FITTED CFG: STABILIZATION BANDS AND CONTRACTION MATCHING

Anonymous authors
Paper under double-blind review

ABSTRACT

Classifier-free guidance (CFG) changes the flow that a diffusion sampler integrates, but the DDIM update was designed for the unguided flow. In an ordered, commuting Gaussian model, guidance strengthens contraction along discriminative directions, and a coarse DDIM step can reverse such a coordinate. On this model, a band of guidance coefficients prevents the reversal and controls discriminative residual amplification without prescribing the exact contraction; it contains the sign cap and a parameter range of deterministic CFG++. Terminal matching fixes the factor left free and selects *fitted CFG*, a one-coefficient change at no extra denoiser cost. Fitted CFG and the caps share first-order absolute accuracy; matching additionally bounds the accumulated log-relative contraction error independently of the noise horizon, for fixed guidance and maximum log step. On CIFAR-10, fitted CFG lowers FID in all cells of two nine-cell grids and in a 50k comparison ($27.87 \rightarrow 20.88$); the sign cap matches this stabilization in paired controls. A calibrated finite-scale extension improves FID, KID and saturation on 30k SDXL samples at some cost in text alignment. The framework separates three design criteria: stability, consistency with the prescribed field, and contraction accuracy.

1 INTRODUCTION

Classifier-free guidance (CFG) [8] can produce oversaturated, high-contrast images at large guidance scales [16]. It changes the flow that the sampler integrates, but the DDIM update [17] is unchanged. In a Gaussian model, guidance strengthens contraction along discriminative directions, and a coarse DDIM step can reverse such a coordinate. We ask which coefficients stabilize the guided update and which reproduce the guided contraction.

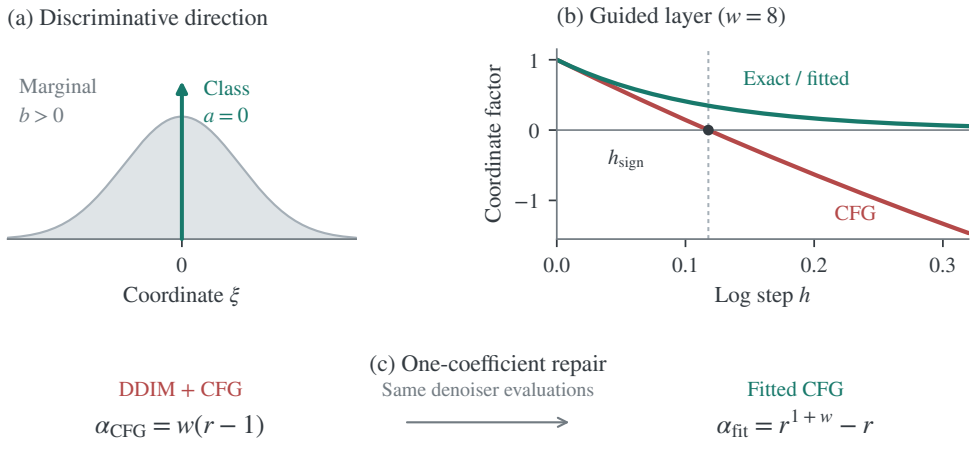

Figure 1: **Guidance–mesh mismatch and contraction matching.** Gaussian coordinate ξ has class variance $a = 0$ and marginal variance $b > 0$. Fitted CFG matches its positive pure-terminal factor r^{1+w}; DDIM+CFG can reverse sign. Here $r = \sigma_{n+1}/\sigma_n$, $w \geq 0$ is the extra guidance weight, $h = -\log r$, and h_{sign} is the sign-reversal threshold.

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.

Following the small-noise analysis of asymptotic-preserving samplers [19], we study positive-noise guided steps on an ordered Gaussian pair. Sign reversal and normalized-residual growth have separate step-size thresholds on this model, and one band of guidance coefficients controls both (Section 3). A coefficient can therefore stabilize the update without reproducing the exact contraction: the band contains the sign cap, which clips the guided factor at zero, and a parameter range of deterministic CFG++ [4].

The band leaves the contraction factor free, and matching fixes it. On the pure terminal mode the guided flow contracts by r^{1+w} per step, where $r = \sigma_{n+1}/\sigma_n$ and $g = 1 + w$ is the guidance scale. Requiring the update to reproduce this factor selects *fitted CFG*, which replaces the coefficient $w(r-1)$ of the denoiser gap by $r^{1+w} - r$ and needs no covariance estimate or extra denoiser evaluation. Fitted CFG and the caps share first-order absolute accuracy, uniformly in the noise floor. Matching adds relative accuracy: the accumulated log-relative contraction error stays bounded independently of the positive-noise horizon, for fixed guidance and maximum log step (Theorem 3), whereas on the pure mode any other fixed positive factor accumulates log-relative distortion linearly over identical steps. A known bound on the marginal variance widens the band and permits a coefficient matched to a finite crossover scale (Section 3.4).

The experiments ask which benefits the band members share and where the specific designs help. On CIFAR-10, fitted CFG lowers FID in every cell of two nine-cell grids and the residual tail in eight of nine; in a 50k comparison FID falls from 27.87 to 20.88. A paired control gives the same finite-step stabilization with the sign cap, whose coefficient is numerically close to fitted CFG at that setting (Section 4.1). Where a crossover scale can be calibrated, a finite-scale coefficient gives a better FID, KID and saturation operating point than CFG on 30k SDXL samples, at some cost in text alignment (Section 4.2). In SDXL we also verify twelve coarse CFG steps at which the local response along a probe direction reverses sign while a converged integration of the same guided field keeps it positive; fitted CFG and the sign cap each repair eleven (Section 4.3).

The analysis thus separates three design criteria for guidance-aware update coefficients: stability, consistency with the prescribed guided field, and contraction accuracy. On the Gaussian model, fitted CFG meets all three using only the guidance weight and the step ratio, and the sign cap meets the first two.

2 Guidance–mesh mismatch

Setup and notation. For clean data X_{data}, observations have the form $X_{\mathrm{data}} + \sigma Z$, $Z \sim \mathcal{N}(0, I)$. Let $D_{\mathrm{c}}(x, \sigma)$ and $D_{\mathrm{u}}(x, \sigma)$ denote conditional and unconditional denoisers. Define the guided denoiser $D_w = (1 + w)D_{\mathrm{c}} - wD_{\mathrm{u}}$ and its normalized residual. CFG samples

$$\frac{\mathrm{d}X}{\mathrm{d}\sigma} = \eta_w(X, \sigma), \qquad \eta_w(x, \sigma) = \frac{x - D_w(x, \sigma)}{\sigma}, \qquad w \geq 0. \tag{1}$$

Here x is a full state, $X(\sigma)$ a continuous sampling trajectory, and x_n a numerical state; x_0 is the noisy initial state. The positive-noise mesh $\sigma_0 > \cdots > \sigma_K > 0$ has K intervals and ends at the noise floor σ_K. For $0 \leq n < K$, $r_n = \sigma_{n+1}/\sigma_n = e^{-h_n}$, where h_n is the increment in $\lambda = -\log\sigma$; the horizon is $\sum_n h_n = \log(\sigma_0/\sigma_K)$. Single-step formulas omit n. The normalized residual is the noise prediction in these additive-noise coordinates, and DDIM's stochasticity parameter is zero throughout.

The mode that DDIM fits. For a scalar mode $\mathrm{d}\xi/\mathrm{d}\sigma = \mu\xi/\sigma$ with constant exponent μ, write $\xi(\sigma)$ for its continuous coordinate and ξ_n for its numerical value. Starting from the same ξ_n, exact integration (superscript ex) and a frozen left-endpoint field give

$$\xi_{n+1}^{\mathrm{ex}} = r^{\mu}\xi_n, \qquad \xi_{n+1}^{\mathrm{DDIM}} = \big[1 - (1 - r)\mu\big]\xi_n. \tag{2}$$

DDIM is exact only at $\mu = 0$ and $\mu = 1$. Its factor is linear in μ, so for $\mu > 1$ it extrapolates past both exact modes and can cross zero, while r^{μ} stays positive; larger μ means stronger contraction. At fixed μ and step ratio r, lowering the noise floor leaves both factors and their mismatch unchanged.

Guided exponents on a Gaussian pair. For fixed condition c, take

$$X_{\mathrm{data}} \mid c \sim \mathcal{N}(0, \Sigma_{\mathrm{c}}), \qquad X_{\mathrm{data}} \sim \mathcal{N}(0, \Sigma_{\mathrm{u}}), \qquad 0 \preceq \Sigma_{\mathrm{c}} \preceq \Sigma_{\mathrm{u}}, \quad \Sigma_{\mathrm{c}}\Sigma_{\mathrm{u}} = \Sigma_{\mathrm{u}}\Sigma_{\mathrm{c}}.$$

With the additive Gaussian noise above, its exact posterior-mean denoisers are

$$D_j(x,\sigma) = \Sigma_j(\Sigma_j + \sigma^2 I)^{-1}x, \qquad j \in \{\mathrm{c},\mathrm{u}\}.$$

Commutativity lets us work one coordinate at a time: choose a common unit eigenvector v with $\Sigma_\mathrm{c} v = av$, $\Sigma_\mathrm{u} v = bv$, $0 \le a \le b$, and set $\xi = v^\top X$. The denoiser components are $D_\mathrm{c} = a\xi/(a+\sigma^2)$ and $D_\mathrm{u} = b\xi/(b+\sigma^2)$. Thus the residual components are $\eta_\mathrm{c} = A\xi/\sigma$ and $\eta_\mathrm{u} = B\xi/\sigma$, and the guided flow gives, for $\xi \neq 0$,

$$A = \frac{\sigma^2}{a+\sigma^2}, \quad B = \frac{\sigma^2}{b+\sigma^2}, \qquad \frac{\mathrm{d}\log|\xi|}{\mathrm{d}\log\sigma} = \mu_w = A + w(A-B). \tag{3}$$

The conditional exponent A gains the guidance term $w(A-B)$; in step factors these noise-dependent weights are evaluated at σ_n. As $\sigma \to 0$, shared normal directions ($a = b = 0$) have exponent 1, whereas tangential directions ($0 < a \le b$) have exponent tending to 0. We focus on *discriminative normal* directions ($a = 0 < b$): conditioning fixes the coordinate while the marginal model retains variation. For $w > 0$, as noise falls through $\sigma^2 \approx b$, $\mu_w = 1 + wb/(b+\sigma^2)$ rises from near 1 toward $1 + w$.

The terminal test mode. The *terminal layer* is the low-noise regime $\sigma^2 \ll b$, where the marginal denoiser nearly returns its input: $D_\mathrm{c} = 0$ and $D_\mathrm{u}/\xi = b/(b+\sigma^2) \to 1$. Taking these coefficient limits gives

$$\frac{\mathrm{d}\xi}{\mathrm{d}\sigma} = (1+w)\frac{\xi}{\sigma}, \qquad \xi(\sigma_{n+1}) = r^{1+w}\xi(\sigma_n).$$

This *pure terminal mode* describes contraction between positive noise levels; the final projection to zero noise is a separate operation, and at finite noise the exponent changes through the crossover.

Write a family of updates using the same left-endpoint denoisers:

$$x_{n+1} = D_\mathrm{c} + r(x_n - D_\mathrm{c}) + \alpha(D_\mathrm{u} - D_\mathrm{c}) = D_\mathrm{c} + r(x_n - D_\mathrm{u}) + s(D_\mathrm{u} - D_\mathrm{c}), \tag{4}$$

where $s = r + \alpha$ is the factor on the pure terminal mode. The update applies to arbitrary denoisers; s is its design parameter, named by its action on the test mode. Ordinary CFG uses $\alpha_\mathrm{CFG} = w(r-1)$, hence $s_\mathrm{CFG} = (1+w)r - w$.

Theorem 1 (Two distinct guided-layer thresholds)**.** *On a nonzero pure discriminative trajectory $\eta_w = (1+w)\xi/\sigma$, $w > 0$, a DDIM+CFG step has factor $G_w(h) = (1+w)e^{-h} - w$. It reverses the coordinate if and only if*

$$h > h_\mathrm{sign} := \log(1 + 1/w), \tag{5}$$

and amplifies the normalized residual magnitude if and only if

$$h > h_\mathrm{amp} := \log(1 + 2/w). \tag{6}$$

For fixed $h > h_\mathrm{amp}$, K steps amplify this residual by $[we^h - (1+w)]^K$, diverging as $\sigma_K = \sigma_0 e^{-Kh} \to 0$. The exact residual contracts by e^{-wh} per step.

The residual multiplier is $|G_w|/r$: division by the new noise level distinguishes residual amplification from coordinate reversal. A step in $(h_\mathrm{sign}, h_\mathrm{amp}]$ reverses the coordinate while its residual magnitude remains nonincreasing. Appendix B.1 gives the absolute-stability threshold and finite-crossover extension. The ordered Gaussian model supplies the test modes; near a class point of a separated finite-label mixture, the terminal exponent can instead be one (Appendix B.4).

3 Stabilization and contraction matching

3.1 A shared stabilization band

A coefficient can stabilize the update without specifying its exact contraction; on the Gaussian model the condition is a band of admissible factors s.

Theorem 2 (Shared Gaussian stabilization)**.** *For the ordered commuting Gaussian pair, the coordinate factor of (4) is*

$$G_s = (1 - A) + rB + s(A - B). \tag{7}$$

If $0 \le s_n \le r_n$ at every positive-noise step, all coordinate factors lie in $(0, 1]$ for finite variances. On a discriminative coordinate $a = 0 < b$, with fixed diagnostic guidance $w \ge 0$, nonzero initial residual and $\mu_n = 1 + w(1 - B_n)$, every prefix satisfies

$$\frac{|\eta_w(\xi_k, \sigma_k)|}{|\eta_w(\xi_0, \sigma_0)|} \le \frac{\mu_k}{\mu_0} \le 1 + w. \tag{8}$$

The wider band $-r_n \le s_n \le r_n$ retains this residual bound while allowing coordinate reversal.

The weights $1 - A, B, A - B$ are nonnegative and sum to one. Thus $0 \le s \le r$ gives a convex combination of $1, r, s$. For $a = 0$, $G_s/r = B + (s/r)(1 - B)$; multiplying residual ratios cancels the successive μ ratios. This proves the common bound (Appendix B.3). Individual crossover steps can still increase the residual through μ_{n+1}/μ_n.

Clipping the CFG coefficient gives two simple members: the sign cap lies in the sign band and the residual cap in the wider band,

$$\alpha_{\text{sign}} = \max\{w(r - 1), -r\}, \qquad \alpha_{\text{res}} = \max\{w(r - 1), -2r\}. \tag{9}$$

Under refinement the caps become inactive and both updates reduce to CFG.

3.2 Selecting fitted CFG

For prescribed w, exact terminal contraction requires $s = r^{1+w}$. Substituting this into (4) gives

$$\alpha_{\text{fit}} = r^{1+w} - r. \tag{10}$$

This is *fitted CFG*. Since $0 < r^{1+w} \le r$, it lies in the sign band at every positive step, and it uses only the guidance weight and the noise ratio. It also reproduces the limiting factors r and 1 of shared-normal and tangential directions, and for fixed w, $\alpha_{\text{fit}} = w(r - 1) + O(h^2)$: fitted CFG is first-order consistent with the original guided field and reduces to DDIM at $w = 0$.

Table 1: **Coefficient choices in ordered Gaussian models,** $w > 0$**.** All five are first-order consistent with the prescribed field. Sign band: $0 \le s \le r$; residual band: $-r \le s \le r$, permitting reversal. The scale band (15) assumes marginal variances $b \le \hat{b}$. Matching identifies the discriminative contraction targeted exactly.

