flow model on $\mathcal{D}_{\text{aux}}$. For each observation (Y_i, C_i, A_i) , conditional on (C_i, A_i) , Y_i is a draw from $\mathbb{P}_{Y|C=C_i,A=A_i}$. Let P_0 be a source distribution, for example $\mathcal{N}(0, I_d)$. The population flow-matching loss is

$$\mathcal{L}_{\text{aux}}(u) = \mathbb{E}_{C,A,Y} \mathbb{E}_{Z \sim P_0, t \sim \text{Unif}[0,1]} [\|u(X_t, t, C, A) - (Y - Z)\|^2], \quad X_t = (1-t)Z + tY.$$

The auxiliary observations (C_i, A_i, Y_i) , $i = n+1, \dots, 2n$, provide samples for training this model. After training, fixing $A = a$ and conditioning on $C = c$ gives the prediction oracle $\hat{\mathbb{P}}_{Y|C=c,A=a}$.

D.2 Score-Field Training Objective

We now explain how the score field in the score-corrected SDE is learned. For each $t \in [0, 1]$, let q_t denote the density of X_t and define $s_t(x) = \nabla_x \log q_t(x)$. The corresponding population score-matching loss is

$$\mathcal{L}_{\text{score,true}}(\phi) = \int_0^1 \mathbb{E} [\|s_\phi(X_t, t) - s_t(X_t)\|^2] dt.$$

This loss cannot be evaluated directly because s_t is unknown. By the Gaussian score identity proved in Lemma 3, $s_t(X_t) = \mathbb{E}[-G/\sigma \mid X_t]$. Therefore, $-G/\sigma = (S_t - X_t)/\sigma^2$ can be used as a regression target for the score. We define the denoising score-matching loss by

$$\mathcal{L}_{\text{score}}(\phi) = \int_0^1 \mathbb{E}_{(Y_0, Y_1) \sim \gamma} \mathbb{E}_{G \sim \mathcal{N}(0, I_d)} \left[\left\| s_\phi(X_t, t) + \frac{G}{\sigma} \right\|^2 \right] dt = \int_0^1 \mathbb{E} \left[\left\| s_\phi(X_t, t) - \frac{S_t - X_t}{\sigma^2} \right\|^2 \right] dt.$$

The equivalence follows from the conditional Pythagorean identity. For each fixed t , set $U = -G/\sigma$. Since $s_t(X_t) = \mathbb{E}[U \mid X_t]$, every square-integrable candidate $s_\phi(t, \cdot)$ satisfies

$$\mathbb{E} \|s_\phi(X_t, t) - U\|^2 = \mathbb{E} \|s_\phi(X_t, t) - s_t(X_t)\|^2 + \mathbb{E} \|U - s_t(X_t)\|^2.$$

The second term is independent of ϕ . Integrating over t proves that the denoising and true-score losses have the same minimizers. Since $X_t - S_t = \sigma G$, the implementable objective may equivalently be written as

$$\mathcal{L}_{\text{score}}(\phi) = \int_0^1 \mathbb{E} \left[\left\| s_\phi(X_t, t) - \frac{S_t - X_t}{\sigma^2} \right\|^2 \right] dt.$$

D.3 Algorithms

Algorithm 1 gives the finite-step score-corrected sampler, while Algorithm 2 gives the doubly robust procedure used to train its velocity and score fields.

Algorithm 1 Score-Corrected SDE Sampler

Input: Learned velocity field $\hat{v}$, learned score field $\hat{s}$, empirical source distribution $\hat{\mathbb{P}}_{Y|A=a}^{\text{aux}}$, smoothing level $\sigma > 0$, SDE noise level $\tau > 0$, number of time steps N .

- 1: Set $h = 1/N$ and $t_k = kh$, $k = 0, \dots, N$.
 - 2: Draw $Y_0^{\text{init}} \sim \hat{\mathbb{P}}_{Y|A=a}^{\text{aux}}$ and $G^{\text{init}} \sim \mathcal{N}(0, I_d)$.
 - 3: Initialize $\hat{Z}_0 = Y_0^{\text{init}} + \sigma G^{\text{init}}$.
 - 4: **for** $k = 0, \dots, N - 1$ **do**
 - 5: Draw $G_k \sim \mathcal{N}(0, I_d)$.
 - 6: Update $\hat{Z}_{k+1} = \hat{Z}_k + h(\hat{v}(\hat{Z}_k, t_k) + \frac{\tau^2}{2}\hat{s}(\hat{Z}_k, t_k)) + \tau\sqrt{h}G_k$.
 - 7: **end for**
 - 8: **return** $\hat{Z}_N$.
-

Algorithm 2 Minibatch Doubly Robust Flow Training for a Fixed Treatment a

Input: Training data $\{(C_i, A_i, Y_i)\}_{i=1}^n$ for the final flow model, and training data $\{(C_i, A_i, Y_i)\}_{i=n+1}^{2n}$ for the following models: propensity estimator $\hat{\mathbb{P}}(A = a \mid C = c)$, conditional outcome sampler $\hat{\mathbb{P}}_{Y \mid C=c, A=a}$, estimated coupling kernel $\hat{Q}_n(dy_0 \mid y_1)$.

Input: Vector-field class $\mathcal{V}_n$, score-field class $\mathcal{S}_n$, noise level $\sigma > 0$, Monte Carlo size M .

- 1: Initialize $v \in \mathcal{V}_n$ and $s \in \mathcal{S}_n$.
- 2: **for** $i \in \{1, \dots, n\}$ **do**
- 3: **for** $k = 1, \dots, M$ **do**
- 4: Draw $\hat{Y}_{i,k} \sim \hat{\mathbb{P}}_{Y \mid C=C_i, A=a}$.
- 5: Draw $\hat{Y}_{0,i,k} \sim \hat{Q}_n(\cdot \mid \hat{Y}_{i,k})$, $t_{i,k} \sim \text{Unif}\{t_0, \dots, t_{N-1}\}$, and $G_{i,k} \sim \mathcal{N}(0, I_d)$.
- 6: **end for**
- 7: **if** $A_i = a$ **then**
- 8: Draw $\hat{Y}_{0,i}^{\text{obs}} \sim \hat{Q}_n(\cdot \mid Y_i)$, $t_i^{\text{obs}} \sim \text{Unif}\{t_0, \dots, t_{N-1}\}$, and $G_i^{\text{obs}} \sim \mathcal{N}(0, I_d)$.
- 9: **end if**
- 10: **end for**
- 11: **for** each training iteration **do**
- 12: Draw a minibatch $\mathcal{B} \subseteq \{1, \dots, n\}$.
- 13: Compute the minibatch doubly robust flow loss

$$\hat{L}_{\mathcal{B}, \text{DR}}(v) = \frac{1}{|\mathcal{B}|} \sum_{i \in \mathcal{B}} \left[\frac{\mathbf{1}\{A_i = a\}}{\hat{\mathbb{P}}(A = a \mid C = C_i)} \left\{ \hat{\ell}_i^{\text{obs}}(v) - \frac{1}{M} \sum_{k=1}^M \hat{\ell}_{i,k}^{\text{plug}}(v) \right\} + \frac{1}{M} \sum_{k=1}^M \hat{\ell}_{i,k}^{\text{plug}}(v) \right],$$

where

$$\begin{aligned} \hat{\ell}_i^{\text{obs}}(v) &= \left\| v \left(\hat{Y}_{0,i}^{\text{obs}} + t_i^{\text{obs}} (Y_i - \hat{Y}_{0,i}^{\text{obs}}) + \sigma G_i^{\text{obs}}, t_i^{\text{obs}} \right) - (Y_i - \hat{Y}_{0,i}^{\text{obs}}) \right\|^2, \\ \hat{\ell}_{i,k}^{\text{plug}}(v) &= \left\| v \left(\hat{Y}_{0,i,k} + t_{i,k} (\hat{Y}_{i,k} - \hat{Y}_{0,i,k}) + \sigma G_{i,k}, t_{i,k} \right) - (\hat{Y}_{i,k} - \hat{Y}_{0,i,k}) \right\|^2. \end{aligned}$$

- 14: Update v by one optimizer step using $\hat{L}_{\mathcal{B}, \text{DR}}(v)$.
- 15: Compute the minibatch doubly robust score loss

$$\hat{L}_{\mathcal{B}, \text{DR}}^{\text{score}}(s) = \frac{1}{|\mathcal{B}|} \sum_{i \in \mathcal{B}} \left[\frac{\mathbf{1}\{A_i = a\}}{\hat{\mathbb{P}}(A = a \mid C = C_i)} \left\{ \hat{r}_i^{\text{obs}}(s) - \frac{1}{M} \sum_{k=1}^M \hat{r}_{i,k}^{\text{plug}}(s) \right\} + \frac{1}{M} \sum_{k=1}^M \hat{r}_{i,k}^{\text{plug}}(s) \right],$$

where

$$\begin{aligned} \hat{r}_i^{\text{obs}}(s) &= \left\| s \left(\hat{Y}_{0,i}^{\text{obs}} + t_i^{\text{obs}} (Y_i - \hat{Y}_{0,i}^{\text{obs}}) + \sigma G_i^{\text{obs}}, t_i^{\text{obs}} \right) + \frac{G_i^{\text{obs}}}{\sigma} \right\|^2, \\ \hat{r}_{i,k}^{\text{plug}}(s) &= \left\| s \left(\hat{Y}_{0,i,k} + t_{i,k} (\hat{Y}_{i,k} - \hat{Y}_{0,i,k}) + \sigma G_{i,k}, t_{i,k} \right) + \frac{G_{i,k}}{\sigma} \right\|^2. \end{aligned}$$

- 16: Update s by one optimizer step using $\hat{L}_{\mathcal{B}, \text{DR}}^{\text{score}}(s)$.
 - 17: **end for**
 - 18: **return** v and s .
-

E Proof of Example 1

Proof of Example 1. In this proof, write $\hat{Z}_k$ for the Euler ODE iterates, so $\hat{Z}_N = \hat{X}_N^{\text{ODE}}$ in the example. Starting from $\hat{Z}_0 \stackrel{d}{=} X_0$, the recursion is

$$\hat{Z}_{k+1} = \hat{Z}_k + hv_{kh}(\hat{Z}_k), \quad k = 0, \dots, N-1.$$

Here N is the number of Euler steps, and $h = 1/N$ is the step size.

Write $D = Y_1 - Y_0$. We have $v_t(x) = \mathbb{E}[D \mid X_t = x]$. Let φ_V denote the density of $\mathcal{N}(0, V)$. Conditional on $B = 0$ we have $X_0 \sim \mathcal{N}(0, 1 + \sigma^2)$ and conditional on $B = 1$ we have $X_0 \sim \mathcal{N}(0, 4 + \sigma^2)$. Bayes' rule implies

$$\mathbb{P}(B = 1 \mid X_0 = x) = \frac{p\varphi_{4+\sigma^2}(x)}{(1-p)\varphi_{1+\sigma^2}(x) + p\varphi_{4+\sigma^2}(x)}.$$

Since $D = 0$ on $\{B = 0\}$,

$$\begin{aligned} v_0(x) &= \mathbb{P}(B = 1 \mid X_0 = x)\mathbb{E}[D \mid X_0 = x, B = 1] + \mathbb{P}(B = 0 \mid X_0 = x)\mathbb{E}[D \mid X_0 = x, B = 0] \\ &= \mathbb{P}(B = 1 \mid X_0 = x)\mathbb{E}[D \mid X_0 = x, B = 1]. \end{aligned}$$

We have that

$$\mathbb{E}[U \mid 2U + \sigma G = x] = \int_{-\infty}^{\infty} u P_{u|2u+\sigma G=x}(u) du.$$

To compute $P_{u|2u+\sigma G=x}(u)$, we use the formula $p_{X|Y=y}(x) = \frac{p_{X,Y}(x,y)}{p_Y(y)}$. We get

$$p_{U|2U+\sigma G=x}(u) = \frac{p_{U,2U+\sigma G}(u,x)}{p_{2U+\sigma G}(x)} \propto p_{U,2U+\sigma G}(u,x) = \frac{1}{\sigma} p_{U,G}(u, \frac{x-2u}{\sigma}),$$

where we used change of variable formula. Then we have

$$p_{U|2U+\sigma G=x}(u) \propto \frac{1}{\sigma} \exp(-u^2/2 - (x-2u)^2/(2\sigma^2)) \propto \exp(-\frac{(u - \frac{2x}{\sigma^2+4})^2}{2\sigma^2/(\sigma^2+4)}).$$

Hence $\mathbb{E}[U \mid 2U + \sigma G = x] = \frac{2}{4+\sigma^2}x$. Using the definition of λ ,

$$\mathbb{E}[D \mid X_0 = x, B = 1] = -2\lambda \frac{2}{4 + \sigma^2}x = -\frac{x}{h},$$

so we have $v_0(x) = -\frac{\mathbb{P}(B=1|X_0=x)}{h}x$.

We next examine the uniform spatial regularity of the population velocity. By the expression above for $\mathbb{P}(B = 1 \mid X_0 = x)$,

$$\begin{aligned} \mathbb{P}(B = 1 \mid X_0 = x) &= \frac{p\varphi_{4+\sigma^2}(x)}{(1-p)\varphi_{1+\sigma^2}(x) + p\varphi_{4+\sigma^2}(x)} = \left(1 + \frac{1-p}{p} \frac{\varphi_{1+\sigma^2}(x)}{\varphi_{4+\sigma^2}(x)}\right)^{-1} \\ &= \left(1 + \frac{1-p}{p} \sqrt{\frac{4+\sigma^2}{1+\sigma^2}} \exp\left\{-\frac{x^2}{2(1+\sigma^2)} + \frac{x^2}{2(4+\sigma^2)}\right\}\right)^{-1} \\ &= \left(1 + \frac{1-p}{p} \sqrt{\frac{4+\sigma^2}{1+\sigma^2}} \exp\left\{-\frac{3x^2}{2(1+\sigma^2)(4+\sigma^2)}\right\}\right)^{-1}. \end{aligned}$$

Therefore,

$$v_0(x) = -\frac{x}{h} \left(1 + \frac{1-p}{p} \sqrt{\frac{4+\sigma^2}{1+\sigma^2}} \exp \left\{ -\frac{3x^2}{2(1+\sigma^2)(4+\sigma^2)} \right\} \right)^{-1}.$$

Differentiating with respect to x gives

$$\begin{aligned} \partial_x v_0(x) = & -\frac{1}{h} \left(1 + \frac{1-p}{p} \sqrt{\frac{4+\sigma^2}{1+\sigma^2}} \exp \left\{ -\frac{3x^2}{2(1+\sigma^2)(4+\sigma^2)} \right\} \right)^{-1} \\ & - \frac{3x^2}{h(1+\sigma^2)(4+\sigma^2)} \frac{\frac{1-p}{p} \sqrt{\frac{4+\sigma^2}{1+\sigma^2}} \exp \left\{ -\frac{3x^2}{2(1+\sigma^2)(4+\sigma^2)} \right\}}{\left(1 + \frac{1-p}{p} \sqrt{\frac{4+\sigma^2}{1+\sigma^2}} \exp \left\{ -\frac{3x^2}{2(1+\sigma^2)(4+\sigma^2)} \right\} \right)^2}. \end{aligned}$$

For all sufficiently small p , we have

$$\frac{1-p}{p} \sqrt{\frac{4+\sigma^2}{1+\sigma^2}} > 1.$$

Choose $x > 0$ such that $x^2 = \frac{2(1+\sigma^2)(4+\sigma^2)}{3} \log \left(\frac{1-p}{p} \sqrt{\frac{4+\sigma^2}{1+\sigma^2}} \right)$, i.e., $\frac{1-p}{p} \sqrt{\frac{4+\sigma^2}{1+\sigma^2}} \exp \left\{ -\frac{3x^2}{2(1+\sigma^2)(4+\sigma^2)} \right\} = 1$. Substituting this value of x into the expression for $\partial_x v_0(x)$ gives

$$\partial_x v_0(x) = -\frac{1}{2h} - \frac{3x^2}{4h(1+\sigma^2)(4+\sigma^2)} = -\frac{1}{2h} \left[1 + \log \left(\frac{1-p}{p} \sqrt{\frac{4+\sigma^2}{1+\sigma^2}} \right) \right].$$

Consequently,

$$\sup_{t \in [0,1]} \sup_{x \in \mathbb{R}} |\partial_x v_t(x)| \geq \sup_{x \in \mathbb{R}} |\partial_x v_0(x)| \geq \frac{1}{2h} \left[1 + \log \left(\frac{1-p}{p} \sqrt{\frac{4+\sigma^2}{1+\sigma^2}} \right) \right].$$

By construction, $D = 0$ on $\{B = 0\}$ and $D = -2\lambda U$ on $\{B = 1\}$. Hence $D_5 = (\mathbb{E}|D|^5)^{1/5} = 2\lambda(\mathbb{E}|U|^5)^{1/5} p^{1/5}$. So we conclude that when $p \rightarrow 0$, $\sup_{t \in [0,1]} \sup_{x \in \mathbb{R}} |\partial_x v_t(x)| \rightarrow +\infty$ but $D_5 \rightarrow 0$.

For the integrated derivative claim, keep h, σ fixed, so that λ is fixed. Since $\mathbb{E}|D|^4 = 16\lambda^4 p \mathbb{E}|U|^4 = 48\lambda^4 p$, Lemma 6 with $d = 1$ gives

$$\int_0^1 \mathbb{E} |\partial_x v_t(X_t)|^2 dt \leq \frac{\sqrt{3}}{\sigma^2} (48\lambda^4 p)^{1/2} = \frac{12\lambda^2}{\sigma^2} \sqrt{p} \rightarrow 0.$$

Next, we show the sampling KL error would blow up. The first Euler update is the deterministic map

$$\hat{Z}_1 = T_p(\hat{Z}_0), \quad T_p(x) = x + h v_0(x) = x(1 - \mathbb{P}(B = 1 \mid X_0 = x)).$$

Moreover,

$$\begin{aligned} 1 - \mathbb{P}(B = 1 \mid X_0 = x) & \leq \frac{1-p}{p} \frac{\varphi_{1+\sigma^2}(x)}{\varphi_{4+\sigma^2}(x)} \\ & = \frac{1-p}{p} \sqrt{\frac{4+\sigma^2}{1+\sigma^2}} \exp \left\{ -\frac{x^2}{2} \left(\frac{1}{1+\sigma^2} - \frac{1}{4+\sigma^2} \right) \right\}. \end{aligned}$$

Since T_p is continuous, its image is bounded. Thus $\hat{Z}_1$ is supported on a compact interval.

For $v_t(x)$, we change the condition $X_0 = 2U + \sigma G = x$ to $X_t = 2U(1 - \lambda t) + \sigma G = x$.

$$\mathbb{E}[D \mid X_t = x, B = 1] = -\frac{4\lambda(1 - \lambda t)}{4(1 - \lambda t)^2 + \sigma^2}x,$$

so

$$v_t(x) = -\frac{4\lambda(1 - \lambda t)}{4(1 - \lambda t)^2 + \sigma^2} \frac{p\varphi_{4(1-\lambda t)^2+\sigma^2}(x)}{(1-p)\varphi_{1+\sigma^2}(x) + p\varphi_{4(1-\lambda t)^2+\sigma^2}(x)}x.$$

Note that $|\frac{4\lambda(1-\lambda t)}{4(1-\lambda t)^2+\sigma^2}| \leq |\frac{4\lambda(1-\lambda t)}{\sigma^2}|$ and $\frac{p\varphi_{4(1-\lambda t)^2+\sigma^2}(x)}{(1-p)\varphi_{1+\sigma^2}(x) + p\varphi_{4(1-\lambda t)^2+\sigma^2}(x)} \leq 1$ are uniformly bounded. Consequently,

$$x \mapsto x + hv_{kh}(x)$$

is continuous for every $k = 0, \dots, N-1$. Since a continuous image of a compact set is compact, $\hat{Z}_N$ is supported on some compact set.

On the other hand, the density $q_1(x)$ of X_1 is strictly positive for every $x \in \mathbb{R}$. Therefore, $\mathcal{L}(X_1)$ is not absolutely continuous with respect to $\mathcal{L}(\hat{Z}_N)$, which proves

$$\text{KL}(\mathcal{L}(X_1) \parallel \mathcal{L}(\hat{Z}_N)) = +\infty.$$

The result on integrated regularity follows from Lemma 6. □

F Proof of Example 2

Proof of Example 2. We first compute the coupling displacement. Conditional on $Y_0 = x$, the displacement $Y_1 - Y_0$ is equal to δ , $-\delta$, or 0, with respective probabilities $p_m(x)$, $p_m(x)$, and $1 - 2p_m(x)$. Therefore

$$\mathbb{E}[|Y_1 - Y_0|^5 \mid Y_0 = x] = 2p_m(x)\delta^5 = \delta^5 \left(\frac{1}{2} + \frac{1}{4} \sin(K_m x) \right).$$

Since Y_0 is symmetric about zero and $x \mapsto \sin(K_m x)$ is odd, $\mathbb{E}[\sin(K_m Y_0)] = 0$. Hence $\mathbb{E}|Y_1 - Y_0|^5 = \delta^5/2$, and therefore

$$D_5 = \frac{\delta}{2^{1/5}} = \varepsilon.$$

In particular, D_5 is independent of m even though $K_m \rightarrow \infty$.

We next verify that the score is well defined despite the absence of smoothing. Let

$$\phi(x) = \frac{1}{\sqrt{2\pi}} e^{-x^2/2}$$

denote the standard Gaussian density. Conditional on the three possible values of $Y_1 - Y_0$, the interpolation satisfies

$$Y_t = \begin{cases} Y_0 + t\delta, & Y_1 - Y_0 = \delta, \\ Y_0 - t\delta, & Y_1 - Y_0 = -\delta, \\ Y_0, & Y_1 - Y_0 = 0. \end{cases}$$

It follows that

$$q_t^{(m)}(y) = \phi(y - t\delta)p_m(y - t\delta) + \phi(y + t\delta)p_m(y + t\delta) + \phi(y)(1 - 2p_m(y)).$$

Since $p_m(y) \leq 3/8$, $1 - 2p_m(y) \geq \frac{1}{4}$, and hence $q_t^{(m)}(y) \geq \frac{1}{4}\phi(y) > 0$. Moreover, $q_t^{(m)}$ is smooth because it is a finite sum of products of smooth functions. Thus $s_t^{(m)} = \partial_x \log q_t^{(m)}$ and $\partial_x s_t^{(m)}$ are well defined for every $t \in [0, 1]$.

Since $\partial_x s_t^{(m)}(x) = \partial_x \frac{(q_t^{(m)})'(x)}{q_t^{(m)}(x)} = \frac{(q_t^{(m)})''(x)q_t^{(m)}(x) - ((q_t^{(m)})'(x))^2}{(q_t^{(m)}(x))^2} = \frac{(q_t^{(m)})''(x)}{q_t^{(m)}(x)} - s_t^{(m)}(x)^2$, we have

$$\begin{aligned} \mathbb{E}[\partial_x s_t^{(m)}(Y_t)] &= \int_{\mathbb{R}} \partial_x s_t^{(m)}(y) q_t^{(m)}(y) dy \\ &= \int_{\mathbb{R}} (q_t^{(m)})''(y) dy - \mathbb{E}[s_t^{(m)}(Y_t)^2] = -\mathbb{E}[s_t^{(m)}(Y_t)^2], \end{aligned}$$

because $(q_t^{(m)})'(y) \rightarrow 0$ as $|y| \rightarrow \infty$. Jensen's inequality therefore yields

$$\mathbb{E}[|\partial_x s_t^{(m)}(Y_t)|^2] \geq \left| \mathbb{E}[\partial_x s_t^{(m)}(Y_t)] \right|^2 = \left(\mathbb{E}[s_t^{(m)}(Y_t)^2] \right)^2.$$

It therefore remains to lower bound $\mathbb{E}[s_t^{(m)}(Y_t)^2]$. Rather than differentiating the explicit expression for $q_t^{(m)}$ and attempting to control the resulting cross terms pointwise, we use a variational consequence of the score identity. For any smooth bounded function f ,

$$\begin{aligned} \mathbb{E}[f'(Y_t)] &= \int_{\mathbb{R}} f'(y) q_t^{(m)}(y) dy \\ &= - \int_{\mathbb{R}} f(y) (q_t^{(m)})'(y) dy \\ &= -\mathbb{E}[f(Y_t) s_t^{(m)}(Y_t)], \end{aligned}$$

where the boundary term vanishes because $q_t^{(m)}$ has Gaussian tails. Hence, by Cauchy-Schwarz,

$$|\mathbb{E}[f'(Y_t)]|^2 \leq \mathbb{E}[f(Y_t)^2] \mathbb{E}[s_t^{(m)}(Y_t)^2],$$

and therefore

$$\mathbb{E}[s_t^{(m)}(Y_t)^2] \geq \frac{|\mathbb{E}[f'(Y_t)]|^2}{\mathbb{E}[f(Y_t)^2]} = \frac{\left(\int q_t^{(m)}(y) f'(y) dy \right)^2}{\int q_t^{(m)}(y) f(y)^2 dy}.$$

We want to choose f so that the lower bound doesn't decrease, as m increases. Take $f(y) = \cos(K_m y)$. Its derivative is $f'(y) = -K_m \sin(K_m y)$, so the preceding inequality gives

$$\mathbb{E}[s_t^{(m)}(Y_t)^2] \geq \frac{K_m^2 |\mathbb{E}[\sin(K_m Y_t)]|^2}{\mathbb{E}[\cos^2(K_m Y_t)]} \geq K_m^2 |\mathbb{E}[\sin(K_m Y_t)]|^2,$$

where the last inequality uses $\cos^2(K_m Y_t) \leq 1$. By construction,

$$\begin{aligned} \mathbb{E}[\sin(K_m Y_t) \mid Y_0 = x] &= p_m(x) \sin(K_m x + tK_m \delta) + p_m(x) \sin(K_m x - tK_m \delta) \\ &\quad + (1 - 2p_m(x)) \sin(K_m x). \end{aligned}$$

Using $\sin(a+b) + \sin(a-b) = 2\sin(a)\cos(b)$ we obtain

$$\begin{aligned}\mathbb{E}[\sin(K_m Y_t) \mid Y_0 = x] &= 2p_m(x) \sin(K_m x) \cos(tK_m \delta) + (1 - 2p_m(x)) \sin(K_m x) \\ &= \sin(K_m x) [1 + 2p_m(x) (\cos(tK_m \delta) - 1)].\end{aligned}$$

Taking expectation with respect to Y_0 and using $\mathbb{E}[\sin(K_m Y_0)] = 0$ gives

$$\mathbb{E}[\sin(K_m Y_t)] = (\cos(tK_m \delta) - 1) \mathbb{E}[2p_m(Y_0) \sin(K_m Y_0)].$$

Recall by definition $2p_m(Y_0) = \frac{1}{2} + \frac{1}{4} \sin(K_m Y_0)$, so we have

$$\mathbb{E}[2p_m(Y_0) \sin(K_m Y_0)] = \frac{1}{2} \mathbb{E}[\sin(K_m Y_0)] + \frac{1}{4} \mathbb{E}[\sin^2(K_m Y_0)] = \frac{1}{4} \mathbb{E}[\sin^2(K_m Y_0)].$$

Using $\sin^2(u) = \frac{1 - \cos(2u)}{2}$ and $\mathbb{E}[\cos(aY_0)] = \int_{\mathbb{R}} \phi(x) \cos(ax) dx = e^{-a^2/2}$, we obtain

$$\mathbb{E}[\sin^2(K_m Y_0)] = \frac{1 - e^{-2K_m^2}}{2}.$$

Consequently,

$$\mathbb{E}[\sin(K_m Y_t)] = -\frac{(1 - \cos(tK_m \delta)) (1 - e^{-2K_m^2})}{8} = -\frac{(1 - \cos((2m+1)\pi t)) (1 - e^{-2K_m^2})}{8},$$

where the second equality uses $K_m \delta = (2m+1)\pi$. Substituting this expression back yields

$$\begin{aligned}\mathbb{E}[s_t^{(m)}(Y_t)^2] &\geq K_m^2 |\mathbb{E}[\sin(K_m Y_t)]|^2 \\ &= \frac{K_m^2 (1 - e^{-2K_m^2})^2}{64} (1 - \cos((2m+1)\pi t))^2.\end{aligned}$$

It follows that

$$\begin{aligned}\mathbb{E}[|\partial_x s_t^{(m)}(Y_t)|^2] &\geq \left(\frac{K_m^2 (1 - e^{-2K_m^2})^2}{64} (1 - \cos((2m+1)\pi t))^2 \right)^2 \\ &= \frac{K_m^4 (1 - e^{-2K_m^2})^4}{4096} (1 - \cos((2m+1)\pi t))^4.\end{aligned}$$

Integrating over $t \in [0, 1]$ therefore gives

$$\int_0^1 \mathbb{E}[|\partial_x s_t^{(m)}(Y_t)|^2] dt \geq \frac{K_m^4 (1 - e^{-2K_m^2})^4}{4096} \int_0^1 (1 - \cos((2m+1)\pi t))^4 dt.$$

It remains only to check that the oscillatory factor does not disappear after integration in time. Expanding,

$$(1 - \cos u)^4 = 1 - 4\cos u + 6\cos^2 u - 4\cos^3 u + \cos^4 u.$$

With the change of variables $u = (2m+1)\pi t$,

$$\int_0^1 (1 - \cos((2m+1)\pi t))^4 dt = \frac{1}{(2m+1)\pi} \int_0^{(2m+1)\pi} (1 - \cos u)^4 du.$$

Since $2m + 1$ is an integer,

$$\begin{aligned}\int_0^{(2m+1)\pi} \cos u \, du &= 0, \\ \int_0^{(2m+1)\pi} \cos^3 u \, du &= 0, \\ \int_0^{(2m+1)\pi} \cos^2 u \, du &= \frac{(2m+1)\pi}{2}, \\ \int_0^{(2m+1)\pi} \cos^4 u \, du &= \frac{3(2m+1)\pi}{8}.\end{aligned}$$

Hence

$$\int_0^1 (1 - \cos((2m+1)\pi t))^4 \, dt = 1 + 6 \cdot \frac{1}{2} + \frac{3}{8} = \frac{35}{8}.$$

Consequently,

$$\int_0^1 \mathbb{E}[|\partial_x s_t^{(m)}(Y_t)|^2] \, dt \geq \frac{35}{32768} K_m^4 (1 - e^{-2K_m^2})^4.$$

Recall that $D_5 = \varepsilon$ for every m , whereas

$$K_m = \frac{(2m+1)\pi}{\delta} \longrightarrow \infty.$$

Therefore,

$$\int_0^1 \mathbb{E}[|\partial_x s_t^{(m)}(Y_t)|^2] \, dt \longrightarrow +\infty$$

while D_5 remains fixed at the arbitrarily prescribed value ε . Thus, without smoothing, a small coupling displacement does not imply small integrated regularity. $\square$

G Proof of Theorem 1

Standard Sampling Conditions. For the sampling analysis, we require that the exact score-corrected drift b_t is jointly continuous on $[0, 1] \times \mathbb{R}^d$, and that the exact SDE initialized from q_0 ,

$$dZ_t = b_t(Z_t) \, dt + \tau \, dB_t,$$

admits a weak solution on $[0, 1]$ satisfying $\mathcal{L}(Z_t) = q_t$ for every $t \in [0, 1]$. We also require each grid-frozen learned drift $\widehat{b}_{t_k}$ to be continuous. The exact-process requirements are verified in Lemma 1. Continuity of the learned drifts follows from that of the neural networks.

Lemma 1 (Sampling well-posedness). *Let $\sigma, \tau > 0$ and $D_5 < \infty$. Then $b_t(x) = v_t(x) + \frac{\tau^2}{2} s_t(x)$ is jointly continuous on $[0, 1] \times \mathbb{R}^d$, and the exact SDE*

$$dZ_t = b_t(Z_t) \, dt + \tau \, dB_t, \quad Z_0 \sim q_0,$$

admits a weak solution on $[0, 1]$ satisfying $\mathcal{L}(Z_t) = q_t$ for every $t \in [0, 1]$.

Proof. Let φ_σ be the density of $\mathcal{N}(0, \sigma^2 I_d)$ and recall that $S_t = Y_0 + tD$, where $D = Y_1 - Y_0$. Then

$$q_t(x) = \mathbb{E}[\varphi_\sigma(x - S_t)], \quad q_t(x)v_t(x) = \mathbb{E}[D\varphi_\sigma(x - S_t)].$$

Fix $(t_n, x_n) \rightarrow (t, x)$. Since $S_{t_n} \rightarrow S_t$ almost surely,

$$\begin{aligned} \varphi_\sigma(x_n - S_{t_n}) &\rightarrow \varphi_\sigma(x - S_t), \\ \nabla_x \varphi_\sigma(x_n - S_{t_n}) &\rightarrow \nabla_x \varphi_\sigma(x - S_t) \end{aligned}$$

almost surely. Both φ_σ and $\nabla_x \varphi_\sigma$ are bounded, so dominated convergence gives joint continuity of $q_t(x)$ and $\nabla_x q_t(x)$. Moreover,

$$D\varphi_\sigma(x_n - S_{t_n}) \rightarrow D\varphi_\sigma(x - S_t)$$

almost surely and

$$\|D\varphi_\sigma(x_n - S_{t_n})\| \leq \|\varphi_\sigma\|_\infty \|D\|.$$

Since $D_5 < \infty$ implies $\mathbb{E}\|D\| < \infty$, dominated convergence also gives joint continuity of $q_t(x)v_t(x)$.

Because $\varphi_\sigma > 0$, we have $q_t(x) > 0$ for every (t, x) . Hence

$$v_t(x) = \frac{\mathbb{E}[D\varphi_\sigma(x - S_t)]}{q_t(x)}, \quad s_t(x) = \frac{\nabla_x q_t(x)}{q_t(x)}$$

are jointly continuous, and therefore so is b_t .

We next verify the Fokker–Planck equation satisfied by $(q_t)_{t \in [0,1]}$. For every $f \in C_c^\infty(\mathbb{R}^d)$,

$$\begin{aligned} \frac{d}{dt} \int_{\mathbb{R}^d} f(x) q_t(x) dx &= \frac{d}{dt} \mathbb{E}[f(X_t)] \\ &= \mathbb{E}[\nabla_x f(X_t)^\top D] \\ &= \int_{\mathbb{R}^d} \nabla_x f(x)^\top v_t(x) q_t(x) dx, \end{aligned}$$

where differentiation under the expectation is justified by $\|\nabla_x f(X_t)^\top D\| \leq \|\nabla_x f\|_\infty \|D\|$. Since $q_t s_t = \nabla_x q_t$, integration by parts gives

$$\int_{\mathbb{R}^d} \nabla_x f(x)^\top s_t(x) q_t(x) dx = - \int_{\mathbb{R}^d} \Delta_x f(x) q_t(x) dx.$$

Thus

$$\frac{d}{dt} \int_{\mathbb{R}^d} f(x) q_t(x) dx = \int_{\mathbb{R}^d} \left(\nabla_x f(x)^\top b_t(x) + \frac{\tau^2}{2} \Delta_x f(x) \right) q_t(x) dx.$$

It remains to verify the integrability required by the superposition principle. By conditional Jensen's inequality,

$$\int_{\mathbb{R}^d} \|v_t(x)\| q_t(x) dx = \mathbb{E}\|\mathbb{E}[D \mid X_t]\| \leq \mathbb{E}\|D\|,$$

while the Gaussian score identity gives

$$\int_{\mathbb{R}^d} \|s_t(x)\| q_t(x) dx = \frac{1}{\sigma} \mathbb{E}\|\mathbb{E}[G \mid X_t]\| \leq \frac{\mathbb{E}\|G\|}{\sigma}.$$

Consequently,

$$\int_0^1 \int_{\mathbb{R}^d} \|b_t(x)\| q_t(x) dx dt \leq \mathbb{E}\|D\| + \frac{\tau^2}{2\sigma} \mathbb{E}\|G\| < \infty.$$

Moreover, $t \mapsto q_t(x) dx$ is narrowly continuous: for every bounded continuous f , almost-sure continuity of $t \mapsto X_t$ and dominated convergence imply that $t \mapsto \mathbb{E}[f(X_t)]$ is continuous.

Therefore, (Trevisan, 2016, Theorem 2.5) yields a probability measure on $C([0, 1]; \mathbb{R}^d)$ solving the martingale problem associated with drift b_t and diffusion matrix $\tau^2 I_d$, with time- t marginal q_t . For the canonical process Z , its martingale part has quadratic covariation $\tau^2 t I_d$; hence, by the Lévy characterization, it can be written as τB_t for a Brownian motion B . Thus Z is a weak solution of the displayed SDE and satisfies $\mathcal{L}(Z_t) = q_t$ for every $t \in [0, 1]$. $\square$

G.1 Path-Space Comparison via Girsanov

We now introduce a framework Chen et al. (2023b), Chen et al. (2023a) that controls the KL divergence between the laws of two SDEs.

