IMPERFECT VERIFIERS, EXACT TARGETS: INFERENCE-TIME SCALING FOR CONTINUOUS GENERATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Verifier-guided inference-time scaling uses additional computation to score intermediate states and steer generation toward better outputs. However, imperfect intermediate verifiers can distort sampling, and repeated refinement alone does not ensure recovery of the desired reward-tilted distribution. We introduce Continuous-State Verifier-Guided Backtracking (CS-VGB), an inference-time scaling method for sampling from reward-tilted distributions with continuous generative models through repeated denoising and renoising. With correctly paired denoising and renoising transitions, the stationary distribution of CS-VGB at the clean checkpoint equals the desired reward-tilted target, even when intermediate verifiers are inaccurate. Under stated regularity conditions, CS-VGB converges exponentially to the target despite verifier error, and the expected number of updates needed to reach a clean sample grows only linearly with the number of checkpoints. In analytically tractable settings, we empirically demonstrate that CS-VGB recovers the desired reward-tilted distribution despite verifier error, whereas the compared verifier-based sampling methods fail to recover the target under the same verifier errors. We also instantiate the framework with diffusion models and one-step stochastic posterior maps. Across prompt-guided ImageNet generation, facial-attribute generation on CelebA, and multi-robot planning, recurrent CS-VGB search improves reward and task success as inference budgets grow, with image-generation rewards continuing beyond the observed plateaus of denoising-only baselines.

1 INTRODUCTION

Inference-time scaling seeks to improve the outputs of pretrained generative models by spending additional computation at sampling time. For generation in continuous spaces, such as images and robot trajectories, this computation can be allocated to Monte Carlo estimation of guidance gradients (Song et al., 2023; Potaptchik et al., 2026; Holderrieth et al., 2026), larger candidate sets for soft value-based decoding (SVDD) (Li et al., 2025), or more particles with intermediate weighting and resampling via sequential Monte Carlo (SMC) (Wu et al., 2023a; Singhal et al., 2025).

Despite these different uses of computation, these methods follow a one-way denoising trajectory and favor intermediate states that appear promising before the terminal reward is available, often relying on a *process verifier*, or *intermediate verifier*, to estimate the reward of future completions. However, errors in these estimates can lead the sampler to discard promising continuations and spend subsequent computation on less favorable ones, while the one-way trajectory prevents it from revisiting earlier noise levels to reconsider those choices.

To address this limitation, recent inference-time scaling methods use renoising to return to earlier noise levels and explore alternative denoising continuations, as in depth-first search (DFS) (Zhang et al., 2026), reward-guided iterative refinement (RERD) (Uehara et al., 2025), and cyclic diffusion search (ABCD) (Lee et al., 2025a). These approaches enable bidirectional exploration across noise levels, but it remains unclear how to preserve a prescribed target distribution despite errors in the process verifier. We therefore seek the *reward-tilted distribution* $\pi_0^\star(dx) \propto r(x)\pi_0(dx)$, where π_0 is the model's clean distribution and r is a nonnegative terminal reward, and ask:

Q Can we construct a bidirectional sampler for continuous-state generation that revisits earlier decisions while preserving the desired reward-tilted target despite errors in process verifiers?

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.

Recent work in discrete generation addresses this question through verifier-guided backtracking: AR-VGB (Rohatgi et al., 2026) and MDM-VGB (Jeon et al., 2026) separate the role of imperfect process verifiers from the target distribution in autoregressive and masked diffusion models, respectively. Motivated by these results, we rigorously extend this sampling principle to continuous generative models through **CS-VGB**, an outer Markov chain Monte Carlo (MCMC) framework for models that admit denoising and renoising transitions between checkpoints. CS-VGB turns a pretrained model's one-way sampler into a bidirectional, recurrent sampler that revisits earlier decisions by alternating denoising and renoising. In this construction, the process verifier guides exploration, while the exact terminal reward determines the stationary distribution at the clean checkpoint.

Our contributions are threefold:

- **CS-VGB with imperfect process verifiers.** We establish that, with correctly paired denoising and renoising transitions, the stationary clean distribution is $\pi_0^\star$ regardless of process-verifier accuracy. To retain this guarantee in practical finite-candidate implementations, we introduce *reverse insertion*, together with *lazy pool evaluation* to reduce the number of process-verifier evaluations without changing the stationary clean target.
- **Convergence guarantees and first-hitting times.** Under stated assumptions, CS-VGB converges exponentially fast despite process-verifier error, while *directional persistence* enables a first clean visit in an expected number of updates linear in the checkpoint count. With sufficiently small transition mismatch, convergence remains still exponential to a potentially biased target, whose bias vanishes as the transition mismatch tends to zero.
- **Inference-time sampling and search.** In analytically tractable settings, we empirically demonstrate that CS-VGB recovers the reward-tilted target despite process-verifier error, whereas the compared verifier-based sampling methods deviate from this target. Building on the CS-VGB framework, we also propose a finite-budget search algorithm for diffusion models and one-step posterior maps (Potaptchik et al., 2026; Holderrieth et al., 2026), demonstrating that recurrent sampling can improve performance beyond the plateaus of denoising-only methods.

We discuss additional related work on reward-guided sampling and search, verifier-guided backtracking, and stochastic transitions and posterior maps in Appendix E.

2 CS-VGB: Recurrent Sampling and Search

We consider inference-time scaling methods for sampling from the *reward-tilted distribution* of a pretrained generator with clean law π_0 and terminal reward $r : \mathcal{X} \to [0, \infty)$, where $\mathcal{X} \subset \mathbb{R}^d$:

$$\pi_0^\star(dx) := \frac{r(x)\pi_0(dx)}{Z}, \qquad Z := \mathbb{E}_{X_0 \sim \pi_0}[r(X_0)] \in (0, \infty). \tag{1}$$

Checkpoints and transition kernels. On a fine time grid $0 = t_0 < \cdots < t_T = 1$, reference generation starts from $X_{t_T} \sim \pi_{t_T} = \mathcal{N}(0, I)$ and follows the pretrained denoising process toward $t_0 = 0$. Along this process, we select *checkpoints* t_{κ_j} with $0 = \kappa_0 < \cdots < \kappa_J \leq T$ and denote the corresponding reference states by $X_j := X_{t_{\kappa_j}}$. These checkpoints are the noise levels at which CS-VGB chooses whether to continue denoising or return to a noisier state. Between adjacent checkpoints, the pretrained generator induces the *denoising kernel* $K_j^{\mathrm{F}}(dy \mid x)$, the conditional law of X_{j-1} given $X_j = x$, while the *renoising kernel* $K_{j-1}^{\mathrm{B}}(dx \mid y)$ proposes a move from checkpoint $j-1$ back to j. Together, the starting law π_J and the denoising kernels determine the reference marginals $\pi_j = \mathrm{Law}(X_j)$ through

$$\pi_{j-1} = \pi_j K_j^{\mathrm{F}}, \qquad j = J, \ldots, 1, \tag{2}$$

where π_J is the reference law at the starting checkpoint and equals $\mathcal{N}(0, I)$ when $\kappa_J = T$. We instantiate these denoising and renoising kernels for diffusion models (Ho et al., 2020; Song et al., 2021) and for two one-step posterior maps: Meta Flow Maps (MFM) (Potaptchik et al., 2026) and Posterior Diamond Maps (PDM) (Holderrieth et al., 2026). The resulting checkpoint transitions are summarized in Table 1.

Process verifier and sampling goal. We assume access to a *process verifier* $\widehat{V}_j(x)$ that estimates the reward of future completions from state x at checkpoint j, whose ideal value is $V_j(x) := \mathbb{E}[r(X_0) \mid X_j = x]$ under the reference denoising process. While $\widehat{V}_j$ may differ from V_j at noisy checkpoints, we assume that $\widehat{V}_0 = r$ is evaluated exactly. Our goal is to construct an outer MCMC kernel P that uses these estimates to select denoising and renoising moves while preserving $\Pi(dx \mid j = 0) = \pi_0^\star(dx)$ under its stationary law Π, without requiring $\widehat{V}_j = V_j$.

Table 1: **Stochastic checkpoint kernels for CS-VGB.** Here $t = t_{\kappa_j}$ and $s = t_{\kappa_{j-1}}$. For the VP SDE with noise-rate schedule β, $\alpha_{t|s} := \exp(-\frac{1}{2}\int_s^t \beta(u)\,du)$ and $\sigma^2_{t|s} := 1 - \alpha^2_{t|s}$. For the linear-interpolant constructions, $a_j := (1-t)/(1-s)$ and $\sigma_j^2 := t^2 - a_j^2 s^2$. Full transitions and endpoint conventions are given in Appendices A.1 and B.4.

Model	Denoising: $K_j^{\mathrm{F}}(dy \mid x)$	Renoising: $K_{j-1}^{\mathrm{B}}(dx \mid y)$
Diffusion models (VP SDE)	$K_j^{\mathrm{F}} = k^{\mathrm{F}}_{\kappa_j} \cdots k^{\mathrm{F}}_{\kappa_{j-1}+1}$, k_k^{F}: reverse diffusion step (Equation (B.23))	$\mathcal{N}(dx; \alpha_{t\|s} y, \sigma^2_{t\|s} I)$
One-step posterior map (MFM/PDM)	$K_j^{\mathrm{F}}(dy \mid x) = \mathrm{Law}(Y \in dy \mid x)$, Y: Gaussian bridge sample (Figure B.6)	$\mathcal{N}(dx; a_j y, \sigma_j^2 I)$

2.1 Verifier-Guided Sampling and Search Framework

Edge-conductance-based construction. To favor states with higher verifier values while preserving the reward-tilted clean target, we construct denoising and renoising transitions from a common edge-conductance measure. Following the weighted-network construction of random walks and its continuous-state counterpart (Doyle & Snell, 1984; Bezuglyi & Jorgensen, 2018), we weight adjacent state pairs and normalize at either endpoint to obtain the two transitions. In this case, the interior edges use process verifiers, whereas the clean edge uses only the terminal reward. Together with the action rule below, this construction preserves the stationary clean target under the exact-kernel conditions of Theorem 2.1.

We first compare the joint distributions of adjacent states (x, y) obtained by denoising from $x \sim \pi_j$ or renoising from $y \sim \pi_{j-1}$, which we call the *edge laws*:

$$\begin{aligned} J_j^{\mathrm{F}}(dx, dy) &:= \pi_j(dx) K_j^{\mathrm{F}}(dy \mid x), \\ J_j^{\mathrm{B}}(dx, dy) &:= \pi_{j-1}(dy) K_{j-1}^{\mathrm{B}}(dx \mid y). \end{aligned} \tag{3}$$

To make the two joint distributions agree, we reweight K_{j-1}^{B} by the *edge ratio* $\rho_j := dJ_j^{\mathrm{F}}/dJ_j^{\mathrm{B}}$, giving $K_{j-1}^{\mathrm{B},\rho} := \rho_j K_{j-1}^{\mathrm{B}}$ and

$$\pi_j(dx) K_j^{\mathrm{F}}(dy \mid x) = \pi_{j-1}(dy) K_{j-1}^{\mathrm{B},\rho}(dx \mid y). \tag{4}$$

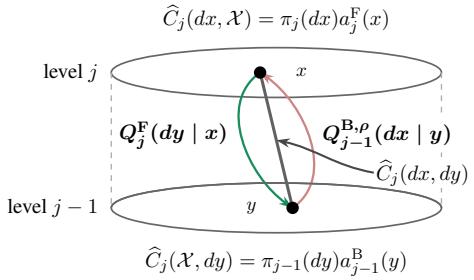

Figure 1: Normalizing $\widehat{C}_j$ gives both sampling transitions, while its endpoint marginals determine the directional masses.

When the joint distributions already coincide, $\rho_j = 1$ and no correction is needed, since $K_{j-1}^{\mathrm{B}}(dx \mid y)$ already gives the conditional law of X_j given $X_{j-1} = y$ under reference denoising. We then weight this common edge law to define the *edge-conductance measure*

$$\widehat{C}_j(dx, dy) := \pi_j(dx) K_j^{\mathrm{F}}(dy \mid x) \begin{cases} r(y), & j = 1, \\ \sqrt{\widehat{V}_j(x) \widehat{V}_{j-1}(y)}, & 2 \le j \le J. \end{cases} \tag{5}$$

Integrating out the noisy endpoint x of the clean-level ($j = 1$) edge gives $\widehat{C}_1(\mathcal{X}, dy) = r(y)\pi_0(dy)$, whose normalization is exactly $\pi_0^\star$, regardless of process verifier errors. At an interior state x, integrating over possible destinations at $j-1$ and $j+1$ similarly gives the *directional masses* $a_j^{\mathrm{F}}(x)$ and $a_j^{\mathrm{B}}(x)$ relative to π_j:

$$\pi_j(dx) a_j^{\mathrm{F}}(x) := \widehat{C}_j(dx, \mathcal{X}), \qquad \pi_j(dx) a_j^{\mathrm{B}}(x) := \widehat{C}_{j+1}(\mathcal{X}, dx), \tag{6}$$

Dividing each weighted reference kernel by this directional mass makes it a probability transition kernel over the next state (Figure 1). Thus, for $2 \le j \le J$,

$$Q_j^{\mathrm{F}}(dy \mid x) = \frac{\sqrt{\widehat{V}_j(x)\widehat{V}_{j-1}(y)}}{a_j^{\mathrm{F}}(x)} K_j^{\mathrm{F}}(dy \mid x), \quad Q_{j-1}^{\mathrm{B},\rho}(dx \mid y) = \frac{\sqrt{\widehat{V}_j(x)\widehat{V}_{j-1}(y)}}{a_{j-1}^{\mathrm{B}}(y)} K_{j-1}^{\mathrm{B},\rho}(dx \mid y). \tag{7}$$

On the clean edge, $Q_1^{\mathrm{F}}(dy \mid x) = r(y)K_1^{\mathrm{F}}(dy \mid x)/a_1^{\mathrm{F}}(x)$ and $Q_0^{\mathrm{B},\rho} = K_0^{\mathrm{B},\rho}$. We can therefore sample the destination of a move using the corresponding Q kernel, and use the relative sizes of a_j^{F} and a_j^{B} to set the move, stay, and flip probabilities below.

Remark 2.1 (When $\rho_j = 1$)**.** For diffusion models, $\rho_j = 1$ when the denoising transitions exactly reverse the reference noising process (Remark B.1). The same identity holds when an exact clean posterior sample is followed by a Gaussian bridge from that noising process (Proposition B.1).

Persistent denoising and renoising. We now describe how the CS-VGB sampler moves between checkpoints while retaining its direction $d \in \{\mathrm{F}, \mathrm{B}\}$ as part of the state (j, x, d), with F moving to $j-1$ and B to $j+1$. At an interior checkpoint, the relative directional masses $p_j^d = a_j^d/(a_j^{\mathrm{F}} + a_j^{\mathrm{B}})$ and their overlap $\ell_j = \min\{p_j^{\mathrm{F}}, p_j^{\mathrm{B}}\}$ determine the action probabilities in Figure 2: a **move** samples from the corresponding Q_j^{F} or $Q_j^{\mathrm{B},\rho}$ kernel, a **stay** leaves the state unchanged, and a **flip** reverses only d. To reduce direction reversals without changing move probabilities, the *directional persistence parameter* $\chi \in [0, 1]$ replaces a larger fraction of flips with stays as χ increases. This follows crossing-flow cancellation in lifted chains (Hayes & Sinclair, 2010), as used in discrete VGB (Rohatgi et al., 2026; Jeon et al., 2026), while boundary moves are forced inward.

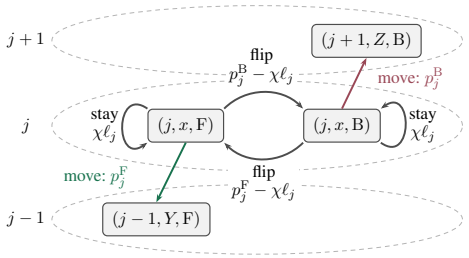

Figure 2: Interior updates: move, stay, or flip. Here $Y \sim Q_j^{\mathrm{F}}$ and $Z \sim Q_j^{\mathrm{B},\rho}$.

Whereas χ controls directional persistence, we may also wish to control the relative preference for denoising and renoising while preserving the clean target. We therefore introduce a *depth weight* $\alpha > 0$ by scaling each edge-conductance measure as $\widehat{C}_j^{\alpha} := \alpha^{j-1}\widehat{C}_j$ (Equation (B.17)). This multiplies the backward-to-forward mass ratio $a_j^{\mathrm{B}}/a_j^{\mathrm{F}}$ by α, so $\alpha < 1$ favors denoising and $\alpha > 1$ favors renoising. We remark that the stationary clean conditional remains still unchanged because the clean edge has factor $\alpha^0 = 1$.

Finite-candidate implementation. As the integrals defining the directional masses in Equation (6) are generally intractable in practice, we approximate them using retained pools of $L_F, L_B \ge 2$ candidates, $Y_i \sim G_j^{\mathrm{F}}(\cdot \mid x)$ and $Z_k \sim G_j^{\mathrm{B}}(\cdot \mid x)$. Here G_j^d is the *proposal kernel* used to generate candidates in direction d. It may replace the reference denoising or renoising kernel K_j^d to propose more promising states, for example through gradient-based guidance in denoising or even renoising.

Combining this depth weighting with the proposal correction gives the candidate weights

$$e_{j,i}^{\mathrm{F}} := \alpha^{j-1}\frac{dK_j^{\mathrm{F}}}{dG_j^{\mathrm{F}}}(Y_i \mid x)\sqrt{\widehat{V}_j(x)\widehat{V}_{j-1}(Y_i)}, \qquad e_{j,k}^{\mathrm{B}} := \alpha^{j}\frac{dK_j^{\mathrm{B},\rho}}{dG_j^{\mathrm{B}}}(Z_k \mid x)\sqrt{\widehat{V}_j(x)\widehat{V}_{j+1}(Z_k)}, \tag{8}$$

where $\sqrt{\widehat{V}_j(x)\widehat{V}_{j-1}(Y_i)}$ is replaced by $r(Y_i)$ at $j = 1$, and $\sqrt{\widehat{V}_j(x)\widehat{V}_{j+1}(Z_k)}$ by $r(x)$ at $j = 0$. With both pools evaluated, use

$$\widehat{a}_j^d := \frac{1}{L_d}\sum_{i=1}^{L_d} e_{j,i}^d, \qquad \widehat{p}_j^d := \frac{\widehat{a}_j^d}{\widehat{a}_j^{\mathrm{F}} + \widehat{a}_j^{\mathrm{B}}}, \qquad \widehat{\ell}_j := \min\{\widehat{p}_j^{\mathrm{F}}, \widehat{p}_j^{\mathrm{B}}\} \tag{9}$$

in the action rule of Figure 2, selecting candidate i on a move with probability $e_{j,i}^d / \sum_h e_{j,h}^d$.

Reverse insertion. However, naively redrawing all candidates after each move can bias the stationary distribution, so we use reverse insertion to preserve the clean target with finite candidate pools. After a

move, insert the previous x into a uniform opposite-pool slot at the destination and draw all remaining entries independently (Figure B.4). For a flip, no move or reverse insertion occurs: the state (j, x) and both pools are retained, with the roles of the current and opposite pools exchanged as d reverses. In the stay case, the state and pools remain unchanged.

Lazy pool evaluation. Opposite-direction candidates determine action probabilities but are not destinations of the current move, so frequently evaluating their full pool can add excessive model calls, especially for verifiers based on denoised clean predictions. We therefore evaluate only the current pool and one *marked* opposite candidate, deferring the remaining evaluations until needed. In direction F, with *marked* Z_K in the B-candidate pool, draw a temporary $I \sim \mathrm{Unif}(\{1, \dots, L_F\})$ and use

$$\mathbb{P}(\mathsf{move}, \mathsf{stay}, \mathsf{flip} \mid I, K) = \frac{(\widehat{a}_j^{\mathrm{F}},\ \chi \min\{e_{j,I}^{\mathrm{F}}, e_{j,K}^{\mathrm{B}}\},\ e_{j,K}^{\mathrm{B}} - \chi \min\{e_{j,I}^{\mathrm{F}}, e_{j,K}^{\mathrm{B}}\})}{\widehat{a}_j^{\mathrm{F}} + e_{j,K}^{\mathrm{B}}}. \tag{10}$$

A move independently redraws Y_i with probability $e_{j,i}^{\mathrm{F}} / \sum_{h=1}^{L_F} e_{j,h}^{\mathrm{F}}$ and marks the reverse-inserted state at an interior destination. A flip uses the preselected index I to mark Y_I and fills only missing B-pool entries, while a stay changes nothing. The B update is symmetric, with evaluated entries reused across flips and stays (Figure B.5).

These updates form the complete CS-VGB algorithm in Algorithm 1, which still preserves the stationary clean distribution $\pi_0^\star$ under reverse insertion and lazy pool evaluation when $e_{j,i}^{\mathrm{F}}$ and $e_{j,k}^{\mathrm{B}}$ are exact (Theorem C.2).

2.2 Theoretical Guarantees of CS-VGB

In this subsection, we study the theoretical guarantees of CS-VGB. In particular, we ask (1) which clean distribution is preserved, (2) how quickly the chain approaches stationarity, and (3) how soon it reaches a clean sample.

Stationary clean target. Since the clean-edge marginal is $\widehat{C}_1(\mathcal{X}, dx) = r(x)\pi_0(dx)$, normalizing it gives the desired target $\pi_0^\star$. The following theorem therefore shows that this target is preserved as the stationary clean conditional of CS-VGB, even when the process verifier is inaccurate.

Theorem 2.1 (Stationary clean target; Theorem C.1)**.** *Under exact terminal reward evaluation (Assumption B.1) and the paired-transition identity in Equation (4), CS-VGB has a stationary law* $\Pi(j, dx, d)$ *satisfying the following without requiring* $\widehat{V}_j = V_j$*:*

$$\Pi(dx, \mathrm{B} \mid j = 0) = \Pi(dx \mid j = 0) = r(x)\pi_0(dx)/Z = \pi_0^\star(dx). \qquad Z = \int r(x)\pi_0(dx), \tag{11}$$

Convergence under verifier error. Preserving the target does not determine how quickly it is reached. We therefore analyze the convergence speed of CS-VGB kernel P defined by the exact directional masses in Equation (6), with depth weight $\alpha = 1$. For the theoretical analysis, we use lazified kernel $P_{\mathrm{L}} = (I + P)/2$ to remove periodicity. We measure initialization error by $D_\nu := \|d\nu/d\Pi - 1\|_{L^2(\Pi)}$ and use $\|\cdot\|_{\mathrm{TV}}$ for total variation distance. With $\pi_j^\star(dx) := V_j(x)\pi_j(dx)/Z$, we quantify the *average-state verifier error* by $\varepsilon_{\mathrm{av}} := \max_{1 \le j \le J} \|\widehat{V}_j/V_j - 1\|_{L^2(\pi_j^\star)}$.

Theorem 2.2 (Convergence under verifier error; Theorem C.4 and Corollary C.1)**.** *Let* $J \ge 2$*,* $0 < r_{\min} \le r \le 1$*, and* $0 < v_{\min} \le \widehat{V}_j \le 1$*. Assume endpoint maximal correlation* $\rho_{\mathrm{end}} < 1$ *(Equation (C.42)) and a finite exact-verifier transition-density moment* H_p *for some* $p > 2$ *(Equation (C.84)). Set* $\gamma_V = (1 - \rho_{\mathrm{end}})r_{\min}/(32J^2)$*. Uniformly over* $\chi \in [0, 1]$*, we can choose*

$$\gamma_{\mathrm{CS}} \ge \max\{\gamma_V r_{\min} v_{\min}^2,\ \gamma_V - O(H_p \varepsilon_{\mathrm{av}}^{1-2/p})\} \ge \gamma_V \left(1 + O\left(\frac{H_p \varepsilon_{\mathrm{av}}^{1-2/p}}{\gamma_V}\right)\right)^{-1} > 0, \tag{12}$$

where the $O(\cdot)$ *hides constants depending only on* $r_{\min}, v_{\min}, p$*. For initial law* $\nu \ll \Pi$ *with* $D_\nu < \infty$*,* $\|\nu P_{\mathrm{L}}^n - \Pi\|_{\mathrm{TV}} = O(D_\nu e^{-\gamma_{\mathrm{CS}} n/4})$*. For* $n > \frac{4}{\gamma_{\mathrm{CS}}} \log(\max\{1, D_\nu/[2\Pi(j=0)]\})$*,*

$$\|(\nu P_{\mathrm{L}}^n)(dx \mid j = 0) - \pi_0^\star(dx)\|_{\mathrm{TV}} = O\left(\frac{D_\nu}{\Pi(j=0)} e^{-\gamma_{\mathrm{CS}} n/4}\right). \tag{13}$$

Remark 2.2 (Sufficient conditions for convergence)**.** For Theorem 2.2, exactly paired Gaussian transitions with bounded clean support ensure $\rho_{\mathrm{end}} < 1$ and $H_p < \infty$ for some $p > 2$ (Lemmas C.4 and C.5). The moment condition also holds under the exponential-tail assumption in Lemma C.5. Under exact pairing, both linear-interpolant renoising at $t_{\kappa_J} = 1$ (Equation (B.29)) and VP renoising with an endpoint-exact schedule (Table A.2) yield an independent noise reset, hence $\rho_{\mathrm{end}} = 0$.

Equation (12) shows that CS-VGB converges exponentially fast to the correct stationary clean target even with inaccurate process verifiers, although the rate depends on verifier accuracy. As the average-state verifier error $\varepsilon_{\mathrm{av}}$ decreases, the guaranteed rate approaches the exact-verifier bound. We remark that, unlike the uniform multiplicative error bound used in MDM-VGB's mixing speed analysis (Jeon et al., 2026), our bound measures error in $L^2(\pi_j^\star)$, following the average-case perspective also studied for AR-VGB (Rohatgi et al., 2026). Importantly, the lower bound rate $\gamma_V r_{\min} v_{\min}^2$ gives a **dimension-free** rate guarantee when $1 - \rho_{\mathrm{end}}, r_{\min}, v_{\min}$ remain uniformly bounded away from zero as dimension d increases.

Remark 2.3 (Clean stationary mass; Proposition C.1)**.** If $J^{-1}\sum_{j=1}^{J} \mathbb{E}_{\pi_j}|\widehat{V}_j - V_j| \le \varepsilon Z$, then $\Pi(j=0) \ge (4J - 2 + 4J\varepsilon)^{-1}$. For bounded ε, the clean level therefore retains stationary mass $\Omega(J^{-1})$, paralleling the inverse-sequence-length leaf mass in MDM-VGB (Jeon et al., 2026).

Remark 2.4 (Transition mismatch; Theorem C.5 and Corollary C.4)**.** Unlike process-verifier error, omitting the edge correction ρ can bias the stationary clean target. Let $\varepsilon_{\rho,\mathrm{av}} := \max_{1\le j\le J} \|\rho_j - 1\|_{L^2(J_j^{\mathrm{B}})}$ quantify this error. Under the assumptions of Theorem C.5, let $\Delta_{\rho,\mathrm{av}} = O(\varepsilon_{\rho,\mathrm{av}}^{(p-2)/(p-1)})$ as in Equation (C.101), for admissible $p > 2$. If $\Delta_{\rho,\mathrm{av}} < \gamma_{\mathrm{CS}}$ and $\Delta_{\rho,\mathrm{av}}/(\gamma_{\mathrm{CS}} - \Delta_{\rho,\mathrm{av}}) < \Pi(j=0)$, the uncorrected CS-VGB kernel admits a stationary law $\bar{\Pi}$ with

$$\|\bar{\Pi}(dx \mid j=0) - \pi_0^\star\|_{\mathrm{TV}} \le \frac{\Delta_{\rho,\mathrm{av}}}{\Pi(j=0)(\gamma_{\mathrm{CS}} - \Delta_{\rho,\mathrm{av}})},$$

The lazified chain still converges exponentially under Corollary C.4, but additional updates alone cannot remove a nonzero stationary bias. Unlike process-verifier error, transition mismatch can therefore affect the limiting clean distribution, with the bias vanishing as $\varepsilon_{\rho,\mathrm{av}} \to 0$.

Directional persistence and the first clean visit. While convergence concerns distributional accuracy, finite-budget generation also requires reaching a clean sample within a limited number of updates. To quantify this, let $\tau_0 := \inf\{n \ge 0 : j_n = 0\}$ be the *first-hitting time* for the original kernel. The following bound gives $\Theta(J)$ expected updates to the first clean checkpoint under an average directional-imbalance condition, paralleling the $\Theta(n)$ bound for momentum MDM-VGB under uniform near-balance, where n is sequence length (Jeon et al., 2026). We remark that this dependence on J is optimal, since at least J moves are necessary.

Proposition 2.1 (First-hitting time; Proposition C.2)**.** *Assume the exact-kernel conditions of Theorem 2.1, $J \ge 2$, $0 < v_{\min} \le \widehat{V}_j \le 1$, and the noise reset $K_{J-1}^{\mathrm{B},\rho}(\cdot \mid y) = \pi_J$ almost surely. Along $Y_J \sim \pi_J$, $Y_{j-1} \sim Q_j^{\mathrm{F}}(\cdot \mid Y_j)$, suppose*

$$\max_{1\le j<J} \mathbb{E}\left[\left(\log \frac{a_j^{\mathrm{B}}(Y_j)}{a_j^{\mathrm{F}}(Y_j)}\right)_+\right] = O(J^{-1}), \qquad 1 - \chi = O(J^{-1}). \tag{14}$$

Then $\mathbb{E}_{(J,x,\mathrm{F})}[\tau_0] = \Theta(J)$ for π_J-almost every x.

2.3 Combining Sequential and Parallel Search

To use parallel compute alongside recurrent exploration, we propose **Multi-chain Synchronized CS-VGB** which shares checkpoints and actions across N particles with local candidate selection, preserving the product clean target under the stated assumptions (Appendix D and Theorem D.2). For practical one-step posterior-map search, particles share the highest-reward endpoint and its parent, then select the highest-weight local candidate. Detailed updates appear in Equations (F.4) and (F.5). Although this heuristic does not inherit the stationary-target guarantee, we propose it as a practical way to extend CS-VGB's value-guided recurrent bidirectional sampling to finite-budget search.

3 EXPERIMENTS

We empirically evaluate (1) whether CS-VGB recovers the reward-tilted target despite process-verifier error, (2) whether recurrent, value-guided denoising and renoising improve search performance at comparable computational cost, and (3) how individual design choices contribute to these gains.

3.1 SAMPLING ACCURACY UNDER IMPERFECT VERIFIERS

We first ask whether repeated denoising and renoising recover $\pi_0^\star$ despite process-verifier error. We use a two-dimensional Gaussian mixture with an analytic reward-tilted target and exactly paired transitions. We then vary only the process-verifier error while keeping the reference process and target fixed. To model verifier error that diminishes toward the clean endpoint, motivated by reward estimates from one-step denoising, we set $\log \widehat{V}_j(x) = \log V_j(x) + \lambda t\, \eta(x,t) - \frac{1}{2}(\lambda t)^2 s_\eta^2$ at $t = t_{\kappa_j}$, where $\eta(x,t)$ is a fixed smooth random Fourier error field and s_η^2 is its variance over random phases (Appendix F.2). In this setup, the scale λ controls the distortion strength, and $\widehat{V}_0 = r$ remains exact.

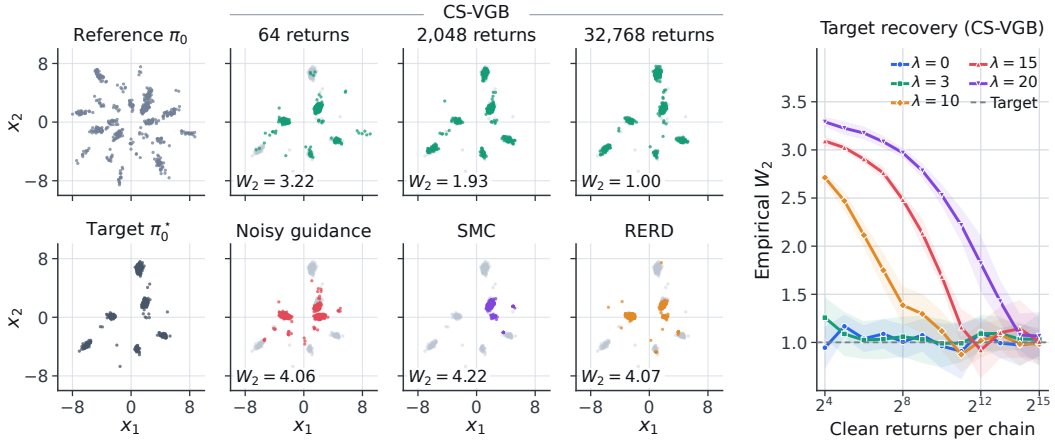

Figure 3: **Target recovery of CS-VGB under process-verifier error.** Left: reference and target samples. Center: CS-VGB at successive clean returns (top) and baseline outputs (bottom), all at $\lambda = 20$, each showing 512 samples overlaid on target samples in gray. Right: empirical W_2 across 10 random seeds for $\lambda \in \{0, 3, 10, 15, 20\}$ (mean ± 1 SD). Each return count is per chain. The dashed line averages W_2 over 100 independent 512-sample target sets against the fixed target reference.

Recovery under verifier error. CS-VGB recovers comparable target accuracy across verifier-error levels (Figure 3). At $\lambda = 20$, empirical W_2 falls from 2.78 at 512 clean returns to 1.07 at 32,768, close to 1.00 with the exact verifier. In contrast, finite-particle SMC and noisy guidance retain distorted mode weights under strong verifier error, as does RERD (Uehara et al., 2025) despite repeated renoising and denoising. This comparison highlights the importance of the target-preserving update rule of CS-VGB, rather than renoising and denoising alone. While CS-VGB reaches comparable target accuracy across error levels, larger λ requires more clean returns, consistent with Theorem 2.2: verifier error can slow convergence without changing the clean target.

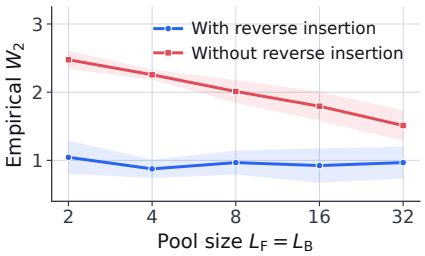

Figure 4: **Small-pool recovery with reverse insertion.** Empirical W_2 after 2^{11} clean returns per chain at $\lambda = 10$.

Why reverse insertion matters. This recovery is not obtained simply by repeatedly drawing fresh candidate pools. Removing reverse insertion increases W_2 from 0.88 to 2.26 at $L_F = L_B = 4$, with 2,048 clean returns per chain (Figure 4). With insertion, error remains near one across pool sizes; without it, larger pools reduce the error but leave a gap even at $L_F = L_B = 32$. Thus, the ablation supports reverse insertion as a mechanism for accurate finite-pool sampling, rather than relying on large pools to approximate the exact masses.

3.2 Compute-Matched Search Performance

The sampling results illustrate how CS-VGB uses imperfect verifiers to guide exploration while preserving the clean target defined by the exact terminal reward. For practical reward-guided generation, a common goal is to find high-reward samples within a limited computational budget. This motivates a complementary question: *can CS-VGB's recurrent, value-guided exploration also improve finite-budget search?* We evaluate this across image generation with posterior maps and multi-robot planning with diffusion models, using the practical search implementations described in Appendix F.1.

To assess search performance, we compare task-specific reward across computational budgets, measured by the number of function evaluations (NFE). We count each generator forward pass as one NFE and each guidance vector–Jacobian product (VJP) through the generator as two additional NFEs. The suffix `-G` denotes gradient-guided variants; guided CS-VGB uses a single chain. Further task descriptions, experimental setups, and hyperparameters are provided in Appendix F; additional ablations in Appendix F.6 examine computational savings from lazy pool evaluation and the effect of directional persistence χ on checkpoint traversal.

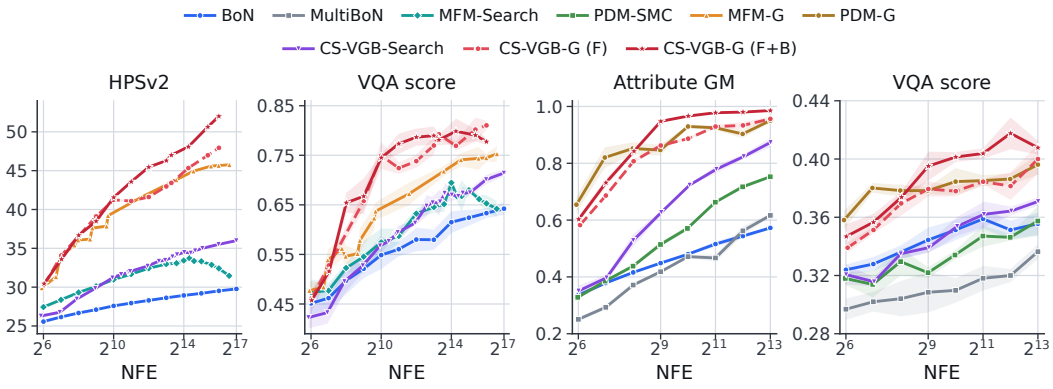

(a) ImageNet with MFM

(b) CelebA with PDM

Figure 5: **Reward-guided image generation with posterior maps.** Optimized reward (ImageNet: HPSv2, CelebA: Attribute GM) and independent VQA scores versus actual NFE. Means over three seeds, each with 100 prompts or conditions; bands show ± 1 SD where available. F and F+B denote gradient guidance in denoising only and both directions, respectively.

Prompt-guided generation with posterior maps. Using Meta Flow Maps (MFM) (Potaptchik et al., 2026) on ImageNet (Deng et al., 2009), we generate images for 100 class–prompt pairs that we construct by pairing each ImageNet class label with an additional visual condition (Appendix F.3). We optimize HPSv2 (Wu et al., 2023b) and use VQAScore (Lin et al., 2024) as a held-out metric of prompt alignment. We compare CS-VGB with best-of-N (BoN), MFM-Search, and MFM-G (Figure 5). CS-VGB-G reaches HPSv2 45.46 at 4,104 NFE, exceeding MFM-G's 44.94 at 23,959 NFE with $5.8\times$ fewer NFEs. It then continues improving beyond MFM-G's plateau, reaching 52.01 versus 45.66 at approximately 2^{16} NFE, with higher held-out VQAScore (0.778 versus 0.745). Without gradients, CS-VGB-Search also surpasses MFM-Search at larger budgets. Thus, recurrent search both reduces the computation needed for comparable reward and makes productive use of additional computation after one-way guidance saturates. Qualitative comparisons are provided in Table F.3.