Choice	Pure-mode factor s	Stability	Matching
CFG	$(1 + w)r - w$	sign: $h \le h_{\text{sign}}$ residual: $h \le h_{\text{amp}}$	none
Sign cap	$\max\{(1 + w)r - w,\ 0\}$	sign band, all r	none
Residual cap	$\max\{(1 + w)r - w,\ -r\}$	residual band, all r	none
Fitted CFG	r^{1+w}	sign band, all r	pure terminal mode
Crossover fit	$r + \alpha_{\text{cross}}$	scale band if $b \le \hat{b}$	coordinate $(0, \hat{b})$

Matching is also what composition requires. Splitting one pure-layer step into two should not change the combined factor, $s(r_1 r_2, w) = s(r_1, w)s(r_2, w)$; with continuity and the correct first-order generator, this selects $s = r^{1+w}$ uniquely within the scalar update family (Proposition 3). Any other fixed positive factor accumulates log-relative distortion $K \log(s/r^{1+w})$ over K identical steps, which ties terminal matching to accumulated accuracy.

Existing update families and effective guidance. Deterministic, unthresholded CFG++ [4] writes

$$\Psi_\beta(x) = D_{\text{u}} + \beta(D_{\text{c}} - D_{\text{u}}) + r(x - D_{\text{u}}), \qquad \beta = 1 - s. \tag{11}$$

In this family the sign band is $\beta \in [1 - r, 1]$, and fitted CFG selects $\beta_{\text{fit}} = 1 - r^{1+w}$; equivalently, it is DDIM+CFG at the effective guidance

$$g_{\text{eff}} = \frac{1 - r^{1+w}}{1 - r}, \qquad w_{\text{eff}} = g_{\text{eff}} - 1 \in [0, w]. \tag{12}$$

The identity holds for arbitrary denoisers. On a uniform log mesh r and g_{eff} are constant, so the positive-noise trajectory is exactly DDIM+CFG at g_{eff}, and complete trajectories agree when the closures match; under refinement $g_{\text{eff}} \to 1 + w$. At $r = 0$ the updates differ: fitted CFG returns D_{c} and standard CFG returns D_w.

3.3 Relative accuracy of matching

Band members can share stability and absolute accuracy yet contract by different amounts. The next result controls the ratio of accumulated coordinate factors across finite Gaussian crossovers.

Theorem 3 (Horizon-independent relative accuracy)**.** *For exact centered Gaussian denoisers with commuting covariances $0 \preceq \Sigma_c \preceq \Sigma_u$, fix $w \geq 0$ and a decreasing positive-noise mesh with $\max_n h_n \leq H \leq 1$. Let $T_{\mathrm{fit}}, T_{\mathrm{ex}}$ be its fitted and exact guided-flow endpoint maps, with positive factors $t_{\mathrm{fit},i}, t_{\mathrm{ex},i}$ in the common eigenbasis. A finite $C(w)$ satisfies*

$$\max_i \left| \log \frac{t_{\mathrm{fit},i}}{t_{\mathrm{ex},i}} \right| \leq C(w)H, \tag{13}$$

and consequently

$$\|T_{\mathrm{fit}} - T_{\mathrm{ex}}\|_{\mathrm{op}} \leq e^{C(w)H} - 1 = O_w(H). \tag{14}$$

The constant is independent of dimension, covariance eigenvalues, step count, total log-noise horizon and terminal floor.

Here t_i is an accumulated factor and the norm is the Euclidean operator norm. The proof localizes log-factor error to the crossover weights $A(1-A), B(1-B)$, each with log-noise integral at most $1/2$, extending the unguided argument of [19], Appendix E.1. The two caps also have absolute $O_w(H)$ accuracy, while their log-relative distortion can grow without bound (Appendix B.3).

The bound depends on guidance and step size: on $a = 0 < b$ with uniform step h, $0 \leq \log(t_{\mathrm{fit}}/t_{\mathrm{ex}}) \leq (1+h)w(w+2)h/4$, so large guidance on coarse steps can leave substantial relative distortion. Keeping H fixed as the floor vanishes takes a growing number of steps; asymptotic preservation in the sense of [19] is a property of the complete configuration, including its clock, mesh and terminal policy.

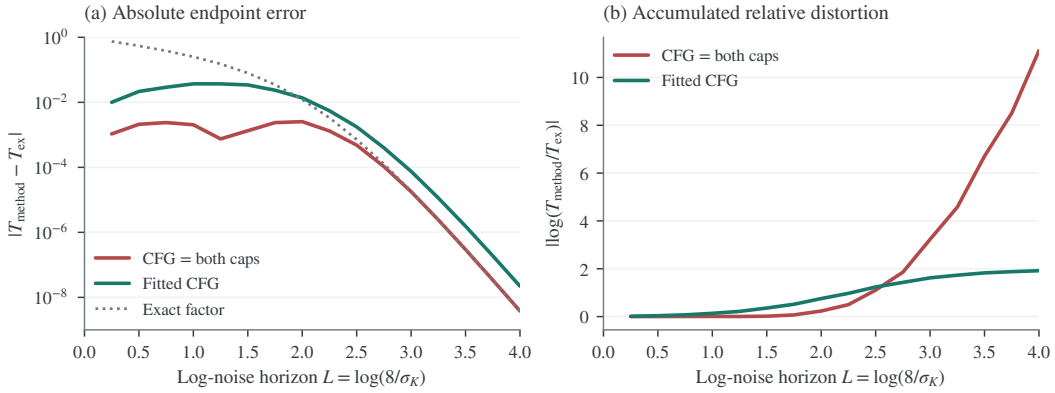

Figure 2: **Shared absolute accuracy and distinct contraction accuracy.** Finite Gaussian crossover $a = 0, b = 1, \sigma_0 = 8, w = 8, h = 0.1$. Both caps are inactive and equal CFG. **Left:** absolute errors and exact factor (dotted). **Right:** accumulated log-relative error; fitted CFG stays bounded, although its absolute error is larger here. Products are evaluated in log space. Appendix B.5 gives the nonuniform-mesh and longer-horizon sweep.

3.4 Scale-informed stabilization and crossover fitting

A known marginal-variance upper bound permits a wider coefficient band. If $0 < b_{\max} < \infty$ and $b \leq b_{\max}$, put $q_{\max}(\sigma) = b_{\max}/(b_{\max} + \sigma^2)$. Since $A - B \leq q_{\max}$ and $(1-A) + rA \geq r$, the condition

$$-r/q_{\max} \leq \alpha \leq 0 \tag{15}$$

gives $0 \leq G_s \leq 1$ and the discriminative residual bound (8). The spectrum-free band is its $b_{\max} \to \infty$ limit (Proposition 4).

To fit a particular coordinate $(a, b) = (0, \hat{b})$ with $\hat{b} > 0$, integrate its exponent $1 + w\hat{b}/(\hat{b} + \sigma^2)$ over the interval:

$$\Phi_n = r_n^{1+w} \left(\frac{\hat{b} + \sigma_n^2}{\hat{b} + \sigma_{n+1}^2} \right)^{w/2}. \tag{16}$$

Writing $q_n = \hat{b}/(\hat{b} + \sigma_n^2)$, its numerical factor is $r_n + \alpha_n q_n$. Exact matching selects

$$\alpha_{\mathrm{cross},n} = \frac{\Phi_n - r_n}{q_n}. \tag{17}$$

Because $0 < \Phi_n \leq r_n$, this coefficient satisfies (15) when $\hat{b}$ bounds the marginal variance. Its pure test factor $r_n + \alpha_n$ can be negative while the specified finite-scale factor remains positive. As $\hat{b} \to \infty$ it reduces to fitted CFG, and both coefficients are first-order consistent with the prescribed field.

The Gaussian scale-band guarantee requires $b \leq \hat{b}$. In learned generation, $\hat{b}$ is a calibrated parameter. Our SDXL study freezes $\hat{b} = 10$ on a separate 512-prompt block (Appendix D.7), with conditional closure and the same number of denoiser evaluations. Theorem 3 concerns fitted CFG; crossover fitting is exact on its specified coordinate.

4 Experiments

Section 4.1 asks whether finite-step stabilization improves generation and whether the benefit is shared within the band; Section 4.2 tests a finite-scale matching target on SDXL [14]; Section 4.3 looks for the discrete response defect in a learned model. CFG, fitted CFG and the caps use identical noise nodes, denoisers and network evaluations, and interval controls match the number of updates. Generation uses $N = K + 1$ updates, K positive-noise intervals and one closure; guidance-off steps and the conditional closure could skip the unconditional evaluation in deployment (Appendix A). Refined integration serves only as a higher-cost reference.

4.1 Shared finite-step stabilization on CIFAR-10

Metrics. CIFAR-10 uses conditional EDM and its unconditional counterpart [10], balanced labels, and matched seeds. For each trajectory, $\mathrm{amp} = \max_{0 \leq n \leq K} \|\eta_w(x_n, \sigma_n)\|_2 / \|\eta_w(x_0, \sigma_0)\|_2$ uses the prescribed $w = g - 1$, including the last positive node but excluding closure. Amp p95 is the 95th percentile of these maxima across trajectories. Clipping is the fraction of output RGB components with magnitude above one before clamping; latent-model saturation counts decoded RGB components below 0.01 or above 0.99 in $[0, 1]$. Both fractions use all image components; their p95 is across images. Distribution quality (FID/KID [1, 7]) and conditioning (classifier accuracy or CLIP) are measured separately. Appendix A specifies implementation details, including the amp denominator safeguards.

Replicated grid and 50k confirmation. In each of two 5k-per-cell blocks, fitted CFG lowers point FID in all nine cells and amp p95 in eight, with no increase in clipping p95 (Figure 3). KID and target accuracy depend on the step count. At $N \in \{16, 32\}$, fitted CFG improves KID in five of six primary cells and keeps target accuracy within 0.04 percentage points of CFG or above it in all six; at the coarsest mesh, $N = 8$, KID favors CFG in all three cells and two cells lose more than two points of accuracy (Appendix D). At $g = 9, N = 8$, an independent 50k comparison confirms the lower FID and residual tails with nearly unchanged target accuracy (Table 2).

Finite-step stabilization at matched closure. Fitted CFG also changes the final projection, returning D_{c} instead of D_w at $r = 0$. To separate the two effects, we cross CFG and fitted finite steps with guided and conditional closures at $g = 9, N \in \{8, 16\}$: two independent balanced-class 5k blocks per setting, 20k paired seed cases in all, with both closures sharing each trajectory's preterminal state and denoiser evaluations. The finite steps carry the stabilization.

Holding either closure fixed, switching from CFG to fitted finite steps reduces final-output clipping p95 from about 0.56 to zero at $N = 8$, and from about 0.65 to at most 0.010 at $N = 16$. All eight

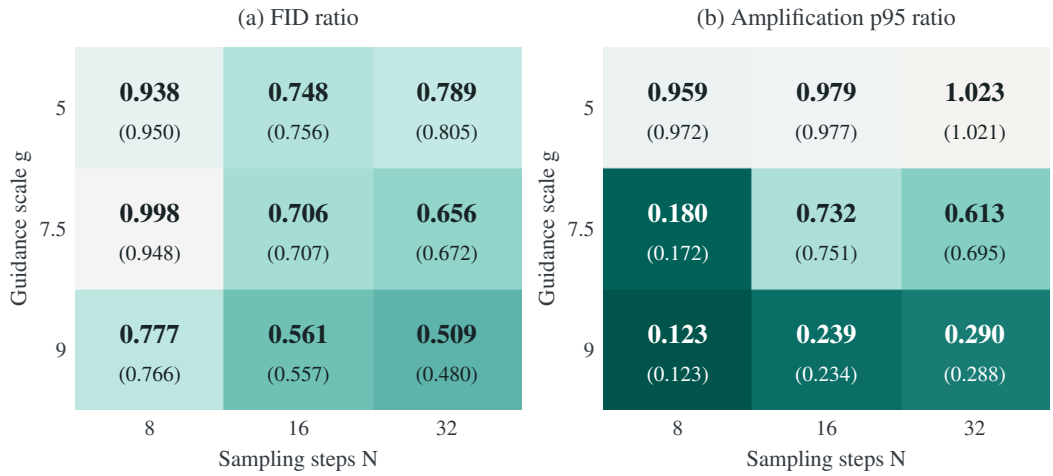

Figure 3: **Fitted CFG across two CIFAR-10 grids.** Left: FID; right: amp p95. Entries are fitted/CFG ratios; values below one favor fitted CFG. Each cell shows the primary 5k block in bold and an independent 5k replicate in parentheses. Color represents the primary ratio, with neutral at one. The amp exception is $g = 5, N = 32$.

Table 2: **CIFAR-10 at fixed update count: 50k images per method,** $g = 9, N = 8$. Accuracy is classifier target accuracy; clipping is measured before output clamping. The exact configuration is in Appendix A.

Method	FID ↓	KID ↓	Acc. (%) ↑	amp p95	Clip p95
CFG	27.87	0.0108	95.12	13.34	0.569
Fitted CFG	20.88	0.0177	94.83	1.65	0.000
Terminal interval	22.08	0.0221	88.95	1.63	0.000

paired clipping-p95 effect intervals are negative. Preterminal amp p95 falls from 13.24–13.31 to 1.64–1.65 at $N = 8$, and from 6.49–7.64 to 1.66–1.67 at $N = 16$.

Changing only CFG's final closure alters point FID by less than 0.04 and leaves preterminal amp identical by construction. On fitted trajectories, conditional closure D_c gives point FID 0.19–0.58 higher than D_w, while slightly reducing clipping at $N = 16$. Under both closures KID favors CFG at $N = 8$ and fitted CFG at $N = 16$. Appendix D.8 gives paired effects, interactions, confidence scores, and shared-subset KID diagnostics.

Coefficient controls. In fresh paired 5k controls at $g = 9, N = 8$ with conditional closure, the sign cap matches fitted CFG: FID 24.735 against 24.733, with zero clipping p95 for both. On this mesh the two coefficients differ by at most $r_n^9 \leq 5.9 \times 10^{-4}$, so the control shows the stabilization shared within the band. Constant CFG at a separately calibrated $g = 2$ reaches comparable FID with higher target accuracy; it samples a weaker guided field, whereas the fitted and capped rules are first-order consistent with the prescribed $g = 9$ field (Appendix D.5). Guidance-off interval controls are in Appendix D.4.

4.2 Finite-scale fitting on SDXL

Section 3.4 replaces the terminal target by a finite crossover scale. We freeze $\hat{b} = 10$ on a separate calibration block and evaluate 30k paired SDXL samples at $g = 12, N = 12$ with the same number of network evaluations. The crossover coefficient gives lower FID and KID than CFG and the tuned interval (Table 3), with the same ordering against CFG in all three prespecified 10k subsets (Appendix D.7), and it lowers saturation p95 by 92.6%. Text alignment is lower: mean CLIP moves from 0.31968 to 0.31681, and the GenEval2 arithmetic-mean score also decreases [9]. Terminal

fitting at this cell removes nearly all saturation but has larger distribution and conditioning costs (Appendix D.2).

Table 3: **SDXL generation: 30k images per method**, $g = 12, N = 12$. Matched prompts, seeds and denoiser-evaluation counts; the crossover scale $\hat{b} = 10$ is frozen before evaluation. KID averages 100 fixed size-1,000 subsets. Saturation is the p95 of per-image decoded RGB fractions.

Method	FID ↓	KID ↓	CLIP ↑	Sat. p95
CFG	19.489	0.006833	0.31968	0.03583
Fitted CFG	23.178	0.008428	0.31105	0.00007
Tuned interval	20.144	0.007488	0.31544	0.01873
Crossover, $\hat{b} = 10$	19.006	0.006682	0.31681	0.00266

The 30k study uses guided closure for CFG and conditional closure for the crossover coefficient. A prespecified paired 5k study separates the finite-step effect, with FP32 updates and a fixed 10k real-reference bank: against conditional-closure CFG, the crossover coefficient lowers FID from 27.815 to 27.179 and KID from 0.007117 to 0.006719, with lower CLIP, and the advantage persists under a common preprocessing pipeline. Absolute values are comparable only within each study; Table 18 adds both CFG closures and fixed-guidance switching controls. Finite-scale fitting is thus a distinct operating point for generation. Figure 4 shows a matched-seed SD1.5 [15] example of the contrast change under fitted CFG.



