CLASSIFIER-FREE GUIDANCE IN FLOW MATCHING: NON-AUTONOMOUS POTENTIALS, OVERSHOOT, AND POSTERIOR-MEAN CONTROL

Anonymous authors
Paper under double-blind review

ABSTRACT

Classifier-free guidance (CFG) improves conditional generation in Flow Matching, but strong guidance can distort the generated distribution and reduce diversity. We provide a geometric account of this behavior by viewing Flow Matching as a time-varying gradient flow and characterizing how CFG reshapes its underlying potential. This view explains how stronger alignment can be accompanied by mean displacement and trajectory concentration, and motivates controlling guidance through the model-implied terminal posterior mean. We therefore propose Posterior-Mean-Capped CFG (PMC-CFG), a training-free, per-sample controller that imposes a pointwise relative bound on the implied terminal mean while retaining the strongest feasible guidance, without extra network evaluations. Experiments on synthetic and large-scale image-generation benchmarks show that PMC-CFG limits guidance-induced distortion and concentration while improving the alignment–diversity trade-off, with particularly strong benefits when nominal guidance is large.

1 INTRODUCTION

Diffusion and flow-based generative models learn to transform simple noise into complex data distributions (Ho et al., 2020; Lipman et al., 2023; Albergo & Vanden-Eijnden, 2022; Liu et al., 2023). For conditional generation, classifier guidance steers this transformation with the gradient of an auxiliary classifier (Dhariwal & Nichol, 2021), whereas classifier-free guidance (CFG) jointly learns conditional and unconditional predictions and combines them during sampling (Ho & Salimans, 2022). CFG is simple and effective: its conditional–unconditional difference acts as an implicit classifier signal, and guided flows have shown substantial gains in conditional generation with Flow Matching (Ho & Salimans, 2022; Zheng et al., 2023). Yet increasing its scale improves condition alignment at the cost of diversity and can cause oversaturation, mean overshoot, and distributional shrinkage (Ho & Salimans, 2022; Karras et al., 2024; Pavasovic et al., 2026).

Existing inference-time remedies schedule or restrict guidance (Wang et al., 2024; Gao et al., 2026; Jiang & Ma, 2026; Kynkäänniemi et al., 2024), constrain or rescale its update (Chung et al., 2025; Xia et al., 2025; Sadat et al., 2025), or apply nonlinear reweighting (Pavasovic et al., 2026). Related control views treat the conditional–unconditional discrepancy as feedback (Wang et al., 2026), adapt the scale to time and state (Azangulov et al., 2026), or project toward an agreement manifold (Cai et al., 2026). In contrast, we retain the original guidance direction but impose a closed-form, per-sample bound on the terminal mean implied by a Flow-Matching velocity. This differs from time-only scheduling and iterative projection, and avoids treating the guided dynamics as a single power-tilted terminal distribution, an identification that is not generally valid (Bradley & Nakkiran, 2024).

We answer this question in posterior-mean space. For the linear Flow-Matching path, every velocity evaluation implies a posterior estimate of the terminal sample, and fixed CFG extrapolates this estimate along the conditional–unconditional direction. This view motivates Posterior-Mean-Capped CFG (PMC-CFG), which intersects the nominal CFG ray with a relative norm ball and selects the farthest feasible point. Thus its distinction is a pointwise feasibility guarantee, not merely a preferred average schedule. As illustrated in Figure 1, PMC-CFG limits excessive extrapolation while retaining nominal guidance whenever it is already feasible. The controller requires no training and adds no network evaluations.

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.

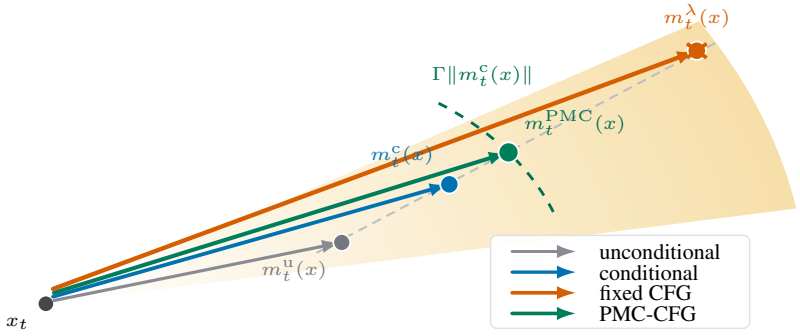


Figure 1: Posterior-mean control. Fixed CFG extrapolates from $m_t^{\mathrm{u}}(x)$ through $m_t^{\mathrm{c}}(x)$ toward the remote $m_t^{\lambda}(x)$, even where condition confidence remains high outside the data region. PMC-CFG instead selects the intersection of this ray with $\|m\| = \Gamma\|m_t^{\mathrm{c}}(x)\|$, limiting excessive extrapolation while preserving the guidance direction.

Our contributions are:

- We characterize linear Flow Matching and CFG as non-autonomous gradient flows. We show that the same posterior force that provably increases the local conditional-alignment rate can also accumulate mean displacement and contract trajectory volume, thereby separating why CFG works from why fixed strong CFG fails.
- We introduce PMC-CFG, a training-free controller that caps the model-implied terminal posterior mean while preserving the CFG direction. Its closed-form, state-dependent scale provides pointwise feasibility, maximal retention of nominal guidance, and non-interference whenever nominal CFG is feasible—properties not supplied by a time-only guidance schedule.
- We evaluate the predicted mechanisms and the resulting alignment–diversity trade-off on an analytic eight-point GMM distribution, ImageNet-256 class-conditional generation, and SD3.5 text-to-image generation.

2 Preliminaries

2.1 Linear Flow Matching

Let the source sample be $X_0 \sim p_0 = \mathcal{N}(0, I_d)$ and the target sample be $X_1 \sim p_1$, independently. We consider the linear probability path

$$X_t = (1-t)X_0 + tX_1, \qquad t \in [0,1], \tag{1}$$

which interpolates from noise at $t = 0$ to data at $t = 1$. Flow Matching learns the marginal velocity $v_t(x) = \mathbb{E}[\dot{X}_t \mid X_t = x]$ and generates samples by solving $\dot{x}_t = v_t(x_t)$. For the linear path, define the model-implied terminal posterior mean $m_t(x) = \mathbb{E}[X_1 \mid X_t = x]$. Then, for $t \in (0,1)$,

$$v_t(x) = \frac{m_t(x) - x}{1-t} = \frac{x}{t} + \frac{1-t}{t}\nabla_x \log p_t(x). \tag{2}$$

The second equality is the linear-path Tweedie identity (Efron, 2011); Appendix A gives a self-contained derivation.

2.2 Classifier-Free Guidance

For a condition c, superscripts u and c denote unconditional and conditional quantities, respectively. Classifier-free guidance (CFG) combines the two learned velocities as

$$v_t^{\gamma} = v_t^{\mathrm{u}} + \gamma(v_t^{\mathrm{c}} - v_t^{\mathrm{u}}), \tag{3}$$

where $\gamma = 0$ is unconditional, $\gamma = 1$ is purely conditional, and $\gamma > 1$ extrapolates beyond the conditional prediction. Equivalently, the guided velocity implies the terminal estimate

$$m_t^\gamma(x) = x + (1-t)v_t^\gamma(x) = m_t^{\mathrm{c}}(x) + (\gamma - 1)\big(m_t^{\mathrm{c}}(x) - m_t^{\mathrm{u}}(x)\big). \tag{4}$$

Thus standard CFG is linear extrapolation in posterior-mean space. This identity supplies the geometric quantity that PMC-CFG will constrain in Section 4.

3 Flow Matching as a Non-Autonomous Gradient Flow

The preliminaries specify the standard Flow-Matching and CFG construction. We now analyze the resulting dynamics: first as a time-varying potential system, then through its displacement and local volume change.

3.1 A time-varying potential

Proposition 3.1 (Non-autonomous gradient representation)**.** *For the independent linear Gaussian path* (1)*, let $t \in (0,1)$ and let the intermediate density p_t be positive and differentiable. Then*

$$v_t(x) = \nabla_x \Phi_t(x), \qquad \Phi_t(x) = \frac{\|x\|^2}{2t} + \frac{1-t}{t}\log p_t(x). \tag{5}$$

Equivalently, with $U_t = -\Phi_t$, the sampler obeys $\dot{x}_t = -\nabla_x U_t(x_t)$, but its energy evolves as

$$\frac{\mathrm{d}}{\mathrm{d}t}U_t(x_t) = \partial_t U_t(x_t) - \|\nabla_x U_t(x_t)\|^2. \tag{6}$$

Because the landscape changes with time, energy need not decrease, and the endpoint transport is determined by the accumulated dynamics rather than by any single instantaneous potential. Appendix A proves Proposition 3.1.

3.2 CFG constructs a new non-autonomous potential

Let Φ_t^{u} and Φ_t^{c} be the potentials obtained from (5), and write $h_t(x) = \log p(c \mid X_t = x)$. Their difference is determined entirely by the conditional-to-unconditional density ratio. Indeed, Bayes' rule gives

$$\begin{aligned}\Phi_t^{\mathrm{c}}(x) - \Phi_t^{\mathrm{u}}(x) &= \frac{1-t}{t}\left[\log p_t^{\mathrm{c}}(x) - \log p_t^{\mathrm{u}}(x)\right] \\ &= \frac{1-t}{t}h_t(x) - \frac{1-t}{t}\log p(c).\end{aligned} \tag{7}$$

The common quadratic term in (5) cancels in the first line, and the last term in the second line is independent of x. When γ is constant or depends only on time, substituting this potential difference into the CFG combination (3) yields

$$v_t^\gamma = (1-\gamma)\nabla_x \Phi_t^{\mathrm{u}} + \gamma \nabla_x \Phi_t^{\mathrm{c}} = \nabla_x \Phi_t^\gamma, \tag{8}$$

$$\Phi_t^\gamma = (1-\gamma)\Phi_t^{\mathrm{u}} + \gamma \Phi_t^{\mathrm{c}} = \Phi_t^{\mathrm{c}} + (\gamma - 1)\frac{1-t}{t}h_t(x) + K_t, \tag{9}$$

where $K_t = -(\gamma-1)(1-t)\log p(c)/t$ is independent of x. Thus CFG superposes a time-weighted condition-posterior potential on the pure-conditional landscape.

The corresponding force identity is the spatial derivative of the same potential difference in (7).

Proposition 3.2 (Posterior force)**.** *For the independent linear Gaussian path* (1)*, let $t \in (0,1)$ and let the conditional and unconditional intermediate densities be positive and differentiable. Then*

$$v_t^{\mathrm{c}}(x) - v_t^{\mathrm{u}}(x) = \nabla_x(\Phi_t^{\mathrm{c}} - \Phi_t^{\mathrm{u}}) = \frac{1-t}{t}\nabla_x \log p(c \mid X_t = x). \tag{10}$$

Consequently, $v_t^\gamma = v_t^{\mathrm{c}} + (\gamma-1)(v_t^{\mathrm{c}} - v_t^{\mathrm{u}})$ adds $(\gamma-1)(1-t)\nabla_x h_t/t$ to the pure-conditional velocity. Equations (9) and (10) are therefore the scalar-potential and vector-field forms of the same

posterior-guidance term; its next spatial derivative changes the velocity Jacobian. Appendix A gives the underlying linear-path derivation.

For a time-only schedule $\gamma(t)$, the representation (9) still holds, but $\partial_t \Phi_t^{\gamma(t)}$ acquires an x-dependent term proportional to $\dot{\gamma}(t) h_t$ (in addition to an x-independent gauge term). A schedule therefore changes both the instantaneous force and the energy-injection term. The frequently used density

$$\widetilde{p}_{t,\gamma}(x) \propto p_t^{\mathrm{u}}(x)^{1-\gamma} p_t^{\mathrm{c}}(x)^{\gamma} \propto p_t^{\mathrm{c}}(x) p(c \mid X_t = x)^{\gamma-1}$$

provides an instantaneous score interpretation, whereas the endpoint density is governed by the full non-autonomous transport.

3.3 Why CFG works—and why fixed guidance can fail

Equation (10) explains the effectiveness of CFG: the conditional–unconditional difference is proportional to the score of the implicit classifier $p(c \mid X_t = x)$. With $h_t(x) = \log p(c \mid X_t = x)$, its value along a guided trajectory satisfies

$$\frac{\mathrm{d}}{\mathrm{d}t} h_t(x_t) = \partial_t h_t(x_t) + \nabla h_t(x_t)^\top v_t^{\mathrm{c}}(x_t) + (\gamma - 1) \frac{1-t}{t} \left\| \nabla h_t(x_t) \right\|^2 . \tag{11}$$

For $\gamma \geq 1$, the last term is nonnegative. Thus CFG increases the instantaneous posterior-growth rate relative to the pure-conditional velocity at the same time–state pair by a squared-gradient term. This explains stronger conditional alignment. Diversity, however, depends on how the velocity field deforms neighboring trajectories, which is governed by posterior covariance.

Proposition 3.3 (Posterior geometry and local volume change)**.** *Assuming differentiation under the posterior integral is valid,*

$$\nabla_x m_t(x) = \frac{t}{(1-t)^2} \operatorname{Cov}(X_1 \mid X_t = x), \tag{12}$$

$$H_t^\gamma - H_t^{\mathrm{c}} = (\gamma - 1) \frac{1-t}{t} \nabla_x^2 h_t(x) = (\gamma - 1) \frac{t}{(1-t)^3} \left(\operatorname{Cov}_t^{\mathrm{c}} - \operatorname{Cov}_t^{\mathrm{u}} \right), \tag{13}$$

$$\nabla \cdot (v_t^\gamma - v_t^{\mathrm{c}}) = (\gamma - 1) \frac{t}{(1-t)^3} \left[\operatorname{tr} \operatorname{Cov}_t^{\mathrm{c}} - \operatorname{tr} \operatorname{Cov}_t^{\mathrm{u}} \right]. \tag{14}$$

where $H_t^\gamma = \nabla_x v_t^\gamma = \nabla_x^2 \Phi_t^\gamma$. For a fixed starting time s and state x_s, let $x_t = X^\gamma(t, s, x_s)$ denote the solution of $\dot{x}_t = v_t^\gamma(x_t)$ with $X^\gamma(s, s, x_s) = x_s$. Define the trajectory Jacobian $J_t = D_{x_s} X^\gamma(t, s, x_s)$. Then $J_s = I_d$ and

$$\dot{J}_t = H_t^\gamma(x_t) J_t, \qquad \frac{\mathrm{d}}{\mathrm{d}t} \log |\det J_t| = \nabla \cdot v_t^\gamma(x_t). \tag{15}$$

The factor $(1-t)^{-3}$ in the covariance form does not by itself imply an endpoint singularity: the equivalent Hessian form in (13) is $O(1-t)$ whenever $\nabla_x^2 h_t$ remains bounded near $t = 1$. We do not assume this endpoint regularity in general; for the circular GMM below, the cancellation is explicit in (21).

When the conditional posterior covariance has smaller trace than the unconditional one, fixed CFG adds negative divergence and locally contracts trajectory volume. Both this contraction and the guided displacement accumulate along the trajectory:

$$x_t - x_s = \int_s^t v_\tau^\gamma(x_\tau) \, \mathrm{d}\tau, \qquad \log |\det J_t| = \int_s^t \nabla \cdot v_\tau^\gamma(x_\tau) \, \mathrm{d}\tau. \tag{16}$$

This trace comparison is condition- and state-dependent. The law of total covariance gives

$$\operatorname{Cov}_t^{\mathrm{u}}(x) = \mathbb{E}_{C \mid X_t = x}[\operatorname{Cov}_t^C(x)] + \operatorname{Cov}_{C \mid X_t = x}(m_t^C(x)),$$

so averaging over the posterior distribution of conditions gives the trace inequality. For a fixed c, the trace comparison above provides a pointwise sufficient condition for contraction. The same posterior signal improves alignment through its gradient while its curvature can amplify displacement and concentration. This motivates retaining the guidance direction but adapting its magnitude to the realized time–state pair. Appendix B gives the derivations.

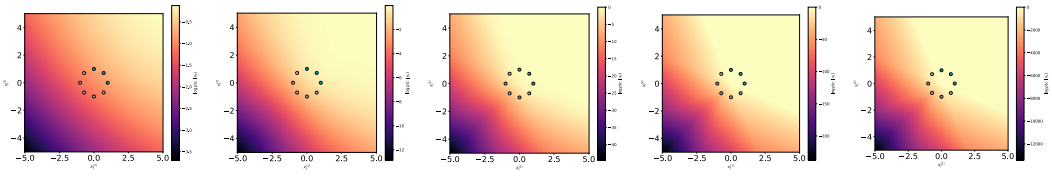


Figure 2: Condition-posterior landscape $\log p(A \mid x_t)$ for the eight-point GMM with $A = \{0, 1, 2\}$ at $t \in \{0.2, 0.4, 0.6, 0.8, 1.0\}$ (left to right). On each nonzero radial shell within the conditioned sector, the posterior peaks on the middle ray and decreases angularly toward the two side branches.

3.4 Guidance effects in a circular GMM

We instantiate posterior ascent, mean extrapolation, local contraction, and branch reweighting in a K-component isotropic GMM in $\mathbb{R}^2$,

$$p_1(x) = \sum_{k=0}^{K-1} w_k \mathcal{N}(x; \mu_k, \sigma^2 I_2), \qquad \mu_k = r u_k, \tag{17}$$

where $u_k = (\cos\theta_k, \sin\theta_k)^\top$. At time t, $X_t \mid k \sim \mathcal{N}(t\mu_k, C_t I_2)$ with $C_t = (1-t)^2 + t^2\sigma^2$.