Facial-attribute control with posterior maps. With Posterior Diamond Maps (PDM) (Holderrieth et al., 2026) on CelebA 64×64 (Liu et al., 2015), we target 100 conditions specifying ten facial attributes each. We optimize the geometric mean (GM) of the target-attribute probabilities predicted by FaRL (Zheng et al., 2022) and use VQAScore as a held-out metric of alignment with the specified attributes. Baselines are BoN, MultiBoN (Lee et al., 2025a; Sabour et al., 2025), PDM-SMC, and PDM gradient guidance (Figure 5). CS-VGB-G reaches attribute GM 0.948 at 514.62 NFE, exceeding PDM-G's 0.929 at 1,023 NFE while using $1.99\times$ fewer NFEs. Without gradients, CS-VGB-Search reaches GM 0.723 at 1,066 NFE, compared with PDM-SMC's 0.717 at 4,068 NFE—a roughly $3.8\times$ reduction in NFE. The gains with and without gradients

suggest that recurrent search can reach strongly attribute-aligned outputs more efficiently, rather than benefiting only from gradient guidance. Qualitative comparisons are provided in Table F.5.

Why does recurrent bidirectional search help? We now examine whether the gains from recurrent search depend on continuing from previously reached clean states, rather than repeatedly restarting from independent noise. Unlike a restart, renoising retains the current clean state as the starting point for exploring alternative continuations, allowing subsequent computation to build on what the search has already found. Consistent with this intuition, recurrence improves both ImageNet HPSv2 and CelebA attribute GM, with larger gains under gradient guidance (Table 2). This pattern suggests that reusing a promising state is especially useful when guidance can steer the resulting exploration toward higher reward, motivating a closer look at how the backward proposals are generated.

Table 2: **Restart versus recurrence.** ImageNet: HPSv2; CelebA: attribute GM. Three-seed means over 100 conditions; target NFE 4,096.

Task / variant	Restart	Recurrent
ImageNet / -G	39.20	**45.57**
ImageNet / Search	32.23	**32.69**
CelebA / -G	0.86	**0.98**
CelebA / Search	0.80	**0.82**

Why does backward guidance help? The next question is therefore whether guidance should act only during denoising or also shape the noisier states from which denoising resumes. The edge-conductance kernels in Equation (7) already favor destinations with larger $\sqrt{\widehat{V}}$ at interior checkpoints, and the candidate weights in Equation (8) implement this preference through selection. However, selection can only exploit the candidates available: forward-only guidance (F) retains value-based backward selection but does not steer backward candidate generation. Bidirectional guidance (F+B) addresses this distinction by guiding the backward proposals themselves toward high-value regions. In Figure 5, F+B reaches ImageNet HPSv2 52.01 versus 47.94 with F at approximately 2^{16} NFE, and improves both attribute reward and VQAScore on CelebA across the evaluated budgets. These results suggest that backward guidance improves recurrent search by steering candidate generation toward high-value regions before value-based selection.

Multi-robot planning with diffusion models. Using pretrained trajectory diffusion models from MMD (Shaoul et al., 2025), we seek collision-free joint trajectories in four environments: Highways, Circle, and two Conveyor variants. Highways and Circle involve 16 robots, while the Conveyor tasks involve four; instances vary robot start and goal positions and obstacle geometry. We compare against BoN-G and the recurrent diffusion-search baselines ABCD-G (Lee et al., 2025a) and DFS-G (Zhang et al., 2026). At the largest NFE cap, CS-VGB-G achieves 84.8% joint success, compared with 71.2% for ABCD-G, 67.6% for BoN-G, and 60.1% for DFS-G, using fewer model evaluations on average (Figure 6). Although CS-VGB-G underperforms ABCD-G and DFS-G at low NFE, it surpasses both at higher budgets. This crossover is consistent with process-verifier evaluation overhead limiting low-budget search, while larger budgets allow value-guided refinement to offset that cost. Environment details and task-wise results are provided in Appendix F.5; Figure F.4 shows qualitative examples of successful CS-VGB-G trajectories.

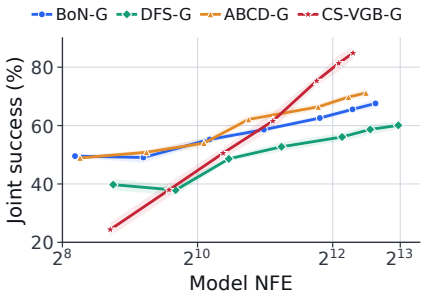

Figure 6: **Multi-robot planning.** Joint success over four tasks with 64 instances each (mean ± 1 SD, five seeds) versus actual model NFE.

4 Conclusion, Limitations, and Future Work

We introduced CS-VGB for recurrent, value-guided sampling and search with pretrained generators in continuous domain. Shared edge conductances and terminal anchoring preserve the clean reward-tilted target despite verifier error under stated assumptions, with convergence, first-hitting, and finite-pool guarantees. Experiments demonstrate target recovery and search gains beyond observed denoising-only plateaus in image generation, alongside improved planning success at larger budgets. As limitations, verifier error can slow mixing, evaluation costs can offset low-budget gains, and sampling guarantees do not automatically extend to search heuristics. A natural future direction is to explore alternatives to geometric-mean conductance weighting, including variationally learned conductances, to improve mixing or computational efficiency while retaining terminal anchoring and the conditions required for target preservation.

AI USE STATEMENT

In this work, we used generative AI tools for editing and polishing the writing, debugging code, and creating TikZ figures. We also used these tools to clarify mathematical background relevant to our derivations and to cross-check proofs. We did not use generative AI tools for developing the core research ideas or drawing scientific conclusions. All AI-assisted text, code, figures, and mathematical arguments were reviewed and verified by the authors. We take responsibility for the final content of this work.

REPRODUCIBILITY STATEMENT

We provide the implementation details needed to reproduce our results in the main paper and appendix, including pretrained model information, inference configurations, evaluation protocols, and hyperparameters. Additional details on datasets, preprocessing, and experimental settings are provided in Appendix F. The theoretical framework, assumptions, derivations, and complete proofs are presented in Appendices B to D. We also provide anonymized source code and scripts for reproducing the main experiments in the supplementary material.

REFERENCES

Michael S. Albergo, Nicholas M. Boffi, and Eric Vanden-Eijnden. Stochastic interpolants: A unifying framework for flows and diffusions. *Journal of Machine Learning Research*, 26(209):1–80, 2025. 16

Brian D. O. Anderson. Reverse-time diffusion equation models. *Stochastic Processes and their Applications*, 12(3):313–326, 1982. doi: 10.1016/0304-4149(82)90051-5. 15

Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibo Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, Humen Zhong, Yuanzhi Zhu, Mingkun Yang, Zhaohai Li, Jianqiang Wan, Pengfei Wang, Wei Ding, Zheren Fu, Yiheng Xu, Jiabo Ye, Xi Zhang, Tianbao Xie, Zesen Cheng, Hang Zhang, Zhibo Yang, Haiyang Xu, and Junyang Lin. Qwen2.5-VL technical report. *arXiv preprint arXiv:2502.13923*, 2025. 67

Sergey Bezuglyi and Palle E. T. Jorgensen. Markov operators generated by symmetric measures. *arXiv preprint arXiv:1812.00081*, 2018. 3, 19

Haim Brezis. *Functional Analysis, Sobolev Spaces and Partial Differential Equations*. Universitext. Springer, 2011. doi: 10.1007/978-0-387-70914-7. 36

Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 248–255. IEEE, 2009. doi: 10.1109/CVPR.2009.5206848. 8, 67

Prafulla Dhariwal and Alexander Nichol. Diffusion models beat GANs on image synthesis. In *Advances in Neural Information Processing Systems*, volume 34, 2021. 64

Peter G. Doyle and J. Laurie Snell. *Random Walks and Electric Networks*, volume 22 of *Carus Mathematical Monographs*. Mathematical Association of America, 1984. 3, 19

James Allen Fill. Eigenvalue bounds on convergence to stationarity for nonreversible Markov chains, with an application to the exclusion process. *The Annals of Applied Probability*, 1(1):62–87, 1991. doi: 10.1214/aoap/1177005981. 32

Thomas P. Hayes and Alistair Sinclair. Liftings of tree-structured Markov chains. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, volume 6302 of *Lecture Notes in Computer Science*, pp. 602–616. Springer, 2010. doi: 10.1007/978-3-642-15369-3_45. 4, 21

Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022. 64

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*, volume 33, pp. 6840–6851, 2020. 2, 65

Peter Holderrieth, Douglas Chen, Luca Eyring, Ishin Shah, Giri Anantharaman, Yutong He, Zeynep Akata, Tommi Jaakkola, Nicholas Matthew Boffi, and Max Simchowitz. Diamond maps: Efficient reward alignment via stochastic flow maps. *arXiv preprint arXiv:2602.05993*, 2026. 1, 2, 8, 16, 19, 25, 65, 68, 69

Vineet Jain, Kusha Sareen, Mohammad Pedramfar, and Siamak Ravanbakhsh. Diffusion tree sampling: Scalable inference-time alignment of diffusion models. In *Advances in Neural Information Processing Systems*, volume 38, 2025. 64

Kijung Jeon, Thuy-Duong Vuong, and Molei Tao. VGB for masked diffusion model: Efficient test-time scaling for reward satisfaction and sample editing. *arXiv preprint arXiv:2606.28301*, 2026. 2, 4, 6, 21, 64

Erwin Kreyszig. *Introductory Functional Analysis with Applications*. John Wiley & Sons, 1978. 39, 57, 59, 60

Gyubin Lee, Bao Truong, Jaesik Yoon, Dongwoo Lee, Minsu Kim, Yoshua Bengio, and Sungjin Ahn. Adaptive inference-time scaling via cyclic diffusion search. In *Advances in Neural Information Processing Systems*, volume 38, pp. 56999–57042, 2025a. doi: 10.52202/085713-1909. 1, 8, 9, 64

Sanghyun Lee, Sunwoo Kim, Seungryong Kim, Jongho Park, and Dongmin Park. Effective test-time scaling of discrete diffusion through iterative refinement. *arXiv preprint arXiv:2511.05562*, 2025b. 64

Xiner Li, Yulai Zhao, Chenyu Wang, Gabriele Scalia, Gokcen Eraslan, Surag Nair, Tommaso Biancalani, Shuiwang Ji, Aviv Regev, Sergey Levine, and Masatoshi Uehara. Derivative-free guidance in continuous and discrete diffusion models with soft value-based decoding. In *Advances in Neural Information Processing Systems*, volume 38, 2025. 1, 17, 21, 64

Han Cheng Lie, Konstantin Fackeldey, and Marcus Weber. A square root approximation of transition rates for a Markov state model. *SIAM Journal on Matrix Analysis and Applications*, 34(2):738–756, 2013. doi: 10.1137/120899959. 19

Zhiqiu Lin, Deepak Pathak, Baiqi Li, Jiayao Li, Xide Xia, Graham Neubig, Pengchuan Zhang, and Deva Ramanan. Evaluating text-to-visual generation with image-to-text generation. In *Computer Vision – ECCV 2024*, volume 15067 of *Lecture Notes in Computer Science*, pp. 366–384. Springer, 2024. doi: 10.1007/978-3-031-72673-6_20. 8, 67, 71

Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. In *International Conference on Learning Representations*, 2023. 16

Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 3730–3738, 2015. URL `https://liuziwei7.github.io/projects/FaceAttributes.html`. 8, 69

Nanye Ma, Shangyuan Tong, Haolin Jia, Hexiang Hu, Yu-Chuan Su, Mingda Zhang, Xuan Yang, Yandong Li, Tommi Jaakkola, Xuhui Jia, and Saining Xie. Scaling inference time compute for diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2025. 64

Sam McCallum, Zander W. Blasingame, Timothy Herschell, Niklas Rindtorff, Alexander Tong, and James Foster. Strong stochastic flow maps. *arXiv preprint arXiv:2606.01086*, 2026. 65

Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. PyTorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*, volume 32, 2019. 65

Xinyu Peng, Ziyang Zheng, Yaoming Wang, Han Li, Nuowen Kan, Wenrui Dai, Chenglin Li, Junni Zou, and Hongkai Xiong. Noise conditional variational score distillation. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 48923–48947. PMLR, 2025. 65

Peter Potaptchik, Adhi Saravanan, Abbas Mammadov, Alvaro Prat, Michael S. Albergo, and Yee Whye Teh. Meta Flow Maps enable scalable reward alignment. *arXiv preprint arXiv:2601.14430*, 2026. 1, 2, 8, 16, 19, 25, 65, 67, 68

Dhruv Rohatgi, Abhishek Shetty, Donya Saless, Yuchen Li, Ankur Moitra, Andrej Risteski, and Dylan J. Foster. Taming imperfect process verifiers: A sampling perspective on backtracking. In *The Fourteenth International Conference on Learning Representations*, 2026. 2, 4, 6, 21, 64

Amirmojtaba Sabour, Michael S. Albergo, Carles Domingo-Enrich, Nicholas M. Boffi, Sanja Fidler, Karsten Kreis, and Eric Vanden-Eijnden. Test-time scaling of diffusions with flow maps. *arXiv preprint arXiv:2511.22688*, 2025. 8, 17

Yorai Shaoul, Itamar Mishani, Shivam Vats, Jiaoyang Li, and Maxim Likhachev. Multi-robot motion planning with diffusion models. In *International Conference on Learning Representations*, 2025. 9, 73

Raghav Singhal, Zachary Horvitz, Ryan Teehan, Mengye Ren, Zhou Yu, Kathleen McKeown, and Rajesh Ranganath. A general framework for inference-time scaling and steering of diffusion models. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 55810–55827. PMLR, 2025. 1, 64

Jiaming Song, Qinsheng Zhang, Hongxu Yin, Morteza Mardani, Ming-Yu Liu, Jan Kautz, Yongxin Chen, and Arash Vahdat. Loss-guided diffusion models for plug-and-play controllable generation. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 32483–32498. PMLR, 2023. 1, 64

Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021. 2, 15, 65

Terence Tao. *An Epsilon of Room, I: Real Analysis: Pages from Year Three of a Mathematical Blog*, volume 117 of *Graduate Studies in Mathematics*. American Mathematical Society, 2010. ISBN 978-0-8218-5278-1. 39

Masatoshi Uehara, Xingyu Su, Yulai Zhao, Xiner Li, Aviv Regev, Shuiwang Ji, Sergey Levine, and Tommaso Biancalani. Reward-guided iterative refinement in diffusion models at test-time with applications to protein and DNA design. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 60515–60529. PMLR, 2025. 1, 7, 64

Jianyi Wang, Kelvin C. K. Chan, and Chen Change Loy. Exploring CLIP for assessing the look and feel of images. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pp. 2555–2563, 2023. doi: 10.1609/aaai.v37i2.25353. URL https://github.com/IceClear/CLIP-IQA. 68, 71

Luhuan Wu, Brian L. Trippe, Christian A. Naesseth, David M. Blei, and John P. Cunningham. Practical and asymptotically exact conditional sampling in diffusion models. In *Advances in Neural Information Processing Systems*, volume 36, 2023a. 1, 64

Xiaoshi Wu, Yiming Hao, Keqiang Sun, Yixiong Chen, Feng Zhu, Rui Zhao, and Hongsheng Li. Human preference score v2: A solid benchmark for evaluating human preferences of text-to-image synthesis. *arXiv preprint arXiv:2306.09341*, 2023b. 8, 67

Giacomo Zanella. Informed proposals for local MCMC in discrete spaces. *Journal of the American Statistical Association*, 115(530):852–865, 2020. doi: 10.1080/01621459.2019.1585255. 19

Xiangcheng Zhang, Haowei Lin, Haotian Ye, James Zou, Jianzhu Ma, Yitao Liang, and Yilun Du. Inference-time scaling of diffusion models through classical search. In *International Conference on Learning Representations*, 2026. 1, 9, 64

Yinglin Zheng, Hao Yang, Ting Zhang, Jianmin Bao, Dongdong Chen, Yangyu Huang, Lu Yuan, Dong Chen, Ming Zeng, and Fang Wen. General facial representation learning in a visual-linguistic manner. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 18697–18709, 2022. 8, 70

Qijie Zhu, Zeqi Ye, Han Liu, Zhaoran Wang, and Minshuo Chen. Training-free adaptation of diffusion models via Doob's h-transform. *arXiv preprint arXiv:2602.16198*, 2026. 64

A NOTATION AND REFERENCE TRANSITIONS

Table A.1: Table of Key Notations

Symbol	Definition	Mathematical description
t_k	Time grid	$0 = t_0 < \cdots < t_T = 1$
κ_j	Checkpoint indices	$0 = \kappa_0 < \cdots < \kappa_J \le T, \quad J \ge 2$
π_j	Reference marginal	$\pi_j = \pi_{t_{\kappa_j}}, \quad \pi_{j-1} = \pi_j K_j^{\mathrm{F}}$
$\pi_j^\star$	Reward-tilted marginal	$\pi_j^\star(dx) = V_j(x)\pi_j(dx)/Z$
r	Terminal reward	$r : \mathcal{X} \to [0, \infty)$
V_j	Ideal process verifier	$V_j(x) := \mathbb{E}_{X_0 \sim (K_j^{\mathrm{F}} \cdots K_1^{\mathrm{F}})(\cdot \mid x)}[r(X_0)]$
$\widehat{V}_j$	Process verifier	$\widehat{V}_j := \widehat{V}_{t_{\kappa_j}}, \quad \widehat{V}_0 = r$
$Z, \mathcal{Z}$	Target and chain normalizers	$\begin{cases} Z = \int r \, d\pi_0, \\ \mathcal{Z} = \mu_0(\mathcal{X}) + 2\sum_{j=1}^{J-1} \mu_j(\mathcal{X}) + \mu_J(\mathcal{X}) \end{cases}$
$k_k^{\mathrm{F}}, k_{k-1}^{\mathrm{B}}$	Fine-grid reference kernels	$t_k \to t_{k-1}, \quad t_{k-1} \to t_k$
K_j^{F}	Checkpoint denoising kernel	$K_j^{\mathrm{F}}(dy \mid x), \quad j \to j-1$
K_{j-1}^{B}	Checkpoint renoising kernel	$K_{j-1}^{\mathrm{B}}(dx \mid y), \quad j-1 \to j$
$G_j^{\mathrm{F}}, G_j^{\mathrm{B}}$	Candidate proposal kernels	$Y_i \sim G_j^{\mathrm{F}}(\cdot \mid x), \quad Z_i \sim G_j^{\mathrm{B}}(\cdot \mid x)$
$J_j^{\mathrm{F}}, J_j^{\mathrm{B}}$	Reference edge laws	$\begin{cases} J_j^{\mathrm{F}}(dx, dy) = \pi_j(dx) K_j^{\mathrm{F}}(dy \mid x), \\ J_j^{\mathrm{B}}(dx, dy) = \pi_{j-1}(dy) K_{j-1}^{\mathrm{B}}(dx \mid y) \end{cases}$
ρ_j	Edge correction ratio	$\rho_j = dJ_j^{\mathrm{F}}/dJ_j^{\mathrm{B}}$
$K_{j-1}^{\mathrm{B},\rho}$	Corrected renoising kernel	$\rho_j(x, y) K_{j-1}^{\mathrm{B}}(dx \mid y)$
$\widehat{C}_j$	Implemented edge-conductance measure	$\begin{cases} \widehat{C}_j(dx, dy) = J_j^{\mathrm{F}}(dx, dy)\sqrt{\widehat{V}_j(x)}\sqrt{\widehat{V}_{j-1}(y)}, \\ \widehat{C}_1(dx, dy) = J_1^{\mathrm{F}}(dx, dy) r(y) \end{cases}$
C_j	Comparison edge-conductance measure	$C_j(dx, dy) = J_j^{\mathrm{F}}(dx, dy) V_{j-1}(y)$
F_j, B_j^ρ	Directional weight integrals	$F_j = K_j^{\mathrm{F}}\sqrt{\widehat{V}_{j-1}}, \quad B_j^\rho = K_j^{\mathrm{B},\rho}\sqrt{\widehat{V}_{j+1}}$
Q_j^{F}	Verifier-weighted denoising	$K_j^{\mathrm{F}}(dy \mid x)\sqrt{\widehat{V}_{j-1}(y)}/F_j(x)$
$Q_j^{\mathrm{B},\rho}$	Verifier-weighted renoising	$K_j^{\mathrm{B},\rho}(dy \mid x)\sqrt{\widehat{V}_{j+1}(y)}/B_j^\rho(x)$
$a_j^{\mathrm{F}}, a_j^{\mathrm{B}}$	Directional incident masses	$\begin{cases} a_j^{\mathrm{F}} = \sqrt{\widehat{V}_j} F_j \ (j \ge 2), \quad a_1^{\mathrm{F}} = F_1, \\ a_j^{\mathrm{B}} = \sqrt{\widehat{V}_j} B_j^\rho \ (1 \le j < J) \end{cases}$
μ_j	Total incident measure	$\begin{cases} \mu_j = \widehat{C}_j(\cdot, \mathcal{X}) + \widehat{C}_{j+1}(\mathcal{X}, \cdot), \\ \mu_0 = r\pi_0, \quad \mu_J = \widehat{C}_J(\cdot, \mathcal{X}) \end{cases}$
Π	Stationary law of CS-VGB	$\Pi(j, dx, d) = \mu_j(dx)/\mathcal{Z}$
P	Exact CS-VGB kernel	$\Pi P = \Pi$
χ	Directional persistence parameter	$\chi \in [0, 1]$
L_F, L_B	Candidate pool sizes	$L_F, L_B \in \mathbb{N}$

The square-root formulas for $\widehat{C}_j, F_j, Q_j^{\mathrm{F}}, a_j^{\mathrm{F}}$ apply to $2 \le j \le J$, and those for $B_j^\rho, Q_j^{\mathrm{B},\rho}, a_j^{\mathrm{B}}$ to $1 \le j < J$. At the clean edge, $F_1 = K_1^{\mathrm{F}} r$, $Q_1^{\mathrm{F}}(dy \mid x) = K_1^{\mathrm{F}}(dy \mid x) r(y)/F_1(x)$, and $Q_0^{\mathrm{B},\rho} = K_0^{\mathrm{B},\rho}$; see Equations (B.10), (B.12) and (B.13).

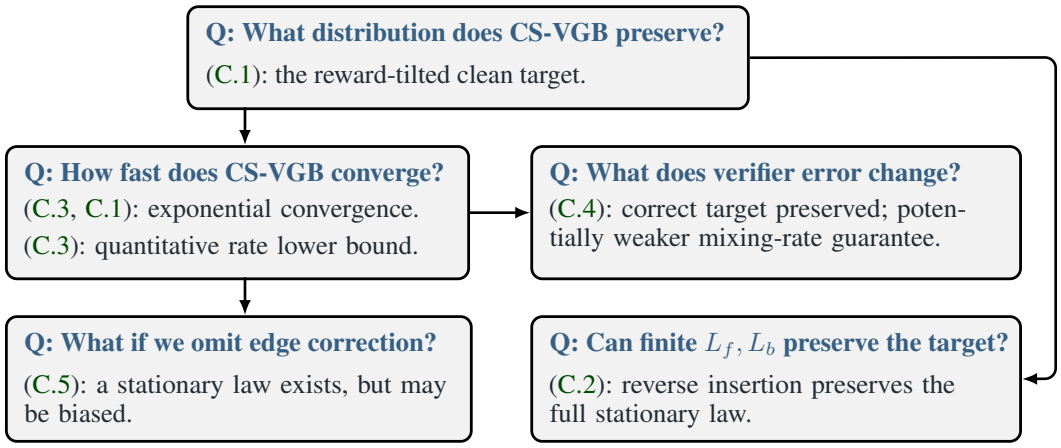

Figure A.1: Roadmap of the theoretical analysis of CS-VGB.

A.1 Reference Transitions across Continuous Generative Models

In this subsection, we review the reference processes underlying the checkpoint kernels in Appendix B.4. We frist fix the same time convention for Diffusion models, flow matching, and one-step posterior map: $t = 0$ is the data endpoint, $X_0 \sim \pi_{\text{data}}$, and $t = 1$ is the noise endpoint, $X_1 \sim \mathcal{N}(0, I)$. Generation runs from 1 to 0, consistent with the time convention in Appendix B.1.

Diffusion models. Score-based diffusion (Song et al., 2021) transports clean data toward noise through the forward SDE

$$dX_t = f(X_t, t)dt + g(t)dW_t, \qquad X_0 \sim \pi_{\text{data}},$$

where W_t is standard Brownian motion, $f(x, t)$ is the drift, and $g(t)$ is a scalar diffusion coefficient depending only on time. We write π_t^{ref} for the marginal law of X_t and $\pi_t^{\text{ref}}(x)$ for its density when it exists. These ideal marginals are distinct from the checkpoint marginals π_j induced by the implemented reference kernels in Appendix B.1. The corresponding reverse-time process is given by (Anderson, 1982)

$$dX_t = \left[f(X_t, t) - g(t)^2 \nabla_x \log \pi_t^{\text{ref}}(X_t)\right] dt + g(t)d\overline{W}_t,$$

where $\overline{W}_t$ is Brownian motion in decreasing time. For the VP (variance-preserving) SDE (Song et al., 2021), the noising marginal can be sampled as

$$X_t = \alpha_t X_0 + \sigma_t \epsilon, \qquad \epsilon \sim \mathcal{N}(0, I),$$

where α_t and σ_t are the signal and noise scales. Table A.2 gives the VP coefficients and a schedule with an exact standard Gaussian prior at $t = 1$.

Table A.2: VP SDE form and an endpoint-exact noise schedule.

General VP form	**Endpoint-exact schedule**
$f(x,t) = -\frac{1}{2}\beta(t)x, \quad g(t) = \sqrt{\beta(t)}$	$\beta(t) = \frac{2}{1-t}$
$\alpha_t = \exp\left(-\frac{1}{2}\int_0^t \beta(u)du\right)$	$\alpha_t = 1 - t$
$\sigma_t^2 = 1 - \alpha_t^2$	$\sigma_t^2 = 1 - (1-t)^2$

Here $\beta(t) > 0$ is the noise rate. An exact Gaussian prior requires $\alpha_1 = 0$ and hence $\int_0^1 \beta(t)dt = \infty$; otherwise some clean signal remains. For the schedule in the table, $\pi_t^{\text{ref}} \to \mathcal{N}(0, I)$ as $t \uparrow 1$, so the prior at $t = 1$ is defined by this limit.

A score network $s_\theta(x, t)$ approximates $\nabla_x \log \pi_t^{\text{ref}}(x)$ through denoising score matching (Song et al., 2021). Substituting it into the reverse process gives

$$dX_t = \left[f(X_t, t) - g(t)^2 s_\theta(X_t, t)\right] dt + g(t)d\overline{W}_t, \tag{A.1}$$

which is integrated from $t = 1$ to $t = 0$ using an SDE integrator such as Euler–Maruyama.

For $0 \le t < s \le 1$, let $q_{s|t}$ denote the forward-SDE transition from the cleaner state X_t to the noisier state X_s. For the VP SDE in Table A.2,

$$q_{s|t}(dx_s \mid x_t) = \mathcal{N}(dx_s; \alpha_{s|t} x_t, \sigma^2_{s|t} I), \qquad \alpha_{s|t} = \frac{\alpha_s}{\alpha_t}, \qquad \sigma^2_{s|t} = \sigma^2_s - \alpha^2_{s|t}\sigma^2_t.$$

Its *Bayes reverse* $q_{t|s}$ is the conditional law of the cleaner state given the noisier state and is determined by

$$\pi^{\mathrm{ref}}_t(dx_t) q_{s|t}(dx_s \mid x_t) = \pi^{\mathrm{ref}}_s(dx_s) q_{t|s}(dx_t \mid x_s).$$

With an exact score and exact integration of Equation (A.1), starting from $X_s = x_s$ and running to $t < s$ gives $X_t \sim q_{t|s}(\cdot \mid x_s)$. Score approximation and numerical discretization generally change this conditional distribution.

Flow matching. Flow matching (Lipman et al., 2023) learns an instantaneous velocity field for a probability path. For the linear interpolant,

$$X_t = (1-t)X_0 + t\epsilon, \qquad X_0 \sim \pi_{\mathrm{data}}, \quad \epsilon \sim \mathcal{N}(0, I).$$

with independent X_0 and ϵ, the sample-conditional velocity is $\epsilon - X_0$. Averaging over the endpoints given $X_t = x$ gives the marginal velocity

$$v_t(x) := \mathbb{E}[\partial_t X_t \mid X_t = x] = \mathbb{E}[\epsilon - X_0 \mid X_t = x],$$

A velocity network $v_\theta(t, x)$ approximates this conditional expectation. With the exact velocity, generation follows the ODE

$$dX_t = v_t(X_t)dt, \tag{A.2}$$

from $t = 1$ to $t = 0$. Its transitions are deterministic, but the same marginals π^{ref}_t can also be realized by the following forward and reverse SDEs with diffusion rate $\beta(t) > 0$ (Albergo et al., 2025):

$$dX_t = \left[v_t(X_t) + \frac{1}{2}\beta(t)\nabla_x \log \pi^{\mathrm{ref}}_t(X_t) \right] dt + \sqrt{\beta(t)} dW_t, \tag{FM-Noising SDE}$$

$$dX_t = \left[v_t(X_t) - \frac{1}{2}\beta(t)\nabla_x \log \pi^{\mathrm{ref}}_t(X_t) \right] dt + \sqrt{\beta(t)} d\overline{W}_t. \tag{FM-Denoising SDE}$$

The first runs in increasing time and the second in decreasing time, with $\overline{W}_t$ denoting backward-time Brownian motion.

Differentiating the Gaussian conditional density of $X_t \mid X_0 \sim \mathcal{N}((1-t)X_0, t^2 I)$ and averaging over $X_0 \mid X_t = x$ expresses the score through the conditional mean of ϵ. Using $\epsilon = X_t + (1-t)(\epsilon - X_0)$ then gives, for $t > 0$,

$$\nabla_x \log \pi^{\mathrm{ref}}_t(x) = -\frac{1}{t}\mathbb{E}[\epsilon \mid X_t = x] = -\frac{x + (1-t)v_t(x)}{t}.$$

Substituting this relation into Equation (FM-Denoising SDE) gives the velocity-only form

$$dX_t = \left[\frac{\beta(t)}{2t} X_t + \left(1 + \frac{(1-t)\beta(t)}{2t} \right) v_t(X_t) \right] dt + \sqrt{\beta(t)} d\overline{W}_t.$$

Replacing v_t by v_θ gives a practical reverse SDE without a separate score model.

One-step posterior sampling. For the linear interpolant above, a one-step posterior sampler transforms an independent Gaussian input into a clean sample conditional on $X_t = x$. In the ideal case,

$$Z = \Phi^\theta_{t,0}(x, \zeta), \qquad \zeta \sim \mathcal{N}(0, I), \qquad \mathrm{Law}(Z \mid x) = \mathrm{Law}(X_0 \mid X_t = x). \tag{A.3}$$

Meta Flow Maps (MFM) and Posterior Diamond Maps (PDM) learn such posterior samplers (Potaptchik et al., 2026; Holderrieth et al., 2026). Under perfect training at the population level, these methods recover the identity in Equation (A.3). Under this identity, independent Gaussian inputs give independent posterior draws, so for integrable reward r,

$$\frac{1}{N}\sum_{i=1}^{N} r(Z_i) \xrightarrow[N\to\infty]{\text{a.s.}} \mathbb{E}[r(X_0) \mid X_t = x], \qquad Z_i \overset{\text{iid}}{\sim} \mathrm{Law}(X_0 \mid X_t = x). \tag{A.4}$$

Thus posterior samples estimate the expected reward rather than the reward of a single deterministic prediction. We combine such a sample with a Gaussian bridge to construct denoising transitions in Equations (B.31) and (B.35).

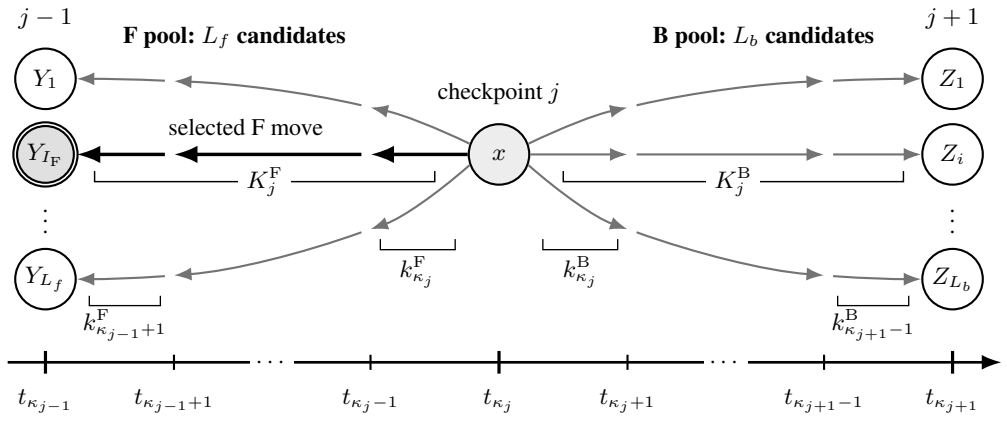

Figure B.1: Fine steps and checkpoint moves. Candidate paths connect neighboring checkpoints; the highlighted F move selects one of L_F candidates. The B pool contains L_B candidates in the opposite direction.

B Full CS-VGB Framework

Many inference-time alignment methods use intermediate reward estimates to guide generation from noise to data, then stop at a clean sample. A common approach evaluates the reward at a clean prediction, using a denoiser or a flow map (Li et al., 2025; Sabour et al., 2025). We call this intermediate reward estimate a *process verifier*. At checkpoint j, its ideal value is

$$V_j(x) := \mathbb{E}_{X_0 \sim (K_j^F \cdots K_1^F)(\cdot|x)}[r(X_0)],$$

where $K_j^F \cdots K_1^F$ composes the reference denoising transitions from checkpoint j to 0. The reward at a single clean prediction need not equal this expectation, so errors can affect subsequent choices and the final distribution. CS-VGB instead repeatedly renoises and denoises the current sample, using additional inference compute to revisit earlier choices rather than only generate independent samples.

Our sampling framework aims to preserve the *reward-tilted terminal distribution*, rather than merely maximize reward. The process verifier guides exploration, while the exact reward determines the target at the clean boundary. Under the conditions of Theorem C.1, the edge-corrected construction preserves the reward-tilted terminal distribution even with an imperfect process verifier: its accuracy affects transitions and mixing, but not this target distribution. We define the target in Appendix B.1 and the repeated denoising and renoising updates in Appendix B.2.

B.1 Problem Setup and Target Distribution

Reference process and target. Let $\mathcal{X} \subseteq \mathbb{R}^d$ be the sample space. We distinguish the *fine denoising time grid* from the *checkpoints* at which CS-VGB chooses an action:

$$0 = t_0 < \cdots < t_T = 1, \qquad 0 = \kappa_0 < \cdots < \kappa_J \le T, \qquad 2 \le J \le T. \tag{B.1}$$

Here k indexes fine-grid steps and j indexes checkpoints. In the CS-VGB construction and analysis, write $X_j := X_{t_{\kappa_j}}$. From checkpoint j, the kernels K_j^F and K_j^B move to checkpoints $j-1$ and $j+1$, respectively. When these moves follow the fine grid, they compose the corresponding steps as

$$K_j^F := k_{\kappa_j}^F k_{\kappa_j-1}^F \cdots k_{\kappa_{j-1}+1}^F, \qquad K_j^B := k_{\kappa_j}^B k_{\kappa_j+1}^B \cdots k_{\kappa_{j+1}-1}^B. \tag{B.2}$$

Here k_k^F is the denoising kernel from t_k to t_{k-1}, and k_k^B is the renoising kernel from t_k to t_{k+1}. Their model-specific forms are given in Appendix B.4. The F composition applies for $1 \le j \le J$ and the B composition for $0 \le j < J$. The composition integrates out all intermediate states, with no direction decision or candidate selection at the intermediate fine-grid points. Starting from the Gaussian prior $\pi_{t_T} := \mathcal{N}(0, I)$ at $t_T = 1$, denoising to t_{κ_J} gives the initial law π_J. Continuing through the checkpoints then defines the remaining marginals, so for $j = J, \ldots, 1$,

$$\pi_J := \pi_{t_T} k_T^F \cdots k_{\kappa_J+1}^F, \qquad \pi_{j-1}(dy) := \int_{\mathcal{X}} \pi_j(dx) K_j^F(dy \mid x). \tag{B.3}$$

The empty kernel product is the identity, so $\pi_J = \mathcal{N}(0, I)$ when $\kappa_J = T$. For a nonnegative reward r, the reward-tilted terminal distribution is

$$\pi_0^\star(dx) := \frac{r(x)\pi_0(dx)}{Z}, \qquad Z := \int_{\mathcal{X}} r(x)\pi_0(dx) \in (0, \infty).$$

Process verifier. The *ideal process verifier* is the expected reward after denoising to $t_0 = 0$:

$$V_j(x) := V_{t_{\kappa_j}}(x) = \mathbb{E}_{X_0 \sim (K_j^{\mathrm{F}} \cdots K_1^{\mathrm{F}})(\cdot|x)}[r(X_0)], \qquad V_0 = r, \qquad V_j = K_j^{\mathrm{F}} V_{j-1}.$$

The same kernels define the *reward-tilted time marginals*,

$$\pi_j^\star(dx) := \frac{V_j(x)\pi_j(dx)}{Z}, \qquad \int V_j \, d\pi_j = \int V_{j-1} \, d\pi_{j-1} = \cdots = \int r \, d\pi_0 = Z.$$

CS-VGB uses a nonnegative, finite approximation $\widehat{V}_j := \widehat{V}_{t_{\kappa_j}}$, obtained from a denoised prediction or a learned verifier. At the clean boundary it uses the exact reward:

Assumption B.1 (Terminal anchoring)**.** At $j = 0$, CS-VGB evaluates r exactly and sets $\widehat{V}_0 := r$.