Figure 4: **Matched-seed SD1.5 samples**, $g = 12, N = 12$. Prompt: *a group of people riding skis on a snowy surface*. The example was selected by ranking the first 64 prompts by CFG-minus-fitted saturation and illustrates the contrast reduction (protocol in Appendix D.3).

4.3 Local responses under the learned guided field

On the Gaussian model the band removes coordinate reversal. To look for the analogous defect in a learned model, we compare coarse SDXL updates with a refined integration of the same guided field, conditioning and interval. At input x, unit direction v and perturbation $\delta = \epsilon \max\{\|x\|_2, 1\}$, define

$$Q_\delta(\Psi) = \frac{\langle v, \Psi(x + \delta v) - \Psi(x - \delta v)\rangle}{2\delta}. \tag{18}$$

A verified reversal pairs negative coarse-CFG response with positive converged-reference response across precision, perturbation and integrator checks, distinguishing it from continuous rotation.

Across two 16-prompt blocks at $g = 12, N = 12$, screening 352 intervals selects 96 reference cases. Twelve reversals pass all checks. Fitted CFG repairs eleven at the original cost; eight-substep CFG gives positive FP32 responses in eleven at eight times the cost, with a different exception. All twelve occur at early finite-noise steps 0, 1 or 3. The sign cap also repairs eleven and the residual cap seven (Table 4). The fitted/sign-cap parity and the residual cap's lower count are qualitatively consistent with the Gaussian band distinction in Theorem 2; here the counts measure a network's local directional response. The crossover coefficient repairs none of these events. Its Gaussian stability condition requires a marginal-variance upper bound; the SDXL scale is calibrated (Section 3.4). The complete protocol and fixed first verified example appear in Appendix C.1 and Figure 5.

For the 64 selected cases whose references converged and passed the cross-checks, a post-hoc replay compares full response vectors $q_\delta(\Psi) = [\Psi(x + \delta v) - \Psi(x - \delta v)]/(2\delta)$. Fitted CFG lowers the vector error against CFG in 63/64 cases, with median paired difference -0.609 and block-stratified prompt-bootstrap 95% interval $[-1.336, -0.357]$. At the same evaluation cost, the sign cap reaches nearly the same errors and the residual cap somewhat larger ones, all well below CFG (Table 4), so this local benefit is shared across coefficient choices.

Table 4: **Selected-case SDXL response accuracy and sign repairs.** Error columns use 64 selected valid cases at $\epsilon = 0.003$ in FP32; repairs use the original twelve verified events. Vector error is $\|q_\delta(\Psi) - q_\delta(\Psi_{\mathrm{ref}})\|_2$. Endpoint error uses the maximum of the two separately normalized endpoint errors (Appendix C.2). Repairs require $Q_\delta \geq 0.02$ at three scales in FP32 and FP16.

Method	Repairs	Vector median	Vector p95	Endpoint median
CFG	0/12	1.4810	4.4084	0.03582
Fitted CFG	11/12	0.6359	1.4452	0.03161
Sign cap	11/12	0.6357	1.4451	0.03138
Residual cap	7/12	0.7467	2.0308	0.02747

An independent follow-up fixes sixteen unscreened prompts, isotropic directions and Heun-prefix states on $\sigma_5 \to \sigma_6$. Fitted CFG lowers the response error against CFG in all sixteen cases (fifteen resolved by the numerical-reference guard). The sign cap and the crossover coefficient give nearly the same medians, and all three raise the error of the unperturbed center state, so response and state accuracy differ at this step (Appendix C.4).

5 Related work

CFG++ supplies the deterministic, unthresholded update family [4]; our analysis gives its Gaussian stability band and matching-based parameter choice. APG modifies the guidance vector [16], interval guidance restricts its noise range [12], and CFG-Zero*/C^{2}FG vary its weight [5, 6]. Guidance also motivates solver design: DPM-Solver++ uses data prediction and multistep integration [13]; RES studies exponential order conditions [18]; DPM-Solver-v3 uses model statistics [20]; ERK-Guid uses embedded error estimates [11]. Analyses of continuous CFG examine the distribution it transports [2, 3]; we take the prescribed guided field as given and study its discretization.

6 Discussion and limitations

The framework separates stabilization, consistency with a prescribed field, and contraction accuracy. For ordered Gaussian denoisers, the nonnegative band controls coordinate reversal and discriminative residual growth. The sign cap and fitted CFG lie in this band, which corresponds to $\beta \in [1-r, 1]$ in deterministic CFG++. Terminal matching specifies the remaining contraction and gives fitted CFG a horizon-independent log-relative-accuracy guarantee without a model-specific scale. In the paired CIFAR control, sign cap and fitted CFG are numerically close and share measured stabilization benefits. Finite-scale fitting extends matching to a calibrated crossover; in the reported SDXL comparisons it improves FID/KID with lower text alignment.

Finite-label mixtures can have a different terminal exponent, and learned response accuracy and image quality require separate measurements; our experiments identify useful operating points with conditioning trade-offs. Asymptotic preservation, which motivates the small-noise perspective, is a configuration-level property that also depends on the clock, the mesh, the terminal rule and floor-independent cost. Within these limits the three criteria give a practical design rule for guidance-aware coefficients: keep the coefficient in the band for stability, and match the terminal contraction when no model scale is available.

Reproducibility Statement

Proofs appear in Appendix B. Implementation, metrics and experimental settings are specified in Appendices A, C–E. The anonymous code supplement includes coefficient and sampler implementations, CPU tests, and frozen scalar records, with commands to reconstruct Tables 2–4 and Figures 2–3. `README.md` and `STUDY_INDEX.md` describe the entry points and dependencies. Model checkpoints and datasets are obtained separately from their public sources.

AI Use Statement

The authors conceived the research, developed the theoretical ideas, designed the experiments, and led the analysis. In this work, we used generative AI tools (ChatGPT/Codex and Claude) to assist with refining theoretical formulations and hypotheses, developing and checking mathematical arguments and proofs, reviewing experimental protocols, implementing methods and data-processing code, interpreting results, and translation. We did not use LLM-generated examples or annotations as experimental data. Thematic coding of interview or survey data is not applicable to this study. Additionally, these tools supported literature retrieval, drafting and polishing sections, reference formatting, and figure and table preparation. The authors reviewed AI-assisted work through mathematical checks, code tests, recorded experimental outputs, independent numerical references, and provenance checks. We take responsibility for the final text, claims, proofs, results, code, and artifacts.

References

[1] Mikołaj Bińkowski, Danica J. Sutherland, Michael Arbel, and Arthur Gretton. Demystifying MMD GANs. In *International Conference on Learning Representations (ICLR)*. OpenReview.net, 2018.

[2] Arwen Bradley and Preetum Nakkiran. Classifier-free guidance is a predictor-corrector. *Transactions on Machine Learning Research (TMLR)*, 2025. arXiv:2408.09000.

[3] Muthu Chidambaram, Khashayar Gatmiry, Sitan Chen, Holden Lee, and Jianfeng Lu. What does guidance do? A fine-grained analysis in a simple setting. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024. arXiv:2409.13074.

[4] 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 (ICLR)*, 2025. arXiv:2406.08070.

[5] Weichen Fan, Amber Yijia Zheng, Raymond A. Yeh, and Ziwei Liu. CFG-Zero*: Improved classifier-free guidance for flow matching models. *arXiv preprint arXiv:2503.18886*, 2025.

[6] 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, 2026. arXiv:2603.08155.

[7] 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.

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

[9] Amita Kamath, Kai-Wei Chang, Ranjay Krishna, Luke Zettlemoyer, Yushi Hu, and Marjan Ghazvininejad. GenEval 2: Addressing benchmark drift in text-to-image evaluation. *arXiv preprint arXiv:2512.16853*, 2025.

[10] Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.

[11] Inho Kong, Sojin Lee, Youngjoon Hong, and Hyunwoo Kim. Error as signal: Stiffness-aware diffusion sampling via embedded Runge–Kutta guidance. In *International Conference on Learning Representations*, volume 2026, pp. 21287–21309, 2026.

[12] 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 (NeurIPS)*, 2024. arXiv:2404.07724.

[13] Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. DPM-Solver++: Fast solver for guided sampling of diffusion probabilistic models. *Machine Intelligence Research*, 22(4):730–751, 2025. doi: 10.1007/s11633-025-1562-4.

[14] Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. SDXL: Improving latent diffusion models for high-resolution image synthesis. In *International Conference on Learning Representations (ICLR)*. OpenReview.net, 2024.

[15] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 10674–10685. IEEE, 2022. doi: 10.1109/CVPR52688.2022.01042.

[16] 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 (ICLR)*, 2025. arXiv:2410.02416.

[17] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations (ICLR)*, 2021.

[18] Qinsheng Zhang, Jiaming Song, and Yongxin Chen. Improved order analysis and design of exponential integrator for diffusion models sampling. *arXiv preprint arXiv:2308.02157*, 2023.

[19] Shiheng Zhang. Asymptotic preservation and uniform accuracy of diffusion and flow-matching samplers. *arXiv preprint arXiv:2607.04113*, 2026.

[20] Kaiwen Zheng, Cheng Lu, Jianfei Chen, and Jun Zhu. DPM-Solver-v3: Improved diffusion ODE solver with empirical model statistics. In *Advances in Neural Information Processing Systems*, 2023.

Organization. Appendix A covers implementation and metrics, Appendix B the proofs of Theorems 1–3 and the band results, Appendix C the same-field protocol behind Section 4.3, Appendix D the generation tables and controls behind Sections 4.1–4.2, and Appendix E the run records.

A Implementation and metrics

All updates are deterministic. We use $g = 1 + w$ for guidance, K for positive-to-positive intervals and $N = K + 1$ when counting the final closure. Standard CFG returns D_w at zero noise; fitted CFG returns D_c. The factorial study separates finite-step and closure choices. Table 25 maps the reported experiments to their update kernels, precisions and closures.

Coordinate conversion and update identities. For native VP state z_t, set $x = z_t/\sqrt{\bar\alpha_t}$ and $\sigma_t = \sqrt{(1-\bar\alpha_t)/\bar\alpha_t}$. A noise prediction gives $D_j = x - \sigma_t \epsilon_j$. Dividing the deterministic CFG++ update by $\sqrt{\bar\alpha_{t-1}}$ yields $rx + \beta D_c + (1 - r - \beta)D_u$, hence $\beta = 1 - r - \alpha = 1 - s$ and

$$\beta_{\mathrm{CFG}} = (1-r)g, \quad \beta_{\mathrm{fit}} = 1 - r^g, \quad \beta_{\mathrm{sign}} = \min\{(1-r)g, 1\}, \quad \beta_{\mathrm{res}} = \min\{(1-r)g, 1+r\}.$$

These identities require deterministic, unthresholded updates. A fixed nonzero β has $g_{\mathrm{eff}} = \beta/(1 - r)$, which grows under refinement. Fitted instead has $\beta = (1-r)g + O(h^2)$. FP32/FP64 tests check these identities and nonlinear-trajectory parity; complete parity requires matching closures.

Precision-aware reference update. The factored execution in (4) preserves the small terminal multiplier directly. The existing P2 sampler uses persistent FP32 states and promotes predictions before denoiser reconstruction. Its fitted branch ($w > 0$, $0 < r < 1$) is

```
x, Dc, Du = x.float(), Dc.float(), Du.float()
s = math.exp((1 + w) * math.log(r))
x_next = Dc + r * (x - Du) + s * (Du - Dc)
```

It is available as `finite_step`, arithmetic `fitted_factored_v1`. The CPU audit covers 60 scalar cases in four input dtypes and three arithmetic paths. P2 retains every tested factor representable in its FP32 output; underflow is classified separately. Twenty actual-wrapper executions with synthetic FP16/BF16 predictions check FP32 branch mixing, states, both closures and CPU autocast on/off. Maximum error against an independent FP64 recurrence is 7.9×10^{-6} with unchanged network-call counts. Two additional CUDA scalar checks retain the pure-mode factors at $w = 8, h \in \{1.32, 4\}$ with relative error below 6×10^{-8}. The 30k SDXL and GenEval2 runs updated states in FP16, which can round prediction products before FP32 reconstruction (Table 25).

Cancellation in the unfactored form (post hoc). Evaluating $D_c + r(x - D_c) + (r^g - r)(D_u - D_c)$ directly loses representable factors to subtraction cancellation (audit of 1,620 scalar steps and 288 crossover trajectories): for $x = 1, w = 8$ on the pure layer it returns zero at $h = 1.32$ in FP16 and $h = 4$ in FP32, although the factors 6.9276×10^{-6} and 2.3195×10^{-16} are representable. The factored update above avoids this. On the finite crossover ($a = 0, b = 1, \sigma_0 = 8, x_0 = 1$, log horizon 12), FP64 log products agree with stable analytic products within 3×10^{-13}, so the distortion in Table 5 is discretization error; FP32 log-product errors stay below 1.18×10^{-6} across horizons $1, 4, 8, 12$.

Table 5: FP64 log-relative distortion on the finite crossover at horizon 12. The bound is Corollary 2; $w = 0$ is exact up to roundoff.

	$w = 4$		$w = 8$	
h	Attained	Bound	Attained	Bound
0.2	1.18641	1.44000	3.96324	4.80000
0.1	0.59212	0.66000	1.97665	2.20000
0.05	0.29574	0.31500	0.98666	1.05000
0.025	0.14778	0.15375	0.49285	0.51250

CIFAR-10 recipe. Table 2 uses public NVIDIA EDM `edm-cifar10-32x32-cond-vp.pkl` and `edm-cifar10-32x32-uncond-vp.pkl` from `https://nvlabs-fi-cdn.nvidia.com/edm/pretrained/`. At $w = 8$ the eight positive nodes are

$$\sigma_i = \left[80^{1/7} + \frac{i}{7}\left(0.002^{1/7} - 80^{1/7}\right)\right]^7, \quad i = 0, \ldots, 7.$$

Appending zero gives seven finite steps and one closure. The terminal interval disables guidance on all seven finite steps. Diagnostics evaluate both branches; sampling alone can omit unconditional evaluations at guidance-off steps and conditional closure. States and inputs are FP32; serialized EMA modules retain their internal precision settings.

Snapshot `standard-endpoint-20260711-v1` uses `image_quality_gate.py`; the study index locates its records. Each method generates 50,000 outputs in batches of 64 and shards of 1,000. At sample offset j, CPU seed $14000 + j$ generates the whole shared batch; labels are global indices modulo ten. Reproduction therefore requires the original batch size. Conversion is $\mathrm{uint8}(\mathrm{clip}(127.5x + 128, 0, 255))$.

Metrics and record conventions. Amp excludes the zero-noise output. Its initial-norm denominator is floored at 10^{-30} in the image gate; the factorial harness checks positivity and divides directly. CIFAR clipping and latent saturation use the component fractions defined in Section 4.1, with percentiles across images. FID uses precomputed Inception reference statistics; KID uses the run-specific feature bank. Shared KID subsets measure resampling stability. Classifier accuracy/probability, CLIP and GenEval2 are conditioning proxies. Replicate blocks remain separate. A zero p95 permits nonzero outliers.

Manifests retain checkpoints, meshes, seeds, dtypes and source hashes. Only positive-endpoint steps enter H; α denotes the update coefficient and $\bar{\alpha}_t$ the native signal coefficient. Superscripts ex and ref identify analytic and numerically refined flows. Record keys `g2`, `fitted` and *terminal fit* all identify fitted CFG.

