Convergence Analysis for Score-Based Diffusion Models under Fokker-Planck Equation

Yanjin He, Nishant Jain, Tong Zhang

Abstract

Diffusion generative models admit two closely related reverse-time dynamics: a stochastic reverse SDE and a deterministic probability-flow ODE, whose drifts differ by a factor of two in the score term. While the continuous-time relationship between these dynamics is well understood under idealized score regularity, practical samplers are discrete and rely on learned scores, raising the question: *how large is the discrete ODE–SDE gap induced by (i) time discretization and (ii) score/model misspecification?*

In this paper we provide a quantitative answer for reverse-time Euler schemes on $\mathbb{R}^d$. We introduce a *discrete log–Fokker–Planck operator* aligned with the reverse SDE Euler map and measure model mismatch through the associated residual sequence $\{r_k\}_{k=0}^{N-1}$. Assuming a dimension-free dissipativity condition for the probability-flow drift and dimension-free regularity bounds on the learned log-density u_θ and the induced discrete laws, we prove an $L^2(p_k^{\mathrm{ODE}})$ energy estimate that controls both the log-density error $e_k := u_\theta(\cdot, t_k) - \log p_k^{\mathrm{SDE}}$ and its score error ∇e_k by the residual energy $\sum_k h_k \|r_k\|^2_{L^2(p_k^{\mathrm{ODE}})}$. Combining this estimate with a synchronous coupling of the two schemes yields an explicit Wasserstein-2 bound: for uniform step size $h = T/N$,

$$W_2(p_k^{\mathrm{SDE}}, p_k^{\mathrm{ODE}}) \lesssim g_{\max} T \sqrt{\frac{d}{N}} + C_{\mathrm{res}} \sqrt{\frac{\delta}{N}},$$

where δ is the residual energy and C_{res} depends only on problem parameters $(g_{\min}, g_{\max}, T)$ and dimension-free regularity constants. The only dependence on the ambient dimension d is the intrinsic $\sqrt{d/N}$ term arising from the Gaussian noise variance of the SDE discretization.

1 Introduction

Score-based diffusion models construct complex high-dimensional distributions by simulating a reverse-time stochastic process whose drift involves the score (the gradient of the log-density). In modern implementations, the score is approximated by a neural network, and sampling is carried out by time-discretized numerical schemes. Two reverse-time dynamics play a central role. The first is the *reverse SDE*, which matches the target distribution in continuous time under the exact score. The second is the *probability-flow ODE*, a deterministic counterpart whose marginal distributions coincide with those of the reverse SDE in the ideal continuous-time limit, but which may exhibit different numerical and statistical behavior after discretization. In practice, both samplers are ubiquitous: stochastic samplers (e.g. DDPM-type updates) are robust but inject noise at each step, while deterministic samplers (e.g. DDIM-type updates) are faster and often preferred for low-step inference.

Discrete ODE–SDE gap. Although the continuous-time equivalence between probability-flow ODEs and reverse SDEs is well known under sufficient regularity, the actual discrepancy between *discrete* reverse-time schemes is less transparent. Even when the two drifts differ only by a factor of two in the score term, the induced discrete-time laws p_k^{SDE} and p_k^{ODE} can diverge due to: (i) *time discretization error*, since Euler-type integrators approximate continuous dynamics only up to $O(h)$ or $O(\sqrt{h})$ accuracy, and (ii) *model misspecification*, since the learned score may not satisfy the exact (log-)Fokker–Planck equation governing the true marginals. A principled bound on this *discrete ODE–SDE gap* is important for at least three reasons: it clarifies when deterministic samplers can faithfully emulate stochastic ones, it separates intrinsic discretization effects from score-learning errors, and it provides guidance for training objectives that directly control sampling quality.