Theorem 3.4 (Confidence–density decoupling). *Fix component* 0 *and the ray* $x = su_0$, $s \geq 0$. *Let* $P_0(s,t) = p(c = 0 \mid X_t = su_0)$ *and assume that* $u_0^\top u_k$ *is uniquely maximized at* $k = 0$ *over* $k \in \{0, \ldots, K-1\}$. *For every* $t > 0$,

$$\begin{aligned} \partial_s \log P_0(s,t) \geq 0, \qquad & \partial_s \log p_t(su_0) < 0 \quad (s > tr), \\ P_0(s,t) \to 1, \qquad & p_t(su_0) \to 0 \quad (s \to \infty). \end{aligned} \tag{18}$$

Thus condition confidence can increase beyond the high-density data region, and its earlier outward force can leave a persistent displacement. Appendix C provides the proof.

For a centered mixture conditioned on component 0, the observation at $t = 0$ is independent of the clean sample, giving

$$m_0^{\mathrm{u}} = 0, \qquad m_0^{\mathrm{c}} = \mu_0, \qquad \Delta_0 = \mu_0, \qquad m_0^{\gamma} = \gamma\mu_0. \tag{19}$$

Thus fixed CFG amplifies the full initial gap; for $r = 1$ and $\gamma = 3$, its implied endpoint has radius 3. Appendix C.1 provides a further time-dependent analysis: it derives the exact conditional and unconditional posterior means, proves that their gap vanishes under a decreasing envelope, and shows when PMC-CFG recovers its nominal scale.

For an equal-weight rotationally symmetric mixture conditioned on component 0, the center field exhibits displacement and contraction simultaneously:

$$(v_t^\gamma - v_t^{\mathrm{c}})(0) = (\gamma - 1)\frac{1-t}{C_t}\mu_0, \tag{20}$$

$$\nabla\cdot(v_t^\gamma - v_t^{\mathrm{c}})(0) = -(\gamma - 1)\frac{t(1-t)r^2}{C_t^2} < 0, \qquad \gamma > 1. \tag{21}$$

As $t \downarrow 0$, $v_t^\gamma(x) \approx \gamma\mu_0 - x$: early CFG targets the extrapolated location $\gamma\mu_0$ while contracting nearby trajectories.

For a concrete multimodal example, take eight equally spaced, equally weighted components and condition on the adjacent triple $A = \{0, 1, 2\}$. Appendix C.3 proves that, on every nonzero radial shell, the condition posterior is reflection-symmetric about the middle ray and strictly decreases toward either side branch. By the posterior-force identity (10), the CFG correction therefore has an angular component pointing toward the middle ray for every $0 < t < 1$ and $\gamma > 1$. Thus CFG can increase the middle branch's weight despite equal component weights.

Figure 2 visualizes this equal-weight, middle-branch reweighting mechanism; Table 1 confirms the resulting endpoint concentration empirically.

Having shown that CFG can create an imbalance among equally weighted branches, we next consider how it amplifies a pre-existing weight imbalance.

Proposition 3.5 (Posterior-potential bias under branch imbalance). *In the eight-component GMM, let $A = \{0, 1\}$ consist of adjacent components with $w_0 > w_1 > 0$, and assume that the complementary weights are reflection symmetric about the angular bisector of u_0 and u_1: $w_2 = w_7$, $w_3 = w_6$, and $w_4 = w_5$. At the interpolated component centers $x_{k,t} = t\mu_k$, define $P_{A,k}(t) = p(A \mid X_t = x_{k,t})$. Then, for $t \in (0, 1]$, $P_{A,0}(t) > P_{A,1}(t)$. Consequently, for $0 < t < 1$ and $\gamma > 1$, the CFG correction to the non-autonomous potential, $(\gamma - 1)(1 - t) \log p(A \mid X_t = x)/t$, is larger at $x_{0,t}$ than at $x_{1,t}$.*

Thus Proposition 3.5 identifies a local potential bias toward the already heavier branch at the two interpolated centers. Table 1 reports the corresponding endpoint behavior empirically. Appendix C contains the derivations and the nonadjacent configuration.

4 Posterior-Mean-Capped Guidance

4.1 A per-sample geometric constraint

Equation (10) shows that CFG moves samples along $\nabla_x \log p(c \mid X_t = x)$. In the circular GMM, Theorem 3.4 shows that this posterior can increase beyond the high-density data region, so its gradient can drive low-density extrapolation. Because the posterior-mean gap is proportional to the same gradient, the limitation of fixed CFG is not its direction but its open-loop magnitude: the same scale is applied without checking the displacement implied at the current sample. PMC-CFG closes this structural gap by turning the fixed gain into pointwise feedback. It keeps the original posterior direction and limits only how far the implied terminal mean may extrapolate, while retaining as much nominal guidance as the constraint permits.

Let $\lambda \geq 1$ denote the nominal CFG scale and $\Gamma \geq 1$ the permitted relative mean amplification. The two usual model evaluations give

$$m_t^{\mathrm{c}}(x) = x + (1 - t)v_t^{\mathrm{c}}(x), m_t^{\mathrm{u}}(x) = x + (1 - t)v_t^{\mathrm{u}}(x), \tag{22}$$

$$\Delta_t(x) = m_t^{\mathrm{c}}(x) - m_t^{\mathrm{u}}(x) = (1 - t)(v_t^{\mathrm{c}}(x) - v_t^{\mathrm{u}}(x)) = \frac{(1 - t)^2}{t} \nabla_x \log p(c \mid X_t = x).$$

Unconstrained nominal CFG implies $m_t^{\lambda} = m_t^{\mathrm{c}} + (\lambda - 1)\Delta_t$. PMC-CFG instead solves

$$\beta_t^*(x) = \max \left\{ \beta : \begin{array}{l} 0 \leq \beta \leq \lambda - 1, \\ \|m_t^{\mathrm{c}} + \beta\Delta_t\| \leq \Gamma \|m_t^{\mathrm{c}}\| \end{array} \right\} \tag{23}$$

and uses

$$v_t^{\mathrm{cap}}(x) = v_t^{\mathrm{c}}(x) + \beta_t^*(x)(v_t^{\mathrm{c}}(x) - v_t^{\mathrm{u}}(x)). \tag{24}$$

Because β_t^* depends on the state, this controlled field is not generally a gradient field; Appendix D characterizes its feedback Jacobian and local integrability condition. The constraint is evaluated independently per sample. Its Euclidean norm is evaluated in the model's native latent coordinates and is therefore coordinate-dependent; Proposition D.2 characterizes its invariances and gives a translation-covariant anchored form.

4.2 Closed-form scale and guarantees

Define

$$a = \|\Delta_t\|^2, \qquad b = \langle m_t^{\mathrm{c}}, \Delta_t \rangle, \qquad q = \|m_t^{\mathrm{c}}\|^2. \tag{25}$$

For $a > 0$, let

$$\beta^{\mathrm{cap}} = \frac{-b + \sqrt{b^2 + (\Gamma^2 - 1)aq}}{a}. \tag{26}$$

Theorem 4.1 (Maximal feasible guidance). *For every $\Gamma, \lambda \geq 1$, problem (23) is feasible. Its largest solution is*

$$\beta_t^* = \min\{\lambda - 1, \beta^{\mathrm{cap}}\} \quad (a > 0). \tag{27}$$

If $a = 0$, all scales produce the same velocity and we set $\beta_t^* = \lambda - 1$. In exact arithmetic,

$$\|x + (1 - t)v_t^{\mathrm{cap}}(x)\| \leq \Gamma \|x + (1 - t)v_t^{\mathrm{c}}(x)\|. \tag{28}$$

Nominal CFG is unchanged whenever it is feasible.

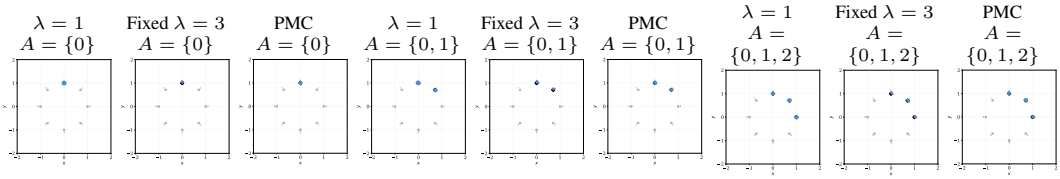


Figure 3: Endpoint samples for the analytic mixture. Gray markers are component means, red diamonds are conditioned means, and blue points are samples. PMC-CFG reduces the overshoot and mode concentration caused by strong fixed guidance.

Table 1: Single-run GMM endpoint statistics ($n = 2 \times 10^5$ per row).

A	Method	λ	Theory occ. (%)	Emp. occ. (%)	Per-label var. ratio	TV↓	Mean error↓
$\{0\}$	Conditional	1	100.00	100.00	0.571	0.00000	0.00005
$\{0\}$	Fixed CFG	3	100.00	100.00	0.135	0.00000	0.04091
$\{0\}$	PMC-CFG	3	100.00	100.00	0.348	0.00000	0.00831
$\{0, 1\}$	Conditional	1	68.00 / 32.00	68.34 / 31.66	0.603 / 0.630	0.00339	0.00247
$\{0, 1\}$	Fixed CFG	3	68.00 / 32.00	83.27 / 16.73	0.315 / 0.358	0.15272	0.11781
$\{0, 1\}$	PMC-CFG	3	68.00 / 32.00	71.85 / 28.15	0.498 / 0.536	0.03851	0.02800
$\{0, 1, 2\}$	Conditional	1	33.33 / 33.33 / 33.33	32.82 / 34.30 / 32.88	0.622 / 0.669 / 0.623	0.00967	0.00270
$\{0, 1, 2\}$	Fixed CFG	3	33.33 / 33.33 / 33.33	15.74 / 68.63 / 15.64	0.368 / 0.526 / 0.368	0.35294	0.12883
$\{0, 1, 2\}$	PMC-CFG	3	33.33 / 33.33 / 33.33	29.26 / 41.38 / 29.37	0.522 / 0.647 / 0.522	0.08046	0.02980

Theorem 4.1 follows from the feasible interval of $a\beta^2 + 2b\beta + (1 - \Gamma^2)q \leq 0$. Its complete proof, including degenerate cases and a cancellation-resistant floating-point evaluation, is in Appendix D.

The construction preserves direction, is always feasible, intervenes minimally, and adapts to each state. Specifically, $\beta = 0$ guarantees feasibility, while maximizing β retains nominal CFG whenever it satisfies the bound; otherwise the scale adapts to the current mean–gap geometry rather than to time alone. The cap is most restrictive for outward-pointing gaps and self-releases as the gap vanishes: in the centered GMM, it limits early extrapolation and recovers the nominal scale near every nonzero endpoint. Appendices D and C.1 give the directional characterization and formal release result.

5 Experiments

5.1 Experimental setup

Within each benchmark, all methods use matched model weights, conditions, initial noise, solver, step count, and decoder. Full implementation and metric details are given in Appendix E.

GMM. We draw 2×10^5 samples from an eight-component circular mixture and compare conditional sampling, fixed CFG, and PMC-CFG ($\lambda = 3$, $\Gamma = 1.1$) across singleton, adjacent-pair, and adjacent-triple conditions.

ImageNet-256. We evaluate class-conditional SiT-XL/2 (Deng et al., 2009; Ma et al., 2024) on 50K samples at 20/50 NFE and $\lambda \in \{1.5, 2.0, 2.5\}$, comparing fixed CFG, C2FG (Gao et al., 2026), APG (Sadat et al., 2025), and PMC-CFG. PMC-CFG uses $\Gamma = 1.10$ at $\lambda = 1.5$ and 1.05 otherwise; metrics are FID, sFID, IS, precision, and recall (Heusel et al., 2017; Salimans et al., 2016; Dhariwal & Nichol, 2021; Kynkäänniemi et al., 2019).

Stable Diffusion 3.5. We evaluate SD3.5 Medium (Esser et al., 2024) at 10 NFE and $\lambda \in \{3, 5, 7, 9\}$ on 5,000 COCO prompts (Lin et al., 2014), with shared prompts and initial latents. We compare fixed CFG, C2FG (Gao et al., 2026), APG (Sadat et al., 2025), and PMC-CFG using CLIP, ImageReward, FID, Recall@3, and mean HSV saturation, the last as a nondirectional diagnostic.

Table 2: ImageNet-256 results on 50K samples. The best FID at each NFE is bold.

		20 NFE					50 NFE				
Method	λ	FID↓	sFID↓	IS↑	Prec.↑	Rec.↑	FID↓	sFID↓	IS↑	Prec.↑	Rec.↑
Fixed CFG	1.50	3.58	5.71	239.49	0.79	0.57	2.39	4.80	258.11	0.80	0.59
APG	1.50	3.61	6.34	245.37	0.78	0.58	2.29	5.20	263.51	0.80	0.60
C2FG	1.50	3.09	5.45	260.62	0.80	0.55	2.21	4.64	278.73	0.81	0.58
PMC-CFG	1.50	3.03	5.51	261.95	0.80	0.57	**2.12**	4.60	283.15	0.82	0.58
Fixed CFG	2.00	3.58	4.74	340.73	0.86	0.50	3.62	4.55	357.64	0.87	0.51
APG	2.00	3.49	4.80	348.26	0.85	0.51	3.31	4.18	365.52	0.86	0.53
C2FG	2.00	3.98	4.67	358.61	0.87	0.49	4.12	4.58	375.54	0.87	0.50
PMC-CFG	2.00	**2.88**	4.79	337.62	0.84	0.53	2.85	4.15	359.72	0.85	0.55
Fixed CFG	2.50	6.22	5.42	402.62	0.89	0.43	6.52	5.63	414.58	0.89	0.45
APG	2.50	6.02	4.33	412.72	0.89	0.44	6.17	4.06	428.46	0.89	0.46
C2FG	2.50	6.84	5.48	417.47	0.90	0.42	7.17	5.76	427.26	0.90	0.44
PMC-CFG	2.50	3.82	4.46	382.53	0.86	0.51	4.11	4.07	403.15	0.87	0.52

Table 3: SD3.5 Medium results on 5,000 COCO 2017 validation prompts at 10 NFE. Bold indicates the best directional metric within each guidance scale.

Scale λ	Method	CLIP↑	ImageReward↑	FID↓	Sat.	Recall@3 ↑
3	Fixed CFG	0.3171	0.6476	26.94	0.301	0.4456
	APG	0.3088	0.2580	25.72	0.197	**0.5020**
	C2FG	**0.3174**	**0.6715**	26.97	0.302	0.4394
	PMC-CFG	0.3148	0.5918	**25.35**	0.223	0.4722
5	Fixed CFG	0.3192	0.7101	29.20	0.407	0.4234
	APG	0.3138	0.5728	**24.68**	0.208	**0.4852**
	C2FG	**0.3193**	0.7271	29.06	0.408	0.4290
	PMC-CFG	0.3168	**0.7538**	25.65	0.227	0.4620
7	Fixed CFG	0.3145	0.4328	31.40	0.492	0.4208
	APG	0.3160	0.7194	**25.21**	0.220	**0.4784**
	C2FG	0.3148	0.4516	31.20	0.493	0.4202
	PMC-CFG	**0.3174**	**0.8033**	25.67	0.229	0.4540
9	Fixed CFG	0.3049	−0.0295	40.83	0.555	**0.4512**
	APG	0.3172	0.8097	25.88	0.234	0.4472
	C2FG	0.3056	−0.0108	40.54	0.556	0.4470
	PMC-CFG	**0.3178**	**0.8262**	**25.61**	0.230	**0.4512**

5.2 RESULTS

GMM. Figure 3 and Table 1 confirm the predicted failure modes: fixed CFG displaces and contracts a single mode and concentrates multimodal mass on heavier or central branches. PMC-CFG reduces these errors and restores within-branch spread. Appendix Figure 7 and Appendix E provide densities and run-level protocols.

ImageNet-256. Table 2 shows that PMC-CFG is most effective under strong guidance: at $\lambda \in \{2.0, 2.5\}$ it improves FID and recall over fixed CFG and APG at both solver budgets, with a small precision trade-off; at $\lambda = 1.5$, it gives the lowest FID while APG gives the highest recall. Appendices F.5, F.1, and F.2 analyze the optimum, per-sample feedback, and the quality–coverage trade-off. Figure 4(a) shows that PMC-CFG better preserves unguided layouts while recovering sharp detail.

Stable Diffusion 3.5. Table 3 shows that APG and PMC-CFG avoid the high-guidance FID and saturation degradation of fixed CFG and C2FG. APG gives the stronger distributional trade-off at intermediate scales, whereas PMC-CFG attains the best FID at $\lambda \in \{3, 9\}$ and the strongest CLIP and ImageReward at $\lambda \in \{7, 9\}$. Appendix F.3 extends the sweep and shows that a tighter cap favors

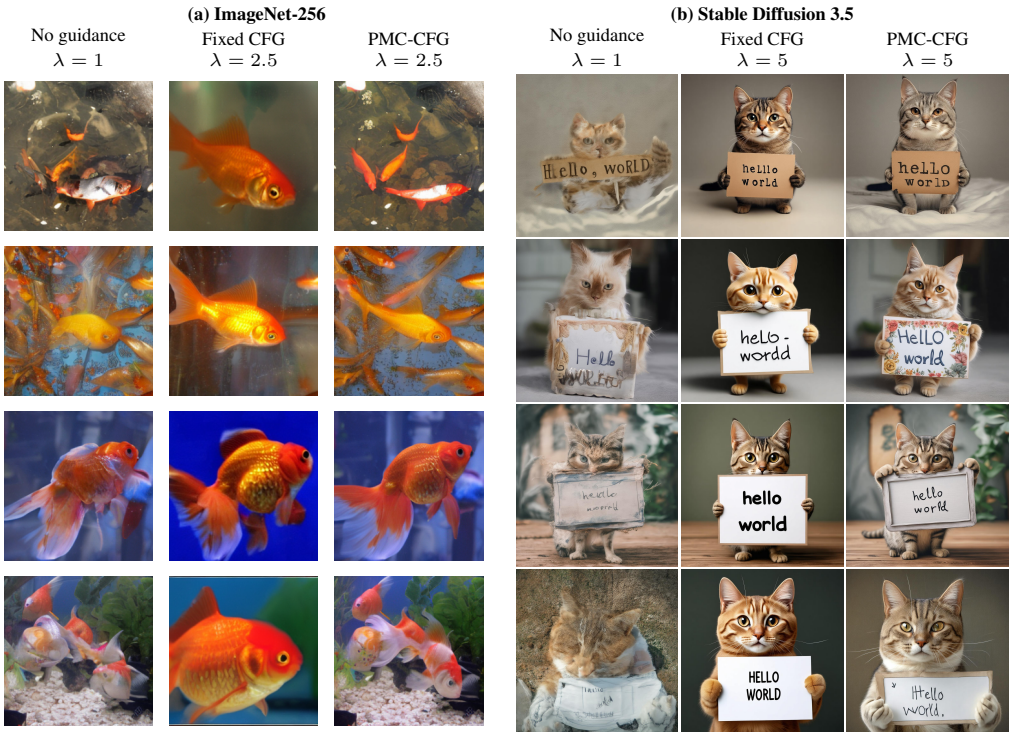


Figure 4: Matched qualitative comparisons. Rows use the same class or prompt and seed across methods. PMC-CFG better retains the unguided composition than fixed CFG on both ImageNet-256 and SD3.5.