Let $T > 0$ (for our flow matching setting, $T = 1$). Let $0 = t_0 < t_1 < \dots < t_N = T$ be a uniform time grid with step size $h = t_{k+1} - t_k$. For every $t \in [t_k, t_{k+1})$, define $\eta_h(t) = t_k$. Recall that our sampling SDE is written as $dZ_t = b_t(Z_t) dt + \tau dB_t$. For the continuous-time Euler discretization, we have

$$d\hat{Z}_t = \hat{b}_{\eta_h(t)}(\hat{Z}_{\eta_h(t)}) dt + \tau dB_t.$$

Equivalently, for $t \in [t_k, t_{k+1})$,

$$d\hat{Z}_t = \hat{b}_{t_k}(\hat{Z}_{t_k}) dt + \tau dB_t.$$

Assume both processes start from the same initial law: $Z_0 \sim \rho_0$ and $\hat{Z}_0 \sim \rho_0$.

Lemma 2 (Girsanov Bound for the Euler Process via Localization). *Let $\tau > 0$, and let b and $\{\hat{b}_{t_k}\}_{k=0}^{N-1}$ be continuous drift fields. Fix a weak solution Z of the exact SDE on $[0, T]$, with $Z_0 \sim \rho_0$, and let $\hat{Z}$ be the full continuous-time Euler process initialized from ρ_0 . Write $\mathbb{P} = \mathcal{L}(Z)$ and $\hat{\mathbb{P}} = \mathcal{L}(\hat{Z})$ for their laws on $C([0, T]; \mathbb{R}^d)$.*

Define $\Delta_t(Z) = \hat{b}_{\eta_h(t)}(Z_{\eta_h(t)}) - b_t(Z_t)$, and assume that $\int_0^T \mathbb{E} [\|\Delta_t(Z)\|^2] dt < \infty$. Also assume

$$\int_0^T \|\Delta_t(\hat{Z})\|^2 dt < \infty, \quad \hat{\mathbb{P}}\text{-a.s.}$$

Then

$$\text{KL}(\mathcal{L}(Z_T) \parallel \mathcal{L}(\hat{Z}_T)) \leq \frac{1}{2\tau^2} \int_0^T \mathbb{E} \left[\left\| b_t(Z_t) - \hat{b}_{\eta_h(t)}(Z_{\eta_h(t)}) \right\|^2 \right] dt.$$

The proof has four steps. First, we define a path-dependent stopping-time construction $\zeta_m(x)$, which gives the stopping times $\zeta_m(Z)$ along the exact process and $\zeta_m(\hat{Z})$ along the full Euler process. Second, for $\zeta_m(Z)$, Novikov's condition holds for each localized approximation, so we can apply Girsanov's theorem and obtain a KL bound for each of them. Third, the additional integrability condition on $\hat{Z}$ is used to show that $\zeta_m(\hat{Z}) \uparrow T$. Finally, we show that the localized Euler laws converge to the full Euler law on every truncated interval $[0, T - \epsilon]$; lower semicontinuity of relative entropy and then $\epsilon \downarrow 0$ complete the argument.

Proof. Let $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{0 \leq t \leq T}, \mathbf{P})$ be a filtered probability space carrying the exact weak solution Z discussed in Lemma 1 and its driving Brownian motion W , so that

$$dZ_t = b_t(Z_t) dt + \tau dW_t.$$

Let $\mathbb{P} = \mathcal{L}_{\mathbf{P}}(Z)$ be the corresponding path law on $C([0, T]; \mathbb{R}^d)$, and let $\hat{\mathbb{P}}$ denote the path law of the full continuous-time Euler process.

For every integer $m \geq 1$, define the stopping time

$$\zeta_m(x) = \inf \left\{ t \in [0, T] : \frac{1}{\tau^2} \int_0^t \|\Delta_s(x)\|^2 ds \geq m \right\} \wedge T.$$

Since $t \mapsto \tau^{-2} \int_0^t \|\Delta_s(Z)\|^2 ds$ is adapted and continuous, its first hitting time of $[m, \infty)$ is an $(\mathcal{F}_t)$ -stopping time. Thus ζ_m is a stopping time. When no confusion arises, we write ζ_m for $\zeta_m(Z)$ in the Girsanov construction, and for $\zeta_m(\hat{Z})$ in the synchronous coupling that will be discussed later. By construction, $\frac{1}{\tau^2} \int_0^{\zeta_m} \|\Delta_t(Z)\|^2 dt = \frac{1}{\tau^2} \int_0^T \|\Delta_t(Z)\|^2 \mathbf{1}_{\{t \leq \zeta_m\}} dt \leq m$. Therefore Novikov's condition holds for the stopped drift perturbation:

$$\mathbb{E}_{\mathbf{P}} \left[\exp \left(\frac{1}{2\tau^2} \int_0^T \|\Delta_t(Z)\|^2 \mathbf{1}_{\{t \leq \zeta_m\}} dt \right) \right] \leq \exp \left(\frac{m}{2} \right) < \infty.$$

Hence the stopped stochastic exponential

$$M_m^T := \exp \left(\frac{1}{\tau} \int_0^T \left(\Delta_t(Z) \mathbf{1}_{\{t \leq \zeta_m\}} \right)^\top dW_t - \frac{1}{2\tau^2} \int_0^T \|\Delta_t(Z)\|^2 \mathbf{1}_{\{t \leq \zeta_m\}} dt \right)$$

is a true martingale with expectation $\mathbb{E}_{\mathbf{P}}[M_m^T] = \mathbb{E}_{\mathbf{P}}[M_m^0] = 1$.

Thus it defines a probability measure $\hat{\mathbf{P}}^{(m)}$ on $(\Omega, \mathcal{F})$ by $\frac{d\hat{\mathbf{P}}^{(m)}}{d\mathbf{P}} = M_m^T$ (note that $\int d\hat{\mathbf{P}}^{(m)} = \int M_m^T d\mathbf{P} = \mathbb{E}_{\mathbf{P}}[M_m^T] = 1$ so $\hat{\mathbf{P}}^{(m)}$ is a probability measure). By Girsanov's theorem, under $\hat{\mathbf{P}}^{(m)}$,

$$W_t^{(m)} = W_t - \frac{1}{\tau} \int_0^t \Delta_s(Z) \mathbf{1}_{\{s \leq \zeta_m\}} ds$$

is a Brownian motion.

Therefore,

$$dW_t = dW_t^{(m)} + \frac{1}{\tau} \Delta_t(Z) \mathbf{1}_{\{t \leq \zeta_m\}} dt.$$

Since, under $\mathbf{P}$,

$$dZ_t = b_t(Z_t) dt + \tau dW_t,$$

we may rewrite the same process under $\hat{\mathbf{P}}^{(m)}$ as

$$\begin{aligned} dZ_t &= b_t(Z_t) dt + \tau \left(dW_t^{(m)} + \frac{1}{\tau} \Delta_t(Z) \mathbf{1}_{\{t \leq \zeta_m\}} dt \right) \\ &= \left(b_t(Z_t) + \Delta_t(Z) \mathbf{1}_{\{t \leq \zeta_m\}} \right) dt + \tau dW_t^{(m)}. \end{aligned}$$

Thus, before the stopping time ζ_m , the drift is

$$b_t(Z_t) + \Delta_t(Z) = \hat{b}_{\eta_h(t)}(Z_{\eta_h(t)}),$$

and after the stopping time, the drift returns to $b_t(Z_t)$. Define $\widehat{\mathbb{P}}^{(m)} := \mathcal{L}_{\widehat{\mathbf{P}}^{(m)}}(Z)$. Thus $\widehat{\mathbb{P}}^{(m)}$ is the path law of the localized Euler approximation. For each m , the path-space KL divergence is well-defined. Since

$$\log \frac{d\mathbf{P}}{d\widehat{\mathbf{P}}^{(m)}} = -\frac{1}{\tau} \int_0^T \left(\Delta_t(Z) \mathbf{1}_{\{t \leq \zeta_m\}} \right)^\top dW_t + \frac{1}{2\tau^2} \int_0^T \left\| \Delta_t(Z) \mathbf{1}_{\{t \leq \zeta_m\}} \right\|^2 dt.$$

Since Z is continuous and adapted, $b_t(Z_t)$ is predictable, while the grid-frozen term admits a predictable modification differing only at the deterministic grid points. Hence $\Delta_t(Z) \mathbf{1}_{\{t \leq \zeta_m\}}$ admits a predictable version. Moreover, by the definition of ζ_m ,

$$\int_0^T \left\| \Delta_t(Z) \mathbf{1}_{\{t \leq \zeta_m\}} \right\|^2 dt \leq m\tau^2,$$

so it is square-integrable and the stochastic integral has mean zero.

Taking expectation under $\mathbf{P}$, we obtain

$$\begin{aligned} \text{KL} \left(\mathbf{P} \parallel \widehat{\mathbf{P}}^{(m)} \right) &= \frac{1}{2\tau^2} \mathbb{E}_{\mathbf{P}} \left[\int_0^{\zeta_m} \left\| \Delta_t(Z) \right\|^2 dt \right] \\ &\leq \frac{1}{2\tau^2} \mathbb{E}_{\mathbf{P}} \left[\int_0^T \left\| \Delta_t(Z) \right\|^2 dt \right]. \end{aligned}$$

Consequently, by the data-processing inequality under the map $\omega \mapsto Z(\omega)$,

$$\text{KL} \left(\mathbb{P} \parallel \widehat{\mathbb{P}}^{(m)} \right) \leq \text{KL} \left(\mathbf{P} \parallel \widehat{\mathbf{P}}^{(m)} \right).$$

Therefore,

$$\begin{aligned} \text{KL} \left(\mathbb{P} \parallel \widehat{\mathbb{P}}^{(m)} \right) &\leq \frac{1}{2\tau^2} \mathbb{E}_{\mathbf{P}} \left[\int_0^T \left\| \Delta_t(Z) \mathbf{1}_{\{t \leq \zeta_m\}} \right\|^2 dt \right] = \frac{1}{2\tau^2} \mathbb{E}_{\mathbf{P}} \left[\int_0^{\zeta_m} \left\| \Delta_t(Z) \right\|^2 dt \right] \\ &\leq \frac{1}{2\tau^2} \mathbb{E}_{\mathbf{P}} \left[\int_0^T \left\| \Delta_t(Z) \right\|^2 dt \right]. \end{aligned}$$

It remains to pass from the localized Euler approximation to the full continuous-time Euler approximation. For each fixed m , since M_m^t is a martingale with $M_m^0 = 1$,

$$\mathbb{E}_{\mathbf{P}}[M_m^T \mid \mathcal{F}_0] = 1,$$

so the change of measure preserves the initial law and $Z_0 \sim \rho_0$ under $\widehat{\mathbf{P}}^{(m)}$.

Under $\widehat{\mathbf{P}}^{(m)}$, on the same probability space, construct the continuous time Euler process $\widehat{Z}$ using the same initial value Z_0 and Brownian motion $W^{(m)}$:

$$\widehat{Z}_t = \widehat{Z}_{t_k} + (t - t_k) \widehat{b}_{t_k}(\widehat{Z}_{t_k}) + \tau \left(W_t^{(m)} - W_{t_k}^{(m)} \right), \quad t \in [t_k, t_{k+1}],$$

with $\widehat{Z}_0 = Z_0$. Because $W^{(m)}$ is Brownian with respect to $(\mathcal{F}_t)$ and Z_0 is $\mathcal{F}_0$ -measurable, $W^{(m)}$ is independent of Z_0 . The change of measure preserves the law ρ_0 of Z_0 , and the grid-by-grid Euler recursion is a measurable function of $(Z_0, W^{(m)})$. Hence $\mathcal{L}_{\widehat{\mathbf{P}}^{(m)}}(\widehat{Z}) = \mathbb{P}$.

Let $\rho_m = \zeta_m(Z) \wedge \zeta_m(\widehat{Z})$. Before ρ_m , both Z and $\widehat{Z}$ satisfy the same continuous-time Euler update with the same initial value and Brownian motion, so they agree grid-by-grid on $[0, \rho_m]$. Hence $\int_0^t \left\| \Delta_s(Z) \right\|^2 ds = \int_0^t \left\| \Delta_s(\widehat{Z}) \right\|^2 ds$ for every $t \leq \rho_m$.

Suppose $\zeta_m(Z) < \zeta_m(\hat{Z})$. Then $\rho_m = \zeta_m(Z) < T$, so continuity of the defining integral gives $\tau^{-2} \int_0^{\rho_m} \|\Delta_s(Z)\|^2 ds = m$. The preceding equality implies the same for $\hat{Z}$, hence $\zeta_m(\hat{Z}) \leq \rho_m = \zeta_m(Z)$, a contradiction. The reverse strict inequality is ruled out in the same way. Therefore $\zeta_m(Z) = \zeta_m(\hat{Z})$.

Fix $\varepsilon \in (0, T)$ and define the path truncation map $\pi_\varepsilon : C([0, T]; \mathbb{R}^d) \rightarrow C([0, T]; \mathbb{R}^d)$ by

$$(\pi_\varepsilon(\omega))(t) = \omega(t \wedge (T - \varepsilon)), \quad 0 \leq t \leq T.$$

On the event $\{\zeta_m(Z) > T - \varepsilon\}$, we have $\pi_\varepsilon(Z) = \pi_\varepsilon(\hat{Z})$. Moreover,

$$\hat{\mathbf{P}}^{(m)}(\zeta_m(Z) \leq T - \varepsilon) = \hat{\mathbf{P}}^{(m)}(\zeta_m(\hat{Z}) \leq T - \varepsilon) = \hat{\mathbb{P}}(\zeta_m(\hat{Z}) \leq T - \varepsilon) \rightarrow 0 \text{ as } m \rightarrow \infty.$$

Indeed, $\int_0^T \|\Delta_t(\hat{Z})\|^2 dt < \infty$, $\hat{\mathbb{P}}$ -almost surely, implies that $\zeta_m(\hat{Z}) = T$ for all sufficiently large m , almost surely.

Therefore, for every bounded continuous $F : C([0, T]; \mathbb{R}^d) \rightarrow \mathbb{R}$,

$$\begin{aligned} & \left| \mathbb{E}_{\hat{\mathbf{P}}^{(m)}}[F(\pi_\varepsilon(Z))] - \mathbb{E}_{\hat{\mathbf{P}}^{(m)}}[F(\pi_\varepsilon(\hat{Z}))] \right| \\ & \leq 2\|F\|_\infty \hat{\mathbf{P}}^{(m)}(\zeta_m(Z) \leq T - \varepsilon) \rightarrow 0. \end{aligned}$$

Since $\mathcal{L}_{\hat{\mathbf{P}}^{(m)}}(\hat{Z}) = \hat{\mathbb{P}}$, we obtain

$$\mathcal{L}_{\hat{\mathbf{P}}^{(m)}}(\pi_\varepsilon(Z)) \xrightarrow{w} (\pi_\varepsilon)_\# \hat{\mathbb{P}}.$$

By lower semicontinuity of KL divergence and the data-processing inequality,

$$\begin{aligned} \text{KL}\left((\pi_\varepsilon)_\# \mathbb{P} \parallel (\pi_\varepsilon)_\# \hat{\mathbb{P}}\right) & \leq \liminf_{m \rightarrow \infty} \text{KL}\left(\mathcal{L}_{\mathbf{P}}(\pi_\varepsilon(Z)) \parallel \mathcal{L}_{\hat{\mathbf{P}}^{(m)}}(\pi_\varepsilon(Z))\right) \\ & \leq \liminf_{m \rightarrow \infty} \text{KL}\left(\mathbf{P} \parallel \hat{\mathbf{P}}^{(m)}\right) \\ & \leq \frac{1}{2\tau^2} \mathbb{E}_{\mathbf{P}} \left[\int_0^T \|\Delta_t(Z)\|^2 dt \right]. \end{aligned}$$

Finally, since $\mathcal{L}(\pi_\varepsilon(Z)) \xrightarrow{w} \mathcal{L}(Z)$ and $\mathcal{L}(\pi_\varepsilon(\hat{Z})) \xrightarrow{w} \mathcal{L}(\hat{Z})$ as $\varepsilon \rightarrow 0$, applying lower semicontinuity of KL divergence once more gives

$$\text{KL}\left(\mathbb{P} \parallel \hat{\mathbb{P}}\right) \leq \liminf_{\varepsilon \downarrow 0} \text{KL}\left(\mathcal{L}(\pi_\varepsilon(Z)) \parallel \mathcal{L}(\pi_\varepsilon(\hat{Z}))\right) \leq \frac{1}{2\tau^2} \mathbb{E}_{\mathbf{P}} \left[\int_0^T \|\Delta_t(Z)\|^2 dt \right].$$

Again, by the data-processing inequality,

$$\text{KL}\left(\mathcal{L}(Z_T) \parallel \mathcal{L}(\hat{Z}_T)\right) \leq \text{KL}\left(\mathbb{P} \parallel \hat{\mathbb{P}}\right) \leq \frac{1}{2\tau^2} \mathbb{E}_{\mathbf{P}} \left[\int_0^T \|\Delta_t(Z)\|^2 dt \right].$$

Equivalently,

$$\text{KL}\left(\mathcal{L}(Z_T) \parallel \mathcal{L}(\hat{Z}_T)\right) \leq \frac{1}{2\tau^2} \int_0^T \mathbb{E}_{\mathbf{P}} \left[\left\| b_t(Z_t) - \hat{b}_{\eta_h(t)}(Z_{\eta_h(t)}) \right\|^2 \right] dt.$$

□

A point concerning the localization limit. The approximation argument in [Chen et al. \(2023b, Section 5.2\)](#) places the localized and full Euler processes on a common probability space and subsequently invokes convergence of the localizing times. The Lipschitz-SDE result used there, cited as Corollary 8.5 in [Chen et al. \(2023b, Section 5.2\)](#), comes from the chapter on stochastic differential equations in [Le Gall \(2016\)](#) and provides a measurable strong-solution representation. Although this result can support a synchronous realization when its hypotheses hold, it does not by itself show that localizing times constructed under the reference law converge to T under the full Euler law. This Euler-side exhaustion is required to obtain agreement of the localized and full Euler paths on every interval $[0, T - \varepsilon]$, and hence the weak convergence of the localized laws. We avoid this issue by using the explicit energy stopping times ζ_m defined above. It remains to verify that

$$\int_0^T \|\Delta_t(\hat{Z})\|^2 dt < \infty, \quad \hat{\mathbb{P}}\text{-a.s.},$$

because this condition implies that $\zeta_m(\hat{Z}) \uparrow T$ almost surely; indeed, $\zeta_m(\hat{Z}) = T$ for all sufficiently large m . The following proposition establishes this property under our continuity assumptions.

Proposition 3 (Almost-Sure Finiteness of the Pathwise Drift Mismatch). *Let $0 = t_0 < t_1 < \dots < t_N = T$ be a finite time grid, and define $\eta_h(t) = t_k$ for $t \in [t_k, t_{k+1})$. Let $\hat{Z} = (\hat{Z}_t)_{0 \leq t \leq T}$ be a continuous-path process under $\hat{\mathbb{P}}$.*

Let $b : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ be continuous. For every $k \in \{0, \dots, N-1\}$, let $\hat{b}_{t_k} : \mathbb{R}^d \rightarrow \mathbb{R}^d$ be continuous. Define

$$\Delta_t(\hat{Z}) = \hat{b}_{\eta_h(t)}(\hat{Z}_{\eta_h(t)}) - b_t(\hat{Z}_t), \quad 0 \leq t \leq T.$$

Then

$$\int_0^T \|\Delta_t(\hat{Z})\|^2 dt < \infty, \quad \hat{\mathbb{P}}\text{-a.s.}$$

Proof. It suffices to show, there exists some Ω_0 with $\hat{\mathbb{P}}(\Omega_0) = 1$ s.t. $\sup_{t \in [0, T]} \|\Delta_t(\hat{Z}(\omega))\| < \infty, \forall \omega \in \Omega_0$. Since $\hat{Z}$ has continuous paths under $\hat{\mathbb{P}}$, there exists Ω_0 such that $\hat{\mathbb{P}}(\Omega_0) = 1$ and for every $\omega \in \Omega_0$, the map $t \mapsto \hat{Z}_t(\omega)$ is continuous on $[0, T]$.

Fix $\omega \in \Omega_0$. Since $[0, T]$ is compact and $t \mapsto \hat{Z}_t(\omega)$ is continuous, the image set $K_\omega := \{\hat{Z}_t(\omega) : 0 \leq t \leq T\}$ is a compact subset of $\mathbb{R}^d$.

Because b and $\hat{b}_{t_k}$ are continuous, and because $[0, T] \times K_\omega$ is compact, we have

$$\begin{aligned} M_b(\omega) &:= \sup_{0 \leq t \leq T} \sup_{x \in K_\omega} \|b_t(x)\| < \infty, \\ M_{\hat{b}_{t_k}}(\omega) &:= \sup_{x \in K_\omega} \|\hat{b}_{t_k}(x)\| < \infty, \\ M_{\hat{b}}(\omega) &:= \max_{0 \leq k \leq N-1} M_{\hat{b}_{t_k}}(\omega) < \infty. \end{aligned}$$

Therefore $\|b_t(\hat{Z}_t(\omega))\| \leq M_b(\omega), \forall t \in [0, T]$ and $\|\hat{b}_{\eta_h(t)}(\hat{Z}_{\eta_h(t)}(\omega))\| \leq M_{\hat{b}}(\omega)$. Hence, for every $t \in [0, T]$,

$$\begin{aligned} \|\Delta_t(\hat{Z}(\omega))\| &= \|b_t(\hat{Z}_t(\omega)) - \hat{b}_{\eta_h(t)}(\hat{Z}_{\eta_h(t)}(\omega))\| \leq \|b_t(\hat{Z}_t(\omega))\| + \|\hat{b}_{\eta_h(t)}(\hat{Z}_{\eta_h(t)}(\omega))\| \\ &\leq M_b(\omega) + M_{\hat{b}}(\omega). \end{aligned}$$

Therefore,

$$\int_0^T \left\| \Delta_t(\widehat{Z}(\omega)) \right\|^2 dt \leq \int_0^T \left(M_b(\omega) + M_{\widehat{b}}(\omega) \right)^2 dt = T \left(M_b(\omega) + M_{\widehat{b}}(\omega) \right)^2 < \infty.$$

Since this holds for every $\omega \in \Omega_0$, and since $\widehat{\mathbb{P}}(\Omega_0) = 1$, we conclude that

$$\int_0^T \left\| \Delta_t(\widehat{Z}) \right\|^2 dt < \infty, \quad \widehat{\mathbb{P}}\text{-a.s.}$$

This proves the claim. $\square$

G.2 Posterior Derivative Identities

This subsection collects the derivative identities and integrated bounds needed for the combined Itô analysis of the score increment

$$s_t(Z_t) - s_{t_k}(Z_{t_k}).$$

The key point is that the term

$$\nabla_x(\nabla_x \cdot v_t)$$

requires control of the full Hessian tensor $D_x^2 v_t$, not merely the componentwise Laplacian $\Delta_x v_t$.

Lemma 3 (Score Identity Under Gaussian Smoothing). *Let Y_0 and Y_1 be random variables taking values in $\mathbb{R}^d$, and fix $t \in [0, T]$. Let $D = Y_1 - Y_0$. Define $S_t = Y_0 + tD$. Let $G \sim \mathcal{N}(0, I_d)$ be independent of (Y_0, D) , and define $X_t = S_t + \sigma G$ with $\sigma > 0$.*

Let q_t be the density of X_t , and define the score $s_t : \mathbb{R}^d \rightarrow \mathbb{R}^d$ by $s_t(x) = \nabla_x \log q_t(x)$. Then

$$s_t(x) = -\frac{1}{\sigma} \mathbb{E}[G \mid X_t = x].$$

Proof. Let $\mathbf{P}_{Y_0, Y_1}$ be the law of (Y_0, Y_1) . For fixed $t \in [0, T]$, define

$$s_t^{\text{path}}(y_0, y_1) = (1 - t)y_0 + ty_1.$$

The superscript “path” is only used in this proof to avoid confusing the interpolation point $s_t^{\text{path}}(y_0, y_1)$ with the score $s_t(x)$.

Since $X_t = S_t + \sigma G$, conditional on $(Y_0, Y_1) = (y_0, y_1)$, the random variable X_t is Gaussian with mean $s_t^{\text{path}}(y_0, y_1)$ and covariance matrix $\sigma^2 I_d$. Therefore the conditional density of $X_t = x$ given $(Y_0, Y_1) = (y_0, y_1)$ is

$$\frac{1}{(2\pi\sigma^2)^{d/2}} \exp \left(-\frac{\left\| x - s_t^{\text{path}}(y_0, y_1) \right\|^2}{2\sigma^2} \right).$$

Hence the marginal density of X_t is

$$q_t(x) = \int_{\mathbb{R}^d \times \mathbb{R}^d} \frac{1}{(2\pi\sigma^2)^{d/2}} \exp \left(-\frac{\left\| x - s_t^{\text{path}}(y_0, y_1) \right\|^2}{2\sigma^2} \right) d\mathbf{P}_{Y_0, Y_1}(y_0, y_1).$$

Because the Gaussian density is strictly positive, $q_t(x) > 0$ for every $x \in \mathbb{R}^d$.

We now compute $\nabla_x q_t(x)$. For each coordinate $j \in \{1, \dots, d\}$,

$$\partial_{x_j} \exp \left(-\frac{\|x - s_t^{\text{path}}(y_0, y_1)\|^2}{2\sigma^2} \right) = -\frac{x_j - s_{t,j}^{\text{path}}(y_0, y_1)}{\sigma^2} \exp \left(-\frac{\|x - s_t^{\text{path}}(y_0, y_1)\|^2}{2\sigma^2} \right).$$

Therefore,

$$\partial_{x_j} q_t(x) = \int_{\mathbb{R}^d \times \mathbb{R}^d} -\frac{x_j - s_{t,j}^{\text{path}}(y_0, y_1)}{\sigma^2} \frac{1}{(2\pi\sigma^2)^{d/2}} \exp \left(-\frac{\|x - s_t^{\text{path}}(y_0, y_1)\|^2}{2\sigma^2} \right) dP_{Y_0, Y_1}(y_0, y_1).$$

Since $q_t(x) > 0$,

$$\begin{aligned} \partial_{x_j} \log q_t(x) &= \frac{\partial_{x_j} q_t(x)}{q_t(x)} \\ &= -\frac{1}{\sigma^2} \frac{\int_{\mathbb{R}^d \times \mathbb{R}^d} (x_j - s_{t,j}^{\text{path}}(y_0, y_1)) \frac{1}{(2\pi\sigma^2)^{d/2}} \exp \left(-\frac{\|x - s_t^{\text{path}}(y_0, y_1)\|^2}{2\sigma^2} \right) dP_{Y_0, Y_1}(y_0, y_1)}{q_t(x)}. \end{aligned}$$

By Bayes' rule, the conditional law of (Y_0, Y_1) given $X_t = x$ is obtained by weighting the prior law P_{Y_0, Y_1} by the same Gaussian likelihood. Thus

$$\partial_{x_j} \log q_t(x) = -\frac{1}{\sigma^2} \mathbb{E}[x_j - S_{t,j} \mid X_t = x].$$

Recall by definition $x = S_t + \sigma G$, so we have $x_j - S_{t,j} = \sigma G_j$ inside the conditional expectation. Therefore,

$$\partial_{x_j} \log q_t(x) = -\frac{1}{\sigma^2} \mathbb{E}[\sigma G_j \mid X_t = x] = -\frac{1}{\sigma} \mathbb{E}[G_j \mid X_t = x].$$

Since this holds for every $j \in \{1, \dots, d\}$, we obtain

$$\nabla_x \log q_t(x) = -\frac{1}{\sigma} \mathbb{E}[G \mid X_t = x].$$

By definition $s_t(x) = \nabla_x \log q_t(x)$, so

$$s_t(x) = -\frac{1}{\sigma} \mathbb{E}[G \mid X_t = x].$$

This proves the lemma. □

Lemma 4 (Posterior Derivative Identities for the Smoothed Velocity and Score). *Let Y_0 and Y_1 be random variables taking values in $\mathbb{R}^d$, and fix $t \in [0, T]$. Let $D = Y_1 - Y_0$. Define $S_t = Y_0 + tD$. Let $G \sim \mathcal{N}(0, I_d)$ be independent of (Y_0, D) , and define $X_t = S_t + \sigma G$ with $\sigma > 0$. Assume $\mathbb{E}\|D\|^3 < \infty$.*

Let q_t be the density of X_t . Define the smoothed velocity $v_t : \mathbb{R}^d \rightarrow \mathbb{R}^d$ by

$$v_t(x) = \mathbb{E}[D \mid X_t = x],$$

and define the score $s_t : \mathbb{R}^d \rightarrow \mathbb{R}^d$ by $s_t(x) = \nabla_x \log q_t(x)$. Under the stated moment condition, all derivatives below exist and the required differentiations under the Gaussian-mixture integral are justified.

Then the score satisfies

$$s_t(x) = -\frac{1}{\sigma} \mathbb{E}[G \mid X_t = x].$$

The velocity derivatives satisfy

$$\begin{aligned} D_x v_t(x) &= -\frac{1}{\sigma} \text{Cov}(D, G \mid X_t = x), \\ \partial_t v_t(x) &= \frac{1}{\sigma} \text{Cov}(D, D^\top G \mid X_t = x), \end{aligned}$$

and for every $i, j, \ell \in \{1, \dots, d\}$,

$$\partial_{x_j x_\ell} v_{t,i}(x) = \frac{1}{\sigma^2} \mathbb{E}[(D_i - \mathbb{E}[D_i \mid X_t = x])(G_j - \mathbb{E}[G_j \mid X_t = x])(G_\ell - \mathbb{E}[G_\ell \mid X_t = x]) \mid X_t = x].$$

In particular,

$$\Delta_x v_t(x) = \frac{1}{\sigma^2} \mathbb{E}[(D - \mathbb{E}[D \mid X_t = x]) \|G - \mathbb{E}[G \mid X_t = x]\|^2 \mid X_t = x].$$

The score derivatives satisfy

$$D_x s_t(x) = \frac{1}{\sigma^2} (\text{Cov}(G, G \mid X_t = x) - I_d),$$

and

$$\Delta_x s_t(x) = -\frac{1}{\sigma^3} \mathbb{E}[(G - \mathbb{E}[G \mid X_t = x]) \|G - \mathbb{E}[G \mid X_t = x]\|^2 \mid X_t = x].$$

Proof. We first record the posterior derivative identity used throughout the proof. Let $F : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}^m$ be differentiable in its second argument and assume that $F(D, G)$ and $\partial_{g_j} F(D, G)$ are integrable against the relevant Gaussian posterior weights. In the expression $F(D, G)$, the first input is D , and the second input is G . For $x \in \mathbb{R}^d$, define

$$m_F(x) = \mathbb{E}[F(D, G) \mid X_t = x].$$

Here $\partial_{x_j} m_F(x)$ denotes the derivative of the function $x \mapsto m_F(x)$ with respect to the j -th coordinate of x . Also, $\partial_{g_j} F(D, G)$ denotes the derivative of the function $g \mapsto F(D, g)$ with respect to the j -th coordinate of g , evaluated at $g = G$. Then, for every $j \in \{1, \dots, d\}$,

$$\partial_{x_j} m_F(x) = \frac{1}{\sigma} \mathbb{E}[\partial_{g_j} F(D, G) \mid X_t = x] - \frac{1}{\sigma} \text{Cov}(F(D, G), G_j \mid X_t = x).$$

If F is vector-valued, the covariance is interpreted componentwise:

$$\text{Cov}(F(D, G), G_j \mid X_t = x) = \mathbb{E}[(F(D, G) - \mathbb{E}[F(D, G) \mid X_t = x])(G_j - \mathbb{E}[G_j \mid X_t = x]) \mid X_t = x].$$

We prove the identity. Since $X_t = S_t + \sigma G$, conditioning on $X_t = x$ implies $G = (x - S_t)/\sigma$. Thus $m_F(x)$ can be written as

$$m_F(x) = \frac{\mathbb{E}_{\Gamma(Y_0, Y_1)} \left[F \left(D, \frac{x - S_t}{\sigma} \right) \exp \left(-\frac{1}{2} \left\| \frac{x - S_t}{\sigma} \right\|^2 \right) \right]}{\mathbb{E}_{\Gamma(Y_0, Y_1)} \left[\exp \left(-\frac{1}{2} \left\| \frac{x - S_t}{\sigma} \right\|^2 \right) \right]}.$$

The normalizing constant of the standard Gaussian density is omitted because it appears in both numerator and denominator and cancels.

Define

$$A_F(x) = \mathbb{E}_{\Gamma(Y_0, Y_1)} \left[F \left(D, \frac{x - S_t}{\sigma} \right) \exp \left(-\frac{1}{2} \left\| \frac{x - S_t}{\sigma} \right\|^2 \right) \right],$$

$$B(x) = \mathbb{E}_{\Gamma(Y_0, Y_1)} \left[\exp \left(-\frac{1}{2} \left\| \frac{x - S_t}{\sigma} \right\|^2 \right) \right].$$

Then $m_F(x) = A_F(x)/B(x)$.

We differentiate $A_F(x)$ with respect to x_j . The derivative of the F -term is

$$\partial_{x_j} F \left(D, \frac{x - S_t}{\sigma} \right) = \frac{1}{\sigma} \partial_{g_j} F(D, G).$$

The derivative of the Gaussian weight is

$$\partial_{x_j} \exp \left(-\frac{1}{2} \left\| \frac{x - S_t}{\sigma} \right\|^2 \right) = -\frac{1}{\sigma} G_j \exp \left(-\frac{\|G\|^2}{2} \right).$$

Therefore,

$$\partial_{x_j} A_F(x) = \frac{1}{\sigma} \mathbb{E}_{\Gamma(Y_0, Y_1)} \left[\partial_{g_j} F(D, G) \exp \left(-\frac{\|G\|^2}{2} \right) \right] - \frac{1}{\sigma} \mathbb{E}_{\Gamma(Y_0, Y_1)} \left[F(D, G) G_j \exp \left(-\frac{\|G\|^2}{2} \right) \right].$$

Similarly,

$$\partial_{x_j} B(x) = -\frac{1}{\sigma} \mathbb{E}_{\Gamma(Y_0, Y_1)} \left[G_j \exp \left(-\frac{\|G\|^2}{2} \right) \right].$$

By the quotient rule,

$$\partial_{x_j} m_F(x) = \frac{\partial_{x_j} A_F(x)}{B(x)} - \frac{A_F(x)}{B(x)^2} \partial_{x_j} B(x).$$

Substituting the formulas for $\partial_{x_j} A_F(x)$ and $\partial_{x_j} B(x)$, we obtain

$$\begin{aligned} \partial_{x_j} m_F(x) &= \frac{1}{\sigma} \frac{\mathbb{E}_{\Gamma(Y_0, Y_1)} \left[\partial_{g_j} F(D, G) \exp \left(-\frac{\|G\|^2}{2} \right) \right]}{B(x)} - \frac{1}{\sigma} \frac{\mathbb{E}_{\Gamma(Y_0, Y_1)} \left[F(D, G) G_j \exp \left(-\frac{\|G\|^2}{2} \right) \right]}{B(x)} \\ &\quad + \frac{1}{\sigma} \frac{A_F(x)}{B(x)} \frac{\mathbb{E}_{\Gamma(Y_0, Y_1)} \left[G_j \exp \left(-\frac{\|G\|^2}{2} \right) \right]}{B(x)}. \end{aligned}$$

Each weighted ratio is a conditional expectation given $X_t = x$. Hence

$$\begin{aligned} \partial_{x_j} m_F(x) &= \frac{1}{\sigma} \mathbb{E} \left[\partial_{g_j} F(D, G) \mid X_t = x \right] - \frac{1}{\sigma} \mathbb{E} [F(D, G) G_j \mid X_t = x] \\ &\quad + \frac{1}{\sigma} \mathbb{E} [F(D, G) \mid X_t = x] \mathbb{E} [G_j \mid X_t = x]. \end{aligned}$$

The last two terms are minus the conditional covariance between $F(D, G)$ and G_j . Therefore,

$$\partial_{x_j} m_F(x) = \frac{1}{\sigma} \mathbb{E} \left[\partial_{g_j} F(D, G) \mid X_t = x \right] - \frac{1}{\sigma} \text{Cov} (F(D, G), G_j \mid X_t = x).$$

We now derive the score identity. Recall we have $s_t(x) = -\frac{1}{\sigma} \mathbb{E}[G \mid X_t = x]$.