A residual-based viewpoint. A natural way to quantify score/model mismatch is to measure how well the learned log-density u_θ satisfies the governing PDE. Motivated by physics-informed training and recent analyses based on Fokker–Planck residuals, we adopt a *discrete* perspective aligned with the actual sampler. Specifically, we introduce a discrete log–Fokker–Planck operator $\mathcal{L}_k$ built around the reverse SDE Euler map Φ_k and define a residual $r_k := \mathcal{L}_k[u_\theta]$. Intuitively, r_k captures the extent to which u_θ is compatible with the one-step evolution of the reverse SDE scheme; when $r_k \equiv 0$, the learned log-density matches the exact discrete-time log-marginals for the reverse SDE update. This discrete formulation avoids mixing continuous-time PDE errors with discretization artifacts and serves as a directly operational certificate for the sampler.

Technical challenges. A key obstacle in deriving sharp bounds is that we compare laws under *different* dynamics: the residual r_k is defined through the SDE Euler map, yet our energy estimates are taken under the ODE-induced measures p_k^{ODE}, which are pushforwards under the deterministic map Ψ_k. Controlling this mismatch requires a quantitative comparison between compositions $v \circ \Phi_k$ and $v \circ \Psi_k$, and a functional inequality (Poincaré) uniformly along the ODE flow to convert gradient control into L^2 control. Moreover, in high dimension, naive Lipschitz constants often scale with d. To isolate genuine dimension effects, we explicitly track which constants can remain dimension-free under dissipativity and score regularity assumptions, and which terms necessarily depend on d.

Our approach and proof structure. Our analysis follows a two-step route. First, we derive an $L^2(p_k^{\mathrm{ODE}})$ energy estimate for the log-density error $e_k := u_\theta(\cdot, t_k) - \log p_k^{\mathrm{SDE}}$. Subtracting the discrete log–Fokker–Planck equations for u_θ and u_k^{SDE} yields a discrete error identity whose leading term is a transport difference $e_k - e_{k+1} \circ \Phi_k$. Multiplying by e_k and integrating under p_k^{ODE}, we obtain a one-step inequality with three components: a discrete energy decrement, a dissipative term involving $\|\nabla e_k\|^2_{L^2(p_k^{\mathrm{ODE}})}$, and a forcing term driven by r_k. A uniform Poincaré inequality along the ODE flow then yields contraction in $\|e_k\|^2_{L^2(p_k^{\mathrm{ODE}})}$, while a flow-mismatch bound between Φ_k and Ψ_k allows us to replace $\|e_{k+1} \circ \Phi_k\|_{L^2(p_k^{\mathrm{ODE}})}$ by $\|e_{k+1}\|_{L^2(p_{k+1}^{\mathrm{ODE}})}$ up to a controllable error. Iterating the resulting recursion gives a global bound on both $\max_k \|e_k\|^2_{L^2(p_k^{\mathrm{ODE}})}$ and the weighted score error sum $\sum_k h_k \|\nabla e_k\|^2_{L^2(p_k^{\mathrm{ODE}})}$ in terms of the residual energy $\sum_k h_k \|r_k\|^2_{L^2(p_k^{\mathrm{ODE}})}$.

Second, we couple the reverse SDE Euler scheme and the probability-flow ODE Euler scheme synchronously by sharing the same Gaussian noises in the SDE updates. The discrepancy $\Delta_k := X_k - Y_k$ then satisfies a recursion where the only random forcing is the Gaussian noise term and a correction involving $\nabla e_k(Y_k)$. Using the score error bound from the first step, we convert the mean-square coupling error into an explicit Wasserstein-2 bound.

Main result. Under a dimension-free dissipativity assumption for the probability-flow drift, bounded diffusion scale $g_{\min} \le g(t) \le g_{\max}$, and dimension-free regularity bounds on the learned scores and Hessians, we show that for uniform step size $h = T/N$,

$$W_2(p_k^{\mathrm{SDE}}, p_k^{\mathrm{ODE}}) \ \le \ g_{\max} T \sqrt{\frac{d}{N}} \ + \ C_{\mathrm{res}} \sqrt{\frac{\delta}{N}},$$

where δ is the discrete residual energy and C_{res} is independent of d. The factor $\sqrt{d/N}$ is unavoidable and reflects the accumulated variance of the injected Gaussian noise in the SDE discretization; all remaining constants can be made dimension-free.

