RÉNYI SAMPLING: NONLINEAR WASSERSTEIN CORRECTIONS FOR DIFFUSION MODELS

Anonymous authors
Paper under double-blind review

ABSTRACT

We introduce Rényi sampling, nonlinear corrections for pretrained denoising diffusion probabilistic models that require no diffusion-network retraining or semantic class guidance. Our construction generalizes Langevin correction from the Wasserstein gradient flow of Kullback–Leibler divergence to that of Rényi divergence. The stochastic corrector combines a pretrained score with an estimated density ratio, with the Rényi order α controlling local correction strength. We establish a moving-target dissipation identity showing that Rényi sampling reduces distributional mismatch throughout denoising while preserving the intended noisy-data marginals. We show that $\alpha < 1$ helps to rectify mode deficits, while $\alpha > 1$ removes excess mode representations. Relative to Langevin correction, $\alpha = 0.7$ recovers $11 - 15\%$ of the rare-mode deficit, and $\alpha = 1.3$ removes $25 - 46\%$ of excess rare mass for imbalanced Gaussian mixtures experiments. Extending our method to a face generation model trained on a 5% cat-95% human data mixture, correction increases the generated cat proportion from 3.8% under ancestral sampling to 4.6%, while improving cat Fréchet Inception Distance from 47.13 to 45.16. These results demonstrate the potential of Rényi correction to improve both distributional fidelity and generation quality during sampling. Our code is available anonymously at `https://github.com/Anonymous-PaperSubmission/RenyiSamplingDDPM`.

1 INTRODUCTION

Denoising diffusion probabilistic models (DDPMs) and score-based generative models (Ho et al., 2020; Song et al., 2021) generate samples by progressively reversing a noise-corruption process. Faithful generation requires recovering both the structure of the data and the probability mass assigned to its different modes. This can be challenging for imbalanced datasets, where underrepresented regions receive limited coverage in generated samples (Sehwag et al., 2022). We investigate whether modifying the sampling dynamics of a pretrained diffusion model can help recover these modes while retaining the data distribution as the intended target.

Guided sampling offers a way to steer generation toward specified classes or conditions. Classifier guidance uses gradients from a classifier (Dhariwal & Nichol, 2021), while classifier-free guidance (CFG) combines conditional and unconditional predictions from a diffusion model trained to support such conditioning (Ho & Salimans, 2022). For a DDPM trained on an arbitrary dataset, a suitable classifier or conditioning capability may be unavailable. Moreover, minority modes may be unknown in advance, and identifying or labeling them can be difficult or expensive. These settings motivate the question: *Can a sampling correction improve the faithful generation of minority modes without requiring their identities, class labels, or explicit guidance toward them?* We consider this question with access to a pretrained diffusion model and unlabeled reference data.

We introduce *Rényi sampling correction*, a family of nonlinear corrections within the predictor–corrector sampling framework (Song et al., 2021). At a fixed noise level t, standard Langevin correction follows the Wasserstein-2 gradient flow of the Kullback–Leibler divergence to the noisy-data target π_t. We replace this correction energy with the more general Rényi divergence of order α, keeping the same Wasserstein geometry. The order α controls a normalized mobility coefficient that depends on the ratio between the current sampling density and π_t, thereby adjusting the correction strength. Different Rényi orders can be explored using the same pre-trained checkpoint. At

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

$\alpha = 1$, where Rényi divergence reduces to KL divergence, we recover standard Langevin correction. Practically, Rényi correction requires only an auxiliary discriminator, whose logits provide the density ratio estimates between the sample distribution and target distribution. No semantic class annotations or enumeration of minority modes is needed.

Our analysis connects this construction to distributional error. Fixed-noise and moving-target dissipation identities show that ideal continuous-time corrections reduce mismatch while preserving the intended denoising marginals. Shared probability-flow transport leaves the Rényi divergence unchanged, allowing us to isolate the corrector's additional dissipation. We quantify how score, log-ratio, and normalization errors affect this balance, motivating empirical evaluation with estimated coefficients and finite sampling budgets.

Experiments span imbalanced Gaussian mixtures, stacked MNIST, and cat–human faces. In GMMs, smaller orders $\alpha < 1$ recover depleted rare mass, while $\alpha > 1$ can remove excess mass for a model deliberately trained to induce overrepresentation of the rare modes, nudging the sampling distribution closer to the true data distribution. We fit the local response relating the sensitivity of log mode-mass ratios to their baseline discrepancies, obtained absolute Pearson correlations of approximately 0.87 in both regimes. Extending the method to image generation, we observe similar improvements for distributional fidelity and generation quality. Correction lowers confidence-filtered total variation (TV) on stacked MNIST, reduces the deficit in generated cat frequency, and improves cat Fr'echet Inception Distance (FID).

Our contributions are threefold:

1. **A Rényi sampling corrector for pretrained DDPMs.** We derive a nonlinear correction from the Wasserstein-2 gradient flow of Rényi divergence and implement it using pretrained scores and discriminator-based density-ratio estimates. The method requires no diffusion-network retraining or semantic class labels and recovers Langevin correction at order one.
2. **Dissipation analysis.** We establish fixed-noise and moving-target dissipation identities that characterize how ideal corrections reduce mismatch while preserving the intended marginals, and derive bounds describing the effects of estimated coefficients.
3. **A controlled study of minority-mode recovery.** We investigate distributional fidelity and generation quality on imbalanced mixtures, to assess how the Rényi order and correction budget affect sampling without explicit class guidance.

2 Related Works

Wasserstein gradient flows in diffusion models. Wasserstein gradient flows provide a variational interpretation of diffusion dynamics through the evolution of probability measures under an energy functional (Jordan et al., 1998; Otto, 2001). Recent work interprets diffusion training as Wasserstein gradient-flow matching (Vuong et al., 2025), while interacting-particle methods use Wasserstein geometry to optimize latent distributions jointly with diffusion-model parameters (Wang et al., 2026). On the sampling side, diffusion-regularized gradient flows have been developed for posterior sampling in inverse problems (Wang & Akyildiz, 2025).

Sampling, guidance, and distribution shaping. Conditional diffusion models are typically trained with classifier guidance (Dhariwal & Nichol, 2021) or classifier-free guidance (Ho & Salimans, 2022). At inference time, classifier guidance modifies the denoising direction using classifier gradients, while classifier-free guidance amplifies the difference between conditional and unconditional denoising predictions. Discriminator guidance (Kim et al., 2023) refines a pretrained diffusion model via the discrepancy between the model score and the data score. Temperature scaling (Shih et al., 2023; Xu et al., 2026) is another sampling approach that reshapes the distribution to be sharper or broader according to a temperature hyperparameter. Whereas guidance and temperature scaling sampling typically modify the target distribution, our approach changes the transport dynamics of Langevin corrections while retaining the prescribed target as a stationary density.

Training and Sampling for Imbalanced Data Imbalanced class distributions pose challenges for both training and sampling. For training (Qin et al., 2023) introduce class-balancing in the form

of a distribution aware regularizer. (Rodriguez et al., 2025) employs supervised contrastive regularization to reduce interference between head and tail classes while (Chen et al., 2025) combines an unsupervised contrastive loss on unconditional latent representations with an alignment loss between conditional and unconditional noise predictions. On the sampling side, (Sehwag et al., 2022) construct a hardness score from class-conditional Gaussian models of feature embeddings and use a contrastive guiding objective to steer diffusion sampling toward high-hardness examples. (Um et al., 2024) introduces a minority metric to steer generation toward selected likelihood levels. (Um et al., 2025) promotes minority generation through variance-boosted initialization and timestep skipping. Our method is orthogonal to these previous works and can be applied together during sampling.

3 METHOD

3.1 PRELIMINARIES: DENOISING AND WASSERSTEIN GRADIENT FLOW CORRECTIONS

Let π_t denote the noisy-data marginals of the variance-preserving (VP) continuous-time DDPM, with noise schedule $\beta(t) > 0$, data distribution π_0, and terminal noise distribution π_T (Ho et al., 2020; Song et al., 2021). Its score $s_t = \nabla \log \pi_t$ is approximated by a pretrained model $s_\theta(x, t)$. For reverse time $u = T - t$, Anderson's theorem gives the denoising stochastic differential equation (SDE) driven by a Wiener process B_u (Anderson, 1982):

$$dY_u = \left[\frac{\beta(t)}{2} Y_u + \beta(t) s_t(Y_u)\right] du + \sqrt{\beta(t)}\, dB_u, \qquad u = T - t. \tag{1}$$

Denote its sampling density by ϱ_u; with $Y_0 \sim \pi_T$, the exact dynamics give $\varrho_u = \pi_{T-u}$. The corresponding Fokker–Planck equation (FPE) is

$$\partial_u \varrho_u = -\nabla \cdot \left[\varrho_u \left(\frac{\beta(t)}{2} x + \beta(t) s_t\right)\right] + \frac{\beta(t)}{2} \Delta \varrho_u. \tag{2}$$

A review of the DDPM forward process and SDE-FPE formulations is given in Appendix A.

We recast the reverse FPE in the Wasserstein gradient-flow picture:

$$\begin{aligned} \partial_u \varrho_u &= \underbrace{-\frac{\beta(t)}{2} \nabla \cdot [\varrho_u (x + s_t(x))]}_{\text{transport}} + \underbrace{\frac{\beta(t)}{2} \nabla \cdot \left[\varrho_u \nabla \log \frac{\varrho_u}{\pi_t}\right]}_{\text{Wasserstein relaxation toward } \pi_t} \\ &\equiv -\nabla \cdot (\varrho_u v_t) + \frac{\beta(t)}{2} \nabla \cdot \left[\varrho_u \nabla \frac{\delta \mathcal{E}_{1,t}}{\delta \varrho}\right], \end{aligned} \tag{3}$$

where $v_t(x) \equiv \frac{\beta(t)}{2}\big(x + s_t(x)\big)$ defines the reverse probability-flow velocity, and

$$\mathcal{E}_{1,t}(\rho) \equiv D_{\mathrm{KL}}(\rho \| \pi_t) \tag{4}$$

defines the Kullback–Leibler (KL) divergence correction energy functional. This decomposition views DDPM sampling as combining transport across the marginals π_t with a Wasserstein gradient flow correction at each fixed noise level.

Predictor–corrector sampling (Song et al., 2021) alternates a reverse-SDE prediction across noise levels with a fixed-noise Langevin correction. During correction, t is held fixed and we denote the sampling density by ρ_τ, evolving in an inner time τ according to

$$\partial_\tau \rho_\tau = \nabla \cdot \left(\rho_\tau \nabla \log \frac{\rho_\tau}{\pi_t}\right) = \nabla \cdot \left[\rho_\tau \nabla \frac{\delta \mathcal{E}_{1,t}}{\delta \rho_\tau}\right]. \tag{5}$$

In the following, we propose a generalization of the correction step, with the more general Rényi divergence taking the place of the KL divergence.

3.2 RÉNYI SAMPLING CORRECTOR

As above, we consider fixed target noised data distribution π_t, and the sampling law ρ_τ evolves along an inner time parameter τ.

For $\alpha > 0$, $\alpha \neq 1$, the Rényi divergence (Rényi, 1961) is defined as

$$D_\alpha(\rho\|\pi_t) = \frac{1}{\alpha - 1} \log \int \rho^\alpha \pi_t^{1-\alpha}\, dx. \tag{6}$$

In the limit $\alpha \to 1$, we recover the KL divergence. We introduce the Rényi divergence correction energy functional as

$$\mathcal{E}_{\alpha,t}(\rho) \equiv \frac{1}{\alpha} D_\alpha(\rho\|\pi_t). \tag{7}$$

We further introduce the following for notational clarity

$$r_\tau = \frac{\rho_\tau}{\pi_t}, \qquad w_\tau = \frac{r_\tau^{\alpha-1}}{\mathbb{E}_{\rho_\tau}[r_\tau^{\alpha-1}]}. \tag{8}$$

Here, r_τ is the ratio of the current sampling density to the fixed noisy-data target at noise level t. The normalized mobility coefficient w_τ controls the local strength of probability transport and satisfies $\mathbb{E}_{\rho_\tau}[w_\tau] = 1$. [1] The variation of the Rényi divergence correction energy functional is

$$\frac{\delta \mathcal{E}_{\alpha,t}}{\delta \rho} = \frac{r^{\alpha-1}}{(\alpha-1)\mathbb{E}_\rho[r^{\alpha-1}]} = \frac{w}{\alpha - 1}, \qquad \nabla \frac{\delta \mathcal{E}_{\alpha,t}}{\delta \rho} = \frac{w \nabla r}{r} = w \nabla \log r. \tag{9}$$

The Wasserstein gradient flow of this Rényi divergence correction energy is

$$\partial_\tau \rho_\tau = \nabla \cdot [\rho_\tau w_\tau \nabla \log r_\tau] = \nabla \cdot \left[\rho_\tau w_\tau \nabla \log \frac{\rho_\tau}{\pi_t} \right]. \tag{10}$$

Compared to the standard Langevin correction, this PDE depends nonlinearly on the current density through w_τ. The relative mobility or correction strength will depend on both r_τ and the order α. Appendix D discusses the limiting behavior as $\alpha \to 0^+$ and $\alpha \to \infty$.