Apply the posterior derivative identity with $F(D, G) = D_i$. Since D_i has no explicit dependence on G , we have $\partial_{g_j} D_i = 0$. Therefore,

$$\partial_{x_j} v_{t,i}(x) = \partial_{x_j} \mathbb{E}[D_i \mid X_t = x] = -\frac{1}{\sigma} \text{Cov}(D_i, G_j \mid X_t = x).$$

This proves

$$D_x v_t(x) = -\frac{1}{\sigma} \text{Cov}(D, G \mid X_t = x).$$

Next, we compute the time derivative of $v_t(x)$. Fix $x \in \mathbb{R}^d$. For $(y_0, y_1) \in \mathbb{R}^d \times \mathbb{R}^d$, write $s_t(y_0, y_1) = y_0 + t(y_1 - y_0)$. Let γ denote the coupling law of (Y_0, Y_1) . Recall the formula for $v_t(x)$ is

$$v_t(x) = \frac{\int_{\mathbb{R}^d \times \mathbb{R}^d} (y_1 - y_0) \exp\left(-\frac{\|x - s_t(y_0, y_1)\|^2}{2\sigma^2}\right) \gamma(dy_0, dy_1)}{\int_{\mathbb{R}^d \times \mathbb{R}^d} \exp\left(-\frac{\|x - s_t(y_0, y_1)\|^2}{2\sigma^2}\right) \gamma(dy_0, dy_1)}.$$

For every $i \in \{1, \dots, d\}$, define the numerator of the i -th coordinate by

$$A_i(t, x) = \int_{\mathbb{R}^d \times \mathbb{R}^d} (y_1 - y_0)_i \exp\left(-\frac{\|x - s_t(y_0, y_1)\|^2}{2\sigma^2}\right) \gamma(dy_0, dy_1),$$

and define the denominator by

$$B(t, x) = \int_{\mathbb{R}^d \times \mathbb{R}^d} \exp\left(-\frac{\|x - s_t(y_0, y_1)\|^2}{2\sigma^2}\right) \gamma(dy_0, dy_1).$$

Then $v_{t,i}(x) = \frac{A_i(t, x)}{B(t, x)}$. We now differentiate the Gaussian weight with respect to t . Since $\partial_t s_t(y_0, y_1) = y_1 - y_0$, we have

$$\partial_t \exp\left(-\frac{\|x - s_t(y_0, y_1)\|^2}{2\sigma^2}\right) = \frac{(y_1 - y_0)^\top (x - s_t(y_0, y_1))}{\sigma^2} \exp\left(-\frac{\|x - s_t(y_0, y_1)\|^2}{2\sigma^2}\right).$$

Using this derivative inside the integral defining $A_i(t, x)$, we get

$$\begin{aligned} \partial_t A_i(t, x) &= \frac{1}{\sigma^2} \int_{\mathbb{R}^d \times \mathbb{R}^d} (y_1 - y_0)_i (y_1 - y_0)^\top (x - s_t(y_0, y_1)) \cdot \exp\left(-\frac{\|x - s_t(y_0, y_1)\|^2}{2\sigma^2}\right) \gamma(dy_0, dy_1), \\ \partial_t B(t, x) &= \frac{1}{\sigma^2} \int_{\mathbb{R}^d \times \mathbb{R}^d} (y_1 - y_0)^\top (x - s_t(y_0, y_1)) \cdot \exp\left(-\frac{\|x - s_t(y_0, y_1)\|^2}{2\sigma^2}\right) \gamma(dy_0, dy_1). \end{aligned}$$

By the quotient rule,

$$\partial_t v_{t,i}(x) = \frac{\partial_t A_i(t, x)}{B(t, x)} - \frac{A_i(t, x)}{B(t, x)^2} \partial_t B(t, x).$$

Substituting the formulas for $\partial_t A_i(t, x)$ and $\partial_t B(t, x)$, and rewriting the weighted ratios as posterior conditional expectations given $X_t = x$, we obtain

$$\partial_t v_{t,i}(x) = \frac{1}{\sigma^2} \mathbb{E}\left[D_i D^\top (x - S_t) \mid X_t = x\right] - \frac{1}{\sigma^2} \mathbb{E}[D_i \mid X_t = x] \mathbb{E}\left[D^\top (x - S_t) \mid X_t = x\right].$$

The right-hand side is exactly the conditional covariance between D_i and $D^\top(x - S_t)$. Hence

$$\partial_t v_{t,i}(x) = \frac{1}{\sigma^2} \text{Cov} \left(D_i, D^\top(x - S_t) \mid X_t = x \right).$$

Since this holds for every $i \in \{1, \dots, d\}$, we get the vector identity

$$\partial_t v_t(x) = \frac{1}{\sigma^2} \text{Cov} \left(D, D^\top(x - S_t) \mid X_t = x \right).$$

Finally, under the conditioning $X_t = x$, we have $x - S_t = \sigma G$. Therefore,

$$\partial_t v_t(x) = \frac{1}{\sigma} \text{Cov} \left(D, D^\top G \mid X_t = x \right).$$

Now apply the posterior derivative identity with $F(D, G) = G_i$. Since $\partial_{g_j} G_i$ equals 1 if $i = j$, and equals 0 otherwise,

$$\partial_{x_j} \mathbb{E}[G_i \mid X_t = x] = \frac{1}{\sigma} \mathbf{1}_{\{i=j\}} - \frac{1}{\sigma} \text{Cov}(G_i, G_j \mid X_t = x).$$

Since $s_{t,i}(x) = -(1/\sigma)\mathbb{E}[G_i \mid X_t = x]$, we get

$$\partial_{x_j} s_{t,i}(x) = \frac{1}{\sigma^2} \left(\text{Cov}(G_i, G_j \mid X_t = x) - \mathbf{1}_{\{i=j\}} \right).$$

Therefore,

$$D_x s_t(x) = \frac{1}{\sigma^2} (\text{Cov}(G, G \mid X_t = x) - I_d).$$

Finally, we compute the Laplacian of the score. Differentiating the last display once more with respect to x_j , and using the same centered-product calculation as above, gives

$$\partial_{x_j x_j} s_{t,i}(x) = -\frac{1}{\sigma^3} \mathbb{E} \left[(G_i - \mathbb{E}[G_i \mid X_t = x]) (G_j - \mathbb{E}[G_j \mid X_t = x])^2 \mid X_t = x \right].$$

Summing over $j = 1, \dots, d$ gives

$$\Delta_x s_{t,i}(x) = -\frac{1}{\sigma^3} \mathbb{E} \left[(G_i - \mathbb{E}[G_i \mid X_t = x]) \|G - \mathbb{E}[G \mid X_t = x]\|^2 \mid X_t = x \right].$$

Since this holds for every $i \in \{1, \dots, d\}$,

$$\Delta_x s_t(x) = -\frac{1}{\sigma^3} \mathbb{E} \left[(G - \mathbb{E}[G \mid X_t = x]) \|G - \mathbb{E}[G \mid X_t = x]\|^2 \mid X_t = x \right].$$

The preceding argument proves all identities in the statement except the full Hessian formula for v_t . We now establish that remaining identity. Fix $i, j, \ell \in \{1, \dots, d\}$. From the Jacobian identity,

$$\partial_{x_j} v_{t,i}(x) = -\frac{1}{\sigma} \text{Cov}(D_i, G_j \mid X_t = x).$$

Therefore,

$$\partial_{x_j x_\ell} v_{t,i}(x) = -\frac{1}{\sigma} \partial_{x_\ell} \text{Cov}(D_i, G_j \mid X_t = x).$$

By definition,

$$\text{Cov}(D_i, G_j \mid X_t = x) = \mathbb{E}[D_i G_j \mid X_t = x] - \mathbb{E}[D_i \mid X_t = x] \mathbb{E}[G_j \mid X_t = x].$$

Here the conditional expectations are taken with respect to the posterior law of (Y_0, Y_1) given $X_t = x$. Explicitly, under this posterior law,

$$\begin{aligned} D_i &= Y_{1,i} - Y_{0,i}, \\ G_j &= \frac{x_j - ((1-t)Y_{0,j} + tY_{1,j})}{\sigma}. \end{aligned}$$

We differentiate each of the three conditional expectations appearing above. First, apply the posterior derivative identity with

$$F(D, G) = D_i G_j.$$

Since D_i does not depend on G , while G_j depends on the j -th coordinate of G , we have

$$\partial_{g_\ell}(D_i G_j) = D_i \mathbf{1}_{\{j=\ell\}}.$$

Therefore,

$$\partial_{x_\ell} \mathbb{E}[D_i G_j \mid X_t = x] = \frac{1}{\sigma} \mathbb{E}[D_i \mathbf{1}_{\{j=\ell\}} \mid X_t = x] - \frac{1}{\sigma} \text{Cov}(D_i G_j, G_\ell \mid X_t = x).$$

Second, apply the posterior derivative identity with $F(D, G) = D_i$, we obtain

$$\partial_{x_\ell} \mathbb{E}[D_i \mid X_t = x] = -\frac{1}{\sigma} \text{Cov}(D_i, G_\ell \mid X_t = x).$$

Third, apply the posterior derivative identity with $F(D, G) = G_j$, we get

$$\partial_{x_\ell} \mathbb{E}[G_j \mid X_t = x] = \frac{1}{\sigma} \mathbf{1}_{\{j=\ell\}} - \frac{1}{\sigma} \text{Cov}(G_j, G_\ell \mid X_t = x).$$

Now differentiate the covariance formula:

$$\begin{aligned} \partial_{x_\ell} \text{Cov}(D_i, G_j \mid X_t = x) &= \partial_{x_\ell} \mathbb{E}[D_i G_j \mid X_t = x] - (\partial_{x_\ell} \mathbb{E}[D_i \mid X_t = x]) \mathbb{E}[G_j \mid X_t = x] \\ &\quad - \mathbb{E}[D_i \mid X_t = x] (\partial_{x_\ell} \mathbb{E}[G_j \mid X_t = x]). \end{aligned}$$

Substituting the three derivative identities above gives

$$\begin{aligned} \partial_{x_\ell} \text{Cov}(D_i, G_j \mid X_t = x) &= \frac{1}{\sigma} \mathbb{E}[D_i \mathbf{1}_{\{j=\ell\}} \mid X_t = x] - \frac{1}{\sigma} \text{Cov}(D_i G_j, G_\ell \mid X_t = x) \\ &\quad + \frac{1}{\sigma} \text{Cov}(D_i, G_\ell \mid X_t = x) \mathbb{E}[G_j \mid X_t = x] \\ &\quad - \frac{1}{\sigma} \mathbb{E}[D_i \mid X_t = x] \mathbf{1}_{\{j=\ell\}} + \frac{1}{\sigma} \mathbb{E}[D_i \mid X_t = x] \text{Cov}(G_j, G_\ell \mid X_t = x). \end{aligned}$$

The two terms containing $\mathbf{1}_{\{j=\ell\}}$ cancel:

$$\frac{1}{\sigma} (\mathbb{E}[D_i \mid X_t = x] - \mathbb{E}[D_i \mid X_t = x]) \mathbf{1}_{\{j=\ell\}} = 0.$$

Therefore,

$$\begin{aligned} & \partial_{x_\ell} \text{Cov}(D_i, G_j \mid X_t = x) \\ &= -\frac{1}{\sigma} \text{Cov}(D_i G_j, G_\ell \mid X_t = x) + \frac{1}{\sigma} \text{Cov}(D_i, G_\ell \mid X_t = x) \mathbb{E}[G_j \mid X_t = x] \\ & \quad + \frac{1}{\sigma} \mathbb{E}[D_i \mid X_t = x] \text{Cov}(G_j, G_\ell \mid X_t = x). \end{aligned}$$

We now rewrite the right-hand side as a centered third moment. Let

$$\begin{aligned} \overline{D}_i(x) &= \mathbb{E}[D_i \mid X_t = x], \\ \overline{G}_j(x) &= \mathbb{E}[G_j \mid X_t = x], \\ \overline{G}_\ell(x) &= \mathbb{E}[G_\ell \mid X_t = x]. \end{aligned}$$

First,

$$\text{Cov}(D_i G_j, G_\ell \mid X_t = x) = \mathbb{E}[D_i G_j G_\ell \mid X_t = x] - \mathbb{E}[D_i G_j \mid X_t = x] \overline{G}_\ell(x).$$

Second,

$$\text{Cov}(D_i, G_\ell \mid X_t = x) = \mathbb{E}[D_i G_\ell \mid X_t = x] - \overline{D}_i(x) \overline{G}_\ell(x).$$

Third,

$$\text{Cov}(G_j, G_\ell \mid X_t = x) = \mathbb{E}[G_j G_\ell \mid X_t = x] - \overline{G}_j(x) \overline{G}_\ell(x).$$

Substituting these three formulas gives

$$\begin{aligned} & \partial_{x_\ell} \text{Cov}(D_i, G_j \mid X_t = x) \\ &= -\frac{1}{\sigma} \mathbb{E}[D_i G_j G_\ell \mid X_t = x] + \frac{1}{\sigma} \mathbb{E}[D_i G_j \mid X_t = x] \overline{G}_\ell(x) \\ & \quad + \frac{1}{\sigma} \mathbb{E}[D_i G_\ell \mid X_t = x] \overline{G}_j(x) - \frac{1}{\sigma} \overline{D}_i(x) \overline{G}_\ell(x) \overline{G}_j(x) \\ & \quad + \frac{1}{\sigma} \overline{D}_i(x) \mathbb{E}[G_j G_\ell \mid X_t = x] - \frac{1}{\sigma} \overline{D}_i(x) \overline{G}_j(x) \overline{G}_\ell(x). \end{aligned}$$

Combining the two identical last-product terms, this becomes

$$\begin{aligned} & \partial_{x_\ell} \text{Cov}(D_i, G_j \mid X_t = x) \\ &= -\frac{1}{\sigma} \left[\mathbb{E}[D_i G_j G_\ell \mid X_t = x] - \mathbb{E}[D_i G_j \mid X_t = x] \overline{G}_\ell(x) \right. \\ & \quad \left. - \mathbb{E}[D_i G_\ell \mid X_t = x] \overline{G}_j(x) - \overline{D}_i(x) \mathbb{E}[G_j G_\ell \mid X_t = x] \right. \\ & \quad \left. + 2 \overline{D}_i(x) \overline{G}_j(x) \overline{G}_\ell(x) \right]. \end{aligned}$$

The bracket is exactly

$$\mathbb{E} \left[\left(D_i - \overline{D}_i(x) \right) \left(G_j - \overline{G}_j(x) \right) \left(G_\ell - \overline{G}_\ell(x) \right) \mid X_t = x \right].$$

Indeed, expanding this centered product gives

$$\mathbb{E} \left[\left(D_i - \overline{D}_i(x) \right) \left(G_j - \overline{G}_j(x) \right) \left(G_\ell - \overline{G}_\ell(x) \right) \mid X_t = x \right]$$

$$\begin{aligned}
&= \mathbb{E}[D_i G_j G_\ell \mid X_t = x] - \mathbb{E}[D_i G_j \mid X_t = x] \overline{G}_\ell(x) \\
&\quad - \mathbb{E}[D_i G_\ell \mid X_t = x] \overline{G}_j(x) - \overline{D}_i(x) \mathbb{E}[G_j G_\ell \mid X_t = x] \\
&\quad + 2 \overline{D}_i(x) \overline{G}_j(x) \overline{G}_\ell(x).
\end{aligned}$$

Therefore,

$$\begin{aligned}
&\partial_{x_\ell} \text{Cov}(D_i, G_j \mid X_t = x) \\
&= -\frac{1}{\sigma} \mathbb{E}[(D_i - \mathbb{E}[D_i \mid X_t = x])(G_j - \mathbb{E}[G_j \mid X_t = x])(G_\ell - \mathbb{E}[G_\ell \mid X_t = x]) \mid X_t = x].
\end{aligned}$$

Returning to the Hessian of v_t , we obtain

$$\begin{aligned}
\partial_{x_j x_\ell} v_{t,i}(x) &= -\frac{1}{\sigma} \partial_{x_\ell} \text{Cov}(D_i, G_j \mid X_t = x) \\
&= \frac{1}{\sigma^2} \mathbb{E}[(D_i - \mathbb{E}[D_i \mid X_t = x])(G_j - \mathbb{E}[G_j \mid X_t = x])(G_\ell - \mathbb{E}[G_\ell \mid X_t = x]) \mid X_t = x].
\end{aligned}$$

Finally, summing over $j = \ell$ gives the componentwise Laplacian:

$$\begin{aligned}
\Delta_x v_{t,i}(x) &= \sum_{j=1}^d \partial_{x_j x_j} v_{t,i}(x) \\
&= \frac{1}{\sigma^2} \mathbb{E} \left[(D_i - \mathbb{E}[D_i \mid X_t = x]) \sum_{j=1}^d (G_j - \mathbb{E}[G_j \mid X_t = x])^2 \mid X_t = x \right] \\
&= \frac{1}{\sigma^2} \mathbb{E} \left[(D_i - \mathbb{E}[D_i \mid X_t = x]) \|G - \mathbb{E}[G \mid X_t = x]\|^2 \mid X_t = x \right].
\end{aligned}$$

Since this holds for every $i \in \{1, \dots, d\}$, we get

$$\Delta_x v_t(x) = \frac{1}{\sigma^2} \mathbb{E} \left[(D - \mathbb{E}[D \mid X_t = x]) \|G - \mathbb{E}[G \mid X_t = x]\|^2 \mid X_t = x \right].$$

□

G.3 Integrated Derivative Bounds

Lemma 5 (Integrated Bounds Required for the Combined Score Analysis). *Let $Y_0, Y_1, D, S_t, G, X_t, q_t, v_t$, and s_t be as in Lemma 4. Assume*

$$\mathbb{E}[\|Y_1 - Y_0\|^4] < \infty,$$

where this expectation is taken under the joint law of (Y_0, Y_1) . Then, for every $t \in [0, T]$, the following bounds hold.

The full Hessian of v_t satisfies

$$\mathbb{E} \left[\|D_x^2 v_t(X_t)\|_{\mathbb{F}}^2 \right] \leq \frac{16 \sqrt{d(d+2)(d+4)(d+6)}}{\sigma^4} \left(\mathbb{E}[\|Y_1 - Y_0\|^4] \right)^{1/2}.$$

Consequently,

$$\mathbb{E} \left[\|\nabla_x (\nabla_x \cdot v_t)(X_t)\|^2 \right] \leq \frac{16d \sqrt{d(d+2)(d+4)(d+6)}}{\sigma^4} \left(\mathbb{E}[\|Y_1 - Y_0\|^4] \right)^{1/2}.$$

The mixed velocity-score term satisfies

$$\mathbb{E} \left[\left\| (D_x v_t(X_t))^\top s_t(X_t) \right\|^2 \right] \leq \frac{\sqrt{d(d+2)(d+4)(d+6)}}{\sigma^4} \left(\mathbb{E} [\|Y_1 - Y_0\|^4] \right)^{1/2}.$$

The same bound holds for $\mathbb{E} [\|(D_x v_t(X_t)) s_t(X_t)\|^2]$. The pure score terms satisfy

$$\begin{aligned} \mathbb{E} [\|D_x s_t(X_t)\|_F^2] &\leq \frac{2d(d+2) + 2d}{\sigma^4}, \\ \mathbb{E} [\|D_x s_t(X_t)\|_F^4] &\leq \frac{8d(d+2)(d+4)(d+6) + 8d^2}{\sigma^8}, \\ \mathbb{E} [\|s_t(X_t)\|^4] &\leq \frac{d(d+2)}{\sigma^4}, \\ \mathbb{E} [\|D_x s_t(X_t) s_t(X_t)\|^2] &\leq \frac{\sqrt{(8d(d+2)(d+4)(d+6) + 8d^2) d(d+2)}}{\sigma^6}, \\ \mathbb{E} [\|\Delta_x s_t(X_t)\|^2] &\leq \frac{64d(d+2)(d+4)}{\sigma^6}. \end{aligned}$$

All expectations in these bounds are taken under the joint law of (Y_0, Y_1, G) .

Proof. We first prove the full Hessian bound for v_t . By Lemma 4, for every $i, j, \ell \in \{1, \dots, d\}$,

$$\partial_{x_j x_\ell} v_{t,i}(X_t) = \frac{1}{\sigma^2} \mathbb{E} [(D_i - \mathbb{E}[D_i | X_t]) (G_j - \mathbb{E}[G_j | X_t]) (G_\ell - \mathbb{E}[G_\ell | X_t]) | X_t].$$

By definition of the Frobenius norm for the full Hessian tensor,

$$\|D_x^2 v_t(X_t)\|_F^2 = \sum_{i=1}^d \sum_{j=1}^d \sum_{\ell=1}^d \left| \partial_{x_j x_\ell} v_{t,i}(X_t) \right|^2.$$

Substituting the componentwise formula for $\partial_{x_j x_\ell} v_{t,i}(X_t)$, we get

$$\|D_x^2 v_t(X_t)\|_F^2 = \frac{1}{\sigma^4} \sum_{i=1}^d \sum_{j=1}^d \sum_{\ell=1}^d \left| \mathbb{E} [(D_i - \mathbb{E}[D_i | X_t]) (G_j - \mathbb{E}[G_j | X_t]) (G_\ell - \mathbb{E}[G_\ell | X_t]) | X_t] \right|^2.$$

Now apply conditional Cauchy–Schwarz to each fixed triple (i, j, ℓ) . We have

$$\begin{aligned} & \left| \mathbb{E} [(D_i - \mathbb{E}[D_i | X_t]) (G_j - \mathbb{E}[G_j | X_t]) (G_\ell - \mathbb{E}[G_\ell | X_t]) | X_t] \right|^2 \\ & \leq \mathbb{E} [|D_i - \mathbb{E}[D_i | X_t]|^2 | X_t] \mathbb{E} [|G_j - \mathbb{E}[G_j | X_t]|^2 | G_\ell - \mathbb{E}[G_\ell | X_t]|^2 | X_t]. \end{aligned}$$

Therefore,

$$\|D_x^2 v_t(X_t)\|_F^2 \leq \frac{1}{\sigma^4} \sum_{i=1}^d \sum_{j=1}^d \sum_{\ell=1}^d \mathbb{E} [|D_i - \mathbb{E}[D_i | X_t]|^2 | X_t] \mathbb{E} [|G_j - \mathbb{E}[G_j | X_t]|^2 | G_\ell - \mathbb{E}[G_\ell | X_t]|^2 | X_t].$$

Since the first conditional expectation does not depend on j or ℓ , and the second conditional expectation does not depend on i , we can rearrange the sums:

$$\|D_x^2 v_t(X_t)\|_F^2 \leq \frac{1}{\sigma^4} \left(\sum_{i=1}^d \mathbb{E} [|D_i - \mathbb{E}[D_i | X_t]|^2 | X_t] \right)$$

$$\cdot \left(\sum_{j=1}^d \sum_{\ell=1}^d \mathbb{E} \left[|G_j - \mathbb{E}[G_j | X_t]|^2 |G_\ell - \mathbb{E}[G_\ell | X_t]|^2 | X_t \right] \right).$$

The first sum is

$$\sum_{i=1}^d \mathbb{E} \left[|D_i - \mathbb{E}[D_i | X_t]|^2 | X_t \right] = \mathbb{E} \left[\|D - \mathbb{E}[D | X_t]\|^2 | X_t \right].$$

The second double sum is

$$\begin{aligned} & \sum_{j=1}^d \sum_{\ell=1}^d \mathbb{E} \left[|G_j - \mathbb{E}[G_j | X_t]|^2 |G_\ell - \mathbb{E}[G_\ell | X_t]|^2 | X_t \right] \\ &= \mathbb{E} \left[\sum_{j=1}^d \sum_{\ell=1}^d |G_j - \mathbb{E}[G_j | X_t]|^2 |G_\ell - \mathbb{E}[G_\ell | X_t]|^2 | X_t \right] \\ &= \mathbb{E} \left[\left(\sum_{j=1}^d |G_j - \mathbb{E}[G_j | X_t]|^2 \right) \left(\sum_{\ell=1}^d |G_\ell - \mathbb{E}[G_\ell | X_t]|^2 \right) | X_t \right] \\ &= \mathbb{E} \left[\|G - \mathbb{E}[G | X_t]\|^4 | X_t \right]. \end{aligned}$$

Combining the two identities gives

$$\left\| D_x^2 v_t(X_t) \right\|_{\mathbb{F}}^2 \leq \frac{1}{\sigma^4} \mathbb{E} \left[\|D - \mathbb{E}[D | X_t]\|^2 | X_t \right] \mathbb{E} \left[\|G - \mathbb{E}[G | X_t]\|^4 | X_t \right].$$

Taking expectation and applying Cauchy-Schwarz gives

$$\mathbb{E} \left[\left\| D_x^2 v_t(X_t) \right\|_{\mathbb{F}}^2 \right] \leq \frac{1}{\sigma^4} \left(\mathbb{E} \left[\left(\mathbb{E} \left[\|D - \mathbb{E}[D | X_t]\|^2 | X_t \right] \right)^2 \right] \right)^{1/2} \left(\mathbb{E} \left[\left(\mathbb{E} \left[\|G - \mathbb{E}[G | X_t]\|^4 | X_t \right] \right)^2 \right] \right)^{1/2}.$$

For the first factor, conditional variance contraction and conditional Jensen give

$$\begin{aligned} \mathbb{E} \left[\|D - \mathbb{E}[D | X_t]\|^2 | X_t \right] &\leq \mathbb{E}[\|D\|^2 | X_t], \\ \left(\mathbb{E} \left[\left(\mathbb{E} \left[\|D - \mathbb{E}[D | X_t]\|^2 | X_t \right] \right)^2 \right] \right)^{1/2} &\leq \left(\mathbb{E}[\|D\|^4] \right)^{1/2}. \end{aligned}$$

For the second factor, Jensen's inequality gives

$$\left(\mathbb{E} \left[\left(\mathbb{E} \left[\|G - \mathbb{E}[G | X_t]\|^4 | X_t \right] \right)^2 \right] \right)^{1/2} \leq \left(\mathbb{E}[\|G - \mathbb{E}[G | X_t]\|^8] \right)^{1/2}.$$

Using

$$\|G - \mathbb{E}[G | X_t]\|^8 \leq 2^7 \left(\|G\|^8 + \|\mathbb{E}[G | X_t]\|^8 \right),$$

and Jensen's inequality $\mathbb{E}[\|\mathbb{E}[G | X_t]\|^8] \leq \mathbb{E}[\|G\|^8]$, we get

$$\mathbb{E} \left[\|G - \mathbb{E}[G | X_t]\|^8 \right] \leq 2^8 \mathbb{E}[\|G\|^8].$$

Since $G \sim \mathcal{N}(0, I_d)$, we know $\mathbb{E}[\|G\|^8] = d(d+2)(d+4)(d+6)$. Therefore,

$$\left(\mathbb{E} \left[\|G - \mathbb{E}[G | X_t]\|^8 \right] \right)^{1/2} \leq 16 \sqrt{d(d+2)(d+4)(d+6)}.$$

Combining the estimates gives

$$\mathbb{E} \left[\left\| D_x^2 v_t(X_t) \right\|_F^2 \right] \leq \frac{16\sqrt{d(d+2)(d+4)(d+6)}}{\sigma^4} \left(\mathbb{E} [\|D\|^4] \right)^{1/2}.$$

Since $D = Y_1 - Y_0$, this is the first displayed bound.

Next, we explain how the bound for the full Hessian $D_x^2 v_t$ implies a bound for $\nabla_x(\nabla_x \cdot v_t)$ (gradient of divergence).

The full Hessian $D_x^2 v_t(x)$ contains all second derivatives

$$\partial_{x_j x_\ell} v_{t,i}(x), \quad i, j, \ell \in \{1, \dots, d\}.$$

Its Frobenius norm is defined by

$$\left\| D_x^2 v_t(x) \right\|_F^2 = \sum_{i=1}^d \sum_{j=1}^d \sum_{\ell=1}^d \left| \partial_{x_j x_\ell} v_{t,i}(x) \right|^2.$$

By contrast, $\nabla_x(\nabla_x \cdot v_t)(x)$ is a vector in $\mathbb{R}^d$. First,

$$\nabla_x \cdot v_t(x) = \sum_{i=1}^d \partial_{x_i} v_{t,i}(x).$$

Therefore, for every $\ell \in \{1, \dots, d\}$, the ℓ -th coordinate of $\nabla_x(\nabla_x \cdot v_t)(x)$ is

$$\begin{aligned} [\nabla_x(\nabla_x \cdot v_t)(x)]_\ell &= \partial_{x_\ell} \left(\sum_{i=1}^d \partial_{x_i} v_{t,i}(x) \right) \\ &= \sum_{i=1}^d \partial_{x_\ell x_i} v_{t,i}(x). \end{aligned}$$

For each fixed ℓ , Cauchy–Schwarz gives

$$\left| \sum_{i=1}^d \partial_{x_\ell x_i} v_{t,i}(x) \right|^2 \leq d \sum_{i=1}^d |\partial_{x_\ell x_i} v_{t,i}(x)|^2.$$

Now sum this inequality over $\ell = 1, \dots, d$. We obtain

$$\left\| \nabla_x(\nabla_x \cdot v_t)(x) \right\|^2 = \sum_{\ell=1}^d |[\nabla_x(\nabla_x \cdot v_t)(x)]_\ell|^2 = \sum_{\ell=1}^d \left| \sum_{i=1}^d \partial_{x_\ell x_i} v_{t,i}(x) \right|^2 \leq d \sum_{\ell=1}^d \sum_{i=1}^d |\partial_{x_\ell x_i} v_{t,i}(x)|^2.$$

The last double sum is only part of the full Hessian norm, because the full Hessian norm sums over all triples (i, j, ℓ) . Hence

$$\sum_{\ell=1}^d \sum_{i=1}^d |\partial_{x_\ell x_i} v_{t,i}(x)|^2 \leq \sum_{i=1}^d \sum_{j=1}^d \sum_{\ell=1}^d \left| \partial_{x_j x_\ell} v_{t,i}(x) \right|^2 = \left\| D_x^2 v_t(x) \right\|_F^2.$$

Therefore,

$$\left\| \nabla_x(\nabla_x \cdot v_t)(x) \right\|^2 \leq d \left\| D_x^2 v_t(x) \right\|_F^2.$$

Evaluating this inequality at $x = X_t$ and taking expectation gives

$$\mathbb{E} \left[\|\nabla_x (\nabla_x \cdot v_t)(X_t)\|^2 \right] \leq d \mathbb{E} \left[\|D_x^2 v_t(X_t)\|_{\text{F}}^2 \right].$$

Using the full Hessian estimate established above in this proof,

$$\mathbb{E} \left[\|D_x^2 v_t(X_t)\|_{\text{F}}^2 \right] \leq \frac{16\sqrt{d(d+2)(d+4)(d+6)}}{\sigma^4} \left(\mathbb{E} [\|D\|^4] \right)^{1/2},$$

we conclude that

$$\mathbb{E} \left[\|\nabla_x (\nabla_x \cdot v_t)(X_t)\|^2 \right] \leq \frac{16d\sqrt{d(d+2)(d+4)(d+6)}}{\sigma^4} \left(\mathbb{E} [\|D\|^4] \right)^{1/2}.$$

We now prove the mixed velocity-score bound directly. From the posterior derivative identity for v_t ,

$$D_x v_t(X_t) = -\frac{1}{\sigma} \text{Cov}(D, G \mid X_t).$$

From the Gaussian score identity,

$$s_t(X_t) = -\frac{1}{\sigma} \mathbb{E}[G \mid X_t].$$

Therefore,

$$(D_x v_t(X_t))^\top s_t(X_t) = \frac{1}{\sigma^2} \text{Cov}(D, G \mid X_t)^\top \mathbb{E}[G \mid X_t].$$

Hence

$$\left\| (D_x v_t(X_t))^\top s_t(X_t) \right\|^2 = \frac{1}{\sigma^4} \left\| \text{Cov}(D, G \mid X_t)^\top \mathbb{E}[G \mid X_t] \right\|^2.$$

We now bound the conditional quantity. Let $\bar{D}(X_t) = \mathbb{E}[D \mid X_t]$ and $\bar{G}(X_t) = \mathbb{E}[G \mid X_t]$. By the definition of conditional covariance,

$$\begin{aligned} \text{Cov}(D, G \mid X_t)^\top \bar{G}(X_t) &= \mathbb{E} \left[\left(G - \bar{G}(X_t) \right) \left(D - \bar{D}(X_t) \right)^\top \mid X_t \right] \bar{G}(X_t) \\ &= \mathbb{E} \left[\left(G - \bar{G}(X_t) \right) \left(D - \bar{D}(X_t) \right)^\top \bar{G}(X_t) \mid X_t \right]. \end{aligned} \quad (6)$$

Since $\bar{G}(X_t)$ is fixed after conditioning on X_t , conditional Cauchy–Schwarz gives

$$\begin{aligned} &\left\| \text{Cov}(D, G \mid X_t)^\top \bar{G}(X_t) \right\|^2 \\ &\leq \mathbb{E} \left[\left\| G - \bar{G}(X_t) \right\|^2 \mid X_t \right] \mathbb{E} \left[\left\| \left(D - \bar{D}(X_t) \right)^\top \bar{G}(X_t) \right\|^2 \mid X_t \right]. \end{aligned}$$

For the second factor,

$$\left\| \left(D - \bar{D}(X_t) \right)^\top \bar{G}(X_t) \right\|^2 \leq \left\| D - \bar{D}(X_t) \right\|^2 \left\| \bar{G}(X_t) \right\|^2.$$

Therefore,

$$\begin{aligned} & \left\| \text{Cov}(D, G \mid X_t)^\top \overline{G}(X_t) \right\|^2 \\ & \leq \mathbb{E} \left[\left\| G - \overline{G}(X_t) \right\|^2 \mid X_t \right] \mathbb{E} \left[\left\| D - \overline{D}(X_t) \right\|^2 \mid X_t \right] \left\| \overline{G}(X_t) \right\|^2. \end{aligned}$$

Using the conditional variance inequality,

$$\begin{aligned} \mathbb{E} \left[\left\| D - \overline{D}(X_t) \right\|^2 \mid X_t \right] & \leq \mathbb{E} \left[\|D\|^2 \mid X_t \right], \\ \mathbb{E} \left[\left\| G - \overline{G}(X_t) \right\|^2 \mid X_t \right] & \leq \mathbb{E} \left[\|G\|^2 \mid X_t \right], \end{aligned}$$

and Jensen's inequality, $\left\| \overline{G}(X_t) \right\|^2 = \left\| \mathbb{E}[G \mid X_t] \right\|^2 \leq \mathbb{E}[\|G\|^2 \mid X_t]$, we get

$$\left\| \text{Cov}(D, G \mid X_t)^\top \overline{G}(X_t) \right\|^2 \leq \mathbb{E}[\|D\|^2 \mid X_t] \left(\mathbb{E}[\|G\|^2 \mid X_t] \right)^2.$$

Taking expectation gives

$$\mathbb{E} \left[\left\| \text{Cov}(D, G \mid X_t)^\top \mathbb{E}[G \mid X_t] \right\|^2 \right] \leq \mathbb{E} \left[\mathbb{E}[\|D\|^2 \mid X_t] \left(\mathbb{E}[\|G\|^2 \mid X_t] \right)^2 \right].$$

By Cauchy-Schwarz,

$$\begin{aligned} & \mathbb{E} \left[\mathbb{E}[\|D\|^2 \mid X_t] \left(\mathbb{E}[\|G\|^2 \mid X_t] \right)^2 \right] \\ & \leq \left(\mathbb{E} \left[\left(\mathbb{E}[\|D\|^2 \mid X_t] \right)^2 \right] \right)^{1/2} \left(\mathbb{E} \left[\left(\mathbb{E}[\|G\|^2 \mid X_t] \right)^4 \right] \right)^{1/2}. \end{aligned}$$

By Jensen's inequality,

$$\begin{aligned} \mathbb{E} \left[\left(\mathbb{E}[\|D\|^2 \mid X_t] \right)^2 \right] & \leq \mathbb{E}[\|D\|^4], \\ \mathbb{E} \left[\left(\mathbb{E}[\|G\|^2 \mid X_t] \right)^4 \right] & \leq \mathbb{E}[\|G\|^8]. \end{aligned}$$

Since $G \sim \mathcal{N}(0, I_d)$,

$$\mathbb{E}[\|G\|^8] = d(d+2)(d+4)(d+6).$$

Therefore,

$$\begin{aligned} & \mathbb{E} \left[\left\| \text{Cov}(D, G \mid X_t)^\top \mathbb{E}[G \mid X_t] \right\|^2 \right] \\ & \leq \sqrt{d(d+2)(d+4)(d+6)} \left(\mathbb{E}[\|D\|^4] \right)^{1/2}. \end{aligned}$$

Finally, recalling the factor $1/\sigma^4$, we conclude that

$$\mathbb{E} \left[\left\| (D_x v_t(X_t))^\top s_t(X_t) \right\|^2 \right] \leq \frac{\sqrt{d(d+2)(d+4)(d+6)}}{\sigma^4} \left(\mathbb{E}[\|D\|^4] \right)^{1/2}.$$

The non-transposed orientation requires a separate argument. Write $C_t = \text{Cov}(D, G \mid X_t)$ and $\bar{G}_t = \mathbb{E}[G \mid X_t]$. Conditional Cauchy–Schwarz gives

$$\|C_t\|_F^2 \leq \mathbb{E}[\|D - \mathbb{E}[D \mid X_t]\|^2 \mid X_t] \mathbb{E}[\|G - \bar{G}_t\|^2 \mid X_t].$$

Consequently,

$$\|C_t \bar{G}_t\|^2 \leq \mathbb{E}[\|D\|^2 \mid X_t] (\mathbb{E}[\|G\|^2 \mid X_t])^2,$$

where we used conditional variance contraction and $\|\bar{G}_t\|^2 \leq \mathbb{E}[\|G\|^2 \mid X_t]$. The same outer Cauchy–Schwarz and Jensen calculation as above therefore proves

$$\mathbb{E}[(D_x v_t(X_t)) s_t(X_t)]^2 \leq \frac{\sqrt{d(d+2)(d+4)(d+6)}}{\sigma^4} (\mathbb{E}\|D\|^4)^{1/2}.$$