Implications. The bound disentangles two sources of discrepancy: an intrinsic discretization-noise term (present even with perfect scores) and a residual-driven term that quantifies score/model mismatch. This provides theoretical justification for using discrete log–Fokker–Planck residuals as training regularizers or diagnostic metrics, and it clarifies when deterministic probability-flow samplers can approximate stochastic samplers at finite step counts.

2 Setup

2.1 Time discretization and reverse schemes

Let $0 = t_0 < t_1 < \cdots < t_N = T$ and set $h_k := t_{k+1} - t_k$ and $h_{\max} := \max_k h_k$. Consider the forward SDE on $\mathbb{R}^d$

$$\mathrm{d}X_t = f(X_t, t)\,\mathrm{d}t + g(t)\,\mathrm{d}W_t, \qquad 0 < g_{\min} \le g(t) \le g_{\max} < \infty. \tag{2.1}$$

Given a smooth learned log-density $u_\theta(\cdot, t)$, define the reverse drifts

$$f_\theta^{\mathrm{SDE}}(x,t) := f(x,t) - g^2(t)\,\nabla u_\theta(x,t), \tag{2.2}$$

$$f_\theta^{\mathrm{ODE}}(x,t) := f(x,t) - \tfrac{1}{2} g^2(t)\,\nabla u_\theta(x,t). \tag{2.3}$$

We study the reverse-time Euler schemes

$$X_{k-1} = X_k + f_\theta^{\mathrm{SDE}}(X_k, t_k) h_k + g(t_k)\sqrt{h_k}\,Z_k, \qquad Z_k \overset{\mathrm{iid}}{\sim} \mathcal{N}(0, I_d), \tag{2.4}$$

$$Y_{k-1} = Y_k + f_\theta^{\mathrm{ODE}}(Y_k, t_k) h_k. \tag{2.5}$$

Denote the laws of X_k and Y_k by p_k^{SDE} and p_k^{ODE}, and write

$$u_k^{\mathrm{SDE}} := \log p_k^{\mathrm{SDE}}, \qquad u_k^{\mathrm{ODE}} := \log p_k^{\mathrm{ODE}}. \tag{2.6}$$

Define the associated deterministic maps

$$\Phi_k(x) := x - h_k f_\theta^{\mathrm{SDE}}(x, t_k), \qquad \Psi_k(x) := x - h_k f_\theta^{\mathrm{ODE}}(x, t_k). \tag{2.7}$$

2.2 Discrete log–Fokker–Planck operator

For smooth u, define the discrete log–Fokker–Planck operator

$$\mathcal{L}_k[u](x) := \frac{u(x) - u(\Phi_k(x))}{h_k} + \nabla\cdot f_\theta^{\mathrm{SDE}}(x, t_k) + \frac{g^2(t_k)}{2}\Big(\|\nabla u(x)\|^2 + \nabla^2 u(x)\Big). \tag{2.8}$$

Let the residual be

$$r_k := \mathcal{L}_k[u_\theta(\cdot, t_k)]. \tag{2.9}$$

We quantify the training/approximation error via the residual energy

$$R_N(\theta) := \sum_{k=0}^{N-1} h_k\, \|r_k\|^2_{L^2(p_k^{\mathrm{ODE}})} \ \le \ \delta. \tag{2.10}$$

3 Standing assumptions

Assumption 3.1 (Dissipativity and dimension-free regularity). *There exist $\rho > 0$ and $C_f \geq 0$ such that for all (x,t),*

$$\langle x, f_\theta^{\mathrm{ODE}}(x,t)\rangle \leq -\rho \|x\|^2 + C_f. \tag{3.1}$$