3.2.1 Implementation

We adopt a SDE formulation to implement the correction algorithm in practice. We note that the FPE 10 can be written as

$$\partial_\tau \rho = \nabla \cdot \left[w \underbrace{\rho \nabla \log \rho}_{=\nabla \rho} - w \rho \underbrace{\nabla \log \pi_t}_{=s_t} \right] = -\nabla \cdot [\rho(w s_t + \nabla w)] + \Delta(w\rho), \tag{11}$$

giving the equivalent SDE as

$$dX_\tau = [w_\tau(X_\tau) s_t(X_\tau) + \nabla w_\tau(X_\tau)]\, d\tau + \sqrt{2 w_\tau(X_\tau)}\, dW_\tau. \tag{12}$$

Estimates of the mobility w_τ and its spatial derivative ∇w_τ We estimate $\log r_\tau$ with an auxiliary discriminator D to distinguish sampler particles ($\sim \rho$) from noisy-data target samples ($\sim \pi$). The optimal discriminator satisfies $D^\star(x,t,\tau) = \frac{\rho_\tau(x)}{\rho_\tau(x)+\pi_t(x)}$, so the discriminator logit $\widehat{\ell}_\tau(x,t) = \log \frac{D^\star(x,t,\tau)}{1-D^\star(x,t,\tau)}$ estimates $\log r_\tau(x,t)$.

We use $\widehat{\ell}_\tau$ to construct $\widehat{w}_\tau$. Let $\{Y_\tau^j\}_{j=1}^N$ be reference particles approximating the current law. We have

$$\widehat{w}_\tau(x) = \frac{\exp((\alpha-1)\widehat{\ell}_\tau(x))}{\frac{1}{N}\sum_{j=1}^N \exp((\alpha-1)\widehat{\ell}_\tau(Y_\tau^j))}, \tag{13}$$

$$\nabla \widehat{w}_\tau(x) = (\alpha-1)\widehat{w}_\tau \nabla \widehat{\ell}_\tau(x). \tag{14}$$

Appendix B.1 details discriminator fitting, logit-offset cancellation, and the convention of holding the reference particles and normalizer fixed when differentiating the mobility.

[1] The factor $1/\alpha$ in Eq. 7 normalizes the mobility, *i.e.* $\mathbb{E}_{\rho_\tau}[w_\tau] = 1$. This keeps the correction budget comparable across different α.

Rényi sampling corrector We thus formulate the Rényi sampling correction recipe with a K-step Euler–Maruyama correction

$$x_t^{(k+1)} = x_t^{(k)} + \epsilon_t \left[\widehat{w}_t^k s_\theta(x^{(k)}, t) + \nabla \widehat{w}_t^k\right](x^{(k)}) + \sqrt{2\epsilon_t \widehat{w}_t^k (x_t^{(k)})}\, z_t^k \tag{15}$$

Here, $\widehat{w} > 0$ denote a smooth estimated mobility, $z_t^k \sim \mathcal{N}(0, I)$ are independently sampled while the step size ϵ_t is shared across the K iterations. Following the signal-to-noise ratio (SNR) heuristic of (Song et al., 2021), we calibrate the step size using the typical drift and noise magnitudes of the standard Langevin corrector

$$\epsilon_t = \frac{2a_t}{K}\left(\frac{\mathrm{snr} \times \mathbb{E}_{Z\sim\mathcal{N}(0,I)}\|Z\|}{\max(S_t, \delta)}\right)^2, \qquad S_t = \frac{1}{N}\sum_{j=1}^{N} \|s_\theta(x_t^{(j)}, t)\|, \tag{16}$$

where $a_t = 1 - \beta_t$ is the noise schedule, $\mathrm{snr} > 0$ is a hyperparameter controlling the correction strength, $\delta > 0$ prevents division by zero, and N denotes the sampling batch size for estimating the average score norm. [2] To summarize, we compute ϵ_t for each noise level, then perform K iterative updates with the estimated mobility and its gradient to correct the realized drift-to-noise ratio. We also note that setting $K = 0$ recovers ancestral DDPM sampling, and setting $\alpha = 1$ recovers the standard Langevin correction. Appendix B.2 derives the step-size calibration and explains the total correction time $K\epsilon_t$ and its relation to the nominal SNR.

Following our Rényi sampling corrector, we alternate back to the standard DDPM predictor for the $t-1$ noise level update.

$$x_{t-1} = \frac{x^{(K)} + \beta_t s_\theta(x^{(K)}, t)}{\sqrt{a_t}} + \sqrt{\beta_t}\, z_t. \tag{17}$$

3.3 Dissipation properties of $\mathcal{E}_{\alpha,t}$

In this section, we show that the $\mathcal{E}_{\alpha,t}$ has favorable dissipation properties, *i.e.*, correction reduces the Rényi divergence between the sampling distribution ρ_τ and target distribution π_t. At positive noise levels, Gaussian corruption gives a smooth positive target. Throughout the following analysis, we assume smooth positive sampling densities and standard regularity conditions, including integrability and sufficient decay. We report only key results and refer interested readers to Appendix C for detailed derivations.

Correction reduces the mismatch with the target. We first consider continuous-time correction in the ideal situation with exact scores and mobility. At a fixed noise level t, the Wasserstein gradient flow in Eqn. 10 satisfies

$$\frac{d}{d\tau}\mathcal{E}_{\alpha,t}(\rho_\tau) = -\int \rho_\tau \left\|\nabla \frac{\delta \mathcal{E}_{\alpha,t}}{\delta \rho}\right\|^2 dx = -\int \rho_\tau w_\tau^2 \|\nabla \log r_\tau\|^2 dx \le 0, \tag{18}$$

which shows that correction reduces Rényi divergence. Its rate depends on $w_\tau^2\|\nabla \log r_\tau\|^2$. At $\rho_\tau = \pi_t$, we have $r_\tau = w_\tau = 1$, and the probability flux vanishes. Conversely, zero dissipation implies that r_τ is constant, giving $\rho_\tau = \pi_t$. This shows that π_t is the unique stationary density within this smooth positive class.

For the convergence rate, if

$$\int \rho_\tau w_\tau^2 \|\nabla \log r_\tau\|^2 dx \ge 2c_{\alpha,t}\mathcal{E}_{\alpha,t}(\rho_\tau) \tag{19}$$

holds along the flow for some $c_{\alpha,t} > 0$, then by Grönwall's inequality, we have

$$\mathcal{E}_{\alpha,t}(\rho_\tau) \le e^{-2c_{\alpha,t}\tau}\mathcal{E}_{\alpha,t}(\rho_0). \tag{20}$$

Appendix C.1 proves the dissipation identity and gives sufficient conditions for Eqn. 19, together with an explicit convergence-rate bound.

[2] The standard Gaussian Z in d dimensions has expected magnitude $\mathbb{E}\|Z\| = \sqrt{2}\,\frac{\Gamma((d+1)/2)}{\Gamma(d/2)}$.

Dissipation along the denoising trajectory. We now advance beyond correction towards a fixed target but the more general case involving prediction where the target evolves as reverse time $u = T - t$ advances. We consider a continuous-time model in which $\lambda_u \geq 0$ specifies the amount of correction per unit (reverse) time. Combining the full reverse dynamics in Eqn. 3 with the Rényi correction, we have

$$\partial_u \varrho_u = -\nabla \cdot (\varrho_u v_t) + \nabla \cdot \left[\varrho_u \left(\frac{\beta(t)}{2} + \lambda_u w_u \right) \nabla \log r_u \right]. \tag{21}$$

The target itself satisfies $\partial_u \pi_t = -\nabla \cdot (\pi_t v_t)$, with $t = T - u$. Consequently, the transport component moves both distributions through the same velocity field and leaves their Rényi divergence unchanged. Only the Wasserstein relaxation terms contribute to its dissipation:

$$\frac{d}{du}\mathcal{E}_{\alpha,t}(\varrho_u) = -\underbrace{\frac{\beta(t)}{2}\int \varrho_u w_u \|\nabla \log r_u\|^2 dx}_{\text{reverse-diffusion contribution}} - \underbrace{\lambda_u \int \varrho_u w_u^2 \|\nabla \log r_u\|^2 dx}_{\text{additional correction contribution}} \leq 0. \tag{22}$$

Here the dissipation accounts for both the evolving sampling density and the changing target. Eqn. 22 shows that the added correction reduces mismatch while preserving the intended denoising trajectory. If sampling is initialized exactly from π_T, the exact reverse dynamics remain on the marginal path, i.e., $\varrho_u = \pi_{T-u}$. In practice, however, sampling is typically initialized from $\mathcal{N}(0, I)$, which only approximates π_T for a finite diffusion horizon. This introduces an initial mismatch, which will not be amplified following the reverse dynamics evolution. Appendix C.2 derives the common-transport cancellation and the integrated dissipation identity.

Removing the exact-score and exact-mobility idealization. As outlined in Section 3.2.1, we use estimates $s_\theta \approx s_t$, and $\widehat{w}_\tau \approx w_\tau$. Writing the true energy derivative in terms of these estimates gives

$$\frac{d}{d\tau}\mathcal{E}_{\alpha,t}(\rho_\tau) = -\int \rho_\tau w_\tau \widehat{w}_\tau \|\nabla \log r_\tau\|^2 dx + \int \rho_\tau w_\tau \widehat{w}_\tau \nabla \log r_\tau \cdot (s_\theta - s_t)\, dx. \tag{23}$$

The first term represents the dissipative contribution of the correction, modulated by the estimated mobility $\widehat{w}_\tau$, while the second term captures the effect of score-estimation error. Since the latter is not sign-definite, score error can weaken or, if sufficiently large and adversely aligned with the true ratio gradient, locally reverse the energy dissipation. Appendix C.3 derives this identity and bounds the accumulated score error along denoising. Appendix C.4 quantifies how log-ratio and normalization errors change the dissipation strength. Finite numerical steps introduce further approximation errors, whose practical effects are evaluated through sampling experiments.

4 Experiments

We investigate three questions using fixed pretrained DDPMs, varying only the sampling configuration. **(1) Distributional fidelity:** can Rényi correction improve agreement with the target distribution beyond ancestral sampling and standard Langevin correction? **(2) Minority-mode imbalance:** how does the order α affect minority modes that are underrepresented or overrepresented by the baseline sampler? We examine both regimes in imbalanced GMMs and evaluate cat frequency in a cat–human mixture, asking whether correction brings minority-mode mass closer to its target without explicit class guidance. **(3) Order and correction budget:** how can the observed sampling response inform the choice of α and the number of correction substeps K? We use local response fits around the Langevin baseline to relate order sensitivity to initial mode-mass discrepancies. We also examine how increasing K affects image fidelity and computational cost.

4.1 Imbalanced Gaussian mixtures

Setup and metrics. We use eight-component isotropic Gaussian mixtures with fixed minimum separation: one component has probability p, and the others share $1 - p$ equally. The underrepresentation sweep uses $p \in \{0.001, 0.003, 0.005, 0.01, 0.02, 0.05\}$; substantial rare-mode deficits occur at the three smallest probabilities. Each condition uses a fixed pretrained DDPM, $\alpha \in \{[0.7, 0.8, 0.9, 1.0, 1.1, 1.2, 1.3\}$, $K = 1$, and five sampling seeds. The snr hyperparameter is

Table 1: GMM mode-frequency TV in the underrepresentation sweep (mean $\pm$ standard deviation across five sampling seeds).

p	$\alpha = 0.7$	$\alpha = 1$	$\alpha = 1.3$
0.001	0.02378 ± 0.0032	0.02553 ± 0.0016	0.02927 ± 0.0012
0.003	0.02338 ± 0.0015	0.02605 ± 0.00094	0.03020 ± 0.0023
0.005	0.05323 ± 0.0027	0.05652 ± 0.0028	0.05986 ± 0.0023
0.01	0.03465 ± 0.0025	0.03782 ± 0.0024	0.04105 ± 0.0015
0.02	0.03171 ± 0.0018	0.03298 ± 0.00089	0.03684 ± 0.0024
0.05	0.04897 ± 0.0035	0.05075 ± 0.0011	0.05190 ± 0.0029

Table 2: Rare-mode mass ratios (seed means) and signed sensitivities (mean $\pm$ standard error). Under/over denote standard/exponentially tilted training. Response fits use $\alpha \in [0.7, 1.3]$ for underrepresentation and $[0.9, 1.1]$ for overrepresentation; see Appendix E.1.