B Proofs and model scope

Hypothesis 1 (Ordered commuting Gaussian pair)**.** The exact denoisers correspond to $\mathcal{N}(0, \Sigma_{\mathrm{c}})$ and $\mathcal{N}(0, \Sigma_{\mathrm{u}})$, with commuting covariances $0 \preceq \Sigma_{\mathrm{c}} \preceq \Sigma_{\mathrm{u}}$. Their common eigenvalues satisfy $0 \le a_i \le b_i$.

This pair has a coherent continuous-condition realization: $X_{\mathrm{data}} = M + \varepsilon$, with independent $M \sim \mathcal{N}(0, \Sigma_{\mathrm{u}} - \Sigma_{\mathrm{c}})$ and $\varepsilon \sim \mathcal{N}(0, \Sigma_{\mathrm{c}})$. Conditioning on $M = 0$ gives the stated class, and marginalizing M gives the unconditional law. Throughout, $w \ge 0$ is fixed and mesh intervals have positive endpoints, except where a separate zero-noise closure is stated.

On each eigen-coordinate, write

$$A = \frac{\sigma^2}{a + \sigma^2}, \quad B = \frac{\sigma^2}{b + \sigma^2}, \quad \eta_{\mathrm{c}} = A\xi/\sigma, \quad \eta_{\mathrm{u}} = B\xi/\sigma, \qquad 0 \le B \le A \le 1. \tag{19}$$

On a nonzero coordinate trajectory, the guided ODE (1) becomes

$$\frac{\mathrm{d}\log|\xi|}{\mathrm{d}\log\sigma} = \mu_w(\sigma) = (1 + w)A - wB. \tag{20}$$

For $a = b = 0$, $\mu_w = 1$; for $0 < a \le b$, $\mu_w \to 0$ as $\sigma \to 0$. For $a = 0 < b$,

$$\mu_w = 1 + \frac{wb}{b + \sigma^2} \longrightarrow 1 + w, \qquad \mu_w(\sqrt{b}) = 1 + w/2. \tag{21}$$

These identities establish the three terminal modes used in the main text.

Lemma 1 (Exact coordinate flow)**.** *On Hypothesis 1, the exact one-step factor is*

$$\Phi_w = \left(\frac{a + \sigma_{n+1}^2}{a + \sigma_n^2}\right)^{(1+w)/2} \left(\frac{b + \sigma_n^2}{b + \sigma_{n+1}^2}\right)^{w/2}. \tag{22}$$

In particular, for $a = 0 < b$ and $r = \sigma_{n+1}/\sigma_n$,

$$\Phi_w = r^{1+w} \left(\frac{b + \sigma_n^2}{b + \sigma_{n+1}^2} \right)^{w/2}. \tag{23}$$

Proof. Integrate (20): $\int \mu_w \, \mathrm{d} \log \sigma = \frac{1+w}{2} \log(a + \sigma^2) - \frac{w}{2} \log(b + \sigma^2)$ up to a constant, since $\mathrm{d} \log(a + \sigma^2) = 2A \, \mathrm{d} \log \sigma$ and $\mathrm{d} \log(b + \sigma^2) = 2B \, \mathrm{d} \log \sigma$. Exponentiate the increment. □

B.1 Barrier and fitted-factor proofs

The frozen update and its Gaussian factor are

$$x_{n+1} = x_n + (\sigma_{n+1} - \sigma_n) \eta_w(x_n, \sigma_n), \tag{24}$$

$$G_{\mathrm{CFG}} = 1 - (1 - r) \mu_w(\sigma_n), \qquad r = e^{-h}. \tag{25}$$

On the pure discriminative layer,

$$G_w(h) = (1 + w)e^{-h} - w. \tag{26}$$

Theorem 1 (Complete barrier statement)**.** *On a nonzero pure discriminative trajectory, for $w > 0$, define $h_{\mathrm{sign}} = \log(1 + 1/w)$, $h_{\mathrm{amp}} = \log(1 + 2/w)$, and $h_{\mathrm{stab}} = \log((1 + w)/(w - 1))$ when $w > 1$. (a) $G_w > 0$ iff $h < h_{\mathrm{sign}}$, is zero at equality, and is negative above it. (b) The pure-layer residual ratio is $|G_w|e^h$, exceeding one iff $h > h_{\mathrm{amp}}$. At fixed such h, after K positive-noise steps on $\sigma_n = \sigma_0 e^{-nh}$, with $\eta_n := \eta_w(\xi_n, \sigma_n)$,*

$$\frac{|\eta_K|}{|\eta_0|} = [we^h - (1 + w)]^K = \left(\frac{\sigma_0}{\sigma_K} \right)^{\kappa_{\mathrm{amp}}}, \qquad \kappa_{\mathrm{amp}} = \frac{\log[we^h - (1 + w)]}{h} > 0. \tag{27}$$

(c) $|G_w| \le 1$ for all h if $w \le 1$, and iff $h \le h_{\mathrm{stab}}$ otherwise. For $w > 1$ the thresholds satisfy $h_{\mathrm{sign}} < h_{\mathrm{amp}} < h_{\mathrm{stab}}$.

For finite $a = 0 < b$, the CFG iterates satisfy the identity

$$|\eta_{n+1}| = \frac{\mu_{n+1}}{\mu_n} \frac{|G_{\mathrm{CFG},n}|}{r} |\eta_n|, \qquad \mu_n = 1 + \frac{wb}{b + \sigma_n^2}.$$

Fix $h > h_{\mathrm{amp}}$, $r = e^{-h}$, and choose $\varepsilon > 0$ so that $\beta_\varepsilon = we^h - (1 + w) - (e^h - 1)w\varepsilon^2 > 1$. Let j be the first node with $\sigma_j \le \varepsilon\sqrt{b}$, and assume $\eta_j \ne 0$. For every subsequent endpoint $K \ge j$, $|\eta_K|/|\eta_j| \ge \beta_\varepsilon^{K-j}$. For fixed $h, \varepsilon, b, \sigma_j$, this lower bound diverges as $K - j \to \infty$ and $\sigma_K \to 0$. Separately, $\log \beta_\varepsilon / h \to \kappa_{\mathrm{amp}}$ as $\varepsilon \to 0$.

Proof. From (26): $G_w > 0 \iff \mathrm{e}^{-h} > w/(1 + w) \iff h < \log(1 + 1/w)$, which is (a). For (b), on the pure layer $\eta_w \propto \xi/\sigma$, so the residual ratio per step is $|\xi_{n+1}/\xi_n| \cdot \sigma_n/\sigma_{n+1} = |G_w|\mathrm{e}^h$. On the branch $G_w \le 0$ this is $(w - (1 + w)\mathrm{e}^{-h})\mathrm{e}^h = w\mathrm{e}^h - (1 + w)$, which exceeds 1 iff $\mathrm{e}^h > 1 + 2/w$; on the branch $G_w \ge 0$ it equals $(1 + w) - w\mathrm{e}^h \le 1$ always. Compounding over K pure-layer steps gives (27) exactly. For the finite-crossover extension, $\eta_w = \mu_n \xi/\sigma$ on the coordinate, so $|\eta_{n+1}| = (\mu_{n+1}/\mu_n) \, |G_{\mathrm{CFG}}| \, |\eta_n|/r$ with $G_{\mathrm{CFG}} = 1 - (1 - r)\mu_n$ (25). On the subwindow, $\sigma_n^2 \le \varepsilon^2 b$ gives $\mu_n \ge 1 + w/(1 + \varepsilon^2)$, hence $1 - (1 - r)\mu_n \le G_w(h) + (1 - r)w\varepsilon^2$; for $h > h_{\mathrm{amp}}$ and $\beta_\varepsilon > 1$ this is negative and $|G_{\mathrm{CFG}}|/r \ge w\mathrm{e}^h - (1 + w) - (\mathrm{e}^h - 1)w\varepsilon^2 = \beta_\varepsilon$. Since μ increases as σ decreases, $\prod_{n=j}^{K-1} \mu_{n+1}/\mu_n = \mu_K/\mu_j \ge 1$, so for $\eta_j \ne 0$ the product of the $K - j$ per-step ratios from node j gives $|\eta_K|/|\eta_j| \ge \beta_\varepsilon^{K-j}$. For (c), $G_w \le 1$ always, and $G_w \ge -1 \iff (1 + w)\mathrm{e}^{-h} \ge w - 1$, vacuous for $w \le 1$ and equivalent to $h \le h_{\mathrm{stab}}$ otherwise. The ordering $h_{\mathrm{amp}} < h_{\mathrm{stab}}$ follows from $\frac{1+w}{w-1} - \left(1 + \frac{2}{w}\right) = \frac{2}{w(w-1)} > 0$. □

For finite crossover, $h_n \le h_{\mathrm{sign}}$ is sufficient but not necessary for nonnegative factors: the actual condition is $(1 - r_n)\mu_n \le 1$. The residual condition also includes μ_{n+1}/μ_n. At a zero endpoint the model-free CFG closure is

$$\lim_{r \to 0} x_{n+1} = D_w = (1 + w)D_{\mathrm{c}} - wD_{\mathrm{u}}. \tag{28}$$

The fitted update, repeated here to fix the proof notation, is

$$x_{n+1} = D_{\mathrm{c}} + r(x_n - D_{\mathrm{c}}) + (r^{1+w} - r)(D_{\mathrm{u}} - D_{\mathrm{c}}). \tag{29}$$

Proposition 1 (Consistency). *For fixed $w \geq 0$, $r^{1+w} - r = w(r-1) + \frac{1}{2}w(w+1)h^2 + O(h^3)$. Thus (29) is first-order consistent with the same guided ODE and equals DDIM at $w = 0$.*

Proof. $r^{1+w} - r = \mathrm{e}^{-(1+w)h} - \mathrm{e}^{-h}$ and $w(r-1) = w(\mathrm{e}^{-h} - 1)$; both are $-wh + O(h^2)$, and expanding to second order, $r^{1+w} - r - w(r-1) = \mathrm{e}^{-(1+w)h} - (1+w)\mathrm{e}^{-h} + w = \frac{1}{2}w(w+1)h^2 + O(h^3)$. □

Proposition 2 (Fitted coordinate factor). *Under Hypothesis 1, for $0 < r < 1$,*

$$G_{\mathrm{fit}} = 1 - (1-r)A + (r^{1+w} - r)(A - B) = (1 - A) + rB + r^{1+w}(A - B) > 0. \tag{30}$$

It equals r, r^{1+w} and 1 at $(A, B) = (1, 1), (1, 0), (0, 0)$, respectively. The full update tends to D_{c} as $r \to 0$.

Proof. On an eigen-coordinate, $D_{\mathrm{c}} = (1-A)\xi$, $x - D_{\mathrm{c}} = A\xi$, and $D_{\mathrm{u}} - D_{\mathrm{c}} = (A-B)\xi$; substituting into (29) gives the first expression in (30), and regrouping $1 - A + rA - r(A-B) + r^{1+w}(A-B)$ gives the second. Nonnegativity of each term is $0 \leq B \leq A \leq 1$ and $0 < r < 1$. The regime values and the $r \to 0$ limit are read off termwise. □

Proposition 3 (semigroup characterization of the terminal coefficient). *On the pure discriminative terminal mode, let the factor of a scalar coefficient update be $F_w(r) = r + \alpha(r, w)$. Suppose $F_w : (0, 1] \to \mathbb{R}$ is continuous,*

$$F_w(r_1 r_2) = F_w(r_1)F_w(r_2), \qquad F_w(1) = 1,$$

and has the correct infinitesimal guided generator,

$$F_w(\mathrm{e}^{-h}) = 1 - (1+w)h + o(h) \qquad (h \downarrow 0).$$

Then uniquely

$$F_w(r) = r^{1+w}, \qquad \alpha(r, w) = r^{1+w} - r.$$

Thus mesh composition and first-order consistency determine the terminal-fitted coefficient without covariance information.

Proof. The local expansion makes $F_w(\mathrm{e}^{-h}) > 0$ for all sufficiently small h. For arbitrary $r = \mathrm{e}^{-t} \in (0, 1]$, choose m so large that t/m lies in this positive neighborhood; multiplicativity gives $F_w(r) = F_w(\mathrm{e}^{-t/m})^m > 0$. Hence $\psi(t) := \log F_w(\mathrm{e}^{-t})$ is well defined, continuous, and additive: $\psi(s+t) = \psi(s) + \psi(t)$. The continuous Cauchy equation gives $\psi(t) = ct$, while the infinitesimal expansion gives $c = -(1+w)$. Therefore $F_w(r) = r^{1+w}$ and $\alpha = F_w - r$. □

The effective-weight identity (12) follows by substitution. Its bound uses $1 - r^w = w\int_r^1 t^{w-1}\,\mathrm{d}t$ and $rt^{w-1} \leq 1$ for $r \leq t \leq 1$.

Corollary 1 (Fitted crossover residual). *On $a = 0 < b$, the fitted factor obeys*

$$\frac{G_{\mathrm{fit},n}}{r_n} = B_n + r_n^w(1 - B_n) \leq 1. \tag{31}$$

Thus its residual satisfies (8).

Proof. Set $s_n = r_n^{1+w}$ in Theorem 2. □

B.2 Horizon-independent relative accuracy

Theorem 3 (floor-uniform accuracy on the ordered Gaussian model, complete statement). *Under Hypothesis 1, fix $w \geq 0$ and $H > 0$, let $\lambda = -\log\sigma$, and consider an arbitrary decreasing noise mesh $\sigma_0 > \cdots > \sigma_K > 0$ with $h_n = \lambda_{n+1} - \lambda_n \leq H$. Let T_{fit} and T_{ex} be the linear endpoint maps of the fitted scheme and exact guided flow, respectively, and set $g = 1 + w$ and*

$$\varepsilon_H(w) := H\left[1 + 2w + \frac{\mathrm{e}^{gH}(1+H)}{4}\left(g^2 + w^2\right)\right]. \tag{32}$$

Let $G_{n,i}$ and $\Phi_{n,i}$ be the fitted and exact one-step factors from (30) and (22) on coordinate i. Then

$$\left|\log \frac{\prod_{n=0}^{K-1} G_{n,i}}{\prod_{n=0}^{K-1} \Phi_{n,i}}\right| \le \varepsilon_H(w) \qquad \textit{for every shared eigen-coordinate } i, \tag{33}$$

and consequently

$$\|T_{\mathrm{fit}} - T_{\mathrm{ex}}\|_{\mathrm{op}} \le \mathrm{e}^{\varepsilon_H(w)} - 1. \tag{34}$$

For every input law ν with finite second moment,

$$W_2\big((T_{\mathrm{fit}})_{\#}\nu, (T_{\mathrm{ex}})_{\#}\nu\big) \le \big(\mathrm{e}^{\varepsilon_H(w)} - 1\big)\left(\mathbb{E}_\nu |X|^2\right)^{1/2}. \tag{35}$$

For $H \le 1$, $\varepsilon_H(w) \le C(w)H$, where one may take $C(w) = 1 + 2w + \frac{1}{2}\mathrm{e}^{1+w}((1+w)^2 + w^2)$. The quantity $\varepsilon_H(w)$ is independent of dimension, covariance spectra, step count, horizon and terminal floor. It is $O_w(H)$ as $H \to 0$, with a constant depending on w.