We next prove the pure score bounds. By Lemma 4,

$$D_x s_t(X_t) = \frac{1}{\sigma^2} (\text{Cov}(G, G \mid X_t) - I_d).$$

Using $\|A - B\|_F^2 \leq 2\|A\|_F^2 + 2\|B\|_F^2$,

$$\|D_x s_t(X_t)\|_F^2 \leq \frac{2}{\sigma^4} \|\text{Cov}(G, G \mid X_t)\|_F^2 + \frac{2}{\sigma^4} \|I_d\|_F^2.$$

Since $\|I_d\|_F^2 = d$, and since the conditional covariance matrix is positive semidefinite,

$$\begin{aligned} \|\text{Cov}(G, G \mid X_t)\|_F^2 &\leq (\text{tr Cov}(G, G \mid X_t))^2 = \left(\mathbb{E}[\|G - \mathbb{E}[G \mid X_t]\|^2 \mid X_t]\right)^2 \\ &= \left(\mathbb{E}[\|G\|^2 \mid X_t] - \|\mathbb{E}[G \mid X_t]\|^2\right)^2 \\ &\leq \left(\mathbb{E}[\|G\|^2 \mid X_t]\right)^2. \end{aligned}$$

Taking expectation and using Jensen's inequality gives

$$\begin{aligned} \mathbb{E}[\|\text{Cov}(G, G \mid X_t)\|_F^2] &\leq \mathbb{E}[\|G\|^4] \\ &= d(d+2). \end{aligned}$$

Therefore,

$$\mathbb{E}[\|D_x s_t(X_t)\|_F^2] \leq \frac{2d(d+2) + 2d}{\sigma^4}.$$

For the fourth moment of $D_x s_t(X_t)$, use $\|A - B\|_F^4 \leq 8\|A\|_F^4 + 8\|B\|_F^4$:

$$\|D_x s_t(X_t)\|_F^4 \leq \frac{8}{\sigma^8} \|\text{Cov}(G, G \mid X_t)\|_F^4 + \frac{8}{\sigma^8} \|I_d\|_F^4.$$

Since $\|I_d\|_F^4 = d^2$, and since the conditional covariance matrix is positive semidefinite,

$$\|\text{Cov}(G, G \mid X_t)\|_F^4 \leq (\text{tr Cov}(G, G \mid X_t))^4 \leq \left(\mathbb{E}[\|G\|^2 \mid X_t]\right)^4.$$

Taking expectation and applying conditional Jensen only at this stage gives

$$\mathbb{E}[\|\text{Cov}(G, G \mid X_t)\|_F^4] \leq \mathbb{E}\left[\left(\mathbb{E}[\|G\|^2 \mid X_t]\right)^4\right]$$

$$\leq \mathbb{E}\|G\|^8 = d(d+2)(d+4)(d+6).$$

Therefore,

$$\mathbb{E} \left[\|D_x s_t(X_t)\|_{\mathbb{F}}^4 \right] \leq \frac{8d(d+2)(d+4)(d+6) + 8d^2}{\sigma^8}.$$

The fourth moment of the score follows directly from $s_t(X_t) = -\sigma^{-1}\mathbb{E}[G \mid X_t]$:

$$\begin{aligned} \mathbb{E} \left[\|s_t(X_t)\|^4 \right] &= \frac{1}{\sigma^4} \mathbb{E} \left[\|\mathbb{E}[G \mid X_t]\|^4 \right] \\ &\leq \frac{1}{\sigma^4} \mathbb{E} \left[\mathbb{E}[\|G\|^4 \mid X_t] \right] = \frac{d(d+2)}{\sigma^4}. \end{aligned}$$

Hence, by Cauchy–Schwarz,

$$\begin{aligned} &\mathbb{E} \left[\|D_x s_t(X_t) s_t(X_t)\|^2 \right] \\ &\leq \left(\mathbb{E} \left[\|D_x s_t(X_t)\|_{\mathbb{F}}^4 \right] \right)^{1/2} \left(\mathbb{E} \left[\|s_t(X_t)\|^4 \right] \right)^{1/2} \\ &\leq \frac{\sqrt{(8d(d+2)(d+4)(d+6) + 8d^2) d(d+2)}}{\sigma^6}. \end{aligned}$$

Finally, Lemma 4 gives

$$\Delta_x s_t(X_t) = -\frac{1}{\sigma^3} \mathbb{E} \left[(G - \mathbb{E}[G \mid X_t]) \|G - \mathbb{E}[G \mid X_t]\|^2 \mid X_t \right].$$

By Jensen’s inequality,

$$\|\Delta_x s_t(X_t)\|^2 \leq \frac{1}{\sigma^6} \mathbb{E} \left[\|G - \mathbb{E}[G \mid X_t]\|^6 \mid X_t \right].$$

Taking expectation gives

$$\mathbb{E} \left[\|\Delta_x s_t(X_t)\|^2 \right] \leq \frac{1}{\sigma^6} \mathbb{E} \left[\|G - \mathbb{E}[G \mid X_t]\|^6 \right].$$

Using $\|u - w\|^6 \leq 2^5 \|u\|^6 + 2^5 \|w\|^6$, and using Jensen’s inequality,

$$\mathbb{E} \left[\|G - \mathbb{E}[G \mid X_t]\|^6 \right] \leq 2^6 \mathbb{E} \left[\|G\|^6 \right].$$

Since $G \sim \mathcal{N}(0, I_d)$,

$$\mathbb{E} \left[\|G\|^6 \right] = d(d+2)(d+4).$$

Therefore,

$$\mathbb{E} \left[\|\Delta_x s_t(X_t)\|^2 \right] \leq \frac{64d(d+2)(d+4)}{\sigma^6}.$$

This completes the proof. □

Lemma 6 (Integrated Derivative Bounds in Terms of Displacement Moments). *Let $Y_0, Y_1, D, S_t, G, X_t, q_t$, and v_t be as in Lemma 4. Assume that*

$$\mathbb{E} \left[\|Y_1 - Y_0\|^5 \right] < \infty,$$

where the expectation is taken under the joint law of (Y_0, Y_1) . Then, for every $t \in [0, T]$,

$$\begin{aligned} \mathbb{E} \left[\|D_x v_t(X_t)\|_{\mathbb{F}}^2 \right] &\leq \frac{\sqrt{d(d+2)}}{\sigma^2} \left(\mathbb{E} \left[\|Y_1 - Y_0\|^4 \right] \right)^{1/2}, \\ \mathbb{E} \left[\|\partial_t v_t(X_t)\|^2 \right] &\leq \frac{\sqrt{d(d+2)}}{\sigma^2} \mathbb{E} \left[\|Y_1 - Y_0\|^4 \right], \\ \mathbb{E} \left[\|\Delta_x v_t(X_t)\|^2 \right] &\leq \frac{16\sqrt{d(d+2)(d+4)(d+6)}}{\sigma^4} \left(\mathbb{E} \left[\|Y_1 - Y_0\|^4 \right] \right)^{1/2}. \end{aligned}$$

Moreover,

$$\begin{aligned} \mathbb{E} \left[\|D_x v_t(X_t) b_t(X_t)\|^2 \right] &\leq \frac{2}{\sigma^2} \left(\mathbb{E} \left[\|Y_1 - Y_0\|^5 \right] \right)^{4/5} (d(d+2)(d+4)(d+6)(d+8))^{1/5} \\ &\quad + \frac{\tau^4}{2\sigma^4} \sqrt{d(d+2)(d+4)(d+6)} \left(\mathbb{E} \left[\|Y_1 - Y_0\|^4 \right] \right)^{1/2}. \end{aligned}$$

All expectations in these bounds are taken under the joint law of (Y_0, Y_1, G) .

Proof. We first bound the Jacobian. By Lemma 4,

$$D_x v_t(X_t) = -\frac{1}{\sigma} \text{Cov}(D, G \mid X_t).$$

By Cauchy-Schwarz, $\|\text{Cov}(D, G \mid X_t)\|_{\mathbb{F}}^2 \leq \mathbb{E}[\|D\|^2 \mid X_t] \mathbb{E}[\|G\|^2 \mid X_t]$. Taking expectation over X_t and apply Cauchy-Schwarz, we get

$$\begin{aligned} \mathbb{E} \left[\|\text{Cov}(D, G \mid X_t)\|_{\mathbb{F}}^2 \right] &\leq \mathbb{E} \left[\mathbb{E}[\|D\|^2 \mid X_t] \mathbb{E}[\|G\|^2 \mid X_t] \right] \\ &\leq \left(\mathbb{E} \left[\left(\mathbb{E}[\|D\|^2 \mid X_t] \right)^2 \right] \right)^{1/2} \left(\mathbb{E} \left[\left(\mathbb{E}[\|G\|^2 \mid X_t] \right)^2 \right] \right)^{1/2} \\ &\leq \left(\mathbb{E}[\|D\|^4] \right)^{1/2} \left(\mathbb{E}[\|G\|^4] \right)^{1/2}. \end{aligned}$$

Since $G \sim \mathcal{N}(0, I_d)$, $\mathbb{E}[\|G\|^4] = d(d+2)$. Therefore,

$$\mathbb{E} \left[\|D_x v_t(X_t)\|_{\mathbb{F}}^2 \right] = \frac{1}{\sigma^2} \mathbb{E} \left[\|\text{Cov}(D, G \mid X_t)\|_{\mathbb{F}}^2 \right] \leq \frac{\sqrt{d(d+2)}}{\sigma^2} \left(\mathbb{E}[\|D\|^4] \right)^{1/2}.$$

For time derivative, by Lemma 4,

$$\partial_t v_t(X_t) = \frac{1}{\sigma} \text{Cov}(D, D^\top G \mid X_t).$$

Similarly, apply Cauchy-Schwarz, we have

$$\left\| \text{Cov}(D, D^\top G \mid X_t) \right\|^2 \leq \mathbb{E}[\|D\|^2 \mid X_t] \mathbb{E}[\|D^\top G\|^2 \mid X_t].$$

Take expectation and apply Cauchy-Schwarz, we have

$$\mathbb{E} \left\| \text{Cov}(D, D^\top G \mid X_t) \right\|^2 \leq \mathbb{E} \left[\left(\mathbb{E}[\|D\|^2 \mid X_t] \right)^2 \right]^{\frac{1}{2}} \mathbb{E} \left[\left(\mathbb{E}[\|D\|^2 \|G\|^2 \mid X_t] \right)^2 \right]^{\frac{1}{2}}$$

$$\leq \mathbb{E} \left[\|D\|^4 \right]^{\frac{1}{2}} \mathbb{E} \left[\|D\|^4 \|G\|^4 \right]^{\frac{1}{2}} = \mathbb{E} \left[\|D\|^4 \right] \mathbb{E} \left[\|G\|^4 \right]^{\frac{1}{2}},$$

where note that D and G are independent. Therefore,

$$\mathbb{E} \left[\|\partial_t v_t(X_t)\|^2 \right] \leq \frac{\sqrt{d(d+2)}}{\sigma^2} \mathbb{E} \left[\|D\|^4 \right].$$

Now consider the Laplacian. By Lemma 4,

$$\Delta_x v_t(X_t) = \frac{1}{\sigma^2} \mathbb{E} \left[(D - \mathbb{E}[D \mid X_t]) \|G - \mathbb{E}[G \mid X_t]\|^2 \mid X_t \right].$$

Conditional Cauchy–Schwarz gives

$$\|\Delta_x v_t(X_t)\|^2 \leq \frac{1}{\sigma^4} \mathbb{E} \left[\|D - \mathbb{E}[D \mid X_t]\|^2 \mid X_t \right] \mathbb{E} \left[\|G - \mathbb{E}[G \mid X_t]\|^4 \mid X_t \right].$$

Following the same proof as we did for $D_x^2 v_t$, we obtain

$$\mathbb{E} \left[\|\Delta_x v_t(X_t)\|^2 \right] \leq \frac{16\sqrt{d(d+2)(d+4)(d+6)}}{\sigma^4} \left(\mathbb{E} \left[\|D\|^4 \right] \right)^{1/2}.$$

We now bound the term

$$\mathbb{E} \left[\|D_x v_t(X_t) b_t(X_t)\|^2 \right].$$

Recall that $s_t(X_t) = -(1/\sigma)\mathbb{E}[G \mid X_t]$ and $b_t(X_t) = v_t(X_t) + \frac{\tau^2}{2}s_t(X_t)$. Therefore,

$$D_x v_t(X_t) b_t(X_t) = D_x v_t(X_t) v_t(X_t) + \frac{\tau^2}{2} D_x v_t(X_t) s_t(X_t).$$

Using $\|a + b\|^2 \leq 2\|a\|^2 + 2\|b\|^2$, we get

$$\|D_x v_t(X_t) b_t(X_t)\|^2 \leq 2 \|D_x v_t(X_t) v_t(X_t)\|^2 + \frac{\tau^4}{2} \|D_x v_t(X_t) s_t(X_t)\|^2.$$

Thus,

$$\begin{aligned} & \mathbb{E} \left[\|D_x v_t(X_t) b_t(X_t)\|^2 \right] \\ & \leq 2 \mathbb{E} \left[\|D_x v_t(X_t) v_t(X_t)\|^2 \right] + \frac{\tau^4}{2} \mathbb{E} \left[\|D_x v_t(X_t) s_t(X_t)\|^2 \right]. \end{aligned}$$

We first bound the velocity part. From the posterior derivative identity,

$$D_x v_t(X_t) = -\frac{1}{\sigma} \text{Cov}(D, G \mid X_t).$$

Since $v_t(X_t) = \mathbb{E}[D \mid X_t]$, we have

$$D_x v_t(X_t) v_t(X_t) = -\frac{1}{\sigma} \text{Cov}(D, G \mid X_t) \mathbb{E}[D \mid X_t].$$

For readability, define

$$\overline{D}(X_t) = \mathbb{E}[D \mid X_t],$$

$$\overline{G}(X_t) = \mathbb{E}[G \mid X_t].$$

By the definition of conditional covariance,

$$\begin{aligned} \text{Cov}(D, G \mid X_t) \overline{D}(X_t) &= \mathbb{E} \left[\left(D - \overline{D}(X_t) \right) \left(G - \overline{G}(X_t) \right)^\top \mid X_t \right] \overline{D}(X_t) \\ &= \mathbb{E} \left[\left(D - \overline{D}(X_t) \right) \left(G - \overline{G}(X_t) \right)^\top \overline{D}(X_t) \mid X_t \right]. \end{aligned}$$

Since $\overline{D}(X_t)$ is fixed after conditioning on X_t , conditional Cauchy–Schwarz gives

$$\begin{aligned} &\left\| \text{Cov}(D, G \mid X_t) \overline{D}(X_t) \right\|^2 \\ &\leq \mathbb{E} \left[\left\| D - \overline{D}(X_t) \right\|^2 \mid X_t \right] \mathbb{E} \left[\left| \left(G - \overline{G}(X_t) \right)^\top \overline{D}(X_t) \right|^2 \mid X_t \right]. \end{aligned}$$

For the second conditional expectation,

$$\left| \left(G - \overline{G}(X_t) \right)^\top \overline{D}(X_t) \right|^2 \leq \left\| G - \overline{G}(X_t) \right\|^2 \left\| \overline{D}(X_t) \right\|^2.$$

Hence

$$\begin{aligned} &\left\| \text{Cov}(D, G \mid X_t) \overline{D}(X_t) \right\|^2 \\ &\leq \mathbb{E} \left[\left\| D - \overline{D}(X_t) \right\|^2 \mid X_t \right] \mathbb{E} \left[\left\| G - \overline{G}(X_t) \right\|^2 \mid X_t \right] \left\| \overline{D}(X_t) \right\|^2. \end{aligned}$$

Using the conditional variance inequalities,

$$\begin{aligned} \mathbb{E} \left[\left\| D - \overline{D}(X_t) \right\|^2 \mid X_t \right] &\leq \mathbb{E} \left[\|D\|^2 \mid X_t \right], \\ \mathbb{E} \left[\left\| G - \overline{G}(X_t) \right\|^2 \mid X_t \right] &\leq \mathbb{E} \left[\|G\|^2 \mid X_t \right], \end{aligned}$$

and Jensen's inequality,

$$\begin{aligned} \left\| \overline{D}(X_t) \right\|^2 &= \left\| \mathbb{E}[D \mid X_t] \right\|^2 \\ &\leq \mathbb{E} \left[\|D\|^2 \mid X_t \right], \end{aligned}$$

we obtain

$$\left\| \text{Cov}(D, G \mid X_t) \overline{D}(X_t) \right\|^2 \leq \left(\mathbb{E} \left[\|D\|^2 \mid X_t \right] \right)^2 \mathbb{E} \left[\|G\|^2 \mid X_t \right].$$

Therefore,

$$\begin{aligned} &\mathbb{E} \left[\left\| D_x v_t(X_t) v_t(X_t) \right\|^2 \right] \\ &\leq \frac{1}{\sigma^2} \mathbb{E} \left[\left(\mathbb{E} \left[\|D\|^2 \mid X_t \right] \right)^2 \mathbb{E} \left[\|G\|^2 \mid X_t \right] \right]. \end{aligned}$$

Apply Hölder's inequality with conjugate exponents 5/4 and 5. Then

$$\mathbb{E} \left[\left(\mathbb{E} \left[\|D\|^2 \mid X_t \right] \right)^2 \mathbb{E} \left[\|G\|^2 \mid X_t \right] \right]$$

$$\leq \left(\mathbb{E} \left[\left(\mathbb{E} [\|D\|^2 \mid X_t] \right)^{5/2} \right] \right)^{4/5} \left(\mathbb{E} \left[\left(\mathbb{E} [\|G\|^2 \mid X_t] \right)^5 \right] \right)^{1/5}.$$

For the first factor, Jensen's inequality gives

$$\mathbb{E} \left[\left(\mathbb{E} [\|D\|^2 \mid X_t] \right)^{5/2} \right] \leq \mathbb{E} [\|D\|^5].$$

For the second factor, Jensen's inequality gives

$$\mathbb{E} \left[\left(\mathbb{E} [\|G\|^2 \mid X_t] \right)^5 \right] \leq \mathbb{E} [\|G\|^{10}].$$

Since $G \sim \mathcal{N}(0, I_d)$,

$$\mathbb{E} [\|G\|^{10}] = d(d+2)(d+4)(d+6)(d+8).$$

Thus,

$$\begin{aligned} \mathbb{E} [\|D_x v_t(X_t) v_t(X_t)\|^2] &\leq \frac{1}{\sigma^2} \left(\mathbb{E} [\|D\|^5] \right)^{4/5} \\ &\quad \cdot (d(d+2)(d+4)(d+6)(d+8))^{1/5}. \end{aligned}$$

Also recall

$$\mathbb{E} [\|D_x v_t(X_t) s_t(X_t)\|^2] \leq \frac{\sqrt{d(d+2)(d+4)(d+6)}}{\sigma^4} \left(\mathbb{E} [\|D\|^4] \right)^{1/2}.$$

Combining the velocity part and the score part, we obtain

$$\begin{aligned} &\mathbb{E} [\|D_x v_t(X_t) b_t(X_t)\|^2] \\ &\leq \frac{2}{\sigma^2} \left(\mathbb{E} [\|D\|^5] \right)^{4/5} (d(d+2)(d+4)(d+6)(d+8))^{1/5} \\ &\quad + \frac{\tau^4}{2\sigma^4} \sqrt{d(d+2)(d+4)(d+6)} \left(\mathbb{E} [\|D\|^4] \right)^{1/2}. \end{aligned}$$

Since $D = Y_1 - Y_0$, this can be written as

$$\begin{aligned} &\mathbb{E} [\|D_x v_t(X_t) b_t(X_t)\|^2] \\ &\leq \frac{2}{\sigma^2} \left(\mathbb{E} [\|Y_1 - Y_0\|^5] \right)^{4/5} (d(d+2)(d+4)(d+6)(d+8))^{1/5} \\ &\quad + \frac{\tau^4}{2\sigma^4} \sqrt{d(d+2)(d+4)(d+6)} \left(\mathbb{E} [\|Y_1 - Y_0\|^4] \right)^{1/2}. \end{aligned}$$

□

G.4 Euler Increment Bounds

Lemma 7 (Itô Control of the Euler Velocity Term). *Let $Y_0, Y_1, D, S_t, G, X_t, q_t, v_t$, and s_t be as in Lemma 4, assume $\mathbb{E} \|Y_1 - Y_0\|^5 < \infty$, and set $b_t = v_t + \frac{\tau^2}{2} s_t$. Let Z_t solve*

$$dZ_t = b_t(Z_t)dt + \tau dB_t,$$

and assume that $Z_t \sim q_t$ for every $t \in [0, T]$. Let $0 = t_0 < t_1 < \dots < t_N = T$ be a uniform grid with step size h , so $t_{k+1} - t_k = h$ and $Nh = T$. Then

$$\begin{aligned}
& \sum_{k=0}^{N-1} \int_{t_k}^{t_{k+1}} \mathbb{E} \left[\|v_t(Z_t) - v_{t_k}(Z_{t_k})\|^2 \right] dt \\
& \leq \frac{2Th^2}{3} \left[\frac{3\sqrt{d(d+2)}}{\sigma^2} \mathbb{E} \left[\|Y_1 - Y_0\|^4 \right] \right. \\
& \quad + \frac{6}{\sigma^2} \left(\mathbb{E} \left[\|Y_1 - Y_0\|^5 \right] \right)^{4/5} (d(d+2)(d+4)(d+6)(d+8))^{1/5} \\
& \quad + \frac{14\tau^4 \sqrt{d(d+2)(d+4)(d+6)}}{\sigma^4} \left(\mathbb{E} \left[\|Y_1 - Y_0\|^4 \right] \right)^{1/2} \left. \right] \\
& \quad + \frac{\tau^2 Th \sqrt{d(d+2)}}{\sigma^2} \left(\mathbb{E} \left[\|Y_1 - Y_0\|^4 \right] \right)^{1/2}.
\end{aligned}$$

All expectations in the displayed bound are taken under the coupling law of (Y_0, Y_1) , except the expectation on the left-hand side, which is taken under the law of the exact SDE path $(Z_t)_{0 \leq t \leq T}$.

Proof. For $t \in [t_k, t_{k+1})$, Itô's formula applied to $v_t(Z_t)$ gives

$$\begin{aligned}
v_t(Z_t) - v_{t_k}(Z_{t_k}) &= \int_{t_k}^t \left(\partial_r v_r(Z_r) + D_x v_r(Z_r) b_r(Z_r) + \frac{\tau^2}{2} \Delta_x v_r(Z_r) \right) dr \\
&\quad + \tau \int_{t_k}^t D_x v_r(Z_r) dB_r.
\end{aligned}$$

Using $\|u + w\|^2 \leq 2\|u\|^2 + 2\|w\|^2$, Jensen's inequality, and Itô isometry, we get

$$\begin{aligned}
& \mathbb{E} \left[\|v_t(Z_t) - v_{t_k}(Z_{t_k})\|^2 \right] \\
& \leq 2(t - t_k) \int_{t_k}^t \mathbb{E} \left[\left\| \partial_r v_r(Z_r) + D_x v_r(Z_r) b_r(Z_r) + \frac{\tau^2}{2} \Delta_x v_r(Z_r) \right\|^2 \right] dr \\
& \quad + 2\tau^2 \int_{t_k}^t \mathbb{E} \left[\|D_x v_r(Z_r)\|_{\mathbb{F}}^2 \right] dr.
\end{aligned}$$

Since $Z_r \sim q_r$ and $X_r \sim q_r$, Lemma 6 applies with X_r replaced by Z_r . Using $\|u + w + z\|^2 \leq 3\|u\|^2 + 3\|w\|^2 + 3\|z\|^2$, we have

$$\begin{aligned}
& \mathbb{E} \left[\left\| \partial_r v_r(Z_r) + D_x v_r(Z_r) b_r(Z_r) + \frac{\tau^2}{2} \Delta_x v_r(Z_r) \right\|^2 \right] \\
& \leq 3\mathbb{E} \left[\|\partial_r v_r(Z_r)\|^2 \right] + 3\mathbb{E} \left[\|D_x v_r(Z_r) b_r(Z_r)\|^2 \right] + 3\mathbb{E} \left[\left\| \frac{\tau^2}{2} \Delta_x v_r(Z_r) \right\|^2 \right] \\
& \leq \frac{3\sqrt{d(d+2)}}{\sigma^2} \mathbb{E} \left[\|Y_1 - Y_0\|^4 \right] \\
& \quad + \frac{6}{\sigma^2} \left(\mathbb{E} \left[\|Y_1 - Y_0\|^5 \right] \right)^{4/5} (d(d+2)(d+4)(d+6)(d+8))^{1/5} \\
& \quad + \frac{3\tau^4}{2\sigma^4} \sqrt{d(d+2)(d+4)(d+6)} \left(\mathbb{E} \left[\|Y_1 - Y_0\|^4 \right] \right)^{1/2}
\end{aligned}$$

$$+ \frac{12\tau^4 \sqrt{d(d+2)(d+4)(d+6)}}{\sigma^4} \left(\mathbb{E} [\|Y_1 - Y_0\|^4] \right)^{1/2}.$$

Also,

$$\mathbb{E} [\|D_x v_r(Z_r)\|_F^2] \leq \frac{\sqrt{d(d+2)}}{\sigma^2} \left(\mathbb{E} [\|Y_1 - Y_0\|^4] \right)^{1/2}.$$

Therefore, for $t \in [t_k, t_{k+1})$,

$$\begin{aligned} & \mathbb{E} [\|v_t(Z_t) - v_{t_k}(Z_{t_k})\|^2] \\ & \leq 2(t - t_k)^2 \left[\frac{3\sqrt{d(d+2)}}{\sigma^2} \mathbb{E} [\|Y_1 - Y_0\|^4] \right. \\ & \quad + \frac{6}{\sigma^2} \left(\mathbb{E} [\|Y_1 - Y_0\|^5] \right)^{4/5} (d(d+2)(d+4)(d+6)(d+8))^{1/5} \\ & \quad + \frac{3\tau^4}{2\sigma^4} \sqrt{d(d+2)(d+4)(d+6)} \left(\mathbb{E} [\|Y_1 - Y_0\|^4] \right)^{1/2} \\ & \quad + \frac{12\tau^4 \sqrt{d(d+2)(d+4)(d+6)}}{\sigma^4} \left(\mathbb{E} [\|Y_1 - Y_0\|^4] \right)^{1/2} \left. \right] \\ & \quad + 2\tau^2(t - t_k) \frac{\sqrt{d(d+2)}}{\sigma^2} \left(\mathbb{E} [\|Y_1 - Y_0\|^4] \right)^{1/2}. \end{aligned}$$

Integrating this estimate over $t \in [t_k, t_{k+1}]$ gives

$$\begin{aligned} & \int_{t_k}^{t_{k+1}} \mathbb{E} [\|v_t(Z_t) - v_{t_k}(Z_{t_k})\|^2] dt \\ & \leq \frac{2h^3}{3} \left[\frac{3\sqrt{d(d+2)}}{\sigma^2} \mathbb{E} [\|Y_1 - Y_0\|^4] \right. \\ & \quad + \frac{6}{\sigma^2} \left(\mathbb{E} [\|Y_1 - Y_0\|^5] \right)^{4/5} (d(d+2)(d+4)(d+6)(d+8))^{1/5} \\ & \quad + \frac{14\tau^4 \sqrt{d(d+2)(d+4)(d+6)}}{\sigma^4} \left(\mathbb{E} [\|Y_1 - Y_0\|^4] \right)^{1/2} \left. \right] \\ & \quad + \frac{\tau^2 h^2 \sqrt{d(d+2)}}{\sigma^2} \left(\mathbb{E} [\|Y_1 - Y_0\|^4] \right)^{1/2}. \end{aligned}$$

Finally, summing over $k = 0, \dots, N-1$ and using $Nh = T$, we obtain

$$\begin{aligned} & \sum_{k=0}^{N-1} \int_{t_k}^{t_{k+1}} \mathbb{E} [\|v_t(Z_t) - v_{t_k}(Z_{t_k})\|^2] dt \\ & \leq \frac{2Th^2}{3} \left[\frac{3\sqrt{d(d+2)}}{\sigma^2} \mathbb{E} [\|Y_1 - Y_0\|^4] \right. \\ & \quad + \frac{6}{\sigma^2} \left(\mathbb{E} [\|Y_1 - Y_0\|^5] \right)^{4/5} (d(d+2)(d+4)(d+6)(d+8))^{1/5} \\ & \quad + \frac{14\tau^4 \sqrt{d(d+2)(d+4)(d+6)}}{\sigma^4} \left(\mathbb{E} [\|Y_1 - Y_0\|^4] \right)^{1/2} \left. \right] \\ & \quad + \frac{\tau^2 Th \sqrt{d(d+2)}}{\sigma^2} \left(\mathbb{E} [\|Y_1 - Y_0\|^4] \right)^{1/2}. \end{aligned}$$

This proves the lemma. □

Lemma 8 (Combined Score Increment Regularity). *Let Y_0 and Y_1 be random variables taking values in $\mathbb{R}^d$. Let $D = Y_1 - Y_0$. Assume that*

$$\mathbb{E} [\|Y_1 - Y_0\|^4] < \infty,$$

where the expectation is taken under the joint law of (Y_0, Y_1) .

For $t \in [0, T]$, define $S_t = Y_0 + tD$. Let $G \sim \mathcal{N}(0, I_d)$ be independent of (Y_0, D) , and define $X_t = S_t + \sigma G$ with $\sigma > 0$. Let q_t be the density of X_t . Define $s_t(x) = \nabla_x \log q_t(x)$ and

$$v_t(x) = \mathbb{E} [Y_1 - Y_0 \mid X_t = x].$$

Let

$$b_t(x) = v_t(x) + \frac{\tau^2}{2} s_t(x),$$

where $\tau > 0$. Let Z_t solve the exact SDE

$$dZ_t = b_t(Z_t)dt + \tau dB_t,$$

and assume that $Z_t \sim q_t$ for every $t \in [0, T]$.

Let $0 = t_0 < t_1 < \dots < t_N = T$ be a uniform grid with step size h , so $t_{k+1} - t_k = h$ and $Nh = T$. Then

$$\begin{aligned} & \sum_{k=0}^{N-1} \int_{t_k}^{t_{k+1}} \mathbb{E} [\|s_t(Z_t) - s_{t_k}(Z_{t_k})\|^2] dt \\ & \leq \frac{\tau^2 T h (2d(d+2) + 2d)}{\sigma^4} \\ & \quad + \frac{2Th^2}{3} \left[\frac{64d\sqrt{d(d+2)(d+4)(d+6)}}{\sigma^4} \left(\mathbb{E} [\|Y_1 - Y_0\|^4] \right)^{1/2} \right. \\ & \quad + \frac{4\sqrt{d(d+2)(d+4)(d+6)}}{\sigma^4} \left(\mathbb{E} [\|Y_1 - Y_0\|^4] \right)^{1/2} \\ & \quad + \frac{\tau^4 \sqrt{(8d(d+2)(d+4)(d+6) + 8d^2) d(d+2)}}{\sigma^6} \\ & \quad \left. + \frac{64\tau^4 d(d+2)(d+4)}{\sigma^6} \right]. \end{aligned}$$

The expectation on the left-hand side is taken under the law of the exact SDE path $(Z_t)_{0 \leq t \leq T}$.

Proof. We control the combined score increment directly. This is different from decomposing it into score-time and score-space pieces.

Fix $k \in \{0, \dots, N-1\}$ and $t \in [t_k, t_{k+1}]$. Applying Itô's formula to the process $s_r(Z_r)$ over $r \in [t_k, t]$ gives

$$\begin{aligned} s_t(Z_t) - s_{t_k}(Z_{t_k}) &= \int_{t_k}^t \left(\partial_r s_r(Z_r) + D_x s_r(Z_r) b_r(Z_r) + \frac{\tau^2}{2} \Delta_x s_r(Z_r) \right) dr \\ &\quad + \tau \int_{t_k}^t D_x s_r(Z_r) dB_r. \end{aligned}$$

Since $b_r(x) = v_r(x) + \frac{\tau^2}{2}s_r(x)$, the drift term inside the integral is

$$\begin{aligned} & \partial_r s_r(Z_r) + D_x s_r(Z_r) b_r(Z_r) + \frac{\tau^2}{2} \Delta_x s_r(Z_r) \\ &= \partial_r s_r(Z_r) + D_x s_r(Z_r) v_r(Z_r) + \frac{\tau^2}{2} D_x s_r(Z_r) s_r(Z_r) + \frac{\tau^2}{2} \Delta_x s_r(Z_r). \end{aligned}$$

Now use the continuity equation for the Gaussian-smoothed interpolation. Since q_r is the density of $X_r = S_r + \sigma G$, and $v_r(x) = \mathbb{E}[D \mid X_r = x]$, the curve $(q_r)_{r \in [0, T]}$ satisfies

$$\partial_r q_r(x) + \nabla_x \cdot (q_r(x) v_r(x)) = 0.$$

Dividing by $q_r(x)$ gives

$$\partial_r \log q_r(x) = -\nabla_x \cdot v_r(x) - v_r(x)^\top s_r(x).$$

Taking ∇_x on both sides gives

$$\partial_r \nabla_x \log q_r(x) = \partial_r s_r(x) = -\nabla_x (\nabla_x \cdot v_r(x)) - (D_x v_r(x))^\top s_r(x) - D_x s_r(x) v_r(x).$$

Substituting $x = Z_r$ gives

$$\partial_r s_r(Z_r) = -\nabla_x (\nabla_x \cdot v_r)(Z_r) - (D_x v_r(Z_r))^\top s_r(Z_r) - D_x s_r(Z_r) v_r(Z_r).$$

Therefore, the term $D_x s_r(Z_r) v_r(Z_r)$ cancels and we have

$$\begin{aligned} & \partial_r s_r(Z_r) + D_x s_r(Z_r) v_r(Z_r) + \frac{\tau^2}{2} D_x s_r(Z_r) s_r(Z_r) + \frac{\tau^2}{2} \Delta_x s_r(Z_r) \\ &= -\nabla_x (\nabla_x \cdot v_r)(Z_r) - (D_x v_r(Z_r))^\top s_r(Z_r) + \frac{\tau^2}{2} D_x s_r(Z_r) s_r(Z_r) + \frac{\tau^2}{2} \Delta_x s_r(Z_r). \end{aligned}$$

Hence

$$\begin{aligned} s_t(Z_t) - s_{t_k}(Z_{t_k}) &= \int_{t_k}^t \left[-\nabla_x (\nabla_x \cdot v_r)(Z_r) - (D_x v_r(Z_r))^\top s_r(Z_r) \right. \\ &\quad \left. + \frac{\tau^2}{2} D_x s_r(Z_r) s_r(Z_r) + \frac{\tau^2}{2} \Delta_x s_r(Z_r) \right] dr \\ &\quad + \tau \int_{t_k}^t D_x s_r(Z_r) dB_r. \end{aligned}$$

Using $\|u + w\|^2 \leq 2\|u\|^2 + 2\|w\|^2$, Jensen's inequality, and Itô isometry, we obtain

$$\begin{aligned} & \mathbb{E} \left[\|s_t(Z_t) - s_{t_k}(Z_{t_k})\|^2 \right] \\ & \leq 2(t - t_k) \int_{t_k}^t \mathbb{E} \left[\left\| -\nabla_x (\nabla_x \cdot v_r)(Z_r) - (D_x v_r(Z_r))^\top s_r(Z_r) \right. \right. \\ & \quad \left. \left. + \frac{\tau^2}{2} D_x s_r(Z_r) s_r(Z_r) + \frac{\tau^2}{2} \Delta_x s_r(Z_r) \right\|^2 \right] dr \\ & \quad + 2\tau^2 \int_{t_k}^t \mathbb{E} \left[\|D_x s_r(Z_r)\|_{\mathbb{F}}^2 \right] dr. \end{aligned}$$

Using $\|u_1 + u_2 + u_3 + u_4\|^2 \leq 4\|u_1\|^2 + 4\|u_2\|^2 + 4\|u_3\|^2 + 4\|u_4\|^2$, we have

$$\begin{aligned} & \mathbb{E} \left[\left\| -\nabla_x (\nabla_x \cdot v_r)(Z_r) - (D_x v_r(Z_r))^\top s_r(Z_r) + \frac{\tau^2}{2} D_x s_r(Z_r) s_r(Z_r) + \frac{\tau^2}{2} \Delta_x s_r(Z_r) \right\|^2 \right] \\ & \leq 4\mathbb{E} \left[\left\| \nabla_x (\nabla_x \cdot v_r)(Z_r) \right\|^2 \right] + 4\mathbb{E} \left[\left\| (D_x v_r(Z_r))^\top s_r(Z_r) \right\|^2 \right] \\ & \quad + \tau^4 \mathbb{E} \left[\left\| D_x s_r(Z_r) s_r(Z_r) \right\|^2 \right] + \tau^4 \mathbb{E} \left[\left\| \Delta_x s_r(Z_r) \right\|^2 \right]. \end{aligned}$$

We now bound each term.