Sweep	p	$r_0(0.7)$	$r_0(1)$	$r_0(1.3)$	$\widehat{\kappa}_0$
Under	0.001	0.417	0.346	0.146	1.703 ± 0.273
Under	0.003	0.467	0.372	0.269	0.889 ± 0.038
Under	0.005	0.600	0.534	0.437	0.466 ± 0.055
Under	0.01	1.053	1.048	0.943	-0.107 ± 0.078
Under	0.02	1.092	1.069	1.038	-0.062 ± 0.030
Under	0.05	1.009	0.991	0.955	0.049 ± 0.043
Over	0.003	4.680	3.492	2.416	-0.960 ± 0.129
Over	0.005	3.230	2.636	1.884	-0.799 ± 0.053
Over	0.01	3.136	2.770	2.323	-0.717 ± 0.056

set to 0.16. We assign samples to the nearest center within half the minimum inter-center separation. For assigned fractions m_d, target weights p_d, and unassigned fraction $m_{\text{unassigned}}$, mode-frequency TV is $\frac{1}{2}(\sum_{d=0}^{7} |m_d - p_d| + m_{\text{unassigned}})$. Lower TV indicates better agreement in mode proportions and less unassigned mass. The rare-mode ratio $r_0 = m_0/p_0$ measures underrepresentation ($r_0 < 1$) or overrepresentation ($r_0 > 1$), using all generated samples in the denominator. Appendix E.1 gives full experimental details and results.

Distributional fidelity. Reducing α generally lowers mode-frequency TV, with $\alpha = 0.7$ attaining the lowest mean in all six conditions (Table 1). Relative to Langevin correction, TV falls by 6.83%, 10.26%, and 5.82% for the three most imbalanced settings.

Rare-mode imbalance and order selection. At a fixed sampling law, $w \propto r^{\alpha-1}$ assigns relatively greater correction strength to smaller pointwise density ratios for $\alpha < 1$ and larger ratios for $\alpha > 1$ (Appendix D). This motivates testing smaller orders for depleted modes and larger orders for oversampled modes. The resulting mode mass also depends on density gradients and transport across mode boundaries, so this heuristic requires empirical validation.

In the three depleted settings, reducing α from 1 to 0.7 recovers 11–15% of the Langevin rare-mass deficit; the remaining settings stay near their target proportions (Table 2). To examine the opposite extreme, we train with exponentially tilted weights (Li et al., 2021) that overrepresent rare modes. Increasing α from 1 to 1.3 removes 25–46% of their excess mass. These percentages measure the reduction in distance from the target mass, rather than the relative change in generated mass. Order selection should check aggregate TV alongside the rare-mode ratio, since improving one mode can redistribute error among the others.

Signed response around Langevin correction. For component d, let $r_d(\alpha) = m_d(\alpha)/p_d$ be its generated-to-target mass ratio. Parameterizing the proposed adjustment as $\alpha_d(h) = 1 + \operatorname{sgn}(\log r_d(1))h$, $h \geq 0$, a smooth positive response satisfies

$$\log r_d(1 + \operatorname{sgn}(\log r_d(1))h) = \log r_d(1) + \kappa_d h + O(h^2), \tag{24}$$

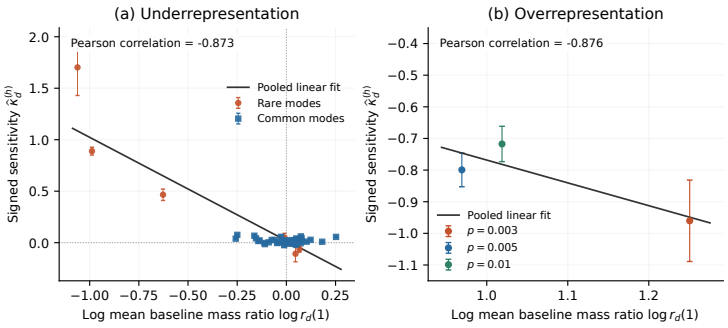


Figure 1: Signed sensitivity versus baseline discrepancy, with separate linear fits: (a) all 48 mode–condition pairs in the underrepresentation sweep; (b) three rare-mode conditions after tilted training. Points show seed means; error bars are standard errors. Fit intervals are given in Table 2.

Table 3: Stacked-MNIST filtered TV with 64 correction anchors (across five seeds).

K	Langevin ($\alpha = 1$)	Rényi ($\alpha = 0.5$)
1	0.4247 ± 0.0043	0.4263 ± 0.0047
2	0.4245 ± 0.0044	0.4214 ± 0.0038
3	0.4242 ± 0.0043	0.4190 ± 0.0021
Ancestral sampling ($K = 0$)		0.4353 ± 0.0026

where $\kappa_d = \operatorname{sgn}(\log r_d(1)) \left.\frac{d \log r_d(\alpha)}{d\alpha}\right|_{\alpha=1}$ is sensitivity along h. Increasing h lowers α for an undersampled mode and raises it for an oversampled mode. A positive κ_d helps a depleted mode, while a negative κ_d helps an oversampled mode; both satisfy $\log r_d(1)\kappa_d < 0$. Sampling uses one shared order, so these mode-specific diagnostics inform a compromise across the distribution.

We estimate $\widehat{\kappa}_d$ by fitting a baseline-anchored slope in α for each seed and applying the sign of its baseline discrepancy. Regressing mean sensitivities against log mean baseline ratios gives $\widehat{\kappa}_d \approx 0.020 - 1.003 \log r_d(1)$ across 48 underrepresentation mode–condition pairs (correlation -0.873), and $\widehat{\kappa}_d \approx -0.047 - 0.721 \log r_d(1)$ across three oversampled rare-mode conditions (correlation -0.876; Fig. 1). The underrepresentation association is largely driven by rare-mode deficits; the common-mode slope is only -0.030. These descriptive fits use different populations and intervals. Their slopes depend on the model and sampling configuration, so the Taylor expansion provides a local picture near $\alpha = 1$: the baseline discrepancy and signed sensitivity jointly indicate a promising adjustment, whose magnitude should be checked empirically.

4.2 Stacked MNIST

Setup and metrics. We stack MNIST (LeCun et al., 1998) digits into three channels, the first two digits are uniform, and the third samples 7 with probability 0.01. A discriminator estimates density ratios from sampler particles and unlabeled references. We use 1,000 predictor steps, 64 correction anchors, and 10,000 images per seed over five seeds, dividing a fixed nominal correction time per anchor across $K \in \{1, 2, 3\}$ substeps. The snr hyperparameter is set to 0.16. A triple passes the confidence filter when all three digit predictions have confidence at least 0.9. For accepted combination frequencies m_c^{acc}, target probabilities p_c, and rejected fraction m_{rej}, filtered TV is $\frac{1}{2}(\sum_{c=1}^{1000} |m_c^{\text{acc}} - p_c| + m_{\text{rej}})$. Frequencies use all generated images as the denominator, retaining rejected images in a zero-target-mass bin. This penalizes both frequency mismatch and low-confidence outputs. Appendix E.2 reports full experimental protocol and results.

Distributional fidelity and correction budget. At $\alpha = 0.5$, $K = 3$, Rényi correction improves filtered TV from the ancestral baseline's 0.4353 ± 0.0026 to 0.4190 ± 0.0021, a 3.74% reduction and a further 1.22% improvement over matched Langevin correction (Table 3). Increasing K lowers TV at this order.

Table 4: Faces: ancestral sampling and matched corrections across five seeds; lower FID is better.

K	α	Cats per 500	Cat FID	Human FID
0	1	19.0 ± 3.4	47.13 ± 1.21	17.37
1	1	18.7 ± 3.4	46.88 ± 1.03	17.38
1	0.7	23.2 ± 3.2	45.32 ± 0.95	17.35
2	1	19.3 ± 3.5	46.53 ± 0.98	17.45
2	0.7	23.2 ± 3.0	45.16 ± 0.84	17.45
3	1	19.2 ± 3.4	46.35 ± 1.11	17.44
3	0.7	23.0 ± 3.3	45.02 ± 0.94	17.40

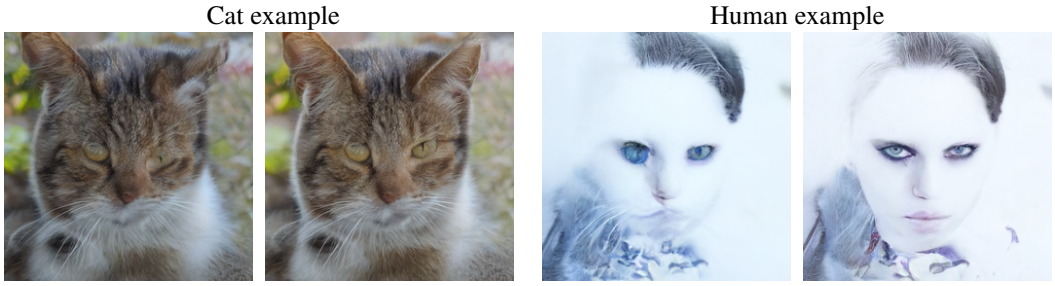


Figure 2: Paired examples at $\alpha = 0.7$, $K = 2$. Correction improves eye details in the cat and reduces cat-like features in the human example.

Sampling time. For $10,000$ generations, sampling takes 47.47s for ancestral, 49.73, 50.35, 53.14s for Langevin corrections and 53.34, 61.97, 70.60s for Rényi corrections at $K = 1, 2, 3$.

4.3 Faces

Setup and metrics. We use a mixture of 5% AFHQ cats (Choi et al., 2020) and 95% CelebA humans (Liu et al., 2015) and a density-ratio discriminator. Across five seeds, we vary $K \in \{1, 2, 3\}$ and $\alpha \in \{0.7, 0.9, 1, 1.1, 1.3\}$. The snr hyperparameter is set to 0.32. We report cat counts per 500 images (target: 25) and class-specific FID (Heusel et al., 2017). The class-specific FID measures fidelity within each class, comparing feature means and covariances of generated and reference images through the squared Wasserstein-2 distance between Gaussian approximations. More details on experiment protocols and results are found in Appendix E.3.

Minority-class fidelity. For $\alpha = 0.7$, correction reduces $\approx 70\%$ of the cat-mass deficit and improves mean cat FID compared to the ancestral sampling and Langevin correction baselines (Table 4). More Rényi correction iterations K also improve FID, consistent with the MNIST experiments. Qualitatively, we illustrate in Fig. 2 how correction can improve cat generation and also reduce artifacts in a problematic hybrid human-cat mix sample.

5 Conclusion and Outlook

We propose Rényi correction, using an estimate of the density ratio between sampling and target distributions to induce a density dependent correction drift, serving to reduce target mismatch during DDPM sampling. Evaluating our approach across controlled Gaussian mixture experiments to more practical face generation experiments, we validate that Rényi corrections can indeed benefit distinct imbalance regimes, whether it be recovering distributional proportions or improving generation quality. A promising next step is extending the approach to diffusion LLM sampling (Nie et al., 2025), adapting ratio estimation and target-preserving corrections to discrete token dynamics.

AI USE STATEMENT

We used generative AI tools as discussion partners during ideation, to polish the manuscript, consolidate experimental results into tables, and assist with formatting. AI tools also assisted with prototyping experiments as well as for monitoring experimental sweeps. In the appendix, we used AI assistance to draft derivations of established results, including Anderson's time-reversal theorem and applying Poincaré inequality for convergence bounds. These materials provide background for completeness and are presented separately from our methodological contributions. The authors retain responsibility for the research design, interpretation of results, and all final content, including the correctness of AI-assisted text, mathematical derivations, code, and experimental reporting.

ETHICS STATEMENT

This work studies sampling corrections for pretrained diffusion models. We do not identify ethical concerns specific to the proposed correction mechanism. Our experiments examine whether sampling can better reproduce imbalanced data distributions without explicit class guidance. We do note that matching a data distribution does not itself ensure fairness, since the reference data may contain societal biases. The method also shares the broader risks of image generation, including deceptive synthetic content and unauthorized depictions of individuals. Any deployment should consider the provenance and permitted uses of the training and reference data, privacy, and the potential consequences of generated outputs.