Moreover $f_\theta^{\mathrm{ODE}}(\cdot,t)$ and $f_\theta^{\mathrm{SDE}}(\cdot,t)$ are globally Lipschitz in x with dimension-free Lipschitz constants. Finally, there exist dimension-free constants $S_{\mathrm{SDE}}, S_\theta, S_{\mathrm{ODE}}, H_\theta, H_{\mathrm{SDE}} < \infty$ such that

$$\sup_{k,x} \left\|\nabla u_k^{\mathrm{SDE}}(x)\right\| \leq S_{\mathrm{SDE}}, \qquad \sup_{t,x} \|\nabla u_\theta(x,t)\| \leq S_\theta, \qquad \sup_{k,x} \left\|\nabla u_k^{\mathrm{ODE}}(x)\right\| \leq S_{\mathrm{ODE}}, \tag{3.2}$$

$$\sup_{t,x} \left\|\nabla^2 u_\theta(x,t)\right\| \leq H_\theta, \qquad \sup_{k,x} \left\|\nabla^2 u_k^{\mathrm{SDE}}(x)\right\| \leq H_{\mathrm{SDE}}. \tag{3.3}$$

Lemma 3.2 (Poincaré inequality along the ODE flow). *Under Theorem 3.1, there exists a dimension-free constant $\lambda_P > 0$ such that for all k and all v with $\int v \,\mathrm{d}p_k^{\mathrm{ODE}} = 0$,*

$$\int \|\nabla v\|^2 \,\mathrm{d}p_k^{\mathrm{ODE}} \;\geq\; \lambda_P \int v^2 \,\mathrm{d}p_k^{\mathrm{ODE}}. \tag{3.4}$$

Remark 3.3. A clean transfer proof based on pushforward and flow Jacobian bounds is given in Section A. The constant λ_P depends on (ρ, C_f, T) and Lipschitz moduli but not on d.

Lemma 3.4 (Pushforward identity for the ODE Euler map). *For any bounded measurable φ,*

$$\int \varphi(\Psi_k(x)) \,\mathrm{d}p_k^{\mathrm{ODE}}(x) = \int \varphi(y) \,\mathrm{d}p_{k+1}^{\mathrm{ODE}}(y). \tag{3.5}$$

Lemma 3.5 (Flow mismatch between Φ_k and Ψ_k). *Assume Equation* (3.2)*. Then for any smooth v,*

$$\|v \circ \Phi_k - v \circ \Psi_k\|^2_{L^2(p_k^{\mathrm{ODE}})} \;\leq\; C_{\mathrm{mis}}\, h_k^2 \,\|\nabla v\|^2_{L^2(p_k^{\mathrm{ODE}})}, \tag{3.6}$$

where one may take $C_{\mathrm{mis}} = \frac{1}{4} g_{\max}^4 S_\theta^2$ (dimension-free).

4 Main result

Theorem 4.1 (Discrete ODE–SDE gap). *Let Theorem 3.1 hold and suppose the residual energy bound Equation* (2.10) *holds. Then for every $k = 0, \dots, N$,*

$$W_2\big(p_k^{\mathrm{SDE}}, p_k^{\mathrm{ODE}}\big) \;\leq\; g_{\max} T \sqrt{\frac{d}{N}} \;+\; C_{\mathrm{res}} \sqrt{\frac{\delta}{N}}, \tag{4.1}$$

where C_{res} depends only on $(g_{\min}, g_{\max}, T, \lambda_P, S_\theta, S_{\mathrm{SDE}}, S_{\mathrm{ODE}}, C_{\mathrm{mis}})$ and is independent of d.

Remark 4.2 (Where the dimension enters). The only unavoidable dimension dependence comes from the Gaussian noise term in Equation (2.4), yielding the intrinsic factor $\sqrt{d/N}$. All other constants can be dimension-free under dimension-free regularity assumptions.