FID and coverage, whereas a looser cap favors CLIP and ImageReward. Figure 4(b) shows that PMC-CFG preserves legible text and greater compositional variation.

6 Discussion and Conclusion

We characterized linear Flow Matching as a non-autonomous gradient flow and showed that CFG tilts its time-varying potential by the condition posterior. The resulting force improves instantaneous conditional alignment but can accumulate mean displacement and local volume contraction, explaining overshoot and diversity loss under strong fixed guidance. PMC-CFG caps this mechanism along the CFG direction. Its distinguishing guarantee is pointwise feasibility under the stated mean-space constraint, while feasible nominal guidance is preserved without training or extra model calls.

Across an analytic GMM, ImageNet-256, and SD3.5, PMC-CFG reduces overshoot and concentration and improves the quality–diversity trade-off most consistently under strong guidance. Its pointwise guarantee is defined in the chosen latent coordinates, motivating extensions that control the full endpoint distribution.

Reproducibility Statement

The theoretical assumptions and complete proofs are provided in Appendices A–D, including stable per-sample implementation details. Appendix E specifies the experimental setups, sampling protocols, baseline configurations, and metric implementations, while Appendix F reports additional ablations, sensitivity analyses, and qualitative results.

AI Use Statement

Generative AI tools were used to assist with drafting and editing portions of the manuscript and with implementing and editing research software. They were not used for any other applicable tasks

requiring disclosure under the ICLR 2027 policy. The authors reviewed and revised all AI-assisted text, tested all AI-assisted code, and independently verified the scientific claims, mathematical derivations, experimental results, and citations. The authors take responsibility for the final content of this work.

REFERENCES

Michael S Albergo and Eric Vanden-Eijnden. Building normalizing flows with stochastic interpolants. *arXiv preprint arXiv:2209.15571*, 2022.

Iskander Azangulov, Peter Potaptchik, Qinyu Li, Eddie Aamari, George Deligiannidis, and Judith Rousseau. Adaptive diffusion guidance via stochastic optimal control. In *Proceedings of the 29th International Conference on Artificial Intelligence and Statistics*, volume 300 of *Proceedings of Machine Learning Research*, pp. 4087–4095, 2026. URL https://proceedings.mlr.press/v300/azangulov26a.html.

Arwen Bradley and Preetum Nakkiran. Classifier-free guidance is a predictor-corrector. *arXiv preprint arXiv:2408.09000*, 2024.

Jian-Feng Cai, Haixia Liu, Zhengyi Su, and Chao Wang. Improving classifier-free guidance of flow matching via manifold projection. *arXiv preprint arXiv:2601.21892*, 2026. doi: 10.48550/arXiv.2601.21892.

Hyungjin Chung, Jeongsol Kim, Geon Yeong Park, Hyelin Nam, and Jong Chul Ye. CFG++: Manifold-constrained classifier free guidance for diffusion models. In *International Conference on Learning Representations*, 2025.

Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pp. 248–255. Ieee, 2009.

Nicki Skafte Detlefsen, Jiri Borovec, Justus Schock, Ananya Harsh Jha, Teddy Koker, Luca Di Liello, Daniel Stancl, Changsheng Quan, Maxim Grechkin, and William Falcon. TorchMetrics: Measuring reproducibility in PyTorch. *Journal of Open Source Software*, 7(70):4101, 2022. doi: 10.21105/joss.04101.

Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34:8780–8794, 2021.

Bradley Efron. Tweedie's formula and selection bias. *Journal of the American Statistical Association*, 106(496):1602–1614, 2011. doi: 10.1198/jasa.2011.tm11181.

Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, Dustin Podell, Tim Dockhorn, Zion English, and Robin Rombach. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first international conference on machine learning*, 2024.

Jiayang Gao, Tianyi Zheng, Jiayang Zou, Fengxiang Yang, Shice Liu, Luyao Fan, Zheyu Zhang, Hao Zhang, Jinwei Chen, Peng-Tao Jiang, Bo Li, and Jia Wang. C^2fg: Control classifier-free guidance via score discrepancy analysis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 34398–34407, June 2026.

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, pp. 6626–6637, 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.

Enze Jiang and Zheng Ma. Analytic distribution of classifier-free guidance for schedule design. *arXiv preprint arXiv:2607.19725*, 2026.

Tero Karras, Miika Aittala, Tero Kynkäänniemi, Jaakko Lehtinen, Timo Aila, and Samuli Laine. Guiding a diffusion model with a bad version of itself. In *Advances in Neural Information Processing Systems*, 2024.

Tuomas Kynkäänniemi, Tero Karras, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Improved precision and recall metric for assessing generative models. In *Advances in Neural Information Processing Systems*, volume 32, 2019.

Tuomas Kynkäänniemi, Miika Aittala, Tero Karras, Samuli Laine, Timo Aila, and Jaakko Lehtinen. Applying guidance in a limited interval improves sample and distribution quality in diffusion models. In *Advances in Neural Information Processing Systems*, volume 37, pp. 122458–122483, 2024. doi: 10.52202/079017-3892.

Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO: Common objects in context. In *European Conference on Computer Vision*, pp. 740–755, 2014.

Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. In *International Conference on Learning Representations*, 2023.

Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. In *International Conference on Learning Representations*, 2023.

Nanye Ma, Mark Goldstein, Michael S. Albergo, Nicholas M. Boffi, Eric Vanden-Eijnden, and Saining Xie. SiT: Exploring flow and diffusion-based generative models with scalable interpolant transformers. In *European Conference on Computer Vision*, pp. 23–40, 2024. doi: 10.1007/978-3-031-72980-5_2.

Gaurav Parmar, Richard Zhang, and Jun-Yan Zhu. On aliased resizing and surprising subtleties in GAN evaluation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11400–11410, 2022. doi: 10.1109/CVPR52688.2022.01112.

Krunoslav Lehman Pavasovic, Jakob Verbeek, Giulio Biroli, and Marc Mézard. Overshoot and shrinkage in classifier-free guidance: From theory to practice. In *International Conference on Learning Representations*, 2026.

Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pp. 8748–8763. PmLR, 2021.

Seyedmorteza Sadat, Otmar Hilliges, and Romann M. Weber. Eliminating oversaturation and artifacts of high guidance scales in diffusion models. In *International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=e2ONKX6qzJ.

Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training GANs. In *Advances in Neural Information Processing Systems*, volume 29, 2016.

Stability AI. Stable diffusion 3.5 medium model card. https://huggingface.co/stabilityai/stable-diffusion-3.5-medium, 2024.

Hanyang Wang, Yiyang Liu, Jiawei Chi, Fangfu Liu, Ran Xue, and Yueqi Duan. Cfg-ctrl: Control-based classifier-free diffusion guidance. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 11437–11447, June 2026.

Xi Wang, Nicolas Dufour, Nefeli Andreou, Marie-Paule Cani, Victoria Fernandez Abrevaya, David Picard, and Vicky Kalogeiton. Analysis of classifier-free guidance weight schedulers. *Transactions on Machine Learning Research*, 2024. ISSN 2835-8856. URL https://openreview.net/forum?id=SUMtDJqicd.

Mengfei Xia, Nan Xue, Yujun Shen, Ran Yi, Tieliang Gong, and Yong-Jin Liu. Rectified diffusion guidance for conditional generation. In *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 13371–13380. IEEE, 2025.

Jiazheng Xu, Xiao Liu, Yuchen Wu, Yuxuan Tong, Qinkai Li, Ming Ding, Jie Tang, and Yuxiao Dong. ImageReward: Learning and evaluating human preferences for text-to-image generation. In *Advances in Neural Information Processing Systems*, volume 36, 2023. doi: 10.52202/075280-0700.

Qinqing Zheng, Matt Le, Neta Shaul, Yaron Lipman, Aditya Grover, and Ricky T. Q. Chen. Guided flows for generative modeling and decision making. *arXiv preprint arXiv:2311.13443*, 2023. doi: 10.48550/arXiv.2311.13443.

A LINEAR FLOW-MATCHING IDENTITIES

This appendix derives the linear-path identities used in Section 2, Proposition 3.2, and the potential representations in Section 3 without invoking a denoising approximation.

Lemma A.1 (Velocity and Tweedie identities)**.** *For the path* (1)*, assume p_t is positive and differentiation under the integral is valid. Then, for $t \in (0, 1)$,*

$$v_t(x) = \frac{\mathbb{E}[X_1 \mid X_t = x] - x}{1 - t}, \tag{29}$$

$$\nabla_x \log p_t(x) = -\frac{x - t\mathbb{E}[X_1 \mid X_t = x]}{(1 - t)^2}, \tag{30}$$

$$v_t(x) = \frac{x}{t} + \frac{1 - t}{t} \nabla_x \log p_t(x). \tag{31}$$

Proof. Fix a pair (X_0, X_1) and differentiate its interpolation:

$$\dot{X}_t = \frac{\mathrm{d}}{\mathrm{d}t}\{(1 - t)X_0 + tX_1\} = X_1 - X_0.$$

For $t < 1$, the path equation gives

$$X_0 = \frac{X_t - tX_1}{1 - t}.$$

Substituting this expression into the paired velocity and collecting the X_1 terms yields

$$\begin{aligned} X_1 - X_0 &= X_1 - \frac{X_t - tX_1}{1 - t} \\ &= \frac{(1 - t)X_1 - X_t + tX_1}{1 - t} \\ &= \frac{X_1 - X_t}{1 - t}. \end{aligned}$$

The exact marginal velocity is the conditional mean of the paired velocity. Therefore, at $X_t = x$,

$$\begin{aligned} v_t(x) &= \mathbb{E}[\dot{X}_t \mid X_t = x] \\ &= \mathbb{E}\left[\frac{X_1 - X_t}{1 - t} \,\middle|\, X_t = x\right] \\ &= \frac{\mathbb{E}[X_1 \mid X_t = x] - x}{1 - t}, \end{aligned}$$

which proves (29).

Conditioning instead on $X_1 = y$ gives the Gaussian kernel

$$p_t(x \mid y) = \frac{1}{[2\pi(1 - t)^2]^{d/2}} \exp\left[-\frac{\|x - ty\|^2}{2(1 - t)^2}\right].$$

The kernel score is

$$\nabla_x \log p_t(x \mid y) = -\frac{x - ty}{(1 - t)^2},$$

and hence

$$\nabla_x p_t(x \mid y) = -p_t(x \mid y)\frac{x - ty}{(1 - t)^2}.$$

Since $p_t(x) = \int p_1(y)p_t(x \mid y)\mathrm{d}y$, differentiation under the integral and Bayes' rule give

$$\begin{aligned} \nabla_x p_t(x) &= \int p_1(y)p_t(x \mid y)\left[-\frac{x - ty}{(1 - t)^2}\right]\mathrm{d}y \\ &= -\frac{1}{(1 - t)^2}\left[xp_t(x) - t\int yp_1(y)p_t(x \mid y)\mathrm{d}y\right] \\ &= -\frac{p_t(x)}{(1 - t)^2}\left[x - t\int y\frac{p_1(y)p_t(x \mid y)}{p_t(x)}\mathrm{d}y\right] \\ &= -\frac{p_t(x)}{(1 - t)^2}\left(x - t\mathbb{E}[X_1 \mid X_t = x]\right). \end{aligned}$$

Dividing by $p_t(x) > 0$ proves (30). Writing $m_t(x) = \mathbb{E}[X_1 \mid X_t = x]$, that identity is equivalent to

$$(1-t)^2 \nabla_x \log p_t(x) = -x + t m_t(x),$$
$$m_t(x) = \frac{x}{t} + \frac{(1-t)^2}{t} \nabla_x \log p_t(x).$$

Substitution into (29) gives

$$\begin{aligned} v_t(x) &= \frac{1}{1-t}\left[\frac{x}{t} + \frac{(1-t)^2}{t}\nabla_x \log p_t(x) - x\right] \\ &= \frac{1}{1-t}\left[\frac{1-t}{t}x + \frac{(1-t)^2}{t}\nabla_x \log p_t(x)\right] \\ &= \frac{x}{t} + \frac{1-t}{t}\nabla_x \log p_t(x), \end{aligned}$$

which is (31). □

Proof of Proposition 3.2. Equation (2) follows from Lemma A.1. Apply (31) separately to $p_t^{\mathrm{c}}(x) = p_t(x \mid c)$ and $p_t^{\mathrm{u}}(x) = p_t(x)$. This gives

$$v_t^{\mathrm{c}}(x) = \frac{x}{t} + \frac{1-t}{t}\nabla_x \log p_t^{\mathrm{c}}(x),$$
$$v_t^{\mathrm{u}}(x) = \frac{x}{t} + \frac{1-t}{t}\nabla_x \log p_t^{\mathrm{u}}(x).$$

Subtracting cancels the common x/t term:

$$\begin{aligned} v_t^{\mathrm{c}}(x) - v_t^{\mathrm{u}}(x) &= \frac{1-t}{t}\left[\nabla_x \log p_t^{\mathrm{c}}(x) - \nabla_x \log p_t^{\mathrm{u}}(x)\right] \\ &= \frac{1-t}{t}\nabla_x \log \frac{p_t^{\mathrm{c}}(x)}{p_t^{\mathrm{u}}(x)}. \end{aligned}$$

Bayes' rule gives

$$p_t(x \mid c) = \frac{p(c \mid X_t = x) p_t(x)}{p(c)}, \qquad \frac{p_t^{\mathrm{c}}(x)}{p_t^{\mathrm{u}}(x)} = \frac{p(c \mid X_t = x)}{p(c)}.$$

Because $p(c)$ is independent of x,

$$\nabla_x \log \frac{p_t^{\mathrm{c}}(x)}{p_t^{\mathrm{u}}(x)} = \nabla_x \log p(c \mid X_t = x),$$

which proves (10).

Finally, the posterior-mean identity in the preliminaries follows directly from the CFG convention $v_t^{\gamma} = v_t^{\mathrm{u}} + \gamma(v_t^{\mathrm{c}} - v_t^{\mathrm{u}})$. Then

$$\begin{aligned} m_t^{\gamma} &= x + (1-t)[v_t^{\mathrm{u}} + \gamma(v_t^{\mathrm{c}} - v_t^{\mathrm{u}})] \\ &= x + (1-t)v_t^{\mathrm{u}} + \gamma(1-t)(v_t^{\mathrm{c}} - v_t^{\mathrm{u}}) \\ &= m_t^{\mathrm{u}} + \gamma(m_t^{\mathrm{c}} - m_t^{\mathrm{u}}) \\ &= (1-\gamma)m_t^{\mathrm{u}} + \gamma m_t^{\mathrm{c}} \\ &= m_t^{\mathrm{c}} + (\gamma - 1)(m_t^{\mathrm{c}} - m_t^{\mathrm{u}}), \end{aligned}$$

proving (4). □

Posterior growth along a guided trajectory. Let x_t solve $\dot{x}_t = v_t^{\gamma}(x_t)$. The chain rule gives

$$\frac{\mathrm{d}}{\mathrm{d}t} h_t(x_t) = \partial_t h_t(x_t) + \nabla_x h_t(x_t)^{\top} v_t^{\gamma}(x_t).$$

Using $v_t^\gamma = v_t^{\mathrm{c}} + (\gamma - 1)\big(v_t^{\mathrm{c}} - v_t^{\mathrm{u}}\big)$ and then (10),

$$\begin{aligned}\frac{\mathrm{d}}{\mathrm{d}t}h_t(x_t) &= \partial_t h_t(x_t) + \nabla_x h_t(x_t)^\top v_t^{\mathrm{c}}(x_t) \\ &\quad + (\gamma-1)\nabla_x h_t(x_t)^\top (v_t^{\mathrm{c}}(x_t) - v_t^{\mathrm{u}}(x_t)) \\ &= \partial_t h_t(x_t) + \nabla_x h_t(x_t)^\top v_t^{\mathrm{c}}(x_t) \\ &\quad + (\gamma-1)\frac{1-t}{t}\nabla_x h_t(x_t)^\top \nabla_x h_t(x_t) \\ &= \partial_t h_t(x_t) + \nabla_x h_t(x_t)^\top v_t^{\mathrm{c}}(x_t) + (\gamma-1)\frac{1-t}{t}\|\nabla_x h_t(x_t)\|^2.\end{aligned}$$

This is (11). Only the last term has a fixed sign for $\gamma \geq 1$; the first two terms remain unrestricted.