REPRODUCIBILITY STATEMENT

Appendix B describes density-ratio estimation and numerical implementation while experimental protocols, evaluation metrics are provided in Appendix E. We also provide an anonymous code repository to facilitate reproduction of the experiments: `https://github.com/Anonymous-PaperSubmission/RenyiSamplingDDPM`.

REFERENCES

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

Dominique Bakry, Ivan Gentil, and Michel Ledoux. *Analysis and Geometry of Markov Diffusion Operators*, volume 348 of *Grundlehren der mathematischen Wissenschaften*. Springer, 2014. doi: 10.1007/978-3-319-00227-9.

Fang Chen, Alex Villa, Gongbo Liang, Xiaoyi Lu, and Meng Tang. Contrastive conditional-unconditional alignment for long-tailed diffusion model. *arXiv preprint arXiv:2507.09052*, 2025.

Yunjey Choi, Youngjung Uh, Jaejun Yoo, and Jung-Woo Ha. StarGAN v2: Diverse image synthesis for multiple domains. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020.

Prafulla Dhariwal and Alex Nichol. Diffusion models beat GANs on image synthesis. In *Advances in Neural Information Processing Systems*, 2021.

Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. GANs trained by a two time-scale update rule converge to a local Nash equilibrium. In *Advances in Neural Information Processing Systems*, volume 30, 2017.

Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*, volume 33, pp. 6840–6851, 2020.

Richard Jordan, David Kinderlehrer, and Felix Otto. The variational formulation of the Fokker–Planck equation. *SIAM Journal on Mathematical Analysis*, 29(1):1–17, 1998.

Dongjun Kim, Yeongmin Kim, Se Jung Kwon, Wanmo Kang, and Il-Chul Moon. Refining generative process with discriminator guidance in score-based diffusion models. In *Proceedings of the 40th International Conference on Machine Learning*, pp. 16567–16598. PMLR, 2023.

Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.

Tian Li, Ahmad Beirami, Maziar Sanjabi, and Virginia Smith. Tilted empirical risk minimization. In *International Conference on Learning Representations*, 2021.

Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of the IEEE International Conference on Computer Vision*, 2015.

Shen Nie, Fengqi Zhu, Zebin You, Xiaolu Zhang, Jingyang Ou, Jun Hu, Jun Zhou, Yankai Lin, Ji-Rong Wen, and Chongxuan Li. Large language diffusion models. *arXiv preprint arXiv:2502.09992*, 2025.

Felix Otto. The geometry of dissipative evolution equations: The porous medium equation. *Communications in Partial Differential Equations*, 26(1-2):101–174, 2001.

Yiming Qin, Huangjie Zheng, Jiangchao Yao, Mingyuan Zhou, and Ya Zhang. Class-balancing diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023.

Alfréd Rényi. On measures of entropy and information. *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Contributions to the Theory of Statistics*, pp. 547–561, 1961.

Esther Rodriguez, Monica Welfert, Samuel McDowell, Nathan Stromberg, Julian Antolin Camarena, and Lalitha Sankar. CORAL: Disentangling latent representations in long-tailed diffusion. *arXiv preprint arXiv:2506.15933*, 2025.

Vikash Sehwag, Caner Hazirbas, Albert Gordo, Firat Ozgenel, and Cristian Canton Ferrer. Generating high fidelity data from low-density regions using diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11492–11501, 2022.

Andy Shih, Dorsa Sadigh, and Stefano Ermon. Long horizon temperature scaling. In *International conference on machine learning*, pp. 31422–31434. PMLR, 2023.

Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021.

Soobin Um, Suhyeon Lee, and Jong Chul Ye. Don't play favorites: Minority guidance for diffusion models. In *International Conference on Learning Representations*, 2024.

Soobin Um, Beomsu Kim, and Jong Chul Ye. Boost-and-skip: A simple guidance-free diffusion for minority generation. In *Proceedings of the 42nd International Conference on Machine Learning*, 2025.

Tim van Erven and Peter Harremoës. Rényi divergence and kullback–leibler divergence. *IEEE Transactions on Information Theory*, 60(7):3797–3820, 2014.

An B. Vuong, Michael T. McCann, Javier E. Santos, and Yen Ting Lin. Are we really learning the score function? reinterpreting diffusion models through Wasserstein gradient flow matching. *arXiv preprint arXiv:2509.00336*, 2025.

Tim Y. J. Wang and O. Deniz Akyildiz. A gradient flow approach to solving inverse problems with latent diffusion models. In *Second Workshop on Frontiers in Probabilistic Inference: Sampling Meets Learning at NeurIPS*, 2025.

Tim Y. J. Wang, Juan Kuntz, and O. Deniz Akyildiz. Training latent diffusion models with interacting particle algorithms. In *Proceedings of the 29th International Conference on Artificial Intelligence and Statistics*, volume 300 of *Proceedings of Machine Learning Research*, pp. 2791–2799. PMLR, 2026.

Yanbo Xu, Yu Wu, Sungjae Park, Zhizhuo Zhou, and Shubham Tulsiani. Temporal score rescaling for temperature sampling in diffusion and flow models. In *International Conference on Machine Learning*, 2026.

A DDPM PRELIMINARIES

A.1 THE FORWARD DDPM PROCESS

The continuous-time limit of the DDPM forward process is the variance-preserving SDE (Ho et al., 2020; Song et al., 2021)

$$dx_t = -\frac{\beta(t)}{2} x_t \, dt + \sqrt{\beta(t)} \, dW_t, \tag{25}$$

where $\beta(t) > 0$ is the noise schedule and W_t is a Wiener process. Its marginal density π_t satisfies

$$\partial_t \pi_t = \nabla \cdot \left(\frac{\beta(t)}{2} x \pi_t \right) + \frac{\beta(t)}{2} \Delta \pi_t. \tag{26}$$

The SDE describes individual stochastic trajectories, while the FPE describes their density. Forward corruption gradually transforms the data distribution π_0 toward a simple noise distribution at time T. In discrete DDPM time k, writing $a_k = 1 - \beta_k$ and $\bar{a}_k = \prod_{j=1}^{k} a_j$ gives the Gaussian corruption marginal

$$\pi_k(x) = \int \mathcal{N}\big(x; \sqrt{\bar{a}_k} x_0, (1 - \bar{a}_k) I\big) \, p_{\mathrm{data}}(dx_0). \tag{27}$$

A.2 THE SDE-FPE CORRESPONDENCE

Consider an Itô SDE $dX_t = b_t(X_t) \, dt + \Sigma_t(X_t) \, dW_t$, and set $A_t = \Sigma_t \Sigma_t^\top$. For a smooth compactly supported test function ϕ, Itô's formula gives

$$\frac{d}{dt} \int \phi p_t \, dx = \int p_t \left[b_t \cdot \nabla \phi + \frac{1}{2} \sum_{i,j} (A_t)_{ij} \partial_i \partial_j \phi \right] dx. \tag{28}$$

The stochastic increment has zero expectation. Integrating the drift term by parts once and the diffusion term twice transfers the derivatives from ϕ to the density. Since the resulting identity holds for every such ϕ,

$$\partial_t p_t = -\nabla \cdot (b_t p_t) + \frac{1}{2} \sum_{i,j} \partial_i \partial_j \big((A_t)_{ij} p_t \big). \tag{29}$$

Under regularity and uniqueness conditions for the SDE and FPE, this identifies the FPE solution with the SDE's marginal law. Substituting $b_t = -\beta(t)x/2$ and $A_t = \beta(t) I$ recovers Eq. 26. The same correspondence gives the corrector SDE in Eq. 12: choosing $b = w s_t + \nabla w$ and $A = 2wI$ yields $-\nabla \cdot [\rho(w s_t + \nabla w)] + \Delta(w\rho) = \nabla \cdot [w \nabla \rho - \rho w s_t]$.

A.3 ANDERSON'S TIME-REVERSAL THEOREM

Under sufficient smoothness, nondegeneracy, and integrability conditions, with positive smooth marginal densities, Anderson's theorem states that the time reversal of a diffusion is again a diffusion (Anderson, 1982; Song et al., 2021). In increasing reverse time $u = T - t$, its SDE is

$$dY_u = \left[-b_t + \nabla \cdot A_t + A_t \nabla \log p_t \right](Y_u) \, du + \Sigma_t(Y_u) \, dB_u, \qquad Y_0 \sim p_T, \tag{30}$$

where $(\nabla \cdot A_t)_i = \sum_j \partial_j (A_t)_{ij}$ and B_u is a Wiener process in reverse time.

For completeness sake, we present a short derivation using Bayes' rule. The backward transition from x at time t to y at $t - h$ is proportional to the forward transition from y to x, weighted by $p_{t-h}(y)$. Let $L_t \phi = b_t \cdot \nabla \phi + \frac{1}{2} \sum_{i,j} (A_t)_{ij} \partial_i \partial_j \phi$, and let L_t^* be the density operator in Eq. 29. Expanding the backward conditional expectation to first order in h gives its generator

$$\overline{L}_u \phi = \frac{L_t^*(p_t \phi) - \phi L_t^* p_t}{p_t} = [-b_t + \nabla \cdot A_t + A_t \nabla \log p_t] \cdot \nabla \phi + \frac{1}{2} \sum_{i,j} (A_t)_{ij} \partial_i \partial_j \phi. \tag{31}$$

Reading off the drift and covariance yields Eq. 30. For the VP process, $A_t = \beta(t) I$ has no spatial derivative and $p_t = \pi_t$, giving Eq. 1. Applying Eq. 29 then gives Eq. 2. With the exact score and terminal law, these dynamics recover $\varrho_u = \pi_{T-u}$.

B Implementation details of the Rényi corrector

B.1 Discriminator fitting and mobility estimation

For fixed noise level t and correction time τ, we train a discriminator $D_\phi(x, t, \tau)$ using equal class weights, with current sampling particles labeled 1 and noisy-data samples labeled 0. The binary cross-entropy objective is

$$\mathcal{L}_{\text{disc}}(\phi) = -\mathbb{E}_{\rho_\tau} \log D_\phi(X, t, \tau) - \mathbb{E}_{\pi_t} \log\big(1 - D_\phi(X, t, \tau)\big). \tag{32}$$

Minimizing the loss at each x gives

$$D^\star(x, t, \tau) = \frac{\rho_\tau(x)}{\rho_\tau(x) + \pi_t(x)}, \qquad \log \frac{D^\star(x, t, \tau)}{1 - D^\star(x, t, \tau)} = \log r_\tau(x). \tag{33}$$

In practice, the fitted logit $\widehat{\ell}_\tau(x, t) = \log[D_\phi(x, t, \tau)/(1 - D_\phi(x, t, \tau))]$ supplies the estimate in Eqn. 14. Target examples are obtained by applying the standard forward corruption to reference data. The particles $\{Y_\tau^j\}_{j=1}^N$ used to normalize the mobility represent the current sampling law ρ_τ.

The discriminator can share parameters across noise levels and correction times, provided its training examples represent the corresponding sampling laws. Online fitting refreshes those examples as sampling proceeds; pretraining uses trajectories generated in advance. Reusing a pretrained discriminator requires checking its accuracy for the sampling configuration in use, including the order α and correction budget. Eqn. 59 below describes how log-ratio error affects the mobility.

We further note that a constant offset in the logit $\widehat{\ell}_\tau = \log r_\tau + c_\tau$ cancels from its normalized exponential

$$\frac{e^{(\alpha-1)(\log r_\tau + c_\tau)}}{\mathbb{E}_{\rho_\tau} e^{(\alpha-1)(\log r_\tau + c_\tau)}} = \frac{r_\tau^{\alpha-1}}{\mathbb{E}_{\rho_\tau}[r_\tau^{\alpha-1}]} = w_\tau. \tag{34}$$

This mitigates potential issues of logit offsets induced by unequal class priors.

When evaluating $\nabla_x \widehat{w}_\tau(x)$, we hold the fitted parameters, reference particles, and normalization factor fixed. Differentiation then acts only on the query x, giving the gradient in Eqn. 14. This convention is essential to the drift–diffusion cancellation proved in Appendix C.3.

B.2 Step size and correction time

For a standard Langevin step of length ϵ, the typical drift and noise magnitudes are ϵS_t and $\sqrt{2\epsilon}\,\mathbb{E}\|Z\|$. Equating their ratio to snr yields $\epsilon = 2(\mathrm{snr}\,\mathbb{E}\|Z\|/S_t)^2$. Eqn. 16 includes the VP factor a_t, a denominator safeguard δ, and a division by K to distribute the correction time across K updates. For $K \geq 1$, the total nominal correction time is therefore

$$K\epsilon_t = 2a_t \left(\frac{\mathrm{snr}\,\mathbb{E}\|Z\|}{\max(S_t, \delta)}\right)^2. \tag{35}$$

Increasing K refines this fixed time budget. The edge case $K = 0$ implies skipping the corrector and its step-size calculation. Since the realized drift-to-noise ratio also depends on the mobility $\widehat{w}_t^k$ and its gradient, the nominal SNR is not the exact SNR of each correction update. Rather, this calibration fixes a common correction-time baseline across Rényi orders, enabling comparisons under the same nominal sampling budget.