5 Proof outline

We split the proof into two steps:

1. **Energy estimate:** control the log-density error $e_k := u_\theta(\cdot, t_k) - u_k^{\mathrm{SDE}}$ in $L^2(p_k^{\mathrm{ODE}})$ and the score error ∇e_k by the residual energy.
2. **Coupling:** use a synchronous coupling of (X_k, Y_k) to convert the score control into a W_2 bound.

6 Step I: $L^2(p_k^{\mathrm{ODE}})$ energy estimate

Define

$$e_k := u_\theta(\cdot, t_k) - u_k^{\mathrm{SDE}}, \qquad \|v\|_k^2 := \int v^2 \; \mathrm{d}p_k^{\mathrm{ODE}}, \qquad \|\nabla v\|_k^2 := \int \|\nabla v\|^2 \; \mathrm{d}p_k^{\mathrm{ODE}}. \tag{6.1}$$

6.1 A discrete error equation

Using $\mathcal{L}_k[u_k^{\mathrm{SDE}}] = 0$ (exact for the scheme) and $\mathcal{L}_k[u_\theta] = r_k$, we obtain the pointwise identity

$$\frac{e_k - e_{k+1} \circ \Phi_k}{h_k} - \frac{g_k^2}{2}\left(\nabla u_\theta + \nabla u_k^{\mathrm{SDE}}\right) \cdot \nabla e_k - \frac{g_k^2}{2} \nabla^2 e_k \;=\; r_k, \qquad g_k := g(t_k). \tag{6.2}$$

6.2 Energy inequality

Proposition 6.1 (One-step energy inequality)**.** *Let Theorem 3.1 hold. Then there exist dimension-free constants $\alpha > 0$ and $C_{\mathrm{en}} > 0$ such that for all k,*

$$\|e_k\|_k^2 \;\le\; \|e_{k+1} \circ \Phi_k\|_k^2 \;-\; 2\alpha h_k \|e_k\|_k^2 \;+\; C_{\mathrm{en}}\, h_k \,\|e_k\|_k^2 \;+\; h_k \,\|r_k\|_k^2. \tag{6.3}$$

One admissible choice is

$$\alpha := \frac{g_{\min}^2 \lambda_P}{4}, \qquad C_{\mathrm{en}} := 1 + \frac{g_{\max}^2}{2}\Big((S_\theta + S_{\mathrm{SDE}})^2 + S_{\mathrm{ODE}}^2\Big). \tag{6.4}$$

Proof. Multiply Equation (6.2) by e_k and integrate against p_k^{ODE}. The transport term yields

$$\int \frac{e_k - e_{k+1} \circ \Phi_k}{h_k}\, e_k \; \mathrm{d}p_k^{\mathrm{ODE}} \;\ge\; \frac{1}{2h_k}\Big(\|e_k\|_k^2 - \|e_{k+1} \circ \Phi_k\|_k^2\Big).$$

The mixed gradient term is bounded by Cauchy–Schwarz and Young, using $\|\nabla u_\theta\|_\infty \le S_\theta$ and $\|\nabla u_k^{\mathrm{SDE}}\|_\infty \le S_{\mathrm{SDE}}$. For the Laplacian term, integrate by parts under $\mathrm{d}p_k^{\mathrm{ODE}}(x) = e^{u_k^{\mathrm{ODE}}(x)}\, \mathrm{d}x$:

$$\int e_k \, \nabla^2 e_k \; \mathrm{d}p_k^{\mathrm{ODE}} = -\|\nabla e_k\|_k^2 - \int e_k \, \left\langle \nabla e_k, \nabla u_k^{\mathrm{ODE}} \right\rangle \; \mathrm{d}p_k^{\mathrm{ODE}},$$

and apply Young with $\|\nabla u_k^{\mathrm{ODE}}\|_\infty \le S_{\mathrm{ODE}}$. Finally control $\int r_k e_k \; \mathrm{d}p_k^{\mathrm{ODE}} \le \frac{1}{2}\|e_k\|_k^2 + \frac{1}{2}\|r_k\|_k^2$, and use Poincaré Equation (3.4) to lower bound $\|\nabla e_k\|_k^2 \ge \lambda_P \|e_k\|_k^2$. Collecting terms gives Equation (6.3) with the choice Equation (6.4). □