Proof of Proposition 3.1. Define

$$\Phi_t(x) = \frac{\|x\|^2}{2t} + \frac{1-t}{t}\log p_t(x).$$

Spatial differentiation gives

$$\begin{aligned}\nabla_x \Phi_t(x) &= \nabla_x \frac{\|x\|^2}{2t} + \frac{1-t}{t}\nabla_x \log p_t(x) \\ &= \frac{x}{t} + \frac{1-t}{t}\nabla_x \log p_t(x) \\ &= v_t(x)\end{aligned}$$

by (31). With $U_t = -\Phi_t$, $\nabla_x U_t = -\nabla_x \Phi_t = -v_t$, and therefore $\dot{x}_t = v_t(x_t) = -\nabla_x U_t(x_t)$. The chain rule along the trajectory yields

$$\begin{aligned}\frac{\mathrm{d}}{\mathrm{d}t}U_t(x_t) &= \partial_t U_t(x_t) + \nabla_x U_t(x_t)^\top \dot{x}_t \\ &= \partial_t U_t(x_t) - \|\nabla_x U_t(x_t)\|^2,\end{aligned}$$

which proves (6). □

CFG potential and instantaneous power tilting. Linearity of the spatial gradient gives

$$v_t^\gamma = (1-\gamma)\nabla\Phi_t^{\mathrm{u}} + \gamma\nabla\Phi_t^{\mathrm{c}} = \nabla[(1-\gamma)\Phi_t^{\mathrm{u}} + \gamma\Phi_t^{\mathrm{c}}].$$

Bayes' rule implies $\log p_t^{\mathrm{c}} - \log p_t^{\mathrm{u}} = h_t - \log p(c)$. To make the substitution explicit,

$$\begin{aligned}\Phi_t^\gamma &= (1-\gamma)\Phi_t^{\mathrm{u}} + \gamma\Phi_t^{\mathrm{c}} \\ &= \Phi_t^{\mathrm{c}} + (\gamma-1)(\Phi_t^{\mathrm{c}} - \Phi_t^{\mathrm{u}}) \\ &= \Phi_t^{\mathrm{c}} + (\gamma-1)\frac{1-t}{t}(\log p_t^{\mathrm{c}} - \log p_t^{\mathrm{u}}) \\ &= \Phi_t^{\mathrm{c}} + (\gamma-1)\frac{1-t}{t}h_t - (\gamma-1)\frac{1-t}{t}\log p(c).\end{aligned}$$

The last term depends on t and c but not on x, so it is absorbed into K_t:

$$K_t = -(\gamma-1)\frac{1-t}{t}\log p(c).$$

Thus

$$\Phi_t^\gamma = \Phi_t^{\mathrm{c}} + (\gamma-1)\frac{1-t}{t}h_t + K_t,$$

proving (9). Equivalently, the spatial score is that of

$$\begin{aligned}\widetilde{p}_{t,\gamma}(x) &\propto \exp\{(1-\gamma)\log p_t^{\mathrm{u}}(x) + \gamma \log p_t^{\mathrm{c}}(x)\} \\ &= p_t^{\mathrm{u}}(x)^{1-\gamma} p_t^{\mathrm{c}}(x)^\gamma \\ &\propto p_t^{\mathrm{c}}(x) p(c \mid X_t = x)^{\gamma-1}.\end{aligned}$$

This equality characterizes the score at each t, while the full non-autonomous ODE determines the transported endpoint density.

B Posterior Curvature and Volume Transport

Proof of Proposition 3.3. Let $q_t(y \mid x)$ denote the posterior density of $X_1 = y$ given $X_t = x$. Bayes' rule writes it as

$$q_t(y \mid x) = \frac{p_1(y)(2\pi(1-t)^2)^{-d/2} \exp[-\|x - ty\|^2/(2(1-t)^2)]}{p_t(x)}.$$

Taking the spatial log-gradient gives

$$\begin{aligned}\nabla_x \log q_t(y \mid x) &= -\frac{x - ty}{(1-t)^2} - \nabla_x \log p_t(x) \\ &= -\frac{x - ty}{(1-t)^2} + \frac{x - tm_t(x)}{(1-t)^2} \\ &= \frac{t}{(1-t)^2}(y - m_t(x)),\end{aligned}$$

where the second line uses (30). Thus

$$\nabla_x \log q_t(y \mid x) = \frac{t}{(1-t)^2}\big(y - m_t(x)\big). \tag{32}$$

For coordinates i, j, differentiate the posterior expectation under the integral sign and use $\partial_{x_j} q_t = q_t \partial_{x_j} \log q_t$:

$$\begin{aligned}\frac{\partial m_{t,i}}{\partial x_j} &= \int y_i \frac{\partial q_t(y \mid x)}{\partial x_j} \mathrm{d}y \\ &= \frac{t}{(1-t)^2}\left(\mathbb{E}[X_{1,i}X_{1,j} \mid X_t = x] - m_{t,j}\mathbb{E}[X_{1,i} \mid X_t = x]\right) \\ &= \frac{t}{(1-t)^2}\left(\mathbb{E}[X_{1,i}X_{1,j} \mid X_t = x] - m_{t,i}m_{t,j}\right) \\ &= \frac{t}{(1-t)^2}\operatorname{Cov}(X_1 \mid X_t = x)_{ij}.\end{aligned}$$

This proves (12). Since t is fixed under the spatial derivative, differentiating $v_t = (m_t - x)/(1-t)$ gives

$$\nabla_x v_t = \frac{t}{(1-t)^3}\operatorname{Cov}(X_1 \mid X_t = x) - \frac{1}{1-t} I_d.$$

Writing the conditional and unconditional posterior covariances as $\Sigma_t^{\mathrm{c}}(x)$ and $\Sigma_t^{\mathrm{u}}(x)$, respectively,

$$\begin{aligned}\nabla_x(v_t^{\mathrm{c}} - v_t^{\mathrm{u}}) &= \left[\frac{t}{(1-t)^3}\Sigma_t^{\mathrm{c}} - \frac{1}{1-t} I_d\right] - \left[\frac{t}{(1-t)^3}\Sigma_t^{\mathrm{u}} - \frac{1}{1-t} I_d\right] \\ &= \frac{t}{(1-t)^3}(\Sigma_t^{\mathrm{c}} - \Sigma_t^{\mathrm{u}}).\end{aligned}$$

Because $v_t^\gamma = v_t^{\mathrm{c}} + (\gamma - 1)(v_t^{\mathrm{c}} - v_t^{\mathrm{u}})$, multiplication by $\gamma - 1$ proves the covariance expression (13). Alternatively, differentiating (10) in space gives

$$\begin{aligned}H_t^\gamma - H_t^{\mathrm{c}} &= (\gamma - 1)\nabla_x(v_t^{\mathrm{c}} - v_t^{\mathrm{u}}) \\ &= (\gamma - 1)\frac{1-t}{t}\nabla_x^2 h_t,\end{aligned}$$

which equals the covariance expression (13) above. Taking traces also gives the corresponding extra divergence explicitly:

$$\nabla \cdot (v_t^\gamma - v_t^{\mathrm{c}}) = (\gamma - 1)\frac{t}{(1-t)^3}\left(\operatorname{tr}\Sigma_t^{\mathrm{c}} - \operatorname{tr}\Sigma_t^{\mathrm{u}}\right).$$

Using the flow-map notation above, differentiate $\partial_t X^\gamma(t, s, x_s) = v_t^\gamma(X^\gamma(t, s, x_s))$ with respect to x_s. The chain rule gives

$$\dot{J}_t = (\nabla_x v_t^\gamma)(x_t) J_t, \qquad x_t = X^\gamma(t, s, x_s).$$

For an invertible J_t, Jacobi's determinant formula and cyclicity of trace yield

$$\frac{\mathrm{d}}{\mathrm{d}t}\log|\det J_t| = \mathrm{tr}(J_t^{-1}\dot{J}_t) = \mathrm{tr}\big(J_t^{-1}(\nabla_x v_t^\gamma)J_t\big) = \mathrm{tr}\big((\nabla_x v_t^\gamma)J_t J_t^{-1}\big) = \mathrm{tr}(\nabla_x v_t^\gamma) = \nabla\cdot v_t^\gamma,$$

where the vector field is evaluated at x_t. Integrating from s to t and using $\log|\det J_s| = 0$ gives

$$\log|\det J_t| = \int_s^t \nabla\cdot v_\tau^\gamma(x_\tau)\,\mathrm{d}\tau,$$

which proves (15). Likewise, integrating $\dot{x}_\tau = v_\tau^\gamma(x_\tau)$ gives the accumulated displacement

$$x_t - x_s = \int_s^t v_\tau^\gamma(x_\tau)\,\mathrm{d}\tau.$$

□

Exact GMM Jacobian. Differentiating the responsibility softmax component by component gives

$$\nabla_x \pi_k = \pi_k\left(\frac{t}{C_t}\mu_k - \sum_j \pi_j \frac{t}{C_t}\mu_j\right) = \frac{t}{C_t}\pi_k(\mu_k - \bar{\mu}_t).$$

Consequently, with $\Sigma_\pi = \sum_k \pi_k(\mu_k - \bar{\mu}_t)(\mu_k - \bar{\mu}_t)^\top$,

$$\begin{aligned}\nabla_x \bar{\mu}_t &= \sum_k \mu_k(\nabla_x \pi_k)^\top \\ &= \frac{t}{C_t}\left(\sum_k \pi_k \mu_k \mu_k^\top - \bar{\mu}_t\bar{\mu}_t^\top\right) = \frac{t}{C_t}\Sigma_\pi.\end{aligned}$$

Now differentiate both terms of (41):

$$\begin{aligned}\nabla_x v_t &= \frac{1-t}{C_t}\nabla_x\bar{\mu}_t + \frac{t\sigma^2 - (1-t)}{C_t} I_2 \\ &= \frac{t(1-t)}{C_t^2}\Sigma_\pi + \frac{t\sigma^2 - (1-t)}{C_t} I_2.\end{aligned}$$

This makes the covariance mechanism in Proposition 3.3 fully explicit for the circular mixture.

C Circular-GMM Derivations

Under component k, write $X_1 = \mu_k + \sigma\varepsilon$ with $\varepsilon \sim \mathcal{N}(0, I_2)$ independent of $X_0 \sim \mathcal{N}(0, I_2)$. The linear path becomes

$$X_t = (1-t)X_0 + t\mu_k + t\sigma\varepsilon.$$

Its conditional mean and covariance are

$$\begin{aligned}\mathbb{E}[X_t \mid k] &= t\mu_k, \\ \mathrm{Cov}(X_t \mid k) &= t^2\sigma^2 I_2 + (1-t)^2 I_2 = C_t I_2,\end{aligned}$$

where $C_t = (1-t)^2 + t^2\sigma^2$. Hence

$$X_t \mid k \sim \mathcal{N}(t\mu_k, C_t I_2).$$

Bayes' rule gives the component responsibility

$$\pi_k(x,t) = p(k \mid X_t = x) = \frac{w_k p_t(x \mid k)}{\sum_j w_j p_t(x \mid j)}.$$

The Gaussian normalizing factors are identical across components and cancel, which yields

$$\pi_k(x,t) = \frac{w_k \exp[-\|x - t\mu_k\|^2/(2C_t)]}{\sum_j w_j \exp[-\|x - t\mu_j\|^2/(2C_t)]}. \tag{33}$$

To obtain the softmax form, expand

$$-\frac{\|x-t\mu_k\|^2}{2C_t} = -\frac{\|x\|^2}{2C_t} + \frac{t}{C_t}x^\top\mu_k - \frac{t^2\|\mu_k\|^2}{2C_t}.$$

All centers satisfy $\|\mu_k\| = r$, so the first and third terms do not depend on k and cancel between numerator and denominator. Therefore

$$\pi_k(x,t) = \mathrm{softmax}_k\left(\log w_k + \frac{t}{C_t}x^\top\mu_k\right).$$

C.1 Posterior means and exact velocity

Define

$$\alpha_t = \frac{(1-t)^2}{C_t}, \qquad \rho_t = \frac{t\sigma^2}{C_t}, \qquad 0 \le t < 1.$$

Within component k, the jointly Gaussian pair (X_1, X_t) has moments

$$\begin{aligned} \mathbb{E}[X_1 \mid k] &= \mu_k, & \mathbb{E}[X_t \mid k] &= t\mu_k, \\ \mathrm{Cov}(X_1 \mid k) &= \sigma^2 I_2, & \mathrm{Cov}(X_t \mid k) &= C_t I_2, \\ \mathrm{Cov}(X_1, X_t \mid k) &= \mathrm{Cov}(X_1, (1-t)X_0 + tX_1 \mid k) = t\sigma^2 I_2. \end{aligned}$$

The Gaussian conditioning formula therefore gives

$$\begin{aligned} m_{t,k}(x) &= \mathbb{E}[X_1 \mid X_t = x, k] \\ &= \mu_k + t\sigma^2 I_2 (C_t I_2)^{-1}(x - t\mu_k) \\ &= \mu_k + \frac{t\sigma^2}{C_t}(x - t\mu_k) \\ &= \left(1 - \frac{t^2\sigma^2}{C_t}\right)\mu_k + \frac{t\sigma^2}{C_t}x \\ &= \frac{C_t - t^2\sigma^2}{C_t}\mu_k + \frac{t\sigma^2}{C_t}x \\ &= \frac{(1-t)^2}{C_t}\mu_k + \frac{t\sigma^2}{C_t}x. \end{aligned} \tag{34}$$

For a nonempty component subset A, define

$$\begin{aligned} \pi_{k|A}(x,t) &= \frac{\pi_k(x,t)}{\sum_{j\in A}\pi_j(x,t)}, \qquad k \in A, \\ \bar{\mu}_t(x) &= \sum_k \pi_k(x,t)\mu_k, \\ \bar{\mu}_A(x) &= \sum_{k\in A} \pi_{k|A}(x,t)\mu_k. \end{aligned} \tag{35}$$

Averaging (34) over all components or over A gives

$$m_t^{\mathrm{u}}(x) = \frac{(1-t)^2}{C_t}\bar{\mu}_t(x) + \frac{t\sigma^2}{C_t}x, \tag{36}$$

$$m_t^{\mathrm{c}}(x) = \frac{(1-t)^2}{C_t}\bar{\mu}_A(x) + \frac{t\sigma^2}{C_t}x. \tag{37}$$

The shared state term cancels, yielding

$$\Delta_t(x) = m_t^{\mathrm{c}}(x) - m_t^{\mathrm{u}}(x) = \alpha_t\big(\bar{\mu}_A(x) - \bar{\mu}_t(x)\big). \tag{38}$$

Proposition C.1 (GMM gap decay and cap release)**.** *For every x and $0 \le t < 1$,*

$$\|\Delta_t(x)\| \le 2r\alpha_t = \frac{2r}{1 + \sigma^2[t/(1-t)]^2} \xrightarrow[t\uparrow 1]{} 0, \tag{39}$$

and the envelope $2r\alpha_t$ decreases strictly on $(0,1)$. If $x_t \to x_1$, then $m_t^{\mathrm{u}}(x_t), m_t^{\mathrm{c}}(x_t) \to x_1$. The nominal PMC-CFG scale is feasible whenever

$$2r(\lambda-1)\alpha_t \le (\Gamma-1)\left\|m_t^{\mathrm{c}}(x)\right\|. \tag{40}$$

Consequently, PMC-CFG recovers $\gamma_t^{\mathrm{eff}} = \lambda$ sufficiently near every nonzero GMM endpoint.

Proof. Both weighted centers lie in the convex hull of $\{\mu_k\}$, so $\|\bar{\mu}_t\| \le r$ and $\|\bar{\mu}_A\| \le r$. The triangle inequality therefore gives

$$\|\Delta_t(x)\| \le \alpha_t\big(\|\bar{\mu}_A(x)\| + \|\bar{\mu}_t(x)\|\big) \le 2r\alpha_t.$$

For $0 \le t < 1$,

$$\frac{(1-t)^2}{C_t} = \left[1+\sigma^2\left(\frac{t}{1-t}\right)^2\right]^{-1}, \qquad \frac{\mathrm{d}\alpha_t}{\mathrm{d}t} = -\frac{2\sigma^2 t(1-t)}{C_t^2} < 0 \quad (0<t<1),$$

so α_t decreases from 1 to 0. Since $\rho_t = t\sigma^2/C_t \to 1$, equations (36) and (37) also give $m_t^{\mathrm{u}}(x_t), m_t^{\mathrm{c}}(x_t) \to x_1$ whenever $x_t \to x_1$. Finally, under (40),

$$\begin{aligned}\left\|m_t^{\mathrm{c}} + (\lambda-1)\Delta_t\right\| &\le \left\|m_t^{\mathrm{c}}\right\| + (\lambda-1)\left\|\Delta_t\right\| \\ &\le \Gamma\left\|m_t^{\mathrm{c}}\right\|,\end{aligned}$$

so the nominal scale is feasible. Along a trajectory converging to $x_1 \neq 0$, $\alpha_t \to 0$ and $m_t^{\mathrm{c}} \to x_1$; therefore (40) holds for all sufficiently large t. □

Writing $\bar{\mu}$ for the weighted center of either field, using (29) gives the corresponding exact velocity

$$\begin{aligned}v_t(x) &= \frac{m_t(x)-x}{1-t} \\ &= \frac{1-t}{C_t}\bar{\mu}(x) + \frac{t\sigma^2 - C_t}{(1-t)C_t}x \\ &= \frac{1-t}{C_t}\bar{\mu}(x) + \frac{t\sigma^2-(1-t)}{C_t}x,\end{aligned} \tag{41}$$

where the last equality uses

$$\begin{aligned}t\sigma^2 - C_t &= t\sigma^2 - (1-t)^2 - t^2\sigma^2 \\ &= (1-t)[t\sigma^2 - (1-t)].\end{aligned}$$

For the singleton condition $A=\{0\}$, $\bar{\mu}_A = \mu_0$. At $t=0$, $X_t = X_0$ is independent of X_1, and a centered mixture therefore satisfies

$$\begin{aligned}m_0^{\mathrm{u}}(x) &= \mathbb{E}[X_1] = \sum_k w_k\mu_k = 0, \\ m_0^{\mathrm{c}}(x) &= \mathbb{E}[X_1 \mid c=0] = \mu_0.\end{aligned}$$

Substitution into (4) gives

$$m_0^{\gamma} = \mu_0 + (\gamma-1)(\mu_0 - 0) = \gamma\mu_0,$$

which proves (19).