C Proofs of dissipation and error bounds

Following Section 3.2, we use ρ_τ for fixed-noise correction and ϱ_u along the schedule $t = T - u$. At positive noise levels, Gaussian corruption gives a smooth positive target π_t. We assume the correction density also remains smooth and positive, with finite second moments and energy, and sufficient integrability and decay for differentiation and integration by parts.

C.1 FIXED-NOISE DISSIPATION AND CONVERGENCE

We introduce the notational shorthand $\psi_\tau = \delta\mathcal{E}_{\alpha,t}/\delta\rho$. By Eqn. 9, $\nabla\psi_\tau = w_\tau \nabla \log r_\tau$. Eqn. 10 gives

$$\begin{aligned}\frac{d}{d\tau}\mathcal{E}_{\alpha,t}(\rho_\tau) &= \int \psi_\tau\, \partial_\tau \rho_\tau\, dx \\ &= \int \psi_\tau \nabla\cdot(\rho_\tau \nabla\psi_\tau)\, dx \\ &= -\int \rho_\tau \|\nabla\psi_\tau\|^2 dx \\ &= -\int \rho_\tau w_\tau^2 \|\nabla \log r_\tau\|^2 dx. \end{aligned} \tag{36}$$

Integration by parts shows that moving mass down the energy gradient lowers the energy by its squared magnitude, proving Eqn. 18.

When the convergence bound holds. Eqn. 19 requires the energy decrease to control the remaining mismatch. The target's mixing properties and the range of the density ratio provide concrete sufficient conditions. For $\alpha = 1$, the bound is exactly the logarithmic Sobolev inequality in density form:

$$D_{\mathrm{KL}}(\rho\|\pi_t) \le \frac{1}{2\kappa_t}\int \rho \left\|\nabla \log \frac{\rho}{\pi_t}\right\|^2 dx. \tag{37}$$

It gives $c_{1,t} = \kappa_t$. A familiar sufficient condition is $\pi_t \propto e^{-V_t}$ with $\nabla^2 V_t \succeq \kappa_t I$ and $\kappa_t > 0$, by the Bakry–Émery criterion (Bakry et al., 2014). For example, a Gaussian target with covariance $\sigma_t^2 I$ admits $\kappa_t = 1/\sigma_t^2$.

For any fixed $\alpha > 0$, a sufficient condition combines a Poincaré inequality for the target with a uniform density-ratio bound along the correction:

$$\mathrm{Var}_{\pi_t}(f) \le C_{P,t}\int \pi_t \|\nabla f\|^2 dx, \qquad R^{-1} \le r_\tau(x) \le R, \quad R \ge 1. \tag{38}$$

Here $C_{P,t} < \infty$ and R are independent of τ on the interval considered. The first condition controls how easily mass moves across the target; the second prevents the correction from becoming arbitrarily weak or singular through extreme density ratios. A uniformly convex V_t also gives $C_{P,t} \le 1/\kappa_t$ (Bakry et al., 2014).

To connect the energy to the Poincaré inequality, we first bound it by the variance of r. Since $\mathbb{E}_{\pi_t} r = 1$, this variance measures the squared departure of the density ratio from its equilibrium value $r = 1$. For $\alpha \neq 1$, define $F_\alpha(z) = [z^\alpha - 1 - \alpha(z-1)]/[\alpha(\alpha-1)]$. Subtracting the constant and linear terms gives $F_\alpha(1) = F'_\alpha(1) = 0$, while $F''_\alpha(z) = z^{\alpha-2}$. On the bounded interval $[R^{-1}, R]$, this curvature is at most $R^{|\alpha-2|}$. Taylor's theorem therefore gives

$$0 \le F_\alpha(r) \le \frac{1}{2}R^{|\alpha-2|}(r-1)^2, \qquad \mathbb{E}_{\pi_t}[F_\alpha(r)] \le \frac{1}{2}R^{|\alpha-2|}\,\mathrm{Var}_{\pi_t}(r).$$

We now relate this average to the Rényi energy. We define $Z_\alpha = \mathbb{E}_\rho[r^{\alpha-1}] = \int \pi_t r^\alpha dx$. The linear term in $F_\alpha(r)$ averages to zero, so $\mathbb{E}_{\pi_t}[F_\alpha(r)] = (Z_\alpha - 1)/[\alpha(\alpha-1)]$. Consequently,

$$\mathcal{E}_{\alpha,t}(\rho) = \frac{\log Z_\alpha}{Z_\alpha - 1}\,\mathbb{E}_{\pi_t}[F_\alpha(r)], \qquad \frac{\log Z_\alpha}{Z_\alpha - 1} \le \max\{1, Z_\alpha^{-1}\} \le R^{|\alpha-1|}.$$

The last inequality follows from $R^{-|\alpha-1|} \le Z_\alpha \le R^{|\alpha-1|}$. Combining the two bounds yields

$$\mathcal{E}_{\alpha,t}(\rho) \le \frac{1}{2}R^{|\alpha-2|+|\alpha-1|}\,\mathrm{Var}_{\pi_t}(r). \tag{39}$$

At $\alpha = 1$, the same argument uses $F_1(z) = z\log z - z + 1$, whose average is exactly the KL energy. Thus the bounded ratio controls both the curvature of F_α and the factor relating its average to the energy.

The same ratio bounds give a lower bound on dissipation:

$$\int \rho w^2 \|\nabla \log r\|^2 dx = \frac{1}{Z_\alpha^2} \int \pi_t r^{2\alpha-3} \|\nabla r\|^2 dx \geq R^{-|2\alpha-3|-2|\alpha-1|} \int \pi_t \|\nabla r\|^2 dx. \tag{40}$$

Applying the Poincaré inequality to r connects this gradient term to Eqn. 39. Thus Eqn. 19 holds with the explicit, conservative constant

$$c_{\alpha,t} = \frac{R^{-(|2\alpha-3|+3|\alpha-1|+|\alpha-2|)}}{C_{P,t}} > 0. \tag{41}$$

For regular fixed-noise solutions satisfying the maximum principle, an initial ratio bound is preserved: dividing Eqn. 10 by π_t gives $\partial_\tau r = \pi_t^{-1} \nabla \cdot (\pi_t w \nabla r)$, whose spatial maxima decrease and minima increase. This provides a way to obtain a time-independent R from the initialization. In strongly imbalanced mixtures, large R and a large Poincaré constant can make the resulting rate small. Smoothness and positivity alone leave these quantitative bounds undetermined, so their validity must be assessed for the target and initialization under study.

Whenever Eqn. 19 holds along the flow, multiplying the dissipation identity by $e^{2c_{\alpha,t}\tau}$ yields

$$\frac{d}{d\tau} \left[e^{2c_{\alpha,t}\tau} \mathcal{E}_{\alpha,t}(\rho_\tau) \right] \leq 0, \qquad \mathcal{E}_{\alpha,t}(\rho_\tau) \leq e^{-2c_{\alpha,t}\tau} \mathcal{E}_{\alpha,t}(\rho_0). \tag{42}$$

The sufficient conditions above make the exponential rate precise while retaining its dependence on the noise level, order, and initial mismatch.

C.2 Dissipation along denoising trajectory

Common transport preserves the density ratio along trajectories. We consider the transport shared by the sampling and target densities:

$$\partial_u \varrho_u = -\nabla \cdot (\varrho_u v_t), \qquad \partial_u \pi_t = -\nabla \cdot (\pi_t v_t), \qquad t = T - u. \tag{43}$$

Along a trajectory with velocity v_t, both densities undergo the same local compression or expansion. Dividing the two equations therefore gives

$$(\partial_u + v_t \cdot \nabla) r_u = 0, \qquad r_u = \frac{\varrho_u}{\pi_t}. \tag{44}$$

The ratio carried by each trajectory stays fixed. Consequently,

$$\partial_u (\pi_t r_u^\alpha) = -\nabla \cdot (\pi_t r_u^\alpha v_t), \qquad \frac{d}{du} \int \pi_t r_u^\alpha \, dx = 0. \tag{45}$$

Since the Rényi energy is a scaled logarithm of this integral, its common-transport contribution is zero. At $\alpha = 1$, applying the same argument to $\pi_t r_u \log r_u$ gives the KL result.

Relaxation supplies the energy decrease. The remaining term in Eqn. 21 changes the sampling density at the instantaneous target. Using $\nabla \psi_u = w_u \nabla \log r_u$, integration by parts gives

$$\begin{aligned} \frac{d}{du} \mathcal{E}_{\alpha,t}(\varrho_u) &= \int \psi_u \nabla \cdot \left[\varrho_u \left(\frac{\beta(t)}{2} + \lambda_u w_u \right) \nabla \log r_u \right] dx \\ &= - \int \varrho_u w_u \left(\frac{\beta(t)}{2} + \lambda_u w_u \right) \|\nabla \log r_u\|^2 dx. \end{aligned} \tag{46}$$

Expanding the last line proves Eqn. 22. Define its total dissipation rate by

$$\mathcal{D}_u = \frac{\beta(t)}{2} \int \varrho_u w_u \|\nabla \log r_u\|^2 dx + \lambda_u \int \varrho_u w_u^2 \|\nabla \log r_u\|^2 dx. \tag{47}$$

Integration over reverse time yields

$$\mathcal{E}_{\alpha,T-U}(\varrho_U) + \int_0^U \mathcal{D}_u \, du = \mathcal{E}_{\alpha,T}(\varrho_0). \tag{48}$$

This shows that the initial discrepancy equals the remaining discrepancy plus the amount removed during denoising. If $\varrho_0 = \pi_T$, nonnegativity of the energy keeps it at zero, so the ideal dynamics follow the intended marginals. Otherwise, the initial mismatch is an upper bound for the subsequent energy.

C.3 Estimated mobility and score error

How score error affects dissipation. Starting from Eqn. 11, the evolution of ρ_τ with estimated mobility and score is

$$\begin{aligned}\partial_\tau \rho_\tau &= \nabla \cdot [\widehat{w}_\tau \nabla \rho_\tau - \rho_\tau \widehat{w}_\tau s_\theta] \\ &= \nabla \cdot \left[\rho_\tau \widehat{w}_\tau \left(\underbrace{\nabla \log \rho_\tau - \nabla \log \pi_t}_{=\nabla \log r_\tau} + \underbrace{\nabla \log \pi_t}_{=s_t} - s_\theta \right) \right] \\ &= \nabla \cdot \left[\rho_\tau \widehat{w}_\tau \left(\nabla \log r_\tau - (s_\theta - s_t) \right) \right]. \end{aligned} \tag{49}$$

Multiplying Eqn. 49 by the true energy variation from Eqn. 9 and integrating by parts gives

$$\frac{d}{d\tau}\mathcal{E}_{\alpha,t}(\rho_\tau) = -\int \rho_\tau w_\tau \widehat{w}_\tau \|\nabla \log r_\tau\|^2 dx + \int \rho_\tau w_\tau \widehat{w}_\tau \nabla \log r_\tau \cdot (s_\theta - s_t)\, dx. \tag{50}$$

This proves Eqn. 23. The first term is nonpositive because the mobility is positive. The second term can oppose descent when the score error points along the true ratio gradient.

Bounding the inner product by half the sum of the squared norms yields

$$\frac{d}{d\tau}\mathcal{E}_{\alpha,t}(\rho_\tau) \le -\frac{1}{2}\int \rho_\tau w_\tau \widehat{w}_\tau \|\nabla \log r_\tau\|^2 dx + \frac{1}{2}\int \rho_\tau w_\tau \widehat{w}_\tau \|s_\theta - s_t\|^2 dx. \tag{51}$$

Energy decreases whenever the remaining correction strength exceeds this weighted score error. Errors in the mobility estimation affect the weight of both terms. We quantify its effect in Appendix C.4.