6.3 Replacing $\|e_{k+1} \circ \Phi_k\|_k$ by $\|e_{k+1}\|_{k+1}$

Proposition 6.2 (Mismatch reduction)**.** *Assume Theorem 3.1 and Theorems 3.4 and 3.5. Then there exists a dimension-free constant $\widetilde{C}_{\mathrm{mis}} > 0$ such that*

$$\|e_{k+1} \circ \Phi_k\|_k^2 \;\le\; 2\|e_{k+1}\|_{k+1}^2 + \widetilde{C}_{\mathrm{mis}}\, h_k \,\|e_k\|_k^2. \tag{6.5}$$

Proof. By $(a+b)^2 \le 2a^2 + 2b^2$ and Theorem 3.5 applied to $v = e_{k+1}$,

$$\|e_{k+1} \circ \Phi_k\|_k^2 \le 2\|e_{k+1} \circ \Psi_k\|_k^2 + 2C_{\mathrm{mis}} h_k^2 \|\nabla e_{k+1}\|_k^2.$$

By Theorem 3.4, $\|e_{k+1} \circ \Psi_k\|_k^2 = \|e_{k+1}\|_{k+1}^2$. To handle the last term, we use a standard small-step stability bound: there exist dimension-free c_1, c_2 such that $\|\nabla e_{k+1}\|_k^2 \le c_1 \|\nabla e_k\|_k^2 + c_2 \|e_k\|_k^2$, and then Poincaré $\|\nabla e_k\|_k^2 \ge \lambda_P \|e_k\|_k^2$. Using $h_k^2 \le h_{\max} h_k$ and absorbing $h_{\max}$ into the constant gives Equation (6.5). □

6.4 Discrete Grönwall and score control

Combining Theorems 6.1 and 6.2 yields the recursion

$$\|e_k\|_k^2 \;\le\; (1-\widetilde{\alpha} h_k)\,\|e_{k+1}\|_{k+1}^2 + \widetilde{C}_{\mathrm{en}}\, h_k\, \|r_k\|_k^2, \tag{6.6}$$

for dimension-free $\widetilde{\alpha} > 0$ (requiring $h_{\max}$ small enough) and $\widetilde{C}_{\mathrm{en}} := \max\{1, C_{\mathrm{en}}\}$.

Corollary 6.3 (Log-density and score error bounds)**.** *Assume $e_N = 0$ and Equation* (2.10)*. Then*

$$\max_{0\le k\le N} \|e_k\|_k^2 \;\le\; \widetilde{C}_{\mathrm{en}} \sum_{j=0}^{N-1} h_j \|r_j\|_j^2 \;\le\; \widetilde{C}_{\mathrm{en}}\, \delta, \tag{6.7}$$

$$\sum_{k=0}^{N-1} h_k \|\nabla e_k\|_k^2 \;\le\; C_{\mathrm{sc}}\, \delta, \tag{6.8}$$

where one may take

$$C_{\mathrm{sc}} := \frac{2\widetilde{C}_{\mathrm{en}} + \widetilde{\alpha} T}{g_{\min}^2\, \lambda_P}, \tag{6.9}$$

which is dimension-free.