When using a learned verifier, one way to train it is to regress on terminal rewards along reference denoising paths:

$$\mathcal{L}_{\mathrm{PV}}(\phi) := \mathbb{E}\Big[(\widehat{V}_\phi(t_{\kappa_j}, X_j) - r(X_0))^2\Big].$$

Here j is sampled over the supported checkpoints, and $(X_J, \ldots, X_0)$ follows $K_J^{\mathrm{F}}, \ldots, K_1^{\mathrm{F}}$ from π_J. In the ideal training, its population minimizer satisfies $\widehat{V}_\phi(t_{\kappa_j}, x) = \mathbb{E}[r(X_0) \mid X_j = x] = V_j(x)$ for π_j-almost every x. The construction below does not require this approximation to be exact.

B.2 CS-VGB at Selected Checkpoints

The state (j, x, d) records the checkpoint index, sample, and direction. Direction F denoises to $j - 1$, and direction B renoises to $j + 1$. At $j = 0$ only B is available, and at $j = J$ only F is available. We first define the exact update, then its candidate-pool implementation.

Corrected renoising. Let $K_{j-1}^{\mathrm{B}}(dx \mid y)$ be the reference renoising kernel from $j - 1$ to j. The reference kernels K_j^{F} and K_{j-1}^{B} need not be exact reverses, meaning that the joint edge laws in Equation (B.4) need not satisfy $J_j^{\mathrm{F}} = J_j^{\mathrm{B}}$. These laws are defined by

$$J_j^{\mathrm{F}}(dx, dy) := \pi_j(dx) K_j^{\mathrm{F}}(dy \mid x), \qquad J_j^{\mathrm{B}}(dx, dy) := \pi_{j-1}(dy) K_{j-1}^{\mathrm{B}}(dx \mid y). \tag{B.4}$$

Assume $J_j^{\mathrm{F}} \ll J_j^{\mathrm{B}}$. The *edge ratio* $\rho_j := dJ_j^{\mathrm{F}}/dJ_j^{\mathrm{B}}$ corrects their mismatch:

$$K_{j-1}^{\mathrm{B},\rho}(dx \mid y) := \rho_j(x, y) K_{j-1}^{\mathrm{B}}(dx \mid y). \tag{B.5}$$

We remark that $K_{j-1}^{\mathrm{B},\rho}$ has total mass one for π_{j-1}-almost every y. Specifically, for every measurable $A \subseteq \mathcal{X}$,

$$\int_A \pi_{j-1}(dy) \int_{\mathcal{X}} K_{j-1}^{\mathrm{B},\rho}(dx \mid y) = \int_{\mathcal{X} \times A} \rho_j \, dJ_j^{\mathrm{B}} = J_j^{\mathrm{F}}(\mathcal{X}, A) = \pi_{j-1}(A). \tag{B.6}$$

Equation (B.6) verifies this normalization. Moreover, the correction in Equation (B.5) makes the two joint edge laws equal, giving

$$\pi_j(dx) K_j^{\mathrm{F}}(dy \mid x) = \pi_{j-1}(dy) K_{j-1}^{\mathrm{B},\rho}(dx \mid y). \tag{B.7}$$

Remark B.1 (Exact pairing across model classes)**.** In this remark, we clarify when $\rho_j = 1$ holds for diffusion models, flow matching, and one-step posterior maps. For VP diffusion, the noising process and learned denoising process are

$$\begin{aligned} \text{Noising:} \quad & dX_t = -\tfrac{1}{2}\beta(t)X_t\,dt + \sqrt{\beta(t)}\,dW_t, \\ \text{Denoising:} \quad & dX_t = \big[-\tfrac{1}{2}\beta(t)X_t - \beta(t)s_\theta(X_t,t)\big]dt + \sqrt{\beta(t)}\,d\overline{W}_t. \end{aligned}$$

For linear flow matching, converting the ODE to the SDEs in Equations (FM-Noising SDE) and (FM-Denoising SDE) with $\beta(t) = 2t/(1-t)$ gives

$$\begin{aligned} \text{Noising:} \quad & dX_t = -\frac{X_t}{1-t}\,dt + \sqrt{\frac{2t}{1-t}}\,dW_t, \\ \text{Denoising:} \quad & dX_t = \left[\frac{X_t}{1-t} + 2v_\theta(t,X_t)\right]dt + \sqrt{\frac{2t}{1-t}}\,d\overline{W}_t. \end{aligned}$$

In both cases, noising runs in increasing time and denoising in decreasing time. Under ideal training, $s_\theta(x,t) = \nabla_x \log \pi_t^{\mathrm{ref}}(x)$ or $v_\theta(t,x) = v_t(x)$, respectively, so exact integration of the denoising and renoising SDEs yields the paired-transition identity below for $\pi_j = \pi^{\mathrm{ref}}_{t_{\kappa_j}}$:

$$\pi_j(dx)K_j^{\mathrm{F}}(dy \mid x) = \pi_{j-1}(dy)K_{j-1}^{\mathrm{B}}(dx \mid y), \qquad \rho_j = 1.$$

For a one-step posterior map, such as MFM or PDM (Potaptchik et al., 2026; Holderrieth et al., 2026), an exact clean posterior draw combined with the Gaussian bridge in Proposition B.1 yields the conditional law $\mathrm{Law}(X_{j-1} \mid X_j = x)$ of the chosen noising process and hence gives $\rho_j = 1$. The implemented kernels are given in Appendix B.4.

Edge-conductance measures and sampling transitions. Starting from the common edge law in Equation (B.7), we weight adjacent state pairs by the process verifiers on interior edges and only the terminal reward on the clean edge. This defines an *edge-conductance measure* $\widehat{C}_j(dx,dy)$, extending the edge weights of a weighted-network random walk to continuous states (Doyle & Snell, 1984; Bezuglyi & Jorgensen, 2018). As illustrated in Figure B.2, normalizing this measure over y with x fixed gives the denoising transition, while normalizing over x with y fixed gives the renoising transition, and the corresponding total weights define the directional masses. Thus, for $j = 2, \dots, J$,

$$\widehat{C}_j(dx,dy) := \pi_j(dx)K_j^{\mathrm{F}}(dy \mid x)\sqrt{\widehat{V}_j(x)\widehat{V}_{j-1}(y)} = \pi_{j-1}(dy)K_{j-1}^{\mathrm{B},\rho}(dx \mid y)\sqrt{\widehat{V}_{j-1}(y)\widehat{V}_j(x)}. \tag{B.8}$$

We remark that this geometric-mean weighting is related to square-root balancing in locally balanced MCMC (Zanella, 2020) and square-root approximations of transition rates in Markov state models (Lie et al., 2013).

On the clean edge, we use only the terminal reward:

$$\widehat{C}_1(dx,dy) := \pi_1(dx)K_1^{\mathrm{F}}(dy \mid x)r(y) = \pi_0(dy)K_0^{\mathrm{B},\rho}(dx \mid y)r(y). \tag{B.9}$$

In particular, $\widehat{C}_1(\mathcal{X},dy) = r(y)\pi_0(dy) = Z\pi_0^\star(dy)$. This choice fixes the target tilted distribution at the clean boundary without requiring $\widehat{V}_j = V_j$ at intermediate grid points.

To obtain a sampling transition, hold the current endpoint fixed and normalize the weights over the other endpoint. For $2 \le j \le J$, this gives the denoising transition

$$F_j(x) := \int_{\mathcal{X}} K_j^{\mathrm{F}}(dy \mid x)\sqrt{\widehat{V}_{j-1}(y)}, \qquad Q_j^{\mathrm{F}}(dy \mid x) := \frac{K_j^{\mathrm{F}}(dy \mid x)\sqrt{\widehat{V}_{j-1}(y)}}{F_j(x)}. \tag{B.10}$$

On the clean edge, the corresponding formulas use the reward itself:

$$F_1(x) := \int_{\mathcal{X}} K_1^{\mathrm{F}}(dy \mid x)r(y), \qquad Q_1^{\mathrm{F}}(dy \mid x) := \frac{K_1^{\mathrm{F}}(dy \mid x)r(y)}{F_1(x)}.$$

On the same edge, the corrected renoising transition goes from y at checkpoint index $j-1$ to x at checkpoint index j. For $2 \le j \le J$, define

$$B^{\rho}_{j-1}(y) := \int_{\mathcal{X}} K^{\mathrm{B},\rho}_{j-1}(dx \mid y)\sqrt{\widehat{V}_j(x)}, \qquad Q^{\mathrm{B},\rho}_{j-1}(dx \mid y) := \frac{K^{\mathrm{B},\rho}_{j-1}(dx \mid y)\sqrt{\widehat{V}_j(x)}}{B^{\rho}_{j-1}(y)}. \tag{B.11}$$

Here $F_j(x)$ and $B^{\rho}_{j-1}(y)$ are the corresponding total masses; we require them to be positive and finite wherever an update is performed. The clean-edge definition instead gives

$$Q^{\mathrm{B},\rho}_0(dx \mid y) := K^{\mathrm{B},\rho}_0(dx \mid y). \tag{B.12}$$

Thus the first renoising move from a clean sample has no process-verifier term, but follows the same construction of conditioning the edge measure on the current endpoint. Figure B.2 illustrates how this construction gives the denoising and renoising transitions on an interior edge.

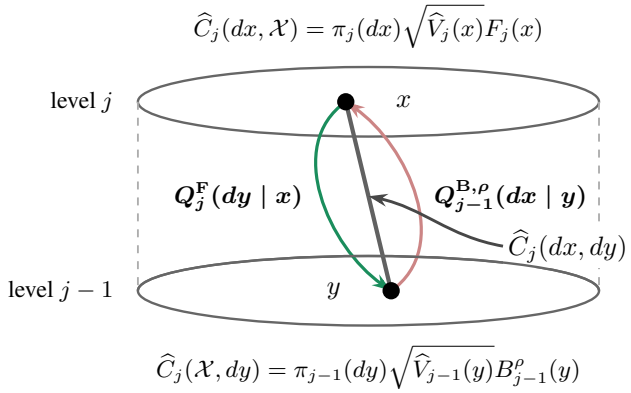

Figure B.2: Reward-tilted edge-conductance measure between adjacent checkpoints ($j \ge 2$). The joint measure $\widehat{C}_j(dx, dy)$ assigns mass to state pairs (x, y) at levels j and $j-1$. Its marginals $\widehat{C}_j(dx, \mathcal{X})$ and $\widehat{C}_j(\mathcal{X}, dy)$ give the incident masses at x and y, while conditioning the same joint measure in the two directions yields the transition kernels $Q^{\mathrm{F}}_j(dy \mid x)$ and $Q^{\mathrm{B},\rho}_{j-1}(dx \mid y)$.

Direction probabilities and directional persistence. At an interior checkpoint $1 \le j < J$, integrating the adjacent edge-conductance measures gives

$$\widehat{C}_j(dx, \mathcal{X}) = \pi_j(dx)a^{\mathrm{F}}_j(x), \qquad \widehat{C}_{j+1}(\mathcal{X}, dx) = \pi_j(dx)a^{\mathrm{B}}_j(x),$$

where

$$a^{\mathrm{F}}_j(x) := \begin{cases} F_1(x), & j = 1, \\ \sqrt{\widehat{V}_j(x)}F_j(x), & j \ge 2, \end{cases} \qquad a^{\mathrm{B}}_j(x) := \sqrt{\widehat{V}_j(x)}B^{\rho}_j(x). \tag{B.13}$$

The kernels Q^{F}_j and $Q^{\mathrm{B},\rho}_j$ specify where to move once a direction is chosen. To choose between them, we use the relative masses of the two adjacent edges, assigning more probability to the direction with larger mass. This gives the probabilities below; their smaller value, $\ell_j(x)$, will be used to control direction changes:

$$p^{\mathrm{F}}_j(x) := \frac{a^{\mathrm{F}}_j(x)}{a^{\mathrm{F}}_j(x) + a^{\mathrm{B}}_j(x)}, \qquad p^{\mathrm{B}}_j(x) := \frac{a^{\mathrm{B}}_j(x)}{a^{\mathrm{F}}_j(x) + a^{\mathrm{B}}_j(x)}, \qquad \ell_j(x) := \min\{p^{\mathrm{F}}_j(x), p^{\mathrm{B}}_j(x)\}. \tag{B.14}$$

These probabilities reflect the local preference for denoising or renoising, but do not account for the previous direction. Choosing a new direction from them at every update can therefore cause frequent reversals, especially when the two masses are comparable. Such reversals repeatedly undo changes in depth and can lead to diffusive progress across checkpoints. Since a practical run has a finite compute budget, the time to reach the clean boundary $j = 0$ matters alongside convergence to stationarity. For example, using the same momentum structure adopted below, MDM-VGB reduces the expected

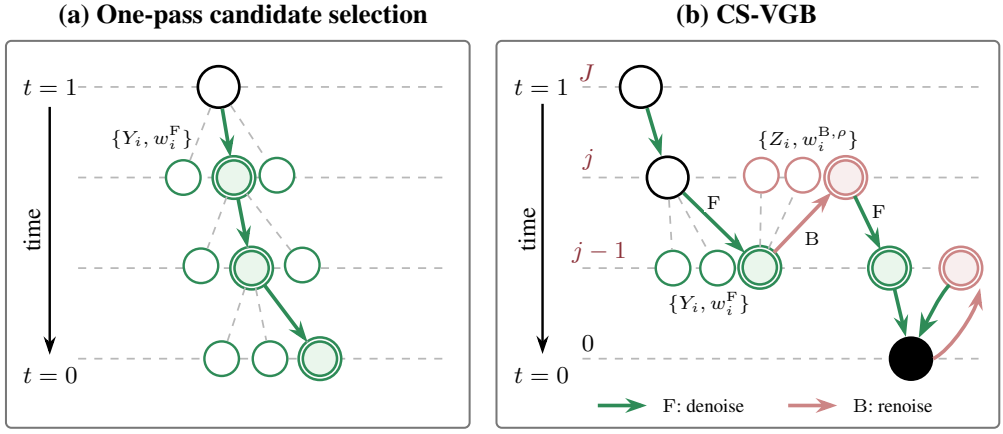

Figure B.3: Algorithm comparison of CS-VGB. (a) One-pass candidate selection, following SVDD (Li et al., 2025), selects one candidate at each denoising step and stops at the first clean output. (b) CS-VGB may instead renoise and repeat the search, while the exact terminal reward anchors every visit to $t = 0$.

first-leaf hitting time from $\Theta(n^2)$ to $O(n)$ when the forward/backward probability imbalance is $O(1/n)$ and $1 - \chi = O(1/n)$, where n is the number of positions to reveal in masked diffusion models (Jeon et al., 2026, Theorem 4.5). In our framework, Proposition C.2 bounds the first hitting time under an average directional-imbalance condition and an exact reset at the top checkpoint.

Motivated by this behavior, CS-VGB retains the current direction d and uses *crossing-flow cancellation* (Hayes & Sinclair, 2010), as in VGB-Momentum (Rohatgi et al., 2026; Jeon et al., 2026). The update in Table B.1 removes the common probability $\chi \ell_j(x)$ from both opposing flips and adds it to the respective stays, leaving the move probabilities unchanged. Increasing the *directional persistence parameter* $\chi \in [0, 1]$ strengthens this cancellation and suppresses more direction changes, while preserving the stationary law established in Theorem C.1.

Table B.1: Momentum update at an interior checkpoint.

Current direction	Action	Probability	Next state
F	Move	$p_j^{\mathrm{F}}(x)$	$(j-1, Y, \mathrm{F}), Y \sim Q_j^{\mathrm{F}}(\cdot \mid x)$
F	Flip	$p_j^{\mathrm{B}}(x) - \chi \ell_j(x)$	(j, x, B)
F	Stay	$\chi \ell_j(x)$	(j, x, F)
B	Move	$p_j^{\mathrm{B}}(x)$	$(j+1, Y, \mathrm{B}), Y \sim Q_j^{\mathrm{B},\rho}(\cdot \mid x)$
B	Flip	$p_j^{\mathrm{F}}(x) - \chi \ell_j(x)$	(j, x, F)
B	Stay	$\chi \ell_j(x)$	(j, x, B)

A move changes the checkpoint index and sample, a flip changes only the direction, and a stay changes neither. The default $\chi = 1$ gives

$$\mathbb{P}(\mathrm{F} \to \mathrm{B}) = \frac{(a_j^{\mathrm{B}} - a_j^{\mathrm{F}})_+}{a_j^{\mathrm{F}} + a_j^{\mathrm{B}}}, \qquad \mathbb{P}(\mathrm{B} \to \mathrm{F}) = \frac{(a_j^{\mathrm{F}} - a_j^{\mathrm{B}})_+}{a_j^{\mathrm{F}} + a_j^{\mathrm{B}}}. \tag{B.15}$$

At the boundaries, exit with probability one using $Q_0^{\mathrm{B},\rho}$ at $j = 0$ and Q_J^{F} at $j = J$. On arrival at $j = 0$, set $d = \mathrm{B}$; on arrival at $j = J$, set $d = \mathrm{F}$. These rules define the *exact one-step kernel* P when the masses and conditional draws are evaluated exactly. Under the conditions of Theorem C.1, P preserves a stationary law Π whose distribution at the clean boundary is $\pi_0^\star$.

B.3 Finite-Candidate Implementation and Practical Refinements

The preceding subsection assumes exact directional masses $a_j^{\mathrm{F}}(x), a_j^{\mathrm{B}}(x)$ and exact draws from the conditional transitions $Q_j^{\mathrm{F}}(\cdot \mid x), Q_j^{\mathrm{B},\rho}(\cdot \mid x)$. In practice, we approximate these quantities using

L_F denoising and L_B renoising candidates at each checkpoint. To avoid the stationary bias that this approximation can introduce, we retain the pools and use *reverse insertion* after each move. Building on the retained pools, *lazy pool evaluation* evaluates candidates in the direction opposite to the current one only as needed, reusing previously computed values to reduce the number of function evaluations (NFE) while remaining compatible with momentum. Beyond this per-update cost, practical performance also depends on the clean samples collected within a finite update budget. Since excessive backtracking can limit clean visits, we introduce *depth-dependent weights* on edge-conductance measures to favor denoising. Together, these refinements preserve the stationary clean target under Theorem C.2.

Candidate weights. The state retains the *candidate pools* $\mathbf{u} = (Y_{1:L_F}, Z_{1:L_B})$ alongside (j, x, d), with fixed $L_F, L_B \geq 2$. Fresh candidates are drawn independently from $G_j^{\mathrm{F}}(\cdot \mid x)$ and $G_j^{\mathrm{B}}(\cdot \mid x)$, with weights

$$w_{j,i}^{\mathrm{F}} := \frac{dK_j^{\mathrm{F}}(\cdot \mid x)}{dG_j^{\mathrm{F}}(\cdot \mid x)}(Y_i) \begin{cases} r(Y_i), & j = 1, \\ \sqrt{\widehat{V}_{j-1}(Y_i)}, & j \geq 2, \end{cases}$$
$$w_{j,i}^{\mathrm{B},\rho} := \rho_{j+1}(Z_i, x) \frac{dK_j^{\mathrm{B}}(\cdot \mid x)}{dG_j^{\mathrm{B}}(\cdot \mid x)}(Z_i) \begin{cases} 1, & j = 0, \\ \sqrt{\widehat{V}_{j+1}(Z_i)}, & j \geq 1. \end{cases} \tag{B.16}$$

At $j = 0$, only the B pool is retained; at $j = J$, only the F pool is retained.

Depth weighting and direction probabilities. For *depth weighting*, fix $\alpha > 0$ and adjust the relative frequency of denoising and renoising through

$$\widehat{C}_j^{\alpha} := \alpha^{j-1} \widehat{C}_j, \qquad j = 1, \ldots, J. \tag{B.17}$$

This leaves the conditional transitions unchanged, but reweights the directional masses. Choosing $\alpha < 1$ favors F and $\alpha > 1$ favors B. Including the current-state factors gives the *candidate contributions*

$$e_{j,i}^{\mathrm{F}} := \alpha^{j-1} w_{j,i}^{\mathrm{F}} \begin{cases} 1, & j = 1, \\ \sqrt{\widehat{V}_j(x)}, & j \geq 2, \end{cases} \qquad e_{j,i}^{\mathrm{B}} := \alpha^{j} w_{j,i}^{\mathrm{B},\rho} \begin{cases} r(x), & j = 0, \\ \sqrt{\widehat{V}_j(x)}, & j \geq 1. \end{cases} \tag{B.18}$$

Their pool averages determine the direction probabilities:

$$\widehat{a}_j^d := \frac{1}{L_d} \sum_{i=1}^{L_d} e_{j,i}^d, \qquad \widehat{p}_j^d := \frac{\widehat{a}_j^d}{\widehat{a}_j^{\mathrm{F}} + \widehat{a}_j^{\mathrm{B}}}, \qquad \widehat{\ell}_j := \min\{\widehat{p}_j^{\mathrm{F}}, \widehat{p}_j^{\mathrm{B}}\}, \tag{B.19}$$

Unavailable masses are zero. Evaluating both pools and substituting these probabilities into Table B.1 determines the move, flip, and stay probabilities.

Reverse insertion. With *reverse insertion*, a move selects a candidate in direction d with probability $e_{j,i}^d / \sum_{h=1}^{L_d} e_{j,h}^d$, then inserts the previous x uniformly into the opposite pool at the destination:

$$\begin{aligned} \mathrm{F}: &\quad (j, x) \longrightarrow (j-1, Y_I), \quad K \sim \mathrm{Unif}([L_B]), \quad Z'_K = x, \\ \mathrm{B}: &\quad (j, x) \longrightarrow (j+1, Z_I), \quad K \sim \mathrm{Unif}([L_F]), \quad Y'_K = x. \end{aligned} \tag{B.20}$$

All other destination entries are drawn independently from their proposals (Figure B.4). When combined with lazy pool evaluation, the inserted candidate is also designated as the opposite *mark*.

Lazy pool evaluation. With *lazy pool evaluation*, an F update uses the full F pool and its B *mark* Z_K; a B update uses the full B pool and its F *mark* Y_I. In direction F, draw $I \sim \mathrm{Unif}([L_F])$ with the B mark K fixed; in direction B, draw $K \sim \mathrm{Unif}([L_B])$ with the F mark I fixed. Conditional on these indices, the probabilities of $(\mathsf{move}, \mathsf{stay}, \mathsf{flip})$ are

$$\begin{aligned} \mathrm{F}: &\quad \frac{(\widehat{a}_j^{\mathrm{F}},\ \chi \min\{e_{j,I}^{\mathrm{F}}, e_{j,K}^{\mathrm{B}}\},\ e_{j,K}^{\mathrm{B}} - \chi \min\{e_{j,I}^{\mathrm{F}}, e_{j,K}^{\mathrm{B}}\})}{\widehat{a}_j^{\mathrm{F}} + e_{j,K}^{\mathrm{B}}}, \\ \mathrm{B}: &\quad \frac{(\widehat{a}_j^{\mathrm{B}},\ \chi \min\{e_{j,K}^{\mathrm{B}}, e_{j,I}^{\mathrm{F}}\},\ e_{j,I}^{\mathrm{F}} - \chi \min\{e_{j,K}^{\mathrm{B}}, e_{j,I}^{\mathrm{F}}\})}{\widehat{a}_j^{\mathrm{B}} + e_{j,I}^{\mathrm{F}}}. \end{aligned} \tag{B.21}$$

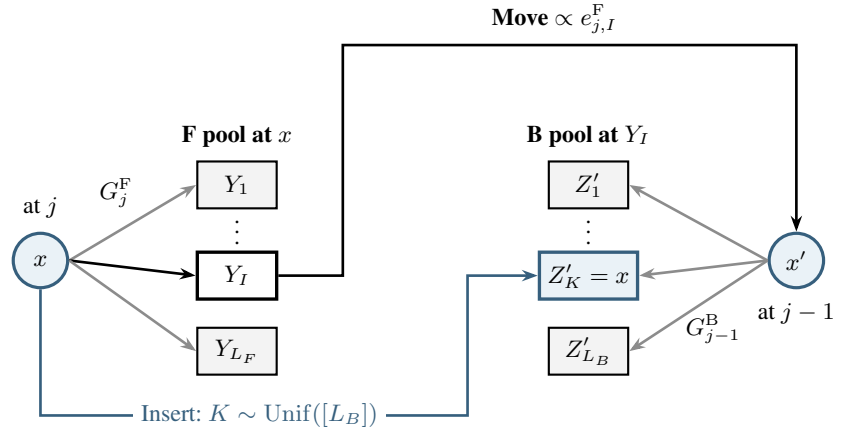

Figure B.4: Reverse insertion for an F move. Select Y_I and insert the previous x into a uniform B slot. All remaining entries are freshly drawn.

On an F-to-B flip, mark Y_I in the retained F pool and complete the B pool by drawing only missing entries from $G^{\mathrm{B}}_j(\cdot \mid x)$, keeping Z_K and all previously generated candidates and values. The B-to-F case is symmetric. On a move, discard the uniform index, independently select a candidate by weight, and apply reverse insertion (Equation (B.20)), marking the inserted state at an interior destination. A stay changes nothing. Figure B.5 illustrates the move and flip cases.

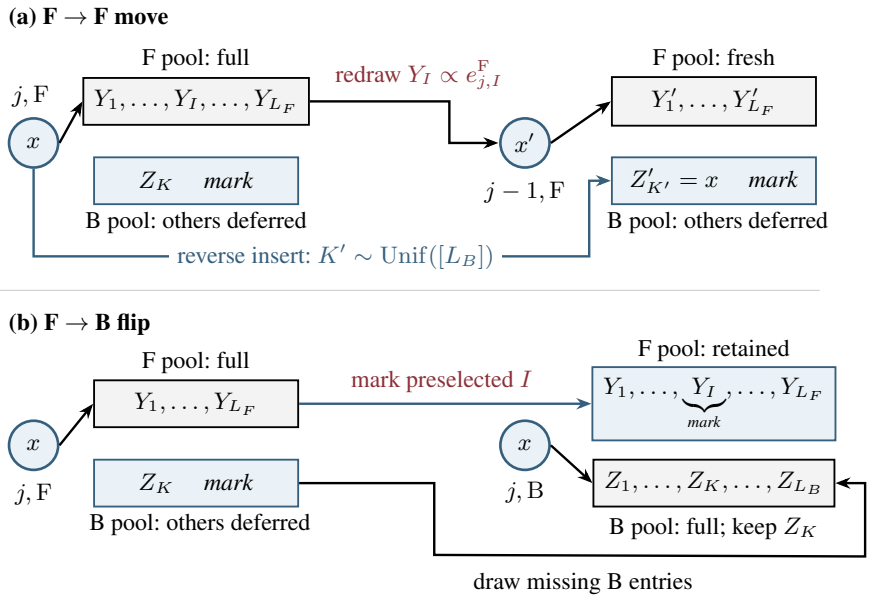

Figure B.5: Lazy pool evaluation from direction F, shown when only the B mark has been generated. First draw $I \sim \mathrm{Unif}(\{1, \ldots, L_F\})$, then choose move, stay, or flip using the probabilities in Equation (B.21), conditioned on I and the index K of the existing B mark. (a) An F move discards that index and independently selects Y_I by weight, creates the destination pools, and reverse-inserts x as the new B mark. Panel (a) assumes an interior destination ($j \geq 2$); for $j = 1$, arrival at checkpoint 0 resets the direction to B and retains only the B pool, with no separate mark. (b) An F-to-B flip leaves (j, x) unchanged, fills missing B entries, and marks the preselected Y_I in the retained F pool. Already generated candidates and values are reused. Stay changes nothing and reuses the same pools and mark.

Omitting the edge correction. When ρ_j is unavailable, an implementation may set $\rho_j = 1$. We measure its effect through the average reference-edge error in Equation (C.98). Under the reward

Algorithm 1: Finite-candidate single-chain CS-VGB.

Input: starting law π_J; checkpoint count J; proposal kernels G_j^d; pool sizes $L_F, L_B \geq 2$; update budget U; depth weight $\alpha > 0$; directional persistence $\chi \in [0, 1]$; lazy evaluation on/off

$x \sim \pi_J$; $(j, d) \leftarrow (J, \mathrm{F})$; $\mathcal{S} \leftarrow \varnothing$
for $h = 0, \ldots, U - 1$ **do**
  if $j = 0$ **then** Append x to $\mathcal{S}$
  Complete the current-direction pool if lazy, or all available pools otherwise; evaluate only new candidates
  if $j \in \{0, J\}$ **then**
    a $\leftarrow$ move
  else if *lazy* **then**
    $I \sim \mathrm{Unif}([L_F])$ if $d = \mathrm{F}$; $K \sim \mathrm{Unif}([L_B])$ if $d = \mathrm{B}$
    Sample a via Equation (B.21) given (I, K); opposite mark fixed
  else
    Sample a via Table B.1 with $(\widehat{p}_j^d, \widehat{\ell}_j)$ from Equation (B.19)
  end
  Apply a:

a = Stay	a = Flip
No change.	**if lazy**: mark Y_I (F) or Z_K (B) using the preselected index; fill missing opposite-pool entries. $d : \mathrm{F} \leftrightarrow \mathrm{B}$; retain pools.
a = F move	**a = B move**
Independently reselect $I \in [L_F]$ using $\mathbb{P}(I = i) = e_{j,i}^{\mathrm{F}} / \sum_h e_{j,h}^{\mathrm{F}}$	Independently reselect $K \in [L_B]$ using $\mathbb{P}(K = k) = e_{j,k}^{\mathrm{B}} / \sum_h e_{j,h}^{\mathrm{B}}$
$(j', x') \leftarrow (j - 1, Y_I)$	$(j', x') \leftarrow (j + 1, Z_K)$
$K' \sim \mathrm{Unif}([L_B])$	$I' \sim \mathrm{Unif}([L_F])$
$Z'_{K'} \leftarrow x$	$Y'_{I'} \leftarrow x$
if lazy and $j' > 0$: mark $Z'_{K'}$	**if lazy and** $j' < J$: mark $Y'_{I'}$

  if a = move **then**
    $(j, x) \leftarrow (j', x')$; reset pools, retaining only the inserted candidate; evaluate its weight
    if $j \in \{0, J\}$ **then** $d \leftarrow \mathrm{B}$ at 0, F at J; inward pool only, no mark
  end
end
if $j = 0$ **then** Append x to $\mathcal{S}$
return $\mathcal{S}$

and verifier bounds in Equation (C.69) and the transition-density condition in Equation (C.100), Theorem C.5 bounds the resulting stationary bias.

B.4 Model-Specific Checkpoint Transitions

We now describe how to construct the reference kernels K_j^{F} and K_j^{B} for diffusion models, flow matching, and one-step posterior samplers. For diffusion and flow matching, denoising composes fine-grid steps as in Equation (B.2), while renoising uses Gaussian transitions across each checkpoint interval. A one-step posterior sampler instead provides a direct denoising transition through a Gaussian bridge.

The reference steps below also serve as unguided proposals, $g_k^{\mathrm{F}} = k_k^{\mathrm{F}}$. For a Gaussian reference step $k_k^{\mathrm{F}}(\cdot \mid x) = \mathcal{N}(m_k(x), \Sigma_k)$, a positive differentiable verifier gives the mean-shift proposal

$$g_k^{\mathrm{F}}(dy \mid x) = \mathcal{N}\Big(dy; m_k(x) + s_k \Sigma_k \nabla \log \widehat{V}_{t_{k-1}}(m_k(x)), \Sigma_k\Big), \tag{B.22}$$

where $s_k \geq 0$ is the guidance strength. The reference kernel remains unchanged; the importance weights account for the proposal change.

Diffusion models. Consider the VP SDE with noise-rate schedule $\beta(t) > 0$, for which $f(x,t) = -\beta(t)x/2$ and $g(t) = \sqrt{\beta(t)}$. Since the coefficients may diverge at $t = 1$, we can choose $\kappa_J < T$ and restrict the checkpoint updates to $[0, t_{\kappa_J}]$, below the fully noisy endpoint. The initial law π_J is then the reference marginal at t_{κ_J} defined in Equation (B.3). The learned reverse drift in Equation (A.1) is $-\beta(t)[x/2 + s_\theta(x,t)]$. Let $\Delta_k := t_k - t_{k-1}$. For $1 \le k \le \kappa_J$, one Euler–Maruyama step from t_k to t_{k-1} gives the fine-grid reference denoising kernel

$$k_k^{\mathrm{F}}(dx_{k-1} \mid x_k) := \mathcal{N}\left(dx_{k-1}; x_k + \Delta_k\beta(t_k)\left[\frac{x_k}{2} + s_\theta(x_k, t_k)\right], \Delta_k\beta(t_k)I\right). \tag{B.23}$$

Recalling the conditional VP scales $\alpha_{t|s} := \exp(-\frac{1}{2}\int_s^t \beta(u)du)$ and $\sigma_{t|s}^2 := 1 - \alpha_{t|s}^2$ for $0 \le s < t \le 1$, the reference renoising kernel over the same time interval is

$$k_{k-1}^{\mathrm{B}}(dx_k \mid x_{k-1}) := \mathcal{N}\left(dx_k; \alpha_{t_k|t_{k-1}}x_{k-1}, \sigma_{t_k|t_{k-1}}^2 I\right). \tag{B.24}$$

Composing the fine-grid noising steps gives the checkpoint renoising kernel

$$K_{j-1}^{\mathrm{B}}(dx \mid y) = \mathcal{N}\left(dx; \alpha_{t_{\kappa_j}|t_{\kappa_{j-1}}}y, \sigma_{t_{\kappa_j}|t_{\kappa_{j-1}}}^2 I\right). \tag{B.25}$$

We remark that composing Gaussian denoising steps generally gives a non-Gaussian checkpoint kernel K_j^{F}. The edge correction ρ_j accounts for any mismatch between the denoising and renoising edge laws.

Flow matching. The flow-matching ODE gives a deterministic transition between grid points, whose matched reverse is the exact inverse solution map. When this inverse is not available as an exact finite-step transition, the two edge laws can have different supports and their density ratio can be singular. We therefore use the stochastic SDE pair in Equations (FM-Noising SDE) and (FM-Denoising SDE). For the linear interpolant $X_t = (1-t)X_0 + t\epsilon$, choosing $\beta(t) = 2t/(1-t)$ gives the renoising SDE (increasing t) and the denoising SDE (decreasing t), respectively,

$$dX_t = -\frac{X_t}{1-t}dt + \sqrt{\frac{2t}{1-t}}dW_t, \qquad dX_t = \left[\frac{X_t}{1-t} + 2v_t(X_t)\right]dt + \sqrt{\frac{2t}{1-t}}d\overline{W}_t. \tag{B.26}$$

Replacing v_t by the pretrained velocity v_θ, one Euler–Maruyama discretization step of the denoising SDE from t_k to t_{k-1} gives the reference denoising kernel

$$k_k^{\mathrm{F}}(dx_{k-1} \mid x_k) := \mathcal{N}\left(dx_{k-1}; x_k - \Delta_k\left[\frac{x_k}{1-t_k} + 2v_\theta(t_k, x_k)\right], \frac{2t_k\Delta_k}{1-t_k}I\right). \tag{B.27}$$

The increasing-time SDE is linear, and its exact one-step noising kernel is

$$k_{k-1}^{\mathrm{B}}(dx_k \mid x_{k-1}) := \mathcal{N}\left(dx_k; \frac{1-t_k}{1-t_{k-1}}x_{k-1}, \left[t_k^2 - \left(\frac{1-t_k}{1-t_{k-1}}\right)^2 t_{k-1}^2\right]I\right). \tag{B.28}$$

A checkpoint F move composes the fine steps in Equation (B.27); the B reference kernel over the whole interval is

$$K_{j-1}^{\mathrm{B}}(dx \mid y) = \mathcal{N}\left(dx; \frac{1-t_{\kappa_j}}{1-t_{\kappa_{j-1}}}y, \left[t_{\kappa_j}^2 - \frac{(1-t_{\kappa_j})^2}{(1-t_{\kappa_{j-1}})^2}t_{\kappa_{j-1}}^2\right]I\right). \tag{B.29}$$

Thus an F proposal can require several model evaluations, while this B reference draw does not require numerical substeps.

One-step posterior sampling. We use the Gaussian reference renoising kernel K_{j-1}^{B} in Equation (B.29) and construct its corresponding denoising kernel K_j^{F} using a one-step posterior sampler. To denoise from t_{κ_j} to $t_{\kappa_{j-1}}$, we first draw a clean endpoint Z conditional on $X_j = x$ using MFM or PDM (Potaptchik et al., 2026; Holderrieth et al., 2026). We then sample X_{j-1} through a Gaussian bridge conditioned on both endpoints (Figure B.6), without evaluating intermediate fine-grid states. For $0 \le t_{\kappa_{j-1}} < t_{\kappa_j} \le 1$, define

$$a_j := \frac{1-t_{\kappa_j}}{1-t_{\kappa_{j-1}}}, \qquad \sigma_j^2 := t_{\kappa_j}^2 - a_j^2 t_{\kappa_{j-1}}^2, \qquad m_j(x,z) := \frac{a_j t_{\kappa_{j-1}}^2}{t_{\kappa_j}^2}x + \frac{(1-t_{\kappa_{j-1}})\sigma_j^2}{t_{\kappa_j}^2}z. \tag{B.30}$$

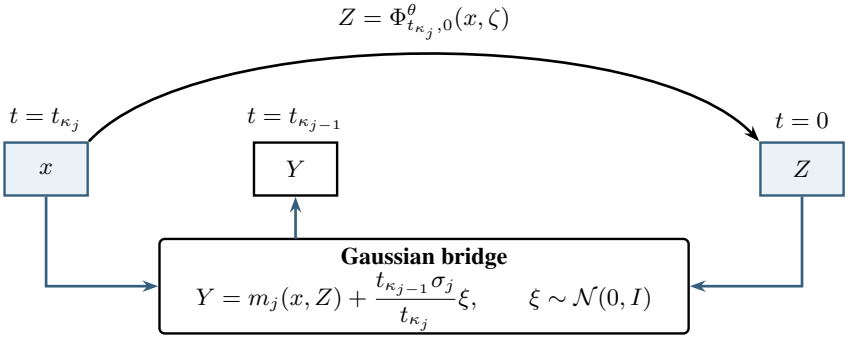

Figure B.6: Posterior sampling followed by a Gaussian bridge conditioned on both endpoints. For an exact posterior sampler, $Y \sim K_j^{\mathrm{F}}(\cdot \mid x)$ (Proposition B.1).