Accumulated score error along the denoising trajectory. Along the denoising schedule $t = T - u$, the shared probability-flow transport still cancels from the energy derivative. Using the learned score in both the reverse process and the additional corrector gives

$$\frac{d}{du}\mathcal{E}_{\alpha,t}(\varrho_u) = -\widehat{\mathcal{D}}_u + \int \varrho_u w_u (\beta(t) + \lambda_u \widehat{w}_u) \nabla \log r_u \cdot (s_\theta - s_t)\, dx, \tag{52}$$

$$\widehat{\mathcal{D}}_u = \int \varrho_u w_u \left(\frac{\beta(t)}{2} + \lambda_u \widehat{w}_u \right) \|\nabla \log r_u\|^2 dx. \tag{53}$$

The score-error coefficient combines the contribution $\beta(t)$ from the reverse SDE with $\lambda_u \widehat{w}_u$ from the corrector. Applying the same squared-norm bound gives

$$\frac{d}{du}\mathcal{E}_{\alpha,t}(\varrho_u) \le -\frac{1}{2}\widehat{\mathcal{D}}_u + \frac{1}{2}\mathcal{B}_u, \tag{54}$$

$$\mathcal{B}_u = \int \varrho_u w_u \frac{(\beta(t) + \lambda_u \widehat{w}_u)^2}{\beta(t)/2 + \lambda_u \widehat{w}_u} \|s_\theta - s_t\|^2 dx. \tag{55}$$

Integrating over reverse time yields

$$\mathcal{E}_{\alpha,T-U}(\varrho_U) + \frac{1}{2}\int_0^U \widehat{\mathcal{D}}_u\, du \le \mathcal{E}_{\alpha,T}(\varrho_0) + \frac{1}{2}\int_0^U \mathcal{B}_u\, du. \tag{56}$$

The remaining discrepancy is therefore controlled by initialization error and accumulated weighted score error.

C.4 How log-ratio error changes the dissipation rate

Fixed noise dissipation We suppose the logit can be decomposed into the form

$$\widehat{\ell}_\tau(x) = \log r_\tau(x) + c_\tau + h_\tau(x), \qquad |h_\tau(x)| \le \varepsilon_\tau, \tag{57}$$

where c_τ is a spatially constant offset, and h_τ represents the spatially varying error in the log-ratio estimation with upper bound denoted as ε_τ. Its population normalizer is given by $\widehat{Z}_\tau = \mathbb{E}_{\rho_\tau} \exp((\alpha - 1)\widehat{\ell}_\tau)$. The empirical normalizer in Eqn. 14 can be written as $\widetilde{Z}_\tau = \widehat{Z}_\tau e^{\xi_\tau}$, where ξ_τ

parametrizes the sampling error for the empirical normalizer with upper bound denoted as $|\xi_\tau| \le \nu_\tau$. This separates the log-ratio error from the normalization error. Both bounds depend on estimator quality and the reference particles used for normalization.

Because $\mathbb{E}_{\rho_\tau} w_\tau = 1$, the measure $\mu_\tau(dx) = \rho_\tau(x) w_\tau(x) dx$ has unit mass. The spatially constant offset c_τ cancels as in Eqn. 34. We have

$$\begin{aligned} \widehat{w}_\tau &= \frac{r_\tau^{\alpha-1} e^{(\alpha-1)h_\tau}}{\mathbb{E}_{\rho_\tau}[r_\tau^{\alpha-1} e^{(\alpha-1)h_\tau}]} e^{-\xi_\tau} \\ w_\tau &= \frac{r_\tau^{\alpha-1}}{\mathbb{E}_{\rho_\tau}[r_\tau^{\alpha-1}]} \\ \frac{\widehat{w}_\tau(x)}{w_\tau(x)} &= \frac{e^{(\alpha-1)h_\tau(x)}}{\mathbb{E}_{\mu_\tau} e^{(\alpha-1)h_\tau}} e^{-\xi_\tau}. \end{aligned} \tag{58}$$

Both the numerator and the averaged denominator lie between $e^{-|\alpha-1|\varepsilon_\tau}$ and $e^{|\alpha-1|\varepsilon_\tau}$. Therefore,

$$e^{-q_\tau} \le \frac{\widehat{w}_\tau}{w_\tau} \le e^{q_\tau}, \qquad q_\tau = 2|\alpha-1|\varepsilon_\tau + \nu_\tau. \tag{59}$$

Multiplying by the nonnegative integrand in the ideal dissipation gives

$$\begin{aligned} e^{-q_\tau} \int \rho_\tau w_\tau^2 \|\nabla \log r_\tau\|^2 dx &\le \int \rho_\tau w_\tau \widehat{w}_\tau \|\nabla \log r_\tau\|^2 dx \\ &\le e^{q_\tau} \int \rho_\tau w_\tau^2 \|\nabla \log r_\tau\|^2 dx. \end{aligned} \tag{60}$$

Thus accurate log ratios keep the dissipation rate close to the ideal rate at the same sampling density. Orders farther from one amplify spatial log-ratio errors through the factor $|\alpha - 1|$.

Dissipation along denoising trajectory The moving-target version follows by replacing ρ_τ with ϱ_u and applying the same bound at each reverse time. For q_u defined analogously,

$$e^{-q_u} \mathcal{D}_u \le \widehat{\mathcal{D}}_u \le e^{q_u} \mathcal{D}_u, \qquad \frac{d}{du} \mathcal{E}_{\alpha,t}(\varrho_u) \le -\frac{e^{-q_u}}{2} \mathcal{D}_u + \frac{1}{2} \mathcal{B}_u. \tag{61}$$

As before, the log-ratio error changes the dissipation strength, while score error contributes through $\mathcal{B}_u$.

D Effects of Rényi orders

Small orders emphasize underrepresented regions. For $0 < \alpha < 1$, Rényi skew symmetry gives $\mathcal{E}_{\alpha,t}(\rho) = D_{1-\alpha}(\pi_t \| \rho)/(1-\alpha)$ (van Erven & Harremoës, 2014). Thus, for smooth positive densities with finite target-first KL divergence,

$$\lim_{\alpha \to 0^+} \mathcal{E}_{\alpha,t}(\rho) = D_{\mathrm{KL}}(\pi_t \| \rho), \qquad \lim_{\alpha \to 0^+} w_\alpha = \frac{\pi_t}{\rho}. \tag{62}$$

Here w_α denotes the mobility at a fixed density as the order varies. The limiting energy generates

$$\partial_\tau \rho_\tau = \nabla \cdot \left[\pi_t \nabla \log \frac{\rho_\tau}{\pi_t} \right], \qquad \frac{d}{d\tau} D_{\mathrm{KL}}(\pi_t \| \rho_\tau) = -\int \frac{\pi_t^2}{\rho_\tau} \|\nabla \log r_\tau\|^2 dx. \tag{63}$$

The factor $1/\alpha$ in the energy retains this nontrivial limit. The mobility increases where the sampler underrepresents the target, providing a motivation for minority-mode experiments. Severe undercoverage also makes the coefficients large. In the extreme case, a zero sampling density for a set of positive target probability gives infinite target-first KL divergence. Transfer between modes still depends on spatial gradients and the regions connecting them.

Large orders concentrate the mobility. Writing $Z_\alpha = \int \pi_t r^\alpha dx$, the mobility-weighted density is $\rho w_\alpha = \pi_t r^\alpha / Z_\alpha$. It increasingly emphasizes the largest density ratios. If $R_\infty = \operatorname{ess\,sup} r < \infty$, then $D_\alpha(\rho \| \pi_t) \to \log R_\infty$ (van Erven & Harremoës, 2014), so $\mathcal{E}_{\alpha,t}(\rho) \to 0$, while ρw_α concentrates near the largest ratios. This concentration can make the limiting correction singular.

E ADDITIONAL EXPERIMENTAL RESULTS

E.1 GMMS WITH UNDERREPRESENTED AND OVERREPRESENTED RARE MODES

The two GMM archives evaluate rare probabilities $\{0.001, 0.003, 0.005, 0.01, 0.02, 0.05\}$ in the underrepresentation sweep and $\{0.003, 0.005, 0.01\}$ in the overrepresentation sweep. Each mixture has eight components, with equal probabilities for the seven common components. We report the seven orders $\{0.7, 0.8, 0.9, 1, 1.1, 1.2, 1.3\}$ and fit sensitivities over $[0.7, 1.3]$ for underrepresentation and $[0.9, 1.1]$ for overrepresentation. With five sampling seeds per setting, this gives 210 reported runs in the underrepresentation sweep and 105 in the overrepresentation sweep. The overrepresentation sweep uses exponentially tilted training weights to investigate the opposite extreme of rare-mode imbalance; in the underrepresentation sweep, the three larger rare probabilities produce near-target baseline ratios.

Training and sampling protocol. Both sweeps use standardized $d = 8$ mixtures with 100,000 training samples, 48,000 generated samples per seed, $T = 250$ predictor steps, and continuous VP rate $\beta(t) = 0.1 + 19.9t$ with per-step variance $\beta(t_i)/T$. Data are centered and divided by a global root-mean-square coordinate standard deviation, preserving isotropic component shapes, and samples are transformed back for evaluation. Corrections precede prediction at every level, with $K = 1$, $\mathrm{snr} = 0.16$, and mobility clipped to $[0.1, 10]$. Density-ratio estimation combines the analytic noisy-target density and Kernel density estimation with Scott's bandwidth rule. To induce overrepresentation, we apply exponential tilt weighting during training while keeping the remaining setup the same as in the underrepresentation sweep. The correction target and evaluation weights are the original mixture probabilities, so these runs test whether sampling corrections can reduce the resulting excess rare-mode mass.

Consistent mode-mass measurements. The stored rare recall is conditional on assignment to a mode. For each seed, we compute the unconditional generated-to-target rare-mass ratio as

$$r_0(\alpha) = \frac{\text{rare recall}(\alpha)\,[1 - \text{unassigned fraction}(\alpha)]}{p_0}. \tag{64}$$

This agrees with the explicitly stored mode-zero ratios in the underrepresentation archive over the reported range. Using the same definition in both archives avoids comparing conditional and unconditional probabilities. In the available GMM implementation, assignment uses the nearest center with a distance cutoff of half the minimum inter-center separation. The reported TV is

$$\mathrm{TV}_{\mathrm{mode}} = \frac{1}{2}\left(\sum_{k=0}^{7} |m_k - p_k| + m_{\mathrm{unassigned}}\right), \tag{65}$$

where m_k is the fraction of all generated samples assigned to mode k. The normalized radius divides the mean distance of samples assigned to a mode by the expected radius $\sigma\sqrt{2}\,\Gamma((d+1)/2)/\Gamma(d/2)$ of its isotropic Gaussian component. Radius and common-mode statistics are supplied only for the underrepresentation sweep. All tabulated uncertainties are sampling-seed standard deviations except the fitted response sensitivities, whose uncertainty is a standard error.

Estimating the response sensitivity. Equation equation 24 motivates a linear approximation anchored at the Langevin baseline. For each sampling seed, we fit a slope with respect to α and apply the baseline sign, using $\mathcal{A} = \{0.7, 0.8, 0.9, 1, 1.1, 1.2, 1.3\}$ for underrepresentation and $\mathcal{A} = \{0.9, 1, 1.1\}$ for overrepresentation:

$$\widehat{\kappa}_d = \mathrm{sgn}(\log r_d(1)) \frac{\sum_{\alpha\in\mathcal{A}} (\alpha - 1)\big[\log r_d(\alpha) - \log r_d(1)\big]}{\sum_{\alpha\in\mathcal{A}} (\alpha - 1)^2}. \tag{66}$$

The fraction is the baseline-anchored least-squares slope with respect to α; multiplying by the sign gives the sensitivity with respect to h. The corresponding approximation is $\log r_d(1) + \widehat{\kappa}_d h$, where $h = \mathrm{sgn}(\log r_d(1))(\alpha - 1)$. Because both order grids are symmetric around one, $\sum_{\alpha\in\mathcal{A}}(\alpha - 1) = 0$: fixing the intercept at the observed baseline gives the same slope with respect to α as ordinary least squares with a free intercept. The seven-order fit pools information across the full underrepresentation sweep. For the overrepresented rare modes, the baseline sign is positive, so the

three-order fit equals the centered difference $[\log r_d(1.1) - \log r_d(0.9)]/0.2$, describing the response closer to one. The reported sensitivities therefore summarize different fitting intervals across the two regimes. The resulting slope estimates sensitivity over a finite interval; its interpretation as the derivative with respect to h at zero depends on the accuracy of the first-order approximation. The Taylor remainder is a local statement and does not guarantee small errors throughout $[0.7, 1.3]$.

For Figure 1 and Table 2, we compute the signed estimate in Eqn. equation 66 separately for each seed, then average these estimates and compute their standard error across seeds. The sign is taken before averaging, which matters for modes whose baseline discrepancy changes sign between seeds. Each plotted point uses the log of the mean baseline ratio horizontally and the mean signed sensitivity vertically. Table 2 uses these same signed estimates and reports their sample standard deviation divided by $\sqrt{5}$ as the sampling standard error. The signed-sensitivity regression across 48 underrepresentation mode–condition pairs has intercept 0.0198, slope -1.0033, and Pearson correlation -0.8730; restricting to the 42 common-mode pairs gives slope -0.0297. The three oversampled rare-mode conditions give intercept -0.0475, slope -0.7209, and correlation -0.8758. These are descriptive associations: modes within each condition share a model and probability constraints, and the overrepresentation archive contains only three retained rare-mode conditions. Common-mode ratios are unavailable for that archive.