C.2 Force and curvature

Using the responsibilities from (35), define

$$\Sigma_A = \sum_{k\in A}\pi_{k|A}(\mu_k - \bar{\mu}_A)(\mu_k - \bar{\mu}_A)^{\top}.$$

We first differentiate the unconditional responsibilities. Writing $\eta_k = \log w_k + (t/C_t)x^\top \mu_k$ and $\pi_k = e^{\eta_k}/\sum_j e^{\eta_j}$, the quotient rule gives

$$\begin{aligned}\nabla_x \pi_k &= \pi_k \left(\nabla_x \eta_k - \sum_j \pi_j \nabla_x \eta_j \right) \\ &= \frac{t}{C_t} \pi_k \left(\mu_k - \sum_j \pi_j \mu_j \right) \\ &= \frac{t}{C_t} \pi_k (\mu_k - \bar{\mu}_t).\end{aligned}$$

Consequently,

$$\begin{aligned}\nabla_x \bar{\mu}_t &= \sum_k \mu_k (\nabla_x \pi_k)^\top \\ &= \frac{t}{C_t} \sum_k \pi_k \mu_k (\mu_k - \bar{\mu}_t)^\top \\ &= \frac{t}{C_t} \left(\sum_k \pi_k \mu_k \mu_k^\top - \bar{\mu}_t \bar{\mu}_t^\top \right) = \frac{t}{C_t} \Sigma_\pi.\end{aligned}$$

The same calculation after renormalizing over A gives $\nabla_x \bar{\mu}_A = (t/C_t)\Sigma_A$.

For completeness, the condition posterior is a ratio of two mixture sums:

$$h_t^A(x) = \log p(A \mid x) = \log \sum_{k \in A} w_k e^{-\|x - t\mu_k\|^2/(2C_t)} - \log \sum_j w_j e^{-\|x - t\mu_j\|^2/(2C_t)}.$$

Differentiating the first log-sum gives $-(x - t\bar{\mu}_A)/C_t$, and differentiating the second gives $-(x - t\bar{\mu}_t)/C_t$. Their common $-x/C_t$ terms cancel. Therefore

$$\nabla_x h_t^A(x) = -\frac{x - t\bar{\mu}_A}{C_t} + \frac{x - t\bar{\mu}_t}{C_t} = \frac{t}{C_t}(\bar{\mu}_A - \bar{\mu}_t), \tag{42}$$

$$\nabla_x^2 h_t^A(x) = \frac{t}{C_t}(\nabla_x \bar{\mu}_A - \nabla_x \bar{\mu}_t) = \frac{t^2}{C_t^2}(\Sigma_A - \Sigma_\pi). \tag{43}$$

By (9), the extra CFG force is

$$v_t^\gamma - v_t^{\mathrm{c}} = (\gamma - 1)\frac{1 - t}{t} \nabla_x h_t^A = (\gamma - 1)\frac{1 - t}{C_t}(\bar{\mu}_A - \bar{\mu}_t),$$

and one more derivative gives

$$\nabla_x (v_t^\gamma - v_t^{\mathrm{c}}) = (\gamma - 1)\frac{t(1 - t)}{C_t^2}(\Sigma_A - \Sigma_\pi).$$

This proves the general GMM force and curvature corrections.

For the eight equally weighted centers, $\mu_k = r u_k$ with angles $\theta_k = 2\pi k/8$. At $x = 0$, all logits are equal, so $\pi_k = 1/8$ and

$$\bar{\mu}_t(0) = \frac{r}{8} \sum_{k=0}^{7} (\cos\theta_k, \sin\theta_k)^\top = 0,$$

$$\Sigma_\pi(0) = \frac{r^2}{8} \sum_{k=0}^{7} u_k u_k^\top = \frac{r^2}{8} \begin{pmatrix} \sum_k \cos^2\theta_k & \sum_k \sin\theta_k \cos\theta_k \\ \sum_k \sin\theta_k \cos\theta_k & \sum_k \sin^2\theta_k \end{pmatrix} = \frac{r^2}{2} I_2,$$

where the three trigonometric sums are $4, 0, 4$. For the singleton condition $A = \{0\}$, $\bar{\mu}_A = \mu_0$ and $\Sigma_A = 0$. Direct substitution yields

$$(v_t^\gamma - v_t^{\mathrm{c}})(0) = (\gamma - 1)\frac{1 - t}{C_t} \mu_0,$$

$$\nabla_x (v_t^\gamma - v_t^{\mathrm{c}})(0) = (\gamma - 1)\frac{t(1 - t)}{C_t^2} \left(0 - \frac{r^2}{2} I_2 \right),$$

which proves (20) and gives

$$\nabla_x(v_t^\gamma - v_t^{\mathrm{c}})(0) = -(\gamma-1)\frac{t(1-t)r^2}{2C_t^2} I_2. \tag{44}$$

For the pure conditional field, substituting $\bar{\mu}_A = \mu_0$ into (41) gives

$$v_t^{\mathrm{c}}(x) = \frac{1-t}{C_t}\mu_0 + \alpha_t x, \qquad \alpha_t = \frac{t\sigma^2 - (1-t)}{C_t} = \frac{t(1+\sigma^2)-1}{C_t}.$$

The extra force has value $(\gamma-1)(1-t)\mu_0/C_t$ and first-order Jacobian from (44). Hence its Taylor expansion at the origin is

$$v_t^\gamma(x) = v_t^\gamma(0) + \nabla_x v_t^\gamma(0)x + o(\|x\|),$$

with $v_t^\gamma(0) = \gamma(1-t)\mu_0/C_t$. Dropping the remainder gives

$$v_t^\gamma(x) \approx \gamma\frac{1-t}{C_t}\mu_0 + \left[\alpha_t - (\gamma-1)\frac{t(1-t)r^2}{2C_t^2}\right] x. \tag{45}$$

At $t=0$, the bracketed coefficient equals -1, while the constant term equals $\gamma\mu_0$. Thus the frozen linear field is exactly $\gamma\mu_0 - x$ and its zero is $x = \gamma\mu_0$. Continuity gives the stated small-t limit.

For $A = \{0,2\}$, define the orthonormal directions $u_+ = (u_0+u_2)/\sqrt{2}$ and $u_- = (u_0-u_2)/\sqrt{2}$. At the origin the two conditional responsibilities equal $1/2$, so

$$\begin{aligned}
\bar{\mu}_A &= \frac{1}{2}(ru_0 + ru_2) = \frac{r}{\sqrt{2}}u_+, \\
\mu_0 - \bar{\mu}_A = \frac{r}{\sqrt{2}}u_-, &\qquad \mu_2 - \bar{\mu}_A = -\frac{r}{\sqrt{2}}u_-, \\
\Sigma_A &= \frac{1}{2}\frac{r^2}{2}u_-u_-^\top + \frac{1}{2}\frac{r^2}{2}u_-u_-^\top = \frac{r^2}{2}u_-u_-^\top.
\end{aligned}$$

Therefore

$$\Sigma_A u_+ = 0, \qquad \Sigma_A u_- = \frac{r^2}{2}u_-,$$

whereas the unconditional covariance at the origin is $\Sigma_\pi = (r^2/2)I_2$ and hence $\Sigma_\pi u_\pm = (r^2/2)u_\pm$. Multiplying the Hessian difference (43) by $(\gamma-1)(1-t)/t$ gives

$$u_+^\top(H_t^\gamma - H_t^{\mathrm{c}})u_+ = -(\gamma-1)\frac{t(1-t)r^2}{2C_t^2}, \tag{46}$$

$$u_-^\top(H_t^\gamma - H_t^{\mathrm{c}})u_- = 0. \tag{47}$$

Finally, differentiating the pure conditional velocity gives

$$\nabla_x v_t^{\mathrm{c}} = \alpha_t I_2 + \frac{t(1-t)}{C_t^2}\Sigma_A.$$

Along u_-, the covariance term equals $t(1-t)r^2/(2C_t^2)$. The CFG correction along this direction is zero, so the complete splitting eigenvalue is

$$\lambda_-(t) = \alpha_t + \frac{t(1-t)r^2}{2C_t^2}.$$

It is independent of γ at the symmetric center.

C.3 Branch reweighting for adjacent conditions

We first prove the continuous equal-weight middle-branch effect. Rotate the coordinate system so that the middle direction u_1 is the positive horizontal axis. The three directions in $A = \{0,1,2\}$ then have angles $-\delta, 0, \delta$, where $\delta = \pi/4$. On a fixed shell, write $x_{\rho,\phi} = \rho(\cos\phi, \sin\phi)^\top$ in these rotated coordinates and set

$$\kappa = \frac{tr\rho}{C_t}, \qquad a = \kappa\cos\phi, \qquad b = \kappa\sin\phi, \qquad q = \frac{1}{\sqrt{2}}.$$

For $0 \le \phi \le \delta$, we have $a \ge b \ge 0$. Common Gaussian factors and equal mixture weights cancel from the condition posterior. Its odds between the conditioned set A and its complement B are therefore

$$\frac{P_t(\rho,\phi)}{1-P_t(\rho,\phi)} = \frac{N(a,b)}{M(a,b)}, \tag{48}$$

where

$$N(a,b) = e^{a} + 2e^{qa}\cosh(qb),$$
$$M(a,b) = e^{-a} + 2e^{-qa}\cosh(qb) + 2\cosh b.$$

Reflection across the middle ray preserves both A and the full eight-point mixture, immediately giving $P_t(\rho,\phi) = P_t(\rho,-\phi)$.

It remains to establish strict angular decrease. Let $L(a,b) = \log N(a,b) - \log M(a,b)$. Its partial derivative with respect to a is strictly positive: every direction in A has horizontal coordinate at least q, whereas every direction in B has horizontal coordinate at most zero. Equivalently, $\partial_a L$ is the difference between the exponentially weighted mean horizontal coordinates in A and B.

For $b > 0$, the corresponding weighted vertical mean in B is strictly larger than that in A, so $\partial_b L < 0$. To verify the sign directly, cross multiplication reduces that mean difference to the positive quantity

$$e^{a}\sinh b + 2e^{qa}\big[\cosh(qb)\sinh b - q\sinh(qb)\cosh b\big] + 2q\sinh(qb)\sinh((1-q)a).$$

The bracket is positive because

$$\tanh b - q\tanh(qb) > 0 \qquad (b>0,\ 0<q<1).$$

Since $\partial_\phi a = -b$ and $\partial_\phi b = a$, we obtain

$$\partial_\phi L = -b\,\partial_a L + a\,\partial_b L < 0 \qquad (0 < \phi \le \delta). \tag{49}$$

The logistic map from odds to probability is strictly increasing, so the same sign holds for $\partial_\phi P_t(\rho,\phi)$. Together with reflection symmetry, this shows that the posterior is symmetric about the middle ray and strictly decreases for $0 < \phi \le \delta$.

For $x_{\rho,\phi}$, define the angular unit vector $e_\phi = (-\sin\phi, \cos\phi)^\top$, so that $e_\phi^\top \nabla_x = \rho^{-1}\partial_\phi$. Together, (10) and (49) show that the angular CFG correction points toward the middle ray for $0 < t < 1$ and $\gamma > 1$. This is an instantaneous force; terminal branch occupancy follows only after integrating the non-autonomous ODE.

We next prove Proposition 3.5. Let $A = \{0,1\}$, assume $w_0 > w_1 > 0$, and impose reflection symmetry on the complement,

$$w_2 = w_7, \qquad w_3 = w_6, \qquad w_4 = w_5.$$

At $x_{0,t}$ and $x_{1,t}$, canceling the same common radial factor gives the unnormalized conditional likelihoods

$$N_0(a) = w_0 e^{a} + w_1 e^{a/\sqrt{2}},$$
$$N_1(a) = w_0 e^{a/\sqrt{2}} + w_1 e^{a}, \tag{50}$$
$$N_0(a) - N_1(a) = (w_0 - w_1)(e^{a} - e^{a/\sqrt{2}}) > 0. \tag{51}$$

To verify the complement contribution, reflect across the angular bisector between u_0 and u_1. This reflection exchanges the component pairs $2 \leftrightarrow 7$, $3 \leftrightarrow 6$, and $4 \leftrightarrow 5$, and maps $x_{0,t}$ to $x_{1,t}$. The assumed equality within every pair makes the sum of complement likelihoods, denoted $E(a)$, equal at the two centers. Hence

$$p_k^A(t) := p(A \mid x_{k,t}) = \frac{N_k(a)}{N_k(a) + E(a)}, \qquad p_0^A(t) > p_1^A(t).$$

The last inequality follows because $z \mapsto z/(z+E)$ is strictly increasing for $E > 0$. For $0 < t < 1$ and $\gamma > 1$, the difference between the CFG potential corrections at the two centers is

$$(\gamma - 1)\frac{1-t}{t}\log\frac{p_0^A(t)}{p_1^A(t)} > 0.$$

Thus, under complement symmetry, the time-varying posterior potential favors the already heavier branch at the two interpolated centers, proving Proposition 3.5. For asymmetric complements, the ordering follows from evaluating the two center contributions jointly.

C.4 Proof of the confidence–density theorem

Proof of Theorem 3.4. On the ray $x = su_0$, define

$$\vartheta_k = \theta_k - \theta_0, \qquad \kappa = \frac{str}{C_t}.$$

Because $\mu_k = ru_k$ and $u_0^\top u_k = \cos\vartheta_k$, the squared distance expands as

$$\|su_0 - t\mu_k\|^2 = s^2 + t^2r^2 - 2str\cos\vartheta_k.$$

Thus every likelihood contains the common factor $\exp[-(s^2 + t^2r^2)/(2C_t)]$, while the component-dependent factor is $e^{\kappa\cos\vartheta_k}$. Canceling the common factor in (33) gives

$$P_0(s,t) = \frac{w_0 e^{\kappa}}{\sum_k w_k e^{\kappa\cos\vartheta_k}}, \qquad \pi_k(s,t) = \frac{w_k e^{\kappa\cos\vartheta_k}}{\sum_j w_j e^{\kappa\cos\vartheta_j}}. \tag{52}$$

Taking logs gives

$$\log P_0 = \log w_0 + \kappa - \log\sum_k w_k e^{\kappa\cos\vartheta_k}.$$

Since $\partial_s\kappa = tr/C_t$, differentiation gives

$$\begin{aligned}
\partial_s \log P_0 &= \frac{tr}{C_t}\left(1 - \frac{\sum_k w_k e^{\kappa\cos\vartheta_k}\cos\vartheta_k}{\sum_j w_j e^{\kappa\cos\vartheta_j}}\right) \\
&= \frac{tr}{C_t}\left(1 - \sum_k \pi_k\cos\vartheta_k\right) \geq 0, \qquad (53) \\
\partial_s^2 \log P_0 &= -\left(\frac{tr}{C_t}\right)^2\left[\sum_k \pi_k\cos^2\vartheta_k - \left(\sum_k \pi_k\cos\vartheta_k\right)^2\right] \\
&= -\left(\frac{tr}{C_t}\right)^2 \mathrm{Var}_\pi(\cos\vartheta_k) \leq 0. \qquad (54)
\end{aligned}$$

The first sign follows from $\cos\vartheta_k \leq 1$. For the second derivative, the intermediate responsibility derivative is

$$\partial_s\pi_k = \frac{tr}{C_t}\pi_k\left(\cos\vartheta_k - \sum_j \pi_j\cos\vartheta_j\right),$$

and summing $\cos\vartheta_k\,\partial_s\pi_k$ produces the variance above.

For the asymptotic posterior, divide numerator and denominator of (52) by w_0e^{κ}:

$$P_0(s,t) = \left[1 + \sum_{j\neq 0}\frac{w_j}{w_0}e^{-\kappa(1-\cos\vartheta_j)}\right]^{-1}.$$

Uniqueness of u_0 means $1 - \cos\vartheta_j > 0$ for $j \neq 0$, hence $P_0(s,t) \to 1$.

The marginal density on the ray is

$$p_t(su_0) = \frac{e^{-(s^2+t^2r^2)/(2C_t)}}{2\pi C_t}\sum_k w_k e^{\kappa\cos\vartheta_k}.$$

Taking the logarithm and differentiating the Gaussian factor and the log-sum separately gives

$$\partial_s \log p_t(su_0) = -\frac{s}{C_t} + \frac{tr}{C_t}\sum_k \pi_k\cos\vartheta_k \leq -\frac{s - tr}{C_t}.$$

The inequality holds because the responsibility-weighted average of the cosines is at most one. This is strictly negative for $s > tr$, while the posterior derivative is nonnegative.

Finally, factor the $k = 0$ term from the mixture sum:

$$p_t(su_0) = \frac{w_0}{2\pi C_t} \exp\left[-\frac{(s-tr)^2}{2C_t}\right] \left[1 + \sum_{j\neq 0} \frac{w_j}{w_0} e^{-\kappa(1-\cos\vartheta_j)}\right].$$

The bracket tends to one by the same strict angular gaps used for the posterior. Therefore

$$p_t(su_0) \sim \frac{w_0}{2\pi C_t} \exp\left[-\frac{(s-tr)^2}{2C_t}\right] \to 0.$$

All claims follow. □

D PMC-CFG Proofs and Implementation

D.1 Feasibility, maximality, and degenerate cases

Proof of Theorem 4.1. Squaring the nonnegative norms in (23) and expanding gives

$$\begin{aligned}\|m_t^c + \beta\Delta_t\|^2 - \Gamma^2\|m_t^c\|^2 &= \left(\|m_t^c\|^2 + 2\beta\langle m_t^c, \Delta_t\rangle + \beta^2\|\Delta_t\|^2\right) - \Gamma^2\|m_t^c\|^2 \\ &= \|\Delta_t\|^2\beta^2 + 2\langle m_t^c, \Delta_t\rangle\beta + (1-\Gamma^2)\|m_t^c\|^2 \\ &= a\beta^2 + 2b\beta + (1-\Gamma^2)q.\end{aligned}$$

Thus the norm constraint is equivalent to $f(\beta) := a\beta^2 + 2b\beta + (1-\Gamma^2)q \leq 0$. Since $\Gamma \geq 1$ and $q \geq 0$, the constant term is nonpositive; therefore $\beta = 0$ is feasible and the program is never empty.