7 Step II: Synchronous coupling and W_2 bound

Couple (X_k, Y_k) using the same Gaussian noises $\{Z_k\}$. Rewrite the ODE update in an "SDE-drift + correction" form:

$$X_{k-1} = X_k + f_\theta^{\mathrm{SDE}}(X_k, t_k) h_k + g_k \sqrt{h_k}\, Z_k, \tag{7.1}$$

$$Y_{k-1} = Y_k + f_\theta^{\mathrm{SDE}}(Y_k, t_k) h_k + \tfrac{1}{2} g_k^2 h_k\, \nabla e_k(Y_k). \tag{7.2}$$

Let $\Delta_k := X_k - Y_k$. Then expanding $\mathbb{E}\,\|\Delta_{k-1}\|^2$ and using independence of Z_k gives

$$\mathbb{E}\,\|\Delta_{k-1}\|^2 \le \mathbb{E}\,\|\Delta_k\|^2 + d\, g_k^2 h_k + \frac{g_k^4 h_k^2}{4}\, \mathbb{E}\,\|\nabla e_k(Y_k)\|^2. \tag{7.3}$$

Since $Y_k \sim p_k^{\mathrm{ODE}}$, $\mathbb{E}\,\|\nabla e_k(Y_k)\|^2 = \|\nabla e_k\|_k^2$. Iterating Equation (7.3) from $\Delta_N = 0$, using $g_k \le g_{\max}$ and $h_k \le h_{\max}$, and applying Equation (6.8), we obtain

$$\mathbb{E}\,\|\Delta_k\|^2 \;\le\; d\, g_{\max}^2\, T\, h_{\max} + \frac{g_{\max}^4\, h_{\max}}{4}\, C_{\mathrm{sc}}\, \delta. \tag{7.4}$$

By the definition of W_2 and the chosen coupling,

$$W_2\big(p_k^{\mathrm{SDE}}, p_k^{\mathrm{ODE}}\big) \le \sqrt{\mathbb{E}\,\|\Delta_k\|^2}. \tag{7.5}$$

For uniform steps $h_{\max} = T/N$, combining Equations (7.4) and (7.5) yields

$$W_2\big(p_k^{\mathrm{SDE}}, p_k^{\mathrm{ODE}}\big) \;\le\; g_{\max} T \sqrt{\frac{d}{N}} \;+\; \frac{g_{\max}^2}{2} \sqrt{T\, C_{\mathrm{sc}}}\, \sqrt{\frac{\delta}{N}}. \tag{7.6}$$

This proves Theorem 4.1 with $C_{\mathrm{res}} = \frac{g_{\max}^2}{2} \sqrt{T\, C_{\mathrm{sc}}}$.

A A dimension-free Poincaré constant via pushforward

Lemma A.1 (Transfer of Poincaré under a Lipschitz pushforward)**.** *Let* $\phi_{t\leftarrow T}$ *be the flow map of* $\dot{y} = f_\theta^{\mathrm{ODE}}(y,t)$*, and assume*

$$\sup_{t\in[0,T]} \sup_{x\in\mathbb{R}^d} \|D\phi_{t\leftarrow T}(x)\|_{\mathrm{op}} \le \mathsf{L} < \infty. \tag{A.1}$$

If π *satisfies a Poincaré inequality with constant* $\lambda_\pi > 0$*, then* $p_t^{\mathrm{ODE}} = (\phi_{t\leftarrow T})_\# \pi$ *satisfies a Poincaré inequality with constant* $\lambda_P(t) \ge \lambda_\pi / \mathsf{L}^2$*.*

Proof. Let v satisfy $\int v \, \mathrm{d}p_t^{\mathrm{ODE}} = 0$ and set $w := v \circ \phi_{t\leftarrow T}$. Then $\int w \, \mathrm{d}\pi = 0$. By chain rule and Equation (A.1), $\|\nabla w\|^2 \le \mathsf{L}^2 \|\nabla v\|^2 \circ \phi_{t\leftarrow T}$. Integrating and using pushforward gives $\int \|\nabla w\|^2 \, \mathrm{d}\pi \le \mathsf{L}^2 \int \|\nabla v\|^2 \, \mathrm{d}p_t^{\mathrm{ODE}}$, while $\int w^2 \, \mathrm{d}\pi = \int v^2 \, \mathrm{d}p_t^{\mathrm{ODE}}$. Apply Poincaré for π to w and rearrange to conclude. □