Interpreting the sign-directed adjustment. For a nonzero baseline discrepancy, substituting $\alpha - 1 = \mathrm{sgn}(\log r_d(1))h$ into the Taylor expansion gives Eqn. equation 24. The signed sensitivity is

$$\kappa_d = \left.\frac{d \log r_d(\alpha_d(h))}{dh}\right|_{h=0} = \mathrm{sgn}(\log r_d(1)) \left.\frac{d \log r_d(\alpha)}{d\alpha}\right|_{\alpha=1}. \tag{67}$$

The upper-left and lower-right quadrants of the sensitivity–discrepancy plane represent a local response toward zero, since $\log r_d(1)\kappa_d < 0$. For the squared log discrepancy,

$$\left.\frac{d}{dh}\frac{1}{2}\left[\log r_d(1 + \mathrm{sgn}(\log r_d(1))h)\right]^2\right|_{h=0} = \log r_d(1)\kappa_d. \tag{68}$$

Thus, $\log r_d(1)\kappa_d < 0$ implies a local reduction in this discrepancy as h increases. If $\log r_d(1) = 0$, the baseline mode mass is already correct and the sign rule proposes no adjustment. The statement concerns the response of a particular mode; a single sampler order can affect different modes differently.

The mobility interpretation concerns pointwise density ratios. At a fixed sampling law, two locations satisfy $w(x)/w(y) = [r(x)/r(y)]^{\alpha-1}$, which explains why smaller and larger orders favor different regions in relative correction strength. Mode masses aggregate over regions, and their changes depend on probability flux across boundaries. The observed signed rare-mode sensitivities support movement toward the target in these experiments; the pointwise weighting by itself does not establish their signs.

Table 5: Underrepresentation GMM sweep, 0.001, 0.003, 0.005 target rare probabilities. All seven orders in $[0.7, 1.3]$ are shown. Entries are means $\pm$ sampling-seed standard deviations; mass ratios are unconditional and unassigned fractions are scaled by 10^4.

p	α	Rare mass / target	Rare radius	TV	Unassigned $\times 10^4$
0.001	0.7	0.417 ± 0.133	1.920 ± 0.198	0.02378 ± 0.00316	0.50 ± 0.28
0.001	0.8	0.400 ± 0.068	1.861 ± 0.153	0.02471 ± 0.00091	0.71 ± 0.32
0.001	0.9	0.396 ± 0.061	1.758 ± 0.057	0.02438 ± 0.00122	0.33 ± 0.32
0.001	1	0.346 ± 0.060	1.565 ± 0.057	0.02553 ± 0.00163	0.04 ± 0.09
0.001	1.1	0.250 ± 0.086	1.454 ± 0.037	0.02882 ± 0.00083	0.04 ± 0.09
0.001	1.2	0.225 ± 0.084	1.512 ± 0.210	0.02650 ± 0.00102	0.08 ± 0.11
0.001	1.3	0.146 ± 0.000	1.256 ± 0.202	0.02927 ± 0.00121	0.12 ± 0.19
0.003	0.7	0.467 ± 0.086	1.340 ± 0.059	0.02338 ± 0.00154	0.79 ± 0.40
0.003	0.8	0.429 ± 0.104	1.371 ± 0.088	0.02462 ± 0.00270	0.62 ± 0.53
0.003	0.9	0.376 ± 0.032	1.295 ± 0.039	0.02455 ± 0.00111	0.21 ± 0.15
0.003	1	0.372 ± 0.031	1.264 ± 0.050	0.02605 ± 0.00094	0.08 ± 0.11
0.003	1.1	0.354 ± 0.042	1.247 ± 0.040	0.02735 ± 0.00284	0.00 ± 0.00
0.003	1.2	0.293 ± 0.058	1.235 ± 0.062	0.02959 ± 0.00052	0.08 ± 0.11
0.003	1.3	0.269 ± 0.067	1.250 ± 0.048	0.03020 ± 0.00227	0.12 ± 0.11
0.005	0.7	0.600 ± 0.038	1.276 ± 0.027	0.05323 ± 0.00269	0.58 ± 0.52
0.005	0.8	0.558 ± 0.029	1.242 ± 0.038	0.05452 ± 0.00224	0.71 ± 0.48
0.005	0.9	0.542 ± 0.026	1.245 ± 0.024	0.05648 ± 0.00195	0.54 ± 0.11
0.005	1	0.534 ± 0.019	1.207 ± 0.015	0.05652 ± 0.00278	0.04 ± 0.09
0.005	1.1	0.528 ± 0.042	1.229 ± 0.042	0.05704 ± 0.00144	0.04 ± 0.09
0.005	1.2	0.474 ± 0.025	1.253 ± 0.024	0.05991 ± 0.00402	0.29 ± 0.24
0.005	1.3	0.437 ± 0.018	1.216 ± 0.060	0.05986 ± 0.00229	0.63 ± 0.39

Table 6: Underrepresentation GMM sweep, 0.01, 0.02, 0.05 target rare probabilities. All seven orders in $[0.7, 1.3]$ are shown. Entries are means $\pm$ sampling-seed standard deviations; mass ratios are unconditional and unassigned fractions are scaled by 10^4.

p	α	Rare mass / target	Rare radius	TV	Unassigned $\times 10^4$
0.01	0.7	1.053 ± 0.066	1.115 ± 0.012	0.03465 ± 0.00251	0.46 ± 0.34
0.01	0.8	1.048 ± 0.010	1.104 ± 0.011	0.03613 ± 0.00168	0.58 ± 0.09
0.01	0.9	1.032 ± 0.030	1.105 ± 0.011	0.03634 ± 0.00124	0.50 ± 0.19
0.01	1	1.048 ± 0.046	1.090 ± 0.008	0.03782 ± 0.00239	0.04 ± 0.09
0.01	1.1	1.008 ± 0.036	1.083 ± 0.011	0.03760 ± 0.00226	0.00 ± 0.00
0.01	1.2	0.974 ± 0.061	1.088 ± 0.015	0.03842 ± 0.00179	0.00 ± 0.00
0.01	1.3	0.943 ± 0.052	1.099 ± 0.013	0.04105 ± 0.00152	0.00 ± 0.00
0.02	0.7	1.092 ± 0.028	1.080 ± 0.010	0.03171 ± 0.00184	0.38 ± 0.34
0.02	0.8	1.057 ± 0.025	1.081 ± 0.007	0.03300 ± 0.00331	0.58 ± 0.40
0.02	0.9	1.059 ± 0.014	1.075 ± 0.009	0.03300 ± 0.00157	0.17 ± 0.17
0.02	1	1.069 ± 0.050	1.069 ± 0.008	0.03298 ± 0.00089	0.04 ± 0.09
0.02	1.1	1.030 ± 0.038	1.073 ± 0.005	0.03390 ± 0.00183	0.00 ± 0.00
0.02	1.2	1.059 ± 0.022	1.068 ± 0.016	0.03698 ± 0.00199	0.00 ± 0.00
0.02	1.3	1.038 ± 0.014	1.073 ± 0.010	0.03684 ± 0.00237	0.00 ± 0.00
0.05	0.7	1.009 ± 0.026	1.087 ± 0.009	0.04897 ± 0.00353	0.67 ± 0.40
0.05	0.8	0.994 ± 0.023	1.094 ± 0.004	0.04995 ± 0.00285	0.54 ± 0.11
0.05	0.9	0.978 ± 0.013	1.085 ± 0.002	0.05288 ± 0.00170	0.38 ± 0.27
0.05	1	0.991 ± 0.011	1.092 ± 0.006	0.05075 ± 0.00114	0.04 ± 0.09
0.05	1.1	0.967 ± 0.022	1.085 ± 0.007	0.05282 ± 0.00135	0.00 ± 0.00
0.05	1.2	0.955 ± 0.021	1.090 ± 0.004	0.05316 ± 0.00255	0.00 ± 0.00
0.05	1.3	0.955 ± 0.015	1.090 ± 0.004	0.05190 ± 0.00292	0.00 ± 0.00

Table 7: Overrepresentation GMM sweep, 0.003, 0.005, 0.01 target rare probabilities. All seven orders in $[0.7, 1.3]$ are shown. Entries are means $\pm$ sampling-seed standard deviations; mass ratios are unconditional and unassigned fractions are scaled by 10^4.

p	α	Rare mass / target	TV	Unassigned $\times 10^4$
0.003	0.7	4.680 ± 0.142	0.04080 ± 0.00244	0.00 ± 0.00
0.003	0.8	4.316 ± 0.221	0.04386 ± 0.00249	0.00 ± 0.00
0.003	0.9	3.919 ± 0.180	0.04272 ± 0.00107	0.00 ± 0.00
0.003	1	3.492 ± 0.130	0.04446 ± 0.00136	0.00 ± 0.00
0.003	1.1	3.232 ± 0.081	0.04329 ± 0.00232	0.00 ± 0.00
0.003	1.2	2.920 ± 0.096	0.04676 ± 0.00191	0.00 ± 0.00
0.003	1.3	2.416 ± 0.080	0.04627 ± 0.00193	0.00 ± 0.00
0.005	0.7	3.230 ± 0.023	0.03214 ± 0.00151	0.04 ± 0.09
0.005	0.8	2.969 ± 0.105	0.03132 ± 0.00148	0.12 ± 0.27
0.005	0.9	2.822 ± 0.066	0.03058 ± 0.00169	0.04 ± 0.09
0.005	1	2.636 ± 0.118	0.03189 ± 0.00146	0.00 ± 0.00
0.005	1.1	2.406 ± 0.065	0.03266 ± 0.00130	0.00 ± 0.00
0.005	1.2	2.124 ± 0.063	0.03363 ± 0.00277	0.00 ± 0.00
0.005	1.3	1.884 ± 0.048	0.03375 ± 0.00221	0.00 ± 0.00
0.01	0.7	3.136 ± 0.103	0.04099 ± 0.00232	0.04 ± 0.09
0.01	0.8	3.034 ± 0.086	0.04117 ± 0.00229	0.00 ± 0.00
0.01	0.9	2.964 ± 0.082	0.04205 ± 0.00059	0.00 ± 0.00
0.01	1	2.770 ± 0.079	0.04180 ± 0.00288	0.00 ± 0.00
0.01	1.1	2.568 ± 0.061	0.04334 ± 0.00211	0.00 ± 0.00
0.01	1.2	2.500 ± 0.014	0.04214 ± 0.00256	0.00 ± 0.00
0.01	1.3	2.323 ± 0.068	0.04351 ± 0.00095	0.00 ± 0.00

E.2 FULL STACKED-MNIST EVALUATION

Data and training. Each image stacks three 28×28 handwritten digits as separate channels. The first two digits are uniform, and the third digit equals 7 with nominal probability 0.01, defining 100 rare combinations among $1{,}000$ modes. The saved composition contains one million digit triples and has rare mass 0.010096; handwriting exemplars are freshly sampled on each access. All configurations use the same 0.81M-parameter diffusion model, trained for $200{,}000$ updates with dataset/training seed 20260912.

Sampling and ratio estimation. The predictor uses $1{,}000$ DDPM steps with a linear variance schedule from 10^{-4} to 0.02 and no x_0 clipping. Corrections precede prediction at 64 evenly spaced noise anchors. Each anchor's nominal correction time is divided across K substeps, with global safeguards able to reduce the realized time. The discriminator is pretrained against baseline noisy trajectories and refreshed against current reference particles at every correction substep. Evaluation outputs are excluded from discriminator training. The sweep contains ancestral sampling and all 18 combinations of $K \in \{1, 2, 3\}$ and $\alpha \in \{0.5, 0.7, 0.9, 1, 1.1, 1.3\}$. Each configuration generates $10{,}000$ outputs per seed for seeds 42–46, giving 95 completed runs with no failed runs or nonfinite samples. All tables report means and standard deviations across these sampling seeds; the checkpoint is fixed.

Metrics and finite-sample coverage. Raw TV compares classifier-assigned combination frequencies with target probabilities. Rare mass is the total generated frequency of combinations ending in 7, and coverage is the fraction of combinations observed at least once. We also report coverage requiring at least five observations per mode. At $10{,}000$ samples, a nominal rare mode of probability 10^{-4} is observed with probability $1 - (1 - 10^{-4})^{10000} \approx 0.63214$. The corresponding nominal expected coverage across all modes is 0.96320. These occupancy references help interpret coverage alongside generated rare mass. Mean raw common-mode coverage ranges from 99.96% to 100% across settings.