Suppose $a > 0$. Applying the quadratic formula to $a\beta^2 + 2b\beta + (1-\Gamma^2)q = 0$ gives the discriminant

$$(2b)^2 - 4a(1-\Gamma^2)q = 4[b^2 + (\Gamma^2-1)aq] \geq 0$$

and roots

$$\beta_\pm = \frac{-2b \pm 2\sqrt{b^2 + (\Gamma^2-1)aq}}{2a} = \frac{-b \pm \sqrt{b^2 + (\Gamma^2-1)aq}}{a}.$$

Because the square root is at least $|b|$, its two numerators satisfy $-b - \sqrt{\cdot} \leq 0$ and $-b + \sqrt{\cdot} \geq 0$; hence $\beta_- \leq 0 \leq \beta_+$. Equivalently, their product is $(1-\Gamma^2)q/a \leq 0$. Since $a > 0$, the quadratic opens upward and its nonpositive set is exactly $[\beta_-, \beta_+]$. Intersecting it with the nominal interval $[0, \lambda - 1]$ gives

$$[0, \lambda-1] \cap [\beta_-, \beta_+] = [0, \min\{\lambda - 1, \beta_+\}].$$

The largest feasible element is therefore (27), with $\beta_+ = \beta^{\text{cap}}$ from (26).

If $a = 0$, then $a = \|\Delta_t\|^2$ implies $\Delta_t = 0$ and consequently $b = \langle m_t^c, \Delta_t\rangle = 0$. Every candidate mean equals m_t^c, and the constraint reduces to $\|m_t^c\| \leq \Gamma\|m_t^c\|$, which holds because $\Gamma \geq 1$. The output velocity is also independent of β, so choosing $\lambda - 1$ preserves the nominal convention.

In both cases the selected β_t^* belongs to the feasible set. Substituting it into the original, unsquared constraint yields

$$\|m_t^c + \beta_t^*\Delta_t\| \leq \Gamma\|m_t^c\|,$$

which is (28). If $\lambda - 1$ is feasible, it is already the largest point in $[0, \lambda - 1]$ and is therefore selected; nominal CFG is then unchanged. □

D.2 Directional form, invariance, and the GMM corollary

Proposition D.1 (Directional cap). *If $m_t^c \neq 0$ and $\Delta_t \neq 0$, and θ is their angle, then*

$$\beta^{\text{cap}} = \frac{\|m_t^c\|}{\|\Delta_t\|}\left[-\cos\theta + \sqrt{\cos^2\theta + \Gamma^2 - 1}\right]. \tag{55}$$

For aligned, orthogonal, and opposing vectors, respectively,

$$\cos\theta = 1: \quad \beta^{\mathrm{cap}} = (\Gamma - 1)\frac{\|m_t^{\mathrm{c}}\|}{\|\Delta_t\|},$$
$$\cos\theta = 0: \quad \beta^{\mathrm{cap}} = \sqrt{\Gamma^2 - 1}\frac{\|m_t^{\mathrm{c}}\|}{\|\Delta_t\|},$$
$$\cos\theta = -1: \quad \beta^{\mathrm{cap}} = (\Gamma + 1)\frac{\|m_t^{\mathrm{c}}\|}{\|\Delta_t\|}.$$

Proof. Substitute $a = \|\Delta_t\|^2$, $q = \|m_t^{\mathrm{c}}\|^2$, and $b = \|m_t^{\mathrm{c}}\|\|\Delta_t\|\cos\theta$ into (26). The numerator becomes

$$\begin{aligned}
&- \|m_t^{\mathrm{c}}\|\|\Delta_t\|\cos\theta \\
&\quad + \sqrt{\|m_t^{\mathrm{c}}\|^2\|\Delta_t\|^2\cos^2\theta + (\Gamma^2 - 1)\|\Delta_t\|^2\|m_t^{\mathrm{c}}\|^2} \\
&= \|m_t^{\mathrm{c}}\|\|\Delta_t\|\left[-\cos\theta + \sqrt{\cos^2\theta + \Gamma^2 - 1}\right].
\end{aligned}$$

Dividing by $a = \|\Delta_t\|^2$ proves (55). For $\cos\theta = 1$, the bracket is $-1 + \sqrt{\Gamma^2} = \Gamma - 1$; for $\cos\theta = 0$, it is $\sqrt{\Gamma^2 - 1}$; and for $\cos\theta = -1$, it is $1 + \sqrt{\Gamma^2} = \Gamma + 1$, because $\Gamma \geq 1$. These are the three displayed cases. □

Proposition D.2 (Coordinate behavior)**.** *The origin-centered constraint is invariant under a common positive scalar rescaling and an orthogonal transformation. It is not generally invariant under translation or anisotropic scaling. If an anchor r_t is transformed together with the means, the constraint*

$$\|m_t^{\mathrm{cap}} - r_t\| \leq \Gamma\,\|m_t^{\mathrm{c}} - r_t\| \tag{56}$$

is translation covariant.

Proof. Let $T(x) = sQx$ with $s > 0$ and $Q^\top Q = I$. The transformed candidate and baseline means are

$$\widetilde{m}^{\mathrm{cap}} = sQ(m^{\mathrm{c}} + \beta\Delta), \qquad \widetilde{m}^{\mathrm{c}} = sQm^{\mathrm{c}}.$$

Therefore

$$\begin{aligned}
\|\widetilde{m}^{\mathrm{cap}}\|^2 &= s^2(m^{\mathrm{c}} + \beta\Delta)^\top Q^\top Q(m^{\mathrm{c}} + \beta\Delta) = s^2\|m^{\mathrm{c}} + \beta\Delta\|^2, \\
\Gamma^2\|\widetilde{m}^{\mathrm{c}}\|^2 &= s^2\Gamma^2\|m^{\mathrm{c}}\|^2.
\end{aligned}$$

The common positive factor s^2 cancels, so the feasible set and its largest β are unchanged.

Under translation by h, the constraint would compare $\|m^{\mathrm{c}} + \beta\Delta + h\|$ with $\Gamma\|m^{\mathrm{c}} + h\|$. The two added cross terms depend differently on β, so this is not generally equivalent to the original constraint. Under an anisotropic linear map A, a squared norm becomes $\|Ay\|^2 = y^\top A^\top Ay$; it is a common scalar multiple of the Euclidean norm for all y only when $A^\top A = s^2 I$.

For the anchored form, translate $m^{\mathrm{cap}}, m^{\mathrm{c}}, r$ by the same h. Then

$$\begin{aligned}
(m^{\mathrm{cap}} + h) - (r + h) &= m^{\mathrm{cap}} - r, \\
(m^{\mathrm{c}} + h) - (r + h) &= m^{\mathrm{c}} - r.
\end{aligned}$$

Both sides of (56) are therefore unchanged, proving translation covariance. □

Corollary D.3 (Exact initial GMM clipping)**.** *For the centered circular GMM conditioned on component* 0*,*

$$m_0^{\mathrm{cap}} = \min\{\lambda, \Gamma\}\mu_0. \tag{57}$$

Proof. Appendix C established $m_0^{\mathrm{c}} = \mu_0$ and $m_0^{\mathrm{u}} = 0$, hence $\Delta_0 = m_0^{\mathrm{c}} - m_0^{\mathrm{u}} = \mu_0$. Every candidate mean is

$$m_0^{\mathrm{c}} + \beta\Delta_0 = (1 + \beta)\mu_0.$$

Because $\beta \geq 0$, $1 + \beta > 0$, and because $\|\mu_0\| = r > 0$, the relative-ball constraint becomes

$$\begin{aligned} \|(1+\beta)\mu_0\| &\leq \Gamma\|\mu_0\| \\ (1+\beta)r &\leq \Gamma r \\ \beta &\leq \Gamma - 1. \end{aligned}$$

The nominal interval independently gives $\beta \leq \lambda - 1$. Thus the feasible interval is $[0, \min\{\lambda - 1, \Gamma - 1\}]$. Maximizing β yields $1 + \beta_0^* = \min\{\lambda, \Gamma\}$, proving the claim. □

D.3 Feedback Jacobian and well-posedness

Let $d_t = v_t^{\mathrm{c}} - v_t^{\mathrm{u}}$. In a region where the active set is fixed and β_t^* is differentiable, the product rule gives

$$\nabla_x v_t^{\mathrm{cap}} = \nabla_x v_t^{\mathrm{c}} + \beta_t^* \nabla_x d_t + d_t (\nabla_x \beta_t^*)^\top. \tag{58}$$

Indeed, for output coordinate i and input coordinate j,

$$\begin{aligned} \partial_{x_j}(v_t^{\mathrm{cap}})_i &= \partial_{x_j}(v_t^{\mathrm{c}})_i + \partial_{x_j}(\beta_t^*(d_t)_i) \\ &= \partial_{x_j}(v_t^{\mathrm{c}})_i + \beta_t^* \partial_{x_j}(d_t)_i + (d_t)_i \partial_{x_j}\beta_t^*, \end{aligned}$$

which is exactly (58). Taking the trace yields

$$\nabla\cdot v_t^{\mathrm{cap}} = \nabla\cdot v_t^{\mathrm{c}} + \beta_t^* \nabla\cdot d_t + (\nabla_x \beta_t^*)^\top d_t.$$

The conditional field and d_t are spatial gradients under fixed CFG, so their Jacobians are symmetric. Writing $g_t = \nabla_x \beta_t^*$, the only possibly nonsymmetric term is $d_t g_t^\top$. Since its transpose is $g_t d_t^\top$, the antisymmetric part of the controlled Jacobian is exactly

$$\frac{1}{2}\left[d_t(\nabla_x\beta_t^*)^\top - (\nabla_x\beta_t^*)d_t^\top\right], \tag{59}$$

It vanishes when $d_t g_t^\top = g_t d_t^\top$. For nonzero vectors this equality holds exactly when d_t and g_t are collinear; it also holds when either is zero. This establishes the stated local integrability condition.

The derivative of the active cap can also be written explicitly without differentiating the square-root formula. Define

$$F(x, \beta) = a(x)\beta^2 + 2b(x)\beta + (1 - \Gamma^2)q(x).$$

On the active boundary, $F(x, \beta_t^*(x)) = 0$. Spatial differentiation gives

$$0 = \nabla_x F = \beta_t^{*2}\nabla_x a + 2\beta_t^* \nabla_x b + (1-\Gamma^2)\nabla_x q + (2a\beta_t^* + 2b)\nabla_x \beta_t^*.$$

If $a\beta_t^* + b \neq 0$, solving for the last term yields

$$\nabla_x \beta_t^* = -\frac{\beta_t^{*2}\nabla_x a + 2\beta_t^*\nabla_x b + (1-\Gamma^2)\nabla_x q}{2(a\beta_t^* + b)}. \tag{60}$$

The denominator is nonzero precisely away from a repeated root. In the inactive region $\beta_t^* = \lambda - 1$ is constant in x, so $\nabla_x \beta_t^* = 0$ and the feedback term disappears.

When the discriminant is positive and the denominator above is nonzero, the active root is smooth by the implicit-function theorem. At the switching surface $\beta^{\mathrm{cap}} = \lambda - 1$, the hard minimum is continuous and piecewise smooth. Under the usual local-Lipschitz assumptions on model fields, this piecewise-smooth feedback field admits the standard Carathéodory interpretation. A smooth approximation to the minimum may improve numerical derivatives, but its feasibility must be checked separately; smoothness alone does not preserve (28).

D.4 Stable per-sample implementation

We log the controlled cap ratio

$$\chi_t = \frac{\|m_t^{\mathrm{c}} + \beta_t^* \Delta_t\|}{\|m_t^{\mathrm{c}}\| + \varepsilon_m}. \tag{61}$$

The positive root admits a cancellation-resistant form when $b \geq 0$. Let $D = b^2 + (\Gamma^2 - 1)aq$. Multiplying numerator and denominator by the conjugate gives

$$\frac{-b + \sqrt{D}}{a} = \frac{(-b + \sqrt{D})(b + \sqrt{D})}{a(b + \sqrt{D})} = \frac{D - b^2}{a(b + \sqrt{D})} = \frac{(\Gamma^2 - 1)q}{b + \sqrt{D}}.$$

For $b \geq 0$, the latter avoids subtracting two nearly equal positive numbers. For $b < 0$, the original numerator $-b + \sqrt{D}$ is an addition and is already stable. If $b + \sqrt{D} = 0$, necessarily $b = 0$ and $D = 0$; the positive root is then zero, so it is assigned directly rather than evaluated as $0/0$. For each sample and ODE evaluation:

1. Evaluate v_t^{c} and v_t^{u}, then compute $m_t^{\mathrm{c}} = x + (1 - t)v_t^{\mathrm{c}}$ and $\Delta_t = (1 - t)(v_t^{\mathrm{c}} - v_t^{\mathrm{u}})$.
2. Test the nominal candidate $m_t^{\mathrm{c}} + (\lambda - 1)\Delta_t$ directly. If it satisfies the bound, set $\beta_t^* = \lambda - 1$.
3. Otherwise compute a, b, q from (25). If $a \leq \varepsilon_a$, use the conservative feasible value $\beta_t^* = 0$.
4. For $a > \varepsilon_a$, form $D = b^2 + (\Gamma^2 - 1)aq$ and clamp it below by zero only to remove negative roundoff. Use the rationalized root above when $b \geq 0$ and $b + \sqrt{D} > 0$, use (26) when $b < 0$, and set the root to zero when $b + \sqrt{D} = 0$. Set $\beta_t^* = \min\{\lambda - 1, \beta^{\mathrm{cap}}\}$.
5. Use (24) and log (61) together with $\mathbf{1}\{\beta_t^* < \lambda - 1\}$.

The direct nominal test avoids unstable division when $\|\Delta_t\|$ is small. The fallback $\beta = 0$ remains feasible for every $\Gamma \geq 1$. Norms and inner products must reduce over each sample's latent dimensions, never over the batch.

E Detailed Experimental Protocol

E.1 Eight-point distribution

The endpoint experiments use a two-dimensional mixture with $K = 8$ component means placed clockwise on a circle of radius $r = 1$. Component 0 is at angle $\pi/2$, and component k is at $\theta_k = \pi/2 - k\pi/4$. Every component has isotropic standard deviation $\sigma_k = 0.03$. For the singleton $A = \{0\}$ and adjacent triple $A = \{0, 1, 2\}$, the unconditional mixture weights are $w_k = 1/8$ for all k. The adjacent-pair experiment $A = \{0, 1\}$ instead uses deliberately imbalanced unconditional mixture weights,

$$(w_0, \ldots, w_7) = (0.17, 0.08, 0.125, 0.125, 0.125, 0.125, 0.125, 0.125),$$

which induce conditional weights $68\%/32\%$ after renormalization over A. This pair is therefore an imbalance stress test rather than an equal-weight counterpart of the other two conditions.

Both conditional and unconditional fields are analytic GMM posterior-mean velocities; no neural network, learned approximation, or checkpoint is used. The conditional weights are obtained by renormalizing the unconditional weights over A, and fixed CFG is evaluated as $v_t^{\mathrm{u}} + \lambda(v_t^{\mathrm{c}} - v_t^{\mathrm{u}})$. Thus $\lambda = 1$ is exactly the analytic conditional field. The PMC-CFG runs use the same fields with nominal $\lambda = 3$ and the relative norm cap $\Gamma = 1.1$ from Section 4.

For each row of Table 1, the implementation independently draws 2×10^5 samples from $X_0 \sim \mathcal{N}(0, I_2)$ in one call and uses PyTorch's default floating-point dtype (float32). It then applies 50 explicit Euler updates on the uniform grid $t_i = i/50$. The field is evaluated at the 50 left endpoints t_i, $i = 0, \ldots, 49$, and the final state is at $t = 1$; there is no $1 - \varepsilon$ endpoint or special last-step rule. Counting one evaluation of the composite guided field as one NFE gives 50 NFEs, equivalently 50 conditional and 50 unconditional analytic subfield evaluations. The nine runs use constant nominal $\lambda \in \{1, 3\}$ and enable no time schedule, cosine-similarity gate, or relative-distance schedule. Three runs enable only the PMC norm cap with $\Gamma = 1.1$. Marginal histograms use 101 bins.

Each sampler call draws new initial noise, so different conditions and guidance scales are not paired sample by sample. Accordingly, Table 1 reports a single $n = 200000$ estimate per setting rather than a mean over repeated runs. For endpoint metrics, samples are assigned by MAP under the conditional GMM, not by an unconditional nearest-center rule. TV is one half of the ℓ_1 distance

between empirical and target conditional occupancies; mean error is $\|\widehat{\mu} - \mu_A\|_2$ before assignment; and each per-label variance ratio is $\mathrm{tr}(\widehat{\Sigma}_k)/(2\sigma_k^2)$, with $\widehat{\Sigma}_k$ the unbiased covariance of samples assigned to branch k.

The $\lambda = 1, A = \{0\}$ variance ratio near 0.571 is the expected 50-step Euler bias: the exact discrete prediction is $R_{50} = 0.5706506$, whereas the exact continuous-time transport has endpoint ratio one.