Proposition B.1 (Exact pairing from posterior sampling)**.** *Let $(X_0, \dots, X_J)$ be the Markov chain initialized from $X_0 \sim \pi_0$ with transitions $X_j \mid X_{0:j-1} \sim K_{j-1}^{\mathrm{B}}(\cdot \mid X_{j-1})$, $j = 1, \dots, J$, given in Equation (B.29), and let π_j denote its marginals. Using an exact sampler of $\mathrm{Law}(X_0 \mid X_j = x)$, draw*

$$Z = \Phi^{\theta}_{t_{\kappa_j},0}(x,\zeta), \qquad Y = m_j(x,Z) + \frac{t_{\kappa_{j-1}}\sigma_j}{t_{\kappa_j}}\xi, \qquad \zeta,\xi \overset{\mathrm{iid}}{\sim} \mathcal{N}(0,I). \tag{B.31}$$

At $j = 1$, $Y = Z$. The kernel $K_j^{\mathrm{F}}(\cdot \mid x) := \mathrm{Law}(Y \mid x)$ satisfies

$$\pi_j(dx)K_j^{\mathrm{F}}(dy \mid x) = \pi_{j-1}(dy)K_{j-1}^{\mathrm{B}}(dx \mid y), \qquad \rho_j = 1. \tag{B.32}$$

Proof. The reference process satisfies $X_j = a_j X_{j-1} + \sigma_j \epsilon_j$, with independent Gaussian noise ϵ_j, and $X_j \mid X_0 = z \sim \mathcal{N}((1 - t_{\kappa_j})z, t_{\kappa_j}^2 I)$. Bayes' rule therefore gives the posterior sampled by Z,

$$\mathrm{Law}(X_0 \in dz \mid X_j = x) = \frac{\varphi_{t_{\kappa_j}^2}(x - (1 - t_{\kappa_j})z)\pi_0(dz)}{\pi_j(x)}, \tag{B.33}$$

where φ_v is the density of $\mathcal{N}(0, vI)$ and $\pi_j(x)$ is the density of the noisy marginal π_j. Given $X_0 = z$, the pair (X_{j-1}, X_j) is jointly Gaussian with covariance $a_j t_{\kappa_{j-1}}^2 I$. Conditioning further on $X_j = x$ gives

$$\mathbb{E}[X_{j-1} \mid X_j = x, X_0 = z] = (1 - t_{\kappa_{j-1}})z + \frac{a_j t_{\kappa_{j-1}}^2}{t_{\kappa_j}^2}\left[x - (1 - t_{\kappa_j})z\right] = m_j(x,z),$$

$$\mathrm{Var}(X_{j-1} \mid X_j = x, X_0 = z) = \left(t_{\kappa_{j-1}}^2 - \frac{a_j^2 t_{\kappa_{j-1}}^4}{t_{\kappa_j}^2}\right) I = \frac{t_{\kappa_{j-1}}^2 \sigma_j^2}{t_{\kappa_j}^2} I.$$

Hence the Gaussian draw in Equation (B.31) has precisely the conditional law

$$\mathrm{Law}(X_{j-1} \mid X_j = x, X_0 = z) = \mathcal{N}\left(m_j(x,z), \frac{t_{\kappa_{j-1}}^2 \sigma_j^2}{t_{\kappa_j}^2} I\right). \tag{B.34}$$

Averaging over Z removes the conditioning on the clean endpoint,

$$\begin{aligned} K_j^{\mathrm{F}}(dy \mid x) &= \int \mathcal{N}\left(dy; m_j(x,z), \frac{t_{\kappa_{j-1}}^2 \sigma_j^2}{t_{\kappa_j}^2} I\right) \mathrm{Law}(X_0 \in dz \mid X_j = x) \\ &= \mathrm{Law}(X_{j-1} \in dy \mid X_j = x). \end{aligned} \tag{B.35}$$

Multiplying by $\pi_j(dx)$ now gives the required pairing,

$$\pi_j(dx)K_j^{\mathrm{F}}(dy \mid x) = \mathrm{Law}(X_j \in dx, X_{j-1} \in dy) = \pi_{j-1}(dy)K_{j-1}^{\mathrm{B}}(dx \mid y).$$

Finally, $t_0 = 0$ gives $m_1(x,Z) = Z$ and zero bridge variance, hence $Y = Z$. □

For $j \geq 2$, one possible guided proposal keeps the posterior draw Z unchanged and shifts the Gaussian bridge mean toward higher verifier values:

$$Y = m_j(x, Z) + s_j \Sigma_j \nabla \log \widehat{V}_{j-1}(m_j(x, Z)) + \Sigma_j^{1/2} \xi, \qquad \Sigma_j := \frac{t_{\kappa_{j-1}}^2 \sigma_j^2}{t_{\kappa_j}^2} I, \tag{B.36}$$

where $s_j \geq 0$ is the guidance strength, $\xi \sim \mathcal{N}(0, I)$ independently of Z, and $\widehat{V}_{j-1}$ is positive and differentiable. At $j = 1$, $\Sigma_1 = 0$ and $Y = Z$, so this bridge guidance has no effect. The ImageNet and CelebA experiments instead use the normalized mean shifts in Equations (F.2) and (F.3), with gradients evaluated at the current state.

C CS-VGB Theoretical Guarantees

We first show that the exact checkpoint update in Appendix B.2 has the desired stationary clean distribution and that reverse insertion preserves this target with finite candidate pools. We then study how the exact chain approaches stationarity and reaches the clean boundary, including the effect of process-verifier error on convergence. Finally, we quantify the bias and convergence error introduced by omitting the edge correction ρ_j.

We retain the notation of Appendices B.1 and B.2: j indexes time t_{κ_j}, π_j is its model marginal, and $\widehat{V}_0 = r$. Throughout, the state space is $\mathcal{X} \subseteq \mathbb{R}^d$, and $J \geq 2$. The process verifiers are nonnegative and finite. In particular, the corrected kernels are taken in versions satisfying Equation (B.7).

C.1 Stationarity Guarantee for CS-VGB

Let P denote the CS-VGB one-step kernel with the exact edge-correction ratios $\rho_j(x, y)$ from Equations (B.4) and (B.5), exact masses F_j, B_j^ρ, and exact draws from $Q_j^{\mathrm{F}}, Q_j^{\mathrm{B},\rho}$. It uses Table B.1 with $\chi \in [0, 1]$ and forced exits at the boundaries. The state of this kernel records the checkpoint index, sample, and direction. Since the direction is fixed at the two boundaries, P acts on

$$\mathsf{S} := \big(\{0\} \times \mathcal{X} \times \{\mathrm{B}\}\big) \cup \bigcup_{j=1}^{J-1} \big(\{j\} \times \mathcal{X} \times \{\mathrm{F}, \mathrm{B}\}\big) \cup \big(\{J\} \times \mathcal{X} \times \{\mathrm{F}\}\big). \tag{C.1}$$

To describe the stationary law of the transitions in Equations (B.10) and (B.11), we define the *incident measure* μ_j at checkpoint j by summing the marginals of its adjacent edge-conductance measures from Equations (B.8) and (B.9). At each boundary only one edge contributes, so for $1 \leq j < J$,

$$\mu_0(dy) := \widehat{C}_1(\mathcal{X}, dy), \qquad \mu_J(dx) := \widehat{C}_J(dx, \mathcal{X}), \qquad \mu_j(dx) := \pi_j(dx) \left(a_j^{\mathrm{F}}(x) + a_j^{\mathrm{B}}(x)\right). \tag{C.2}$$

Assuming the edge-conductance measures have positive finite total mass, we define the *normalizing constant* $\mathcal{Z}$ by summing these measures over all checkpoint-direction states, with two directions at each interior checkpoint and one at each boundary,

$$\mathcal{Z} := \mu_0(\mathcal{X}) + 2 \sum_{j=1}^{J-1} \mu_j(\mathcal{X}) + \mu_J(\mathcal{X}) \in (0, \infty). \tag{C.3}$$

Theorem C.1 (Stationary law of exact CS-VGB)**.** *Under terminal anchoring (Assumption B.1) and the corrected edge identity in Equation (B.7), the CS-VGB kernel P has stationary distribution Π, given by the following expressions with $1 \leq j < J$ and $d \in \{\mathrm{F}, \mathrm{B}\}$ for the interior checkpoints,*

$$\Pi(0, dx, \mathrm{B}) := \frac{\mu_0(dx)}{\mathcal{Z}}, \quad \Pi(J, dx, \mathrm{F}) := \frac{\mu_J(dx)}{\mathcal{Z}}, \quad \Pi(j, dx, d) := \frac{\mu_j(dx)}{\mathcal{Z}}. \tag{C.4}$$

At the clean boundary $j = 0$,

$$\Pi(0, \mathcal{X}, \mathrm{B}) = \frac{\mu_0(\mathcal{X})}{\mathcal{Z}}, \qquad \frac{\Pi(0, dx, \mathrm{B})}{\Pi(0, \mathcal{X}, \mathrm{B})} = \frac{r(x)\pi_0(dx)}{Z} = \pi_0^\star(dx). \tag{C.5}$$

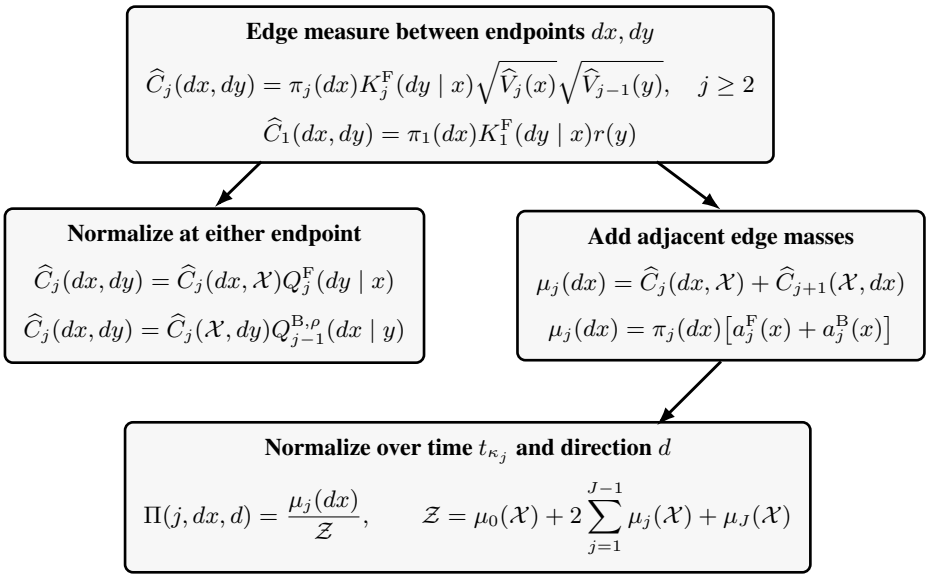

Figure C.1: Edge-conductance measures define both the CS-VGB update and its stationary law. For $1 \leq j < J$, μ_j sums the marginals of the two adjacent edge-conductance measures; at the boundaries, $\mu_0 = \widehat{C}_1(\mathcal{X}, \cdot)$ and $\mu_J = \widehat{C}_J(\cdot, \mathcal{X})$. The resulting law and clean target are established in Theorem C.1.

Figure C.1 summarizes how the same edge-conductance measures determine the sampling transitions and stationary law. The factor 2 in $\mathcal{Z}$ counts the two interior direction states. At the clean boundary, conditioning cancels $\mathcal{Z}$ and gives $r\pi_0/Z = \pi_0^\star$.

Proof. For $2 \leq j < J$, the definitions of $a_j^{\mathrm{F}}, a_j^{\mathrm{B}}$ and $p_j^{\mathrm{F}}, p_j^{\mathrm{B}}$ give

$$\mu_j(dx)p_j^{\mathrm{F}}(x) = \pi_j(dx)a_j^{\mathrm{F}}(x) = \pi_j(dx)\sqrt{\widehat{V}_j(x)}F_j(x) = \int_{\mathcal{X}} \widehat{C}_j(dx,dy) = \widehat{C}_j(dx,\mathcal{X}),$$
$$\mu_j(dy)p_j^{\mathrm{B}}(y) = \pi_j(dy)a_j^{\mathrm{B}}(y) = \pi_j(dy)\sqrt{\widehat{V}_j(y)}B_j^{\rho}(y) = \int_{\mathcal{X}} \widehat{C}_{j+1}(dx,dy) = \widehat{C}_{j+1}(\mathcal{X},dy). \tag{C.6}$$

For $j = 1$, the same identities follow by using $a_1^{\mathrm{F}} = F_1$ and the clean-edge definition of $\widehat{C}_1$ in the forward relation. For an interior edge, the two moving fluxes, before division by $\mathcal{Z}$, are therefore

$$\mu_j(dx)p_j^{\mathrm{F}}(x)Q_j^{\mathrm{F}}(dy \mid x) = \widehat{C}_j(dx,dy) = \mu_{j-1}(dy)p_{j-1}^{\mathrm{B}}(y)Q_{j-1}^{\mathrm{B},\rho}(dx \mid y), \qquad 2 \leq j < J. \tag{C.7}$$

At the clean boundary, the right-hand side is instead $\mu_0(dy)Q_0^{\mathrm{B},\rho}(dx \mid y) = \widehat{C}_1(dx,dy)$. At the noise boundary, the left-hand side is $\mu_J(dx)Q_J^{\mathrm{F}}(dy \mid x) = \widehat{C}_J(dx,dy)$.

Fix an interior checkpoint j and recall from Equation (B.14) that $\ell_j(x) := \min\{p_j^{\mathrm{F}}(x), p_j^{\mathrm{B}}(x)\}$. Moving arrivals in direction F have measure $\widehat{C}_{j+1}(\mathcal{X},dx) = \mu_j(dx)p_j^{\mathrm{B}}(x)$. The other arrivals are flips from B and stays in F. Their sum is

$$\mu_j(dx)p_j^{\mathrm{B}}(x) + \mu_j(dx)\left(p_j^{\mathrm{F}}(x) - \chi\ell_j(x)\right) + \mu_j(dx)\chi\ell_j(x) = \mu_j(dx). \tag{C.8}$$

Similarly, the incoming measure in direction B is

$$\mu_j(dx)p_j^{\mathrm{F}}(x) + \mu_j(dx)\left(p_j^{\mathrm{B}}(x) - \chi\ell_j(x)\right) + \mu_j(dx)\chi\ell_j(x) = \mu_j(dx). \tag{C.9}$$

Equations (C.8) and (C.9) show that the unnormalized incoming measure at each interior lifted state (j,dx,d) is $\mu_j(dx)$. The boundary relations give the corresponding incoming measures $\mu_0(dy)$ and $\mu_J(dx)$. Dividing each incoming measure by $\mathcal{Z}$ therefore recovers the components in Equation (C.4), which proves $\Pi P = \Pi$. Finally, terminal anchoring and normalization of $K_0^{\mathrm{B},\rho}$ give

$$\mu_0(dy) = r(y)\pi_0(dy)\int_{\mathcal{X}} K_0^{\mathrm{B},\rho}(dx \mid y) = r(y)\pi_0(dy), \qquad \mu_0(\mathcal{X}) = \int_{\mathcal{X}} r(y)\pi_0(dy) = Z.$$

Hence $\Pi(0, dy, \mathrm{B})/\Pi(0, \mathcal{X}, \mathrm{B}) = r(y)\pi_0(dy)/Z = \pi_0^\star(dy)$. □

Thus, changing the interior verifier changes the invariant law over intermediate states, but not its clean conditional. We next bound the stationary probability of being at the clean checkpoint in terms of the average verifier error under the reference marginals.

Proposition C.1 (Clean stationary mass under average verifier error)**.** *Under the assumptions of Theorem C.1, let $0 < Z := \int_{\mathcal{X}} r(x)\pi_0(dx) < \infty$ and $J \geq 2$. Suppose the nonnegative process verifiers satisfy, for some $\varepsilon \geq 0$,*

$$\frac{1}{J}\sum_{j=1}^{J} \mathbb{E}_{\pi_j}\left[|\widehat{V}_j - V_j|\right] \leq \varepsilon Z. \tag{C.10}$$

Then

$$\Pi(j=0) := \Pi(0, \mathcal{X}, \mathrm{B}) = \frac{Z}{\mathcal{Z}} \geq \frac{1}{4J - 2 + 4J\varepsilon}. \tag{C.11}$$

For exact verifiers, this recovers $\Pi(j=0) \geq 1/(4J-2)$, with equality for any positive constant reward.

Proof. Write $M_j := \widehat{C}_j(\mathcal{X}, \mathcal{X})$. Terminal anchoring gives $M_1 = Z$. For $j \geq 2$, taking expectations under $\pi_j(dx)K_j^{\mathrm{F}}(dy \mid x)$ yields

$$\begin{aligned} M_j = \mathbb{E}\sqrt{\widehat{V}_j(X_j)\widehat{V}_{j-1}(X_{j-1})} &\overset{(\circ)}{\leq} \frac{1}{2}\left(\mathbb{E}_{\pi_j}\widehat{V}_j + \mathbb{E}_{\pi_{j-1}}\widehat{V}_{j-1}\right) \\ &\overset{(\triangle)}{\leq} Z + \frac{1}{2}\left(\mathbb{E}_{\pi_j}|\widehat{V}_j - V_j| + \mathbb{E}_{\pi_{j-1}}|\widehat{V}_{j-1} - V_{j-1}|\right), \end{aligned}$$

where $(\circ)$ uses $\sqrt{ab} \leq (a+b)/2$, and $(\triangle)$ uses $\mathbb{E}_{\pi_j} V_j = Z$. By Equations (C.2) and (C.3), each verifier error enters the normalizer with coefficient at most 4, so

$$\mathcal{Z} = 3Z + 4\sum_{j=2}^{J-1} M_j + 3M_J \leq (4J-2)Z + 4\sum_{j=1}^{J} \mathbb{E}_{\pi_j}|\widehat{V}_j - V_j| \leq (4J - 2 + 4J\varepsilon)Z.$$

Since $\Pi(j=0) = Z/\mathcal{Z}$, this proves the bound. If $\widehat{V}_j = V_j$ and $r \equiv c > 0$, then $V_j \equiv c$ and $M_j = Z = c$ for every edge, giving equality in the exact-verifier bound. □

For bounded average verifier error $\varepsilon = O(1)$, the clean checkpoint retains stationary probability $\Omega(J^{-1})$, ruling out exponential dilution of its mass as the number of checkpoints grows. Therefore, this implies clean samples are not exponentially rare under the stationary law.

C.2 Exact Finite-Candidate Updates

In this subsection, we verify that reverse insertion and lazy pool evaluation preserve the clean target despite using finite pools. We keep the weights $e_{j,i}^{\mathrm{F}}, e_{j,k}^{\mathrm{B}}$ and averages $\widehat{a}_j^{\mathrm{F}}, \widehat{a}_j^{\mathrm{B}}$ from Appendix B.3. For $\mathbf{u} = (Y_{1:L_F}, Z_{1:L_B})$, the law of fresh pools is

$$\mathsf{G}_j(d\mathbf{u} \mid x) := \prod_{i=1}^{L_F} G_j^{\mathrm{F}}(dY_i \mid x) \prod_{k=1}^{L_B} G_j^{\mathrm{B}}(dZ_k \mid x),$$

omitting the unavailable pool at each boundary. Let $\mathcal{Z}_\alpha \in (0, \infty)$ be the normalizer in Equation (C.3) with $\widehat{C}_j$ replaced by $\alpha^{j-1}\widehat{C}_j$. We assume the updates below are well-defined.

For the CS-VGB kernel in Appendix B.3 with reverse insertion, lazy pool evaluation, depth weighting, and exact edge corrections ρ_j, define $\Pi_{\mathrm{mark}}^\alpha$ at $1 \leq j < J$ by

$$\begin{aligned} \Pi_{\mathrm{mark}}^\alpha(j, dx, d\mathbf{u}, \mathrm{F}, k) &:= \frac{\pi_j(dx)\mathsf{G}_j(d\mathbf{u} \mid x)}{\mathcal{Z}_\alpha} \frac{\widehat{a}_j^{\mathrm{F}} + e_{j,k}^{\mathrm{B}}}{L_B}, \\ \Pi_{\mathrm{mark}}^\alpha(j, dx, d\mathbf{u}, \mathrm{B}, i) &:= \frac{\pi_j(dx)\mathsf{G}_j(d\mathbf{u} \mid x)}{\mathcal{Z}_\alpha} \frac{\widehat{a}_j^{\mathrm{B}} + e_{j,i}^{\mathrm{F}}}{L_F}. \end{aligned}$$

At the boundaries, use Equation (C.12) without a mark.

Theorem C.2 (Stationary clean target with finite candidates)**.** *For $L_F, L_B \geq 2$, the above kernel has stationary law $\Pi^{\alpha}_{\mathrm{mark}}$, with marginal*

$$\Pi^{\alpha}_{\mathrm{pool}}(j, dx, d\mathbf{u}, d) = \frac{\pi_j(dx)\mathsf{G}_j(d\mathbf{u} \mid x)}{\mathcal{Z}_\alpha}(\widehat{a}^{\mathrm{F}}_j + \widehat{a}^{\mathrm{B}}_j), \tag{C.12}$$

after summing over the mark. Without lazy pool evaluation (i.e., using only reverse insertion and depth weighting), the CS-VGB kernel instead has stationary law $\Pi^{\alpha}_{\mathrm{pool}}$. Both satisfy

$$\Pi^{\alpha}_{\mathrm{pool}}(dx \mid j = 0) = \Pi^{\alpha}_{\mathrm{mark}}(dx \mid j = 0) = \pi^{\star}_0(dx).$$

For $\alpha = 1$, both stationary laws have marginal Π on (j, x, d), as defined in Theorem C.1.

Proof. We first establish an identity for the incoming move measures used below. Let $\mathsf{T}^{\mathrm{F}}_{j,i}$ and $\mathsf{T}^{\mathrm{B}}_{j,k}$ fix one candidate and draw all remaining entries independently:

$$\mathsf{T}^{\mathrm{F}}_{j,i}(d\mathbf{u} \mid x; y) := \delta_y(dY_i) \prod_{h \neq i} G^{\mathrm{F}}_j(dY_h \mid x) \prod_{\ell=1}^{L_B} G^{\mathrm{B}}_j(dZ_\ell \mid x),$$

$$\mathsf{T}^{\mathrm{B}}_{j,k}(d\mathbf{u} \mid x; y) := \delta_y(dZ_k) \prod_{h=1}^{L_F} G^{\mathrm{F}}_j(dY_h \mid x) \prod_{\ell \neq k} G^{\mathrm{B}}_j(dZ_\ell \mid x).$$

Unavailable pools are omitted. For fixed i, k, we compute the destination measure after moving to Y_i and inserting the previous state x into B slot k. The exact weights in Equation (B.16) give

$$\begin{aligned} &\int \pi_j(dx)\mathsf{G}_j(d\mathbf{u} \mid x) e^{\mathrm{F}}_{j,i} \delta_{Y_i}(dy) \mathsf{T}^{\mathrm{B}}_{j-1,k}(d\mathbf{u}' \mid y; x) \\ &\qquad = \alpha^{j-1} \int \widehat{C}_j(dx, dy) \mathsf{T}^{\mathrm{B}}_{j-1,k}(d\mathbf{u}' \mid y; x) = \pi_{j-1}(dy)\mathsf{G}_{j-1}(d\mathbf{u}' \mid y) e^{\mathrm{B}}_{j-1,k}. \end{aligned} \tag{C.13}$$

On the left, $\pi_j(dx)\mathsf{G}_j(d\mathbf{u} \mid x)$ is the source reference measure, weighted by $e^{\mathrm{F}}_{j,i}$; $\delta_{Y_i}(dy)$ sets $y = Y_i$, and $\mathsf{T}^{\mathrm{B}}_{j-1,k}$ inserts x into slot k and generates the remaining destination entries. Integrating out $(x, \mathbf{u})$ leaves a measure on $(y, \mathbf{u}')$. The first equality integrates the source pools $\mathbf{u}$; the second uses the exact B weight at $Z'_k = x$. The B move is symmetric.

We first verify Equation (C.12) without lazy pool evaluation by summing the measures entering an interior F state $(j, x, \mathbf{u}, \mathrm{F})$. At the source of an incoming F move,

$$\Pi^{\alpha}_{\mathrm{pool}}(j+1, dx', d\mathbf{u}', \mathrm{F})\widehat{p}^{\mathrm{F}}_{j+1} = \frac{\pi_{j+1}(dx')\mathsf{G}_{j+1}(d\mathbf{u}' \mid x')}{\mathcal{Z}_\alpha}\widehat{a}^{\mathrm{F}}_{j+1}.$$

Selecting candidate i with probability $e^{\mathrm{F}}_{j+1,i}/(L_F\widehat{a}^{\mathrm{F}}_{j+1})$ and insertion slot k uniformly, then applying Equation (C.13), gives

$$\begin{aligned} &\int \Pi^{\alpha}_{\mathrm{pool}}(j+1, dx', d\mathbf{u}', \mathrm{F})\mathbb{P}(\mathsf{move}, dx, d\mathbf{u} \mid j+1, x', \mathbf{u}', \mathrm{F}) \\ &\qquad = \frac{\pi_j(dx)\mathsf{G}_j(d\mathbf{u} \mid x)}{\mathcal{Z}_\alpha} \sum_{i=1}^{L_F} \sum_{k=1}^{L_B} \frac{e^{\mathrm{B}}_{j,k}}{L_F L_B} = \frac{\pi_j(dx)\mathsf{G}_j(d\mathbf{u} \mid x)}{\mathcal{Z}_\alpha}\widehat{a}^{\mathrm{B}}_j. \end{aligned}$$

A B-to-F flip and an F stay leave $(j, x, \mathbf{u})$ unchanged, so their incoming measures are

$$\Pi^{\alpha}_{\mathrm{pool}}(j, dx, d\mathbf{u}, \mathrm{B})(\widehat{p}^{\mathrm{F}}_j - \chi\widehat{\ell}_j) = \frac{\pi_j(dx)\mathsf{G}_j(d\mathbf{u} \mid x)}{\mathcal{Z}_\alpha}\left(\widehat{a}^{\mathrm{F}}_j - \chi \min\{\widehat{a}^{\mathrm{F}}_j, \widehat{a}^{\mathrm{B}}_j\}\right),$$

$$\Pi^{\alpha}_{\mathrm{pool}}(j, dx, d\mathbf{u}, \mathrm{F})\chi\widehat{\ell}_j = \frac{\pi_j(dx)\mathsf{G}_j(d\mathbf{u} \mid x)}{\mathcal{Z}_\alpha}\chi \min\{\widehat{a}^{\mathrm{F}}_j, \widehat{a}^{\mathrm{B}}_j\}.$$

Thus their sum is

$$\frac{\pi_j(dx)\mathsf{G}_j(d\mathbf{u} \mid x)}{\mathcal{Z}_\alpha}\left[\widehat{a}^{\mathrm{B}}_j + \widehat{a}^{\mathrm{F}}_j - \chi \min\{\widehat{a}^{\mathrm{F}}_j, \widehat{a}^{\mathrm{B}}_j\} + \chi \min\{\widehat{a}^{\mathrm{F}}_j, \widehat{a}^{\mathrm{B}}_j\}\right] = \Pi^{\alpha}_{\mathrm{pool}}(j, dx, d\mathbf{u}, \mathrm{F}).$$

Interchanging F and B proves stationarity of the B copy.

With lazy pool evaluation, fix the destination state $(j, x, \mathbf{u}, \mathrm{F}, k)$. We first compute the total outgoing F-move measure from checkpoint $j+1 < J$, summing over its possible B marks ℓ:

$$\sum_{\ell=1}^{L_B} \Pi^{\alpha}_{\mathrm{mark}}(j+1, dx', d\mathbf{u}', \mathrm{F}, \ell) \frac{\widehat{a}^{\mathrm{F}}_{j+1}}{\widehat{a}^{\mathrm{F}}_{j+1} + e^{\mathrm{B}}_{j+1,\ell}} = \frac{\pi_{j+1}(dx')\mathsf{G}_{j+1}(d\mathbf{u}' \mid x')}{\mathcal{Z}_\alpha} \widehat{a}^{\mathrm{F}}_{j+1}.$$

The same selection and insertion rule, keeping destination mark k fixed, gives by Equation (C.13)

$$\sum_{\ell=1}^{L_B} \int \Pi^{\alpha}_{\mathrm{mark}}(j+1, dx', d\mathbf{u}', \mathrm{F}, \ell) \mathbb{P}(\mathsf{move}, dx, d\mathbf{u}, K = k \mid j+1, x', \mathbf{u}', \mathrm{F}, \ell)$$

$$= \frac{\pi_j(dx)\mathsf{G}_j(d\mathbf{u} \mid x)}{\mathcal{Z}_\alpha} \sum_{i=1}^{L_F} \frac{e^{\mathrm{B}}_{j,k}}{L_F L_B} = \frac{\pi_j(dx)\mathsf{G}_j(d\mathbf{u} \mid x)}{\mathcal{Z}_\alpha} \frac{e^{\mathrm{B}}_{j,k}}{L_B}.$$

For a B-to-F flip, the uniform draw of $K = k$ and the conditional flip probability in Equation (B.21) give

$$\sum_{i=1}^{L_F} \Pi^{\alpha}_{\mathrm{mark}}(j, dx, d\mathbf{u}, \mathrm{B}, i) \mathbb{P}(\mathrm{B} \to \mathrm{F}, K = k \mid j, x, \mathbf{u}, I = i)$$

$$= \frac{\pi_j(dx)\mathsf{G}_j(d\mathbf{u} \mid x)}{\mathcal{Z}_\alpha} \sum_{i=1}^{L_F} \frac{\widehat{a}^{\mathrm{B}}_j + e^{\mathrm{F}}_{j,i}}{L_F} \frac{1}{L_B} \frac{e^{\mathrm{F}}_{j,i} - \chi \min\{e^{\mathrm{B}}_{j,k}, e^{\mathrm{F}}_{j,i}\}}{\widehat{a}^{\mathrm{B}}_j + e^{\mathrm{F}}_{j,i}}$$

$$= \frac{\pi_j(dx)\mathsf{G}_j(d\mathbf{u} \mid x)}{\mathcal{Z}_\alpha} \frac{1}{L_F L_B} \sum_{i=1}^{L_F} \left[e^{\mathrm{F}}_{j,i} - \chi \min\{e^{\mathrm{F}}_{j,i}, e^{\mathrm{B}}_{j,k}\} \right].$$

An F stay retains mark k; averaging over its temporary uniform index gives

$$\Pi^{\alpha}_{\mathrm{mark}}(j, dx, d\mathbf{u}, \mathrm{F}, k) \mathbb{P}(\mathsf{stay} \mid j, x, \mathbf{u}, \mathrm{F}, k) = \frac{\pi_j(dx)\mathsf{G}_j(d\mathbf{u} \mid x)}{\mathcal{Z}_\alpha} \frac{1}{L_F L_B} \sum_{i=1}^{L_F} \chi \min\{e^{\mathrm{F}}_{j,i}, e^{\mathrm{B}}_{j,k}\}.$$

The minimum terms cancel, so adding the incoming move measure gives

$$\frac{\pi_j(dx)\mathsf{G}_j(d\mathbf{u} \mid x)}{\mathcal{Z}_\alpha} \left[\frac{e^{\mathrm{B}}_{j,k}}{L_B} + \frac{1}{L_F L_B} \sum_{i=1}^{L_F} e^{\mathrm{F}}_{j,i} \right] = \Pi^{\alpha}_{\mathrm{mark}}(j, dx, d\mathbf{u}, \mathrm{F}, k). \tag{C.14}$$

The B case is symmetric. At a boundary, the forced move and the same insertion calculation give Equation (C.12) without a mark. Summing over the mark gives

$$\sum_{k=1}^{L_B} \Pi^{\alpha}_{\mathrm{mark}}(j, dx, d\mathbf{u}, \mathrm{F}, k) = \frac{\pi_j(dx)\mathsf{G}_j(d\mathbf{u} \mid x)}{\mathcal{Z}_\alpha} \left(\widehat{a}^{\mathrm{F}}_j + \frac{1}{L_B} \sum_{k=1}^{L_B} e^{\mathrm{B}}_{j,k} \right) = \Pi^{\alpha}_{\mathrm{pool}}(j, dx, d\mathbf{u}, \mathrm{F}).$$

The B case is symmetric. Finally, the exact weights satisfy

$$\pi_j(dx) \int \mathsf{G}_j(d\mathbf{u} \mid x) \widehat{a}^{\mathrm{F}}_j = \alpha^{j-1} \widehat{C}_j(dx, \mathcal{X}), \qquad \pi_j(dx) \int \mathsf{G}_j(d\mathbf{u} \mid x) \widehat{a}^{\mathrm{B}}_j = \alpha^{j} \widehat{C}_{j+1}(\mathcal{X}, dx),$$

for each available direction. Thus integrating out the pools recovers the depth-weighted law with normalizer $\mathcal{Z}_\alpha$. Since the clean edge has weight $\alpha^0 = 1$,

$$\Pi^{\alpha}_{\mathrm{pool}}(0, dx, \mathrm{B}) = \frac{\widehat{C}_1(\mathcal{X}, dx)}{\mathcal{Z}_\alpha} = \frac{r(x)\pi_0(dx)}{\mathcal{Z}_\alpha}, \qquad \Pi^{\alpha}_{\mathrm{pool}}(dx \mid j = 0) = \frac{r(x)\pi_0(dx)}{Z} = \pi^{\star}_0(dx).$$

The marked law has the same clean conditional. For $\alpha = 1$, integrating out the pools gives Π from Theorem C.1. □

C.3 Convergence Guarantee for CS-VGB

Having identified the stationary law, we now ask how quickly the chain approaches it. We relate movement between states to the decay of distributional error and derive rate bounds for the exact kernel P, then quantify how process-verifier error affects these bounds. The first hitting time of the clean boundary is treated separately in Appendix C.4.

From a Poincaré inequality to convergence. For probability measures $\nu \ll \nu'$ on S, define their *total variation distance* by

$$\|\nu - \nu'\|_{\mathrm{TV}} := \sup_{E \subseteq \mathsf{S}} |\nu(E) - \nu'(E)| = \frac{1}{2} \int_{\mathsf{S}} \left| \frac{d\nu}{d\nu'} - 1 \right| d\nu'. \tag{C.15}$$

For $f \in L^2(\Pi)$, let $\Pi(f)$, $\|f\|_{2,\Pi}$, and $\mathrm{Var}_\Pi(f)$ denote its mean, $L^2(\Pi)$ norm, and variance under Π, respectively. The *Dirichlet form* $\mathcal{E}_P(f, f)$ measures its mean squared change over one stationary transition:

$$\begin{aligned} \Pi(f) &:= \int f \, d\Pi, \qquad\qquad \|f\|_{2,\Pi}^2 := \Pi(f^2), \\ \mathrm{Var}_\Pi(f) &:= \|f - \Pi(f)\|_{2,\Pi}^2, \quad \mathcal{E}_P(f, f) := \frac{1}{2} \iint [f(z') - f(z)]^2 \, \Pi(dz) P(z, dz'). \end{aligned} \tag{C.16}$$

We say that P satisfies a *Poincaré inequality* with constant $\gamma > 0$ if

$$\mathcal{E}_P(f, f) \geq \gamma \, \mathrm{Var}_\Pi(f), \qquad f \in L^2(\Pi). \tag{C.17}$$

This condition requires the one-step variation to control the stationary variance uniformly over all $f \in L^2(\Pi)$. For the theoretical convergence analysis, we introduce a *lazy kernel* to remove possible periodicity and turn this inequality into a contraction bound:

$$P_{\mathrm{L},\eta} := \frac{\eta I + P}{1 + \eta}, \qquad \eta > 0, \qquad P_{\mathrm{L}} := P_{\mathrm{L},1}. \tag{C.18}$$

Here I is the identity kernel, so the chain stays at its current state with probability $\eta/(1+\eta)$ and otherwise applies P. This additional self-loop preserves Π and is separate from the momentum update's stay action. The mixing bounds below apply to the $1/2$-lazy kernel P_{L}; this analytical modification does not require adding lazification to the practical algorithm.