Proof. Fix one common eigen-coordinate and write $A_n = \sigma_n^2/(a+\sigma_n^2)$, $B_n = \sigma_n^2/(b+\sigma_n^2)$, and $\mu_n = gA_n - wB_n$, with $g = 1 + w$ fixed. In this proof, $\mu(\lambda) := \mu_w(\mathrm{e}^{-\lambda})$. Separate the frozen-exponent error from variation of the exponent:

$$\log \frac{\prod_n G_n}{\prod_n \Phi_n} = \sum_n (\log G_n + h_n \mu_n) + \left[\int_{\lambda_0}^{\lambda_K} \mu(\lambda)\,\mathrm{d}\lambda - \sum_n h_n \mu_n\right].$$

Represent the deterministic fitted factor by an auxiliary three-point random variable

$$\kappa_n = \begin{cases} 0, & \text{with probability } 1 - A_n, \\ 1, & \text{with probability } B_n, \\ g, & \text{with probability } A_n - B_n. \end{cases}$$

Then $\mathbb{E}\kappa_n = \mu_n$ and the fitted factor (30) is $G_n = \mathbb{E}\mathrm{e}^{-h_n\kappa_n}$. Directly from the three-point law,

$$\mathrm{Var}(\kappa_n) = g^2 A_n(1-A_n) + w^2 B_n(1-B_n) - 2gwB_n(1-A_n) \le g^2 A_n(1-A_n) + w^2 B_n(1-B_n).$$

Let $\mathbb{E}_t$ denote expectation under the normalized tilt by $\mathrm{e}^{-t\kappa_n}$, $0 \le t \le h_n$. Its density ratio is at most e^{gh_n}. Since variance minimizes squared distance to a constant,

$$\mathrm{Var}_t(\kappa_n) \le \mathbb{E}_t[(\kappa_n - \mu_n)^2] \le \mathrm{e}^{gh_n}\,\mathrm{Var}(\kappa_n).$$

The second derivative of $\log \mathbb{E}\mathrm{e}^{-t\kappa_n}$ is this tilted variance. Taylor's formula gives

$$0 \le \log G_n + h_n\mu_n \le \frac{1}{2}\mathrm{e}^{gh_n} h_n^2\,\mathrm{Var}(\kappa_n). \tag{36}$$

For any bounded-variation function f, let $\mathrm{TV}(f)$ denote its total variation on $[\lambda_0, \lambda_K]$. The left-rectangle rule on this nonuniform mesh satisfies

$$\left|\sum_n h_n f(\lambda_n) - \int_{\lambda_0}^{\lambda_K} f(\lambda)\,\mathrm{d}\lambda\right| \le H\,\mathrm{TV}(f).$$

Viewing A and B as functions of λ through $\sigma = \mathrm{e}^{-\lambda}$, $A' = -2A(1-A)$ and $B' = -2B(1-B)$. Integrating these identities bounds each crossover-profile integral by $1/2$. Each profile is unimodal with height at most $1/4$, so its total variation is at most $1/2$ as well. The rectangle estimate gives

$$\sum_n h_n A_n(1-A_n) \le \frac{1+H}{2}, \qquad \sum_n h_n B_n(1-B_n) \le \frac{1+H}{2}.$$

Summing (36) hence bounds the accumulated cumulant defect by

$$\frac{H\mathrm{e}^{gH}(1+H)}{4}(g^2 + w^2).$$

Meanwhile the bounded-variation rectangle rule gives

$$\left|\sum_n h_n\mu_n - \int_{\lambda_0}^{\lambda_K} \mu(\lambda)\,\mathrm{d}\lambda\right| \le H\,\mathrm{TV}(\mu) \le H(g+w) = H(1+2w).$$

Since the exact log-factor is $-\int_{\lambda_0}^{\lambda_K} \mu\,\mathrm{d}\lambda$, these two estimates prove (33). The exact factor lies in $(0,1]$ because $\mu \ge 0$, so the scalar endpoint error is at most $\mathrm{e}^{\varepsilon_H(w)} - 1$. Simultaneous diagonalization proves (34); coupling the two pushforwards with the same input $X \sim \nu$ proves (35). □

Corollary 2 (Sharper discriminative-crossover bound)**.** *For $a = 0 < b$, $w \geq 0$ and uniform log step $0 < h \leq 1$, define $\Theta_K = \prod_{n=0}^{K-1} G_n / \prod_{n=0}^{K-1} \Phi_n$ using the fitted and exact coordinate factors. From the same input ξ_0, $\xi_K^{\mathrm{fit}} = \Theta_K \xi_K^{\mathrm{ex}}$. Then $0 \leq \log \Theta_K \leq (1+h)w(w+2)h/4$, independently of b, K, σ_K. For an input law with finite second moment, synchronous coupling gives $W_2(\mathrm{Law}(\xi_K^{\mathrm{fit}}), \mathrm{Law}(\xi_K^{\mathrm{ex}})) \leq (e^{w(w+2)h/2} - 1)(\mathbb{E}|\xi_0|^2)^{1/2}$. The constant depends on w.*

Proof. On $a = 0$ the exact one-step factor (Lemma 1) is $\Phi_n = r^{1+w}\big((1+u_n)/(1+r^2 u_n)\big)^{w/2}$, with $u_n = \sigma_n^2/b$, $r = \mathrm{e}^{-h}$; the fitted factor (Proposition 2, $A = 1$) is $G_n = \mathrm{e}^{-h}(u_n + \mathrm{e}^{-wh})/(1+u_n)$. Writing $y_n = \log u_n$, $\delta(h, y) := \log(G_n/\Phi_n) = wh + \log(\mathrm{e}^y + \mathrm{e}^{-wh}) - \big(1 + \frac{w}{2}\big)\log(1+\mathrm{e}^y) + \frac{w}{2}\log(1+\mathrm{e}^{y-2h})$. Then $\delta(0, y) = \partial_h \delta(0, y) = 0$ (first-order consistency, Proposition 1), and with $F(y) = \mathrm{e}^y/(1+\mathrm{e}^y)^2$ one has $\partial_h^2 \delta = w^2 F(y + wh) + 2wF(y - 2h) \geq 0$; hence $\delta \geq 0$, so $\Theta_K = \prod_n \mathrm{e}^{\delta(h, y_n)} \geq 1$, and Taylor's formula gives $\delta(h, y_n) = \int_0^h (h - s)\big[w^2 F(y_n + ws) + 2wF(y_n - 2s)\big]\, ds$. On the uniform mesh $y_{n+1} = y_n - 2h$. For every shift c, apply the rectangle estimate on the K intervals $[y_{n+1} + c, y_n + c]$. Since $\int_{\mathbb{R}} F = 1$ and $\mathrm{TV}_{\mathbb{R}}(F) = 1/2$, it gives $2h \sum_{n=0}^{K-1} F(y_n + c) \leq 1 + h$, so $\log \Theta_K \leq \int_0^h (h - s)\frac{1+h}{2h}(w^2 + 2w)\, ds = \frac{1+h}{4} w(w+2)h$. Since $h \leq 1$, the last bound is at most $w(w+2)h/2$. Synchronous coupling with $|\xi_K^{\mathrm{ex}}| \leq |\xi_0|$ gives the W_2 bound. □

B.3 Shared and scale-informed stability proofs

Proof of Theorem 2. Substituting $D_{\mathrm{c}} = (1 - A)\xi$ and $D_{\mathrm{u}} = (1 - B)\xi$ gives $G_s = (1 - A) + rB + s(A - B)$. The weights $1 - A, B, A - B$ are nonnegative and sum to one. Thus $0 \leq s \leq r < 1$ gives $0 \leq G_s \leq 1$. For finite variances and positive noise, $a > 0$ implies $1 - A > 0$; $a = 0$ implies $B > 0$, giving strict positivity. On $a = 0 < b$, put $L_n = B_n + (s_n/r_n)(1 - B_n)$. For the wider band $-r_n \leq s_n \leq r_n$, $|L_n| \leq 1$ and

$$|\eta_w(\xi_{n+1}, \sigma_{n+1})| = \frac{\mu_{n+1}}{\mu_n} |L_n|\, |\eta_w(\xi_n, \sigma_n)|.$$

Multiplication yields (8), including trajectories that reach zero; $1 \leq \mu_n \leq 1 + w$. □

Proposition 4 (A finite variance bound)**.** *Under Hypothesis 1, suppose $0 < b_{\max} < \infty$ bounds all marginal eigenvalues. For each positive-noise step let $q_{\max} = b_{\max}/(b_{\max} + \sigma^2)$. Every coefficient in $[-r/q_{\max}, 0]$ has $0 \leq G_s \leq 1$ on all coordinates and satisfies (8) on $a = 0 < b$. The crossover coefficient with $\hat{b} = b_{\max}$ has strictly positive factors.*

Proof. Since $A - B \leq 1 - B = b/(b + \sigma^2) \leq q_{\max}$ and $\alpha \leq 0$,

$$G_s = (1 - A) + rA + \alpha(A - B) \geq r + \alpha q_{\max} \geq 0, \qquad G_s \leq (1 - A) + rA \leq 1.$$

For $a = 0$, $G_s = r + \alpha(1 - B) \leq r$; the preceding residual product proof applies. For crossover matching, $0 < \Phi \leq r$ since the integrated exponent is at least one, so $\alpha_{\mathrm{cross}} = (\Phi - r)/q_{\max} \in (-r/q_{\max}, 0]$ and $G_s \geq \Phi > 0$. □

The assumption is an upper bound in the Gaussian model. A scale chosen from learned-model calibration supplies an operating parameter. For example, $r = 1/2, w = 2, \sigma^2 = \hat{b} = 1$ gives $\Phi = 1/5$, $\alpha_{\mathrm{cross}} = -3/5$ and pure test factor $s = -1/10$, while the matched finite-scale factor is positive.

The sign and residual caps in (9) have artifact names `clip_noflip` and `clip_e1`. They clip coefficients, recover the prescribed ODE under refinement, and use conditional closure at $r = 0$. On the pure layer, sign cap has factor $\max\{(1 + w)r - w, 0\}$. When active, its coefficient differs from fitted CFG by r^{1+w}. At $w = 11, r = 11/12$ these factors are zero and about 0.352; two half-log capped steps instead give $[(1 + w)\sqrt{r} - w]^2 > 0$. Fitted factors compose exactly to r^{1+w}.

Proposition 5 (Absolute floor-uniform accuracy of the caps)**.** *For exact centered Gaussian denoisers with commuting covariances $0 \preceq \Sigma_{\mathrm{c}} \preceq \Sigma_{\mathrm{u}}$, fix $w \geq 0$ and any finite decreasing positive noise mesh, with $H = \max_n \log(\sigma_n/\sigma_{n+1}) > 0$. For either cap above, its linear endpoint map satisfies*

$$\|T_{\mathrm{cap}} - T_{\mathrm{ex}}\|_{\mathrm{op}} \leq (4w + 6)H.$$

The constant is independent of dimension, covariance spectra, step count, total log-noise horizon, and terminal floor. It is not optimized or uniform in w.

Proof. Work in a common eigen-coordinate and set $g = 1 + w$, $\lambda = -\log\sigma$, and $\mu = gA - wB$. Since A, B are nonincreasing in λ with total variation at most one,

$$0 \le \mu \le g, \qquad \mathrm{TV}(\mu) \le 2g - 1.$$

First suppose $H \le H_0 = (2g+1)^{-1}$. For $q_n = 1 - e^{-h_n}$, we have $gq_n \le gh_n < 1/2$ and $r_n + w(r_n - 1) = 1 - gq_n > 1/2$; neither cap is active. Both factors are $G_n = 1 - z_n$, where $z_n = q_n\mu_n < 1/2$. Using $0 \le h_n - q_n \le h_n^2/2$ and $0 \le -\log(1-z) - z \le z^2$ for $0 \le z \le 1/2$ gives

$$|\log G_n + h_n\mu_n| \le \tfrac{1}{2}h_n^2\mu_n + h_n^2\mu_n^2 \le (g + \tfrac{1}{2})h_n^2\mu_n.$$

Put $S = \sum_n h_n\mu_n$ and $d = (g + 1/2)H \le 1/2$. Then $|\log\prod_n G_n + S| \le dS$, so

$$\left|\prod_n G_n - e^{-S}\right| \le dSe^{-(1-d)S} \le \frac{d}{(1-d)e} \le 2d.$$

This estimate is uniform in $S \ge 0$, even when S grows with the horizon; at $S = 0$ both compared factors equal one. The bounded-variation rectangle rule gives $|S - \int \mu\, d\lambda| \le H(2g-1)$. Since e^{-s} is 1-Lipschitz for $s \ge 0$, the absolute endpoint error is at most $2d + H(2g-1) = 4gH$.

For $H > H_0$, bound the factor magnitudes directly. The decomposition is a convex combination of 1, r, and $r + \alpha$, including zero or negative values of the last factor. The sign cap has $r + \alpha \in [0, r]$; the residual cap has $r + \alpha \in [-r, r]$. Thus both numerical endpoint factors have magnitude at most one, and the exact factor lies in $(0, 1]$. Their absolute difference is at most two. When $H > H_0$, $2 < (4g+2)H = (4w+6)H$. Combining the two cases and taking the maximum over the common orthonormal eigenbasis proves the operator bound. □

For any common input law ν with finite second moment, synchronous coupling also yields

$$W_2((T_{\mathrm{cap}})_{\#}\nu, (T_{\mathrm{ex}})_{\#}\nu) \le (4w+6)H\left(\mathbb{E}_\nu\|X\|_2^2\right)^{1/2}.$$

This absolute bound does not rank the methods by accuracy, and is not a fixed-cost AP result: controlling H can require more steps as the log-noise horizon grows.

Absolute error versus accumulated multiplicative distortion. The fitted bound (33) controls the endpoint log-factor ratio independently of the horizon. Neither cap has the same property. Fix $w > 0$ and $0 < h < \log(1 + 1/w)$, and let $r = e^{-h}$. Both caps are inactive. On the pure discriminative layer their common factor $P = (1+w)r - w$ and exact factor $E = r^{1+w}$ satisfy $0 < P < E < 1$, by strict convexity of $r \mapsto r^{1+w}$. Thus $\log(P^K/E^K) = K\log(P/E) \to -\infty$, whereas $|P^K - E^K| \le E^K \to 0$.

This counterexample also holds for finite marginal variance. Fix $a = 0 < b < \infty$ and $\sigma_n = \sigma_0 r^n$, with w, h, b, σ_0 fixed. Then $B_n = O(r^{2n})$ and

$$G_{\mathrm{cap},n} = P + w(1-r)B_n = P + O(r^{2n}), \qquad \Phi_n = E(1 + O(r^{2n})).$$

Since $P > 0$, the log corrections are summable, and

$$\log\frac{\prod_{n<K} G_{\mathrm{cap},n}}{\prod_{n<K}\Phi_n} = K\log(P/E) + O(1) \longrightarrow -\infty.$$

The $O(1)$ remainder is bounded independently of K for these fixed parameters. Both endpoint products still tend to zero, while their signed log ratio tends to $-\infty$. This finite-variance example separates fitted CFG's log-relative guarantee from the two caps despite their shared absolute accuracy.

B.4 A FINITE-CONDITION BOUNDARY CASE

Conditional collapse does not by itself imply exponent $1 + w$. Let X_{data} equal $-m$ or m with equal probability, $m > 0$, and let the condition reveal its sign. For the positive condition, exact denoisers are

$$D_{\text{c}}(x, \sigma) = m, \qquad D_{\text{u}}(x, \sigma) = m \tanh(mx/\sigma^2).$$

Consequently, the local residual derivative is

$$\sigma \partial_x \eta_w = 1 + w \frac{m^2}{\sigma^2} \operatorname{sech}^2(mx/\sigma^2) \longrightarrow 1$$

uniformly on $|x - m| \leq m/2$. Indeed, on this neighborhood, $0 \leq \sigma \partial_x \eta_w - 1 \leq 4wm^2\sigma^{-2} \exp(-m^2/\sigma^2) \to 0$ for fixed $m > 0$ and $w \geq 0$. The posterior already identifies the class there, so the marginal denoiser becomes locally constant too. This is a statement about local linearization, not a globally homogeneous scalar flow. Near the ambiguous point $x = 0$ the same derivative instead equals $1 + wm^2/\sigma^2$. The ordered Gaussian model's persistent discriminative terminal mode is therefore one geometry, not a universal description of discrete labels or text conditions. Its coherent realization in Appendix B uses a continuous Gaussian condition.