E.2 ImageNet-256 protocol

We use the released class-conditional SiT-XL/2 checkpoint for 256×256 ImageNet generation (Deng et al., 2009; Ma et al., 2024), together with its native latent normalization. If the implementation parameterizes time or noise in a variable other than the linear-path t, we convert the model output to the velocity and remaining-time factor used in $m_t = x + (1 - t)v_t$ before applying PMC-CFG. This conversion is verified on synthetic inputs; silently reusing a diffusion noise prediction as a velocity would invalidate the cap.

The evaluation draws 50 samples for each of 1,000 ImageNet classes, forming the standard 50K set. We use the official ADM/guided-diffusion Inception implementation and its reference statistics for FID, sFID, IS, precision, and recall (Dhariwal & Nichol, 2021; Heusel et al., 2017; Salimans et al., 2016; Kynkäänniemi et al., 2019). These five metrics are computed from the same 50K sample set for every method.

E.3 Stable Diffusion 3.5 protocol

We evaluate the official Stable Diffusion 3.5 Medium checkpoint (Stability AI, 2024) through the public Diffusers pipeline. We generate one 1024×1024 image for each of 5,000 prompts from the COCO 2017 validation split (Lin et al., 2014), following their fixed insertion order in the prompt file. Each run uses one persistent CUDA generator initialized with seed 42, and all methods share the same scheduler, text encoders, VAE, prompt order, and initial latent sequence. Generation uses batch size one; under CFG, the conditional and unconditional inputs are concatenated into a transformer batch of two. We use empty-text unconditional conditioning unless a baseline requires a different negative prompt, in which case that difference is reported.

We use the `FlowMatchEulerDiscreteScheduler` with 10 deterministic Euler steps, `num_train_timesteps=1000`, and a fixed flow shift of 3.0. Dynamic shifting, stochastic sampling, and Karras, exponential, beta sigma schedules are disabled. The SD3.5 pipeline runs in bfloat16, while scheduler timesteps and sigmas, PMC-CFG norms and inner products, and clean-latent estimates are computed in FP32. We test nominal guidance scales $\lambda \in \{3, 5, 7, 9\}$. Fixed CFG applies the nominal scale at every step; PMC-CFG uses $\Gamma = 1.15$ and caps each sample's guided clean-latent displacement without exceeding the nominal scale. C2FG uses its method-specific coefficient 0.2 and the normalized sampling progress p on the shifted sigma grid, giving the schedule $w(p) = w \exp(0.2p)$. APG applies its projected, rescaled, momentum-guidance update (Sadat et al., 2025) to the same conditional and unconditional predictions; its evaluated sample sets use the same nominal scales, initial latent sequence, and solver configuration.

CLIP is the mean raw cosine similarity between matched images and prompts using OpenAI CLIP ViT-B/32 (Radford et al., 2021); it is not multiplied by 100. ImageReward is the mean ImageReward-v1.0 score (Xu et al., 2023) and may be negative. CLIP and FID use evaluation batch size eight, whereas ImageReward is evaluated one image at a time. For FID, generated and COCO validation images are resized directly to 299×299, converted to 8-bit tensors, and evaluated with the default 2048-dimensional Inception features in TorchMetrics (Detlefsen et al., 2022); these values should therefore not be mixed with Clean-FID results (Parmar et al., 2022). Mean saturation is computed from the normalized 8-bit HSV saturation channel and is reported only to diagnose over- or under-saturation. Finally, diversity is improved Recall@3 (Kynkäänniemi et al., 2019): using Clean-FID's clean Inception features, a real sample is covered if it lies inside the third-nearest-neighbor ball of at least one generated sample. FID and Recall@3 use the 5,000 COCO validation images as the reference set; the remaining metrics use only each generated image and its matched prompt.

F Additional Diagnostics and Qualitative Results

F.1 Ablating per-sample control on ImageNet-256

We isolate whether PMC-CFG benefits from adapting to the current sample, rather than merely inducing a favorable time-dependent guidance schedule. For each NFE, nominal scale, and cap used in Table 2, we first run PMC-CFG on a separate calibration set of $N_{\mathrm{cal}} = 1024$ class–noise pairs, using a random seed and class/noise draws independent of the 50K evaluation set. At Euler step i, we average the effective per-sample scales,

$$\bar{\gamma}_i = \frac{1}{N_{\mathrm{cal}}} \sum_{j=1}^{N_{\mathrm{cal}}} \gamma_{i,j}^{\mathrm{eff}}, \tag{62}$$

and freeze the resulting curve. We call this time-only baseline PMC-Mean. It uses the single scalar $\bar{\gamma}_i$ for every sample at step i, $v_i^{\mathrm{mean}} = v_i^{\mathrm{u}} + \bar{\gamma}_i(v_i^{\mathrm{c}} - v_i^{\mathrm{u}})$, with no dependence on the current sample. It therefore preserves the average schedule learned from PMC-CFG while removing its per-sample feedback.

The ablation separates the benefit of scheduling from that of per-sample feedback. First, compared with fixed CFG in Table 2, PMC-Mean reduces FID in all six settings, improves sFID in five and ties in one, and improves recall in five of six. Thus the low-to-high schedule inherited from PMC-CFG—weak guidance during early structure formation and stronger guidance later—is already more effective than applying the nominal scale throughout sampling. This agrees with the effective-scale dynamics in Figure 8.

Second, Table 4 shows that the average temporal profile does not explain the full benefit of PMC-CFG. Relative to PMC-Mean, per-sample control has lower FID in five of the six settings and ties after rounding in the remaining setting, while increasing IS for every scale and solver budget. The individual sFID, precision, and recall comparisons are mixed, as expected for competing quality and coverage metrics; nevertheless, the per-sample controller improves the overall quality–diversity trade-off most consistently. For example, at 20 NFE and $\lambda = 1.5$, adapting to each sample reduces FID from 3.23 to 3.03 while increasing IS from 252.97 to 261.95. At the stronger $\lambda = 2.5$, it improves FID and IS at both 20 and 50 NFE. These results show that the increasing time schedule is useful on its own, and that state-dependent feedback provides a further gain beyond this schedule.

Table 4: Per-sample PMC-CFG versus its frozen mean schedule on ImageNet-256 (50K samples). Both variants use the same nominal scale λ and cap Γ; the frozen schedule is estimated from an independent 1,024-sample calibration run. Bold marks the better directional metric within each pair.

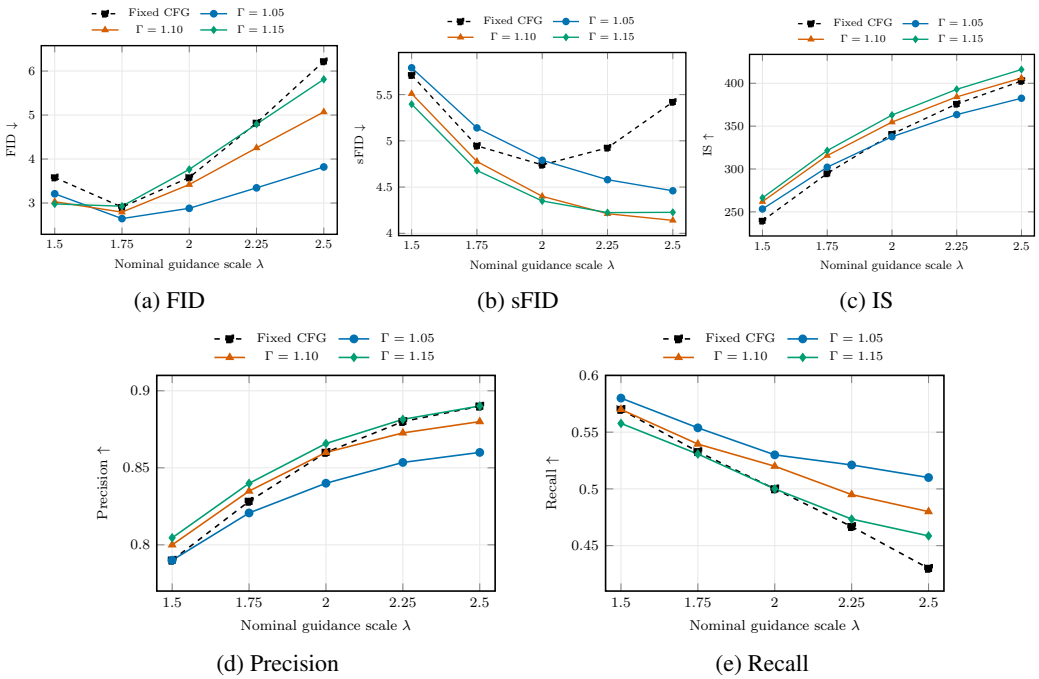
NFE	λ	Controller	FID↓	sFID↓	IS↑	Prec.↑	Rec.↑
20	1.50	PMC-Mean	3.23	5.71	252.97	0.79	**0.58**
		Per-sample PMC	**3.03**	**5.51**	**261.95**	**0.80**	0.57
20	2.00	PMC-Mean	2.91	**4.64**	334.42	**0.84**	**0.53**
		Per-sample PMC	**2.88**	4.79	**337.62**	**0.84**	**0.53**
20	2.50	PMC-Mean	3.86	**4.32**	377.22	**0.86**	0.50
		Per-sample PMC	**3.82**	4.46	**382.53**	**0.86**	**0.51**
50	1.50	PMC-Mean	**2.12**	4.70	273.77	0.81	**0.58**
		Per-sample PMC	**2.12**	**4.60**	**283.15**	**0.82**	**0.58**
50	2.00	PMC-Mean	2.89	**4.13**	354.83	**0.85**	0.54
		Per-sample PMC	**2.85**	4.15	**359.72**	**0.85**	**0.55**
50	2.50	PMC-Mean	4.15	4.10	396.06	**0.87**	**0.52**
		Per-sample PMC	**4.11**	**4.07**	**403.15**	**0.87**	**0.52**

F.2 Sensitivity to the posterior-mean cap on ImageNet-256

We further vary the PMC-CFG cap over $\Gamma \in \{1.05, 1.10, 1.15\}$ and densify the nominal guidance-scale sweep to $\lambda \in \{1.5, 1.75, 2.0, 2.25, 2.5\}$. All runs use the same SiT-XL/2 checkpoint, 20-step Euler solver, 50K-sample evaluation protocol, and official ADM/guided-diffusion Inception statistics described above. Fixed CFG provides the uncapped reference at each nominal scale.

Figure 5 reveals a consistent cap-controlled trade-off between distributional quality and fidelity-oriented metrics. Under strong nominal guidance ($\lambda \geq 2.0$), the tightest cap, $\Gamma = 1.05$, achieves the lowest FID and highest recall among the three PMC-CFG variants, while degrading much more slowly than fixed CFG as λ increases. Relaxing the cap shifts the balance in the opposite direction: $\Gamma = 1.15$ yields the highest IS and precision at nearly every scale, but also produces higher FID and lower recall at large λ. sFID shows a related but less pronounced trend, generally improving with λ and with a looser cap before saturating—and slightly reversing—around $\lambda = 2.25$. Together, these curves show that Γ controls a genuine quality–coverage trade-off rather than uniformly improving every metric.

This trade-off also explains the cap values used in the main experiments. At $\lambda = 1.5$, $\Gamma = 1.10$ improves FID from 3.58 to 3.03 and IS from 239.49 to 261.95 relative to fixed CFG, while preserving recall at 0.57. At stronger guidance, the tighter $\Gamma = 1.05$ setting is more robust: at $\lambda = 2.5$, it reduces FID from 6.22 to 3.82 and increases recall from 0.43 to 0.51. We therefore use $\Gamma = 1.10$ at $\lambda = 1.5$ to retain more guidance benefit without losing coverage, and $\Gamma = 1.05$ at larger scales to suppress the rapidly growing FID and recall degradation.



(a) FID (b) sFID (c) IS

(d) Precision (e) Recall

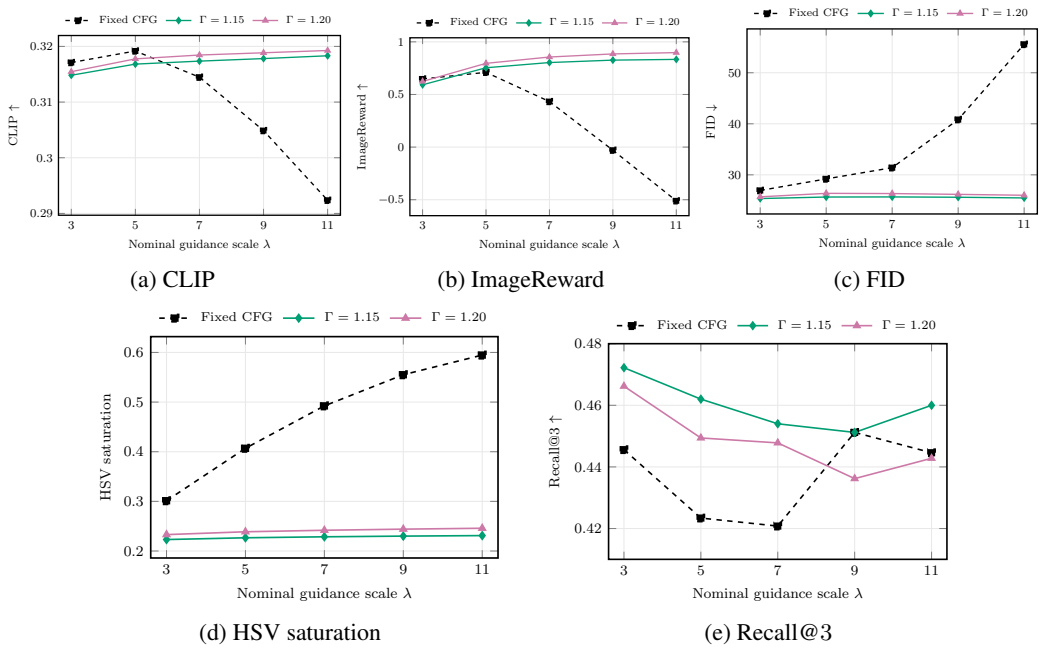
Figure 5: ImageNet-256 sensitivity to the posterior-mean cap at 20 NFE and 50K samples. Each panel plots one metric against the nominal guidance scale λ for fixed CFG and PMC-CFG with three cap values Γ. Lower is better for FID and sFID; higher is better for IS, precision, and recall.

F.3 Sensitivity to the posterior-mean cap on SD3.5

We extend the SD3.5 evaluation to nominal guidance scale $\lambda = 11$ and compare fixed CFG with PMC-CFG at $\Gamma \in \{1.15, 1.20\}$. All settings use the same SD3.5 Medium checkpoint, 10-step solver, 5,000 COCO validation prompts, initial latent sequence, and metric implementations described in Appendix E. Thus the sweep isolates how the cap changes sensitivity to increasingly strong nominal guidance.

Figure 6 shows that posterior-mean capping makes SD3.5 substantially less sensitive to large nominal guidance scales. Fixed CFG begins to deteriorate beyond $\lambda = 5$: by $\lambda = 11$, CLIP falls to 0.2924, ImageReward to -0.5107, and FID rises to 55.57, while mean saturation increases to 0.595. In contrast, both capped variants remain stable through $\lambda = 11$, keeping FID near 25–26 and saturation near 0.23–0.25 while CLIP and ImageReward continue to improve. This widening gap indicates that the cap prevents nominal-scale increases from translating directly into the oversaturation and quality collapse observed with fixed CFG.

The two cap values expose the same fidelity–coverage trade-off observed on ImageNet-256. Relative to $\Gamma = 1.15$, the looser $\Gamma = 1.20$ cap improves CLIP and ImageReward at every tested scale; at $\lambda = 11$, the gains are $0.3183 \to 0.3193$ and $0.8333 \to 0.8978$, respectively. The tighter $\Gamma = 1.15$ cap, however, gives lower FID at every scale and higher Recall@3 at every scale, while also maintaining slightly lower saturation. At $\lambda = 11$, its Recall@3 remains 0.4600, compared with 0.4428 for $\Gamma = 1.20$ and 0.4446 for fixed CFG. Therefore $\Gamma = 1.20$ favors prompt-aligned reward, whereas $\Gamma = 1.15$ provides the more conservative distributional trade-off.



(a) CLIP (b) ImageReward (c) FID

(d) HSV saturation (e) Recall@3

Figure 6: SD3.5 Medium sensitivity to the posterior-mean cap on 5,000 COCO validation prompts at 10 NFE. Each panel plots one metric against nominal guidance scale λ for fixed CFG and PMC-CFG with two cap values Γ. Saturation is diagnostic rather than directional.

F.4 GMM MARGINAL ENDPOINT DENSITIES

Figure 7 complements the endpoint samples and statistics in the main text by showing how guidance changes the endpoint marginal densities. The conditional baseline ($\lambda = 1$, top row) closely follows the target shape, apart from the known finite-step contraction. Standard fixed-scale CFG ($\lambda = 3$, middle row) produces a clear distributional shift: for the singleton condition, the marginal peak moves away from the target and becomes substantially narrower; for the pair and triple conditions, the generated mass is additionally reweighted toward the heavier or geometrically central branch. Consequently, the generated histograms no longer align with either the location, spread, or relative mode mass of the conditional target.

PMC-CFG uses the same nominal scale $\lambda = 3$ but limits the early posterior-mean extrapolation responsible for this displacement. The bottom row therefore places the marginal peaks closer to the target, restores within-mode spread, and reduces branch reweighting. Table 1 quantifies the same effect: relative to fixed CFG, PMC-CFG reduces mean error from 0.04091 to 0.00831 for the singleton, from 0.11781 to 0.02800 for the pair, and from 0.12883 to 0.02980 for the triple. For the multimodal cases, occupancy TV also decreases from 0.15272 to 0.03851 and from 0.35294 to 0.08046, respectively. Thus the cap mitigates both global displacement and guidance-induced concentration rather than merely translating the generated distribution.