Corollary C.1 (Exponential convergence of the outer MCMC)**.** *Suppose Equation (C.17) holds, and let ν be an initial distribution satisfying $\nu \ll \Pi$ and $D_\nu := \|d\nu/d\Pi - 1\|_{2,\Pi} < \infty$. For every integer $n \geq 0$,*

$$\left\| \frac{d(\nu P_{\mathrm{L}}^n)}{d\Pi} - 1 \right\|_{2,\Pi}^2 \leq (1 - \gamma/2)^n D_\nu^2, \qquad \|\nu P_{\mathrm{L}}^n - \Pi\|_{\mathrm{TV}} \leq \frac{D_\nu}{2} (1 - \gamma/2)^{n/2} \leq \frac{D_\nu}{2} e^{-\gamma n/4}. \tag{C.19}$$

Whenever $D_\nu (1 - \gamma/2)^{n/2}/2 < \Pi(0, \mathcal{X}, \mathrm{B})$, the clean conditional is well-defined and satisfies

$$\left\| \frac{(\nu P_{\mathrm{L}}^n)(0, dx, \mathrm{B})}{(\nu P_{\mathrm{L}}^n)(0, \mathcal{X}, \mathrm{B})} - \pi_0^\star(dx) \right\|_{\mathrm{TV}} \leq \frac{D_\nu (1 - \gamma/2)^{n/2}}{2\Pi(0, \mathcal{X}, \mathrm{B})} \leq \frac{D_\nu e^{-\gamma n/4}}{2\Pi(0, \mathcal{X}, \mathrm{B})}. \tag{C.20}$$

Thus clean-conditional error at most $\delta \in (0, 1)$ is guaranteed by

$$n \geq \left\lceil \frac{4}{\gamma} \log \left(\max \left\{ 1, \frac{D_\nu}{2\Pi(0, \mathcal{X}, \mathrm{B})\delta} \right\} \right) \right\rceil. \tag{C.21}$$

Proof. Write $Pf(z) := \int f(z') P(z, dz')$, and let P^* be the adjoint of P with respect to the $L^2(\Pi)$ inner product $\langle f, g \rangle_\Pi := \Pi(fg)$, so that $\langle f, Pg \rangle_\Pi = \langle P^* f, g \rangle_\Pi$. Stationarity and Jensen's inequality give

$$\mathcal{E}_P(f, f) = \|f\|_{2,\Pi}^2 - \langle f, Pf \rangle_\Pi = \langle f, (I - P) f \rangle_\Pi, \qquad \|Pf\|_{2,\Pi}^2 \leq \Pi(P(f^2)) = \|f\|_{2,\Pi}^2. \tag{C.22}$$

The adjoint P^* corresponds to the time-reversed kernel under Π, characterized by $\Pi(dz) P(z, dz') = \Pi(dz') P^*(z', dz)$. It therefore also preserves Π, and the same Jensen argument gives $\|P^* g\|_{2,\Pi} \leq \|g\|_{2,\Pi}$. Using $P_{\mathrm{L}}^* = (I + P^*)/2$ and $\langle g, P^* g \rangle_\Pi = \langle g, Pg \rangle_\Pi$, the lazy-chain argument (Fill, 1991) gives, for $\Pi(g) = 0$,

$$\|P_{\mathrm{L}}^* g\|_{2,\Pi}^2 = \frac{1}{4} \big(\|g\|_{2,\Pi}^2 + 2\langle g, P^* g \rangle_\Pi + \|P^* g\|_{2,\Pi}^2 \big) \leq \|g\|_{2,\Pi}^2 - \frac{1}{2} \mathcal{E}_P(g, g) \overset{(\circ)}{\leq} (1 - \gamma/2) \|g\|_{2,\Pi}^2. \tag{C.23}$$

Here $(\circ)$ follows from the Poincaré inequality in Equation (C.17), since $\Pi(g) = 0$ gives $\mathrm{Var}_\Pi(g) = \|g\|_{2,\Pi}^2$. By duality, $\int f\, d(\nu P_{\mathrm{L}}) = \langle P_{\mathrm{L}}^*(d\nu/d\Pi), f\rangle_\Pi$. Since $P_{\mathrm{L}}^*\mathbf{1} = \mathbf{1}$, iteration yields

$$\frac{d(\nu P_{\mathrm{L}}^n)}{d\Pi} - 1 = (P_{\mathrm{L}}^*)^n \left(\frac{d\nu}{d\Pi} - 1\right). \tag{C.24}$$

Applying Equation (C.23) and Cauchy–Schwarz gives

$$\|\nu P_{\mathrm{L}}^n - \Pi\|_{\mathrm{TV}} \le \frac{1}{2}\left\|\frac{d(\nu P_{\mathrm{L}}^n)}{d\Pi} - 1\right\|_{2,\Pi} \le \frac{D_\nu}{2}(1-\gamma/2)^{n/2} \le \frac{D_\nu}{2} e^{-\gamma n/4}. \tag{C.25}$$

For the clean bound, put $a := \Pi(0, \mathcal{X}, \mathrm{B})$ and $b := (\nu P_{\mathrm{L}}^n)(0, \mathcal{X}, \mathrm{B})$. The stated condition gives $|a - b| < a$, hence $b > 0$. For any two laws with these positive boundary masses,

$$\left\|\frac{(\nu P_{\mathrm{L}}^n)(0, dx, \mathrm{B})}{b} - \frac{\Pi(0, dx, \mathrm{B})}{a}\right\|_{\mathrm{TV}} \le \frac{\|\nu P_{\mathrm{L}}^n - \Pi\|_{\mathrm{TV}}}{\max\{a, b\}}. \tag{C.26}$$

Indeed, when $b \ge a$ and the conditional difference on $E \subseteq \mathcal{X}$ is nonnegative,

$$\begin{aligned} b\left(\frac{(\nu P_{\mathrm{L}}^n)(0, E, \mathrm{B})}{b} - \frac{\Pi(0, E, \mathrm{B})}{a}\right) &= (\nu P_{\mathrm{L}}^n)(0, E, \mathrm{B}) - \Pi(0, E, \mathrm{B}) - \frac{b-a}{a}\Pi(0, E, \mathrm{B}) \\ &\le \|\nu P_{\mathrm{L}}^n - \Pi\|_{\mathrm{TV}}. \end{aligned}$$

Use E^c for a negative difference and exchange the two laws when $a > b$. Combining Equation (C.26) with Equation (C.5) proves Equation (C.20); Equation (C.21) follows by solving the exponential bound. □

A positive CS-VGB Poincaré constant. Corollary C.1 expresses the convergence rate in terms of the Poincaré constant. We now show when this constant is positive for CS-VGB and how to bound it using the model transitions and process-verifier weights.

We compare the edge-conductance measures $\widehat{C}_j$ used by CS-VGB with the reward-tilted reference edges $C_j(dx, dy) := \pi_j(dx) K_j^{\mathrm{F}}(dy \mid x) V_{j-1}(y)$, $j = 1, \dots, J$. Both use the same denoising kernels K_j^{F}. Recall that the exact continuation values $V_j := V_{t_{\kappa_j}}$ satisfy

$$V_0 = r, \qquad V_j = K_j^{\mathrm{F}} V_{j-1}, \qquad j = 1, \dots, J. \tag{C.27}$$

Weighting a reference trajectory by its terminal reward and normalizing gives the path law $\pi_{0:J}^\star$ and its marginal $\overline{C}_j$ for the pair (X_j, X_{j-1}) at neighboring checkpoints:

$$\pi_{0:J}^\star(dx_{0:J}) := \frac{r(x_0)}{Z}\pi_J(dx_J)\prod_{j=1}^{J} K_j^{\mathrm{F}}(dx_{j-1} \mid x_j), \tag{C.28}$$

$$\overline{C}_j(dx, dy) := \frac{C_j(dx, dy)}{Z} = \frac{\pi_j(dx) K_j^{\mathrm{F}}(dy \mid x) V_{j-1}(y)}{Z}. \tag{C.29}$$

To see why $\overline{C}_j$ is the joint law of (X_j, X_{j-1}), integrate the path law in Equation (C.28) over all other states. The part from noise to checkpoint index j and the reward-weighted part from checkpoint index $j-1$ to clean give, respectively,

$$\int_{\mathcal{X}^{J-j}} \pi_J(dx_J)\prod_{\ell=j+1}^{J} K_\ell^{\mathrm{F}}(dx_{\ell-1} \mid x_\ell) = \pi_j(dx_j), \int_{\mathcal{X}^{j-1}} r(x_0)\prod_{\ell=1}^{j-1} K_\ell^{\mathrm{F}}(dx_{\ell-1} \mid x_\ell) = V_{j-1}(x_{j-1}). \tag{C.30}$$

At $j = J$ the first expression is simply $\pi_J(dx_J)$; at $j = 1$ the second is $r(x_0) = V_0(x_0)$. The remaining transition is $K_j^{\mathrm{F}}(dy \mid x)$, so the joint marginal is $\pi_j(dx) K_j^{\mathrm{F}}(dy \mid x) V_{j-1}(y)/Z = \overline{C}_j(dx, dy)$.

To use C_j as a reference for the convergence analysis, we verify that its two marginals are the reward-tilted measures $Z\pi_j^\star$ and $Z\pi_{j-1}^\star$. Integrating out either endpoint and using $V_j = K_j^{\mathrm{F}} V_{j-1}$ and $\pi_j K_j^{\mathrm{F}} = \pi_{j-1}$ gives

$$\begin{aligned} C_j(dx, \mathcal{X}) &= \pi_j(dx)\int_{\mathcal{X}} K_j^{\mathrm{F}}(dy \mid x) V_{j-1}(y) = \pi_j(dx) V_j(x) = Z\pi_j^\star(dx), \\ C_j(\mathcal{X}, dy) &= V_{j-1}(y)\int_{\mathcal{X}} \pi_j(dx) K_j^{\mathrm{F}}(dy \mid x) = V_{j-1}(y)\pi_{j-1}(dy) = Z\pi_{j-1}^\star(dy). \end{aligned} \tag{C.31}$$

To measure how the verifier changes these edge weights, assume $\widehat{C}_j \ll C_j$ and define the *edge-weight ratio*

$$R_1 := 1, \qquad R_j(x,y) := \frac{d\widehat{C}_j}{dC_j}(x,y), \qquad j = 2, \dots, J. \tag{C.32}$$

Terminal anchoring gives $R_1 = 1$. For $j \geq 2$, recall that

$$\widehat{C}_j(dx,dy) = \pi_j(dx)K_j^{\mathrm{F}}(dy \mid x)\sqrt{\widehat{V}_j(x)\widehat{V}_{j-1}(y)}, \qquad C_j(dx,dy) = \pi_j(dx)K_j^{\mathrm{F}}(dy \mid x)V_{j-1}(y).$$

Both measures have the same factor $\pi_j(dx)K_j^{\mathrm{F}}(dy \mid x)$, so their density ratio is the ratio of the remaining weights. Writing the verifier ratio as $\xi_j(x) := \widehat{V}_j(x)/V_j(x)$ wherever $V_j(x) > 0$ gives

$$R_j(x,y) = \frac{\sqrt{\widehat{V}_j(x)}\sqrt{\widehat{V}_{j-1}(y)}}{V_{j-1}(y)} = \sqrt{\frac{V_j(x)}{V_{j-1}(y)}}\sqrt{\xi_j(x)\xi_{j-1}(y)}, \qquad j \geq 2. \tag{C.33}$$

From edge weights to the stationary measure. To compare the stationary law with the reward-tilted reference marginals, we translate the edge ratios R_j into weights on each checkpoint. We first verify that these ratios give a finite normalizing constant. For this purpose, assume

$$0 \leq R_j < \infty \quad \overline{C}_j\text{-almost surely}, \qquad \max_{1\leq j\leq J} \mathbb{E}_{\overline{C}_j}[R_j] < \infty. \tag{C.34}$$

Recall from Equation (C.2) that $\mu_j(dx) := \pi_j(dx)\left(a_j^{\mathrm{F}}(x) + a_j^{\mathrm{B}}(x)\right)$ for $1 \leq j < J$: it adds the incident masses in the two directions. By Equation (C.6),

$$\pi_j(dx)a_j^{\mathrm{F}}(x) = \widehat{C}_j(dx, \mathcal{X}), \qquad \pi_j(dy)a_j^{\mathrm{B}}(y) = \widehat{C}_{j+1}(\mathcal{X}, dy).$$

The first term collects edges toward checkpoint index $j-1$, and the second collects edges toward $j+1$. Adding them, with only one incident edge at each boundary, gives

$$\mu_0 = \widehat{C}_1(\mathcal{X}, \cdot), \qquad \mu_J = \widehat{C}_J(\cdot, \mathcal{X}), \qquad \mu_j = \widehat{C}_j(\cdot, \mathcal{X}) + \widehat{C}_{j+1}(\mathcal{X}, \cdot), \quad 1 \leq j < J. \tag{C.35}$$

Recall that $Z = \int_{\mathcal{X}} r(x)\pi_0(dx)$ normalizes the reward-tilted target, while $\mathcal{Z}$ normalizes the full stationary measure Π. Since $\widehat{C}_j = ZR_j\overline{C}_j$, each edge has total mass

$$\widehat{C}_j(\mathcal{X}, \mathcal{X}) = Z\int R_j\, d\overline{C}_j = Z\mathbb{E}_{\overline{C}_j}[R_j] \leq Z \max_{1\leq \ell\leq J} \mathbb{E}_{\overline{C}_\ell}[R_\ell].$$

Each boundary has one incident edge; each interior checkpoint has two, counted once for each direction. There are therefore $2 + 4(J-1) = 4J-2$ edge contributions, giving

$$0 < \mathcal{Z} = \mu_0(\mathcal{X}) + 2\sum_{j=1}^{J-1}\mu_j(\mathcal{X}) + \mu_J(\mathcal{X}) \leq (4J-2)Z \max_{1\leq j\leq J} \mathbb{E}_{\overline{C}_j}[R_j] < \infty. \tag{C.36}$$

This proves that Π has a finite normalizing constant. We now express the stationary measure Π as a weighted version of the checkpoint index targets $\pi_j^\star$. This allows us to analyze Π using the reward-tilted reference path. To do so, define the density w_j of μ_j relative to $Z\pi_j^\star$:

$$w_j := \frac{d\mu_j}{Z\, d\pi_j^\star}. \tag{C.37}$$

To compute w_j, recall that $\overline{C}_j$ is the joint law of (X_j, X_{j-1}) and has marginals $\pi_j^\star$ and $\pi_{j-1}^\star$. Thus, writing $R_j = R_j(X_j, X_{j-1})$ and taking conditional expectations under $\pi_{0:J}^\star$,

$$\widehat{C}_j(dx, \mathcal{X}) = Z\int_{\mathcal{X}} R_j(x,y)\overline{C}_j(dx,dy) = Z\pi_j^\star(dx)\mathbb{E}_{\overline{C}_j}[R_j \mid X_j = x],$$

$$\widehat{C}_{j+1}(\mathcal{X}, dy) = Z\int_{\mathcal{X}} R_{j+1}(x,y)\overline{C}_{j+1}(dx,dy) = Z\pi_j^\star(dy)\mathbb{E}_{\overline{C}_{j+1}}[R_{j+1} \mid X_j = y].$$

At an interior checkpoint $1 \leq j < J$, μ_j is the sum of these two marginals. Dividing by $Z\pi_j^\star$ adds their conditional expectations. At the boundaries there is only one term, and $R_1 = 1$, giving

$$w_0 = 1, \quad w_J(X_J) = \mathbb{E}_{\overline{C}_J}[R_J \mid X_J], \quad w_j(X_j) = \mathbb{E}_{\overline{C}_j}[R_j \mid X_j] + \mathbb{E}_{\overline{C}_{j+1}}[R_{j+1} \mid X_j]. \tag{C.38}$$

Substituting $\mu_j = Zw_j\pi_j^\star$ into $\Pi(j, dx, d) = \mu_j(dx)/\mathcal{Z}$ gives

$$\Pi(j, dx, d) = \frac{Zw_j(x)}{\mathcal{Z}}\pi_j^\star(dx).$$

This expresses Π as a weighted version of the reward-tilted marginals $\pi_j^\star$, with weights w_j. We use this representation below to relate the variance under Π to the reward-tilted reference path in the Poincaré analysis.

A condition ensuring that states are connected. Finite total mass ensures that Π is a probability measure, but does not ensure convergence to it. If the chain cannot leave a region A with $0 < \Pi(A) < 1$, a chain started in A cannot approach Π. We therefore require the denoising and renoising moves to cross every such region. For a measurable set $A \subseteq \mathsf{S}$, let $A_j^d := \{x : (j, x, d) \in A\}$, using $A_0^{\mathrm{F}} = A_0^{\mathrm{B}}$ and $A_J^{\mathrm{B}} = A_J^{\mathrm{F}}$ at the boundaries. The condition is

$$0 < \Pi(A) < 1 \quad \Longrightarrow \quad \sum_{j=1}^{J} \sum_{d \in \{\mathrm{F},\mathrm{B}\}} \iint [\mathbf{1}_{A_j^d}(x) - \mathbf{1}_{A_{j-1}^d}(y)]^2 \, \widehat{C}_j(dx, dy) > 0. \tag{C.39}$$

The integral counts the edge mass crossing between A and its complement. The following lemma gives a sufficient condition in terms of the reference path.

Lemma C.1 (Connectedness from the reference endpoints)**.** *Suppose*

$$0 < R_j < \infty \qquad \overline{C}_j\text{-almost surely}, \qquad j = 1, \dots, J, \tag{C.40}$$

and, for every $u \in L^2(\pi_0^\star)$ and $v \in L^2(\pi_J^\star)$,

$$u(X_0) = v(X_J) \quad \pi_{0:J}^\star\text{-almost surely} \quad \Longrightarrow \quad \mathrm{Var}_{\pi_0^\star}(u) = 0. \tag{C.41}$$

Then Equation (C.39) holds for every measurable $A \subseteq \mathsf{S}$ with $0 < \Pi(A) < 1$.

Proof. Suppose the sum in Equation (C.39) vanishes. Each nonnegative term then vanishes, so $\widehat{C}_j = ZR_j\overline{C}_j$ gives

$$0 = \int \left(\mathbf{1}_{A_j^d}(x) - \mathbf{1}_{A_{j-1}^d}(y)\right)^2 \widehat{C}_j(dx, dy) = Z \int \left(\mathbf{1}_{A_j^d}(x) - \mathbf{1}_{A_{j-1}^d}(y)\right)^2 R_j(x, y)\overline{C}_j(dx, dy).$$

Since $R_j > 0$ $\overline{C}_j$-almost surely by Equation (C.40), the two indicators agree $\overline{C}_j$-almost surely. Moreover, $\overline{C}_j$ is the law of (X_j, X_{j-1}) under $\pi_{0:J}^\star$, hence

$$\pi_{0:J}^\star\left(\mathbf{1}_{A_j^d}(X_j) \neq \mathbf{1}_{A_{j-1}^d}(X_{j-1})\right) = \int \mathbf{1}_{\{\mathbf{1}_{A_j^d}(x) \neq \mathbf{1}_{A_{j-1}^d}(y)\}} \overline{C}_j(dx, dy) = 0.$$

Therefore,

$$\mathbf{1}_{A_j^d}(X_j) = \mathbf{1}_{A_{j-1}^d}(X_{j-1}) \quad \pi_{0:J}^\star\text{-almost surely}$$

for every j and d. Following these equalities along the path shows that $\mathbf{1}_{A_0}(X_0) = \mathbf{1}_{A_J}(X_J)$ almost surely. By Equation (C.41), the variance of $\mathbf{1}_{A_0}(X_0)$ is zero. This indicator is therefore either zero almost surely or one almost surely. Equality along each edge gives that same value at every checkpoint index and in both directions. Since $\Pi(j, dx, d) = Zw_j(x)\pi_j^\star(dx)/\mathcal{Z}$, a set of zero $\pi_j^\star$-measure also has zero stationary mass at that checkpoint index. Thus, in the first case $\Pi(A) = 0$, and in the second case $\Pi(\mathsf{S} \setminus A) = 0$, so $\Pi(A) = 1$. Consequently, Equation (C.39) holds for every measurable $A \subseteq \mathsf{S}$ with $0 < \Pi(A) < 1$. □

The next lemma gives a sufficient condition for Equation (C.41) in terms of the maximal correlation between the clean and noise endpoints.

Lemma C.2 (An endpoint correlation condition)**.** *Suppose that, for all $u \in L^2(\pi_0^\star)$ and $v \in L^2(\pi_J^\star)$,*

$$\left|\mathrm{Cov}_{\pi_{0:J}^\star}\big(u(X_0), v(X_J)\big)\right| \leq \rho_{\mathrm{end}} \sqrt{\mathrm{Var}_{\pi_0^\star}(u)\,\mathrm{Var}_{\pi_J^\star}(v)}, \qquad 0 \leq \rho_{\mathrm{end}} < 1. \tag{C.42}$$

Then Equation (C.41) holds. In particular, Equation (C.42) holds with $\rho_{\mathrm{end}} = 0$ if the clean and noise endpoints are independent under the reference path $\pi_J(dx_J) \prod_{j=1}^{J} K_j^{\mathrm{F}}(dx_{j-1} \mid x_j)$.

Proof. Let $u(X_0) = v(X_J)$ $\pi_{0:J}^\star$-almost surely. The two random variables then have the same mean and variance, and

$$\mathrm{Cov}_{\pi_{0:J}^\star}\big(u(X_0), v(X_J)\big) = \mathrm{Var}_{\pi_0^\star}(u) = \mathrm{Var}_{\pi_J^\star}(v).$$

Substituting these identities into Equation (C.42) gives

$$\mathrm{Var}_{\pi_0^\star}(u) \le \rho_{\mathrm{end}} \mathrm{Var}_{\pi_0^\star}(u), \qquad (1-\rho_{\mathrm{end}})\mathrm{Var}_{\pi_0^\star}(u) \le 0.$$

Since $\rho_{\mathrm{end}} < 1$ and variance is nonnegative, $\mathrm{Var}_{\pi_0^\star}(u) = 0$, proving Equation (C.41).

For the last claim, independence under this reference path gives the endpoint law $\pi_0(dx_0)\pi_J(dx_J)$. Weighting by $r(x_0)/Z$ therefore gives

$$\begin{aligned}\pi^\star_{0:J}(X_0 \in dx_0, X_J \in dx_J) &= \frac{r(x_0)}{Z}\int_{\mathcal{X}^{J-1}} \pi_J(dx_J)\prod_{j=1}^{J} K_j^{\mathrm{F}}(dx_{j-1} \mid x_j) \\ &= \frac{r(x_0)}{Z}\pi_0(dx_0)\pi_J(dx_J) = \pi_0^\star(dx_0)\pi_J(dx_J).\end{aligned}$$

Thus the endpoints remain independent under $\pi^\star_{0:J}$, so their covariance is zero for every such u, v. □

Condition Equation (C.42) bounds the correlation between all square-integrable functions of the clean state X_0 and the top-checkpoint state X_J strictly below one under $\pi^\star_{0:J}$. In a noising interpretation, this rules out preserving any nonconstant square-integrable function of the clean state without error at the top checkpoint.

Finally, we require the squared edge density to have a finite integral after accounting for the mass at both endpoints. For $j = 1, \dots, J$, assume

$$\widehat{C}_j \ll \mu_j \otimes \mu_{j-1}, \qquad h_j := \frac{d\widehat{C}_j}{d(\mu_j \otimes \mu_{j-1})}, \qquad \mathcal{I}_j := \iint h_j(x,y)^2\, \mu_j(dx)\mu_{j-1}(dy) < \infty. \tag{C.43}$$

In the proof below, write $f_j^d(x) := f(j, x, d)$, with the boundary convention

$$f_0^{\mathrm{F}} = f_0^{\mathrm{B}} = f_0 := f(0, \cdot, \mathrm{B}), \qquad f_J^{\mathrm{F}} = f_J^{\mathrm{B}} = f_J := f(J, \cdot, \mathrm{F}).$$

The next lemma is a consequence of the Fredholm alternative (Brezis, 2011); we give a direct proof for completeness.

Lemma C.3 (From compactness to a positive Poincaré constant)**.** *Let M be a bounded, self-adjoint, compact operator on $L^2(\Pi)$, and let c satisfy $0 < c_0 \le c \le c_1 < \infty$ Π-almost surely. Define $Lf := cf - Mf$. Suppose*

$$\langle f, Lf\rangle_\Pi \ge 0 \quad \textit{for every } f \in L^2(\Pi), \qquad \ker L = \mathrm{span}\{\mathbf{1}\}.$$

Then $\langle f, Lf\rangle_\Pi \ge \gamma\, \mathrm{Var}_\Pi(f)$ for some $\gamma > 0$ and every $f \in L^2(\Pi)$.

Proof. Since $L\mathbf{1} = 0$, it suffices to consider mean-zero unit vectors f, i.e., $\Pi(f) = 0$ and $\|f\|_{2,\Pi} = 1$. If the conclusion failed, there would be f_n with

$$\Pi(f_n) = 0, \qquad \|f_n\|_{2,\Pi} = 1, \qquad \langle f_n, Lf_n\rangle_\Pi \longrightarrow 0.$$

Positivity of $\langle h + tf, L(h+tf)\rangle_\Pi$ for every real t gives Cauchy–Schwarz for the form $\langle h, Lf\rangle_\Pi$, i.e., $|\langle h, Lf\rangle_\Pi|^2 \le \langle h, Lh\rangle_\Pi \langle f, Lf\rangle_\Pi$. Taking $h = Lf$ yields

$$\|Lf\|_{2,\Pi}^4 \le \langle Lf, L(Lf)\rangle_\Pi \langle f, Lf\rangle_\Pi \le \|L\|\|Lf\|_{2,\Pi}^2 \langle f, Lf\rangle_\Pi, \qquad \|Lf_n\|_{2,\Pi} \longrightarrow 0.$$

Now, we recall that compactness gives a subsequence with $Mf_{n_h} \to v$. Since $c^{-1} \le c_0^{-1}$,

$$f_{n_h} = c^{-1}(Mf_{n_h} + Lf_{n_h}) \longrightarrow f := c^{-1}v \quad \text{in } L^2(\Pi).$$

By continuity, $\Pi(f) = 0$, $\|f\|_{2,\Pi} = 1$, and $Lf = 0$. The first two identities imply that f is nonconstant, since a mean-zero constant has norm zero. Together with $Lf = 0$, this contradicts $\ker L = \mathrm{span}\{\mathbf{1}\}$. □

We now apply this lemma to the exact CS-VGB kernel.

Theorem C.3 (A positive Poincaré constant for CS-VGB)**.** *Let P be the exact CS-VGB kernel with $0 < \mathcal{Z} < \infty$. Suppose Equations (C.39) and (C.43) hold. Then there is a $\gamma_{\mathrm{CS}} \in (0, 1]$, independent of χ, such that*

$$\mathcal{E}_P(f, f) \geq \gamma_{\mathrm{CS}} \operatorname{Var}_\Pi(f), \qquad f \in L^2(\Pi), \qquad \chi \in [0, 1]. \tag{C.44}$$

Consequently, the full-state and clean-conditional convergence bounds in Corollary C.1 hold with $\gamma = \gamma_{\mathrm{CS}}$.

Proof. We apply Lemma C.3 to the movement part of P. Let P_{move} retain only checkpoint moves with their original probabilities, and let P^*_{move} be its adjoint in $L^2(\Pi)$. Define

$$M := \frac{P_{\mathrm{move}} + P^*_{\mathrm{move}}}{2}, \qquad c := M\mathbf{1}, \qquad \mathcal{L}_{\mathrm{move}} f := cf - Mf. \tag{C.45}$$

For $1 \leq j < J$, the outgoing movement mass is

$$(P_{\mathrm{move}}\mathbf{1})(j, x, d) = P_{\mathrm{move}}((j, x, d), \mathsf{S}) = p^d_j(x).$$

By the adjoint identity $\langle P^*_{\mathrm{move}}\mathbf{1}, f\rangle_\Pi = \langle \mathbf{1}, P_{\mathrm{move}} f\rangle_\Pi$, the incoming mass at an interior state $z = (j, x, d)$, $1 \leq j < J$, can be written as

$$(P^*_{\mathrm{move}}\mathbf{1})(z)\Pi(dz) = (\Pi P_{\mathrm{move}})(dz) = \int \Pi(dz') P_{\mathrm{move}}(z', dz) = \big(1 - p^d_j(x)\big)\Pi(dz).$$

Indeed, arrivals in directions F and B come from $j + 1$ and $j - 1$, respectively. With $p^{\mathrm{F}}_J = p^{\mathrm{B}}_0 = 1$ for forced boundary moves, Equations (C.6) and (C.7) gives

$$\begin{aligned}
(P^*_{\mathrm{move}}\mathbf{1})(j, x, \mathrm{F})\Pi(j, dx, \mathrm{F}) &= \int_{\mathcal{X}} \underbrace{\Pi(j+1, dy, \mathrm{F})}_{\mu_{j+1}(dy)/\mathcal{Z}} \underbrace{p^{\mathrm{F}}_{j+1}(y)}_{\text{move probability}} \underbrace{Q^{\mathrm{F}}_{j+1}(dx \mid y)}_{\text{destination law}} \\
&= \frac{1}{\mathcal{Z}} \int_{\mathcal{X}} \mu_{j+1}(dy) p^{\mathrm{F}}_{j+1}(y) Q^{\mathrm{F}}_{j+1}(dx \mid y) \\
&= \frac{\widehat{C}_{j+1}(\mathcal{X}, dx)}{\mathcal{Z}} = p^{\mathrm{B}}_j(x) \frac{\mu_j(dx)}{\mathcal{Z}}, \\
(P^*_{\mathrm{move}}\mathbf{1})(j, x, \mathrm{B})\Pi(j, dx, \mathrm{B}) &= \int_{\mathcal{X}} \underbrace{\Pi(j-1, dy, \mathrm{B})}_{\mu_{j-1}(dy)/\mathcal{Z}} \underbrace{p^{\mathrm{B}}_{j-1}(y)}_{\text{move probability}} \underbrace{Q^{\mathrm{B},\rho}_{j-1}(dx \mid y)}_{\text{destination law}} \\
&= \frac{1}{\mathcal{Z}} \int_{\mathcal{X}} \mu_{j-1}(dy) p^{\mathrm{B}}_{j-1}(y) Q^{\mathrm{B},\rho}_{j-1}(dx \mid y) \\
&= \frac{\widehat{C}_j(dx, \mathcal{X})}{\mathcal{Z}} = p^{\mathrm{F}}_j(x) \frac{\mu_j(dx)}{\mathcal{Z}}.
\end{aligned}$$

Since $\Pi(j, dx, d) = \mu_j(dx)/\mathcal{Z}$, these identities give, almost surely on each stationary component,

$$(P_{\mathrm{move}}\mathbf{1})(j, x, d) = p^d_j(x), \qquad (P^*_{\mathrm{move}}\mathbf{1})(j, x, d) = 1 - p^d_j(x).$$

At the boundaries both quantities equal one, with the direction reset inward after arrival. Therefore

$$c(j, x, d) = \begin{cases} 1, & j \in \{0, J\}, \\ 1/2, & 1 \leq j < J. \end{cases} \tag{C.46}$$

The kernel M is nonnegative and symmetric with respect to Π, with $M\mathbf{1} = c$. Thus

$$|Mf(z)|^2 = \left| \int f(z') M(z, dz') \right|^2 \overset{(\circ)}{\leq} \left(\int M(z, dz') \right) \left(\int f(z')^2 M(z, dz') \right) = c(z) M(f^2)(z),$$

where $(\circ)$ is Cauchy–Schwarz. Integrating with respect to $\Pi(dz)$ gives

$$\|Mf\|^2_{2,\Pi} \leq \Pi\big(cM(f^2)\big) \overset{(\triangle)}{\leq} \Pi\big(M(f^2)\big) \overset{(\square)}{=} \Pi(cf^2) \overset{(\triangle)}{\leq} \Pi(f^2) = \|f\|^2_{2,\Pi},$$

where $(\triangle)$ uses $c \le 1$, and $(\square)$ follows from the symmetry of M and $M\mathbf{1} = c$. Thus M is bounded and self-adjoint. To compute its movement energy, use $M\mathbf{1} = c$:

$$\iint f(z)^2\,\Pi(dz)M(z,dz') = \int f(z)^2 c(z)\Pi(dz) = \Pi(cf^2).$$

Symmetry and the definition of Mf give the other two terms:

$$\iint f(z')^2\,\Pi(dz)M(z,dz') = \iint f(z)^2\,\Pi(dz)M(z,dz') = \Pi(cf^2),$$
$$\iint f(z)f(z')\,\Pi(dz)M(z,dz') = \int f(z)Mf(z)\Pi(dz) = \langle f, Mf\rangle_\Pi.$$

On each edge, Equation (C.7) gives

$$\Pi(j,dx,\mathrm{F})p_j^{\mathrm{F}}(x)Q_j^{\mathrm{F}}(dy\mid x) = \frac{\mu_j(dx)p_j^{\mathrm{F}}(x)Q_j^{\mathrm{F}}(dy\mid x)}{\mathcal{Z}} = \frac{\widehat{C}_j(dx,dy)}{\mathcal{Z}},$$
$$\Pi(j-1,dy,\mathrm{B})p_{j-1}^{\mathrm{B}}(y)Q_{j-1}^{\mathrm{B},\rho}(dx\mid y) = \frac{\mu_{j-1}(dy)p_{j-1}^{\mathrm{B}}(y)Q_{j-1}^{\mathrm{B},\rho}(dx\mid y)}{\mathcal{Z}} = \frac{\widehat{C}_j(dx,dy)}{\mathcal{Z}}.$$

Since $[f(z)-f(z')]^2$ is symmetric, its integrals against the stationary edge measures of P_{move} and P_{move}^* agree. Hence

$$\begin{aligned}
\langle f, \mathcal{L}_{\mathrm{move}} f\rangle_\Pi &= \Pi(cf^2) - \langle f, Mf\rangle_\Pi = \frac{1}{2}\iint \big(f(z)^2 + f(z')^2 - 2f(z)f(z')\big)\,\Pi(dz)M(z,dz') \\
&= \frac{1}{2}\iint \big(f(z)-f(z')\big)^2\,\Pi(dz)M(z,dz') \\
&= \frac{1}{2}\iint \big(f(z)-f(z')\big)^2\,\Pi(dz)P_{\mathrm{move}}(z,dz') \\
&= \frac{1}{2\mathcal{Z}}\sum_{j=1}^{J}\sum_{d\in\{\mathrm{F},\mathrm{B}\}}\iint \big(f_j^d(x) - f_{j-1}^d(y)\big)^2\,\widehat{C}_j(dx,dy) \\
&\overset{(\circ)}{\le} \mathcal{E}_P(f,f),
\end{aligned} \tag{C.47}$$

where $(\circ)$ drops the nonnegative flip contributions from Equation (C.16); stays contribute zero.

To establish compactness, write $M(z,dz') = m(z,z')\Pi(dz')$. For $z = (j,x,d)$ and $z' = (j-1,y,d)$, averaging P_{move} and its adjoint gives

$$\Pi(dz)M(z,dz') = \frac{\widehat{C}_j(dx,dy)}{2\mathcal{Z}} = m(z,z')\Pi(dz)\Pi(dz') = m\big((j,x,d),(j-1,y,d)\big)\frac{\mu_j(dx)\mu_{j-1}(dy)}{\mathcal{Z}^2}. \tag{C.48}$$

Using the definition of h_j in Equation (C.43), we obtain

$$m\big((j,x,d),(j-1,y,d)\big)\frac{\mu_j(dx)\mu_{j-1}(dy)}{\mathcal{Z}^2} = \frac{\widehat{C}_j(dx,dy)}{2\mathcal{Z}} = \frac{h_j(x,y)}{2\mathcal{Z}}\mu_j(dx)\mu_{j-1}(dy).$$

Comparing densities with respect to $\mu_j \otimes \mu_{j-1}$, and using the symmetry of M, gives almost everywhere

$$m\big((j,x,d),(j-1,y,d)\big) = \frac{\mathcal{Z}}{2}h_j(x,y), \qquad m(z,z') = m(z',z),$$

with the boundary directions identified as before. Using $\mathcal{I}_j$ defined in Equation (C.43), each direction and orientation contributes

$$\iint \left(\frac{\mathcal{Z}}{2}h_j(x,y)\right)^2 \frac{\mu_j(dx)}{\mathcal{Z}}\frac{\mu_{j-1}(dy)}{\mathcal{Z}} = \frac{1}{4}\mathcal{I}_j.$$

Summing over the two direction labels and both orientations of each edge in M gives

$$\iint m(z,z')^2\,\Pi(dz)\Pi(dz') = \sum_{j=1}^{J} 4\left(\frac{\mathcal{I}_j}{4}\right) = \sum_{j=1}^{J}\mathcal{I}_j < \infty. \tag{C.49}$$

Finite sums of products are dense in $L^2(\Pi \otimes \Pi)$ (Tao, 2010), so we can choose finite-rank approximations

$$m_N(z,z') = \sum_{\ell=1}^{r_N} a_{N,\ell}(z) b_{N,\ell}(z'), \qquad \|m - m_N\|_{L^2(\Pi\otimes\Pi)} \longrightarrow 0,$$

$$M_N f(z) := \int m_N(z,z') f(z') \Pi(dz') = \sum_{\ell=1}^{r_N} a_{N,\ell}(z) \int b_{N,\ell}(z') f(z') \Pi(dz').$$

The range of M_N lies in $\operatorname{span}\{a_{N,1}, \ldots, a_{N,r_N}\}$, so M_N has finite rank. For every $f \in L^2(\Pi)$, Cauchy–Schwarz gives

$$\|(M - M_N)f\|_{2,\Pi}^2 = \int \left| \int (m(z,z') - m_N(z,z')) f(z') \Pi(dz') \right|^2 \Pi(dz) \leq \|m - m_N\|_{L^2(\Pi\otimes\Pi)}^2 \|f\|_{2,\Pi}^2.$$

Hence $\|M - M_N\| \to 0$. For any sequence with $\sup_n \|f_n\|_{2,\Pi} \leq B$, finite dimensionality gives a convergent subsequence of $(M_N f_n)$ for each N (Kreyszig, 1978). A diagonal subsequence makes $M_N f_{n_h}$ converge for every fixed N, and

$$\|M f_{n_h} - M f_{n_\ell}\|_{2,\Pi} \leq 2B \|M - M_N\| + \|M_N f_{n_h} - M_N f_{n_\ell}\|_{2,\Pi}.$$

Choosing N large and then h, ℓ large makes both terms arbitrarily small. Thus $(M f_{n_h})$ is Cauchy in $L^2(\Pi)$, which is complete, and hence converges, proving compactness.

It remains to show that $\mathcal{L}_{\text{move}} f = 0$ holds only for constant functions. Suppose $\mathcal{L}_{\text{move}} f = 0$. Then $\langle f, \mathcal{L}_{\text{move}} f \rangle_\Pi = 0$, so Equation (C.47) implies

$$f_j^d(x) = f_{j-1}^d(y) \qquad \widehat{C}_j\text{-almost everywhere} \tag{C.50}$$

for every j, d. If f is not constant Π-almost surely, there is an $a \in \mathbb{R}$ such that $A := \{z : f(z) > a\}$ satisfies $0 < \Pi(A) < 1$. But Equation (C.50) then gives

$$\sum_{j=1}^{J} \sum_{d \in \{\mathrm{F},\mathrm{B}\}} \iint \left(\mathbf{1}_{A_j^d}(x) - \mathbf{1}_{A_{j-1}^d}(y) \right)^2 \widehat{C}_j(dx, dy) = 0,$$

contradicting Equation (C.39). Since $\mathcal{L}_{\text{move}} \mathbf{1} = c - M\mathbf{1} = 0$, we conclude

$$\ker \mathcal{L}_{\text{move}} = \operatorname{span}\{\mathbf{1}\}.$$

Now Lemma C.3 applies with $1/2 \leq c \leq 1$. Hence there exists $\gamma_{\text{CS}} > 0$, reduced to at most one if necessary, such that

$$\mathcal{E}_P(f, f) \geq \langle f, \mathcal{L}_{\text{move}} f \rangle_\Pi \geq \gamma_{\text{CS}} \operatorname{Var}_\Pi(f).$$

Neither P_{move} nor Π depends on χ, so the same γ_{CS} works for every $\chi \in [0, 1]$. Applying Corollary C.1 completes the proof.

□

Sufficient conditions for a positive Poincaré constant. Theorem C.3 guarantees convergence under the connectedness and integrability conditions in Equations (C.39) and (C.43). To see when these conditions hold for the kernels used by CS-VGB, we consider Gaussian renoising paired with its exact denoising reversal. The following result verifies both conditions using an exponential-moment bound on the clean distribution and terminal reward, together with bounds on the verifier values.