Reverse covariance ordering. Ordering is also substantive. For $w > 0$ and $b = 0 < a$, we have $A \to 0$, $B = 1$ and $\mu_w = (1+w)A - w \to -w$: the exact guided coordinate expands as noise tends to zero. At fixed $0 < r < 1$, the fitted factor tends to $1 + r - r^{1+w} > 1$. Neither the contraction nor relative-accuracy theorem above asserts a guarantee in this reversed geometry.

B.5 ANALYTIC CALIBRATION PROTOCOLS

Relative-accuracy comparison. The post-hoc comparison in Figure 2 uses the exact integrated flow below, not a frozen exponent. The full sweep uses $(a, b) \in \{(0, 1), (10^{-4}, 10^{-2}), (1, 4), (0, 0)\}$, $w \in \{0, 4, 8\}$, $H \in \{0.2, 0.1, 0.05, 0.025\}$, uniform and alternating $H, H/2$ log steps, $\sigma_0 = 8$, and horizons $L = 0.25, 0.5, \ldots, 4, 8, 12, 20$. The final step is truncated as needed. The CSV contains all four methods' signed factors, nonpositive-step counts, absolute errors, and log-relative errors. The latter are left undefined when any step factor is nonpositive. This is an illustration of the finite-variance counterexample in Appendix B.3, not an empirical claim about learned denoisers.

Accuracy calibration. The additional fixed pair is $(a_1, \ldots, a_4) = (0, 0, 10^{-4}, 1)$, $(b_1, \ldots, b_4) = (0, 1, 10^{-2}, 4)$, with $w \in \{0, 4, 8\}$, $\sigma_0 = 80$, floors $10^{-1}, 10^{-3}, 10^{-6}$, and $H \in \{0.2, 0.1, 0.05, 0.025\}$. Uniform log steps and alternating $H, H/2$ steps end with a truncated step when needed. We compare diagonal products of the fitted factors with the exact Gaussian flow

$$(T_{\text{ex}})_{ii} = \left(\frac{a_i + \sigma_K^2}{a_i + \sigma_0^2} \right)^{(1+w)/2} \left(\frac{b_i + \sigma_K^2}{b_i + \sigma_0^2} \right)^{-w/2}.$$

The operator error is the maximum absolute coordinate-factor difference. All 72 cases are retained in the machine-readable CSV. Table 6 reports error/H across every tested guidance weight and mesh family. Its boundedness as the floor falls illustrates the theorem on a fixed model; it does not prove its uniformity over all covariance spectra.

Table 6: Gaussian accuracy calibration: $10^3 \, \|T_{\text{fit}} - T_{\text{ex}}\|_{\text{op}}/H$. Each interval spans the three fixed floors, not sampling uncertainty.

w	Mesh	$H = .2$	.1	.05	.025
0	uniform	3.030–3.047	3.070–3.086	3.089–3.105	3.099–3.115
0	alternating	2.535–2.550	2.564–2.577	2.577–2.590	2.584–2.597
4	uniform	5.212–5.274	4.652–4.715	4.337–4.397	4.174–4.233
4	alternating	4.126–4.176	3.762–3.814	3.557–3.607	3.450–3.499
8	uniform	3.366–3.427	1.951–1.997	1.374–1.410	1.137–1.168
8	alternating	2.270–2.314	1.424–1.459	1.067–1.096	0.914–0.939

C SAME-FIELD NUMERICAL EXPERIMENTS

C.1 SAME-INTERVAL PROTOCOL AND COMPLETE COUNTS

Design and selection. SDXL base 1.0 revision `462165984030` uses 1024×1024, deterministic DDIM, $g = 12, N = 12$, empty negative prompts, COCO val2017 indices 3000–3031 and seeds 80,000,000–80,000,031 in two blocks. States come from production FP16 CFG. Diagnostics reload original weights in FP32 and promote stored conditioning; integration/pair differences use FP64 with TF32 disabled.

All eleven positive intervals are screened. Two-pass Gram–Schmidt builds a 24-dimensional basis from the guidance gap, conditional residual, state, channel-restricted gaps/residuals and seeded Rademacher vectors. Central differences at $\epsilon = 0.003$ form the symmetric projected CFG differential; its lowest eigenvector is fixed across candidates and $\epsilon \in \{0.01, 0.003, 0.001\}$, with $\delta = \epsilon \max\{\|x\|, 1\}$. The two lowest-response intervals plus fixed step 5 select three distinct cases per prompt: 96 of 352 screened intervals. Of 26 negative screen values, 25 are selected. Selection precedes reference evaluation.

References and verification. The primary field inverts the training log-sigma table to a fractional timestep. Heun integrates $\mathrm{d}X/\mathrm{d}\lambda = -\sigma\eta_w$ at subdivisions 16, 32, 64, 128, 256. Two successive refinements must satisfy

$$\max_{s=\pm 1} \frac{\|z_{m,s} - z_{\mathrm{prev},s}\|}{\max\{1, \|z_{m,s}\|\}} \le 10^{-4}, \quad \frac{\|q_m - q_{\mathrm{prev}}\|}{\max\{1, \|q_m\|\}} \le 0.02, \quad |Q_m - Q_{\mathrm{prev}}| \le 0.01 + 0.02|Q_m|.$$

Here $q_m = (z_{m,+} - z_{m,-})/(2\delta)$ and $Q_m = \langle v, q_m \rangle$. All three scales must converge. At $\epsilon = 0.003$, RK4 at the accepted Heun count and an independently refined alternate field must agree under the same tolerances, treating the cross-check as current. The alternate field interpolates adjacent integer-timestep epsilon predictions using training log-sigma weights. Reloaded FP16 diagnostics include batched prediction and latent quantization.

A verified reversal requires CFG $Q \le -0.02$ in both precisions at every scale, converged-reference $Q \ge 0.02$ at every scale, and both agreeing cross-checks with $Q \ge 0.02$ at the primary scale. Repair requires candidate $Q \ge 0.02$ at all scales in both precisions. Eight-substep CFG is a higher-cost mesh control; its positivity count uses all scales under the FP32 field.

Table 7: Selected-case SDXL counts by prompt block.

	Block 1	Block 2	Total
Prompts with a verified event	6/16	5/16	11/32
Verified events	7	5	12
Fitted CFG repairs	7/7	4/5	11/12
Eight-substep CFG positive	6/7	5/5	11/12
Crossover repairs	0/7	0/5	0/12
Primary reference unresolved	16/48	16/48	32/96

Of the 25 selected robust CFG negatives, 12 pass all reference checks and 13 have unresolved primary references. Fitted's exception is prompt 3022, step 0 ($Q = -0.0776/-0.0624$ in FP32/FP16 at $\epsilon = 0.001$). Eight-substep CFG's exception is prompt 3001, step 0 ($Q = -0.00041$ at the smallest perturbation). The registered primary criterion combined replication with crossover repair and is not met, since crossover repairs 0/12; fitted CFG's 11/12 is a separate outcome. Events occur at steps 0, 1 and 3; steps 4–10 have no robust screen negative. An earlier eight-dimensional late-step pilot (0/12) is recorded separately and not pooled.

C.2 RESPONSE-VECTOR AND ENDPOINT ERRORS

A post-hoc replay of all 96 selected cases keeps states, directions, checkpoint, conditioning and endpoints unchanged. Eligibility requires all-scale convergence, RK4 agreement, and alternate-extension convergence/agreement, independently of reference sign or candidate performance. All 64

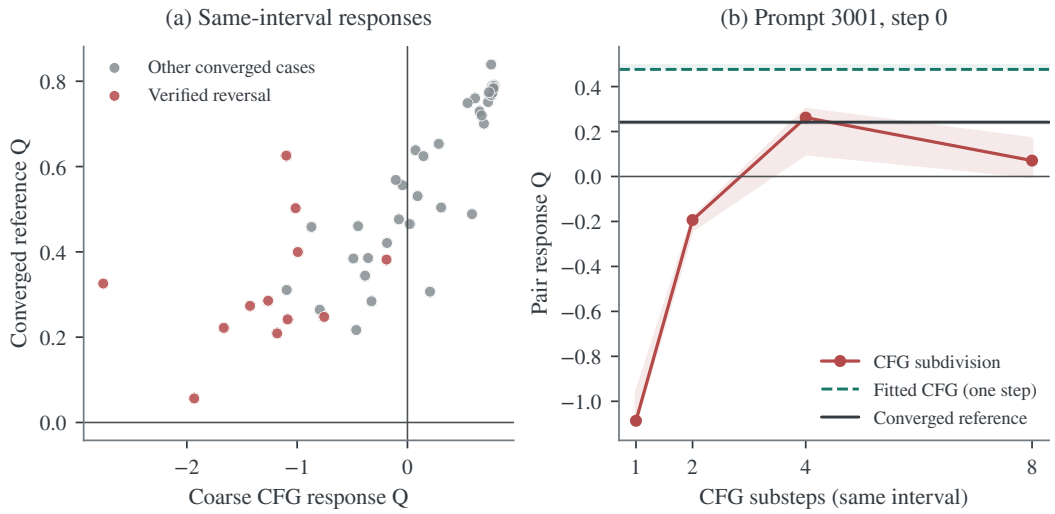

Figure 5: **Same-state, same-interval learned evidence.** Left: selected cases with converged primary references at the middle perturbation size; red marks the twelve verified events. Right: the first verified case in prompt/step order, unchanged from the original analysis. Ranges span perturbation sizes, not confidence intervals. Fitted CFG restores a positive response in one step. CFG subdivision approaches the reference nonmonotonically; eight substeps fail the all-size positivity criterion for this case.

originally converged cases pass; 32 remain unresolved. The replay reproduces all 1,728 previously logged projected responses.

For $q_\Psi = [\Psi(x+\delta v) - \Psi(x-\delta v)]/(2\delta)$, projected, vector and scaled-vector errors are $|Q_\delta(\Psi) - Q_\delta(\Psi_{\text{ref}})|$, $\|q_\Psi - q_{\text{ref}}\|$, and this norm divided by $\max\{1, \|q_{\text{ref}}\|\}$. Endpoint error is

$$\max_{s=\pm 1} \frac{\|\Psi(x+s\delta v) - \Psi_{\text{ref}}(x+s\delta v)\|}{\max\{1, \|\Psi_{\text{ref}}(x+s\delta v)\|\}}.$$

Both candidate precisions use the FP32-field reference. Primary $\epsilon = 0.003$ is fixed; scales remain paired. Intervals use 2,000 whole-prompt bootstrap draws stratified by the original blocks, seed 20260918; tie tolerance is 10^{-12}.

Table 8: Replay on 64 eligible cases, $\epsilon = 0.003$; median / p95.

Precision	Method	Projected	Vector	Scaled vector	Endpoint
FP32	CFG	0.1726/1.7204	1.4810/4.4084	1.2618/3.8662	0.0358/0.0754
FP32	Fitted CFG	0.0883/0.3658	0.6359/1.4452	0.6359/1.1987	0.0316/0.0509
FP32	Sign cap	0.0879/0.3601	0.6357/1.4451	0.6357/1.2049	0.0314/0.0508
FP32	Residual cap	0.0252/0.4181	0.7467/2.0308	0.7467/1.9078	0.0275/0.0472
FP16	CFG	0.1670/1.7456	1.5183/4.4248	1.3239/3.8858	0.0359/0.0755
FP16	Fitted CFG	0.0851/0.3650	0.6409/1.4491	0.6409/1.2013	0.0316/0.0509
FP16	Sign cap	0.0865/0.3596	0.6408/1.4492	0.6408/1.2075	0.0314/0.0508
FP16	Residual cap	0.0346/0.4185	0.7673/2.0419	0.7673/1.9143	0.0275/0.0473

Table 9: FP32 fitted-minus-comparator paired differences, 64 cases. Medians follow per-case subtraction; 95% intervals use block-stratified prompt bootstrap. Negative differences favor fitted.

Comparator	Error	Paired median	95% interval	W/T/L
CFG	Projected	-0.030953	[-0.288939, -0.003862]	45/0/19
CFG	Vector	-0.608753	[-1.336357, -0.357047]	63/0/1
CFG	Scaled vector	-0.608753	[-1.268093, -0.329314]	63/0/1
CFG	Endpoint	-0.001363	[-0.005382, 0.002245]	35/0/29
Sign cap	Projected	0.001675	[0.000799, 0.002817]	7/0/57
Sign cap	Vector	-0.000282	[-0.001249, -0.000022]	41/0/23
Sign cap	Scaled vector	-0.000282	[-0.000958, -0.000016]	41/0/23
Sign cap	Endpoint	0.000187	[0.000116, 0.000225]	0/0/64
Residual cap	Projected	0.016542	[0.006100, 0.048074]	17/0/47
Residual cap	Vector	-0.084243	[-0.156785, -0.048073]	54/0/10
Residual cap	Scaled vector	-0.076749	[-0.147704, -0.044486]	54/0/10
Residual cap	Endpoint	0.004051	[0.003491, 0.004453]	9/0/55

Fitted and sign cap each repair 11/12 events across three scales and both precisions; residual cap repairs 7/12 (8/12 in FP32 alone). Vector improvement over CFG is more consistent than endpoint improvement (63/64 versus 35/64). Projected-only wins are 45/64, 37/64 and 47/64 at $\epsilon = 0.003, 0.01, 0.001$; the paired interval crosses zero at 0.01.

Reference sensitivity. Re-evaluation at accepted Heun count m and $m/2$ passes the original tolerances in all 64 cases and matches logged values within $10^{-5} + 10^{-4}|\text{value}|$. Let $\rho = \|q_m - q_{m/2}\|$; median/p95/maximum are $0.000739/0.001834/0.003715$. The reverse triangle inequality bounds changes in paired vector-error differences by 2ρ.

Table 10: Ordering sensitivity to the preceding Heun reference. W/L/U denotes win/loss/undetermined. Enclosures bound this numerical reference switch; sampling intervals are in Table 9.

Precision	Comparator	Wins	Guaranteed W/L/U	Median enclosure
FP32	CFG	63/64	63/1/0	[-0.611357, -0.606149]
FP32	Sign cap	41/64	22/8/34	[-0.001975, 0.000510]
FP32	Residual cap	54/64	53/10/1	[-0.085590, -0.082896]
FP16	CFG	64/64	64/0/0	[-0.637845, -0.632638]
FP16	Sign cap	52/64	37/1/26	[-0.002766, -0.000307]
FP16	Residual cap	64/64	64/0/0	[-0.096370, -0.093676]

Against CFG, fitted CFG retains 63 wins and one loss; against the sign cap the median enclosure crosses zero; at least 53 of 54 wins against the residual cap survive. The checks bound reference sensitivity; they do not certify the exact-flow error.

Error decomposition. For endpoint defects $e_\pm$, put $b_\delta = (e_+ + e_-)/2$ and $d_\delta = (e_+ - e_-)/2$. Then $\frac{1}{2}(\|e_+\|^2 + \|e_-\|^2) = \|b_\delta\|^2 + \|d_\delta\|^2$ and $q_\Psi - q_{\text{ref}} = d_\delta/\delta$. Bounded second derivatives give $b_\delta = e_0 + O(\delta^2)$. Orthogonal projection onto v likewise separates projected and transverse response errors. Appendix C.4 measures the center defect directly.

C.3 Selection strata and unresolved references

At FP32, $\epsilon = 0.003$, the fixed control step contributes 32/32 valid cases and screen-selected intervals 32/64. Both populations use screen-selected directions within the fixed subspace.

Table 11: Post-hoc fitted-versus-CFG strata. W/L counts smaller/larger error. Δ is the median paired difference; 95% intervals resample prompts within original blocks.

Population	Valid/selected	Vector W/L	Δ vector	Endpoint W/L	Δ endpoint
Fixed step 5	32/32	31/1	-0.115	5/27	0.0050
Screen-selected	32/64	32/0	-1.936	30/2	-0.0168

Vector errors favor fitted in both strata; endpoint ordering changes. Bootstrap intervals quantify variation within these selected populations.