A triple passes the confidence filter when all three digit predictions have confidence at least 0.9. Confidence-filtered TV retains rejected triples in a separate bin with target mass zero and uses the original sample denominator. Confident coverage counts modes represented by accepted triples, and confident rare mass counts accepted rare triples relative to all generated triples. Confidence provides an auxiliary classifier-based fidelity diagnostic. The lowest mean raw TV occurs at $\alpha = 0.5$, $K = 1$, and the lowest mean filtered TV at the same order with $K = 3$; Tables 8–10 report the complete sweep.

Correction time. At $K = 2$, $\alpha = 0.5$ and Langevin have nearly identical mean realized correction times (0.588972 and 0.589192). Rényi correction clips 7.81% of weights and caps 0.47% of steps. At $K = 3$, $\alpha = 0.5$, the correction time is 0.588277, compared with 0.589247 for Langevin, while 29.07% of weights are clipped and 2.60% of steps are capped. The lowest-raw-TV setting, $\alpha = 0.5$, $K = 1$, caps 90.94% of steps and has a slightly shorter correction time of 0.574414. Thus, even with similar integrated correction time, substep count and clipping affect the implemented dynamics. Reported sampling takes 62.09 versus 50.35 seconds at $K = 2$, and 70.52 versus 53.14 seconds at $K = 3$, for $\alpha = 0.5$ versus Langevin.

Table 8: Full stacked-MNIST mode-frequency results with 64 correction anchors (mean $\pm$ sampling-seed standard deviation). Rare mass and coverage are percentages; the saved target rare mass is 1.0096%. $K = 0$ is ancestral sampling and $\alpha = 1$ with $K > 0$ is Langevin correction.

K	α	Raw TV	Rare mass (%)	All coverage (%)	Rare coverage (%)
0	1	0.13916 ± 0.00417	2.398 ± 0.159	99.16 ± 0.26	91.60 ± 2.61
1	0.5	0.13572 ± 0.00411	2.388 ± 0.213	98.94 ± 0.32	89.40 ± 3.21
1	0.7	0.13814 ± 0.00440	2.416 ± 0.213	99.02 ± 0.35	90.20 ± 3.49
1	0.9	0.13826 ± 0.00570	2.418 ± 0.217	99.16 ± 0.34	91.80 ± 3.56
1	1	0.13884 ± 0.00665	2.414 ± 0.234	99.16 ± 0.39	91.80 ± 4.15
1	1.1	0.13970 ± 0.00619	2.436 ± 0.230	99.02 ± 0.29	90.60 ± 2.88
1	1.3	0.13822 ± 0.00435	2.416 ± 0.267	99.24 ± 0.26	92.40 ± 2.61
2	0.5	0.13758 ± 0.00079	2.372 ± 0.207	98.86 ± 0.31	88.60 ± 3.13
2	0.7	0.13828 ± 0.00511	2.430 ± 0.180	99.02 ± 0.23	90.20 ± 2.28
2	0.9	0.13796 ± 0.00624	2.392 ± 0.196	99.00 ± 0.34	90.20 ± 3.77
2	1	0.13878 ± 0.00650	2.404 ± 0.224	99.16 ± 0.39	91.80 ± 4.15
2	1.1	0.13944 ± 0.00525	2.440 ± 0.254	99.20 ± 0.25	92.20 ± 2.77
2	1.3	0.13990 ± 0.00291	2.446 ± 0.219	99.18 ± 0.18	92.00 ± 2.00
3	0.5	0.13888 ± 0.00534	2.478 ± 0.158	99.00 ± 0.12	90.00 ± 1.22
3	0.7	0.13938 ± 0.00217	2.386 ± 0.206	99.04 ± 0.11	90.40 ± 1.14
3	0.9	0.13888 ± 0.00526	2.386 ± 0.196	98.94 ± 0.23	89.40 ± 2.30
3	1	0.13892 ± 0.00656	2.412 ± 0.221	99.16 ± 0.34	91.80 ± 3.70
3	1.1	0.13874 ± 0.00414	2.474 ± 0.223	99.20 ± 0.21	92.20 ± 2.17
3	1.3	0.13852 ± 0.00403	2.436 ± 0.173	99.12 ± 0.22	91.40 ± 2.30

Table 9: Full confidence-filtered results with 64 correction anchors (mean $\pm$ sampling-seed standard deviation). Rejection, accepted rare mass, and coverage are percentages. Filtered TV includes the rejection bin.

K	α	Filtered TV	Rejected (%)	Rare mass (%)	All coverage (%)	Rare coverage (%)
0	1	0.43530 ± 0.00259	42.42 ± 0.30	1.030 ± 0.041	96.12 ± 0.44	64.40 ± 5.27
1	0.5	0.42634 ± 0.00465	41.48 ± 0.42	1.048 ± 0.085	95.90 ± 0.51	61.60 ± 5.13
1	0.7	0.42504 ± 0.00421	41.24 ± 0.48	1.052 ± 0.125	95.70 ± 0.27	59.40 ± 2.70
1	0.9	0.42326 ± 0.00672	41.11 ± 0.72	1.070 ± 0.090	95.90 ± 0.62	61.80 ± 4.97
1	1	0.42470 ± 0.00429	41.26 ± 0.43	1.070 ± 0.110	96.00 ± 0.47	63.00 ± 4.18
1	1.1	0.42306 ± 0.00473	41.12 ± 0.38	1.090 ± 0.122	96.24 ± 0.35	65.20 ± 3.90
1	1.3	0.42674 ± 0.00311	41.49 ± 0.30	1.100 ± 0.123	96.32 ± 0.24	65.40 ± 2.88
2	0.5	0.42140 ± 0.00383	40.93 ± 0.29	1.012 ± 0.116	95.80 ± 0.19	60.80 ± 1.79
2	0.7	0.42184 ± 0.00313	40.98 ± 0.23	1.068 ± 0.128	96.00 ± 0.28	63.00 ± 2.92
2	0.9	0.42370 ± 0.00448	41.13 ± 0.44	1.062 ± 0.107	95.92 ± 0.27	62.20 ± 3.03
2	1	0.42450 ± 0.00438	41.25 ± 0.42	1.056 ± 0.113	95.94 ± 0.47	62.60 ± 3.97
2	1.1	0.42408 ± 0.00428	41.27 ± 0.29	1.060 ± 0.134	96.08 ± 0.44	63.60 ± 4.28
2	1.3	0.42614 ± 0.00539	41.50 ± 0.49	1.076 ± 0.102	96.22 ± 0.45	65.80 ± 4.32
3	0.5	0.41900 ± 0.00207	40.69 ± 0.19	1.042 ± 0.071	96.06 ± 0.22	63.80 ± 2.59
3	0.7	0.41956 ± 0.00275	40.66 ± 0.18	1.018 ± 0.101	95.94 ± 0.21	62.40 ± 2.41
3	0.9	0.42250 ± 0.00432	41.08 ± 0.29	1.046 ± 0.099	95.80 ± 0.32	61.80 ± 3.83
3	1	0.42418 ± 0.00429	41.20 ± 0.43	1.054 ± 0.109	95.90 ± 0.42	62.40 ± 3.65
3	1.1	0.42502 ± 0.00442	41.38 ± 0.30	1.112 ± 0.161	96.28 ± 0.45	65.40 ± 4.16
3	1.3	0.42562 ± 0.00383	41.39 ± 0.30	1.096 ± 0.136	96.00 ± 0.38	64.20 ± 3.63

Table 10: Timing diagnostics with 64 correction anchors (mean ± sampling-seed standard deviation). Correction time is accumulated numerical integration time. Sampling seconds reproduce the reported timer.

K	α	Sampling (s)	Correction time
0	1	47.47 ± 11.70	0.00000 ± 0.00000
1	0.5	53.44 ± 0.49	0.57441 ± 0.00034
1	0.7	53.38 ± 0.12	0.57781 ± 0.00019
1	0.9	53.26 ± 0.20	0.58903 ± 0.00016
1	1	49.73 ± 0.02	0.58903 ± 0.00016
1	1.1	53.26 ± 0.13	0.58903 ± 0.00016
1	1.3	53.38 ± 0.07	0.58902 ± 0.00016
2	0.5	62.09 ± 0.24	0.58897 ± 0.00034
2	0.7	62.03 ± 0.19	0.58916 ± 0.00017
2	0.9	62.05 ± 0.07	0.58919 ± 0.00016
2	1	50.35 ± 0.05	0.58919 ± 0.00016
2	1.1	61.80 ± 0.22	0.58919 ± 0.00016
2	1.3	61.88 ± 0.29	0.58919 ± 0.00015
3	0.5	70.52 ± 0.45	0.58828 ± 0.00091
3	0.7	70.57 ± 0.26	0.58920 ± 0.00018
3	0.9	70.74 ± 0.19	0.58924 ± 0.00017
3	1	53.14 ± 0.25	0.58925 ± 0.00016
3	1.1	70.55 ± 0.45	0.58925 ± 0.00016
3	1.3	70.60 ± 0.23	0.58924 ± 0.00015

E.3 FACES: MIXTURE PROPORTIONS AND CLASS-SPECIFIC FIDELITY

The target mixture contains 5% AFHQ cat faces and 95% CelebA human faces. We use the same pretrained diffusion model for ancestral sampling and all correction settings, with a discriminator estimating density ratios. The reported sweep uses five sampling seeds and contains ancestral sampling and 15 corrected configurations: $K \in \{1, 2, 3\}$ and $\alpha \in \{0.7, 0.9, 1, 1.1, 1.3\}$. We report generated cat counts per 500 images and separate FIDs for cats and humans. A cat count of 25 corresponds to the target proportion.

Full results. Table 11 reproduces the supplied summary. The $\pm$ values describe variability across five sampling seeds for cat counts and cat FID; no uncertainty estimates were provided for human FID. The summary does not specify whether the variability denotes standard deviations or standard errors, or the generated and reference sample counts used for each FID. In particular, "cats per 500" specifies the frequency metric's scale and does not establish the sample count used to compute FID. These comparisons summarize observed means.

Table 11: Complete faces sweep for the 5% cat mixture. Cat count and cat FID include the reported variability across five seeds; human FID is reported without uncertainty. The target is 25 cats per 500 images.

K	α	Cats per 500	Cat FID	Human FID
0	1	19.0 ± 3.4	47.13 ± 1.21	17.37
1	0.7	23.2 ± 3.2	45.32 ± 0.95	17.35
1	0.9	21.6 ± 3.1	45.96 ± 0.92	17.45
1	1	18.7 ± 3.4	46.88 ± 1.03	17.38
1	1.1	18.5 ± 3.2	46.28 ± 1.33	17.48
1	1.3	17.3 ± 2.8	46.54 ± 1.11	17.45
2	0.7	23.2 ± 3.0	45.16 ± 0.84	17.45
2	0.9	22.6 ± 3.2	45.39 ± 0.93	17.45
2	1	19.3 ± 3.5	46.53 ± 0.98	17.45
2	1.1	19.2 ± 3.3	46.17 ± 1.03	17.38
2	1.3	17.2 ± 3.2	46.66 ± 0.99	17.45
3	0.7	23.0 ± 3.3	45.02 ± 0.94	17.40
3	0.9	22.7 ± 3.4	45.31 ± 0.83	17.54
3	1	19.2 ± 3.4	46.35 ± 1.11	17.44
3	1.1	18.8 ± 3.2	45.90 ± 1.23	17.45
3	1.3	17.0 ± 3.3	46.10 ± 1.18	17.43

Ancestral sampling underrepresents cats, producing 19.0 per 500 images (3.80%) against the target of 25 (5%). At $\alpha = 0.7$, mean counts increase to 23.2, 23.2, and 23.0 for $K = 1$, 2, and 3, recovering approximately 67–70% of the ancestral rare-class deficit. Langevin correction yields 18.7–19.3 cats, while $\alpha = 1.3$ yields 17.0–17.3, increasing the deficit. Thus smaller orders improve agreement with the target mixture proportion in this experiment.

At $\alpha = 0.7$, mean cat FIDs are 45.32, 45.16, and 45.02 for $K = 1$, 2, and 3, compared with 47.13 for ancestral sampling. All corrected configurations have lower mean cat FID than the ancestral baseline, while human FID spans 17.35–17.54. The $K = 2$, $\alpha = 0.7$ setting used for the qualitative examples increases the cat proportion from 3.80% to 4.64% and lowers mean cat FID by 4.18%; $K = 3$ gives the lowest reported mean cat FID.