First, by the posterior-velocity Hessian bound in Lemma 5,

$$\mathbb{E} \left[\left\| \nabla_x (\nabla_x \cdot v_r)(Z_r) \right\|^2 \right] \leq \frac{16d\sqrt{d(d+2)(d+4)(d+6)}}{\sigma^4} \left(\mathbb{E} \left[\|Y_1 - Y_0\|^4 \right] \right)^{1/2}.$$

Also,

$$\mathbb{E} \left[\left\| (D_x v_r(Z_r))^\top s_r(Z_r) \right\|^2 \right] \leq \frac{\sqrt{d(d+2)(d+4)(d+6)}}{\sigma^4} \left(\mathbb{E} \left[\|Y_1 - Y_0\|^4 \right] \right)^{1/2}.$$

Third,

$$\mathbb{E} \left[\left\| D_x s_r(Z_r) s_r(Z_r) \right\|^2 \right] \leq \frac{\sqrt{(8d(d+2)(d+4)(d+6) + 8d^2) d(d+2)}}{\sigma^6}.$$

Fourth, we use the already proved Gaussian score-geometry bound

$$\mathbb{E} \left[\left\| \Delta_x s_r(Z_r) \right\|^2 \right] \leq \frac{64d(d+2)(d+4)}{\sigma^6}.$$

Finally, using the already proved Gaussian score-geometry bound

$$\mathbb{E} \left[\left\| D_x s_r(Z_r) \right\|_F^2 \right] \leq \frac{2d(d+2) + 2d}{\sigma^4}.$$

Therefore, for every $r \in [0, T]$,

$$\begin{aligned} & \mathbb{E} \left[\left\| -\nabla_x (\nabla_x \cdot v_r)(Z_r) - (D_x v_r(Z_r))^\top s_r(Z_r) + \frac{\tau^2}{2} D_x s_r(Z_r) s_r(Z_r) + \frac{\tau^2}{2} \Delta_x s_r(Z_r) \right\|^2 \right] \\ & \leq \frac{64d\sqrt{d(d+2)(d+4)(d+6)}}{\sigma^4} \left(\mathbb{E} \left[\|Y_1 - Y_0\|^4 \right] \right)^{1/2} \\ & \quad + \frac{4\sqrt{d(d+2)(d+4)(d+6)}}{\sigma^4} \left(\mathbb{E} \left[\|Y_1 - Y_0\|^4 \right] \right)^{1/2} \\ & \quad + \frac{\tau^4 \sqrt{(8d(d+2)(d+4)(d+6) + 8d^2) d(d+2)}}{\sigma^6} \\ & \quad + \frac{64\tau^4 d(d+2)(d+4)}{\sigma^6}. \end{aligned}$$

Plugging this and the martingale bound into the Itô estimate gives, for every $t \in [t_k, t_{k+1}]$,

$$\mathbb{E} \left[\left\| s_t(Z_t) - s_{t_k}(Z_{t_k}) \right\|^2 \right]$$

$$\begin{aligned}
&\leq 2(t - t_k)^2 \left[\frac{64d\sqrt{d(d+2)(d+4)(d+6)}}{\sigma^4} \left(\mathbb{E} [\|Y_1 - Y_0\|^4] \right)^{1/2} \right. \\
&\quad + \frac{4\sqrt{d(d+2)(d+4)(d+6)}}{\sigma^4} \left(\mathbb{E} [\|Y_1 - Y_0\|^4] \right)^{1/2} \\
&\quad + \frac{\tau^4 \sqrt{(8d(d+2)(d+4)(d+6) + 8d^2) d(d+2)}}{\sigma^6} \\
&\quad \left. + \frac{64\tau^4 d(d+2)(d+4)}{\sigma^6} \right] \\
&\quad + 2\tau^2(t - t_k) \frac{2d(d+2) + 2d}{\sigma^4}.
\end{aligned}$$

Integrating over $t \in [t_k, t_{k+1}]$ gives

$$\begin{aligned}
&\int_{t_k}^{t_{k+1}} \mathbb{E} [\|s_t(Z_t) - s_{t_k}(Z_{t_k})\|^2] dt \\
&\leq \frac{2h^3}{3} \left[\frac{64d\sqrt{d(d+2)(d+4)(d+6)}}{\sigma^4} \left(\mathbb{E} [\|Y_1 - Y_0\|^4] \right)^{1/2} \right. \\
&\quad + \frac{4\sqrt{d(d+2)(d+4)(d+6)}}{\sigma^4} \left(\mathbb{E} [\|Y_1 - Y_0\|^4] \right)^{1/2} \\
&\quad + \frac{\tau^4 \sqrt{(8d(d+2)(d+4)(d+6) + 8d^2) d(d+2)}}{\sigma^6} \\
&\quad \left. + \frac{64\tau^4 d(d+2)(d+4)}{\sigma^6} \right] \\
&\quad + \tau^2 h^2 \frac{2d(d+2) + 2d}{\sigma^4}.
\end{aligned}$$

Summing over $k = 0, \dots, N-1$ and using $Nh = T$, we get

$$\begin{aligned}
&\sum_{k=0}^{N-1} \int_{t_k}^{t_{k+1}} \mathbb{E} [\|s_t(Z_t) - s_{t_k}(Z_{t_k})\|^2] dt \\
&\leq \frac{2Th^2}{3} \left[\frac{64d\sqrt{d(d+2)(d+4)(d+6)}}{\sigma^4} \left(\mathbb{E} [\|Y_1 - Y_0\|^4] \right)^{1/2} \right. \\
&\quad + \frac{4\sqrt{d(d+2)(d+4)(d+6)}}{\sigma^4} \left(\mathbb{E} [\|Y_1 - Y_0\|^4] \right)^{1/2} \\
&\quad + \frac{\tau^4 \sqrt{(8d(d+2)(d+4)(d+6) + 8d^2) d(d+2)}}{\sigma^6} \\
&\quad \left. + \frac{64\tau^4 d(d+2)(d+4)}{\sigma^6} \right] \\
&\quad + \frac{\tau^2 Th (2d(d+2) + 2d)}{\sigma^4}.
\end{aligned}$$

This proves the lemma. □

G.5 Completion of the KL Sampling Bound

Proof of Theorem 1. Set $T = t_N = 1$, and let Z be the exact weak solution with $\mathcal{L}(Z_t) = q_t$ supplied by Lemma 1. Let $\hat{Z}$ denote the continuous-time Euler interpolation, so that $\hat{Z}_{t_k} = \hat{Z}_k$.

Write

$$\begin{aligned} V_h &= \sum_{k=0}^{N-1} \int_{t_k}^{t_{k+1}} \mathbb{E} \|v_t(Z_t) - v_{t_k}(Z_{t_k})\|^2 dt, \\ S_h &= \sum_{k=0}^{N-1} \int_{t_k}^{t_{k+1}} \mathbb{E} \|s_t(Z_t) - s_{t_k}(Z_{t_k})\|^2 dt. \end{aligned}$$

We prove the theorem in four steps.

Step 1: reduction to V_h , S_h , and the estimation errors. For $t \in [t_k, t_{k+1})$, decompose the drift mismatch as

$$\begin{aligned} b_t(Z_t) - \hat{b}_{t_k}(Z_{t_k}) &= \{v_t(Z_t) - v_{t_k}(Z_{t_k})\} + \{v_{t_k}(Z_{t_k}) - \hat{v}_{t_k}(Z_{t_k})\} \\ &\quad + \frac{\tau^2}{2} \{s_t(Z_t) - s_{t_k}(Z_{t_k})\} + \frac{\tau^2}{2} \{s_{t_k}(Z_{t_k}) - \hat{s}_{t_k}(Z_{t_k})\}. \end{aligned}$$

Using the inequality $\|a_1 + a_2 + a_3 + a_4\|^2 \leq 4 \sum_{j=1}^4 \|a_j\|^2$, and Assumption 1 we obtain

$$\int_0^1 \mathbb{E} \|b_t(Z_t) - \hat{b}_{\eta_h(t)}(Z_{\eta_h(t)})\|^2 dt \leq 4V_h + 4\varepsilon_{\text{vel}}^2 + \tau^4 S_h + \tau^4 \varepsilon_{\text{score}}^2.$$

Steps 2 and 3 below show that $V_h, S_h < \infty$; hence the condition $\int_0^T \mathbb{E}_{\mathbb{P}} [\|\Delta_t(Z)\|^2] dt < \infty$ in Lemma 2 holds. The corresponding condition for $\hat{Z}$ is justified by Proposition 3. Applying that lemma and dividing by $2\tau^2$ then gives

$$\text{KL}(\mathcal{L}(Z_{t_N}) \parallel \mathcal{L}(\hat{Z}_N)) \leq \frac{2}{\tau^2} V_h + \frac{2}{\tau^2} \varepsilon_{\text{vel}}^2 + \frac{\tau^2}{2} S_h + \frac{\tau^2}{2} \varepsilon_{\text{score}}^2. \quad (7)$$

Step 2: explicit velocity-increment bound. Lyapunov's inequality yields

$$\mathbb{E} \|Y_1 - Y_0\|^4 \leq D_5^4, \quad \left(\mathbb{E} \|Y_1 - Y_0\|^4 \right)^{1/2} \leq D_5^2.$$

For $d \geq 1$,

$$\sqrt{d(d+2)} \leq d+1, \quad \{d(d+2)(d+4)(d+6)(d+8)\}^{1/5} \leq d+8,$$

and

$$\sqrt{d(d+2)(d+4)(d+6)} \leq (d+6)^2.$$

Applying these inequalities to Lemma 7 with $T = 1$ gives

$$V_h \leq \frac{2(9d+51)}{3} \frac{h^2 D_5^4}{\sigma^2} + \frac{28}{3} \frac{h^2 \tau^4 (d+6)^2 D_5^2}{\sigma^4} + \frac{h \tau^2 (d+1) D_5^2}{\sigma^2}. \quad (8)$$

Step 3: explicit score-increment bound. The two displacement-dependent terms in Lemma 8 satisfy

$$(64d+4) \sqrt{d(d+2)(d+4)(d+6)} \leq 68(d+6)^3.$$

For the two pure-score drift terms, letting $g_8 = d(d+2)(d+4)(d+6)$, we have $g_8 \geq d^2$ and hence

$$\begin{aligned} \sqrt{(8g_8 + 8d^2)d(d+2)} &\leq 4(d+6)^3, \\ 64d(d+2)(d+4) &\leq 64(d+6)^3. \end{aligned}$$

Finally, $2d(d+2) + 2d = 2d(d+3)$. Therefore

$$S_h \leq \frac{136}{3} \frac{h^2(d+6)^3 D_5^2}{\sigma^4} + \frac{2h\tau^2 d(d+3)}{\sigma^4} + \frac{136}{3} \frac{h^2\tau^4(d+6)^3}{\sigma^6}. \quad (9)$$

Step 4: collect constants and balance the leading terms. Substituting Equation (8) and Equation (9) into Equation (7) gives

$$\begin{aligned} \text{KL}(\mathcal{L}(Z_{t_N}) \parallel \mathcal{L}(\hat{Z}_N)) &\leq \frac{1}{\tau^2} \left\{ \frac{4(9d+51)}{3} \frac{h^2 D_5^4}{\sigma^2} + 2\varepsilon_{\text{vel}}^2 \right\} \\ &\quad + \tau^2 \left\{ \frac{56}{3} \frac{h^2(d+6)^2 D_5^2}{\sigma^4} + \frac{68}{3} \frac{h^2(d+6)^3 D_5^2}{\sigma^4} + \frac{1}{2} \varepsilon_{\text{score}}^2 \right\} \\ &\quad + \frac{h\tau^4 d(d+3)}{\sigma^4} + \frac{68}{3} \frac{h^2\tau^6(d+6)^3}{\sigma^6} + \frac{2h(d+1)D_5^2}{\sigma^2}. \end{aligned}$$

Because $(d+6)^2 \leq (d+6)^3$, the coefficient in braces on the second line is at most B_h . The first line is A_h/τ^2 , and the remaining three terms are exactly the last three terms on the right-hand side of Equation (2). This proves that bound.

Since $A_h, B_h > 0$, the choice $\tau = (A_h/B_h)^{1/4}$ satisfies

$$\frac{A_h}{\tau^2} + B_h\tau^2 = 2\sqrt{A_h B_h}.$$

Thus this choice balances the first two terms, as claimed. $\square$

G.6 Wasserstein Conversion to the Unsmoothed Target

The following two corollaries convert the terminal KL guarantee in Theorem 1 into W_1 bounds relative to the unsmoothed target. Corollary 1 applies to the unprojected Euler output, whereas Corollary 2 gives a sharper conversion for the terminally projected output when the target is supported on $[-1, 1]^d$.

Throughout this subsection, set $T = 1$ and condition on the auxiliary data, the EOT target reservoir, and the resulting coupling kernels.

By Lemma 10, the population-valid comparison coupling satisfies

$$\int Q(\cdot \mid y_1) \int \mathbb{P}_{Y|C=c, A=a} d\mathbb{P}_C(c)(y_1) = \hat{\mathbb{P}}_{Y|A=a}^{\text{aux}}.$$

Consequently, the exact interpolation associated with Q has $Y_0 \sim \hat{\mathbb{P}}_{Y|A=a}^{\text{aux}}$. Algorithm 1 is initialized from the same empirical source distribution. Therefore,

$$\mathcal{L}(Z_0) = \hat{\mathbb{P}}_{Y|A=a}^{\text{aux}} * N(0, \sigma^2 I_d) = \mathcal{L}(\hat{Z}_0).$$

Thus $\text{KL}(\mathcal{L}(Z_0) \parallel \mathcal{L}(\hat{Z}_0)) = 0$.

At terminal time, $\mathcal{L}(Z_{t_N}) = \mathcal{L}(Z_1) = \mathcal{L}(Y_1 + \sigma G) = \int \mathbb{P}_{Y|C=c, A=a} d\mathbb{P}_C(c) * N(0, \sigma^2 I_d)$. We treat the discrepancy between $\mathcal{L}(Y_1 + \sigma G)$ and $\int \mathbb{P}_{Y|C=c, A=a} d\mathbb{P}_C(c)$ solely as Gaussian-smoothing error. No stopping time $T < 1$ is introduced.

Corollary 1 (Wasserstein Error Relative to the Unsmoothed Target). *Assume the conditions of Theorem 1 and Assumption 2. Write the learned score as $\hat{s}_t(x) = \{\hat{u}_t(x) - x\}/\sigma^2$, and assume the network class bounds*

$$\sup_{x,k} \|\hat{v}_{t_k}(x)\| \leq 2\sqrt{d}, \quad \sup_{x,k} \|\hat{u}_{t_k}(x)\| \leq \sqrt{d}.$$

Assume also $h\tau^2/(2\sigma^2) \leq 1$. Then

$$\begin{aligned} W_1 \left(\int \mathbb{P}_{Y|C=c, A=a} d\mathbb{P}_C(c), \mathcal{L}(\widehat{Z}_N) \right) \\ \leq \sigma\sqrt{d} + 2^{1/4}\sqrt{d} \left[1 + \sigma^2 + \left(\sqrt{1 + \sigma^2} + 3 + \min\{\tau, \sqrt{2}\sigma\} \right)^2 \right]^{1/2} \cdot \text{KL} \left(\mathcal{L}(Y_1 + \sigma G) \parallel \mathcal{L}(\widehat{Z}_N) \right)^{1/4}. \end{aligned}$$

Proof. Let $Y_1 \sim \int \mathbb{P}_{Y|C=c, A=a} d\mathbb{P}_C(c)$ and $G \sim \mathcal{N}(0, I_d)$ be independent. Coupling Y_1 with $Y_1 + \sigma G$ gives

$$W_1 \left(\int \mathbb{P}_{Y|C=c, A=a} d\mathbb{P}_C(c), \mathcal{L}(Y_1 + \sigma G) \right) \leq \sigma \mathbb{E} \|G\| \leq \sigma\sqrt{d}.$$

Let $(X, \tilde{X})$ be a maximal coupling of $\mathcal{L}(Y_1 + \sigma G)$ and $\mathcal{L}(\widehat{Z}_N)$. Then

$$\mathbb{P}(X \neq \tilde{X}) = \text{TV} \left(\mathcal{L}(Y_1 + \sigma G), \mathcal{L}(\widehat{Z}_N) \right).$$

Therefore

$$\begin{aligned} W_1 \left(\mathcal{L}(Y_1 + \sigma G), \mathcal{L}(\widehat{Z}_N) \right) &\leq \mathbb{E} \|X - \tilde{X}\| = \mathbb{E} \left[\|X - \tilde{X}\| \mathbf{1}\{X \neq \tilde{X}\} \right] \\ &\leq \left(\mathbb{E} \|X - \tilde{X}\|^2 \right)^{1/2} \text{TV} \left(\mathcal{L}(Y_1 + \sigma G), \mathcal{L}(\widehat{Z}_N) \right)^{1/2} \\ &\leq \left(2\mathbb{E} \|X\|^2 + 2\mathbb{E} \|\tilde{X}\|^2 \right)^{1/2} \left[\frac{1}{2} \text{KL} \left(\mathcal{L}(Y_1 + \sigma G) \parallel \mathcal{L}(\widehat{Z}_N) \right) \right]^{1/4} \\ &= 2^{1/4} \left(\mathbb{E} \|X\|^2 + \mathbb{E} \|\tilde{X}\|^2 \right)^{1/2} \text{KL} \left(\mathcal{L}(Y_1 + \sigma G) \parallel \mathcal{L}(\widehat{Z}_N) \right)^{1/4}. \end{aligned}$$

Since $Y_1 \in [-1, 1]^d$,

$$\mathbb{E} \|X\|^2 = \mathbb{E} \|Y_1 + \sigma G\|^2 = \mathbb{E} \|Y_1\|^2 + \sigma^2 \mathbb{E} \|G\|^2 \leq d(1 + \sigma^2).$$

It remains to bound the second moment of $\widehat{Z}_N$. By the definition of the learned score class,

$$\widehat{s}_t(x) = \frac{\widehat{u}_t(x) - x}{\sigma^2},$$

the Euler recursion can be written as

$$a_h := 1 - \frac{h\tau^2}{2\sigma^2} \in [0, 1], \quad \xi_k := \widehat{v}_{t_k}(\widehat{Z}_k) + \frac{\tau^2}{2\sigma^2} \widehat{u}_{t_k}(\widehat{Z}_k).$$

$$\widehat{Z}_{k+1} = a_h \widehat{Z}_k + h\xi_k + \tau\sqrt{h} G_k.$$

Iterating this recursion gives

$$\begin{aligned} \widehat{Z}_N &= a_h^N \widehat{Z}_0 + h \sum_{k=0}^{N-1} a_h^{N-1-k} \xi_k \\ &\quad + \tau\sqrt{h} \sum_{k=0}^{N-1} a_h^{N-1-k} G_k. \end{aligned}$$

Applying Minkowski's inequality in L^2 , we obtain

$$\left(\mathbb{E}[\|\widehat{Z}_N\|^2]\right)^{1/2} \leq a_h^N \left(\mathbb{E}[\|\widehat{Z}_0\|^2]\right)^{1/2} + h \sum_{k=0}^{N-1} a_h^{N-1-k} \left(\mathbb{E}[\|\xi_k\|^2]\right)^{1/2} + \tau\sqrt{h} \left(\mathbb{E} \left[\left\| \sum_{k=0}^{N-1} a_h^{N-1-k} G_k \right\|^2 \right] \right)^{1/2}.$$

Since we have

$$\left\| \widehat{v}_{t_k}(\widehat{Z}_k) + \frac{\tau^2}{2\sigma^2} \widehat{u}_{t_k}(\widehat{Z}_k) \right\| \leq \left(2 + \frac{\tau^2}{2\sigma^2}\right) \sqrt{d}, \quad \mathbb{E} \left[\left\| \sum_{k=0}^{N-1} \left(1 - \frac{h\tau^2}{2\sigma^2}\right)^{N-1-k} G_k \right\|^2 \right] = d \sum_{k=0}^{N-1} \left(1 - \frac{h\tau^2}{2\sigma^2}\right)^{2k}$$

where G_k are independent, $\mathbb{E}[\|G_k\|^2] = d$, and all cross terms vanish. Therefore,

$$\left(\mathbb{E}[\|\widehat{Z}_N\|^2]\right)^{1/2} \leq \left(\mathbb{E}[\|\widehat{Z}_0\|^2]\right)^{1/2} + h \left(2 + \frac{\tau^2}{2\sigma^2}\right) \sqrt{d} \sum_{k=0}^{N-1} \left(1 - \frac{h\tau^2}{2\sigma^2}\right)^k + \tau \left[h d \sum_{k=0}^{N-1} \left(1 - \frac{h\tau^2}{2\sigma^2}\right)^{2k} \right]^{1/2}.$$

Since $Nh = 1$, $h \sum_{k=0}^{N-1} \left(1 - \frac{h\tau^2}{2\sigma^2}\right)^k \leq hN = 1$, so we have

$$h \sum_{k=0}^{N-1} \left(1 - \frac{h\tau^2}{2\sigma^2}\right)^k = h \frac{2\sigma^2}{h\tau^2} \left[1 - \left(1 - \frac{h\tau^2}{2\sigma^2}\right)^N \right] \leq \min \left\{ 1, \frac{2\sigma^2}{\tau^2} \right\},$$

and

$$h \sum_{k=0}^{N-1} \left(1 - \frac{h\tau^2}{2\sigma^2}\right)^{2k} \leq h \sum_{k=0}^{N-1} \left(1 - \frac{h\tau^2}{2\sigma^2}\right)^k \leq \min \left\{ 1, \frac{2\sigma^2}{\tau^2} \right\}.$$

Furthermore, observe that $\left(2 + \frac{\tau^2}{2\sigma^2}\right) \min \left\{ 1, \frac{2\sigma^2}{\tau^2} \right\} \leq 3$ and $\tau \left(\min \left\{ 1, \frac{2\sigma^2}{\tau^2} \right\} \right)^{1/2} = \min\{\tau, \sqrt{2}\sigma\}$. It follows that

$$\left(\mathbb{E}[\|\widehat{Z}_N\|^2]\right)^{1/2} \leq \left(\mathbb{E}[\|\widehat{Z}_0\|^2]\right)^{1/2} + 3\sqrt{d} + \sqrt{d} \min\{\tau, \sqrt{2}\sigma\}.$$

Since $\mathbb{E}[\|\widehat{Z}_0\|^2] = \mathbb{E}[\|\widehat{Y}_0\|^2] + \sigma^2 d \leq d(1 + \sigma^2)$, we obtain

$$\mathbb{E}[\|\widehat{Z}_N\|^2] \leq d \left(\sqrt{1 + \sigma^2} + 3 + \min\{\tau, \sqrt{2}\sigma\} \right)^2.$$

It follows that

$$\mathbb{E}\|\widehat{Z}_N\|^2 \leq d \left(\sqrt{1 + \sigma^2} + 3 + \min\{\tau, \sqrt{2}\sigma\} \right)^2.$$

Combining the preceding moment bounds yields

$$\begin{aligned} W_1 \left(\mathcal{L}(Y_1 + \sigma G), \mathcal{L}(\widehat{Z}_N) \right) &\leq 2^{1/4} \sqrt{d} \left[1 + \sigma^2 + \left(\sqrt{1 + \sigma^2} + 3 + \min\{\tau, \sqrt{2}\sigma\} \right)^2 \right]^{1/2} \\ &\quad \cdot \text{KL} \left(\mathcal{L}(Y_1 + \sigma G) \parallel \mathcal{L}(\widehat{Z}_N) \right)^{1/4}. \end{aligned}$$

The triangle inequality through $\mathcal{L}(Y_1 + \sigma G)$ proves the result. $\square$

Corollary 2 (Wasserstein Error After Output Clipping). *Let $Y_1 \sim \mu_1$ be supported on $K = [-1, 1]^d$, let $G \sim \mathcal{N}(0, I_d)$ be independent of Y_1 , and let $\tilde{Z}_N$ be any $\mathbb{R}^d$ -valued random vector. Define the Euclidean projection*

$$\Pi = \text{clip}_1 : \mathbb{R}^d \rightarrow K, \quad \tilde{Z}_N = \Pi(\hat{Z}_N).$$

Then

$$W_1(\mu_1, \mathcal{L}(\tilde{Z}_N)) \leq \sigma\sqrt{d} + \sqrt{2d \text{KL}(\mathcal{L}(Y_1 + \sigma G) \parallel \mathcal{L}(\hat{Z}_N))}. \quad (10)$$

Under the conditions of Theorem 1, its KL bound may be substituted into the right-hand side of Equation (10). Clipping is applied only after the final Euler step.

Proof. Write $P_\sigma = \mathcal{L}(Y_1 + \sigma G)$ and $\hat{P} = \mathcal{L}(\hat{Z}_N)$. Since $\Pi(Y_1) = Y_1$ and Euclidean projection onto K is nonexpansive,

$$W_1(\mu_1, \Pi_\# P_\sigma) \leq \mathbb{E} \|\Pi(Y_1) - \Pi(Y_1 + \sigma G)\| \leq \sigma \mathbb{E} \|G\| \leq \sigma\sqrt{d}.$$

Both $\Pi_\# P_\sigma$ and $\Pi_\# \hat{P}$ are supported on K , whose diameter is $2\sqrt{d}$. A maximal coupling, Pinsker's inequality, and data processing for KL divergence therefore give

$$\begin{aligned} W_1(\Pi_\# P_\sigma, \Pi_\# \hat{P}) &\leq 2\sqrt{d} \text{TV}(\Pi_\# P_\sigma, \Pi_\# \hat{P}) \\ &\leq \sqrt{2d \text{KL}(\Pi_\# P_\sigma \parallel \Pi_\# \hat{P})} \\ &\leq \sqrt{2d \text{KL}(P_\sigma \parallel \hat{P})}. \end{aligned}$$

The triangle inequality proves Equation (10). $\square$

This is a guarantee for the clipped output $\tilde{Z}_N$; the KL bound in Theorem 1 continues to concern the Euler output $\hat{Z}_N$ without clipping.

H Coupling Construction and Replacement

We first introduce some notations. For any distribution ζ and any conditional distribution $K(dy_0 \mid y_1)$, write

$$\Gamma_{\zeta, K}(dy_0, dy_1) = K(dy_0 \mid y_1) \zeta(dy_1).$$

This is the joint law obtained by drawing $Y_1 \sim \zeta$, and then drawing Y_0 from $K(\cdot \mid Y_1)$. On $\mathbb{R}^d \times \mathbb{R}^d$, define

$$\rho((y_0, y_1), (\tilde{y}_0, \tilde{y}_1)) = \|y_0 - \tilde{y}_0\| + \|y_1 - \tilde{y}_1\|,$$

and let $W_{1, \rho}$ denote the 1-Wasserstein distance induced by ρ .

We first construct $\hat{Q}_n$ using finitely many samples from the prediction oracle. Define the empirical auxiliary covariate distribution by $\hat{\mathbb{P}}_C := n^{-1} \sum_{i=n+1}^{2n} \delta_{C_i}$. Conditional on $\mathcal{D}_{\text{aux}}$, for each $i = n+1, \dots, 2n$, draw

$$\hat{Y}_{i,1}^Q, \dots, \hat{Y}_{i,M_Q}^Q \stackrel{\text{ind.}}{\sim} \hat{\mathbb{P}}_{Y \mid C=c, A=a} \text{ evaluated at } c = C_i.$$

We then define

$$\hat{\mathbb{P}}_{Y_1}^Q = \frac{1}{nM_Q} \sum_{i=n+1}^{2n} \sum_{r=1}^{M_Q} \delta_{\hat{Y}_{i,r}^Q},$$

which approximates $\int \hat{\mathbb{P}}_{Y|C=c,A=a} d\hat{\mathbb{P}}_C(c)$. To keep the coupling arguments readable, write

$$\mu_1 := \int \mathbb{P}_{Y|C=c,A=a} d\mathbb{P}_C(c), \quad \bar{\mu}_Q := \int \hat{\mathbb{P}}_{Y|C=c,A=a} d\hat{\mathbb{P}}_C(c), \quad \mu_Q^{\text{MC}} := \hat{\mathbb{P}}_{Y_1}^Q.$$

Thus μ_1 is the population counterfactual target, $\bar{\mu}_Q$ is the fitted target averaged over auxiliary covariates, and μ_Q^{MC} is its finite- M_Q empirical approximation.

In this section, we focus on the entropic optimal transport kernel between $\hat{\mathbb{P}}_{Y|A=a}^{\text{aux}}$ and μ_Q^{MC} :

$$\hat{\Gamma}_{\text{EOT}} \in \arg \min_{\Gamma \in \Pi(\hat{\mathbb{P}}_{Y|A=a}^{\text{aux}}, \mu_Q^{\text{MC}})} \left\{ \int \|y_1 - y_0\|^2 \Gamma(dy_0, dy_1) + \varepsilon_{\text{OT}} \text{KL} \left(\Gamma \parallel \hat{\mathbb{P}}_{Y|A=a}^{\text{aux}} \otimes \mu_Q^{\text{MC}} \right) \right\}.$$

When ε_{OT} is close to zero, the EOT objective is close to the unregularized optimal transport problem between $\hat{\mathbb{P}}_{Y|A=a}^{\text{aux}}$ and μ_Q^{MC} .

For a generic kernel, one would need to assume regularity of $y_1 \mapsto \hat{Q}_n(\cdot | y_1)$. This regularity is needed because we need to control the coupling replacement error. The proof evaluates $\hat{Q}_n$ not only at samples from μ_Q^{MC} , where the EOT marginal identity $\int \hat{Q}_n(\cdot | y_1) \mu_Q^{\text{MC}}(dy_1) = \hat{\mathbb{P}}_{Y|A=a}^{\text{aux}}$ holds, but also at y_1 's from the population target. For the EOT kernel, this needed regularity can be proved directly: Lemma 9 shows that $y_1 \mapsto \hat{Q}_n(\cdot | y_1)$ is Lipschitz in W_1 , with constant proportional to $1/\varepsilon_{\text{OT}}$. We then establish a coupling replacement error bound, see Lemma 10 and Corollary 3.

H.1 The Learned EOT Kernel

The entropic optimal transport problem defining $\hat{\Gamma}_{\text{EOT}}$ has a unique optimizer. By (Genevay et al., 2016, Proposition 2.1), there exist $\hat{\varphi}_n$ and $\hat{\psi}_n$ such that

$$d\hat{\Gamma}_{\text{EOT}}(y_0, y_1) = \exp \left(\frac{\hat{\varphi}_n(y_0) + \hat{\psi}_n(y_1) - \|y_0 - y_1\|^2}{\varepsilon_{\text{OT}}} \right) d\hat{\mathbb{P}}_{Y|A=a}^{\text{aux}}(y_0) d\mu_Q^{\text{MC}}(y_1)$$

for almost every (y_0, y_1) . Consequently, for μ_Q^{MC} -almost every y , the conditional distribution of Y_0 given $Y_1 = y$ under $\hat{\Gamma}_{\text{EOT}}$ is

$$\hat{Q}_n(dy_0 | y) = \frac{\exp \left(\frac{\hat{\varphi}_n(y_0) - \|y_0 - y\|^2}{\varepsilon_{\text{OT}}} \right)}{\int \exp \left(\frac{\hat{\varphi}_n(z) - \|z - y\|^2}{\varepsilon_{\text{OT}}} \right) \hat{\mathbb{P}}_{Y|A=a}^{\text{aux}}(dz)} \hat{\mathbb{P}}_{Y|A=a}^{\text{aux}}(dy_0).$$

Motivated by the out-of-sample entropic extension discussed before Equation (13) of Pooladian and Niles-Weed (2021), we use the displayed Gibbs formula to define $\hat{Q}_n(\cdot | y)$ for every $y \in \mathbb{R}^d$. We remark that this is an extension of the conditional kernel. The denominator is finite and strictly positive, so the right-hand side is a probability distribution for every y . It holds that

$$\hat{\Gamma}_{\text{EOT}}(dy_0, dy_1) = \hat{Q}_n(dy_0 | y_1) \mu_Q^{\text{MC}}(dy_1).$$

Moreover, under Assumption 2, we have $\hat{Q}_n([-1, 1]^d | y) = 1, \forall y \in \mathbb{R}^d$. The same kernel can be evaluated at target queries drawn from the training-fold fitted distribution or from the population target distribution. The resulting joint distribution remains well defined, but its Y_0 -marginal need not equal $\hat{\mathbb{P}}_{Y|A=a}^{\text{aux}}$, because the exact source-marginal identity holds only when $\hat{Q}_n$ is integrated against μ_Q^{MC} .

By construction, we have the marginal identity

$$\int \widehat{Q}_n(\cdot | y_1) \mu_Q^{\text{MC}}(dy_1) = \widehat{\mathbb{P}}_{Y|A=a}^{\text{aux}}.$$

This identity is not an identity under $\int \widehat{\mathbb{P}}_{Y|C=c,A=a} d\widehat{\mathbb{P}}_C(c)$, nor under the training-fold fitted target distribution $\frac{1}{n} \sum_{i=1}^n \widehat{\mathbb{P}}_{Y|C=C_i,A=a}$.