Corollary C.2 (Ideal paired Gaussian transitions). *Take $\mathcal{X} = \mathbb{R}^d$ and a fixed grid with $J \geq 2$. Suppose the reference kernels and their marginals satisfy*

$$\pi_j(dx)K_j^{\mathrm{F}}(dy \mid x) = \pi_{j-1}(dy)K_{j-1}^{\mathrm{B}}(dx \mid y), \qquad K_{j-1}^{\mathrm{B}}(dx \mid y) = \mathcal{N}(dx; a_j y, \sigma_j^2 I), \quad \text{(C.51)}$$

where $a_j \in \mathbb{R}$ and $0 < \sigma_j < \infty$ for $j = 1, \ldots, J$. Write $A_j := \prod_{\ell=1}^{j} a_\ell$ for the cumulative signal scale. Assume $r \geq 0$ and $Z = \int r \, d\pi_0 \in (0, \infty)$, and suppose that, for some $c > 0$,

$$\mathbb{E}_{\pi_0}\left((1 + r(X_0))e^{c\|X_0\|^2}\right) < \infty, \qquad c \geq \max_{1 \leq j \leq J} \frac{2A_j^2}{\sigma_j^2}. \quad \text{(C.52)}$$

This permits unbounded clean support and controls the tails of both π_0 and $\pi_0^\star$. For some $v_{\min} \in (0, 1]$, also assume

$$v_{\min} \leq \widehat{V}_j(x) \leq 1 \quad \pi_j\text{-almost surely}, \qquad j = 1, \ldots, J, \qquad \widehat{V}_0 = r. \quad \text{(C.53)}$$

Then the conditions in Equations (C.39) and (C.43) hold. Therefore, there exists $\gamma_{\mathrm{CS}} \in (0, 1]$, independent of the directional persistence parameter χ, such that

$$\mathcal{E}_P(f, f) \geq \gamma_{\mathrm{CS}} \operatorname{Var}_\Pi(f), \qquad \|\nu P_{\mathrm{L}}^n - \Pi\|_{\mathrm{TV}} \leq \frac{D_\nu}{2} e^{-\gamma_{\mathrm{CS}} n/4}, \quad \text{(C.54)}$$

for $f \in L^2(\Pi)$, $\chi \in [0, 1]$, and every initial law $\nu \ll \Pi$ with $D_\nu := \|d\nu/d\Pi - 1\|_{2,\Pi} < \infty$. In particular, the initialization $\nu_J = \delta_J \otimes \pi_J \otimes \delta_{\mathrm{F}}$ satisfies

$$D_{\nu_J}^2 \leq \frac{\mathcal{Z}}{v_{\min}} - 1 < \infty. \quad \text{(C.55)}$$

Proof. We verify Equations (C.39) and (C.43). By exact pairing, the reference path can be generated from $X_0 \sim \pi_0$ by

$$X_j = a_j X_{j-1} + \sigma_j \varepsilon_j.$$

Here the $\varepsilon_j \sim \mathcal{N}(0, I)$ are independent of one another and of X_0. Writing the cumulative signal scale A_j and noise scale τ_j as

$$A_0 = 1, \qquad \tau_0^2 = 0, \qquad A_j = a_j A_{j-1}, \qquad \tau_j^2 = a_j^2 \tau_{j-1}^2 + \sigma_j^2. \quad \text{(C.56)}$$

gives $X_j \mid X_0 = u \sim \mathcal{N}(A_j u, \tau_j^2 I)$. Let φ_v denote the density of $\mathcal{N}(0, vI)$. For $j \geq 1$, define the marginal density of π_j by

$$\pi_j(x) := \int \varphi_{\tau_j^2}(x - A_j u)\pi_0(du). \quad \text{(C.57)}$$

To bound $\mathcal{I}_j$, we use $K_{j-1}^{\mathrm{B}}(x \mid y) := \varphi_{\sigma_j^2}(x - a_j y)$ for the density of $K_{j-1}^{\mathrm{B}}(dx \mid y)$. We show that the following integral H_j is finite by the Gaussian calculation that follows:

$$H_j := \int \pi_{j-1}(dy) \int \frac{K_{j-1}^{\mathrm{B}}(x \mid y)^2}{\pi_j(x)} \, dx \leq \frac{1}{b_{j,v}} \int \pi_{j-1}(dy) \int \frac{\varphi_{\sigma_j^2}(x - a_j y)^2}{\varphi_v(x)} \, dx. \quad \text{(C.58)}$$

The inequality uses the lower bound $\pi_j(x) \geq b_{j,v}\varphi_v(x)$ established in Equation (C.59). Choose a finite B such that $\pi_0(\{u : \|u\| \leq B\}) > 0$; such a ball exists because π_0 is a probability measure. For $0 < v < \tau_j^2$ and $\|u\| \leq B$, expanding the square gives

$$\begin{aligned}\|x - A_j u\|^2 &= \|x\|^2 - 2A_j\langle x, u\rangle + A_j^2\|u\|^2 \overset{(\circ)}{\leq} \|x\|^2 + \frac{\tau_j^2 - v}{v}\|x\|^2 + \frac{vA_j^2}{\tau_j^2 - v}\|u\|^2 + A_j^2\|u\|^2 \\ &= \frac{\tau_j^2}{v}\|x\|^2 + \frac{\tau_j^2 A_j^2}{\tau_j^2 - v}\|u\|^2 \leq \frac{\tau_j^2}{v}\|x\|^2 + \frac{\tau_j^2 A_j^2 B^2}{\tau_j^2 - v}.\end{aligned}$$

Here $(\circ)$ uses $2|A_j\langle x,u\rangle| \le \epsilon\|x\|^2 + \epsilon^{-1}A_j^2\|u\|^2$ with $\epsilon = (\tau_j^2 - v)/v > 0$. Restricting the mixture integral in Equation (C.57) to this ball gives

$$\pi_j(x) \ge \int_{\|u\|\le B} \varphi_{\tau_j^2}(x - A_j u)\pi_0(du) \ge b_{j,v}\varphi_v(x),$$

$$b_{j,v} := \pi_0(\{u : \|u\| \le B\}) \left(\frac{v}{\tau_j^2}\right)^{d/2} \exp\left(-\frac{A_j^2 B^2}{2(\tau_j^2 - v)}\right) > 0. \tag{C.59}$$

We now evaluate the inner Gaussian integral in Equation (C.58). For $v > \sigma_j^2/2$,

$$\int \frac{\varphi_{\sigma_j^2}(x - a_j y)^2}{\varphi_v(x)}\, dx = \frac{(2\pi v)^{d/2}}{(2\pi\sigma_j^2)^d} \int \exp\left(-\frac{\|x - a_j y\|^2}{\sigma_j^2} + \frac{\|x\|^2}{2v}\right) dx$$

$$= \left(\frac{v}{\sigma_j\sqrt{2v - \sigma_j^2}}\right)^d \exp\left(\frac{a_j^2\|y\|^2}{2v - \sigma_j^2}\right).$$

Choose

$$v_j := \tau_j^2 - \frac{\sigma_j^2}{4}, \qquad \eta_j := \frac{a_j^2}{2v_j - \sigma_j^2}. \tag{C.60}$$

The variance recursion implies

$$\frac{\sigma_j^2}{2} < v_j < \tau_j^2, \qquad 2v_j - \sigma_j^2 = 2a_j^2\tau_{j-1}^2 + \frac{\sigma_j^2}{2}, \qquad 2\eta_j\tau_{j-1}^2 = \frac{2a_j^2\tau_{j-1}^2}{2a_j^2\tau_{j-1}^2 + \sigma_j^2/2} < 1. \tag{C.61}$$

Consequently,

$$H_j \le \frac{1}{b_{j,v_j}} \left(\frac{v_j}{\sigma_j\sqrt{2v_j - \sigma_j^2}}\right)^d \mathbb{E}_{\pi_{j-1}}[e^{\eta_j\|X_{j-1}\|^2}]. \tag{C.62}$$

Recall that the moment-generating function of the noncentral chi-squared distribution gives, for $Y \sim \mathcal{N}(m, s^2 I)$ and $2\eta s^2 < 1$,

$$\mathbb{E}[e^{\eta\|Y\|^2}] = (1 - 2\eta s^2)^{-d/2} \exp\left(\frac{\eta\|m\|^2}{1 - 2\eta s^2}\right).$$

Applying this formula to $X_{j-1} \mid X_0 = u \sim \mathcal{N}(A_{j-1}u, \tau_{j-1}^2 I)$ and averaging over X_0 gives

$$\begin{aligned}\mathbb{E}_{\pi_{j-1}}[e^{\eta_j\|X_{j-1}\|^2}] &= \int \mathbb{E}\left[e^{\eta_j\|X_{j-1}\|^2} \mid X_0 = u\right] \pi_0(du) \\ &= (1 - 2\eta_j\tau_{j-1}^2)^{-d/2}\mathbb{E}_{\pi_0}\left[\exp\left(\frac{\eta_j A_{j-1}^2\|X_0\|^2}{1 - 2\eta_j\tau_{j-1}^2}\right)\right] \\ &\overset{(\circ)}{\le} (1 - 2\eta_j\tau_{j-1}^2)^{-d/2}\mathbb{E}_{\pi_0}[e^{c\|X_0\|^2}] \overset{(\triangle)}{<} \infty.\end{aligned} \tag{C.63}$$

Here $(\circ)$ follows from Equation (C.60) and the variance recursion, which give

$$\frac{\eta_j A_{j-1}^2}{1 - 2\eta_j\tau_{j-1}^2} = \frac{a_j^2 A_{j-1}^2}{2v_j - \sigma_j^2 - 2a_j^2\tau_{j-1}^2} = \frac{2A_j^2}{\sigma_j^2} \le c,$$

and $(\triangle)$ uses $2\eta_j\tau_{j-1}^2 < 1$ and Equation (C.52). This also holds for $j = 1$, with $\tau_0 = 0$, and proves $H_j < \infty$. For the reward-tilted clean law,

$$\mathbb{E}_{\pi_0^\star}[e^{c\|X_0\|^2}] = \frac{1}{Z}\mathbb{E}_{\pi_0}[r(X_0)e^{c\|X_0\|^2}] < \infty.$$

Repeating the calculation with $\pi_0^\star$ as the initial law gives $\pi_j^\star(x) \ge b_{j,v}^\star\varphi_v(x)$ for some $b_{j,v}^\star > 0$. Together with the exponential moment above, this proves the same integrability bound with π_{j-1} and π_j replaced by $\pi_{j-1}^\star$ and $\pi_j^\star$, respectively.

Recalling the renoising joint edge law J_j^{B} from Equation (B.4), exact pairing and the verifier bounds give, for $j \geq 2$,

$$v_{\min} J_j^{\mathrm{B}} \leq \widehat{C}_j \leq J_j^{\mathrm{B}}, \qquad J_j^{\mathrm{B}}(dx, dy) = \pi_{j-1}(dy) K_{j-1}^{\mathrm{B}}(x \mid y) dx. \tag{C.64}$$

Since $\widehat{C}_1(\mathcal{X}, \mathcal{X}) = Z$ and $\widehat{C}_j(\mathcal{X}, \mathcal{X}) \leq 1$ for $j \geq 2$,

$$0 < \mathcal{Z} = 3Z + 4 \sum_{j=2}^{J-1} \widehat{C}_j(\mathcal{X}, \mathcal{X}) + 3\widehat{C}_J(\mathcal{X}, \mathcal{X}) \leq (4J - 2) \max\{1, Z\} < \infty.$$

Each incident measure contains the corresponding marginal of $\widehat{C}_j$, so Equation (C.64) gives

$$\mu_j(dx) \geq \widehat{C}_j(dx, \mathcal{X}) \geq v_{\min} J_j^{\mathrm{B}}(dx, \mathcal{X}) = v_{\min} \pi_j(dx),$$
$$\mu_{j-1}(dy) \geq \widehat{C}_j(\mathcal{X}, dy) \geq v_{\min} J_j^{\mathrm{B}}(\mathcal{X}, dy) = v_{\min} \pi_{j-1}(dy).$$

For $j \geq 2$, we use $\widehat{C}_j(x, y)$ and $\mu_j(x)$ for the corresponding Lebesgue densities. By definition,

$$h_j(x, y) = \frac{\widehat{C}_j(x, y)}{\mu_j(x) \mu_{j-1}(y)}, \;\; \mathcal{I}_j = \iint h_j(x, y)^2 \mu_j(x) \mu_{j-1}(y) \, dx \, dy = \iint \frac{\widehat{C}_j(x, y)^2}{\mu_j(x) \mu_{j-1}(y)} \, dx \, dy.$$

For $j \geq 2$, using $\widehat{C}_j(x, y) = \pi_{j-1}(y) K_{j-1}^{\mathrm{B}}(x \mid y) \sqrt{\widehat{V}_j(x)} \sqrt{\widehat{V}_{j-1}(y)} \leq \pi_{j-1}(y) K_{j-1}^{\mathrm{B}}(x \mid y)$ and the two marginal bounds gives

$$\mathcal{I}_j \leq \iint \frac{\pi_{j-1}(y)^2 K_{j-1}^{\mathrm{B}}(x \mid y)^2}{v_{\min}^2 \pi_j(x) \pi_{j-1}(y)} \, dx \, dy = v_{\min}^{-2} \int \pi_{j-1}(dy) \int \frac{K_{j-1}^{\mathrm{B}}(x \mid y)^2}{\pi_j(x)} \, dx = v_{\min}^{-2} H_j < \infty. \tag{C.65}$$

For the clean edge, let $\pi_1^\star(x) := \int K_0^{\mathrm{B}}(x \mid y) \pi_0^\star(dy)$. Then

$$\widehat{C}_1(dx, dy) = Z \pi_0^\star(dy) K_0^{\mathrm{B}}(x \mid y) dx, \qquad \mu_0 = Z \pi_0^\star, \qquad \mu_1(dx) \geq Z \pi_1^\star(x) dx.$$

Substituting these identities in the same normalized integral yields

$$\mathcal{I}_1 \leq \int \pi_0^\star(dy) \int \frac{K_0^{\mathrm{B}}(x \mid y)^2}{\pi_1^\star(x)} \, dx < \infty. \tag{C.66}$$

The last inequality follows from the same Gaussian integral bound with clean law $\pi_0^\star$. Thus Equation (C.43) holds for every edge.

To verify Equation (C.39), note that reward tilting replaces π_0 by $\pi_0^\star$ without changing the noising kernels. Hence

$$\pi_{0:J}^\star(X_0 \in du, X_J \in dx) = \pi_0^\star(du) \varphi_{\tau_J^2}(x - A_J u) dx. \tag{C.67}$$

The Gaussian density $\varphi_{\tau_J^2}(x - A_J u)$ is positive for every u, x, and hence so is the marginal density $\pi_J^\star(x)$. Thus this joint law and $\pi_0^\star \otimes \pi_J^\star$ have the same null sets. If $u_0(X_0) = u_J(X_J)$ under the joint law, the equality also holds under the product law. For square-integrable u_0, u_J, integration under that product gives

$$\begin{aligned} 0 = \mathbb{E}_{\pi_0^\star \otimes \pi_J^\star} \left((u_0(X_0) - u_J(X_J))^2 \right) &= \mathbb{E}_{\pi_0^\star}[u_0^2] + \mathbb{E}_{\pi_J^\star}[u_J^2] - 2\mathbb{E}_{\pi_0^\star}[u_0] \mathbb{E}_{\pi_J^\star}[u_J] \\ &= \mathrm{Var}_{\pi_0^\star}(u_0) + \mathrm{Var}_{\pi_J^\star}(u_J) + \left(\mathbb{E}_{\pi_0^\star}[u_0] - \mathbb{E}_{\pi_J^\star}[u_J] \right)^2. \end{aligned}$$

Since all three terms are nonnegative, this implies $\mathrm{Var}_{\pi_0^\star}(u_0) = 0$, so Equation (C.41) holds.

Moreover, $V_j(x) = Z \pi_j^\star(x) / \pi_j(x) \in (0, \infty)$ for $j \geq 1$, where $\pi_j^\star(x)$ is the Gaussian-mixture density obtained from $\pi_0^\star$. The verifier bounds therefore imply Equation (C.40), and Lemma C.1 gives Equation (C.39).

Both required conditions now hold, so Theorem C.3 gives Equation (C.54). Finally, recalling the incident measures and stationary law from Equations (C.2) and (C.4), at the noise boundary we have

$$\Pi(J, dx, \mathrm{F}) = \frac{\mu_J(dx)}{\mathcal{Z}} = \frac{\widehat{C}_J(dx, \mathcal{X})}{\mathcal{Z}} = \frac{\sqrt{\widehat{V}_J(x)} F_J(x)}{\mathcal{Z}} \pi_J(dx).$$

Since $J \geq 2$, both factors in the numerator are at least $\sqrt{v_{\min}}$. The density $d\nu_J/d\Pi$ equals $\mathcal{Z}/(\sqrt{\widehat{V}_J F_J})$ at (J, x, F) and zero elsewhere. Since $\int (d\nu_J/d\Pi)\, d\Pi = 1$,

$$\begin{aligned} D^2_{\nu_J} &= \int \left(\frac{d\nu_J}{d\Pi} - 1\right)^2 d\Pi = \int \left(\frac{d\nu_J}{d\Pi}\right)^2 d\Pi - 1 \\ &= \int \left(\frac{\mathcal{Z}}{\sqrt{\widehat{V}_J(x) F_J(x)}}\right)^2 \frac{\sqrt{\widehat{V}_J(x)} F_J(x)}{\mathcal{Z}} \pi_J(dx) - 1 \\ &= \mathcal{Z} \int \frac{\pi_J(dx)}{\sqrt{\widehat{V}_J(x) F_J(x)}} - 1 \leq \frac{\mathcal{Z}}{v_{\min}} - 1 < \infty. \end{aligned} \tag{C.68}$$

□

We remark that Corollary C.2 establishes connectedness under Gaussian pairing. Here we show that bounded clean support also ensures the endpoint correlation bound in Equation (C.42), which is used in the explicit rate bound below.

Lemma C.4 (Endpoint correlation with bounded clean support)**.** *Assume the Gaussian pairing in Equation (C.51), $\|X_0\| \leq R < \infty$ almost surely, and $0 < Z = \mathbb{E}_{\pi_0} r(X_0) < \infty$ for $r \geq 0$. Then the clean and noise endpoints have maximal correlation $\rho_{\mathrm{end}} < 1$ under $\pi^\star_{0:J}$ (Equation (C.42)).*

Proof. Exact pairing gives $X_J = A_J X_0 + \tau_J \xi$, where $\xi \sim \mathcal{N}(0, I)$ is independent of X_0, $A_J = \prod_{j=1}^J a_j$, and $\tau_J^2 = \sum_{k=1}^J \sigma_k^2 \left(\prod_{\ell=k+1}^J a_\ell^2\right) > 0$. Weighting by $r(X_0)/Z$ changes only the clean law, so this conditional Gaussian law also holds under $\pi^\star_{0:J}$. For every $\|x\| \leq R$, its density satisfies

$$\varphi_{\tau_J^2}(z - A_J x) \geq h(z) := (2\pi\tau_J^2)^{-d/2} \exp\left[-\frac{(\|z\| + |A_J| R)^2}{2\tau_J^2}\right].$$

Since $(\|z\| + |A_J|R)^2 \geq \|z\|^2$, we have $0 < h(z) \leq \varphi_{\tau_J^2}(z)$, so $\epsilon := \int h(z)\, dz \in (0, 1]$. By centering and rescaling, we may assume without loss of generality that $U = u(X_0)$ and $V = v(X_J)$ have mean zero and variance one under $\pi^\star_{0:J}$; zero-variance cases satisfy the correlation bound trivially. The density bound gives

$$2 - 2\mathbb{E}[UV] = \mathbb{E}[(U - V)^2] \geq \int h(z) \int (u(x) - v(z))^2\, \pi_0^\star(dx)\, dz \overset{(\circ)}{=} \int h(z)(1 + v(z)^2)\, dz \geq \epsilon,$$

where $(\circ)$ uses $\int u\, d\pi_0^\star = 0$ and $\int u^2\, d\pi_0^\star = 1$. Applying the same bound to $-V$ yields $|\mathbb{E}[UV]| \leq 1 - \epsilon/2$. Taking the supremum over these normalized functions gives $\rho_{\mathrm{end}} \leq 1 - \epsilon/2 < 1$. □

An explicit bound under bounded reward and verifier weights. Theorem C.3 establishes a positive Poincaré constant under its connectivity and integrability assumptions. We now derive an explicit constant by additionally bounding the reward and implemented verifier weights above and away from zero. These bounds concern the ranges of the weights, not the accuracy of $\widehat{V}_j/V_j$.

Corollary C.3 (An explicit Poincaré bound under bounded weights)**.** *Let P be the exact CS-VGB one-step kernel in Theorem C.3, with $J \geq 2$, terminal anchoring $\widehat{V}_0 = r$, and the continuation recursion in Equation (C.27). Suppose that, for constants $r_{\min}, v_{\min} \in (0,1]$ and $j = 1, \dots, J$,*

$$r_{\min} \leq r(x) \leq 1 \quad \pi_0\text{-almost surely}, \qquad v_{\min} \leq \widehat{V}_j(x) \leq 1 \quad \pi_j\text{-almost surely}. \tag{C.69}$$

Assume also that the reward-tilted reference path satisfies

$$\left|\mathrm{Cov}_{\pi^\star_{0:J}}\big(u(X_0), v(X_J)\big)\right| \leq \rho_{\mathrm{end}} \sqrt{\mathrm{Var}_{\pi^\star_0}(u)\,\mathrm{Var}_{\pi^\star_J}(v)}, \qquad 0 \leq \rho_{\mathrm{end}} < 1,$$

for all $u \in L^2(\pi^\star_0)$ and $v \in L^2(\pi^\star_J)$, as in Equation (C.42). Under these assumptions, we can choose the same Poincaré constant for every directional persistence parameter $\chi \in [0,1]$ as follows:

$$\mathcal{E}_P(f,f) \geq \gamma_{\mathrm{CS}}\,\mathrm{Var}_\Pi(f), \qquad \gamma_{\mathrm{CS}} := \frac{(1-\rho_{\mathrm{end}}) r_{\min}^2 v_{\min}^2}{32 J^2} > 0, \qquad f \in L^2(\Pi). \tag{C.70}$$

Consequently, for every $\nu \ll \Pi$ with $D_\nu := \|d\nu/d\Pi - 1\|_{2,\Pi} < \infty$,

$$\|\nu P_{\mathrm{L}}^n - \Pi\|_{\mathrm{TV}} \leq \frac{D_\nu}{2} \exp\left[-\frac{(1-\rho_{\mathrm{end}}) r_{\min}^2 v_{\min}^2}{128 J^2} n\right]. \tag{C.71}$$

The corresponding clean-conditional bound follows from Equation (C.20).

Proof. Recall that $\pi^\star_j(dx) = V_j(x)\pi_j(dx)/Z$ is the reward-tilted marginal at checkpoint index j. We use these marginals to define a comparison law $\Pi_\star$ on the full state (j,x,d). Write $f_j^d(x) := f(j,x,d)$, identifying the boundary copies as

$$f_0^{\mathrm{F}} = f_0^{\mathrm{B}} = f_0 := f(0,\cdot,\mathrm{B}), \qquad f_J^{\mathrm{F}} = f_J^{\mathrm{B}} = f_J := f(J,\cdot,\mathrm{F}).$$

Recall that $\Pi(j,dx,d) = \mu_j(dx)/\mathcal{Z}$, where μ_j sums the incident marginals of $\widehat{C}_j$. Similarly, we define the reference incident measures using C_j, whose marginals are $Z\pi^\star_j$ and $Z\pi^\star_{j-1}$:

$$\begin{aligned} \mu_{\star,0}(dx) &:= C_1(\mathcal{X}, dx) = Z\pi^\star_0(dx), \qquad \mu_{\star,J}(dx) := C_J(dx, \mathcal{X}) = Z\pi^\star_J(dx), \\ \mu_{\star,j}(dx) &:= C_j(dx,\mathcal{X}) + C_{j+1}(\mathcal{X}, dx) = 2Z\pi^\star_j(dx), \qquad 1 \leq j < J. \end{aligned}$$

Thus $\mu_{\star,j} = Z\omega_j\pi^\star_j$, where $\omega_0 = \omega_J = 1$ and $\omega_j = 2$ in the interior. Counting both interior directions gives

$$\mathcal{Z}_\star := \mu_{\star,0}(\mathcal{X}) + 2\sum_{j=1}^{J-1} \mu_{\star,j}(\mathcal{X}) + \mu_{\star,J}(\mathcal{X}) = Z + 4(J-1)Z + Z = (4J-2)Z.$$

Normalizing these incident measures, with one direction at each boundary, yields

$$\Pi_\star(j,dx,d) := \frac{\mu_{\star,j}(dx)}{\mathcal{Z}_\star} = \frac{Z\omega_j\pi^\star_j(dx)}{(4J-2)Z} = \frac{\omega_j}{4J-2}\pi^\star_j(dx). \tag{C.72}$$

Recall from Equations (C.16) and (C.47) that the time-step contribution to $\mathcal{E}_P$ sums squared changes across $\widehat{C}_j$ with factor $1/(2\mathcal{Z})$. Similarly, we define $\mathcal{E}_\star$ using C_j and $\mathcal{Z}_\star$. With $\overline{C}_j = C_j/Z$, write

$$S(f) := \sum_{j=1}^{J} \sum_{d\in\{\mathrm{F},\mathrm{B}\}} \mathbb{E}_{\overline{C}_j}\left[(f_j^d(X_j) - f_{j-1}^d(X_{j-1}))^2\right], \tag{C.73}$$

so that

$$\mathcal{E}_\star(f,f) := \frac{1}{2\mathcal{Z}_\star}\sum_{j=1}^{J}\sum_{d\in\{\mathrm{F},\mathrm{B}\}} \iint [f_j^d(x) - f_{j-1}^d(y)]^2 C_j(dx,dy) = \frac{Z}{2\mathcal{Z}_\star} S(f) = \frac{S(f)}{2(4J-2)}. \tag{C.74}$$

First compare the measures. The recursion $V_j = K_j^{\mathrm{F}} V_{j-1}$ with $V_0 = r$ gives $r_{\min} \leq V_j \leq 1$. For $j \geq 2$,

$$v_{\min} \leq \frac{d\widehat{C}_j}{dC_j}(x,y) = \frac{\sqrt{\widehat{V}_j(x)}\sqrt{\widehat{V}_{j-1}(y)}}{V_{j-1}(y)} \leq r_{\min}^{-1}.$$

Set $a := v_{\min}$ and $b := r_{\min}^{-1}$. Since $\widehat{C}_1 = C_1$ and $a \le 1 \le b$,

$$aC_j \le \widehat{C}_j \le bC_j, \qquad 1 \le j \le J. \tag{C.75}$$

The marginals of C_j are $Z\pi_j^\star$ and $Z\pi_{j-1}^\star$. For $1 \le j < J$, summing the two incident marginals gives

$$\mu_j(dx) = \widehat{C}_j(dx, \mathcal{X}) + \widehat{C}_{j+1}(\mathcal{X}, dx) \ge a\big[C_j(dx, \mathcal{X}) + C_{j+1}(\mathcal{X}, dx)\big] = 2aZ\pi_j^\star(dx).$$

The upper bound follows in the same way with b. Each boundary has only one incident edge. Thus, for $0 \le j \le J$, these bounds and their sums over time-direction states yield

$$aZ\omega_j\pi_j^\star \le \mu_j \le bZ\omega_j\pi_j^\star, \qquad a\mathcal{Z}_\star \le \mathcal{Z} \le b\mathcal{Z}_\star, \qquad \frac{a}{b}\Pi_\star \le \Pi \le \frac{b}{a}\Pi_\star. \tag{C.76}$$

Thus every $f \in L^2(\Pi)$ also belongs to $L^2(\Pi_\star)$, and $\mathrm{Var}_\Pi(f) \le (b/a)\,\mathrm{Var}_{\Pi_\star}(f)$.

We now bound the reference variance. All following expectations, variances, and covariances are under $\pi_{0:J}^\star$. The endpoint correlation bound gives

$$\begin{aligned}
\mathbb{E}[(f_J(X_J) - f_0(X_0))^2] &= \mathrm{Var}(f_0(X_0)) + \mathrm{Var}(f_J(X_J)) - 2\,\mathrm{Cov}(f_0(X_0), f_J(X_J)) \\
&\quad + (\mathbb{E}[f_J(X_J)] - \mathbb{E}[f_0(X_0)])^2 \\
&\ge \mathrm{Var}(f_0(X_0)) + \mathrm{Var}(f_J(X_J)) - 2\rho_{\mathrm{end}}\sqrt{\mathrm{Var}(f_0(X_0))\,\mathrm{Var}(f_J(X_J))} \\
&\ge (1-\rho_{\mathrm{end}})[\mathrm{Var}(f_0(X_0)) + \mathrm{Var}(f_J(X_J))] \ge (1-\rho_{\mathrm{end}})\,\mathrm{Var}(f_0(X_0)).
\end{aligned}$$

The last step uses $2\sqrt{uv} \le u + v$. For either direction, telescoping along the path and applying Cauchy–Schwarz yields

$$\begin{aligned}
\mathbb{E}[(f_j^d(X_j) - f_0(X_0))^2] &= \mathbb{E}\left[\left(\sum_{\ell=1}^{j}[f_\ell^d(X_\ell) - f_{\ell-1}^d(X_{\ell-1})]\right)^2\right] \\
&\le j\sum_{\ell=1}^{j}\mathbb{E}[(f_\ell^d(X_\ell) - f_{\ell-1}^d(X_{\ell-1}))^2] \le JS(f).
\end{aligned}$$

At $j = J$, this gives $\mathrm{Var}(f_0(X_0)) \le \mathbb{E}[(f_J(X_J) - f_0(X_0))^2]/(1-\rho_{\mathrm{end}}) \le JS(f)/(1-\rho_{\mathrm{end}})$. Consequently, for every available (j, d),

$$\begin{aligned}
\mathbb{E}[(f_j^d(X_j) - \mathbb{E}[f_0(X_0)])^2] &\le 2\mathbb{E}[(f_j^d(X_j) - f_0(X_0))^2] + 2\,\mathrm{Var}(f_0(X_0)) \\
&\le 2JS(f) + \frac{2JS(f)}{1-\rho_{\mathrm{end}}} \le \frac{4JS(f)}{1-\rho_{\mathrm{end}}}.
\end{aligned} \tag{C.77}$$

Observing that

$$\Pi_\star[(f - \mathbb{E}[f_0(X_0)])^2] = \frac{1}{4J-2}\sum_{j=0}^{J}\omega_j\sum_d \mathbb{E}_{\pi_j^\star}\left[(f_j^d(X_j) - \mathbb{E}[f_0(X_0)])^2\right],$$

where d ranges over both directions at interior checkpoints and only the prescribed direction at each boundary, we apply Equation (C.77) to each term. The weights sum to $1 + 4(J-1) + 1 = 4J - 2$, so

$$\mathrm{Var}_{\Pi_\star}(f) \le \Pi_\star[(f - \mathbb{E}[f_0(X_0)])^2] \le \frac{4JS(f)}{1-\rho_{\mathrm{end}}} = \frac{8J(4J-2)}{1-\rho_{\mathrm{end}}}\mathcal{E}_\star(f,f) \le \frac{32J^2}{1-\rho_{\mathrm{end}}}\mathcal{E}_\star(f,f). \tag{C.78}$$

Finally, retain only the moves between adjacent checkpoints in $\mathcal{E}_P$. By Equations (C.75) and (C.76),

$$\begin{aligned}
\mathcal{E}_P(f,f) &\overset{(\circ)}{\ge} \frac{1}{2\mathcal{Z}}\sum_{j=1}^{J}\sum_{d\in\{\mathrm{F},\mathrm{B}\}}\iint [f_j^d(x) - f_{j-1}^d(y)]^2\widehat{C}_j(dx,dy) \\
&\ge \frac{aZ}{2\mathcal{Z}}S(f) = \frac{a\mathcal{Z}_\star}{\mathcal{Z}}\mathcal{E}_\star(f,f) \ge \frac{a}{b}\mathcal{E}_\star(f,f).
\end{aligned} \tag{C.79}$$

Here $(\circ)$ follows from Equation (C.47) by dropping nonnegative flip terms; stays contribute zero. Combining this with Equation (C.78) and the variance bound above gives

$$\mathcal{E}_P(f,f) \ge \frac{a}{b}\frac{1-\rho_{\mathrm{end}}}{32J^2}\,\mathrm{Var}_{\Pi_\star}(f) \ge \left(\frac{a}{b}\right)^2\frac{1-\rho_{\mathrm{end}}}{32J^2}\,\mathrm{Var}_\Pi(f) = \frac{(1-\rho_{\mathrm{end}})r_{\min}^2 v_{\min}^2}{32J^2}\,\mathrm{Var}_\Pi(f).$$

The bound is independent of χ. Since $0 < \gamma_{\mathrm{CS}} \le 1$, Corollary C.1 gives Equation (C.71). □

If the clean and noise endpoints are independent under the unweighted reference path, weighting by $r(X_0)/Z$ preserves their independence, so $\rho_{\mathrm{end}} = 0$. In this case, Equation (C.70) gives $\gamma_{\mathrm{CS}} = r_{\min}^2 v_{\min}^2/(32J^2)$.

The convergence bound in Equation (C.71) also applies to the standard noise initialization $\nu_J := \delta_J \otimes \pi_J \otimes \delta_{\mathrm{F}}$. Indeed, Equation (C.69) gives $\sqrt{\widehat{V}_J} \geq \sqrt{v_{\min}}$ and $F_J = K_J^{\mathrm{F}}\sqrt{\widehat{V}_{J-1}} \geq \sqrt{v_{\min}}$ for $J \geq 2$, so the same calculation as in Equation (C.68) yields

$$D_{\nu_J}^2 = \int \left(\frac{d\nu_J}{d\Pi} - 1\right)^2 d\Pi = \mathcal{Z} \int_{\mathcal{X}} \frac{\pi_J(dx)}{\sqrt{\widehat{V}_J(x)}F_J(x)} - 1 \leq \frac{\mathcal{Z}}{v_{\min}} - 1 < \infty. \tag{C.80}$$

Effect of process-verifier error. The preceding convergence bounds do not explicitly quantify the error in $\widehat{V}_j$ relative to the true process verifier V_j. We now incorporate this error into a lower bound on the Poincaré constant under an average verifier-error assumption.

Let P^V be the exact CS-VGB kernel with $\widehat{V}_j$ replaced by the true process verifier V_j in Equation (C.27), and let Π^V be its stationary law. Both chains use the same reference transitions and terminal anchoring $\widehat{V}_0 = V_0 = r$. The corresponding exact-verifier edge-conductance measures are $C_1^V := C_1$ and, for $2 \leq j \leq J$,

$$C_j^V(dx, dy) := \pi_j(dx)K_j^{\mathrm{F}}(dy \mid x)\sqrt{V_j(x)V_{j-1}(y)} = \sqrt{\frac{V_j(x)}{V_{j-1}(y)}}C_j(dx, dy). \tag{C.81}$$

To quantify the discrepancy from this exact-verifier construction, we measure the relative error in $\widehat{V}_j$ in mean square under the reward-tilted marginal $\pi_j^\star$:

$$\varepsilon_{\mathrm{av}} := \max_{1\leq j\leq J} \left(\mathbb{E}_{\pi_j^\star}\left[\left(\frac{\widehat{V}_j(X_j)}{V_j(X_j)} - 1\right)^2\right]\right)^{1/2}. \tag{C.82}$$

To relate this average error to the convergence rate, we also impose an integrability condition on the exact-verifier transitions. As in Equation (C.45), let P_{move}^V retain only moves between adjacent checkpoints and define

$$M^V := \frac{P_{\mathrm{move}}^V + (P_{\mathrm{move}}^V)^*}{2}, \qquad M^V(z, dz') = m^V(z, z')\Pi^V(dz'), \tag{C.83}$$

where the adjoint is taken in $L^2(\Pi^V)$. For each finite $p > 2$, define

$$H_p := \left(\iint m^V(z, z')^p \, \Pi^V(dz)\Pi^V(dz')\right)^{1/p}. \tag{C.84}$$

Lemma C.5 gives a sufficient condition for $H_p < \infty$ under paired Gaussian transitions. We assume this moment is finite for at least one such p and minimize the error bound over all exponents satisfying this condition:

$$\Delta_{\mathrm{av}} := \frac{6}{r_{\min}v_{\min}^2} \inf_{\substack{2<p<\infty \\ H_p<\infty}} \left\{\frac{H_p}{r_{\min}^{1/p}}\varepsilon_{\mathrm{av}}^{1-2/p}\right\}. \tag{C.85}$$

Theorem C.4 (Average verifier error and the Poincaré constant)**.** *Let P be the exact CS-VGB one-step kernel with $J \geq 2$. Assume the reward and verifier bounds in Equation (C.69), the endpoint correlation bound in Equation (C.42), and $H_p < \infty$ in Equation (C.84) for at least one finite $p > 2$. For every directional persistence parameter $\chi \in [0, 1]$ and every $f \in L^2(\Pi)$, we have*

$$\mathcal{E}_P(f, f) \geq \gamma_{\mathrm{CS}} \operatorname{Var}_\Pi(f), \qquad \gamma_V := \frac{(1 - \rho_{\mathrm{end}}) r_{\min}}{32 J^2}.$$

Here γ_V is the exact-verifier lower bound, and γ_{CS} can be chosen independently of χ to satisfy

$$\gamma_{\mathrm{CS}} \geq \max\left\{\gamma_V r_{\min} v_{\min}^2,\ \gamma_V - \Delta_{\mathrm{av}}\right\} \geq \gamma_V \left(1 + \frac{\Delta_{\mathrm{av}}}{\gamma_V r_{\min} v_{\min}^2}\right)^{-1} > 0. \tag{C.86}$$

The convergence and clean-conditional bounds in Corollary C.1 apply with $\gamma = \gamma_{\mathrm{CS}}$; the clean stationary law remains $\pi_0^\star$.