All 32 unresolved cases fail primary all-scale refinement: 14 at step 0, 13 at step 1, three at step 2 and two at step 10. Thirteen of 25 selected robust screen negatives are unresolved. Overlapping failures include 21 alternate-extension refinement, two RK4 agreement and two alternate-extension agreement failures. Cases are retained without replacement.

On every selected interval the sign cap is active. Its same-state difference from fitted is $r^{1+w}(D_{\mathrm{u}} - D_{\mathrm{c}})$; scalar multipliers are 0.00645–0.02573 at steps 0–3, 0.03760 at step 5 and 3.46×10^{-11} at step 10. State and trajectory differences additionally depend on the denoiser gap and propagation.

C.4 Independent local state/response study

Sixteen prompts in two blocks of eight were fixed before sampling by a seeded caption permutation, excluding prior mechanism/smoke cases. Generation benchmarks may share captions; noise and direction seeds are fresh. Directions are normalized isotropic Gaussians without response screening. SDXL uses 1024×1024, $g = 12, N = 12$, interval $\sigma_5 \to \sigma_6$, and anchors from eight Heun subdivisions on each prefix interval. Candidates share each input's denoiser predictions. The field uses FP32 network evaluations and FP64 integration/branch mixing, with separate branches, TF32/autocast disabled. Primary $\epsilon = 0.003$ and sensitivity 0.001 set $\delta = \epsilon \max\{1, \|x\|\}$.

References restart from center and perturbed inputs at Heun counts $4, 8, 16, 32, 64, 128, 256$. Acceptance requires two successive refinements, then RK4 at half the accepted count, with gaps

$$\max_{s\in\{0,+,-\}} \frac{\|z_s - z'_s\|}{\max\{1, \|z_s\|, \|z'_s\|\}} \le 10^{-4}, \qquad \frac{\|q - q'\|}{\max\{1, \|q\|, \|q'\|\}} \le 0.02.$$

Every case passes at both scales: one at Heun32, thirteen at Heun64, two at Heun128. Each scale stops independently at first acceptance.

Let $e_0 = \Psi(x) - \Psi_{\mathrm{ref}}(x)$, $b_\delta = (e_+ + e_-)/2$, $d_\delta = (e_+ - e_-)/2$, and $S = \max\{1, \|\Psi_{\mathrm{ref}}(x)\|\}$. We measure $C = \|e_0\|/S$, $B = \|b_\delta\|^2/S^2$, $D = \|d_\delta\|^2/S^2$, $E = (\|e_+\|^2 + \|e_-\|^2)/(2S^2) = B + D$, and unscaled response error $V = \|d_\delta\|/\delta$. The endpoint maximum in the preceding study uses a different normalization.

Table 12: Independent study, sixteen cases, $\epsilon = 0.003$. V, C are medians; B, D, E are means. Gaussian-b=10 is the prespecified finite-scale comparator, $\hat{b} = 10$.

Method	Median V	Median C	Mean B	Mean D	Mean E
CFG	0.066505	0.014917	2.3939×10^{-4}	1.3557×10^{-7}	2.3953×10^{-4}
Fitted CFG	0.055499	0.021821	4.7268×10^{-4}	7.1697×10^{-8}	4.7275×10^{-4}
Sign cap	0.055258	0.021514	4.5948×10^{-4}	7.1671×10^{-8}	4.5955×10^{-4}
Gaussian-b=10	0.054595	0.020559	4.1974×10^{-4}	7.1932×10^{-8}	4.1981×10^{-4}

Table 13: Fitted-minus-comparator paired differences, with marginal 95% intervals from 2,000 block-stratified prompt draws. W/T/L uses tie tolerance 10^{-12} and retains unresolved numerical orderings.

Comparator	Error	Median paired difference [95% interval]	W/T/L
CFG	V	-0.00909348 $[-0.0158271, -0.00621956]$	16/0/0
CFG	C	0.00553561 $[0.00449846, 0.00769511]$	0/0/16
Sign cap	V	0.000225059 $[0.000140135, 0.00027228]$	3/0/13
Sign cap	C	0.000270943 $[0.000247463, 0.000326211]$	0/0/16

Fitted lowers V and increases C, B, E against CFG in all sixteen cases. Mean changes $\Delta B = 2.33285 \times 10^{-4}$ and $\Delta D = -6.38764 \times 10^{-8}$ give $\Delta E = 2.33221 \times 10^{-4}$. Against the sign cap, fitted CFG has 3/16 response wins and larger center error in all sixteen; its mean response difference, 1.31335×10^{-4}, has interval $[-5.68735 \times 10^{-5}, 2.75504 \times 10^{-4}]$, while the table's positive interval is for the median.

Reference and perturbation sensitivity. Let ρ be the maximum response gap to the preceding Heun and RK4 references. The 2ρ contrast guard resolves fifteen fitted wins against CFG, leaving one unresolved; against cap it resolves five losses, leaving eleven unresolved. All cases remain in primary summaries. Previous-Heun/RK4 references and $\epsilon = 0.001$ preserve the 16/16 and 3/16 response-win counts; at the smaller scale three cap-relative losses are resolved. No candidate-dependent refinement is added; the error of the anchor trajectory itself is not quantified.

Magnitude and direction (post hoc). For response norms $M = \|q_{\text{method}}\|$, $R = \|q_{\text{ref}}\|$ and angle θ, $V^2 = (M - R)^2 + 2MR(1 - \cos\theta)$. All cases have $R > 2\rho$; empirical norm-sensitivity intervals put fitted, cap and Gaussian-b=10 below the reference norm in all sixteen, and CFG below in fourteen. Fitted improves directional agreement against CFG in all cases, while $|M - R|$ increases in fourteen. Mean squared-magnitude and directional changes are -2.7062×10^{-5} and -0.0046601, summing to $\Delta V^2 = -0.0046872$; median $\Delta|M - R| = 0.001916$. Counts persist at the smaller scale and alternative references. Against cap, fitted improves both components in three cases. These descriptive decompositions establish the direction/magnitude distinction; the 2ρ guard applies to V, not its individual components.

D Generation and controls

D.1 CIFAR grid robustness

Table 14 reports the primary grid's finite-sample estimates. An independently seeded 5k grid reproduces the counts: $9/9$ lower FID, $9/9$ non-increased clipping p95 and $8/9$ lower amp p95 for fitted CFG. KID and classifier-proxy conditionality are mixed. The terminal interval wins FID at $N = 16, 32$ and loses at $N = 8$, with lower target accuracy than fitted throughout. The 50k recipe is in Appendix A.

Table 14: Primary 5k CIFAR grid. Each sampler includes its terminal $r = 0$ step. FID and target-accuracy entries are CFG/Fitted CFG/terminal interval; the other entries are CFG/Fitted CFG.

g	N	FID ↓	KID ↓	amp p95 ↓	clip p95 ↓	target acc (%) ↑
5	8	27.06 / 25.39 / 26.78	0.0104 / 0.0173	1.09 / 1.05	0.165 / 0	99.14 / 94.66 / 89.38
5	16	21.58 / 16.15 / 13.07	0.0069 / 0.0068	1.12 / 1.09	0.226 / 0.005	99.22 / 99.18 / 94.08
5	32	20.15 / 15.90 / 10.70	0.0061 / 0.0050	1.04 / 1.07	0.261 / 0.068	99.00 / 99.36 / 96.94
7.5	8	25.66 / 25.60 / 27.07	0.0058 / 0.0176	7.83 / 1.41	0.387 / 0	98.04 / 94.56 / 89.04
7.5	16	24.36 / 17.21 / 12.78	0.0068 / 0.0072	1.94 / 1.42	0.479 / 0.007	96.98 / 99.16 / 94.92
7.5	32	27.20 / 17.84 / 8.84	0.0103 / 0.0057	2.29 / 1.40	0.565 / 0.104	95.86 / 99.46 / 95.22
9	8	32.45 / 25.22 / 26.56	0.0108 / 0.0176	13.31 / 1.64	0.573 / 0	95.42 / 94.84 / 89.58
9	16	30.42 / 17.06 / 12.77	0.0113 / 0.0070	6.97 / 1.67	0.650 / 0.006	94.50 / 99.26 / 94.38
9	32	36.69 / 18.67 / 8.63	0.0179 / 0.0059	5.68 / 1.65	0.734 / 0.128	90.38 / 99.54 / 94.88

D.2 SDXL GENERATION AND GENEVAL2

Table 15 compares fitted CFG and CFG within each fixed-size SDXL block and reference protocol.

Table 15: **SDXL comparison at matched update count.** Hard: 30k images, $g = 12, N = 12$; ordinary: 10k, $g = 5, N = 30$. Saturation is measured on decoded RGB. Dashes denote unavailable metrics.

Setting	Method	FID ↓	KID ↓	CLIP ↑	Sat. p95
Hard	CFG	19.489	0.00683	0.3197	0.03583
	Fitted CFG	23.178	0.00843	0.3110	0.0000715
	Tuned interval	20.144	0.00749	0.3154	0.01873
Ordinary	CFG	20.578	0.00591	–	–
	Fitted CFG	20.596	0.00585	–	–

Implementation and data. SDXL-base-1.0 uses 1024×1024, float16 deterministic DDIM and no negative prompt. The shared U-Net takes conditional/unconditional text embeddings with the conversion in Appendix A. Fixed COCO train2017 captions are paired with reference images and recorded seeds. For FID/KID, real images use 256-pixel LANCZOS `ImageOps.fit`; generated images are saved at 256 pixels by antialiased bilinear resizing and uint8 rounding. All methods share this metric pipeline. Saturation is the decoded RGB fraction below 0.01 or above 0.99. The ordinary-cell interval has FID 20.578 and KID 0.00594, with CLIP/saturation unmeasured.

GenEval2. All 800 official prompts use published order, seeds $40{,}000{,}000$–$40{,}000{,}799$ and the official Qwen Soft-TIFA evaluator. We average each prompt's question probabilities arithmetically/geometrically, then average prompts equally and multiply by 100 (6,012 questions; zero probability gives zero geometric mean). The first eight CFG outputs reproduce the stored baseline pixels exactly. Evaluator excursions above one (maximum 1.4×10^{-7}) are clipped within a fixed 10^{-6} tolerance; larger violations are rejected. Recomputed AM/GM match exactly. Guidance-step norm p95 is the sample 95th percentile of $\max_{0 \le n < N} \|\alpha_n(D_\mathrm{u} - D_\mathrm{c})\|_2$, including closure, in the additive-noise coordinates of Appendix A. The interval uses $\alpha_n = 0$ when $h_n > \log(1 + 1/w)$, including conditional closure; on this hard-cell mesh all finite-step corrections are disabled.

Table 16: GenEval2 at the SDXL hard cell. Soft-TIFA AM/GM are higher-is-better; the last two columns are sampling diagnostics.

scheme	Soft-TIFA AM ↑	Soft-TIFA GM ↑	saturation p95 ↓	Guidance-step norm p95
CFG	**50.65**	**7.98**	0.03327	161.31
Fitted CFG	44.82	6.31	0.00008	27.81
interval	27.36	3.16	**0.00000**	0.00

Late switching. A switched rule applies CFG on positive intervals $n < j$ and fitted CFG afterwards. On 512 calibration prompts with fresh paired seeds ($g = 12, N = 12$; FP16 networks, FP32 states and update), switches at $j \in \{3, 5, 7\}$ (noise levels 2.90712, 1.61289, 0.93236) interpolate between the two rules: mean CLIP rises toward CFG and saturation p95 from 0.0141 to 0.0282 (Table 17). The prespecified rule required CLIP within 0.002 below each CFG baseline and at least 0.002 above fitted CFG with D_c closure (met by all), and saturation p95 at most 0.75 times CFG with D_c closure; switches 3 and 5 meet the saturation target and switch 7 the CLIP target, and none meets all three criteria.

Table 17: SDXL switch calibration, $g = 12, N = 12$: 512 paired cases, 2,560 trajectories, 5,120 outputs from two closures. Shown: conditional closures and native CFG. p95 aggregates image fractions; clipping is pre-clamp.

Positive-step rule	Closure	Mean CLIP	Saturation p95	Clipping p95
CFG	D_{c}	0.32179700	0.03378800	0.00420012
CFG	D_w	0.32193384	0.03383667	0.00421139
Fitted CFG	D_{c}	0.31341063	0.00010414	0.00002236
CFG then fitted, $j = 3$	D_{c}	0.31798009	0.01411327	0.00149166
CFG then fitted, $j = 5$	D_{c}	0.31971593	0.02288125	0.00194958
CFG then fitted, $j = 7$	D_{c}	0.32075017	0.02818497	0.00247345

D.3 Terminal fitting and crossover comparisons

The subsequent study fixes $g = 12, N = 12$, FP32 state/update arithmetic, one fresh block of 5,000 paired COCO caption–seed cases and a 10,000-image real bank. CF7 uses seven CFG intervals then four fitted intervals and D_{c} closure. CS7 uses the sign cap on the same suffix; B10 uses the calibrated $\hat{b} = 10$ coefficient (Appendix D.7). CF7 is primary; CS7/B10 are secondary, fixed before held-out generation.

Each trajectory uses 24 branch-equivalent evaluations. Paired CFG closures give five displayed rows from four trajectories per case. Real images use 256-pixel LANCZOS `ImageOps.fit`; generated images use bilinear antialiased 256-pixel resizing before uint8 rounding. CLIP/pixel diagnostics use decoded 1024-pixel images; clipping is pre-clamp. FID is a point estimate; KID uses the full-set unbiased polynomial kernel. Comparisons are within block.

Table 18: Fixed-guidance SDXL comparison: 5k paired cases per row, common 10k real bank, nominal $g = 12$. These rows isolate coefficient and closure choices at the prescribed guidance. Saturation and clipping columns are p95 of per-image fractions.

Rule	Closure	FID ↓	KID ↓	CLIP ↑	Sat. p95	Clip p95
CFG, $g = 12$	D_w	27.590	0.006988	0.320152	0.035251	0.004000
CFG, $g = 12$	D_{c}	27.815	0.007117	0.319976	0.034971	0.003934
CF7	D_{c}	28.176	0.007392	0.319464	0.028462	0.002848
CS7	D_{c}	28.171	0.007388	0.319467	0.028519	0.002864
B10	D_{c}	27.179	0.006719	0.317151	0.002574	0.000318

CF7 reduces saturation/clipping p95 by $19.3\%/28.8\%$ versus native CFG, with CLIP difference -0.000688. Higher FID/KID leave the combined criterion unmet under either closure. CF7/CS7 nearly coincide; B10 improves FID/KID with lower CLIP.

Post-hoc common preprocessing. Reprocessing the same 10k original real images with the generated-image function matches direct-square bilinear antialiased resizing and uint8 rounding. Generated outputs remain fixed. In Table 18 row order, FID is 24.980, 25.208, 25.341, 25.339, 24.148; KID is 0.005172, 0.005294, 0.005418, 0.005416, 0.004784. Crossover retains lower FID/KID than both CFG closures. This check changes reference cropping and interpolation together; CLIP, saturation and clipping retain the table's values.

On the first 512 held-out cases, full fitted and reverse-order fitting (first seven intervals fitted, last four CFG) have mean CLIP 0.312126 and 0.314732, compared with CF7's 0.320472 and matched-closure CFG's 0.320815. The CF7-minus-full-fitted paired difference is 0.008346 with a descriptive 2.5th–97.5th bootstrap percentile range $[0.006428, 0.010273]$. Delaying fitting improves alignment here, with fitted-step exposure also changing. This single block separates alignment recovery from distribution quality, with the real bank fixed.

Qualitative example protocol. Figure 4 uses Stable Diffusion 1.5 at 512×512, float16 deterministic DDIM, $g = 12, N = 12$, and no negative prompt. The three methods share the COCO val2017

caption, initial latent, noise schedule and seed. Selection ranks the first 64 prompts by CFG-minus-fitted decoded saturation. The example illustrates the contrast change; the aggregate SD1.5 study, with lower saturation and lower CLIP than CFG, is kept in the reproducibility archive.