Lemma 9 (Stability of the Finite-Sample Entropic Coupling). *Assume Assumption 2. Then, for every $y_1, y'_1 \in \mathbb{R}^d$,*

$$W_1 \left(\widehat{Q}_n(\cdot | y_1), \widehat{Q}_n(\cdot | y'_1) \right) \leq \frac{2d}{\varepsilon_{\text{OT}}} \|y_1 - y'_1\|.$$

Proof. Recall that

$$\widehat{Q}_n(dy_0 | y) = \frac{\exp \left(\frac{\widehat{\varphi}_n(y_0) - \|y_0 - y\|^2}{\varepsilon_{\text{OT}}} \right)}{\int \exp \left(\frac{\widehat{\varphi}_n(z) - \|z - y\|^2}{\varepsilon_{\text{OT}}} \right) \widehat{\mathbb{P}}_{Y|A=a}^{\text{aux}}(dz)} \widehat{\mathbb{P}}_{Y|A=a}^{\text{aux}}(dy_0).$$

Fix $y_1, y'_1 \in \mathbb{R}^d$ and a 1-Lipschitz function $f : \mathbb{R}^d \rightarrow \mathbb{R}$. For $t \in [0, 1]$, substituting $y = (1-t)y_1 + ty'_1$ into the preceding display gives

$$\int f(y_0) \widehat{Q}_n(dy_0 | (1-t)y_1 + ty'_1) = \frac{\int f(y_0) \exp \left(\frac{\widehat{\varphi}_n(y_0) - \|y_0 - ((1-t)y_1 + ty'_1)\|^2}{\varepsilon_{\text{OT}}} \right) \widehat{\mathbb{P}}_{Y|A=a}^{\text{aux}}(dy_0)}{\int \exp \left(\frac{\widehat{\varphi}_n(z) - \|z - ((1-t)y_1 + ty'_1)\|^2}{\varepsilon_{\text{OT}}} \right) \widehat{\mathbb{P}}_{Y|A=a}^{\text{aux}}(dz)}.$$

For every $y_0 \in [-1, 1]^d$,

$$\frac{d}{dt} \left[-\frac{\|y_0 - ((1-t)y_1 + ty'_1)\|^2}{\varepsilon_{\text{OT}}} \right] = \frac{2}{\varepsilon_{\text{OT}}} \langle y_0 - ((1-t)y_1 + ty'_1), y'_1 - y_1 \rangle.$$

Applying the quotient rule to the preceding ratio yields

$$\begin{aligned} & \frac{d}{dt} \int f(y_0) \widehat{Q}_n(dy_0 | (1-t)y_1 + ty'_1) \\ &= \frac{2}{\varepsilon_{\text{OT}}} \int f(y_0) \langle y_0 - ((1-t)y_1 + ty'_1), y'_1 - y_1 \rangle \widehat{Q}_n(dy_0 | (1-t)y_1 + ty'_1) \\ & \quad - \frac{2}{\varepsilon_{\text{OT}}} \int f(y_0) \widehat{Q}_n(dy_0 | (1-t)y_1 + ty'_1) \int \langle z - ((1-t)y_1 + ty'_1), y'_1 - y_1 \rangle \widehat{Q}_n(dz | (1-t)y_1 + ty'_1) \\ &= \frac{2}{\varepsilon_{\text{OT}}} \text{Cov}_{Y_0 \sim \widehat{Q}_n(\cdot | (1-t)y_1 + ty'_1)} (f(Y_0), \langle Y_0, y'_1 - y_1 \rangle). \end{aligned}$$

Assumption 2 implies that $\widehat{Q}_n(\cdot | (1-t)y_1 + ty'_1)$ is supported on $[-1, 1]^d$. Since f is 1-Lipschitz,

$$\sup_{y_0, z \in [-1, 1]^d} |f(y_0) - f(z)| \leq 2\sqrt{d}, \quad \text{and} \quad \sup_{y_0, z \in [-1, 1]^d} |\langle y_0 - z, y'_1 - y_1 \rangle| \leq 2\sqrt{d} \|y'_1 - y_1\|.$$

Therefore, Popoviciu's variance inequality gives

$$\text{Var}_{Y_0 \sim \widehat{Q}_n(\cdot | (1-t)y_1 + ty'_1)} (f(Y_0)) \leq d,$$

$$\text{Var}_{Y_0 \sim \hat{Q}_n(\cdot | (1-t)y_1 + ty'_1)} (\langle Y_0, y'_1 - y_1 \rangle) \leq d \|y'_1 - y_1\|^2.$$

By Cauchy–Schwarz,

$$\left| \text{Cov}_{Y_0 \sim \hat{Q}_n(\cdot | (1-t)y_1 + ty'_1)} (f(Y_0), \langle Y_0, y'_1 - y_1 \rangle) \right| \leq d \|y'_1 - y_1\|.$$

Hence, for every $t \in [0, 1]$,

$$\left| \frac{d}{dt} \int f(y_0) \hat{Q}_n(dy_0 | (1-t)y_1 + ty'_1) \right| \leq \frac{2d}{\varepsilon_{\text{OT}}} \|y'_1 - y_1\|.$$

Integrating over $t \in [0, 1]$ yields

$$\left| \int f(y_0) \hat{Q}_n(dy_0 | y'_1) - \int f(y_0) \hat{Q}_n(dy_0 | y_1) \right| \leq \frac{2d}{\varepsilon_{\text{OT}}} \|y'_1 - y_1\|.$$

Taking the supremum over all 1-Lipschitz f and applying Kantorovich–Rubinstein duality gives

$$W_1(\hat{Q}_n(\cdot | y_1), \hat{Q}_n(\cdot | y'_1)) \leq \frac{2d}{\varepsilon_{\text{OT}}} \|y_1 - y'_1\|.$$

□

Proposition 4 (Target-Query Discrepancy Under Empirical C). *Assume Assumptions 2 and 5. Then*

$$\begin{aligned} & W_1 \left(\int \mathbb{P}_{Y|C=c, A=a} d\mathbb{P}_C(c), \frac{1}{n} \sum_{i=n+1}^{2n} \hat{\mathbb{P}}_{Y|C=C_i, A=a} \right) \\ & \leq W_1 \left(\int \mathbb{P}_{Y|C=c, A=a} d\mathbb{P}_C(c), \frac{1}{n} \sum_{i=n+1}^{2n} \mathbb{P}_{Y|C=C_i, A=a} \right) + \varepsilon_{\text{pred, train}}. \end{aligned}$$

Consequently, on the event of Lemma 12,

$$W_1 \left(\int \mathbb{P}_{Y|C=c, A=a} d\mathbb{P}_C(c), \frac{1}{n} \sum_{i=n+1}^{2n} \hat{\mathbb{P}}_{Y|C=C_i, A=a} \right) \lesssim n^{-1/(d+2)} + \varepsilon_{\text{pred, train}}.$$

Proof. The triangle inequality gives

$$\begin{aligned} & W_1 \left(\int \mathbb{P}_{Y|C=c, A=a} d\mathbb{P}_C(c), \frac{1}{n} \sum_{i=n+1}^{2n} \hat{\mathbb{P}}_{Y|C=C_i, A=a} \right) \\ & \leq W_1 \left(\int \mathbb{P}_{Y|C=c, A=a} d\mathbb{P}_C(c), \frac{1}{n} \sum_{i=n+1}^{2n} \mathbb{P}_{Y|C=C_i, A=a} \right) \\ & \quad + W_1 \left(\frac{1}{n} \sum_{i=n+1}^{2n} \mathbb{P}_{Y|C=C_i, A=a}, \frac{1}{n} \sum_{i=n+1}^{2n} \hat{\mathbb{P}}_{Y|C=C_i, A=a} \right). \end{aligned}$$

For each $i = n+1, \dots, 2n$, take an optimal coupling of $\mathbb{P}_{Y|C=C_i, A=a}$ and $\hat{\mathbb{P}}_{Y|C=C_i, A=a}$, and mix these couplings uniformly over i . This gives

$$W_1 \left(\frac{1}{n} \sum_{i=n+1}^{2n} \mathbb{P}_{Y|C=C_i, A=a}, \frac{1}{n} \sum_{i=n+1}^{2n} \hat{\mathbb{P}}_{Y|C=C_i, A=a} \right)$$

$$\begin{aligned}
&\leq \frac{1}{n} \sum_{i=n+1}^{2n} W_1 \left(\mathbb{P}_{Y|C=C_i, A=a}, \widehat{\mathbb{P}}_{Y|C=C_i, A=a} \right) \\
&\leq \left[\frac{1}{n} \sum_{i=n+1}^{2n} W_1 \left(\mathbb{P}_{Y|C=C_i, A=a}, \widehat{\mathbb{P}}_{Y|C=C_i, A=a} \right)^2 \right]^{1/2} \leq \varepsilon_{\text{pred, train}}.
\end{aligned}$$

Combining this bound with Lemma 12 proves the claim. $\square$

Conditional on $\mathcal{D}_{\text{aux}}$, $\mathcal{D}_{\text{tr}}$, and $\{\widehat{Y}_{i,r}^Q\}$, the offline final-flow Monte Carlo array is drawn independently of the coupling reservoir:

$$\widehat{Y}_{i,1}, \dots, \widehat{Y}_{i,M} \stackrel{\text{ind.}}{\sim} \widehat{\mathbb{P}}_{Y|C=c, A=a} \text{ evaluated at } c = C_i, \quad i = 1, \dots, n. \quad (11)$$

For each i and k , draw independently

$$\widehat{Y}_{0,i,k} \sim \widehat{Q}_n(\cdot | \widehat{Y}_{i,k}), \quad K_{i,k} \sim \text{Unif}\{0, \dots, N-1\}, \quad G_{i,k} \sim \mathcal{N}(0, I_d), \quad (12)$$

and set $t_{i,k} = t_{K_{i,k}}$. For every i with $A_i = a$, draw independently

$$\widehat{Y}_{0,i}^{\text{obs}} \sim \widehat{Q}_n(\cdot | Y_i), \quad K_i^{\text{obs}} \sim \text{Unif}\{0, \dots, N-1\}, \quad G_i^{\text{obs}} \sim \mathcal{N}(0, I_d), \quad (13)$$

and set $t_i^{\text{obs}} = t_{K_i^{\text{obs}}}$. These variables are drawn once before minimizing the final empirical objective and are held fixed during optimization. The M -sample empirical-process error for this fixed objective is separate from the M_Q -sample coupling-reservoir error above.

Assumption 7 (Nonempty Treatment- a Auxiliary Sample). *The auxiliary split contains at least one observation with treatment $A = a$:*

$$\sum_{i=n+1}^{2n} \mathbf{1}\{A_i = a\} \geq 1.$$

Under the overlap condition $\mathbb{P}(A = a | C = c) \geq \pi_{\min}$, the event in this assumption has probability at least

$$1 - (1 - \pi_{\min})^n \geq 1 - e^{-n\pi_{\min}}.$$

Thus the assumption may equivalently be replaced by this additional exponentially small failure probability.

Lemma 10 (Existence of an Ideal Coupling Q With the Correct Population Marginals). *Under Assumptions 2 and 7, let λ^* be an optimal coupling of μ_1 and μ_Q^{MC} , and write $\lambda^*(d\widehat{y}_1 | y_1)$ for the conditional law of its second coordinate given the first. Define*

$$Q(dy_0 | y_1) := \int \widehat{Q}_n(dy_0 | \widehat{y}_1) \lambda^*(d\widehat{y}_1 | y_1).$$

Then Q satisfies the marginal identity

$$\int Q(\cdot | y_1) \mu_1(dy_1) = \widehat{\mathbb{P}}_{Y|A=a}^{\text{aux}}$$

and the coupling-replacement bound

$$W_{1,\rho}(\Gamma_{\mu_1, Q}, \Gamma_{\mu_1, \widehat{Q}_n}) \leq \frac{2d}{\varepsilon_{\text{OT}}} W_1(\mu_1, \mu_Q^{\text{MC}}).$$

Proof. Under the construction in the statement, draw $(Y_1, \hat{Y}_1) \sim \lambda^*$, followed by

$$Y_0 \sim \hat{Q}_n(\cdot \mid \hat{Y}_1).$$

The EOT marginal identity implies

$$Y_0 \sim \int \hat{Q}_n(\cdot \mid y) \mu_Q^{\text{MC}}(dy) = \hat{\mathbb{P}}_{Y|A=a}^{\text{aux}}.$$

Thus Q has the displayed marginal identity.

Conditional on $(Y_1, \hat{Y}_1)$, couple $\tilde{Y}_0 \sim \hat{Q}_n(\cdot \mid Y_1)$ and $Y_0 \sim \hat{Q}_n(\cdot \mid \hat{Y}_1)$ optimally. Then $((Y_0, Y_1), (\tilde{Y}_0, Y_1))$ couples $\Gamma_{\mu_1, Q}$ and $\Gamma_{\mu_1, \hat{Q}_n}$, so

$$\begin{aligned} W_{1,\rho}(\Gamma_{\mu_1, Q}, \Gamma_{\mu_1, \hat{Q}_n}) &\leq \mathbb{E} \|Y_0 - \tilde{Y}_0\| \\ &= \mathbb{E} W_1(\hat{Q}_n(\cdot \mid \hat{Y}_1), \hat{Q}_n(\cdot \mid Y_1)) \\ &\leq \frac{2d}{\varepsilon_{\text{OT}}} \mathbb{E} \|\hat{Y}_1 - Y_1\| \\ &= \frac{2d}{\varepsilon_{\text{OT}}} W_1(\mu_1, \mu_Q^{\text{MC}}), \end{aligned}$$

where the penultimate line uses Lemma 9. $\square$

Remark 2. Under Assumption 2, the coupling constructed in Lemma 10 is supported on $[-1, 1]^d \times [-1, 1]^d$. Hence its conditional kernel $Q(dy_0 \mid y_1)$ may be chosen so that

$$Q([-1, 1]^d \mid y_1) = 1, \quad y_1 \in [-1, 1]^d.$$

Corollary 3 (Coupling-Replacement Error). Under Assumptions 2, 5, and 7, with probability at least $1 - \frac{2}{n}$,

$$W_{1,\rho}(\Gamma_{\mu_1, Q}, \Gamma_{\mu_1, \hat{Q}_n}) \lesssim \frac{1}{\varepsilon_{\text{OT}}} \left(n^{-1/(d+2)} + \varepsilon_{\text{pred,train}} + (nM_Q)^{-1/(d+2)} \right),$$

where d is treated as fixed.

Proof. Apply triangle inequality and Lemma 10,

$$\begin{aligned} W_{1,\rho}(\Gamma_{\mu_1, Q}, \Gamma_{\mu_1, \hat{Q}_n}) &\leq \frac{2d}{\varepsilon_{\text{OT}}} W_1(\mu_1, \mu_Q^{\text{MC}}) \\ &\leq \frac{2d}{\varepsilon_{\text{OT}}} \left\{ W_1(\mu_1, \bar{\mu}_Q) + W_1(\bar{\mu}_Q, \mu_Q^{\text{MC}}) \right\}. \end{aligned}$$

By Proposition 4, with probability at least $1 - \frac{1}{n}$,

$$W_1(\mu_1, \bar{\mu}_Q) \lesssim n^{-1/(d+2)} + \varepsilon_{\text{pred,train}}.$$

Taking $\delta_Q = 1/n$ in Lemma 11 gives, with probability at least $1 - \frac{1}{n}$,

$$W_1(\bar{\mu}_Q, \mu_Q^{\text{MC}}) \lesssim (nM_Q)^{-1/(d+2)},$$

For fixed d , substituting the preceding two bounds gives

$$W_{1,\rho}(\Gamma_{\mu_1, Q}, \Gamma_{\mu_1, \hat{Q}_n}) \lesssim \frac{1}{\varepsilon_{\text{OT}}} \left(n^{-1/(d+2)} + \varepsilon_{\text{pred,train}} + (nM_Q)^{-1/(d+2)} \right).$$

A union bound over the two events gives probability at least $1 - \frac{2}{n}$. $\square$

H.2 Finite-Sample Target Approximation

In this subsection, we establish two high-probability W_1 approximation bounds used in the coupling analysis. Bounds of this type are standard in empirical Wasserstein theory; see, for example, [Weed and Bach \(2019\)](#). We provide a proof for completeness.

Lemma 11 (Finite- M_Q Approximation of the Fitted Target Distribution). *Assume Assumption 2. Conditional on $\mathcal{D}_{\text{aux}}$, for every $\delta_Q \in (0, 1)$, with probability at least $1 - \delta_Q$,*

$$W_1(\bar{\mu}_Q, \mu_Q^{\text{MC}}) \leq \sqrt{d} \left[\left(2 + 2^{d/2}\right) (nM_Q)^{-1/(d+2)} + \sqrt{\frac{2 \log(1/\delta_Q)}{nM_Q}} \right].$$

In particular, for fixed d and $\delta_Q = 1/n$, the right-hand side is $O_d((nM_Q)^{-1/(d+2)})$.

Proof. Set $m = nM_Q$ and $K = \lceil m^{1/(d+2)} \rceil$. Partition $[-1, 1]^d$ into K^d cubes $A_1, \dots, A_{K^d}$ of side length $2/K$, and let a_j be the center of A_j . If $\zeta^\#$ denotes the law obtained by moving the mass in each A_j to a_j , then

$$W_1(\zeta, \zeta^\#) \leq \frac{\sqrt{d}}{K}$$

for every ζ supported on the cube.

Define the conditional cell probabilities and their empirical counterparts by

$$\begin{aligned} p_j &:= \bar{\mu}_Q(A_j) = \frac{1}{n} \sum_{i=n+1}^{2n} \hat{\mathbb{P}}_{Y|C=C^i, A=a}(A_j), \\ \hat{p}_j &:= \mu_Q^{\text{MC}}(A_j) = \frac{1}{m} \sum_{i=n+1}^{2n} \sum_{r=1}^{M_Q} \mathbf{1}\{\hat{Y}_{i,r}^Q \in A_j\}. \end{aligned}$$

The discretized measures have masses p_j and $\hat{p}_j$ at a_j . Since the cube has diameter $2\sqrt{d}$,

$$W_1(\bar{\mu}_Q^\#, \mu_Q^{\text{MC}\#}) \leq 2\sqrt{d} \text{TV}(\bar{\mu}_Q^\#, \mu_Q^{\text{MC}\#}) = \sqrt{d} \sum_{j=1}^{K^d} |p_j - \hat{p}_j|.$$

Consequently, with $F := \sum_{j=1}^{K^d} |p_j - \hat{p}_j|$,

$$W_1(\bar{\mu}_Q, \mu_Q^{\text{MC}}) \leq \frac{2\sqrt{d}}{K} + \sqrt{d} F. \tag{14}$$

We next control F conditionally on $\mathcal{D}_{\text{aux}}$. Replacing one of the m independent draws moves empirical mass $1/m$ from at most one cell to at most one other cell. Hence the empirical histogram changes by at most $2/m$ in ℓ_1 , and therefore F changes by at most $2/m$. McDiarmid's inequality gives

$$\mathbb{P}(F - \mathbb{E}[F | \mathcal{D}_{\text{aux}}] \geq t | \mathcal{D}_{\text{aux}}) \leq \exp\left(-\frac{mt^2}{2}\right).$$

Thus, with conditional probability at least $1 - \delta_Q$,

$$F \leq \mathbb{E}[F | \mathcal{D}_{\text{aux}}] + \sqrt{\frac{2 \log(1/\delta_Q)}{m}}. \tag{15}$$

It remains to bound the conditional mean. Cauchy–Schwarz followed by Jensen gives

$$\mathbb{E}[F \mid \mathcal{D}_{\text{aux}}] \leq \sqrt{K^d} \left[\sum_{j=1}^{K^d} \mathbb{E}[(\hat{p}_j - p_j)^2 \mid \mathcal{D}_{\text{aux}}] \right]^{1/2}.$$

For $p_{i,j} := \widehat{\mathbb{P}}_{Y|C=C^i, A=a}(A_j)$, conditional independence and unbiasedness imply

$$\begin{aligned} \sum_{j=1}^{K^d} \mathbb{E}[(\hat{p}_j - p_j)^2 \mid \mathcal{D}_{\text{aux}}] &= \frac{1}{m^2} \sum_{i=n+1}^{2n} \sum_{r=1}^{M_Q} \sum_{j=1}^{K^d} p_{i,j}(1 - p_{i,j}) \\ &\leq \frac{1}{m^2} \sum_{i=n+1}^{2n} \sum_{r=1}^{M_Q} \sum_{j=1}^{K^d} p_{i,j} = \frac{1}{m}. \end{aligned}$$

Therefore

$$\mathbb{E}[F \mid \mathcal{D}_{\text{aux}}] \leq \sqrt{\frac{K^d}{m}}.$$

Combining this inequality with Equation (14) and Equation (15) yields

$$W_1(\bar{\mu}_Q, \mu_Q^{\text{MC}}) \leq \sqrt{d} \left[\frac{2}{K} + \sqrt{\frac{K^d}{m}} + \sqrt{\frac{2 \log(1/\delta_Q)}{m}} \right].$$

Finally, $m \geq 1$ implies

$$\frac{1}{K} \leq m^{-1/(d+2)} \quad \text{and} \quad K \leq 2m^{1/(d+2)},$$

so $\sqrt{K^d/m} \leq 2^{d/2} m^{-1/(d+2)}$. Substituting $m = nM_Q$ proves the stated bound. $\square$

Lemma 12 (Empirical Covariate Standardization Error). *Assume Assumption 2 and $n \geq 2$. Define*

$$\mu_{1,n} := \frac{1}{n} \sum_{i=n+1}^{2n} \mathbb{P}_{Y|C=C^i, A=a}, \quad K_n := \lceil n^{1/(d+2)} \rceil.$$

Then, with probability at least $1 - 1/n$ over $C^{n+1}, \dots, C^{2n}$,

$$W_1(\mu_1, \mu_{1,n}) \leq \frac{4\sqrt{d}}{K_n} + 2\sqrt{d} \left(\frac{K_n^d}{n} \right)^{1/2} + 2\sqrt{d} \left(\frac{2 \log n}{n} \right)^{1/2}.$$

Consequently, for fixed d , this error is $O_d(n^{-1/(d+2)})$.

Proof. Partition $[-1, 1]^d$ into K_n^d cubes $A_1, \dots, A_{K_n^d}$ of side length $2/K_n$. Moving every point to the center of its cell changes any law supported on the cube by at most $2\sqrt{d}/K_n$ in W_1 .

For each cell, define

$$q_j(c) := \mathbb{P}(Y \in A_j \mid C = c, A = a), \quad p_j := \mathbb{E}_C[q_j(C)], \quad \hat{p}_j := \frac{1}{n} \sum_{i=n+1}^{2n} q_j(C^i).$$

The discretizations of μ_1 and $\mu_{1,n}$ put masses p_j and $\hat{p}_j$ on the same cell centers. Therefore, with

$$F := \sum_{j=1}^{K_n^d} |p_j - \hat{p}_j|,$$

the triangle inequality and the diameter bound give

$$W_1(\mu_1, \mu_{1,n}) \leq \frac{4\sqrt{d}}{K_n} + 2\sqrt{d} F. \quad (16)$$

Changing one covariate C^i changes the probability vector $(q_1(C^i), \dots, q_{K_n^d}(C^i))$ to another probability vector. Its ℓ_1 -change is at most 2, so F changes by at most $2/n$. McDiarmid's inequality and $\delta = 1/n$ yield, with probability at least $1 - 1/n$,

$$F \leq \mathbb{E}F + \sqrt{\frac{2 \log n}{n}}. \quad (17)$$

For the expectation, Cauchy–Schwarz and Jensen imply

$$\mathbb{E}F \leq \sqrt{K_n^d} \left[\sum_{j=1}^{K_n^d} \text{Var}(\hat{p}_j) \right]^{1/2}.$$

Because the covariates are i.i.d.,

$$\sum_{j=1}^{K_n^d} \text{Var}(\hat{p}_j) = \frac{1}{n} \sum_{j=1}^{K_n^d} \text{Var}(q_j(C)) \leq \frac{1}{n} \mathbb{E}_C \left[\sum_{j=1}^{K_n^d} q_j(C)^2 \right] \leq \frac{1}{n},$$

where the last inequality uses $0 \leq q_j(C) \leq 1$ and $\sum_j q_j(C) = 1$. Hence

$$\mathbb{E}F \leq \sqrt{\frac{K_n^d}{n}}. \quad (18)$$

Combining Equation (16), Equation (17), and Equation (18) proves the explicit bound.

Finally, $K_n \leq 2n^{1/(d+2)}$, and for fixed $d \geq 1$, $\sqrt{\log n/n} = O_d(n^{-1/(d+2)})$. This proves the stated order. $\square$

I Proof of Proposition 1

In this section, we work on $\mathcal{E}_a$, conditioning on the auxiliary split, coupling reservoir, and fitted nuisance objects. We use the kernels $Q, \hat{Q}_n$ and the offline sampling scheme from Equation (11), Equation (12) and Equation (13) in Appendix H. Population expectations below are taken over fresh observations or interpolation variables, with the fitted objects fixed.

To align the training analysis with the discrete-time sampling analysis, we evaluate the loss on the same time grid. Let $0 = t_0 < t_1 < \dots < t_N = 1$, with $h = 1/N$, and let $K \sim \text{Unif}\{0, \dots, N-1\}$.

Let Θ be a parameter class with metric d_Θ and finite proper covering numbers $\mathcal{N}(r, \Theta, d_\Theta)$; every cover center belongs to Θ . The class and loss family are fixed under the preceding conditioning. For $\theta \in \Theta$, let $\Phi_\theta(y_0, y_1, k, g)$ be a measurable nonnegative loss on $[-1, 1]^d \times [-1, 1]^d \times \{0, \dots, N-1\} \times \mathbb{R}^d$. With $K \sim \text{Unif}\{0, \dots, N-1\}$ and $G \sim \mathcal{N}(0, I_d)$ mutually independent and independent of the endpoints, define

$$\begin{aligned} H_\theta(y_0, y_1) &= \mathbb{E}_{K,G}[\Phi_\theta(y_0, y_1, K, G)], \\ \ell_{\hat{Q}_{n,\theta}}(y_1) &= \int H_\theta(y_0, y_1) \hat{Q}_n(dy_0 \mid y_1), \end{aligned}$$

$$L_{a,Q}^\Phi(\theta) = \int H_\theta(y_0, y_1) \Gamma_{\mu_1, Q}(dy_0, dy_1).$$

Define $L_{a, \hat{Q}_n}^\Phi$ by replacing Q with $\hat{Q}_n$. For $O = (C, A, Y)$ and $O_i = (C_i, A_i, Y_i)$, set

$$\begin{aligned} m_\theta(c) &= \mathbb{E}_{Y \sim \mathbb{P}_{Y|C=c, A=a}}[\ell_{\hat{Q}_n, \theta}(Y)], \\ \hat{m}_\theta(c) &= \mathbb{E}_{\hat{Y} \sim \hat{\mathbb{P}}_{Y|C=c, A=a}}[\ell_{\hat{Q}_n, \theta}(\hat{Y})], \\ \hat{\ell}_{\text{DR}, \theta}(O) &= \mathbf{1}\{A = a\} \hat{\alpha}(C) \{\ell_{\hat{Q}_n, \theta}(Y) - \hat{m}_\theta(C)\} + \hat{m}_\theta(C). \end{aligned}$$

For the offline draws, write

$$\begin{aligned} \Phi_{i, \theta}^{\text{obs}} &= \Phi_\theta(\hat{Y}_{0,i}^{\text{obs}}, Y_i, K_i^{\text{obs}}, G_i^{\text{obs}}), \\ \bar{\Phi}_{i, \theta}^{\text{aux}} &= \frac{1}{M} \sum_{m=1}^M \Phi_\theta(\hat{Y}_{0,i,m}, \hat{Y}_{i,m}, K_{i,m}, G_{i,m}), \end{aligned}$$

with $\Phi_{i, \theta}^{\text{obs}} = 0$ for $A_i \neq a$. The implemented empirical loss is

$$\hat{L}_{\text{actual}}^\Phi(\theta) = \frac{1}{n} \sum_{i=1}^n \left[\mathbf{1}\{A_i = a\} \hat{\alpha}(C_i) \{\Phi_{i, \theta}^{\text{obs}} - \bar{\Phi}_{i, \theta}^{\text{aux}}\} + \bar{\Phi}_{i, \theta}^{\text{aux}} \right]. \quad (19)$$

The following lemma shows concretely how a covering number enters the concentration argument: we first approximate every function in the class by one of J cover elements, apply Hoeffding's inequality and a union bound to these J representatives, and then incur at most an additional $2r$ error when passing back to the full class.

Lemma 13 (A finite-cover empirical-process bound). *Conditional on a sigma-field $\mathcal{F}$, let $W_1, \dots, W_m$ be independent. Let $\mathcal{Z}$ be a separable class of measurable functions taking values in a common interval of length $B_{\text{env}} > 0$, with a proper uniform r -cover $\{z_1, \dots, z_J\}$. Conditional on $\mathcal{F}$, the function class $\mathcal{Z}$ and the chosen cover $\{z_1, \dots, z_J\}$ are treated as fixed. For $0 < \delta < 1$, with conditional probability at least $1 - \delta$,*

$$\sup_{z \in \mathcal{Z}} \left| \frac{1}{m} \sum_{i=1}^m \{z(W_i) - \mathbb{E}[z(W_i) | \mathcal{F}]\} \right| \leq 2r + B_{\text{env}} \sqrt{\frac{\log(2J/\delta)}{2m}}.$$

Proof. For each z_j , conditional Hoeffding's inequality gives

$$\mathbb{P} \left(\left| \frac{1}{m} \sum_{i=1}^m \{z_j(W_i) - \mathbb{E}[z_j(W_i) | \mathcal{F}]\} \right| > u \mid \mathcal{F} \right) \leq 2 \exp \left(-\frac{2mu^2}{B_{\text{env}}^2} \right).$$

Take $u = B_{\text{env}} \sqrt{\log(2J/\delta)/(2m)}$ and apply a union bound over the cover. Replacing a function by its cover center changes its empirical average and its conditional expected average by at most r each, giving the additional $2r$. $\square$

We now turn minimization of the empirical doubly robust loss into a high-probability bound on the population excess risk, in the following proposition. The error separates into three sources: replacing the ideal coupling Q by the fitted coupling $\hat{Q}_n$, estimation error in the nuisance functions, and finite-sample error from the observed data and offline Monte Carlo draws. The last part is controlled through the covering-number concentration argument above.

Proposition 5 (Doubly Robust ERM Bound). *Under Assumptions 2, 3, and 4, let the loss satisfy the following bounds for constants $B_\Phi, L_Y, L_\Theta > 0$ and all admissible arguments:*

$$0 \leq \Phi_\theta(y_0, y_1, k, g) \leq B_\Phi, \quad (20)$$

$$|H_\theta(y_0, y_1) - H_\theta(y'_0, y'_1)| \leq L_Y(\|y_0 - y'_0\| + \|y_1 - y'_1\|), \quad (21)$$

$$|\Phi_\theta(y_0, y_1, k, g) - \Phi_{\theta'}(y_0, y_1, k, g)| \leq L_\Theta d_\Theta(\theta, \theta'). \quad (22)$$

Define

$$N_{\text{DR}}^\Phi = \mathcal{N}\left(\frac{B_\Phi(1 + 2\pi_{\min}^{-1})}{L_\Theta(1 + 4\pi_{\min}^{-1})\sqrt{n}}, \Theta, d_\Theta\right), N_{\text{obs}}^\Phi = \mathcal{N}\left(\frac{B_\Phi}{L_\Theta\sqrt{n}}, \Theta, d_\Theta\right), N_{\text{aux}}^\Phi = \mathcal{N}\left(\frac{B_\Phi}{L_\Theta\sqrt{nM}}, \Theta, d_\Theta\right).$$

For $0 < \delta < 1$, set

$$\begin{aligned} \mathfrak{R}_\Phi(\delta) &= 4B_\Phi(1 + 2\pi_{\min}^{-1}) \left[\frac{1}{\sqrt{n}} + \sqrt{\frac{2\log(6N_{\text{DR}}^\Phi/\delta)}{n}} \right] \\ &\quad + 8B_\Phi\pi_{\min}^{-1} \left[\frac{1}{\sqrt{n}} + \sqrt{\frac{2\log(6N_{\text{obs}}^\Phi/\delta)}{n}} \right] \\ &\quad + 4B_\Phi(1 + 2\pi_{\min}^{-1}) \left[\frac{1}{\sqrt{nM}} + \sqrt{\frac{2\log(6N_{\text{aux}}^\Phi/\delta)}{nM}} \right]. \end{aligned} \quad (23)$$

Write $\Delta_Q = W_{1,\rho}(\Gamma_{\mu_1,Q}, \Gamma_{\mu_1,\widehat{Q}_n})$. Let $\widehat{\theta}_M$ be an exact minimizer of $\widehat{L}_{\text{actual}}^\Phi$ over Θ . With conditional probability at least $1 - \delta$ over the training observations and offline Monte Carlo draws,

$$L_{a,Q}^\Phi(\widehat{\theta}_M) - \inf_{\theta \in \Theta} L_{a,Q}^\Phi(\theta) \leq 2L_Y\Delta_Q + 2L_Y \left(1 + \frac{2d}{\varepsilon_{\text{OT}}}\right) \varepsilon_\alpha \varepsilon_{\text{pred}} + \mathfrak{R}_\Phi(\delta). \quad (24)$$

Proof. Empirical optimality gives

$$L_{a,Q}^\Phi(\widehat{\theta}_M) - \inf_{\theta \in \Theta} L_{a,Q}^\Phi(\theta) \leq 2 \sup_{\theta \in \Theta} |L_{a,Q}^\Phi(\theta) - \widehat{L}_{\text{actual}}^\Phi(\theta)|.$$

Inserting the fitted-kernel population loss, the population DR loss, and its empirical average gives

$$\begin{aligned} L_{a,Q}^\Phi(\widehat{\theta}_M) - \inf_{\theta \in \Theta} L_{a,Q}^\Phi(\theta) &\leq 2 \sup_{\theta \in \Theta} |L_{a,Q}^\Phi(\theta) - L_{a,\widehat{Q}_n}^\Phi(\theta)| \\ &\quad + 2 \sup_{\theta \in \Theta} |L_{a,\widehat{Q}_n}^\Phi(\theta) - \mathbb{E}[\widehat{\ell}_{\text{DR},\theta}(O)]| \\ &\quad + 2 \sup_{\theta \in \Theta} \left| \mathbb{E}[\widehat{\ell}_{\text{DR},\theta}(O)] - \frac{1}{n} \sum_{i=1}^n \widehat{\ell}_{\text{DR},\theta}(O_i) \right| \\ &\quad + 2 \sup_{\theta \in \Theta} \left| \frac{1}{n} \sum_{i=1}^n \widehat{\ell}_{\text{DR},\theta}(O_i) - \widehat{L}_{\text{actual}}^\Phi(\theta) \right|. \end{aligned} \quad (25)$$

Coupling replacement and nuisance error. By Equation (21) and Kantorovich–Rubinstein duality,

$$\sup_{\theta \in \Theta} |L_{a,Q}^\Phi(\theta) - L_{a,\widehat{Q}_n}^\Phi(\theta)| \leq L_Y\Delta_Q.$$

Conditioning on C and using $\mathbb{P}(A = a \mid C) = 1/\alpha_a(C)$ gives

$$\mathbb{E}[\widehat{\ell}_{\text{DR},\theta}(O) \mid C] = \frac{\widehat{\alpha}(C)}{\alpha_a(C)} \{m_\theta(C) - \widehat{m}_\theta(C)\} + \widehat{m}_\theta(C).$$

Since $L_{a,\widehat{Q}_n}^\Phi(\theta) = \mathbb{E}_C[m_\theta(C)]$,

$$L_{a,\widehat{Q}_n}^\Phi(\theta) - \mathbb{E}[\widehat{\ell}_{\text{DR},\theta}(O)] = \mathbb{E}_C \left[\left(1 - \frac{\widehat{\alpha}(C)}{\alpha_a(C)}\right) \{m_\theta(C) - \widehat{m}_\theta(C)\} \right]. \quad (26)$$

Coupling $\widehat{Q}_n(\cdot \mid y)$ and $\widehat{Q}_n(\cdot \mid y')$ optimally, Equation (21) and Lemma 9 yield

$$\begin{aligned} |\ell_{\widehat{Q}_n,\theta}(y) - \ell_{\widehat{Q}_n,\theta}(y')| &\leq L_Y \{W_1(\widehat{Q}_n(\cdot \mid y), \widehat{Q}_n(\cdot \mid y')) + \|y - y'\|\} \\ &\leq L_Y \left(1 + \frac{2d}{\varepsilon_{\text{OT}}}\right) \|y - y'\|. \end{aligned}$$

Kantorovich–Rubinstein duality and Assumption 4 therefore give

$$\sup_{\theta \in \Theta} \|m_\theta - \widehat{m}_\theta\|_{L^2(\mathbb{P}_C)} \leq L_Y \left(1 + \frac{2d}{\varepsilon_{\text{OT}}}\right) \varepsilon_{\text{pred}}.$$

Cauchy–Schwarz in Equation (26), followed by Assumption 3, bounds the second term of Equation (25) by $2L_Y(1 + 2d/\varepsilon_{\text{OT}})\varepsilon_\alpha\varepsilon_{\text{pred}}$.

Generalization. The training observations remain independent under the initial conditioning. Integration of Equation (20) and Equation (22) gives

$$\begin{aligned} |\widehat{\ell}_{\text{DR},\theta}(O)| &\leq B_\Phi(1 + 2\pi_{\min}^{-1}), \\ |\widehat{\ell}_{\text{DR},\theta}(O) - \widehat{\ell}_{\text{DR},\theta'}(O)| &\leq L_\Theta(1 + 4\pi_{\min}^{-1})d_\Theta(\theta, \theta'). \end{aligned}$$

Thus the parameter cover defining N_{DR}^Φ induces a proper uniform cover of the DR loss class with radius $B_\Phi(1 + 2\pi_{\min}^{-1})/\sqrt{n}$. Lemma 13, with $m = n$, interval length $2B_\Phi(1 + 2\pi_{\min}^{-1})$, and failure probability $\delta/3$, bounds the third term of Equation (25) by the first term of Equation (23).

Monte Carlo approximation. Condition additionally on $O_1, \dots, O_n$. All training observations and fitted objects are now fixed. The observed tuples are independent across i , and the auxiliary tuples are independent across (i, m) , as specified in Appendix H. Dependence within each tuple, including the conditional draw from $\widehat{Q}_n$, is retained. Under this additional conditioning,

$$\begin{aligned} \mathbb{E}[\mathbf{1}\{A_i = a\}\widehat{\alpha}(C_i)\Phi_{i,\theta}^{\text{obs}} \mid O_1, \dots, O_n] &= \mathbf{1}\{A_i = a\}\widehat{\alpha}(C_i)\ell_{\widehat{Q}_n,\theta}(Y_i), \\ \mathbb{E}[\Phi_\theta(\widehat{Y}_{0,i,m}, \widehat{Y}_{i,m}, K_{i,m}, G_{i,m}) \mid O_1, \dots, O_n] &= \widehat{m}_\theta(C_i). \end{aligned}$$

Expanding Equation (19) gives

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^n \widehat{\ell}_{\text{DR},\theta}(O_i) - \widehat{L}_{\text{actual}}^\Phi(\theta) &= \frac{1}{n} \sum_{i=1}^n \mathbf{1}\{A_i = a\}\widehat{\alpha}(C_i)\{\ell_{\widehat{Q}_n,\theta}(Y_i) - \Phi_{i,\theta}^{\text{obs}}\} \\ &\quad + \frac{1}{n} \sum_{i=1}^n [1 - \mathbf{1}\{A_i = a\}\widehat{\alpha}(C_i)]\{\widehat{m}_\theta(C_i) - \overline{\Phi}_{i,\theta}^{\text{aux}}\}. \end{aligned}$$

Apply Lemma 13 to these two averages, including i as a deterministic argument of each integrand. For the observed term, the weighted losses lie in $[0, 2B_\Phi\pi_{\min}^{-1}]$ and are $2\pi_{\min}^{-1}L_\Theta$ -Lipschitz in θ . The

cover defining N_{obs}^Φ therefore has induced radius $2B_\Phi\pi_{\min}^{-1}/\sqrt{n}$. With n draws and failure probability $\delta/3$, twice this discrepancy is bounded by the second term of Equation (23).

For the auxiliary term, apply the lemma to the nM individual weighted losses, not to their n within-observation averages. Their absolute values are at most $B_\Phi(1 + 2\pi_{\min}^{-1})$, and their parameter Lipschitz constant is $(1 + 2\pi_{\min}^{-1})L_\Theta$. The cover defining N_{aux}^Φ has induced radius $B_\Phi(1 + 2\pi_{\min}^{-1})/\sqrt{nM}$. With interval length $2B_\Phi(1 + 2\pi_{\min}^{-1})$ and failure probability $\delta/3$, twice this discrepancy is bounded by the third term of Equation (23). The loss classes used above are separable because they are Lipschitz images of the totally bounded class Θ .

Each Monte Carlo failure probability is at most $\delta/3$ conditional on the training observations and the initially fixed objects. Taking conditional expectations removes the additional conditioning. A union bound with the generalization event therefore gives total conditional failure probability at most δ . Substituting all four bounds into Equation (25) proves Equation (24). $\square$

We remark that Proposition 5 applies to a generic loss class. Its conditions will be verified separately for the velocity and score fields below.

I.1 Velocity-Field Bound

We now specialize the generic doubly robust framework above to velocity-field estimation. Define the squared velocity loss

$$\Phi_v^{\text{vel}}(y_0, y_1, k, g) = \|v(y_0 + t_k(y_1 - y_0) + \sigma g, t_k) - (y_1 - y_0)\|^2.$$

Each parameter θ induces a velocity field v_θ . To simplify notation in this specialization, we write $v = v_\theta$, suppress the abstract index θ , and use the velocity field itself to index the loss, writing Φ_v^{vel} in place of the generic Φ_θ .