Proof. Use a superscript V also for the exact-verifier incident measures and normalizer. The reward bounds and the recursion for V_j give $r_{\min} \leq V_j \leq 1$, so

$$\sqrt{r_{\min}}\, C_j \leq C_j^V \leq r_{\min}^{-1/2} C_j, \qquad 1 \leq j \leq J.$$

Summing incident marginals and normalizing therefore gives

$$\sqrt{r_{\min}}\, \mathcal{Z}_\star \leq \mathcal{Z}^V \leq r_{\min}^{-1/2} \mathcal{Z}_\star, \qquad \Pi^V \leq \frac{\mathcal{Z}_\star}{\sqrt{r_{\min}}\, \mathcal{Z}^V} \Pi_\star.$$

Since the mean minimizes the expected squared deviation,

$$\operatorname{Var}_{\Pi^V}(f) \leq \frac{\mathcal{Z}_\star}{\sqrt{r_{\min}}\, \mathcal{Z}^V} \operatorname{Var}_{\Pi_\star}(f).$$

Let $\mathcal{L}_{\mathrm{move}}^V = cI - M^V$, where $c = M^V \mathbf{1}$ equals $1/2$ in the interior and 1 at the boundaries. Applying the lower edge-conductance measure bound to the movement energy gives

$$\begin{aligned}
\langle f, \mathcal{L}_{\mathrm{move}}^V f \rangle_{\Pi^V} &= \frac{1}{2\mathcal{Z}^V} \sum_{j=1}^{J} \sum_{d \in \{\mathrm{F}, \mathrm{B}\}} \iint [f_j^d(x) - f_{j-1}^d(y)]^2 \, C_j^V(dx, dy) \\
&\geq \frac{\sqrt{r_{\min}}\, \mathcal{Z}_\star}{\mathcal{Z}^V} \mathcal{E}_\star(f, f) \geq \frac{(1 - \rho_{\mathrm{end}}) r_{\min}}{32 J^2} \operatorname{Var}_{\Pi^V}(f),
\end{aligned}$$

where the last step uses Equation (C.78) and the preceding variance comparison. Thus

$$\langle f, \mathcal{L}_{\mathrm{move}}^V f \rangle_{\Pi^V} \geq \gamma_V \operatorname{Var}_{\Pi^V}(f), \qquad \Pi^V \leq r_{\min}^{-1} \Pi_\star. \tag{C.87}$$

Write $\Delta(j, x, d) := \widehat{V}_j(x)/V_j(x) - 1$, with $\Delta = 0$ at $j = 0$, and fix $p > 2$ with $H_p < \infty$. Since $\Pi_\star$ averages the time marginals $\pi_j^\star$, Equations (C.82) and (C.87) gives

$$\|\Delta\|_{2,\Pi^V}^2 \leq r_{\min}^{-1} \|\Delta\|_{2,\Pi_\star}^2 \leq r_{\min}^{-1} \max_{1 \leq j \leq J} \mathbb{E}_{\pi_j^\star}\left[\left(\frac{\widehat{V}_j}{V_j} - 1\right)^2\right] = r_{\min}^{-1} \varepsilon_{\mathrm{av}}^2.$$

Also, $0 < \widehat{V}_j \leq 1$ and $V_j \geq r_{\min}$ imply

$$0 < \frac{\widehat{V}_j}{V_j} \leq r_{\min}^{-1} \quad \Longrightarrow \quad |\Delta| \leq \max\{1, r_{\min}^{-1} - 1\} \leq r_{\min}^{-1}.$$

Thus, for $q := 2p/(p-2)$,

$$\|\Delta\|_{2,\Pi^V} \leq r_{\min}^{-1/2} \varepsilon_{\mathrm{av}}, \quad \|\Delta\|_\infty \leq r_{\min}^{-1} \implies \|\Delta\|_{q,\Pi^V} \leq \|\Delta\|_\infty^{1-2/q} \|\Delta\|_{2,\Pi^V}^{2/q} \leq r_{\min}^{-(1-1/q)} \varepsilon_{\mathrm{av}}^{2/q}.$$

For each edge, define its ratio relative to the exact-verifier edge-conductance measure by $J_1 := 1$ and

$$J_j(x, y) := \frac{d\widehat{C}_j}{dC_j^V}(x, y) = \sqrt{\frac{\widehat{V}_j(x)\widehat{V}_{j-1}(y)}{V_j(x) V_{j-1}(y)}} = \sqrt{\xi_j(x)\xi_{j-1}(y)}, \qquad j \geq 2, \qquad \xi_j := \widehat{V}_j / V_j.$$

For lifted states z, z' on edge j, write $J_j(z, z')$ for this ratio evaluated at the noisier sample x and cleaner sample y, in either orientation. The notation $\iint_{j\leftrightarrow j-1}$ means integrating over $z = (j, x, d)$ and $z' = (j-1, y, d)$ and also over the reversed pair, for $d \in \{\mathrm{F}, \mathrm{B}\}$. At the boundaries, both direction labels refer to the single state $(0, x, \mathrm{B})$ or (J, x, F). The verifier bounds give $v_{\min} \le J_j \le r_{\min}^{-1}$. For $j \ge 2$, the identity $|\sqrt{\xi} - 1| = |\xi - 1|/(\sqrt{\xi} + 1) \le |\xi - 1|$ gives

$$|J_j(z, z') - 1| \le \sqrt{\xi_j(x)}\,|\sqrt{\xi_{j-1}(y)} - 1| + |\sqrt{\xi_j(x)} - 1| \le r_{\min}^{-1/2}\big(|\Delta(z)| + |\Delta(z')|\big).$$

The same inequality holds on the clean edge since $J_1 = 1$.

Define the incident-mass ratio $s(j, x, d) := d\mu_j/d\mu_j^V(x)$. For an interior state $z = (\ell, x, d)$, $1 \le \ell < J$, the two incident edges give

$$\mu_\ell^V(dx) = C_\ell^V(dx, \mathcal{X}) + C_{\ell+1}^V(\mathcal{X}, dx),$$
$$\mu_\ell(dx) = \int J_\ell(x, y) C_\ell^V(dx, dy) + \int J_{\ell+1}(y, x) C_{\ell+1}^V(dy, dx).$$

Applying the symmetric edge identity in Equation (C.48), established in the proof of Theorem C.3, to the exact-verifier chain gives

$$\mu_\ell^V(dx) M^V\big(z, (\ell-1, dy, d)\big) = \tfrac{1}{2} C_\ell^V(dx, dy), \quad \mu_\ell^V(dx) M^V\big(z, (\ell+1, dy, d)\big) = \tfrac{1}{2} C_{\ell+1}^V(dy, dx).$$

Multiplying by the corresponding J_j and adding therefore yields

$$\mu_\ell^V(dx) \sum_{j=1}^{J} \int_{j\leftrightarrow j-1} J_j(z, z') M^V(z, dz') = \tfrac{1}{2}\mu_\ell(dx) = \tfrac{1}{2} s(z) \mu_\ell^V(dx).$$

Comparing densities with respect to μ_ℓ^V gives the sum as $c(z)s(z)$, since $c(z) = 1/2$ in the interior. At a boundary, the two direction branches of its sole incident edge together give the same identity with $c(z) = 1$. Dividing by $c(z)$ and normalizing the incident measures gives

$$s(z) = \frac{1}{c(z)} \sum_{j=1}^{J} \int_{j\leftrightarrow j-1} J_j(z, z') M^V(z, dz'), \qquad \frac{d\Pi}{d\Pi^V} = \frac{\mathcal{Z}^V}{\mathcal{Z}} s. \tag{C.88}$$

In particular, since $s(z)$ is a weighted average of the edge ratios J_j satisfying $v_{\min} \le J_j \le r_{\min}^{-1}$, we have $v_{\min} \le s \le r_{\min}^{-1}$. Since $M^V(z, \cdot)/c(z)$ is a probability kernel,

$$\|s - 1\|_{q,\Pi^V}^q \overset{(\circ)}{\le} 2 \sum_{j=1}^{J} \iint_{j\leftrightarrow j-1} |J_j(z, z') - 1|^q \, \Pi^V(dz) M^V(z, dz')$$
$$\overset{(\triangle)}{\le} 2^{q+1} r_{\min}^{-q/2} \int c(z) |\Delta(z)|^q \Pi^V(dz) \le 2^{q+1} r_{\min}^{-q/2} \|\Delta\|_{q,\Pi^V}^q.$$

Here $(\circ)$ uses Jensen's inequality and $1/c(z) \le 2$; $(\triangle)$ uses the bound on $|J_j - 1|$, $(a+b)^q \le 2^{q-1}(a^q + b^q)$ for $a, b \ge 0$, and the symmetry of M^V with respect to Π^V; the last step uses $c \le 1$. Taking qth roots gives $\|s - 1\|_{q,\Pi^V} \le 4 r_{\min}^{-1/2} \|\Delta\|_{q,\Pi^V}$.

To compare the operators despite their different stationary laws, use the norm-preserving map $Uf := \sqrt{d\Pi/d\Pi^V} f$. On edge j, the identities $\widehat{C}_j = J_j C_j^V$ and $\mu_\ell = s\mu_\ell^V$ give $M(z, dz') = [J_j(z, z')/s(z)] M^V(z, dz')$ by Equation (C.48). Substituting $d\Pi/d\Pi^V = (\mathcal{Z}^V/\mathcal{Z})s$ from Equation (C.88), the constant $\mathcal{Z}^V/\mathcal{Z}$ cancels between U and U^{-1}:

$$(UMU^{-1} f)(z) = \sqrt{\frac{d\Pi}{d\Pi^V}(z)} \int \frac{f(z')}{\sqrt{\frac{d\Pi}{d\Pi^V}(z')}} M(z, dz') = \int \frac{\sqrt{s(z)}}{\sqrt{s(z')}} f(z') M(z, dz')$$
$$= \sum_{j=1}^{J} \int_{j\leftrightarrow j-1} \frac{J_j(z, z') f(z')}{\sqrt{s(z)s(z')}} M^V(z, dz') = \sum_{j=1}^{J} \int_{j\leftrightarrow j-1} \frac{J_j(z, z') m^V(z, z') f(z')}{\sqrt{s(z)s(z')}} \Pi^V(dz').$$

Thus the density on edge j is $m^V(z, z') J_j(z, z')/\sqrt{s(z)s(z')}$. For $F(a, b) := (ab)^{-1/2}$, both partial derivatives have absolute value at most $1/(2v_{\min}^2)$ when $a, b \ge v_{\min}$. Since $s(z), s(z') \ge v_{\min}$ and

$v_{\min} \le 1$,

$$\left|\frac{J_j(z,z')}{\sqrt{s(z)s(z')}} - 1\right| \le \frac{|J_j(z,z')-1|}{\sqrt{s(z)s(z')}} + \left|\frac{1}{\sqrt{s(z)s(z')}} - 1\right| \overset{(\circ)}{\le} v_{\min}^{-1}|J_j(z,z')-1| + \frac{|s(z)-1|+|s(z')-1|}{2v_{\min}^2}.$$

Here $(\circ)$ applies the mean value theorem to F between $(s(z), s(z'))$ and $(1,1)$. The edge domains are disjoint: each state pair (z,z') belongs to at most one checkpoint index edge. Thus, summing a nonnegative integrand over these domains gives at most its integral over all state pairs. In particular,

$$\sum_{j=1}^{J} \iint_{j\leftrightarrow j-1} |\Delta(z)|^q \, \Pi^V(dz)\Pi^V(dz') \le \iint |\Delta(z)|^q \, \Pi^V(dz)\Pi^V(dz') = \|\Delta\|_{q,\Pi^V}^q.$$

Using the same bound for $\Delta(z')$ and $s-1$, the preceding estimates give

$$\left(\sum_{j=1}^{J} \iint_{j\leftrightarrow j-1} \left|\frac{J_j(z,z')}{\sqrt{s(z)s(z')}} - 1\right|^q \Pi^V(dz)\Pi^V(dz')\right)^{1/q} \overset{(\circ)}{\le} \frac{6}{\sqrt{r_{\min}}v_{\min}^2}\|\Delta\|_{q,\Pi^V},$$

where $(\circ)$ uses the triangle inequality, $|J_j - 1| \le r_{\min}^{-1/2}(|\Delta(z)| + |\Delta(z')|)$, and $\|s-1\|_{q,\Pi^V} \le 4r_{\min}^{-1/2}\|\Delta\|_{q,\Pi^V}$, giving a coefficient $2/(\sqrt{r_{\min}}v_{\min}) + 4/(\sqrt{r_{\min}}v_{\min}^2) \le 6/(\sqrt{r_{\min}}v_{\min}^2)$ since $v_{\min} \le 1$. Since $1/2 = 1/p + 1/q$, recalling the operator norm $\|A\|_{2\to 2} := \sup_{\|f\|_{2,\Pi^V}=1} \|Af\|_{2,\Pi^V}$ on $L^2(\Pi^V)$, we obtain

$$\|UMU^{-1} - M^V\|_{2\to 2} \overset{(\circ)}{\le} \left(\sum_{j=1}^{J} \iint_{j\leftrightarrow j-1} m^V(z,z')^2 \left|\frac{J_j(z,z')}{\sqrt{s(z)s(z')}} - 1\right|^2 \Pi^V(dz)\Pi^V(dz')\right)^{1/2} \overset{(\triangle)}{\le} \frac{6H_p}{\sqrt{r_{\min}}v_{\min}^2}\|\Delta\|_{q,\Pi^V} \le \frac{6H_p}{r_{\min}^{1+1/p}v_{\min}^2}\varepsilon_{\mathrm{av}}^{1-2/p}. \tag{C.89}$$

Here $(\circ)$ is Cauchy–Schwarz for integral kernels and $(\triangle)$ applies Hölder's inequality over the disjoint edge domains. Taking the infimum over admissible p gives $\|UMU^{-1} - M^V\|_{2\to 2} \le \Delta_{\mathrm{av}}$.

Finally, set $L := U\mathcal{L}_{\mathrm{move}}U^{-1}$ and $\psi := U\mathbf{1}$. Then L is nonnegative and self-adjoint, $L\psi = 0$, and $\|\psi\|_{2,\Pi^V} = 1$. Multiplication by c commutes with U, so

$$L = U(cI - M)U^{-1} = cI - UMU^{-1}, \qquad \mathcal{L}_{\mathrm{move}}^V = cI - M^V.$$

Consequently,

$$L - \mathcal{L}_{\mathrm{move}}^V = M^V - UMU^{-1}, \qquad \|L - \mathcal{L}_{\mathrm{move}}^V\|_{2\to 2} = \|UMU^{-1} - M^V\|_{2\to 2} \le \Delta_{\mathrm{av}}.$$

For a function $g \in L^2(\Pi^V)$ such that $g \perp \psi$ and $\|g\|_{2,\Pi^V} = 1$, use the orthogonality of g and ψ to choose scalars α, β so that $h = \alpha\psi + \beta g$ satisfies $\Pi^V(h) = 0$ and $\|h\|_{2,\Pi^V} = 1$. Then $\alpha^2 + \beta^2 = 1$, and

$$\gamma_V \overset{(\circ)}{\le} \langle h, \mathcal{L}_{\mathrm{move}}^V h\rangle_{\Pi^V} \overset{(\triangle)}{\le} \langle h, Lh\rangle_{\Pi^V} + \Delta_{\mathrm{av}} \overset{(\square)}{=} \beta^2\langle g, Lg\rangle_{\Pi^V} + \Delta_{\mathrm{av}} \le \langle g, Lg\rangle_{\Pi^V} + \Delta_{\mathrm{av}},$$

where $(\circ)$ uses Equation (C.87) and $\mathrm{Var}_{\Pi^V}(h) = 1$; $(\triangle)$ uses $|\langle h, (\mathcal{L}_{\mathrm{move}}^V - L)h\rangle_{\Pi^V}| \le \|\mathcal{L}_{\mathrm{move}}^V - L\|_{2\to 2}\|h\|_{2,\Pi^V}^2 \le \Delta_{\mathrm{av}}$; and $(\square)$ uses $L\psi = 0$ and self-adjointness, so $\langle \psi, Lg\rangle_{\Pi^V} = \langle L\psi, g\rangle_{\Pi^V} = 0$. The last inequality uses $\beta^2 \le 1$ and $\langle g, Lg\rangle_{\Pi^V} \ge 0$.

For nonconstant f, take $g = U(f - \Pi(f))/\sqrt{\mathrm{Var}_\Pi(f)}$. Since U preserves inner products, $g \perp \psi$ and $\|g\|_{2,\Pi^V} = 1$. Moreover, $L = U\mathcal{L}_{\mathrm{move}}U^{-1}$ and the self-adjoint operator $\mathcal{L}_{\mathrm{move}}$ satisfies $\mathcal{L}_{\mathrm{move}}\mathbf{1} = 0$, so

$$\langle g, Lg\rangle_{\Pi^V} = \frac{\langle f - \Pi(f), \mathcal{L}_{\mathrm{move}}(f - \Pi(f))\rangle_\Pi}{\mathrm{Var}_\Pi(f)} = \frac{\langle f, \mathcal{L}_{\mathrm{move}}f\rangle_\Pi}{\mathrm{Var}_\Pi(f)}.$$

Multiplying the preceding bound by $\mathrm{Var}_\Pi(f)$ and using Equation (C.47) gives

$$\mathcal{E}_P(f,f) \ge \langle f, \mathcal{L}_{\mathrm{move}}f\rangle_\Pi \ge (\gamma_V - \Delta_{\mathrm{av}})\,\mathrm{Var}_\Pi(f).$$

For constant f, both sides vanish. Combining this with the error-independent bound $\gamma_V r_{\min} v_{\min}^2$ from Equation (C.70) gives the maximum in Equation (C.86), uniformly in χ. The multiplicative bound follows from $\max\{a, 1-x\} \geq a/(a+x)$ for $0 < a \leq 1$ and $x \geq 0$, with $a = r_{\min} v_{\min}^2$ and $x = \Delta_{\mathrm{av}}/\gamma_V$. □

As $\varepsilon_{\mathrm{av}} \to 0$, the lower bound in Equation (C.86) recovers γ_V; for larger errors, it retains the positive bound from Corollary C.3.

Lastly, we remark that the moment assumption $H_p < \infty$ in Equation (C.84) holds for paired Gaussian transitions under a slightly stronger clean-tail condition than in Corollary C.2, as stated below.

Lemma C.5 (Higher-order moments for paired Gaussian transitions)**.** *Assume the Gaussian pairing in Equation (C.51) and the reward and verifier bounds in Equation (C.69). If*

$$\mathbb{E}_{\pi_0} e^{c\|X_0\|^2} < \infty \quad \textit{for some} \quad c > \max_{1 \leq j \leq J} \frac{2A_j^2}{\sigma_j^2}, \tag{C.90}$$

where A_j and σ_j are the cumulative signal scale and renoising noise scale from Corollary C.2, respectively, then $H_p < \infty$ in Equation (C.84) for some finite $p > 2$. In particular, bounded support of π_0 ensures this tail condition.

Proof. We repeat the Gaussian calculation in the proof of Corollary C.2 for the pth moment of $\frac{dJ_j^{\mathrm{F}}}{d(\pi_j \otimes \pi_{j-1})}(x,y) = \frac{\varphi_{\sigma_j^2}(x - a_j y)}{\pi_j(x)}$. Using $\pi_j \geq b_{j,v_j}\varphi_{v_j}$ and $v_j = \tau_j^2 - \sigma_j^2/4$ from Equations (C.59) and (C.60), the integrations in Equations (C.62) and (C.63) bound this moment by a finite multiple of $\mathbb{E}_{\pi_0} e^{c_j(p)\|X_0\|^2}$, where

$$\eta_j(p) := \frac{p(p-1)a_j^2}{2[pv_j - (p-1)\sigma_j^2]}, \qquad c_j(p) := \frac{\eta_j(p) A_{j-1}^2}{1 - 2\eta_j(p)\tau_{j-1}^2},$$

provided both denominators are positive. At $p = 2$, we have $\eta_j(2) = \eta_j$, so both denominators are strictly positive by Equation (C.61), and $c_j(2) = 2A_j^2/\sigma_j^2 < c$ by Equation (C.90). Continuity with respect to p and the finiteness of J therefore give a common $p > 2$ with

$$\max_{1 \leq j \leq J} \left\| \frac{dJ_j^{\mathrm{F}}}{d(\pi_j \otimes \pi_{j-1})} \right\|_{L^p(\pi_j \otimes \pi_{j-1})} < \infty.$$

Since $J_j^{\mathrm{B}} = J_j^{\mathrm{F}}$ and $r_{\min} \leq V_j \leq 1$ by Equation (C.27), applying Lemma C.7 to the exact-verifier kernel P^V and its stationary law Π^V gives $p_{\mathrm{move}}^V \in L^p(\Pi^V \otimes \Pi^V)$, where p_{move}^V is the density of P_{move}^V relative to Π^V. By Equations (C.83) and (C.84),

$$m^V(z,z') = \frac{p_{\mathrm{move}}^V(z,z') + p_{\mathrm{move}}^V(z',z)}{2}, \qquad H_p^p \leq \iint p_{\mathrm{move}}^V(z,z')^p \, \Pi^V(dz)\Pi^V(dz') < \infty,$$

where the inequality uses $((a+b)/2)^p \leq (a^p + b^p)/2$ for $a, b \geq 0$ and exchanges z, z' in the second integral. □

C.4 Momentum and First Hitting Time of the Clean Boundary

In practice, CS-VGB runs under a finite computational budget and collects samples at the clean checkpoint. Visits to this checkpoint are therefore a practical concern distinct from mixing time. We analyze the *first hitting time* $\tau_0 := \inf\{n \geq 0 : j_n = 0\}$, where j_n is the checkpoint index after n updates, under an average directional-imbalance condition.

Let $Y_J \sim \pi_J$ and $Y_{j-1} \mid Y_j \sim Q_j^{\mathrm{F}}(\cdot \mid Y_j)$ for $j = J, \ldots, 1$. The expectations below are taken along this denoising path. Write $(u)_+ := \max\{u, 0\}$.

Proposition C.2 (First hitting time under average directional imbalance)**.** *Assume the conditions of Theorem C.1, with $J \geq 2$, and use the exact CS-VGB kernel P. Suppose*

$$K_{J-1}^{\mathrm{B},\rho}(dx \mid y) = \pi_J(dx) \quad \text{for } \pi_{J-1}\text{-almost every } y. \tag{C.91}$$

Assume that, for some $C_{\mathrm{HT}} \geq 0$,

$$\max_{1\leq j<J} \mathbb{E}\left[\left(\log \frac{a_j^{\mathrm{B}}(Y_j)}{a_j^{\mathrm{F}}(Y_j)}\right)_+\right] \leq \frac{C_{\mathrm{HT}}}{J-1}, \qquad 1-\chi \leq \frac{1}{J-1}, \quad \chi \in [0,1]. \tag{C.92}$$

If $v_{\min} \leq \widehat{V}_j \leq 1$ π_j-almost surely for $1 \leq j \leq J$, with $v_{\min} > 0$, then, for π_J-almost every x,

$$J \leq \mathbb{E}_{(J,x,\mathrm{F})}[\tau_0] \leq e^{C_{\mathrm{HT}}+1}\frac{3Z+4J-5}{v_{\min}}. \tag{C.93}$$

For fixed C_{HT}, $Z > 0$, and $v_{\min} > 0$, this gives $\mathbb{E}_{(J,x,\mathrm{F})}[\tau_0] = \Theta(J)$.

Proof. Write $M_j := \widehat{C}_j(\mathcal{X}, \mathcal{X})$. We first bound the ratio $\mathcal{Z}/M_J$. Terminal anchoring and Equations (C.2) and (C.3) give

$$M_1 = Z, \qquad \mathcal{Z} = 3Z + 4\sum_{j=2}^{J-1} M_j + 3M_J.$$

The verifier bounds give $v_{\min} \leq \sqrt{\widehat{V}_j \widehat{V}_{j-1}} \leq 1$, hence $v_{\min} \leq M_j \leq 1$ for $j \geq 2$ and

$$\frac{\mathcal{Z}}{M_J} \leq \frac{3Z+4J-5}{v_{\min}}.$$

In particular, $M_J \geq v_{\min} > 0$. We next bound the time needed to reach checkpoint 0 by considering repeated attempts from checkpoint J.

The corrected edge identity and Equation (C.91) give $\pi_J(dx)K_J^{\mathrm{F}}(dy \mid x) = \pi_J(dx)\pi_{J-1}(dy)$, hence $K_J^{\mathrm{F}}(dy \mid x) = \pi_{J-1}(dy)$ for π_J-almost every x. Substituting this and Equation (C.91) into the definitions of the normalized transitions gives

$$Q_J^{\mathrm{F}}(dy \mid x) = \frac{\sqrt{\widehat{V}_{J-1}(y)}\,\pi_{J-1}(dy)}{\int \sqrt{\widehat{V}_{J-1}}\,d\pi_{J-1}}, \qquad Q_{J-1}^{\mathrm{B},\rho}(dx \mid y) = \frac{\sqrt{\widehat{V}_J(x)}\,\pi_J(dx)}{\int \sqrt{\widehat{V}_J}\,d\pi_J}. \tag{C.94}$$

The F transition does not depend on the current sample x. Thus the first update from J always produces the same distribution of states $(J-1, Y, \mathrm{F})$; denote it by ζ. Since the returning B transition has a density with respect to π_J, the next F update follows the same distribution ζ after every return to J, independently of the returned sample.

We next bound the mean time between visits to J. Starting from checkpoint J, let $\tau_J^+ := \inf\{n \geq 1 : j_n = J\}$ be the first return time to J. Write $D := \{(j,x,d) \in \mathsf{S} : j < J\}$ for the states below J. We use $\Pi_D(A) := \Pi(A \cap D)$ and $P_D(z, A) := P(z, A \cap D)$ for $z \in D$. Thus P_D keeps the original transition probabilities but excludes any move reaching J. For any measurable $A \subseteq D$, stationarity splits the probability of entering A into arrivals from J and arrivals from D:

$$\Pi_D(A) = \int_{\{j=J\}} P(z,A)\Pi(dz) + \int_D P(z,A)\Pi(dz) = \Pi(j=J)\zeta(A) + (\Pi_D P_D)(A).$$

Applying this identity repeatedly gives

$$\Pi_D = \Pi(j=J)\zeta + \Pi_D P_D = \Pi(j=J)\sum_{m=0}^{n-1}\zeta P_D^m + \Pi_D P_D^n.$$

Taking total masses and dropping the nonnegative remainder $(\Pi_D P_D^n)(D)$ gives

$$\Pi(D) \geq \Pi(j=J)\sum_{m=0}^{n-1}(\zeta P_D^m)(D), \qquad \sum_{m\geq 0}(\zeta P_D^m)(D) \leq \frac{\Pi(D)}{\Pi(j=J)},$$

After the first update from J, the state has distribution ζ. For $m \geq 1$, remaining in D for the next m updates has probability

$$(\zeta P_D^m)(D) = \int_D \zeta(dz_0) \int_D P(z_0, dz_1) \cdots \int_D P(z_{m-1}, dz_m) = \mathbb{P}_{(J,x,\mathrm{F})}(\tau_J^+ > m+1).$$

Every intermediate state is integrated over D, so a path that reaches J and then leaves it is not counted. Using $\mathbb{E}[T] = \sum_{n \geq 0} \mathbb{P}(T > n)$ for a nonnegative integer-valued time T, we obtain

$$\mathbb{E}_{(J,x,\mathrm{F})}[\tau_J^+] = 1 + \sum_{m \geq 0} (\zeta P_D^m)(D) \leq 1 + \frac{\Pi(D)}{\Pi(j=J)} = \frac{1}{\Pi(j=J)} = \frac{\mathcal{Z}}{M_J} < \infty. \tag{C.95}$$

Having bounded the return time, we now bound the probability $\mathbb{P}_{(J,x,\mathrm{F})}(\tau_0 < \tau_J^+)$ of visiting 0 before returning to J. We obtain a lower bound by counting only paths that move in direction F without a flip; stays are allowed. At an interior state (j, x, F), each stay leaves the state unchanged, so

$$\mathbb{P}_{(j,x,\mathrm{F})}(\text{move before flip}) = \sum_{n=0}^{\infty} (\chi \ell_j(x))^n p_j^{\mathrm{F}}(x) = \frac{p_j^{\mathrm{F}}(x)}{1 - \chi \ell_j(x)},$$

For $u = a_j^{\mathrm{B}}(x)/a_j^{\mathrm{F}}(x)$, the directional probabilities satisfy

$$p_j^{\mathrm{F}}(x) = \frac{1}{1+u}, \qquad \ell_j(x) = \frac{\min\{1,u\}}{1+u}.$$

Substituting these expressions gives

$$\frac{p_j^{\mathrm{F}}(x)}{1 - \chi \ell_j(x)} = \frac{1}{1 + u - \chi \min\{1,u\}} \geq \frac{e^{-(\log u)_+}}{2 - \chi},$$

because $1 + u - \chi \min\{1,u\} \leq (2-\chi)\max\{1,u\} = (2-\chi)e^{(\log u)_+}$. Along the denoising path generated by $Y_J \sim \pi_J$ and $Y_{j-1} \mid Y_j \sim Q_j^{\mathrm{F}}(\cdot \mid Y_j)$ for $j = J, \ldots, 1$, multiply these move-before-flip probabilities and average over the samples. This counts paths reaching 0 without a flip and therefore gives

$$\begin{aligned}
\mathbb{P}_{(J,x,\mathrm{F})}(\tau_0 < \tau_J^+) &\geq \mathbb{E}\left[\prod_{j=1}^{J-1} \frac{p_j^{\mathrm{F}}(Y_j)}{1 - \chi \ell_j(Y_j)}\right] \\
&\geq (2-\chi)^{-(J-1)} \mathbb{E}\left[\exp\left\{-\sum_{j=1}^{J-1} \left(\log \frac{a_j^{\mathrm{B}}(Y_j)}{a_j^{\mathrm{F}}(Y_j)}\right)_+\right\}\right] \\
&\overset{(\circ)}{\geq} (2-\chi)^{-(J-1)} \exp\left[-\sum_{j=1}^{J-1} \mathbb{E}\left[\left(\log \frac{a_j^{\mathrm{B}}(Y_j)}{a_j^{\mathrm{F}}(Y_j)}\right)_+\right]\right] \overset{(\triangle)}{\geq} e^{-(C_{\mathrm{HT}}+1)},
\end{aligned}$$

where $(\circ)$ is Jensen's inequality, and $(\triangle)$ uses Equation (C.92): the sum is at most C_{HT}, and $(J-1)\log(2-\chi) \leq (J-1)(1-\chi) \leq 1$.

We combine these two bounds by repeating the attempt after each return to J. Let L_i be the number of updates from the start of attempt i to the next visit to J, and let N be the first attempt that visits 0. Let $\mathcal{F}_{i-1}$ be the sigma-algebra generated by the chain up to the start of attempt i. By the strong Markov property, conditional on $\mathcal{F}_{i-1}$, the next attempt succeeds with the same probability $\mathbb{P}_{(J,x,\mathrm{F})}(\tau_0 < \tau_J^+)$, so $\mathbb{P}(N \geq i) = [1 - \mathbb{P}_{(J,x,\mathrm{F})}(\tau_0 < \tau_J^+)]^{i-1}$. The same property and Equation (C.95) give

$$\mathbb{E}[L_i \mid \mathcal{F}_{i-1}] \leq \frac{\mathcal{Z}}{M_J} \qquad \text{almost surely.}$$

The event $\{N \geq i\}$ is determined by the previous attempts and is therefore $\mathcal{F}_{i-1}$-measurable. Therefore, the tower property yields

$$\begin{aligned}
\mathbb{E}[\mathbf{1}_{\{N \geq i\}} L_i] &= \mathbb{E}\left[\mathbb{E}[\mathbf{1}_{\{N \geq i\}} L_i \mid \mathcal{F}_{i-1}]\right] = \mathbb{E}\left[\mathbf{1}_{\{N \geq i\}} \mathbb{E}[L_i \mid \mathcal{F}_{i-1}]\right] \\
&\leq \mathbb{E}\left[\mathbf{1}_{\{N \geq i\}} \frac{\mathcal{Z}}{M_J}\right] = \mathbb{P}(N \geq i) \frac{\mathcal{Z}}{M_J}.
\end{aligned}$$

Finally, the first visit to 0 occurs before the successful attempt ends, giving

$$\mathbb{E}_{(J,x,\mathrm{F})}[\tau_0] \le \sum_{i\ge 1} \mathbb{E}[\mathbf{1}_{\{N\ge i\}} L_i] \le \frac{\mathcal{Z}}{M_J} \sum_{i\ge 1} \big[1 - \mathbb{P}_{(J,x,\mathrm{F})}(\tau_0 < \tau_J^+)\big]^{i-1}$$
$$= \frac{\mathcal{Z}}{M_J\, \mathbb{P}_{(J,x,\mathrm{F})}(\tau_0 < \tau_J^+)} \le e^{C_{\mathrm{HT}}+1} \frac{\mathcal{Z}}{M_J}. \tag{C.96}$$

Substituting the bound on $\mathcal{Z}/M_J$ and using $\tau_0 \ge J$ proves the claim. □

First hitting time under exact balance. If $p_j^{\mathrm{F}} = p_j^{\mathrm{B}} = 1/2$ at every reachable interior state, then

$$\mathbb{E}_{(J,x,\mathrm{F})}[\tau_0] = (2J-1)\frac{1+(1-\chi)J}{2-\chi}. \tag{C.97}$$

At $\chi = 1$, the mean is $2J-1$, compared with J^2 when the direction is resampled at each update (Figure C.2). To obtain this formula, define $f_j := \mathbb{E}_{(j,x,\mathrm{F})}[\tau_0]$ and $b_j := \mathbb{E}_{(j,x,\mathrm{B})}[\tau_0]$, which are independent of x under exact balance. Conditioning on the first update gives, for $1 \le j < J$,

$$f_j = 1 + \frac{1}{2} f_{j-1} + \frac{1-\chi}{2} b_j + \frac{\chi}{2} f_j, \qquad b_j = 1 + \frac{1}{2} b_{j+1} + \frac{1-\chi}{2} f_j + \frac{\chi}{2} b_j.$$

Solving these recurrences with $f_0 = 0$ and $b_J = f_J = 1 + f_{J-1}$ yields Equation (C.97).

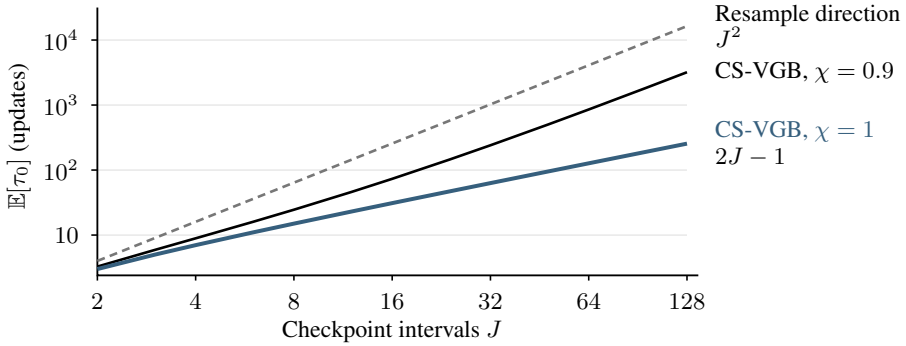

Figure C.2: Exact expected first hitting times from $j = J$ under balanced masses (Equation (C.97)). Resample direction moves directly to $j-1$ or $j+1$ with probabilities p_j^{F} or p_j^{B}, respectively; here both equal $1/2$ at interior checkpoints, and the move from J is forced inward. Its mean is J^2, compared with $2J-1$ for CS-VGB at $\chi = 1$.

C.5 Stationarity and Convergence under Edge-Correction Error

We now study CS-VGB with the edge correction omitted, establishing conditions for a stationary distribution, bounding its deviation from the exact stationary law, and analyzing convergence to this potentially biased limit.

Average edge-correction error. Let P be the exact CS-VGB kernel and $\bar{P}$ the generally biased kernel obtained by replacing the edge correction $\rho_j = dJ_j^{\mathrm{F}}/dJ_j^{\mathrm{B}}$ with 1, keeping $\widehat{V}$ and χ fixed. Assume the reward and verifier bounds in Equation (C.69), and measure the correction error by

$$\varepsilon_{\rho,\mathrm{av}} := \max_{1\le j\le J} \Big(\mathbb{E}_{J_j^{\mathrm{B}}}[(\rho_j - 1)^2]\Big)^{1/2} < \infty. \tag{C.98}$$

To turn this average edge error into an operator bound, let P_{move} and $\bar{P}_{\mathrm{move}}$ retain only checkpoint moves, with their original probabilities. Assume they have densities with respect to the exact stationary law Π:

$$P_{\mathrm{move}}(z, dz') = p_{\mathrm{move}}(z, z')\Pi(dz'), \qquad \bar{P}_{\mathrm{move}}(z, dz') = \bar{p}_{\mathrm{move}}(z, z')\Pi(dz'), \tag{C.99}$$

for Π-almost every z. For at least one finite $p > 2$, suppose

$$H_{\rho,p} := \left(\iint [p_{\mathrm{move}}(z, z') + \bar{p}_{\mathrm{move}}(z, z')]^p \, \Pi(dz)\Pi(dz') \right)^{1/p} < \infty. \tag{C.100}$$

Combining this density moment with the average correction error $\varepsilon_{\rho,\mathrm{av}}$, define

$$\Delta_{\rho,\mathrm{av}} := \inf_{\substack{2<p<\infty \\ H_{\rho,p}<\infty}} H_{\rho,p}^{\,p/[2(p-1)]} \left(\frac{\varepsilon_{\rho,\mathrm{av}}}{\sqrt{\min\{r_{\min}, v_{\min}\}}} \right)^{(p-2)/(p-1)}. \tag{C.101}$$

where $r_{\min}$ and $v_{\min}$ are the positive lower bounds on the reward and process verifier, respectively. Here, $\| \cdot \|_{2\to 2}$ and all adjoints are taken in $L^2(\Pi)$.

Lemma C.6 (Average edge error controls the movement operator)**.** *Under Equations (C.69), (C.98) and (C.100), $\bar{P}$ is bounded on $L^2(\Pi)$ and*

$$\|\bar{P}_{\mathrm{move}} - P_{\mathrm{move}}\|_{2\to 2} \le \|\bar{p}_{\mathrm{move}} - p_{\mathrm{move}}\|_{L^2(\Pi\otimes\Pi)} \le \Delta_{\rho,\mathrm{av}}. \tag{C.102}$$

The same bound holds for $\|\bar{M} - M\|_{2\to 2}$, where $M = (P_{\mathrm{move}} + P_{\mathrm{move}}^*)/2$ and $\bar{M} = (\bar{P}_{\mathrm{move}} + \bar{P}_{\mathrm{move}}^*)/2$.