F.5 IMAGENET-256 MECHANISM DIAGNOSTICS

Figure 8 exhibits the posterior-mean evolution derived for the GMM in Proposition C.1. Early in sampling, $\|m_t^{\mathrm{u}}\|$ is substantially smaller than $\|m_t^{\mathrm{c}}\|$ and $\|\Delta_t\|$ is large, resembling the centered-GMM initialization $m_0^{\mathrm{u}} = 0$ and $m_0^{\mathrm{c}} = \mu_0$. The unconditional prediction averages over many possible classes, whereas conditioning supplies a directed endpoint estimate. As x_t becomes informative, the two predictions converge and $\|\Delta_t\|$ approaches zero, matching the GMM limit in (39).

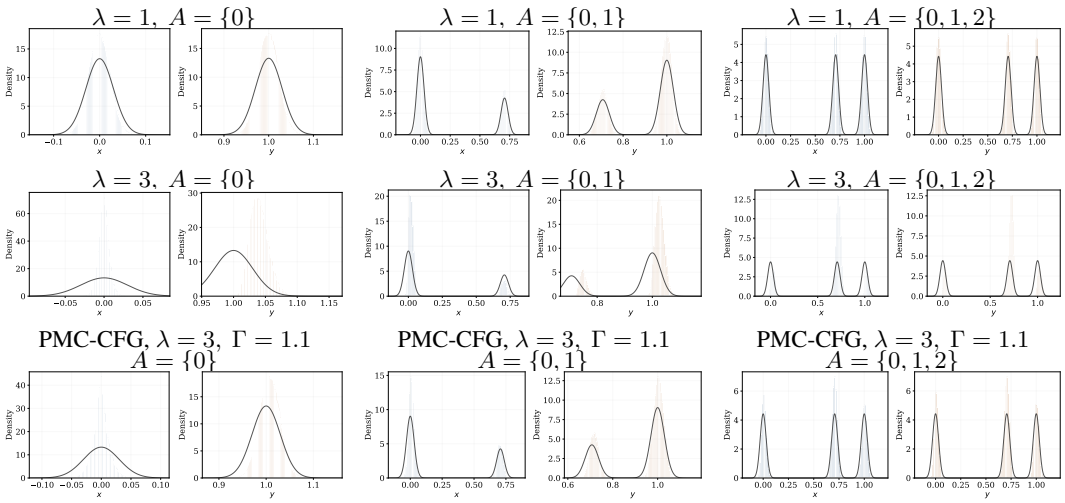


Figure 7: Marginal endpoint densities for the same runs as Figure 3. Rows show the conditional baseline ($\lambda = 1$), fixed CFG ($\lambda = 3$), and PMC-CFG ($\lambda = 3, \Gamma = 1.1$). Histograms are generated marginals and dark curves are the theoretical conditional target. Fixed CFG shifts and concentrates the distribution, whereas PMC-CFG recovers substantially more of its location, spread, and relative mode mass.

This evolution of the posterior-mean gap explains the increasing effective scale in the right panel. Writing $\gamma_t^{\text{eff}} = 1 + \beta_t^*$, the actual endpoint prediction used by PMC-CFG is

$$m_t^{\text{PMC}} = m_t^{\text{c}} + (\gamma_t^{\text{eff}} - 1)\Delta_t, \qquad \gamma_t^{\text{eff}} = 1 + \min\left\{\lambda - 1, \frac{\|m_t^{\text{c}}\|}{\|\Delta_t\|}\left[-\cos\theta_t + \sqrt{\cos^2\theta_t + \Gamma^2 - 1}\right]\right\}, \tag{63}$$

for $\Delta_t \neq 0$, where θ_t is the angle between m_t^{c} and Δ_t; if $\Delta_t = 0$, we set $\gamma_t^{\text{eff}} = \lambda$. A large relative gap activates the cap and keeps the early scale close to conditional guidance. As the gap contracts, the admissible scale increases. Once

$$\|m_t^{\text{c}} + (\lambda - 1)\Delta_t\| \leq \Gamma \|m_t^{\text{c}}\|, \tag{64}$$

nominal CFG is feasible, so the controller reaches and retains the upper limit $\gamma_t^{\text{eff}} = \lambda$.

This produces a coarse-to-fine guidance schedule. During the early, pattern-forming phase, the smaller scale limits posterior-mean overshoot and extra trajectory contraction, preserving variation while the global layout is established. Later, Δ_t is small, so PMC-CFG can use the nominal scale with a small absolute correction $(\lambda - 1)\Delta_t$ to sharpen class-specific boundaries and texture. This mechanism agrees with the improved recall in Table 2 and the matched samples in Figure 9.

The endpoint metrics also exhibit a solver-budget-dependent trend. As NFE increases from 20 to 50, the FID-optimal scale for APG shifts from $\lambda = 2.0$ to $\lambda = 1.5$, and the best reported PMC-CFG configuration likewise changes from $(\lambda, \Gamma) = (2.0, 1.05)$ to $(1.5, 1.10)$. Fixed CFG changes from a tie between these scales to a clear preference for $\lambda = 1.5$, while C2FG favors $\lambda = 1.5$ at both budgets. Moreover, $\lambda = 2.5$ is suboptimal for every method and budget.

A plausible explanation is a change in the dominant source of error. Increasing NFE reduces numerical integration error for a weakly guided flow, but it does not remove guidance-induced mean displacement and local contraction because these are properties of the modified vector field itself. Once discretization error is reduced, increasing the scale can expose more of this structural bias rather than improve integration accuracy. The endpoint metrics and mechanism diagnostics above support this explanation as a plausible mechanism.

F.6 Additional qualitative results

To complement the aggregate metrics, we provide matched qualitative results that isolate how each guidance rule changes the composition and appearance of individual samples. Within each

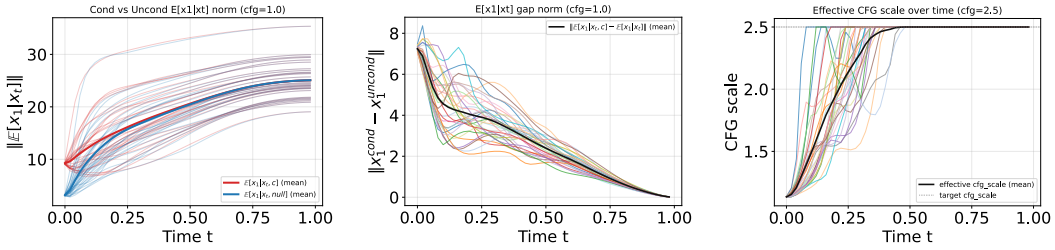


Figure 8: ImageNet-256 mechanism diagnostics. From left to right: the conditional and unconditional posterior-mean norms $\|m_t^{\mathrm{c}}\|$ and $\|m_t^{\mathrm{u}}\|$ evaluated along the $\lambda = 1$ trajectory; their gap $\|\Delta_t\|$ on the same trajectory; and the effective PMC-CFG scale $\gamma_t^{\mathrm{eff}} = 1 + \beta_t^*$ for nominal scale $\lambda = 2.5$. The initially smaller unconditional norm and large gap resemble the centered-GMM geometry; as the gap vanishes, the cap relaxes and the effective scale reaches its nominal upper limit. Thin curves denote individual samples and thick curves their averages.

comparison, no guidance, fixed CFG, and PMC-CFG use the same class or text prompt and the same initial latent seed. Figure 9 reports additional class-conditional results on ImageNet-256, while Figure 10 covers four text-prompt categories on SD3.5 Medium. The latter uses the following prompts:

- **Dog:** *A dog is shown sitting on rocks by the beach, photorealistic*.
- **Landscape:** *A beautiful mountain and river landscape, photorealistic*.
- **Bedroom:** *A bedroom with a large bed sitting under a painting, photorealistic*.
- **People:** *a person on a small boat with another boat in the background, photorealistic*.

Across both ImageNet-256 and SD3.5, the matched samples show that PMC-CFG largely preserves the objects, layout, and overall spatial pattern of the unguided counterpart while producing clearer structures and finer details. Compared with fixed CFG, PMC-CFG also avoids much of the excessive saturation, improving image quality without replacing the composition established by the unguided trajectory.

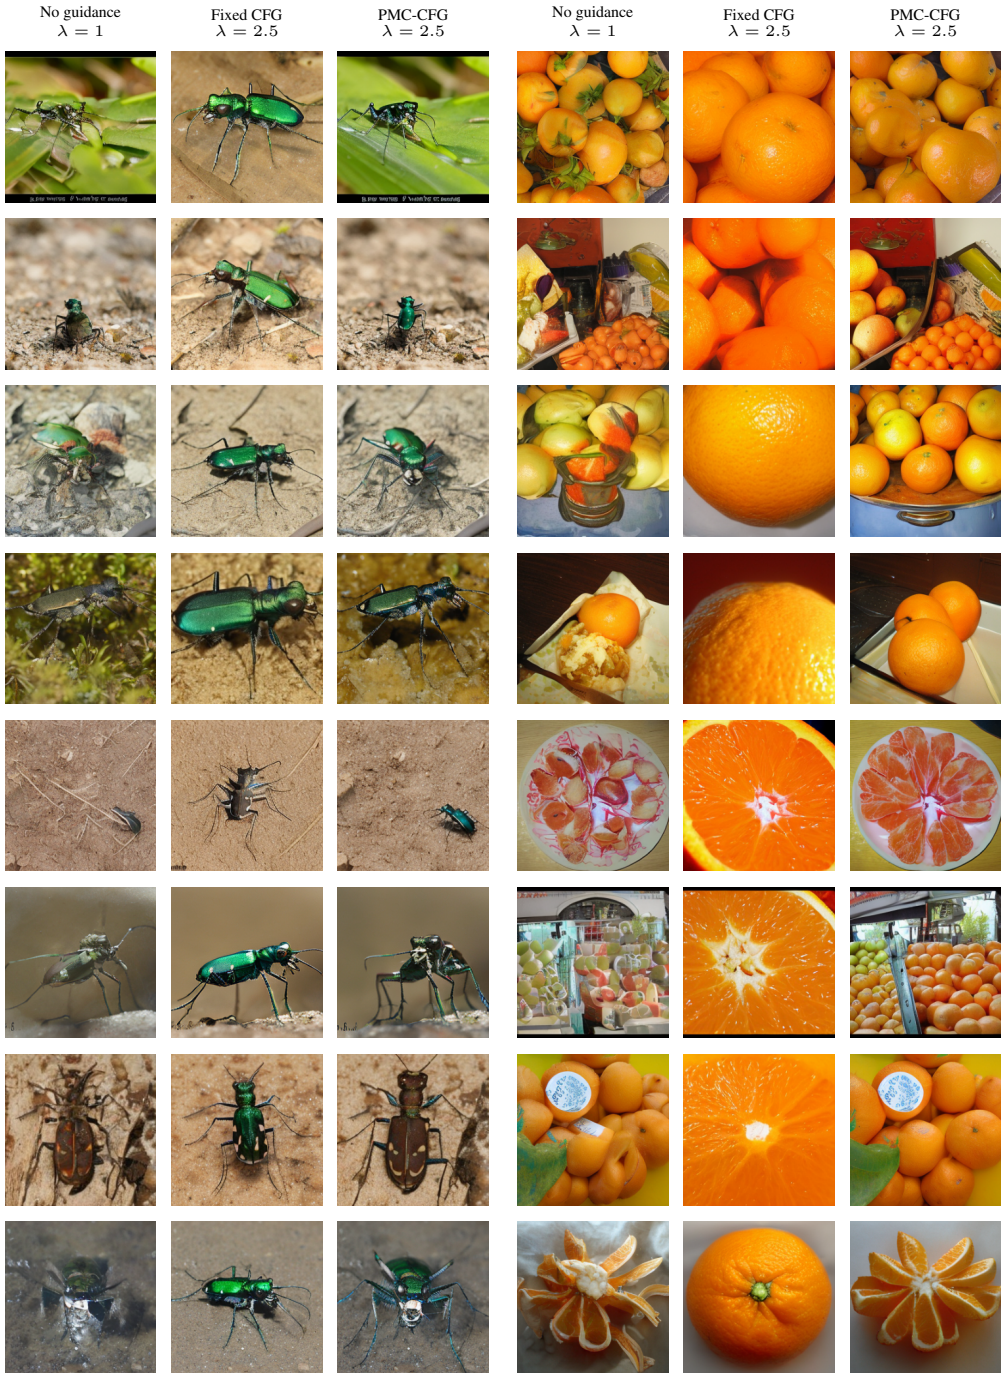


Figure 9: Additional matched qualitative comparisons on ImageNet-256. Each three-column group compares no guidance, fixed CFG, and PMC-CFG; corresponding images within each group use the same class and initial latent seed. Fixed CFG and PMC-CFG both use nominal guidance scale $\lambda = 2.5$.

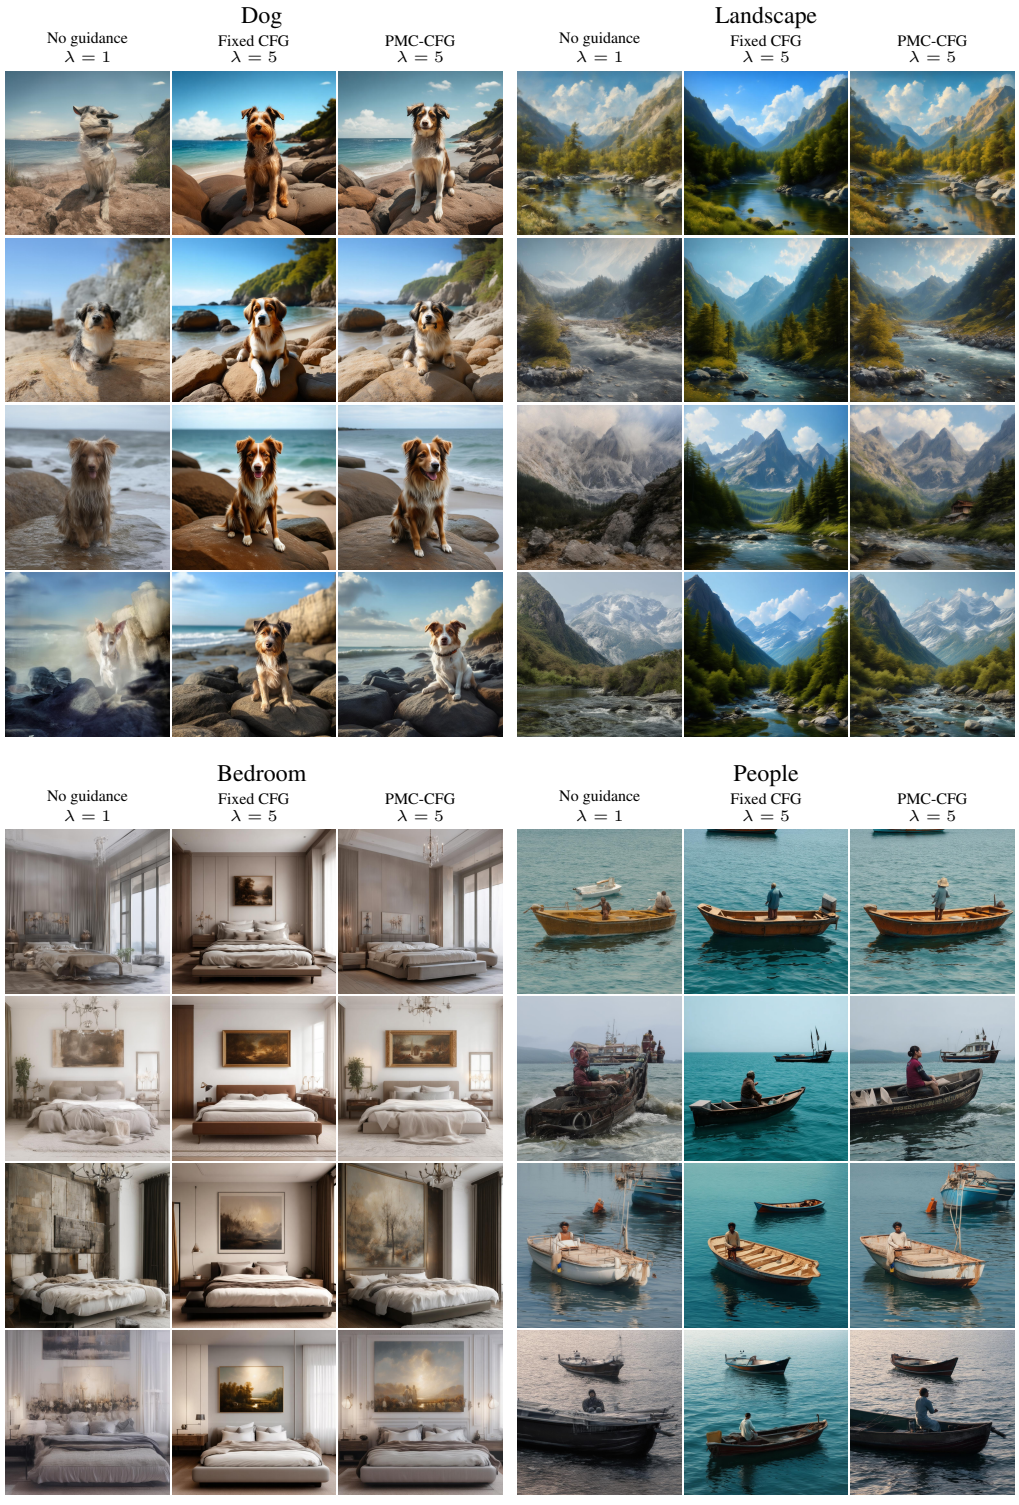


Figure 10: Additional matched qualitative comparisons on SD3.5 Medium across dog, landscape, bedroom, and people prompts. Each four-row category compares no guidance, fixed CFG, and PMC-CFG using matched prompts and initial latent seeds. Fixed CFG and PMC-CFG both use nominal guidance scale $\lambda = 5$.