The generic population risks $L_{a,Q}^\Phi$ and $L_{a,\hat{Q}_n}^\Phi$ then become the velocity losses below; we write them as $L_{a,Q}(v)$ and $L_{a,\hat{Q}_n}(v)$ for brevity. Then the previous abstract loss becomes

$$\begin{aligned}\ell_Q(v; y_1) &= \mathbb{E}_{K, Y_0 \sim Q(\cdot|y_1), G} \left[\|v(Y_0 + t_K(y_1 - Y_0) + \sigma G, t_K) - (y_1 - Y_0)\|^2 \right], \\ L_{a,Q}(v) &= \mathbb{E}_{Y_1 \sim \int \mathbb{P}_{Y|C=c, A=a} d\mathbb{P}_C(c)} [\ell_Q(v; Y_1)], \\ \ell_{\hat{Q}_n}(v; y_1) &= \mathbb{E}_{K, Y_0 \sim \hat{Q}_n(\cdot|y_1), G} \left[\|v(Y_0 + t_K(y_1 - Y_0) + \sigma G, t_K) - (y_1 - Y_0)\|^2 \right], \\ L_{a,\hat{Q}_n}(v) &= \mathbb{E}_{Y_1 \sim \int \mathbb{P}_{Y|C=c, A=a} d\mathbb{P}_C(c)} [\ell_{\hat{Q}_n}(v; Y_1)].\end{aligned}$$

Since $v_Q^*(X_{t_K}, t_K) = \mathbb{E}[D \mid X_{t_K}, K]$, the residual $D - v_Q^*(X_{t_K}, t_K)$ is orthogonal in L^2 to every measurable function of (X_{t_K}, K) . Expanding the two squared losses gives

$$\begin{aligned}L_{a,Q}(v) - L_{a,Q}(v_Q^*) &= \mathbb{E} \left[\|v(X_{t_K}, t_K) - v_Q^*(X_{t_K}, t_K)\|^2 \right] \\ &= \frac{1}{N} \sum_{k=0}^{N-1} \mathbb{E} \left[\|v(X_{t_k}, t_k) - v_Q^*(X_{t_k}, t_k)\|^2 \right].\end{aligned}\tag{27}$$

The right-hand side corresponds to the quantity used in the sampling analysis.

I.1.1 Velocity Network Class and Approximation

Throughout this subsection, Q denotes the ideal coupling kernel, and we write v_Q^* for the population velocity field associated with this coupling.

We approximate the velocity field by the following scalar neural network. For integers $L \geq 2$ and $p, J \geq 1$, and constants $K, \kappa, \gamma > 0$, define the scalar sparse ReLU class

$$\begin{aligned} & \mathfrak{N}_{\text{scal}}(q, L, p, J, K, \kappa, \gamma) \\ &= \left\{ h : [0, 1]^q \rightarrow \mathbb{R} : \exists p_0, \dots, p_L \in \mathbb{N}, p_0 = q, p_L = 1, \max_{1 \leq \ell \leq L-1} p_\ell \leq p, \right. \\ & \quad \exists W_\ell \in \mathbb{R}^{p_\ell \times p_{\ell-1}}, b_\ell \in \mathbb{R}^{p_\ell}, \ell = 1, \dots, L, \text{ such that } h_0(z) = z, \\ & \quad h_\ell(z) = \text{ReLU}(W_\ell h_{\ell-1}(z) + b_\ell), \quad \ell = 1, \dots, L-1, \\ & \quad h(z) = W_L h_{L-1}(z) + b_L, \\ & \quad \sum_{\ell=1}^L (\|W_\ell\|_0 + \|b_\ell\|_0) \leq J, \max_{1 \leq \ell \leq L} (\|W_\ell\|_{\max} \vee \|b_\ell\|_{\max}) \leq \kappa, \sup_{z \in [0, 1]^q} |h(z)| \leq K, \\ & \quad \left. |h(z) - h(z')| \leq \gamma \|z - z'\|_\infty \text{ for every } z, z' \in [0, 1]^q \right\}. \end{aligned}$$

Thus L counts all affine layers, including the output layer. Here $\|W_\ell\|_0$ and $\|b_\ell\|_0$ denote the number of nonzero entries, and $\|W_\ell\|_{\max}$ and $\|b_\ell\|_{\max}$ denote the largest absolute value of their entries.

For $u, \omega \in \mathbb{R}$, define

$$\text{clip}_\omega(u) = -\omega + \text{ReLU}(u + \omega) - \text{ReLU}(u - \omega).$$

For $x = (x_1, \dots, x_d)$, define

$$\text{clip}_R(x) = (\text{clip}_R(x_1), \dots, \text{clip}_R(x_d)), \quad S_R(x, t) = \left(\frac{\text{clip}_R(x_1) + R}{2R}, \dots, \frac{\text{clip}_R(x_d) + R}{2R}, t \right).$$

We first define some constants that will be used to construct a network class. Let $R > 1$, and define $\Lambda = \frac{16d(2R+1)}{\sigma^2}$ and

$$T_R(z) = (-R + 2Rz_1, \dots, -R + 2Rz_d, z_{d+1}), \quad z \in [0, 1]^{d+1}.$$

Lemma 14 shows that $v_Q^* \circ T_R$ is Λ -Lipschitz jointly in the rescaled spatial and time coordinates.

Lemma 14 (Lipschitz Regularity of the True Vector Field). *Assume Assumption 2. Then v_Q^* admits a version such that, for every $i = 1, \dots, d$, the function*

$$g_i(z) = v_{Q,i}^*(T_R(z)), \quad z \in [0, 1]^{d+1},$$

satisfies

$$|g_i(z) - g_i(z')| \leq \Lambda \|z - z'\|_\infty, \quad z, z' \in [0, 1]^{d+1}.$$

Proof of Lemma 14. All expectations and conditional expectations below are taken under the joint law induced by Q . By Lemma 4, for every $i, j \in \{1, \dots, d\}$,

$$\partial_{x_j} v_{Q,i}^*(x, t) = -\frac{1}{\sigma} \text{Cov}(D_i, G_j \mid X_t = x),$$

and

$$\partial_t v_{Q,i}^*(x, t) = \frac{1}{\sigma} \text{Cov}(D_i, D^\top G \mid X_t = x).$$

Under the conditional law given $X_t = x$, we have

$$G = \frac{x - S_t}{\sigma}.$$

Since x is fixed under this conditional law,

$$\partial_{x_j} v_{Q,i}^*(x, t) = \frac{1}{\sigma^2} \text{Cov}(D_i, S_{t,j} \mid X_t = x),$$

and

$$\partial_t v_{Q,i}^*(x, t) = \frac{1}{\sigma^2} \text{Cov}(D_i, D^\top(x - S_t) \mid X_t = x).$$

Since $Y_0, Y_1 \in [-1, 1]^d$, we have

$$D_i \in [-2, 2], \quad S_{t,j} \in [-1, 1].$$

Hence, conditional on $X_t = x$,

$$\begin{aligned} |D_i - \mathbb{E}[D_i \mid X_t = x]| &\leq 4, \\ |S_{t,j} - \mathbb{E}[S_{t,j} \mid X_t = x]| &\leq 2. \end{aligned}$$

Therefore,

$$\left| \partial_{x_j} v_{Q,i}^*(x, t) \right| \leq \frac{8}{\sigma^2}.$$

For the time derivative, if $x \in [-R, R]^d$, then $\|D\| \leq 2\sqrt{d}$ and $\|x - S_t\| \leq (R+1)\sqrt{d}$. Thus

$$\left| D^\top(x - S_t) \right| \leq 2d(R+1),$$

and so, conditional on $X_t = x$,

$$\left| D^\top(x - S_t) - \mathbb{E}[D^\top(x - S_t) \mid X_t = x] \right| \leq 4d(R+1).$$

Together with $|D_i - \mathbb{E}[D_i \mid X_t = x]| \leq 4$, this gives

$$\left| \partial_t v_{Q,i}^*(x, t) \right| \leq \frac{16d(R+1)}{\sigma^2}.$$

For $g_i(z) = v_{Q,i}^*(T_R(z))$, the preceding derivative bounds imply

$$\begin{aligned} \left| \partial_{z_j} g_i(z) \right| &\leq \frac{16R}{\sigma^2}, \quad j = 1, \dots, d, \\ \left| \partial_{z_{d+1}} g_i(z) \right| &\leq \frac{16d(R+1)}{\sigma^2}. \end{aligned}$$

Therefore, for every $z, z' \in [0, 1]^{d+1}$,

$$|g_i(z) - g_i(z')| \leq \left[\frac{16Rd}{\sigma^2} + \frac{16d(R+1)}{\sigma^2} \right] \|z - z'\|_\infty = \Lambda \|z - z'\|_\infty.$$

□

Based on the scalar neural network defined above, we use it to construct a neural network class that we use to approximate our vector field. Let $\varepsilon > 0$, recall $\Lambda = \frac{16d(2R+1)}{\sigma^2}$. Define $\eta = \frac{\varepsilon}{\sqrt{d}\Lambda}$ as an auxiliary variable that control the approximation accuracy of each scalar network.

Proposition 6 (Approximation Error Under the Noisy Interpolation Law). *Assume Assumption 2. Let $R > 1$, $q = d + 1$ and $\varepsilon > 0$ satisfy $0 < \eta = \varepsilon/(\sqrt{d}\Lambda) < 1$ and*

$$32d^2 \exp\left(-\frac{(R-1)^2}{2\sigma^2}\right) \leq \varepsilon^2.$$

Choose the width and parameter bounds so that $p_{\text{Lip}}(q, \eta) \geq q$ and $\kappa_{\text{Lip}}(q, \eta) \geq 2$. The scalar-network architecture can then be chosen with

$$\begin{aligned} K_{\text{Lip}}(q, \eta) &= \sqrt{q}, \quad L_{\text{Lip}}(q, \eta) = O\left(\log \frac{1}{\eta} + q\right), \quad p_{\text{Lip}}(q, \eta) = O(q + \eta^{-q}), \\ J_{\text{Lip}}(q, \eta) &= O\left(\eta^{-q} \left[\log \frac{1}{\eta} + q\right]\right), \quad \kappa_{\text{Lip}}(q, \eta) = O(1). \end{aligned}$$

These are the scalar approximation parameters from (Chen et al., 2020, Lemma 10), with L_{Lip} denoting the total number of affine layers. Enlarging the width and parameter bounds to enforce $p_{\text{Lip}} \geq q$ and $\kappa_{\text{Lip}} \geq 2$ does not alter these orders for fixed q . We construct a vector valued neural network coordinate-wisely by

$$v_i^{\text{app}}(x, t) = \text{clip}_2\left(\Lambda \hat{h}_i(S_R(x, t)) + g_i(0)\right), \quad i = 1, \dots, d.$$

where $\hat{h}_i \in \mathfrak{N}_{\text{scal}}(q, L_{\text{Lip}}(q, \eta), p_{\text{Lip}}(q, \eta), J_{\text{Lip}}(q, \eta), K_{\text{Lip}}(q, \eta), \kappa_{\text{Lip}}(q, \eta), 10q)$ and $g_i(0)$ are constants that will be specified in the proof. Then we have

$$\frac{1}{N} \sum_{k=0}^{N-1} \mathbb{E} \left[\left\| v^{\text{app}}(X_{t_k}, t_k) - v_Q^*(X_{t_k}, t_k) \right\|^2 \right] \leq 2\varepsilon^2.$$

Proof. By Lemma 14, for each $i = 1, \dots, d$, the function $g_i(z) = v_{Q,i}^*(T_R(z))$, $z \in [0, 1]^{d+1}$ is Λ -Lipschitz with respect to $\|\cdot\|_\infty$. Define

$$f_i(z) = \frac{g_i(z) - g_i(0)}{\Lambda}.$$

Then $f_i(0) = 0$ and f_i is 1-Lipschitz with respect to $\|\cdot\|_\infty$. Since $\|z - z'\|_\infty \leq \|z - z'\|_2$, it is also 1-Lipschitz with respect to the Euclidean norm required by (Chen et al., 2020, Lemma 10). Consequently, that lemma provides

$$\hat{h}_i \in \mathfrak{N}_{\text{scal}}(q, L_{\text{Lip}}(q, \eta), p_{\text{Lip}}(q, \eta), J_{\text{Lip}}(q, \eta), K_{\text{Lip}}(q, \eta), \kappa_{\text{Lip}}(q, \eta), 10q)$$

such that

$$\sup_{z \in [0, 1]^{d+1}} \left| \hat{h}_i(z) - f_i(z) \right| \leq \eta.$$

Now recall we constructed

$$v_i^{\text{app}}(x, t) = \text{clip}_2\left(\Lambda \hat{h}_i(S_R(x, t)) + g_i(0)\right), \quad i = 1, \dots, d.$$

If $x \in [-R, R]^d$, then $T_R(S_R(x, t)) = (x, t)$. Since $\eta = \varepsilon/(\sqrt{d}\Lambda)$,

$$\left| \Lambda \hat{h}_i(S_R(x, t)) + g_i(0) - v_{Q,i}^*(x, t) \right| \leq \frac{\varepsilon}{\sqrt{d}}.$$

Because $v_{Q,i}^*(x, t) \in [-2, 2]$ and clip_2 is the projection onto $[-2, 2]$, clipping does not increase the error. Therefore our constructed vector field approximation satisfies

$$\sup_{x \in [-R, R]^d, t \in [0, 1]} \left\| v^{\text{app}}(x, t) - v_Q^*(x, t) \right\|^2 \leq \varepsilon^2.$$

It remains to control the event $X_t \notin [-R, R]^d$. Since $S_t \in [-1, 1]^d$, for every $t \in [0, 1]$,

$$\mathbb{P} \left(X_t \notin [-R, R]^d \right) \leq 2d \exp \left(-\frac{(R-1)^2}{2\sigma^2} \right).$$

On this event, $\left\| v^{\text{app}}(X_t, t) - v_Q^*(X_t, t) \right\|^2 \leq 16d$. Thus, for every $t \in [0, 1]$,

$$\mathbb{E} \left[\left\| v^{\text{app}}(X_t, t) - v_Q^*(X_t, t) \right\|^2 \right] \leq \varepsilon^2 + 32d^2 \exp \left(-\frac{(R-1)^2}{2\sigma^2} \right) \leq 2\varepsilon^2.$$

The bound holds at every grid time t_k ; averaging over $k = 0, \dots, N-1$ proves the claim. $\square$

The preceding proposition motivates the following coordinatewise and vector-valued network classes:

$$\mathcal{F}_{\text{coord}}(R, \Lambda, L, p, J, K, \kappa, \gamma) = \left\{ f : \mathbb{R}^d \times [0, 1] \rightarrow \mathbb{R} : f(x, t) = \text{clip}_2(\Lambda h(S_R(x, t)) + b), \right. \\ \left. h \in \mathfrak{N}_{\text{scal}}(d+1, L, p, J, K, \kappa, \gamma), \quad |b| \leq 2 \right\}$$

and

$$\mathcal{V}(R, \Lambda, L, p, J, K, \kappa, \gamma) = \left\{ v : \mathbb{R}^d \times [0, 1] \rightarrow \mathbb{R}^d : v(x, t) = (v_1(x, t), \dots, v_d(x, t)), \right. \\ \left. v_i \in \mathcal{F}_{\text{coord}}(R, \Lambda, L, p, J, K, \kappa, \gamma), \quad i = 1, \dots, d \right\}.$$

Concretely, we instantiate the class $\mathcal{V}$ as

$$\mathcal{V}_n = \left\{ v : \mathbb{R}^d \times [0, 1] \rightarrow \mathbb{R}^d : v(x, t) = (v_1(x, t), \dots, v_d(x, t)), \right. \\ v_i(x, t) = \text{clip}_2(\Lambda h_i(S_R(x, t)) + b_i), \quad i = 1, \dots, d, \\ h_i \in \mathfrak{N}_{\text{scal}}(q, L_{\text{Lip}}(q, \eta), p_{\text{Lip}}(q, \eta), J_{\text{Lip}}(q, \eta), K_{\text{Lip}}(q, \eta), \kappa_{\text{Lip}}(q, \eta), 10q), \quad i = 1, \dots, d, \\ \left. |b_i| \leq 2, \quad i = 1, \dots, d \right\}.$$

Then Proposition 6 implies

$$\inf_{v \in \mathcal{V}_n} \frac{1}{N} \sum_{k=0}^{N-1} \mathbb{E} \left[\left\| v(X_{t_k}, t_k) - v_Q^*(X_{t_k}, t_k) \right\|^2 \right] \leq 2\varepsilon^2.$$

We transfer a cover of the scalar subnetworks directly to the clipped vector class.

Lemma 15 (Direct Covering-Number Transfer for Clipped Vector Networks). *Let $\mathcal{H} = \mathfrak{N}_{\text{scal}}(q, L, p, J, K, \kappa, \gamma)$, where L is the total number of affine layers, $p \geq q$, and $\kappa > 1$. Let*

$$\mathcal{V} = \{v : v_i(x, t) = \text{clip}_\omega(\Lambda h_i(S_R(x, t)) + b_i), h_i \in \mathcal{H}, |b_i| \leq \omega, i = 1, \dots, d\}.$$

Then, for every $\delta_{\text{cov}} > 0$,

$$\begin{aligned} \log \mathcal{N}(\delta_{\text{cov}}, \mathcal{V}, \|\cdot\|_{\infty, [-R, R]^d \times [0, 1]}) &\leq dJ \log \left(1 + \frac{8L^2(p+2)\kappa^L p^{L+1} \Lambda \sqrt{d}}{\delta_{\text{cov}}} \right) \\ &\quad + d \log \left(1 + \frac{4\omega \sqrt{d}}{\delta_{\text{cov}}} \right). \end{aligned} \quad (28)$$

Proof. Set

$$\eta_h = \frac{\delta_{\text{cov}}}{2\Lambda\sqrt{d}}, \quad \eta_b = \frac{\delta_{\text{cov}}}{2\sqrt{d}}.$$

After zero-padding hidden layers when necessary, $\mathcal{H}$ is a subclass of the sparse-ReLU class in (Chen et al., 2022, Lemma 5.3) with total depth L , maximum layer width p , input bound $B = 1$, sparsity J , and parameter bound $\kappa > 1$. The cited result therefore supplies an $\eta_h/2$ -cover, not necessarily with centers in $\mathcal{H}$, whose log-cardinality is at most

$$J \log \left(1 + \frac{8L^2(p+2)\kappa^L p^{L+1} \Lambda \sqrt{d}}{\delta_{\text{cov}}} \right).$$

Discard balls that do not intersect $\mathcal{H}$, and select one element of $\mathcal{H}$ from each remaining ball. By the triangle inequality, the selected elements form a proper η_h -cover with no larger cardinality. The interval $[-\omega, \omega]$ has an η_b -cover of cardinality at most $1 + 4\omega\sqrt{d}/\delta_{\text{cov}}$ (this is by definition of cover).

For a coordinate v_i , choose h'_i and b'_i from these covers. Because clip_ω is 1-Lipschitz,

$$\sup_{x, t} |v_i(x, t) - v'_i(x, t)| \leq \Lambda \|h_i - h'_i\|_{\infty, [0, 1]^q} + |b_i - b'_i| \leq \frac{\delta_{\text{cov}}}{\sqrt{d}}.$$

Taking the product cover over the d coordinates gives $\sup_{x, t} \|v(x, t) - v'(x, t)\| \leq \delta_{\text{cov}}$. Multiplying the scalar and bias cover cardinalities and then taking logarithms proves Equation (28). $\square$

For $\mathcal{V}_n$, take $\omega = 2$, Λ as above, and substitute the scalar approximation parameters $L_{\text{Lip}}(q, \eta)$, $p_{\text{Lip}}(q, \eta)$, $J_{\text{Lip}}(q, \eta)$, and $\kappa_{\text{Lip}}(q, \eta)$ into Equation (28). To keep the later bounds readable, define, for $r > 0$,

$$\begin{aligned} \mathfrak{C}_v(r) &:= dJ_{\text{Lip}}(q, \eta) \log \left(1 + \frac{8L_{\text{Lip}}(q, \eta)^2(p_{\text{Lip}}(q, \eta) + 2)\kappa_{\text{Lip}}(q, \eta)^{L_{\text{Lip}}(q, \eta)} p_{\text{Lip}}(q, \eta)^{L_{\text{Lip}}(q, \eta)+1} \Lambda \sqrt{d}}{r} \right) \\ &\quad + d \log \left(1 + \frac{8\sqrt{d}}{r} \right). \end{aligned} \quad (29)$$

Thus, for every $\delta_{\text{cov}} > 0$,

$$\log \mathcal{N}(\delta_{\text{cov}}, \mathcal{V}_n, \|\cdot\|_{\infty, [-R, R]^d \times [0, 1]}) \leq \mathfrak{C}_v(\delta_{\text{cov}}). \quad (30)$$

The function $r \mapsto \mathfrak{C}_v(r)$ is nonincreasing, a fact used below when a loss-class cover is transferred to a cover of $\mathcal{V}_n$.

Lemma 16 (Uniform Bounds for the Vector Network Class). *Every $v \in \mathcal{V}_n$ satisfies*

$$\sup_{x \in \mathbb{R}^d, t \in [0,1]} \|v(x, t)\| \leq 2\sqrt{d}, \quad (31)$$

$$\|v(x, t) - v(x', t)\| \leq \gamma_{x,R,\varepsilon} \|x - x'\|_\infty, \quad x, x' \in \mathbb{R}^d, \quad t \in [0, 1], \quad (32)$$

$$\|v(x, t) - v(x, s)\| \leq \gamma_{t,R,\varepsilon} |t - s|, \quad x \in \mathbb{R}^d, \quad t, s \in [0, 1], \quad (33)$$

where

$$\gamma_{x,R,\varepsilon} = \frac{5(d+1)\sqrt{d}\Lambda}{R}, \quad \gamma_{t,R,\varepsilon} = 10(d+1)\sqrt{d}\Lambda.$$

Proof. From the coordinatewise clipping, we have $|v_i(x, t)| \leq 2, \forall i = 1, \dots, d$. Hence $\|v(x, t)\| \leq 2\sqrt{d}$.

Next, fix $x, x' \in \mathbb{R}^d$ and $t \in [0, 1]$. Since clip_2 is 1-Lipschitz and each h_i is $10q$ -Lipschitz with respect to $\|\cdot\|_\infty$,

$$|v_i(x, t) - v_i(x', t)| \leq \Lambda |h_i(S_R(x, t)) - h_i(S_R(x', t))| \leq 10q\Lambda \|S_R(x, t) - S_R(x', t)\|_\infty.$$

By the definition of S_R , $\|S_R(x, t) - S_R(x', t)\|_\infty \leq \frac{1}{2R} \|x - x'\|_\infty$. Therefore, $|v_i(x, t) - v_i(x', t)| \leq \frac{5(d+1)\Lambda}{R} \|x - x'\|_\infty$. Summing over coordinates gives

$$\|v(x, t) - v(x', t)\| \leq \frac{5(d+1)\sqrt{d}\Lambda}{R} \|x - x'\|_\infty.$$

For variation in the time coordinate, $t, s \in [0, 1]$,

$$|v_i(x, t) - v_i(x, s)| \leq \Lambda |h_i(S_R(x, t)) - h_i(S_R(x, s))| \leq 10q\Lambda \|S_R(x, t) - S_R(x, s)\|_\infty = 10(d+1)\Lambda |t - s|.$$

Summing over coordinates gives $\|v(x, t) - v(x, s)\| \leq 10(d+1)\sqrt{d}\Lambda |t - s|$. $\square$

I.1.2 Completion of the Velocity Bound

Recall that for $v \in \mathcal{V}_n$, we defined the pointwise velocity loss

$$\Phi_v^{\text{vel}}(y_0, y_1, k, g) = \|v(y_0 + t_k(y_1 - y_0) + \sigma g, t_k) - (y_1 - y_0)\|^2. \quad (34)$$

Equip $\mathcal{V}_n$ with

$$d_v(v, \tilde{v}) = \sup_{x \in [-R, R]^d, t \in [0,1]} \|v(x, t) - \tilde{v}(x, t)\|.$$

Lemma 17 (Constants for the velocity loss). *The loss in Equation (34) satisfies Equation (20)–Equation (22) with*

$$\begin{aligned} B_{\text{vel}} &= 16d, \\ L_{Y,\text{vel}} &= 8\sqrt{d}(1 + \gamma_{x,R,\varepsilon}), \\ L_{\Theta,\text{vel}} &= 8\sqrt{d}. \end{aligned} \quad (35)$$

Moreover,

$$\log \mathcal{N}(r, \mathcal{V}_n, d_v) \leq \mathfrak{C}_v(r). \quad (36)$$

Proof. Write

$$S_t = (1 - t)y_0 + ty_1, \quad D = y_1 - y_0,$$

and define S'_t and D' from y'_0, y'_1 . Bounded support and Equation (31) give

$$\|v(S_t + \sigma g, t) - D\| \leq \|v(S_t + \sigma g, t)\| + \|D\| \leq 4\sqrt{d},$$

so the squared loss is at most $16d$.

For two endpoint pairs,

$$\|S_t - S'_t\|_\infty \leq \|y_0 - y'_0\| + \|y_1 - y'_1\|, \quad \|D - D'\| \leq \|y_0 - y'_0\| + \|y_1 - y'_1\|.$$

Using $\|a\|^2 - \|b\|^2 \leq (\|a\| + \|b\|)\|a - b\|$ and Equation (32), we obtain

$$\begin{aligned} & \left| \Phi_v^{\text{vel}}(y_0, y_1, k, g) - \Phi_v^{\text{vel}}(y'_0, y'_1, k, g) \right| \\ & \leq 8\sqrt{d} \left[\gamma_{x, R, \varepsilon} \|S_{t_k} - S'_{t_k}\|_\infty + \|D - D'\| \right] \\ & \leq L_{Y, \text{vel}} (\|y_0 - y'_0\| + \|y_1 - y'_1\|). \end{aligned}$$

Averaging over K, G proves the required endpoint bound for H_v^{vel} .

For $v, \tilde{v} \in \mathcal{V}_n$, both residuals have norm at most $4\sqrt{d}$, and hence

$$\left| \Phi_v^{\text{vel}}(y_0, y_1, k, g) - \Phi_{\tilde{v}}^{\text{vel}}(y_0, y_1, k, g) \right| \leq 8\sqrt{d} \|v(S_{t_k} + \sigma g, t_k) - \tilde{v}(S_{t_k} + \sigma g, t_k)\|.$$

Every network in $\mathcal{V}_n$ depends on x through $\text{clip}_R(x)$, so the last norm is bounded by $d_v(v, \tilde{v})$. This proves the parameter bound. Finally, Equation (36) is exactly Equation (30) under the metric d_v . $\square$

Proof of the velocity bound in Proposition 1. Choose

$$\varepsilon = n^{-1/(d+5)}, \quad R = 1 + \sigma \left[2 \log \left(\frac{32d^2}{\varepsilon^2} \right) \right]^{1/2}.$$

This choice satisfies the tail condition in Proposition 6. By Equation (27) and that proposition,

$$\inf_{v \in \mathcal{V}_n} \left\{ L_{a, Q}(v) - L_{a, Q}(v_Q^*) \right\} \leq 2\varepsilon^2. \quad (37)$$

For Φ_v^{vel} , $L_{a, Q}^\Phi(v) = L_{a, Q}(v)$, and $\hat{L}_{\text{actual}}^\Phi(v)$ is the velocity training objective. Lemma 17 verifies every loss hypothesis of Proposition 5, with

$$\Theta = \mathcal{V}_n, \quad d_\Theta = d_v, \quad B_\Phi = B_{\text{vel}}, \quad L_Y = L_{Y, \text{vel}}, \quad L_\Theta = L_{\Theta, \text{vel}}.$$

We now evaluate the quantities in that proposition. Since $\Lambda = 16d(2R + 1)/\sigma^2$,

$$\eta^{-1} = \frac{\sqrt{d}\Lambda}{\varepsilon} = \frac{16d^{3/2} \left[3 + 2\sigma\sqrt{2\log(32d^2/\varepsilon^2)} \right]}{\sigma^2\varepsilon}. \quad (38)$$

The architecture orders in Proposition 6, substituted into $\mathfrak{C}_v$, imply for the three radii in Proposition 5

$$\log N_{\text{DR}}^{\text{vel}} + \log N_{\text{obs}}^{\text{vel}} + \log N_{\text{aux}}^{\text{vel}} = \tilde{O} \left(\eta^{-(d+1)} \right). \quad (39)$$

Here $\tilde{O}$ hides logarithms in $n, M, \sigma^{-1}, \varepsilon_{\text{OT}}^{-1}$, and $\pi_{\min}^{-1}$, and constants depending on fixed d . Combining Equation (23), Equation (38), and Equation (39), with $\delta = n^{-1}$, give

$$\mathfrak{R}_{\text{vel}}(n^{-1}) \lesssim \frac{1 + 2\pi_{\min}^{-1}}{\sigma^{d+1}} n^{-2/(d+5)}. \quad (40)$$

Indeed, before substituting ε , the leading term is

$$\tilde{O}\left((1 + 2\pi_{\min}^{-1})\sigma^{-(d+1)}\varepsilon^{-(d+1)/2}n^{-1/2}\right),$$

and $\varepsilon^{-(d+1)/2}n^{-1/2} = n^{-2/(d+5)}$.

The endpoint constant has the explicit order

$$L_{Y,\text{vel}} = 8\sqrt{d}\left[1 + \frac{5(d+1)\sqrt{d}\Lambda}{R}\right] = O(\sigma^{-2}), \quad (41)$$

because $R \geq 1$ and $\Lambda/R = 16d(2 + R^{-1})/\sigma^2$. On the event in Corollary 3,

$$\Delta_Q \lesssim \varepsilon_{\text{OT}}^{-1} \left[n^{-1/(d+2)} + \varepsilon_{\text{pred,train}} + (nM_Q)^{-1/(d+2)} \right].$$

Substitute this display, Equation (37), Equation (40), and Equation (41) into Equation (24). Since $M_Q \geq 1$, $(nM_Q)^{-1/(d+2)} \leq n^{-1/(d+2)}$, and since $0 < \sigma \leq 1$, the approximation term is absorbed by Equation (40). Thus

$$\begin{aligned} L_{a,Q}(\hat{v}_M) - L_{a,Q}(v_Q^*) &\lesssim \frac{\varepsilon_{\text{OT}}^{-1}}{\sigma^2} \left[n^{-1/(d+2)} + \varepsilon_{\text{pred,train}} \right] \\ &\quad + \frac{1 + \varepsilon_{\text{OT}}^{-1}}{\sigma^2} \varepsilon_\alpha \varepsilon_{\text{pred}} + \frac{1 + 2\pi_{\min}^{-1}}{\sigma^{d+1}} n^{-2/(d+5)}. \end{aligned}$$

Applying Equation (27) with $v = \hat{v}_M$ proves Equation (4). The coupling event has probability at least $1 - 2/n$, and the conditional ERM bound with $\delta = 1/n$ has probability at least $1 - 1/n$. Finally, $\mathbb{P}(\mathcal{E}_a^c) \leq e^{-n\pi_{\min}}$, which gives the stated failure probability. $\square$

I.2 Score-Field Bound

Recall that $S_t = Y_0 + t(Y_1 - Y_0)$ and $X_t = S_t + \sigma G$. We define the oracle by $u_Q^*(x, t) = \mathbb{E}[S_t \mid X_t = x]$. By the Gaussian score identity,

$$s_Q^*(x, t) = \frac{u_Q^*(x, t) - x}{\sigma^2}.$$

Instead of approximating s_Q^* directly, we approximate u_Q^* . For every $u : \mathbb{R}^d \times [0, 1] \rightarrow \mathbb{R}^d$, define the induced score field by

$$s_u(x, t) = \frac{u(x, t) - x}{\sigma^2}.$$

I.2.1 Score Loss and Exact Risk Identity

We now specialize the generic doubly robust bound to the score estimator. Here the abstract parameter θ indexes a conditional-mean field $u_\theta \in \mathcal{U}_n$, where $\mathcal{U}_n$ is defined later in Equation (52). This in turn induces the score field

$$s_{u_\theta}(x, t) = \frac{u_\theta(x, t) - x}{\sigma^2}.$$

For brevity, we suppress the abstract index θ and write u for u_θ and s_u for the induced score field. For $u \in \mathcal{U}_n$, define

$$\Phi_u^{\text{score}}(y_0, y_1, k, g) = \left\| s_u(y_0 + t_k(y_1 - y_0) + \sigma g, t_k) + \frac{g}{\sigma} \right\|^2.$$

Recall population minimizer of the denoising score-matching loss satisfies

$$s_Q^*(x, t) = \nabla_x \log q_{t,Q}(x) = \mathbb{E} \left[-\frac{G}{\sigma} \mid Y_0 + t(Y_1 - Y_0) + \sigma G = x \right],$$

where $Y_1 \sim \int \mathbb{P}_{Y|C=c, A=a} d\mathbb{P}_C(c)$ and $Y_0 \sim Q(\cdot \mid Y_1)$.

For a score field $s : \mathbb{R}^d \times [0, 1] \rightarrow \mathbb{R}^d$, let $K \sim \text{Unif}\{0, \dots, N-1\}$, independently of all endpoint and Gaussian variables, we have

$$\begin{aligned} r_Q(s; y_1) &= \mathbb{E}_{K, Y_0 \sim Q(\cdot \mid y_1), G \sim \mathcal{N}(0, I_d)} \left[\left\| s(Y_0 + t_K(y_1 - Y_0) + \sigma G, t_K) + \frac{G}{\sigma} \right\|^2 \right], \\ L_{a, Q}^{\text{score}}(s) &= \mathbb{E}_{Y_1 \sim \int \mathbb{P}_{Y|C=c, A=a} d\mathbb{P}_C} [r_Q(s; Y_1)], \\ r_{\hat{Q}_n}(s; y_1) &= \mathbb{E}_{K, Y_0 \sim \hat{Q}_n(\cdot \mid y_1), G \sim \mathcal{N}(0, I_d)} \left[\left\| s(Y_0 + t_K(y_1 - Y_0) + \sigma G, t_K) + \frac{G}{\sigma} \right\|^2 \right], \\ L_{a, \hat{Q}_n}^{\text{score}}(s) &= \mathbb{E}_{Y_1 \sim \int \mathbb{P}_{Y|C=c, A=a} d\mathbb{P}_C} [r_{\hat{Q}_n}(s; Y_1)]. \end{aligned}$$

Exact Excess-Risk Identity. Because $s_Q^*(X_{t_K}, t_K) = \mathbb{E}[-G/\sigma \mid X_{t_K}, K]$, conditional projection gives, for every score field s with finite population loss,

$$\begin{aligned} L_{a, Q}^{\text{score}}(s) - L_{a, Q}^{\text{score}}(s_Q^*) &= \mathbb{E} \left[\left\| s(X_{t_K}, t_K) - s_Q^*(X_{t_K}, t_K) \right\|^2 \right] \\ &= \frac{1}{N} \sum_{k=0}^{N-1} \mathbb{E} \left[\left\| s(X_{t_k}, t_k) - s_Q^*(X_{t_k}, t_k) \right\|^2 \right]. \end{aligned} \quad (42)$$

Thus it remains to control the excess risk of the empirical minimizer.

I.2.2 Score Network Class and Approximation

The denoising score-matching loss is evaluated under the joint sampling procedure in which the same Gaussian draw G is used both to construct $X_t = S_t + \sigma G$ and to form the denoising target $-G/\sigma$. Consequently,

$$\left\| s_u(S_t + \sigma G, t) + \frac{G}{\sigma} \right\|^2 = \left\| \frac{u(S_t + \sigma G, t) - S_t - \sigma G}{\sigma^2} + \frac{G}{\sigma} \right\|^2 = \frac{1}{\sigma^4} \|u(S_t + \sigma G, t) - S_t\|^2. \quad (43)$$

This identity is not a statement about arbitrary independently chosen $x \in \mathbb{R}^d$ and $g \in \mathbb{R}^d$. It applies to the jointly constructed pair $x = S_t + \sigma g$, with the same g appearing in the denoising target. Consequently, the population score loss can be written as

$$L_{a,Q}^{\text{score}}(s_u) = \frac{1}{\sigma^4} \mathbb{E}_K \mathbb{E} \left[\|u(X_{t_K}, t_K) - S_{t_K}\|^2 \right].$$

Since $u_Q^*(X_t, t) = \mathbb{E}[S_t | X_t]$, the conditional projection identity gives

$$L_{a,Q}^{\text{score}}(s_u) - L_{a,Q}^{\text{score}}(s_Q^*) = \frac{1}{\sigma^4} \mathbb{E}_K \mathbb{E} \left[\|u(X_{t_K}, t_K) - u_Q^*(X_{t_K}, t_K)\|^2 \right]. \quad (44)$$

Moreover, $s_u(X_t, t) - s_Q^*(X_t, t) = \{u(X_t, t) - u_Q^*(X_t, t)\}/\sigma^2$, so

$$L_{a,Q}^{\text{score}}(s_u) - L_{a,Q}^{\text{score}}(s_Q^*) = \mathbb{E}_K \mathbb{E} \left[\|s_u(X_{t_K}, t_K) - s_Q^*(X_{t_K}, t_K)\|^2 \right]. \quad (45)$$

Under Assumption 2, $S_t \in [-1, 1]^d$ almost surely. Therefore,

$$u_Q^*(x, t) = \mathbb{E}[S_t | X_t = x] \in [-1, 1]^d \quad (46)$$

for every $x \in \mathbb{R}^d$ and $t \in [0, 1]$. Thus u_Q^* is globally bounded, although the induced score $s_Q^*(x, t) = \{u_Q^*(x, t) - x\}/\sigma^2$ is not globally bounded.