D.4 CIFAR-10 GUIDANCE-OFF INTERVAL CONTROLS

The terminal-interval baseline disables guidance when $h_n > \log(1 + 1/w)$ and uses conditional closure. On $(g, N) = (9, 8)$ this disables all finite-step guidance. For the tuned baseline, a guidance-off band is selected separately for each setting on 1,024 validation samples and evaluated on a fresh 5k block. CFG is active outside the band; its correction is zero at left endpoints inside the closed band. The nine bands are $[0.002, u]$ for $u \in \{1, 2, 4, 8, 16, 80\}$, and $[1, 8]$, $[2, 16]$, $[4, 32]$. Eligible candidates have target accuracy and mean target probability each no more than 0.02 below fitted CFG (two percentage points for accuracy). We select the lexicographic minimum of clipping p95, amp p95, FID, KID, then off-step count. The selected bands are $[0.002, 16]$ for both $N = 8$ settings and $[0.002, 8]$ for $g = 9, N = 16$. The tuned interval has higher FID than fitted in these settings. The terminal interval wins FID at $g = 9, N = 16$ with lower target accuracy. This classifier proxy can also reflect reduced diversity, limiting its fidelity interpretation.

Table 19: Guidance-off interval controls: 5k images per setting. "tuned int." uses the conditionality floors above; "terminal int." uses the step threshold. Clipping is measured before clamping; target accuracy is in percent.

setting	scheme	FID ↓	KID ↓	amp p95 ↓	clip p95 ↓	target acc ↑
g=9, N=8	CFG	32.23	0.0105	13.31	0.574	95.5
	Fitted CFG	25.70	0.0180	1.66	0	94.9
	tuned int.	34.25	0.0205	1.71	0.018	95.3
	terminal int.	26.68	0.0225	1.63	0	89.1
g=9, N=16	CFG	30.82	0.0117	7.30	0.664	93.8
	Fitted CFG	17.12	0.0070	1.66	0.006	99.3
	tuned int.	30.18	0.0120	1.75	0.257	98.8
	terminal int.	12.63	0.0072	1.66	0.000	94.8
g=7.5, N=8	CFG	25.38	0.0057	8.08	0.395	98.4
	Fitted CFG	25.36	0.0175	1.40	0	94.8
	tuned int.	32.60	0.0208	1.43	0.002	94.7
	terminal int.	26.57	0.0216	1.38	0	89.8

D.5 CIFAR-10 CLOSE COEFFICIENT CONTROLS

We compare CFG, fitted CFG and the sign cap at $g = 9, N = 8$ on a fresh balanced block of 5,000 paired seeds, using the checkpoints, mesh, FP32 arithmetic and image conversion of the closure factorial. The sign-cap coefficient is $\max\{w(r - 1), -r\}$. An independent balanced 1k calibration block compares fitted CFG with constant-scale CFG at $g \in \{1.5, 2, 3, 5\}$. Among candidates whose classifier accuracy is within one percentage point below fitted CFG, the locked rule selects minimum point FID (differences within 0.01 are tied by accuracy, then smaller g). All candidates qualify; $g = 2$ is selected before formal generation. Selection and primary comparisons use conditional closure. Each finite trajectory also supplies a guided closure from the same final positive-noise state, giving 40,000 outputs from 20,000 formal trajectories.

Table 20: CIFAR-10 close controls: 5k paired cases per row, $N = 8$. KID averages 100 shared size-1,000 subsets against the existing 8,192-image real-feature bank; FID uses the common CIFAR reference statistics. Accuracy is target-label accuracy. Both closure choices are retained.

Rule	Closure	FID ↓	KID ↓	Acc. (%)	amp p95	Clip p95
CFG, $g = 9$	D_c	32.559	0.011160	95.04	13.413	0.58216
Fitted, $g = 9$	D_c	24.733	0.017289	94.92	1.658	0
Sign cap, $g = 9$	D_c	24.735	0.017290	95.00	1.658	0
Calibrated CFG, $g = 2$	D_c	24.430	0.013927	98.04	1.000	0
CFG, $g = 9$	D_w	32.559	0.011166	95.00	13.413	0.58208
Fitted, $g = 9$	D_w	24.166	0.016295	94.88	1.658	0
Sign cap, $g = 9$	D_w	24.164	0.016293	94.90	1.658	0
Calibrated CFG, $g = 2$	D_w	24.423	0.013920	98.08	1.000	0

At $g = 9$, the cap is active on all seven positive intervals, so its coefficient differs from fitted CFG by exactly r_n^9. The corresponding conditional-closure images differ by at most one uint8 level, averaging 0.00304 levels per component. Fitted-minus-cap point FID is -0.00134 and KID is -1.16×10^{-6}. The paired accuracy difference is -0.08 percentage points, with a descriptive 95% class-stratified bootstrap interval $[-0.18, -0.0195]$; the confidence-difference interval includes zero. Both rules strongly reduce residual amplification and clipping relative to CFG at $g = 9$.

With conditional closure, calibrated CFG at $g = 2$ has lower point FID/KID and higher accuracy than fitted CFG. Under guided closure, fitted CFG has slightly lower FID, while calibrated CFG retains lower KID and higher accuracy. These are generation comparisons between different guidance fields. Each arm's amp uses its own w and therefore measures amplification relative to that field's initial residual. Per-image comparisons use 2,000 paired, class-stratified bootstrap draws; KID subset ranges describe resampling stability. The model's relative-contraction guarantee supplies the accuracy criterion for selecting fitted CFG at a prescribed guidance.

D.6 SDXL COEFFICIENT-CAP CONTROLS

A 1k-prompt SDXL control at $g = 12, N = 12$ with paired prompts and seeds compares the caps with fitted CFG on CLIP and decoded saturation.

Table 21: SDXL coefficient controls, 1k paired prompts: mean CLIP and decoded-pixel saturation p95.

Method	CLIP mean	Saturation p95
CFG	0.32069	0.04090
Fitted CFG	0.31099	0.00007342
Sign cap	0.31147	0.00007254
Residual cap	0.31788	0.00073349

The sign cap and fitted CFG reach similar saturation, with slightly higher mean CLIP for the cap. The residual cap retains the highest mean CLIP of the three, with saturation p95 about ten times theirs and still more than fifty times below CFG.

D.7 FINITE-SCALE CROSSOVER CALIBRATION AND EVALUATION

Section 3.4 derives the Gaussian-shell coefficient (17), exact on $(a, b) = (0, \hat{b})$, with conditional closure. The following calibration selects its variance scale. On a disjoint 512-prompt SDXL block, the prespecified rule considered $\hat{b} \in \{0.01, 0.1, 1, 10, 100\}$, required saturation p95 below 0.0204, then maximized mean CLIP. Only 10 and 100 were eligible; their CLIP values were 0.3185 and 0.3148. These calibration-block values select $\hat{b} = 10$, which is frozen for the subsequent evaluations.

The frozen 30k comparison is reported in Table 3. Its three consecutive 10k subsets were fixed before evaluation; crossover has lower FID and KID in each (Table 22). These subsets partition the same generated block.

Table 22: Prespecified 10k subsets of the SDXL 30k block. Entries are CFG/crossover; subset numbers follow the fixed prompt order. Each row uses its corresponding real-image subset and gives an ordering check.

Subset	FID ↓	KID ↓
1	23.387 / 22.748	0.006776 / 0.006521
2	23.589 / 23.221	0.006847 / 0.006628
3	23.769 / 23.421	0.006842 / 0.006621

The crossover GenEval2 AM/GM are $47.94/8.13$. However, it repairs 0/12 verified reversals versus fitted's 11/12: the original combined replication-and-crossover-repair criterion remains unmet. The independent local study and subsequent 5k quality comparison prespecify it as a secondary comparator.

The 30k crossover-minus-CFG KID difference is negative on 92/100 shared subsets, with 5th–95th percentiles $[-4.79\times10^{-4}, 2.52\times10^{-5}]$ crossing zero. These percentiles describe subset-resampling stability within the fixed generated set. The scale search, separate 5k confirmation and mixed residual-cap comparison remain in the archive.

D.8 Finite-step and terminal-closure factorial experiment

The frozen protocol uses $g = 9$ ($w = 8$), $N \in \{8, 16\}$ and two independent balanced 5k blocks per N. CFG/fitted trajectories each share their last denoiser pair across D_w/D_c closures: 20,000 seed cases, 40,000 trajectories, 80,000 outputs. Disjoint blocks start at 90,000,000, 90,005,000, 90,010,000 and 90,015,000; smoke seeds are separate. All 80 shards/output identities pass audit. The classifier's held-out accuracy is 0.9437 in each run.

Primary endpoints are preterminal amp median/p95, pre-clamp clipping median/p95 and target accuracy/confidence. We use 2,000 class-stratified paired bootstrap draws within each block, retaining all four outputs per seed. Contrasts cover finite rules at fixed closure, closures at fixed finite rule and their interaction. Auxiliary FID is a point estimate; 100 shared size-1,000 KID subsets measure resampling stability.

Table 23: Complete factorial summaries; each row summarizes 5,000 outputs. Amp and clipping entries are median / p95; conditioning entries are accuracy (%) / target confidence.

N	Block	Finite	Closure	amp	Clipping	Conditioning	FID / KID
8	1	CFG	D_w	2.264/13.314	0.150/0.558	94.70/0.9419	32.29/0.01125
8	2	CFG	D_w	2.239/13.239	0.147/0.564	95.22/0.9475	32.27/0.01083
8	1	CFG	D_c	2.264/13.314	0.150/0.559	94.72/0.9420	32.29/0.01124
8	2	CFG	D_c	2.239/13.239	0.147/0.564	95.22/0.9477	32.26/0.01081
8	1	Fitted CFG	D_w	1.348/1.639	0.000/0.000	94.70/0.9428	24.63/0.01654
8	2	Fitted CFG	D_w	1.348/1.645	0.000/0.000	95.22/0.9484	25.50/0.01734
8	1	Fitted CFG	D_c	1.348/1.639	0.000/0.000	94.84/0.9435	25.20/0.01756
8	2	Fitted CFG	D_c	1.348/1.645	0.000/0.000	95.26/0.9481	26.09/0.01835
16	1	CFG	D_w	1.607/6.494	0.207/0.650	94.14/0.9350	30.75/0.01165
16	2	CFG	D_w	1.629/7.641	0.214/0.647	94.30/0.9370	31.11/0.01152
16	1	CFG	D_c	1.607/6.494	0.207/0.648	94.20/0.9352	30.79/0.01165
16	2	CFG	D_c	1.629/7.641	0.213/0.646	94.26/0.9373	31.15/0.01154
16	1	Fitted CFG	D_w	1.407/1.670	0.000/0.009	99.34/0.9930	17.01/0.00709
16	2	Fitted CFG	D_w	1.404/1.660	0.000/0.008	99.12/0.9900	16.63/0.00654
16	1	Fitted CFG	D_c	1.407/1.670	0.000/0.007	99.36/0.9928	17.20/0.00734
16	2	Fitted CFG	D_c	1.404/1.660	0.000/0.007	99.10/0.9897	16.86/0.00682

Table 24 reports finite-step contrasts. Clipping-p95 intervals are negative throughout. Accuracy intervals cross zero at $N = 8$ and show 4.82–5.20 percentage-point gains at $N = 16$. KID is worse for fitted at $N = 8$ and better at $N = 16$, under either closure.

Changing from D_w to D_c on fitted trajectories leaves preterminal amp unchanged, gives no clipping-p95 change at $N = 8$, and reduces it by 0.00195 and 0.00163 at $N = 16$ (paired intervals $[-0.00260, -0.00098]$ and $[-0.00262, -0.00098]$). Accuracy effects are -0.02 to 0.14 percentage points; all closure and interaction intervals for accuracy cross zero, as do clipping-p95 interaction intervals. Complete estimates are in `effects.csv`. Conditional closure raises fitted's point FID by $0.57/0.58$ at $N = 8$ and $0.19/0.23$ at $N = 16$. Float64 recomputation differs from raw float32 quantiles by less than 10^{-5}, leaving displayed precision unchanged.

Table 24: Fitted-CFG-minus-CFG finite-step effects at each fixed closure. Brackets are paired, class-stratified 95% bootstrap intervals; accuracy differences are percentage points.

N	Block	Closure	Δ amp p95	Δ clip p95	Δ accuracy
8	1	D_w	-11.68 $[-11.90, -11.47]$	-0.558 $[-0.576, -0.543]$	0.00 $[-0.88, 0.82]$
8	1	D_c	-11.68 $[-11.90, -11.47]$	-0.559 $[-0.576, -0.542]$	0.12 $[-0.76, 0.92]$
8	2	D_w	-11.59 $[-11.79, -11.42]$	-0.564 $[-0.580, -0.551]$	0.00 $[-0.78, 0.80]$
8	2	D_c	-11.59 $[-11.79, -11.42]$	-0.564 $[-0.580, -0.551]$	0.04 $[-0.74, 0.86]$
16	1	D_w	-4.82 $[-5.30, -4.37]$	-0.641 $[-0.656, -0.624]$	5.20 $[4.50, 5.88]$
16	1	D_c	-4.82 $[-5.30, -4.37]$	-0.641 $[-0.658, -0.626]$	5.16 $[4.50, 5.84]$
16	2	D_w	-5.98 $[-6.81, -5.06]$	-0.639 $[-0.658, -0.624]$	4.82 $[4.16, 5.50]$
16	2	D_c	-5.98 $[-6.81, -5.06]$	-0.640 $[-0.659, -0.625]$	4.84 $[4.18, 5.52]$

D.9 Comparison with a 512-step CFG numerical reference

The CIFAR pair uses $g = 9$, $N = 8, 16, 512$ samples each (starting seeds 109000, 110000), cycling labels and paired initial states. Karras meshes use $\rho = 7$, endpoints $80, 0.002$ and a final zero step. Coarse CFG and the 512-update reference return D_w; fitted returns D_c, so errors include closure differences. The reference has no further convergence check.

Per-sample relative L_2 is $\|z - z_{\text{ref}}\|_2 / \max\{\|z_{\text{ref}}\|_2, 10^{-30}\}$; images use $z = \text{clip}(127.5x + 128, 0, 255)/255$ without rounding. Residuals are measured before closure at $\sigma = 0.002$. Columns divide each fitted statistic by its CFG counterpart (denominators floored at 10^{-30}); the first column is a ratio of medians. Fitted reduces amp while increasing image error in both cells:

setting	image rel-L_2 med	image rel-L_2 p95	residual rel-L_2 p95	amp p95
g=9, N=8	1.155	1.074	0.277	0.122
g=9, N=16	1.290	1.160	1.021	0.257

E Implementation and record index

`STUDY_INDEX.md` links results to protocols, frozen sources and manifests. Table 25 summarizes their execution settings; Appendix A gives the separate reference-kernel audit.

Table 25: **Execution map.** Precisions are bits for persistent state, returned predictions and update accumulation. Add/Fact evaluate the left/right forms of (4); P2 factors its fitted intervals. Native closure is D_w for CFG and D_c for fitted/interval/crossover. Local studies end at positive noise.

Experiment	Kernel	State	Pred.	Update	Closure
CIFAR (Fig. 3, Table 2; D.1)	Add	32	32	32	Native
CIFAR factorial / controls (D.8, D.5)	Add	32	32	32	Both
SDXL 30k / GenEval2 (Tables 3, 16)	Add	16	16	32	Native
SD1.5 illustration (Fig. 4)	Add	16	16	16	Native
SDXL 5k P2 (Table 18)	Add/Fact	32	16	32	Specified
SDXL local FP32 (Table 4; C.4)	Add	64	32	64	–
SDXL production FP16 checks (C.1)	Add	16	16	32	–