Recall that for $R > 1$, we defined

$$T_R(z) = (-R + 2Rz_1, \dots, -R + 2Rz_d, z_{d+1}), \quad z \in [0, 1]^{d+1}.$$

Lemma 18 (Lipschitz Regularity of the Conditional Mean Oracle). *Assume Assumption 2. For $i = 1, \dots, d$, define $g_i(z) = u_{Q,i}^*(T_R(z))$. Then*

$$|g_i(z) - g_i(z')| \leq \Lambda_u \|z - z'\|_\infty, \quad z, z' \in [0, 1]^{d+1},$$

where

$$\Lambda_u = 2 + \frac{8d(2R+1)}{\sigma^2}. \quad (47)$$

Proof. By the score identity, $u_Q^*(x, t) = x + \sigma^2 s_Q^*(x, t)$. By Lemma 4,

$$D_x s_Q^*(x, t) = \frac{1}{\sigma^2} (\text{Cov}(G, G | X_t = x) - I_d).$$

Therefore,

$$D_x u_Q^*(x, t) = I_d + \sigma^2 D_x s_Q^*(x, t) = \text{Cov}(G, G | X_t = x).$$

Under the conditioning $X_t = x$, we have $G = (x - S_t)/\sigma$, and hence

$$G - \mathbb{E}[G | X_t = x] = -\frac{S_t - \mathbb{E}[S_t | X_t = x]}{\sigma}.$$

Consequently, for every $i, j \in \{1, \dots, d\}$,

$$\partial_{x_j} u_{Q,i}^*(x, t) = \text{Cov}(G_i, G_j | X_t = x) = \frac{1}{\sigma^2} \text{Cov}(S_{t,i}, S_{t,j} | X_t = x).$$

Since $S_t \in [-1, 1]^d$,

$$|S_{t,i} - \mathbb{E}[S_{t,i} \mid X_t = x]| \leq 2, \quad |S_{t,j} - \mathbb{E}[S_{t,j} \mid X_t = x]| \leq 2.$$

It follows that

$$\left| \partial_{x_j} u_{Q,i}^*(x, t) \right| \leq \frac{1}{\sigma^2} \mathbb{E}[|S_{t,i} - \mathbb{E}[S_{t,i} \mid X_t = x]| |S_{t,j} - \mathbb{E}[S_{t,j} \mid X_t = x]| \mid X_t = x] \leq \frac{4}{\sigma^2}. \quad (48)$$

We next control the time derivative. Write $D = Y_1 - Y_0$ and, for fixed x , define the unnormalized posterior weight

$$w_t(Y_0, Y_1; x) = \exp\left(-\frac{\|x - S_t\|^2}{2\sigma^2}\right).$$

Bounded support permits differentiation under the endpoint expectation, and

$$\partial_t w_t(Y_0, Y_1; x) = \frac{D^\top(x - S_t)}{\sigma^2} w_t(Y_0, Y_1; x).$$

Therefore, for any differentiable scalar $F_t = F_t(Y_0, Y_1)$, the quotient rule applied to

$$\mathbb{E}[F_t \mid X_t = x] = \frac{\mathbb{E}[F_t w_t(Y_0, Y_1; x)]}{\mathbb{E}[w_t(Y_0, Y_1; x)]}$$

gives the posterior differentiation identity

$$\partial_t \mathbb{E}[F_t \mid X_t = x] = \mathbb{E}[\partial_t F_t \mid X_t = x] + \frac{1}{\sigma^2} \text{Cov}\left(F_t, D^\top(x - S_t) \mid X_t = x\right).$$

Taking $F_t = S_{t,i}$, for which $\partial_t F_t = D_i$, yields

$$\partial_t u_{Q,i}^*(x, t) = \mathbb{E}[Y_{1,i} - Y_{0,i} \mid X_t = x] + \frac{1}{\sigma^2} \text{Cov}\left(S_{t,i}, (Y_1 - Y_0)^\top(x - S_t) \mid X_t = x\right).$$

The first term satisfies

$$|\mathbb{E}[Y_{1,i} - Y_{0,i} \mid X_t = x]| \leq 2.$$

If $x \in [-R, R]^d$, then

$$\left| (Y_1 - Y_0)^\top(x - S_t) \right| \leq \sum_{j=1}^d |Y_{1,j} - Y_{0,j}| |x_j - S_{t,j}| \leq 2d(R + 1).$$

Therefore,

$$\begin{aligned} & \left| (Y_1 - Y_0)^\top(x - S_t) - \mathbb{E}\left[(Y_1 - Y_0)^\top(x - S_t) \mid X_t = x\right] \right| \\ & \leq \left| (Y_1 - Y_0)^\top(x - S_t) \right| + \mathbb{E}\left[\left| (Y_1 - Y_0)^\top(x - S_t) \right| \mid X_t = x\right] \\ & \leq 4d(R + 1). \end{aligned}$$

Together with $|S_{t,i} - \mathbb{E}[S_{t,i} \mid X_t = x]| \leq 2$, this gives

$$\left| \partial_t u_{Q,i}^*(x, t) \right| \leq 2 + \frac{1}{\sigma^2} \mathbb{E}[|S_{t,i} - \mathbb{E}[S_{t,i} \mid X_t = x]|]$$

$$\begin{aligned} & \cdot \left| (Y_1 - Y_0)^\top (x - S_t) - \mathbb{E} \left[(Y_1 - Y_0)^\top (x - S_t) \mid X_t = x \right] \right| \mid X_t = x \Big] \\ & \leq 2 + \frac{8d(R+1)}{\sigma^2}. \end{aligned} \quad (49)$$

For $j = 1, \dots, d$, the chain rule and Equation (48) give $|\partial_{z_j} g_i(z)| = 2R \left| \partial_{x_j} u_{Q,i}^*(T_R(z)) \right| \leq \frac{8R}{\sigma^2}$, so we get

$$\begin{aligned} |g_i(z) - g_i(z')| & \leq \left[\sum_{j=1}^{d+1} \sup_{\tilde{z} \in [0,1]^{d+1}} |\partial_{z_j} g_i(\tilde{z})| \right] \|z - z'\|_\infty \leq \left[\frac{8Rd}{\sigma^2} + 2 + \frac{8d(R+1)}{\sigma^2} \right] \|z - z'\|_\infty \\ & = \left[2 + \frac{8d(2R+1)}{\sigma^2} \right] \|z - z'\|_\infty = \Lambda_u \|z - z'\|_\infty. \end{aligned}$$

□

Network Construction and Approximation Guarantee. Let $\varepsilon > 0$, and define

$$\eta_u = \frac{\sigma^2 \varepsilon}{\sqrt{d} \Lambda_u}. \quad (50)$$

Proposition 7 (Approximation Error for the Score Oracle Under the Conditional Mean Parametrization). *Assume Assumption 2. Let $R > 1$, $q = d + 1$, and $\varepsilon > 0$ satisfy $0 < \eta_u = \sigma^2 \varepsilon / (\sqrt{d} \Lambda_u) < 1$ and*

$$\frac{8d^2}{\sigma^4} \exp \left(-\frac{(R-1)^2}{2\sigma^2} \right) \leq \varepsilon^2. \quad (51)$$

Use the same architecture bounds $K_{\text{Lip}}(q, \eta_u)$, $L_{\text{Lip}}(q, \eta_u)$, $p_{\text{Lip}}(q, \eta_u)$, $J_{\text{Lip}}(q, \eta_u)$, and $\kappa_{\text{Lip}}(q, \eta_u)$ as in Proposition 6, now evaluated at η_u . In particular, L_{Lip} counts total affine layers, $p_{\text{Lip}} \geq q$, and $\kappa_{\text{Lip}} \geq 2$. For every $i = 1, \dots, d$, construct

$$u_i^{\text{app}}(x, t) = \text{clip}_1 \left(\Lambda_u \hat{h}_i(S_R(x, t)) + g_i(0) \right),$$

where

$$\hat{h}_i \in \mathfrak{N}_{\text{scal}}(q, L_{\text{Lip}}(q, \eta_u), p_{\text{Lip}}(q, \eta_u), J_{\text{Lip}}(q, \eta_u), K_{\text{Lip}}(q, \eta_u), \kappa_{\text{Lip}}(q, \eta_u), 10q),$$

and $g_i(z) = u_{Q,i}^*(T_R(z))$. Define $u^{\text{app}} = (u_1^{\text{app}}, \dots, u_d^{\text{app}})$ and $s^{\text{app}}(x, t) = \{u^{\text{app}}(x, t) - x\} / \sigma^2$. Then

$$\frac{1}{N} \sum_{k=0}^{N-1} \mathbb{E} \left[\left\| s^{\text{app}}(X_{t_k}, t_k) - s_Q^*(X_{t_k}, t_k) \right\|^2 \right] \leq 2\varepsilon^2.$$

Proof. The proof has two subclaims: uniform approximation on the cube $[-R, R]^d \times [0, 1]$, followed by control of the Gaussian tail outside the cube.

For the first subclaim, define

$$f_i(z) = \frac{g_i(z) - g_i(0)}{\Lambda_u}, \quad z \in [0, 1]^{d+1}.$$

Lemma 18 implies that $f_i(0) = 0$ and that f_i is 1-Lipschitz in the ℓ_∞ norm. Since $\|z - z'\|_\infty \leq \|z - z'\|_2$, it is also 1-Lipschitz in the Euclidean norm required by the scalar approximation result invoked in Proposition 6. That result provides $\hat{h}_i$ in the displayed network class such that

$$\sup_{z \in [0,1]^{d+1}} |\hat{h}_i(z) - f_i(z)| \leq \eta_u.$$

If $x \in [-R, R]^d$, then $T_R(S_R(x, t)) = (x, t)$, and hence

$$\left| \Lambda_u \widehat{h}_i(S_R(x, t)) + g_i(0) - u_{Q,i}^*(x, t) \right| \leq \Lambda_u \eta_u = \frac{\sigma^2 \varepsilon}{\sqrt{d}}.$$

Both $u_{Q,i}^*(x, t)$ and the clipped approximation lie in $[-1, 1]$, so projection onto this interval cannot increase the error. Consequently,

$$\sup_{x \in [-R, R]^d, t \in [0, 1]} \|s^{\text{app}}(x, t) - s_Q^*(x, t)\|^2 \leq \varepsilon^2.$$

For the second subclaim, $S_t \in [-1, 1]^d$ almost surely, and a coordinatewise Gaussian tail bound followed by a union bound gives, uniformly in t ,

$$\mathbb{P}\{X_t \notin [-R, R]^d\} \leq 2d \exp\left(-\frac{(R-1)^2}{2\sigma^2}\right).$$

On this event, the two conditional mean remain in $[-1, 1]^d$, so

$$\|s^{\text{app}}(X_t, t) - s_Q^*(X_t, t)\|^2 = \frac{\|u^{\text{app}}(X_t, t) - u_Q^*(X_t, t)\|^2}{\sigma^4} \leq \frac{4d}{\sigma^4}.$$

The tail contribution is therefore at most

$$\frac{8d^2}{\sigma^4} \exp\left(-\frac{(R-1)^2}{2\sigma^2}\right) \leq \varepsilon^2$$

by Equation (51). Adding the cube and tail contributions proves the claimed $2\varepsilon^2$ bound at every grid time, and averaging over $k = 0, \dots, N-1$ completes the proof. $\square$

The preceding proposition motivates

$$\begin{aligned} \mathcal{U}_n = \Big\{ & u : \mathbb{R}^d \times [0, 1] \rightarrow \mathbb{R}^d : u(x, t) = (u_1(x, t), \dots, u_d(x, t)), \\ & u_i(x, t) = \text{clip}_1(\Lambda_u h_i(S_R(x, t)) + b_i), \quad i = 1, \dots, d, \\ & h_i \in \mathfrak{N}_{\text{scal}}(q, L_{\text{Lip}}(q, \eta_u), p_{\text{Lip}}(q, \eta_u), J_{\text{Lip}}(q, \eta_u), K_{\text{Lip}}(q, \eta_u), \kappa_{\text{Lip}}(q, \eta_u), 10q), \\ & |b_i| \leq 1, \quad i = 1, \dots, d \Big\}. \end{aligned} \tag{52}$$

The induced score class is

$$\mathcal{S}_n = \left\{ s_u : \mathbb{R}^d \times [0, 1] \rightarrow \mathbb{R}^d : s_u(x, t) = \frac{u(x, t) - x}{\sigma^2}, \quad u \in \mathcal{U}_n \right\}. \tag{53}$$

Proposition 7 implies

$$\inf_{s \in \mathcal{S}_n} \frac{1}{N} \sum_{k=0}^{N-1} \mathbb{E} \left[\left\| s(X_{t_k}, t_k) - s_Q^*(X_{t_k}, t_k) \right\|^2 \right] \leq 2\varepsilon^2.$$

Covering Number Evaluation. The conditional-mean class is an instance of Lemma 15 with clipping level $\omega = 1$, scale Λ_u , and scalar-network accuracy $\eta_u = \sigma^2 \varepsilon / (\sqrt{d} \Lambda_u)$. Its scalar architecture orders are

$$\begin{aligned} L_{\text{Lip}}(q, \eta_u) &= O\left(\log \frac{\sqrt{d} \Lambda_u}{\sigma^2 \varepsilon} + d\right), \\ p_{\text{Lip}}(q, \eta_u) &= O\left(d + 1 + \left(\frac{\sqrt{d} \Lambda_u}{\sigma^2 \varepsilon}\right)^{d+1}\right), \\ J_{\text{Lip}}(q, \eta_u) &= O\left(\left(\frac{\sqrt{d} \Lambda_u}{\sigma^2 \varepsilon}\right)^{d+1} \left[\log \frac{\sqrt{d} \Lambda_u}{\sigma^2 \varepsilon} + d\right]\right). \end{aligned} \quad (54)$$

For $r > 0$, define

$$\begin{aligned} \mathfrak{C}_u(r) &:= d J_{\text{Lip}}(q, \eta_u) \log \left(1 + \frac{8 L_{\text{Lip}}(q, \eta_u)^2 (p_{\text{Lip}}(q, \eta_u) + 2) \kappa_{\text{Lip}}(q, \eta_u)^{L_{\text{Lip}}(q, \eta_u)} p_{\text{Lip}}(q, \eta_u)^{L_{\text{Lip}}(q, \eta_u)+1} \Lambda_u \sqrt{d}}{r}\right) \\ &\quad + d \log \left(1 + \frac{4\sqrt{d}}{r}\right). \end{aligned} \quad (55)$$

Therefore, for every $\delta_{\text{cov}} > 0$,

$$\log \mathcal{N}(\delta_{\text{cov}}, \mathcal{U}_n, \|\cdot\|_{\infty, [-R, R]^d \times [0, 1]}) \leq \mathfrak{C}_u(\delta_{\text{cov}}). \quad (56)$$

For $u, \tilde{u} \in \mathcal{U}_n$,

$$s_u(x, t) - s_{\tilde{u}}(x, t) = \frac{u(x, t) - \tilde{u}(x, t)}{\sigma^2}.$$

Therefore,

$$\sup_{x \in [-R, R]^d, t \in [0, 1]} \|s_u(x, t) - s_{\tilde{u}}(x, t)\| = \frac{1}{\sigma^2} \sup_{x \in [-R, R]^d, t \in [0, 1]} \|u(x, t) - \tilde{u}(x, t)\|.$$

It follows that

$$\mathcal{N}(\delta_{\text{cov}}, \mathcal{S}_n, \|\cdot\|_{\infty, [-R, R]^d \times [0, 1]}) \leq \mathcal{N}(\sigma^2 \delta_{\text{cov}}, \mathcal{U}_n, \|\cdot\|_{\infty, [-R, R]^d \times [0, 1]}).$$

Hence

$$\log \mathcal{N}(\delta_{\text{cov}}, \mathcal{S}_n, \|\cdot\|_{\infty, [-R, R]^d \times [0, 1]}) \leq \mathfrak{C}_u(\sigma^2 \delta_{\text{cov}}). \quad (57)$$

The fixed affine term $-x/\sigma^2$ does not increase the covering number, because it is identical for every element of $\mathcal{S}_n$ and cancels when two induced score fields are subtracted.

Uniform Bounds Used in the Error Analysis.

Lemma 19 (Uniform Bounds for the Conditional-Mean Class and Induced Score Field). *Every $u \in \mathcal{U}_n$ satisfies*

$$\sup_{x \in \mathbb{R}^d, t \in [0, 1]} \|u(x, t)\| \leq \sqrt{d}, \quad (58)$$

$$\|u(x, t) - u(x', t)\| \leq \gamma_{x, R, \varepsilon}^u \|x - x'\|_{\infty}, \quad (59)$$

$$\|u(x, t) - u(x, r)\| \leq \gamma_{t,R,\varepsilon}^u |t - r|, \quad (60)$$

where

$$\gamma_{x,R,\varepsilon}^u = \frac{5(d+1)\sqrt{d}\Lambda_u}{R}, \quad \gamma_{t,R,\varepsilon}^u = 10(d+1)\sqrt{d}\Lambda_u.$$

The induced score field satisfies

$$\|s_u(x, t) - s_u(x', t)\| \leq \frac{\gamma_{x,R,\varepsilon}^u + \sqrt{d}}{\sigma^2} \|x - x'\|_\infty, \quad (61)$$

$$\|s_u(x, t) - s_u(x, r)\| \leq \frac{\gamma_{t,R,\varepsilon}^u}{\sigma^2} |t - r|. \quad (62)$$

No finite uniform bound on $\sup_{x \in \mathbb{R}^d, t \in [0,1]} \|s_u(x, t)\|$ is asserted.

Proof. The proof has two parts. First, clipping each coordinate to $[-1, 1]$ immediately gives $\|u(x, t)\| \leq \sqrt{d}$.

For regularity, each scalar subnetwork h_i is $10q = 10(d+1)$ -Lipschitz in the ℓ_∞ norm. The spatial part of S_R is $1/(2R)$ -Lipschitz, while its time coordinate is 1-Lipschitz. Since the outer clipping map is 1-Lipschitz,

$$\begin{aligned} |u_i(x, t) - u_i(x', t)| &\leq \frac{5(d+1)\Lambda_u}{R} \|x - x'\|_\infty, \\ |u_i(x, t) - u_i(x, r)| &\leq 10(d+1)\Lambda_u |t - r|. \end{aligned}$$

Summing the squared coordinate bounds proves Equation (59) and Equation (60).

Finally,

$$s_u(x, t) - s_u(x', t) = \frac{u(x, t) - u(x', t) - (x - x')}{\sigma^2}.$$

Using $\|x - x'\| \leq \sqrt{d} \|x - x'\|_\infty$ gives Equation (61); the time bound follows because the affine term $-x/\sigma^2$ is independent of time. The score itself is not uniformly bounded because of this affine term, which is why the subsequent loss analysis uses its exact cancellation with the denoising target. $\square$

Lemma 20 (Uniform Bounds and Lipschitz Properties of the Denoising Score Loss). *For every $u \in \mathcal{U}_n$, $y_t \in [-1, 1]^d$, $t \in [0, 1]$, and $g \in \mathbb{R}^d$,*

$$0 \leq \left\| s_u(y_t + \sigma g, t) + \frac{g}{\sigma} \right\|^2 \leq \frac{4d}{\sigma^4}. \quad (63)$$

For every $y_t, y'_t \in [-1, 1]^d$, using the same g in both evaluation points,

$$\left| \left\| s_u(y_t + \sigma g, t) + \frac{g}{\sigma} \right\|^2 - \left\| s_u(y'_t + \sigma g, t) + \frac{g}{\sigma} \right\|^2 \right| \leq \frac{4d}{\sigma^4} \left(1 + \frac{5(d+1)\Lambda_u}{R} \right) \|y_t - y'_t\|_\infty. \quad (64)$$

Finally, for every $u, \tilde{u} \in \mathcal{U}_n$,

$$\left| \left\| s_u(y_t + \sigma g, t) + \frac{g}{\sigma} \right\|^2 - \left\| s_{\tilde{u}}(y_t + \sigma g, t) + \frac{g}{\sigma} \right\|^2 \right| \leq \frac{4\sqrt{d}}{\sigma^4} \|u(y_t + \sigma g, t) - \tilde{u}(y_t + \sigma g, t)\|. \quad (65)$$

Proof. By $s_u(x, t) = \{u(x, t) - x\}/\sigma^2$, and because the same g appears in the evaluation point $y_t + \sigma g$ and in the denoising target g/σ ,

$$s_u(y_t + \sigma g, t) + \frac{g}{\sigma} = \frac{u(y_t + \sigma g, t) - (y_t + \sigma g)}{\sigma^2} + \frac{g}{\sigma} = \frac{u(y_t + \sigma g, t) - y_t}{\sigma^2}.$$

Since $u(y_t + \sigma g, t) \in [-1, 1]^d$ and $y_t \in [-1, 1]^d$,

$$\left\| s_u(y_t + \sigma g, t) + \frac{g}{\sigma} \right\| = \frac{\|u(y_t + \sigma g, t) - y_t\|}{\sigma^2} \leq \frac{2\sqrt{d}}{\sigma^2}.$$

Squaring proves Equation (63).

For $y_t, y'_t \in [-1, 1]^d$, applying the cancellation identity at each evaluation point gives

$$s_u(y_t + \sigma g, t) + \frac{g}{\sigma} = \frac{u(y_t + \sigma g, t) - y_t}{\sigma^2}, \quad s_u(y'_t + \sigma g, t) + \frac{g}{\sigma} = \frac{u(y'_t + \sigma g, t) - y'_t}{\sigma^2}.$$

Using $||a||^2 - ||b||^2 \leq (||a|| + ||b||)||a - b||$,

$$\begin{aligned} & \left| \left\| s_u(y_t + \sigma g, t) + \frac{g}{\sigma} \right\|^2 - \left\| s_u(y'_t + \sigma g, t) + \frac{g}{\sigma} \right\|^2 \right| \\ & \leq \frac{4\sqrt{d}}{\sigma^4} [\|u(y_t + \sigma g, t) - u(y'_t + \sigma g, t)\| + \|y_t - y'_t\|] \\ & \leq \frac{4\sqrt{d}}{\sigma^4} (\gamma_{x, R, \varepsilon}^u + \sqrt{d}) \|y_t - y'_t\|_\infty \\ & = \frac{4d}{\sigma^4} \left(1 + \frac{5(d+1)\Lambda_u}{R} \right) \|y_t - y'_t\|_\infty. \end{aligned}$$

where note that

$$\|(y_t + \sigma g) - (y'_t + \sigma g)\|_\infty = \|y_t - y'_t\|_\infty.$$

This proves Equation (64).

Finally, for $u, \tilde{u} \in \mathcal{U}_n$,

$$s_u(y_t + \sigma g, t) + \frac{g}{\sigma} = \frac{u(y_t + \sigma g, t) - y_t}{\sigma^2}, \quad s_{\tilde{u}}(y_t + \sigma g, t) + \frac{g}{\sigma} = \frac{\tilde{u}(y_t + \sigma g, t) - y_t}{\sigma^2}.$$

Therefore,

$$\left| \left\| s_u(y_t + \sigma g, t) + \frac{g}{\sigma} \right\|^2 - \left\| s_{\tilde{u}}(y_t + \sigma g, t) + \frac{g}{\sigma} \right\|^2 \right| \leq \frac{4\sqrt{d}}{\sigma^4} \|u(y_t + \sigma g, t) - \tilde{u}(y_t + \sigma g, t)\|,$$

which proves Equation (65). $\square$

Using the same g in the two evaluation points is valid in the later population comparison because both losses are integrated with respect to the same Gaussian measure. In particular,

$$\begin{aligned} & \mathbb{E}_G \left[\left\| s_u(y_t + \sigma G, t) + \frac{G}{\sigma} \right\|^2 \right] - \mathbb{E}_G \left[\left\| s_u(y'_t + \sigma G, t) + \frac{G}{\sigma} \right\|^2 \right] \\ & = \mathbb{E}_G \left[\left\| s_u(y_t + \sigma G, t) + \frac{G}{\sigma} \right\|^2 - \left\| s_u(y'_t + \sigma G, t) + \frac{G}{\sigma} \right\|^2 \right]. \end{aligned}$$

This is only a common-variable representation of two expectations under the same Gaussian law. It does not impose any additional probabilistic relationship between y_t and y'_t . For each term, the same Gaussian G is used in its noisy evaluation point and its denoising target.

I.2.3 Completion of the Score Bound

Recall that for $u \in \mathcal{U}_n$, we defined

$$\Phi_u^{\text{score}}(y_0, y_1, k, g) = \left\| s_u(y_0 + t_k(y_1 - y_0) + \sigma g, t_k) + \frac{g}{\sigma} \right\|^2. \quad (66)$$

We equip $\mathcal{U}_n$ with

$$d_u(u, \tilde{u}) = \sup_{x \in [-R, R]^d, t \in [0, 1]} \|u(x, t) - \tilde{u}(x, t)\|.$$

Lemma 21 (Constants for the score loss). *The loss in Equation (66) satisfies Equation (20)–Equation (22) with*

$$\begin{aligned} B_{\text{score}} &= \frac{4d}{\sigma^4}, \\ L_{Y, \text{score}} &= \frac{4d}{\sigma^4} \left(1 + \frac{5(d+1)\Lambda_u}{R} \right), \\ L_{\Theta, \text{score}} &= \frac{4\sqrt{d}}{\sigma^4}. \end{aligned} \quad (67)$$

Moreover,

$$\log \mathcal{N}(r, \mathcal{U}_n, d_u) \leq \mathfrak{C}_u(r). \quad (68)$$

Proof. Let $S_t = (1-t)y_0 + ty_1$. The cancellation

$$s_u(S_t + \sigma g, t) + \frac{g}{\sigma} = \frac{u(S_t + \sigma g, t) - S_t}{\sigma^2}$$

and the bounds $u(S_t + \sigma g, t), S_t \in [-1, 1]^d$ give $\Phi_u^{\text{score}} \leq 4d/\sigma^4$.

For a second endpoint pair, write $S'_t = (1-t)y'_0 + ty'_1$. The pointwise bound Equation (64) gives

$$\begin{aligned} & |\Phi_u^{\text{score}}(y_0, y_1, k, g) - \Phi_u^{\text{score}}(y'_0, y'_1, k, g)| \\ & \leq \frac{4d}{\sigma^4} \left(1 + \frac{5(d+1)\Lambda_u}{R} \right) \|S_{t_k} - S'_{t_k}\|_{\infty} \\ & \leq L_{Y, \text{score}} (\|y_0 - y'_0\| + \|y_1 - y'_1\|). \end{aligned}$$

Averaging this inequality over K, G proves the endpoint condition. The parameter perturbation bound Equation (65) gives

$$\left| \Phi_u^{\text{score}}(y_0, y_1, k, g) - \Phi_u^{\text{score}}(y_0, y_1, k, g) \right| \leq \frac{4\sqrt{d}}{\sigma^4} \|u(S_{t_k} + \sigma g, t_k) - \tilde{u}(S_{t_k} + \sigma g, t_k)\|.$$

Every $u \in \mathcal{U}_n$ depends on its spatial argument through clip_R , so the last norm is at most $d_u(u, \tilde{u})$. Finally, Equation (68) follows from Equation (56). $\square$

Proof of the score bound in Proposition 1. Choose

$$\varepsilon = n^{-1/(d+5)}, \quad R = 1 + \sqrt{2} \sigma \left[\log \left(\frac{8d^2}{\sigma^4 \varepsilon^2} \right) \right]^{1/2}.$$

The tail condition in Proposition 7 then holds. Write $\hat{s}_M = s_{\hat{u}_M}$ for the empirical minimizer, with $\hat{u}_M \in \mathcal{U}_n$. By Equation (45) and the approximation proposition,

$$\inf_{u \in \mathcal{U}_n} \left\{ L_{a,Q}^{\text{score}}(s_u) - L_{a,Q}^{\text{score}}(s_Q^*) \right\} \leq 2\varepsilon^2. \quad (69)$$

For Φ_u^{score} , $L_{a,Q}^\Phi(u) = L_{a,Q}^{\text{score}}(s_u)$, and $\hat{L}_{\text{actual}}^\Phi(u)$ is the implemented score objective. Lemma 21 verifies the hypotheses of Proposition 5 with

$$\Theta = \mathcal{U}_n, \quad d_\Theta = d_u, \quad B_\Phi = B_{\text{score}}, \quad L_Y = L_{Y,\text{score}}, \quad L_\Theta = L_{\Theta,\text{score}}.$$

For the selected R ,

$$\begin{aligned} \frac{\Lambda_u}{R} &= \frac{2}{R} + \frac{8d(2R+1)}{R\sigma^2} = O(\sigma^{-2}), \\ \Lambda_u &= \tilde{O}(\sigma^{-2}), \quad \eta_u^{-1} = \frac{\sqrt{d}\Lambda_u}{\sigma^2\varepsilon} = \tilde{O}\left(\frac{1}{\sigma^4\varepsilon}\right). \end{aligned} \quad (70)$$

Substituting the architecture orders Equation (54) into $\mathfrak{C}_u$, at the three radii in Proposition 5, yields

$$\log N_{\text{DR}}^{\text{score}} + \log N_{\text{obs}}^{\text{score}} + \log N_{\text{aux}}^{\text{score}} = \tilde{O}\left[\left(\frac{1}{\sigma^4\varepsilon}\right)^{d+1}\right]. \quad (71)$$

Since $B_{\text{score}} = 4d/\sigma^4$, combining Equation (23) and Equation (71), with $\delta = n^{-1}$, give

$$\mathfrak{R}_{\text{score}}(n^{-1}) \lesssim \frac{1 + 2\pi_{\min}^{-1}}{\sigma^{2d+6}} n^{-2/(d+5)}. \quad (72)$$

Before the substitution $\varepsilon = n^{-1/(d+5)}$, the leading term is

$$\tilde{O}\left((1 + 2\pi_{\min}^{-1})\sigma^{-(2d+6)}\varepsilon^{-(d+1)/2}n^{-1/2}\right).$$

The bounds Equation (67) and Equation (70) give

$$L_{Y,\text{score}} = O(\sigma^{-6}). \quad (73)$$

On the event in Corollary 3,

$$\Delta_Q \lesssim \varepsilon_{\text{OT}}^{-1} \left[n^{-1/(d+2)} + \varepsilon_{\text{pred,train}} + (nM_Q)^{-1/(d+2)} \right].$$

Substitute this display, Equation (69), Equation (72), and Equation (73) into Equation (24). The inequalities $(nM_Q)^{-1/(d+2)} \leq n^{-1/(d+2)}$ and $0 < \sigma \leq 1$ give

$$\begin{aligned} L_{a,Q}^{\text{score}}(\hat{s}_M) - L_{a,Q}^{\text{score}}(s_Q^*) &\lesssim \frac{\varepsilon_{\text{OT}}^{-1}}{\sigma^6} \left[n^{-1/(d+2)} + \varepsilon_{\text{pred,train}} \right] \\ &\quad + \frac{1 + \varepsilon_{\text{OT}}^{-1}}{\sigma^6} \varepsilon_\alpha \varepsilon_{\text{pred}} + \frac{1 + 2\pi_{\min}^{-1}}{\sigma^{2d+6}} n^{-2/(d+5)}. \end{aligned}$$

Finally, Equation (42) converts this excess risk into the grid-averaged score error in Equation (5). The coupling and conditional statistical events fail with probability at most $2/n$ and $1/n$, respectively, while $\mathbb{P}(\mathcal{E}_a^c) \leq e^{-n\pi_{\min}}$. A final union bound with the velocity event proves that both conclusions of Proposition 1 hold with probability $1 - \mathcal{O}(n^{-1}) - e^{-n\pi_{\min}}$. $\square$

J Proof of Theorem 2

Proof. Work on the event where the population coupling Q is defined and the two bounds in Proposition 1 hold. Since $D_5 \leq \overline{D}_5 < \infty$, Lemma 1 supplies a weak solution Z satisfying $Z_t \sim q_t$ for every $t \in [0, 1]$. Moreover, the grid-frozen learned drifts are continuous by continuity of the fitted neural networks. Thus the [standard sampling conditions](#) hold.

Since $h = 1/N$, each grid-averaged field bound in Proposition 1 is exactly the corresponding integrated estimation error in Assumption 1. Thus that proposition may be substituted into Theorem 1. In the stated regime, the velocity part of Proposition 1 gives

$$\varepsilon_{\text{vel}}^2 = \tilde{O} \left(\frac{2 + \varepsilon_{\text{OT}}^{-1}}{\sigma^2} n^{-1/(d+2)} + \frac{1 + 2\pi_{\min}^{-1}}{\sigma^{d+1}} n^{-2/(d+5)} \right),$$

while its score part gives

$$\varepsilon_{\text{score}}^2 = \tilde{O} \left(\frac{1}{\sigma^6 \varepsilon_{\text{OT}}} n^{-1/(d+2)} + \frac{1 + 2\pi_{\min}^{-1}}{\sigma^{2d+6}} n^{-2/(d+5)} \right).$$

Since $\tau^2 = \sigma \overline{D}_5 / d$,

$$\frac{\varepsilon_{\text{vel}}^2}{\tau^2} = \tilde{O} \left(\frac{d(2 + \varepsilon_{\text{OT}}^{-1})}{\sigma^3 \overline{D}_5} n^{-1/(d+2)} + \frac{d(1 + 2\pi_{\min}^{-1})}{\sigma^{d+2} \overline{D}_5} n^{-2/(d+5)} \right), \quad (74)$$

and

$$\tau^2 \varepsilon_{\text{score}}^2 = \tilde{O} \left(\frac{\overline{D}_5}{d \varepsilon_{\text{OT}} \sigma^5} n^{-1/(d+2)} + \frac{(1 + 2\pi_{\min}^{-1}) \overline{D}_5}{d \sigma^{2d+5}} n^{-2/(d+5)} \right). \quad (75)$$

By Corollary 1, under $\frac{h\tau^2}{2\sigma^2} \leq 1$ and $\sigma \leq 1$,

$$W_1 \left(\int \mathbb{P}_{Y|C=c, A=a} d\mathbb{P}_C(c), \mathcal{L}(\widehat{Z}_N) \right) \lesssim \sigma + \text{KL} \left(\mathcal{L}(Y_1 + \sigma G) \parallel \mathcal{L}(\widehat{Z}_N) \right)^{1/4}.$$

Taking $\sigma = \Theta(\varepsilon_{\text{target}})$, it is sufficient to require

$$\text{KL} \left(\mathcal{L}(Y_1 + \sigma G) \parallel \mathcal{L}(\widehat{Z}_N) \right) \lesssim \varepsilon_{\text{target}}^4.$$

To make each estimation contribution in Equation (74) and Equation (75) at most of order $\varepsilon_{\text{target}}^4$, it is sufficient that

$$\begin{aligned} n &\gtrsim \left[\frac{(2 + \varepsilon_{\text{OT}}^{-1})}{\overline{D}_5} \right]^{d+2} \varepsilon_{\text{target}}^{-7(d+2)}, & n &\gtrsim \left[\frac{(1 + 2\pi_{\min}^{-1})}{\overline{D}_5} \right]^{(d+5)/2} \varepsilon_{\text{target}}^{-(d+6)(d+5)/2}, \\ n &\gtrsim \left[\frac{\overline{D}_5}{\varepsilon_{\text{OT}}} \right]^{d+2} \varepsilon_{\text{target}}^{-9(d+2)}, & n &\gtrsim \left[(1 + 2\pi_{\min}^{-1}) \overline{D}_5 \right]^{(d+5)/2} \varepsilon_{\text{target}}^{-(2d+9)(d+5)/2}, \end{aligned}$$

where $\sigma = \Theta(\varepsilon_{\text{target}})$ has been substituted. Note that the last condition determines the asymptotic dependence on $\varepsilon_{\text{target}}$.

It remains to control the discretization terms in Theorem 1. The choice

$$h \leq c_h \left(\frac{d\overline{D_5}^2}{\varepsilon_{\text{target}}^4 \sigma^2} + \frac{d\overline{D_5}}{\sigma} \right)^{-1}$$

implies $\frac{hd\overline{D_5}^2}{\sigma^2} \lesssim \varepsilon_{\text{target}}^4$ and $\frac{hd\overline{D_5}}{\sigma} \lesssim 1$. Consequently, $\text{KL}(\mathcal{L}(Z_{t_N}) \parallel \mathcal{L}(\hat{Z}_N)) \lesssim \varepsilon_{\text{target}}^4$. Moreover,

$$\frac{h\tau^2}{2\sigma^2} = \frac{h\overline{D_5}}{2d\sigma} \leq \frac{1}{2d^2} \frac{hd\overline{D_5}}{\sigma} = \frac{h\overline{D_5}}{2d\sigma} \leq c_h \frac{\frac{\overline{D_5}}{2d\sigma}}{\frac{d\overline{D_5}^2}{\varepsilon_{\text{target}}^4 \sigma^2} + \frac{d\overline{D_5}}{\sigma}} \leq \frac{c_h}{2d^2} \leq 1,$$

so the required stability condition holds. □