Proof. Write $b_* := \min\{r_{\min}, v_{\min}\} > 0$ and define the uncorrected weighted edge measures

$$\widehat{C}_1^{\rho=1}(dx, dy) := r(y) J_1^{\mathrm{B}}(dx, dy), \quad \widehat{C}_j^{\rho=1}(dx, dy) := \sqrt{\widehat{V}_j(x)\widehat{V}_{j-1}(y)}\, J_j^{\mathrm{B}}(dx, dy), \ 2 \le j \le J. \tag{C.103}$$

Then $\widehat{C}_j = \rho_j \widehat{C}_j^{\rho=1}$ and $\widehat{C}_j^{\rho=1} \le J_j^{\mathrm{B}}$. Writing $\mathcal{E}_{\rho,j} := \int (\rho_j - 1)^2 \, d\widehat{C}_j^{\rho=1}$, the reward and verifier bounds give

$$\mathcal{E}_{\rho,j} \le \varepsilon_{\rho,\mathrm{av}}^2, \qquad \mathcal{Z} = 3\widehat{C}_1(\mathcal{X}, \mathcal{X}) + 4\sum_{j=2}^{J-1} \widehat{C}_j(\mathcal{X}, \mathcal{X}) + 3\widehat{C}_J(\mathcal{X}, \mathcal{X}) \ge (4J - 2)b_*.$$

Hence

$$\frac{\mathcal{E}_{\rho,1} + 4\sum_{j=2}^{J} \mathcal{E}_{\rho,j}}{\mathcal{Z}} \le \frac{4J - 3}{(4J - 2)b_*} \, \varepsilon_{\rho,\mathrm{av}}^2 \le \frac{\varepsilon_{\rho,\mathrm{av}}^2}{b_*}. \tag{C.104}$$

For probability measures R, S, define the *triangular discrimination* and *squared Hellinger distance*, respectively, by

$$\mathcal{T}(R, S) := \int \frac{(dR - dS)^2}{dR + dS}, \qquad \mathcal{H}^2(R, S) := \int (\sqrt{dR} - \sqrt{dS})^2,$$

For densities u, v with respect to a common dominating measure, $(u - v)^2/(u + v) = (\sqrt{u} - \sqrt{v})^2(\sqrt{u} + \sqrt{v})^2/(u + v) \le 2(\sqrt{u} - \sqrt{v})^2$, since $(\sqrt{u} + \sqrt{v})^2 \le 2(u + v)$. Integrating gives $\mathcal{T}(R, S) \le 2\mathcal{H}^2(R, S)$. For finite nonzero measures A, B with total masses a, b, respectively, let $c_0 := \int \sqrt{d(A/a)d(B/b)} \in [0, 1]$. Then

$$a\,\mathcal{T}(A/a, B/b) \overset{(\circ)}{\le} 2a\mathcal{H}^2(A/a, B/b) = 4a(1 - c_0) \overset{(\triangle)}{\le} 4\int (\sqrt{dA} - \sqrt{dB})^2, \tag{C.105}$$

where $(\circ)$ uses $\mathcal{T} \le 2\mathcal{H}^2$, the equality follows from $\mathcal{H}^2(A/a, B/b) = 2(1 - c_0)$, and $(\triangle)$ uses $\int (\sqrt{dA} - \sqrt{dB})^2 = a + b - 2\sqrt{ab}c_0 = (\sqrt{b} - \sqrt{a}\, c_0)^2 + a(1 - c_0^2) \ge a(1 - c_0)$.

Fix an interior source x at checkpoint j. On disjoint destination copies labelled F and B, define

$$A_{j,x}^{\rho}(dD, dy) := a_j^{\mathrm{F}}(x)\delta_{\mathrm{F}}(dD)Q_j^{\mathrm{F}}(dy \mid x) + \sqrt{\widehat{V}_j(x)}\,\delta_{\mathrm{B}}(dD)K_j^{\mathrm{B}}(dy \mid x)\sqrt{\widehat{V}_{j+1}(y)}\,\rho_{j+1}(y, x),$$

$$A_{j,x}^{\rho=1}(dD, dy) := a_j^{\mathrm{F}}(x)\delta_{\mathrm{F}}(dD)Q_j^{\mathrm{F}}(dy \mid x) + \sqrt{\widehat{V}_j(x)}\,\delta_{\mathrm{B}}(dD)K_j^{\mathrm{B}}(dy \mid x)\sqrt{\widehat{V}_{j+1}(y)}.$$

For this fixed (j, x), write $a := A_{j,x}^{\rho}(\{\mathrm{F}, \mathrm{B}\} \times \mathcal{X}) = a_j^{\mathrm{F}}(x) + a_j^{\mathrm{B}}(x)$ for the total mass of $A_{j,x}^{\rho}$ and $b := A_{j,x}^{\rho=1}(\{\mathrm{F}, \mathrm{B}\} \times \mathcal{X})$. Dividing by these total masses gives the probabilities of moving in each

current direction:

$$P_{\mathrm{move}}((j,x,\mathrm{F}),(j-1,dy,\mathrm{F})) = \frac{A^{\rho}_{j,x}(\mathrm{F},dy)}{a},$$
$$P_{\mathrm{move}}((j,x,\mathrm{B}),(j+1,dy,\mathrm{B})) = \frac{A^{\rho}_{j,x}(\mathrm{B},dy)}{a}.$$

The corresponding probabilities for $\bar{P}$ are obtained by replacing $A^{\rho}_{j,x}/a$ with $A^{\rho=1}_{j,x}/b$. Moreover, each current direction has stationary measure

$$\Pi(j,dx,\mathrm{F}) = \Pi(j,dx,\mathrm{B}) = \frac{a\,\pi_j(dx)}{\mathcal{Z}}.$$

For each fixed (j,x), set $R := A^{\rho}_{j,x}/a$ and $S := A^{\rho=1}_{j,x}/b$. Splitting the defining integral of $\mathcal{T}(R,S)$ over the F and B destinations gives

$$\begin{aligned}\mathcal{T}(R,S) &= \sum_{d\in\{\mathrm{F},\mathrm{B}\}}\int_{\mathcal{X}} \frac{[dR(d,y)-dS(d,y)]^2}{dR(d,y)+dS(d,y)} \\ &= \sum_{d\in\{\mathrm{F},\mathrm{B}\}}\int_{\mathsf{S}} \frac{[\bar{p}_{\mathrm{move}}((j,x,d),z')-p_{\mathrm{move}}((j,x,d),z')]^2}{\bar{p}_{\mathrm{move}}((j,x,d),z')+p_{\mathrm{move}}((j,x,d),z')}\,\Pi(dz'),\end{aligned}$$

where the second equality uses the moving densities p_{move} and $\bar{p}_{\mathrm{move}}$ with respect to the common dominating measure Π, with $0/0 := 0$. Multiplying by $\Pi(j,dx,d) = a\,\pi_j(dx)/\mathcal{Z}$, which is the same for both directions, and integrating over x therefore gives

$$\begin{aligned}&\sum_{d\in\{\mathrm{F},\mathrm{B}\}}\int_{\mathcal{X}}\Pi(j,dx,d)\int_{\mathsf{S}}\frac{[\bar{p}_{\mathrm{move}}((j,x,d),z')-p_{\mathrm{move}}((j,x,d),z')]^2}{\bar{p}_{\mathrm{move}}((j,x,d),z')+p_{\mathrm{move}}((j,x,d),z')}\,\Pi(dz') \\ &\quad= \frac{1}{\mathcal{Z}}\int_{\mathcal{X}} a\,\mathcal{T}(R,S)\,\pi_j(dx) = \frac{1}{\mathcal{Z}}\int_{\mathcal{X}} a\,\mathcal{T}(A^{\rho}_{j,x}/a, A^{\rho=1}_{j,x}/b)\,\pi_j(dx) \\ &\quad\le \frac{4}{\mathcal{Z}}\int_{\mathcal{X}}\pi_j(dx)\int_{\{\mathrm{F},\mathrm{B}\}\times\mathcal{X}}\left(\sqrt{dA^{\rho}_{j,x}}-\sqrt{dA^{\rho=1}_{j,x}}\right)^2,\end{aligned}$$

where the inequality is Equation (C.105). The F terms agree, so only the B terms remain in the last integral. Since $(\sqrt{\rho}-1)^2 \le (\rho-1)^2$,

$$\begin{aligned}&\frac{4}{\mathcal{Z}}\int_{\mathcal{X}}\pi_j(dx)\int_{\{\mathrm{F},\mathrm{B}\}\times\mathcal{X}}\left(\sqrt{dA^{\rho}_{j,x}}-\sqrt{dA^{\rho=1}_{j,x}}\right)^2 \\ &\quad= \frac{4}{\mathcal{Z}}\iint\sqrt{\widehat{V}_j(x)\widehat{V}_{j+1}(y)}\left(\sqrt{\rho_{j+1}(y,x)}-1\right)^2\pi_j(dx)K^{\mathrm{B}}_j(dy\mid x) \\ &\quad\le \frac{4}{\mathcal{Z}}\iint[\rho_{j+1}(y,x)-1]^2\,\widehat{C}^{\rho=1}_{j+1}(dy,dx) = \frac{4\mathcal{E}_{\rho,j+1}}{\mathcal{Z}}.\end{aligned}$$

At the clean boundary, fix a clean sample y and compare the two renoising distributions $R := K^{\mathrm{B},\rho}_0(\cdot\mid y)$ and $S := K^{\mathrm{B}}_0(\cdot\mid y)$. Both are probability measures, and $dR/dS = \rho_1(\cdot,y)$ gives

$$\mathcal{T}(R,S) = \int_{\mathcal{X}}\frac{[\rho_1(x,y)-1]^2}{\rho_1(x,y)+1}\,K^{\mathrm{B}}_0(dx\mid y) \le \int_{\mathcal{X}}[\rho_1(x,y)-1]^2\,K^{\mathrm{B}}_0(dx\mid y).$$

Integrating against $\Pi(0,dy,\mathrm{B}) = r(y)\pi_0(dy)/\mathcal{Z}$ gives

$$\begin{aligned}\int\mathcal{T}(R,S)\,\Pi(0,dy,\mathrm{B}) &\le \frac{1}{\mathcal{Z}}\iint[\rho_1(x,y)-1]^2 r(y)\pi_0(dy)K^{\mathrm{B}}_0(dx\mid y) \\ &= \frac{1}{\mathcal{Z}}\int(\rho_1-1)^2\,d\widehat{C}^{\rho=1}_1 = \frac{\mathcal{E}_{\rho,1}}{\mathcal{Z}}.\end{aligned}$$

At checkpoint J, both kernels force an F move with the same transition Q^{F}_J, so this boundary contributes no error. Consequently,

$$\iint\frac{[\bar{p}_{\mathrm{move}}(z,z')-p_{\mathrm{move}}(z,z')]^2}{\bar{p}_{\mathrm{move}}(z,z')+p_{\mathrm{move}}(z,z')}\,\Pi(dz)\Pi(dz') \le \frac{\mathcal{E}_{\rho,1}+4\sum_{j=2}^{J}\mathcal{E}_{\rho,j}}{\mathcal{Z}} \le \frac{\varepsilon^2_{\rho,\mathrm{av}}}{b_*}. \tag{C.106}$$

Since $|\bar{p}_{\mathrm{move}} - p_{\mathrm{move}}| \le \bar{p}_{\mathrm{move}} + p_{\mathrm{move}}$, Hölder's inequality gives, for $p > 2$,

$$\begin{aligned}
&\iint |\bar{p}_{\mathrm{move}} - p_{\mathrm{move}}|^2 \, d(\Pi \otimes \Pi) \\
&\le \iint \left(\frac{|\bar{p}_{\mathrm{move}} - p_{\mathrm{move}}|^2}{\bar{p}_{\mathrm{move}} + p_{\mathrm{move}}} \right)^{\frac{p-2}{p-1}} (\bar{p}_{\mathrm{move}} + p_{\mathrm{move}})^{\frac{p}{p-1}} \, d(\Pi \otimes \Pi) \\
&\le \left(\iint \frac{|\bar{p}_{\mathrm{move}} - p_{\mathrm{move}}|^2}{\bar{p}_{\mathrm{move}} + p_{\mathrm{move}}} \, d(\Pi \otimes \Pi) \right)^{\frac{p-2}{p-1}} \left(\iint (\bar{p}_{\mathrm{move}} + p_{\mathrm{move}})^p \, d(\Pi \otimes \Pi) \right)^{\frac{1}{p-1}}.
\end{aligned}$$

Taking square roots proves $\|\bar{p}_{\mathrm{move}} - p_{\mathrm{move}}\|_2 \le H_{\rho,p}^{p/[2(p-1)]} (\varepsilon_{\rho,\mathrm{av}} / \sqrt{b_*})^{(p-2)/(p-1)}$. Applying Cauchy–Schwarz to the integral defining $(\bar{P}_{\mathrm{move}} - P_{\mathrm{move}})f$ gives $\|(\bar{P}_{\mathrm{move}} - P_{\mathrm{move}})f\|_{2,\Pi} \le \|\bar{p}_{\mathrm{move}} - p_{\mathrm{move}}\|_{L^2(\Pi\otimes\Pi)} \|f\|_{2,\Pi}$, which bounds the operator norm by the same L^2 norm. Taking the infimum over admissible p gives Equation (C.102). Since $\|A^*\|_{2\to 2} = \|A\|_{2\to 2}$, symmetrization gives $\|\bar{M} - M\|_{2\to 2} \le \frac{1}{2}(\|\bar{P}_{\mathrm{move}} - P_{\mathrm{move}}\|_{2\to 2} + \|(\bar{P}_{\mathrm{move}} - P_{\mathrm{move}})^*\|_{2\to 2}) = \|\bar{P}_{\mathrm{move}} - P_{\mathrm{move}}\|_{2\to 2} \le \Delta_{\rho,\mathrm{av}}$, so the same bound holds for $\bar{M} - M$.

Finally, write $\bar{P} = \bar{P}_{\mathrm{move}} + \bar{S}$, where $\bar{S}$ contains flips and stays. At each interior (j, x), it acts on the two directional values of f through

$$\bar{S}_j(x) = \begin{pmatrix} \chi \bar{\ell}_j(x) & \bar{p}_j^{\mathrm{B}}(x) - \chi \bar{\ell}_j(x) \\ \bar{p}_j^{\mathrm{F}}(x) - \chi \bar{\ell}_j(x) & \chi \bar{\ell}_j(x) \end{pmatrix}, \qquad \bar{\ell}_j(x) := \min\{\bar{p}_j^{\mathrm{F}}(x), \bar{p}_j^{\mathrm{B}}(x)\}, \tag{C.107}$$

where $\bar{p}_j^d$ denotes the directional probability for $\bar{P}$. All entries are nonnegative, and the row sums are $(\bar{p}_j^{\mathrm{B}}, \bar{p}_j^{\mathrm{F}})$, while the column sums are $(\bar{p}_j^{\mathrm{F}}, \bar{p}_j^{\mathrm{B}})$. Thus the Euclidean operator norm satisfies $\|\bar{S}_j(x)\|_2 \le \sqrt{\|\bar{S}_j(x)\|_1 \|\bar{S}_j(x)\|_\infty} \le 1$. Since $\Pi(j, dx, \mathrm{F}) = \Pi(j, dx, \mathrm{B})$ and $\bar{S}$ vanishes at the boundaries, integrating this inequality gives $\|\bar{S}f\|_{2,\Pi} \le \|f\|_{2,\Pi}$. Consequently, $\|\bar{P}\|_{2\to 2} \le \|\bar{P}_{\mathrm{move}}\|_{2\to 2} + 1 < \infty$, which proves the boundedness of $\bar{P}$. □

Stationary bias under average edge error. We now use the movement error bound in Lemma C.6 to control the stationary bias caused by omitting ρ_j. To do so, we compare $\Delta_{\rho,\mathrm{av}}$ with the lower bound on the exact chain's movement form established in the proof of Theorem C.3:

$$\langle f, (cI - M)f \rangle_\Pi \ge \gamma_{\mathrm{CS}} \operatorname{Var}_\Pi(f), \qquad c(j, x, d) := \begin{cases} 1, & j \in \{0, J\}, \\ 1/2, & 1 \le j < J, \end{cases} \tag{C.108}$$

Here $M = (P_{\mathrm{move}} + P_{\mathrm{move}}^*)/2$, as above, and we take $0 < \gamma_{\mathrm{CS}} \le 1/2$ by reducing the lower bound if necessary. Under their respective assumptions, Corollary C.3 and Theorem C.4 provide alternative lower bounds for γ_{CS} in Equation (C.108). The result below shows that if the movement error is smaller than this lower bound, the uncorrected chain has a stationary law close to Π.

Theorem C.5 (Stability under average omitted-edge error)**.** *Assume Equations (C.69), (C.98), (C.100) and (C.108) and $\Delta_{\rho,\mathrm{av}} < \gamma_{\mathrm{CS}}$. Then $\bar{P}$ has a unique invariant probability $\bar{\Pi} \ll \Pi$ among those with an $L^2(\Pi)$ density, and*

$$\|\bar{\Pi} - \Pi\|_{\mathrm{TV}} \le \frac{\Delta_{\rho,\mathrm{av}}}{\gamma_{\mathrm{CS}} - \Delta_{\rho,\mathrm{av}}}. \tag{C.109}$$

If $\Delta_{\rho,\mathrm{av}}/(\gamma_{\mathrm{CS}} - \Delta_{\rho,\mathrm{av}}) < \Pi(j = 0) = Z/\mathcal{Z}$, then the stationary clean conditional $\bar{\pi}_0$ is well-defined and

$$\|\bar{\pi}_0 - \pi_0^\star\|_{\mathrm{TV}} \le \frac{\Delta_{\rho,\mathrm{av}}}{\Pi(j = 0)\,(\gamma_{\mathrm{CS}} - \Delta_{\rho,\mathrm{av}})}. \tag{C.110}$$

These bias bounds hold for every $\chi \in [0, 1]$, with the same error budget and movement-gap constant.

Proof. For the proof, write $g_\rho := \gamma_{\mathrm{CS}} - \Delta_{\rho,\mathrm{av}} > 0$. Recall $\bar{P} = \bar{P}_{\mathrm{move}} + \bar{S}$ and the matrix in Equation (C.107). Since the F and B states at each (j, x) have equal Π weight, the adjoint $\bar{S}^*$ acts

there through $\bar{S}_j(x)^{\mathsf{T}}$. By the definition of the inner product,

$$\mathrm{Re}\langle f,(I-\bar{S})f\rangle_\Pi = \int_{\mathsf{S}} \mathrm{Re}\Big[\overline{f(z)}\big(f(z)-(\bar{S}f)(z)\big)\Big]\,\Pi(dz).$$

At an interior (j,x), recall $f_j^d(x)=f(j,x,d)$. With the argument x suppressed, the flip/stay probabilities give

$$(\bar{S}f)_j^{\mathrm{F}} = \chi\bar{\ell}_j f_j^{\mathrm{F}} + (\bar{p}_j^{\mathrm{B}} - \chi\bar{\ell}_j) f_j^{\mathrm{B}}, \qquad (\bar{S}f)_j^{\mathrm{B}} = (\bar{p}_j^{\mathrm{F}} - \chi\bar{\ell}_j) f_j^{\mathrm{F}} + \chi\bar{\ell}_j f_j^{\mathrm{B}}.$$

Substituting these expressions into the inner-product integrand and using $\bar{p}_j^{\mathrm{F}} + \bar{p}_j^{\mathrm{B}} = 1$ yields

$$\begin{aligned}\sum_{d\in\{\mathrm{F},\mathrm{B}\}} \mathrm{Re}\Big[\overline{f_j^d}\big(f_j^d - (\bar{S}f)_j^d\big)\Big] &= (1-\chi\bar{\ell}_j)(|f_j^{\mathrm{F}}|^2 + |f_j^{\mathrm{B}}|^2) - (1-2\chi\bar{\ell}_j)\,\mathrm{Re}(\overline{f_j^{\mathrm{F}}} f_j^{\mathrm{B}}) \\ &= \tfrac{1}{2}(|f_j^{\mathrm{F}}|^2 + |f_j^{\mathrm{B}}|^2) + (\tfrac{1}{2} - \chi\bar{\ell}_j)|f_j^{\mathrm{F}} - f_j^{\mathrm{B}}|^2.\end{aligned}$$

At the boundaries, $\bar{S}=0$ leaves $|f|^2$. Thus the first terms combine into $\langle f, cf\rangle_\Pi$, with $c=1/2$ in the interior and $c=1$ at the boundaries. Integrating the two interior directional terms against their common measure $\mu_j(dx)/\mathcal{Z}$ yields

$$\mathrm{Re}\langle f,(I-\bar{S})f\rangle_\Pi = \langle f,cf\rangle_\Pi + \sum_{j=1}^{J-1}\int_{\mathcal{X}} \big(\tfrac{1}{2} - \chi\bar{\ell}_j(x)\big)\,|f(j,x,\mathrm{F}) - f(j,x,\mathrm{B})|^2\,\frac{\mu_j(dx)}{\mathcal{Z}} \geq \langle f,cf\rangle_\Pi, \tag{C.111}$$

since $\chi\bar{\ell}_j \leq 1/2$. Thus Equations (C.102) and (C.108) give, for $\Pi(f)=0$,

$$\begin{aligned}\mathrm{Re}\langle f,(I-\bar{P})f\rangle_\Pi &\geq \langle f,cf\rangle_\Pi - \mathrm{Re}\langle f,\bar{P}_{\mathrm{move}}f\rangle_\Pi \\ &= \langle f,cf\rangle_\Pi - \left\langle f, \frac{\bar{P}_{\mathrm{move}}+\bar{P}_{\mathrm{move}}^*}{2} f\right\rangle_\Pi = \langle f,(cI-\bar{M})f\rangle_\Pi \\ &= \langle f,(cI-M)f\rangle_\Pi - \langle f,(\bar{M}-M)f\rangle_\Pi \\ &\geq (\gamma_{\mathrm{CS}} - \Delta_{\rho,\mathrm{av}})\|f\|_{2,\Pi}^2 = g_\rho\|f\|_{2,\Pi}^2.\end{aligned} \tag{C.112}$$

The same inequality holds with $\bar{P}^*$ in place of $\bar{P}$.

Also, $\bar{S}^*\mathbf{1} = \bar{P}_{\mathrm{move}}\mathbf{1} + \mathbf{1} - 2c$ and $M\mathbf{1}=c$, so

$$b_\rho := \bar{P}^*\mathbf{1} - \mathbf{1} = \bar{P}_{\mathrm{move}}^*\mathbf{1} + \bar{S}^*\mathbf{1} - \mathbf{1} = 2(\bar{M}-M)\mathbf{1}, \quad \Pi(b_\rho)=0, \quad \|b_\rho\|_{2,\Pi} \leq 2\Delta_{\rho,\mathrm{av}}. \tag{C.113}$$

To construct an invariant density, set

$$T_t := e^{t(\bar{P}-I)} = e^{-t}\sum_{n=0}^{\infty}\frac{t^n}{n!}\bar{P}^n.$$

Because T_t is positive and $T_t\mathbf{1}=\mathbf{1}$, T_t^* maps the closed set of nonnegative, unit-integral $L^2(\Pi)$ densities into itself, i.e.,

$$h\geq 0, \quad \Pi(h)=1, \quad h\in L^2(\Pi) \quad\Longrightarrow\quad T_t^*h \geq 0, \quad \Pi(T_t^*h)=1, \quad T_t^*h \in L^2(\Pi).$$

The set of these densities is complete as a closed subset of $L^2(\Pi)$. To show contraction, consider any $f\in L^2(\Pi)$ with $\Pi(f)=0$ and put $f_t := T_t^*f$. Then $\Pi(f_t)=0$, and Equation (C.112) gives

$$\frac{d}{dt}\|f_t\|_{2,\Pi}^2 = 2\,\mathrm{Re}\langle f_t,(\bar{P}^*-I)f_t\rangle_\Pi \leq -2g_\rho\|f_t\|_{2,\Pi}^2, \quad \|T_t^*f\|_{2,\Pi} = \|f_t\|_{2,\Pi} \leq e^{-g_\rho t}\|f\|_{2,\Pi}. \tag{C.114}$$

Applying this bound to the difference of two probability densities shows that, for fixed $t>0$, T_t^* is a strict contraction on the complete set of densities above. The contraction mapping theorem (Kreyszig, 1978, Section 5.1) therefore gives a unique fixed density $\bar{h}_t$. For every $s\geq 0$, commutativity gives

$$T_t^*(T_s^*\bar{h}_t) = T_s^*(T_t^*\bar{h}_t) = T_s^*\bar{h}_t.$$

Thus $T_s^*\bar{h}_t$ is also a fixed density of T_t^*, so uniqueness gives $T_s^*\bar{h}_t = \bar{h}_t$. For $s>0$, uniqueness for T_s^* then gives $\bar{h}_t = \bar{h}_s$. Denote this common density by $\bar{h}$. Differentiating $T_s^*\bar{h}=\bar{h}$ at $s=0$ yields $(\bar{P}^*-I)\bar{h}=0$, hence $\bar{P}^*\bar{h}=\bar{h}$.

Define $\bar{\Pi}(dz) := \bar{h}(z)\Pi(dz)$. Then $\bar{P}^*\bar{h} = \bar{h}$ implies that $\bar{\Pi}$ is stationary for $\bar{P}$, with $d\bar{\Pi}/d\Pi = \bar{h}$.

Set $f := \bar{h} - 1$. The invariant equation gives $(I - \bar{P}^*)f = b_\rho$. Since $\Pi(f) = 0$, applying Equation (C.112) and then Cauchy–Schwarz with Equation (C.113) gives

$$g_\rho \|f\|_{2,\Pi}^2 \le \mathrm{Re}\langle f, b_\rho\rangle_\Pi \le 2\Delta_{\rho,\mathrm{av}}\|f\|_{2,\Pi}.$$

Consequently,

$$\|\bar{h} - 1\|_{2,\Pi} \le \frac{2\Delta_{\rho,\mathrm{av}}}{g_\rho}, \qquad \|\bar{\Pi} - \Pi\|_{\mathrm{TV}} = \frac{1}{2}\Pi(|\bar{h} - 1|) \le \frac{1}{2}\|\bar{h} - 1\|_{2,\Pi} \le \frac{\Delta_{\rho,\mathrm{av}}}{g_\rho}.$$

The condition $\Delta_{\rho,\mathrm{av}}/g_\rho < \Pi(j = 0)$ ensures $\bar{\Pi}(j = 0) \ge \Pi(j = 0) - \Delta_{\rho,\mathrm{av}}/g_\rho > 0$, and Equation (C.26) yields

$$\|\bar{\pi}_0 - \pi_0^\star\|_{\mathrm{TV}} \le \frac{\|\bar{\Pi} - \Pi\|_{\mathrm{TV}}}{\Pi(j = 0)} \le \frac{\Delta_{\rho,\mathrm{av}}}{\Pi(j = 0)\, g_\rho}.$$

□

Convergence under edge-correction omission. Following the lazification in Equation (C.18), we denote by $\bar{P}_{\mathrm{L}} := (I + \bar{P})/2$ the $1/2$-lazy version of $\bar{P}$. We now show geometric convergence to $\bar{\Pi}$ and combine it with the stationary-bias bound to control the total error relative to Π.

Corollary C.4 (Exponential convergence under average edge error)**.** *Under the assumptions of Theorem C.5, there is a finite constant $C_{\rho,\chi}$ such that, for every initial law $\nu \ll \Pi$ with $d\nu/d\Pi \in L^2(\Pi)$ and $n \ge 0$,*

$$\|\nu\bar{P}_{\mathrm{L}}^n - \bar{\Pi}\|_{\mathrm{TV}} \le \frac{C_{\rho,\chi}}{2} e^{-(\gamma_{\mathrm{CS}} - \Delta_{\rho,\mathrm{av}})n/8} \left\| \frac{d\nu}{d\Pi} - \frac{d\bar{\Pi}}{d\Pi} \right\|_{2,\Pi}. \tag{C.115}$$

In particular, for $D_\nu := \|d\nu/d\Pi - 1\|_{2,\Pi} < \infty$,

$$\|\nu\bar{P}_{\mathrm{L}}^n - \Pi\|_{\mathrm{TV}} \le \underbrace{\frac{C_{\rho,\chi}}{2} e^{-(\gamma_{\mathrm{CS}} - \Delta_{\rho,\mathrm{av}})n/8} \left(D_\nu + \frac{2\Delta_{\rho,\mathrm{av}}}{\gamma_{\mathrm{CS}} - \Delta_{\rho,\mathrm{av}}} \right) + \frac{\Delta_{\rho,\mathrm{av}}}{\gamma_{\mathrm{CS}} - \Delta_{\rho,\mathrm{av}}}}_{(\star)}. \tag{C.116}$$

If $(\star) < \Pi(j = 0)$, the clean conditional of $\nu\bar{P}_{\mathrm{L}}^n$ is well-defined, and Equation (C.26) bounds its total-variation distance from $\pi_0^\star$ by $(\star)/\Pi(j = 0)$.

Proof. For the proof, write $\bar{h} := d\bar{\Pi}/d\Pi$ and $g_\rho := \gamma_{\mathrm{CS}} - \Delta_{\rho,\mathrm{av}} > 0$, and recall $\bar{P} = \bar{P}_{\mathrm{move}} + \bar{S}$, where $\bar{S}$ is the flip/stay operator from Equation (C.107), vanishing at the boundaries. We first bound the spectrum of $\bar{P}$, then obtain a strict bound for the lazy kernel on mean-zero functions. Since $\Pi \otimes \Pi$ is a probability measure and $p > 2$,

$$\|\bar{p}_{\mathrm{move}}\|_{L^2(\Pi\otimes\Pi)} \le \|\bar{p}_{\mathrm{move}}\|_{L^p(\Pi\otimes\Pi)} < \infty.$$

To see compactness, approximate this density in $L^2(\Pi \otimes \Pi)$ by finite sums $\sum_{i=1}^N u_i(z)v_i(z')$ with real-valued u_i, v_i. The corresponding operators A_N act as

$$(A_N f)(z) := \int \left[\sum_{i=1}^N u_i(z)v_i(z') \right] f(z')\,\Pi(dz') = \sum_{i=1}^N u_i(z) \int v_i(z')f(z')\,\Pi(dz').$$

Their range lies in $\mathrm{span}\{u_1, \dots, u_N\}$ and is therefore finite-dimensional. Cauchy–Schwarz gives

$$\left\|\bar{P}_{\mathrm{move}} - A_N\right\|_{2\to 2} \le \left\| \bar{p}_{\mathrm{move}} - \sum_{i=1}^N u_i \otimes v_i \right\|_{L^2(\Pi\otimes\Pi)} \longrightarrow 0. \tag{C.117}$$

Thus $\bar{P}_{\mathrm{move}}$ is compact as an operator-norm limit of finite-rank operators, and so is $\bar{P}_{\mathrm{move}}^*$: its approximants satisfy $(A_N^* f)(z) = \sum_{i=1}^N v_i(z) \int u_i(z')f(z')\,\Pi(dz')$ and hence have finite rank, while $\|\bar{P}_{\mathrm{move}}^* - A_N^*\|_{2\to 2} = \|\bar{P}_{\mathrm{move}} - A_N\|_{2\to 2} \to 0$ by Equation (C.117).

We use the following consequence of the Fredholm alternative for compact operators (Kreyszig, 1978, Theorem 8.7-2). If $A = B + K$, with K compact, then every spectral value λ of A with $|\lambda| > \|B\|$ is an eigenvalue. Indeed, $\lambda I - B$ is invertible by its Neumann series, and

$$\lambda I - A = (\lambda I - B)\left[I - (\lambda I - B)^{-1}K\right].$$

The operator $(\lambda I - B)^{-1}K$ is compact. By the Fredholm alternative, noninvertibility of the bracket implies a nonzero vector in its kernel, hence in the kernel of $\lambda I - A$.

The proof of Theorem C.5 gives $T_t^*\bar{h} = \bar{h}$ and the exponential contraction on mean-zero functions in Equation (C.114). For any $f \in L^2(\Pi)$, write

$$f = \Pi(f)\bar{h} + \big(f - \Pi(f)\bar{h}\big), \qquad \Pi\big(f - \Pi(f)\bar{h}\big) = 0.$$

Since $T_t^*\bar{h} = \bar{h}$, applying the contraction bound to $f - \Pi(f)\bar{h}$ gives

$$\|T_t^* f\|_{2,\Pi} \le |\Pi(f)|\|\bar{h}\|_{2,\Pi} + e^{-g_\rho t}\|f - \Pi(f)\bar{h}\|_{2,\Pi} \le \|f\|_{2,\Pi} + 2|\Pi(f)|\|\bar{h}\|_{2,\Pi} \le (1 + 2\|\bar{h}\|_{2,\Pi})\|f\|_{2,\Pi},$$

where we used $e^{-g_\rho t} \le 1$, the triangle inequality, and $|\Pi(f)| \le \|f\|_{2,\Pi}$. Taking the supremum over $\|f\|_{2,\Pi} = 1$ and $t \ge 0$ yields

$$\sup_{t \ge 0} \|T_t\|_{2\to 2} = \sup_{t \ge 0} \|T_t^*\|_{2\to 2} \le 1 + 2\|\bar{h}\|_{2,\Pi} < \infty. \tag{C.118}$$

Suppose $\bar{P}f = \lambda f$ for $f \neq 0$ and $|\lambda| > 1$. Positivity of $\bar{P}$ and the triangle inequality imply

$$\bar{P}|f| \ge |\bar{P}f| = |\lambda||f|, \qquad \bar{P}^n|f| \ge |\lambda|^n|f|.$$

Recalling $T_t = e^{-t}\sum_{n=0}^{\infty}(t^n/n!)\bar{P}^n$, we obtain

$$T_t|f| = e^{-t}\sum_{n=0}^{\infty}\frac{t^n}{n!}\bar{P}^n|f| \ge e^{-t}\sum_{n=0}^{\infty}\frac{(t|\lambda|)^n}{n!}|f| = e^{(|\lambda|-1)t}|f|.$$

Consequently $\|T_t\|_{2\to 2} \ge e^{(|\lambda|-1)t}$, contradicting the uniform bound in Equation (C.118) as $t \to \infty$. Since $\bar{P} = \bar{S} + \bar{P}_{\mathrm{move}}$, with $\|\bar{S}\|_{2\to 2} \le 1$ and $\bar{P}_{\mathrm{move}}$ compact, the Fredholm argument above excludes all spectral values of $\bar{P}$ outside the unit disk:

$$\sigma(\bar{P}) \subseteq \{\lambda \in \mathbb{C} : |\lambda| \le 1\}, \qquad \bar{P}\mathbf{1} = \mathbf{1} \implies 1 \in \sigma(\bar{P}), \qquad r(\bar{P}) := \sup_{\lambda \in \sigma(\bar{P})}|\lambda| = 1 = r(\bar{P}^*).$$

We next bound the part of the lazy kernel that does not move. By Equation (C.111) and $c \ge 1/2$,

$$\mathrm{Re}\langle f, \bar{S}f\rangle_\Pi \le \|f\|_{2,\Pi}^2 - \langle f, cf\rangle_\Pi \le \tfrac{1}{2}\|f\|_{2,\Pi}^2.$$

Together with $\|\bar{S}\|_{2\to 2} \le 1$, this yields

$$\left\|\frac{I + \bar{S}}{2}f\right\|_{2,\Pi}^2 = \frac{\|f\|_{2,\Pi}^2 + 2\,\mathrm{Re}\langle f, \bar{S}f\rangle_\Pi + \|\bar{S}f\|_{2,\Pi}^2}{4} \le \frac{3}{4}\|f\|_{2,\Pi}^2.$$

Let $L_0^2(\Pi) := \{f \in L^2(\Pi) : \Pi(f) = 0\}$. This subspace is invariant under $\bar{P}_{\mathrm{L}}^*$ because

$$\Pi(\bar{P}_{\mathrm{L}}^* f) = \langle \bar{P}_{\mathrm{L}}\mathbf{1}, f\rangle_\Pi = \Pi(f).$$

For its orthogonal projection $Qf := f - \Pi(f)\mathbf{1}$, $\|Q\|_{2\to 2} \le 1$, and hence on $L_0^2(\Pi)$,

$$\bar{P}_{\mathrm{L}}^*\big|_{L_0^2(\Pi)} = Q\frac{I + \bar{S}^*}{2}\bigg|_{L_0^2(\Pi)} + \frac{1}{2}Q\bar{P}_{\mathrm{move}}^*\bigg|_{L_0^2(\Pi)}.$$

The first term has norm at most $\sqrt{3}/2$ and the second is compact. The same Fredholm argument therefore shows that every spectral value of this restriction outside $|\lambda| \le \sqrt{3}/2$ is an eigenvalue. For such an eigenvalue, choose $0 \neq f \in L_0^2(\Pi)$ with $\bar{P}_{\mathrm{L}}^* f = \lambda f$. Then

$$\bar{P}^* f = (2\lambda - 1)f, \qquad \mu := 2\lambda - 1, \qquad |\mu| \le r(\bar{P}^*) = 1.$$

Moreover, Equation (C.112) gives

$$g_\rho\|f\|_{2,\Pi}^2 \le \mathrm{Re}\langle f, (I - \bar{P}^*)f\rangle_\Pi = (1 - \mathrm{Re}\,\mu)\|f\|_{2,\Pi}^2, \qquad \mathrm{Re}\,\mu \le 1 - g_\rho.$$

Thus

$$|\lambda|^2 = \frac{1 + 2\,\mathrm{Re}\,\mu + |\mu|^2}{4} \le \frac{1 + 2(1 - g_\rho) + 1}{4} = 1 - \frac{g_\rho}{2}.$$

Since $0 < g_\rho \le 1/2$, we also have $3/4 \le 1 - g_\rho/2$. This bounds the entire spectrum of the restriction, giving

$$r\left(\bar{P}_{\mathrm{L}}^*\big|_{L_0^2(\Pi)}\right) \le \sqrt{1 - g_\rho/2} \le e^{-g_\rho/4} < e^{-g_\rho/8}.$$

By Lemma C.6, $\bar{P}$ is bounded on $L^2(\Pi)$, so $\bar{P}_{\mathrm{L}}^*\big|_{L_0^2(\Pi)}$ is a bounded linear operator on $L_0^2(\Pi)$. The spectral-radius formula (Kreyszig, 1978, Theorem 7.5-5) therefore gives

$$\lim_{n\to\infty} \left\|(\bar{P}_{\mathrm{L}}^*)^n\big|_{L_0^2(\Pi)}\right\|_{2\to 2}^{1/n} = r\left(\bar{P}_{\mathrm{L}}^*\big|_{L_0^2(\Pi)}\right) < e^{-g_\rho/8}.$$

Hence the norm is at most $e^{-g_\rho n/8}$ for all sufficiently large n. Enlarging a finite constant $C_{\rho,\chi}$ to cover the remaining values of n gives

$$\left\|(\bar{P}_{\mathrm{L}}^*)^n f\right\|_{2,\Pi} \le C_{\rho,\chi} e^{-g_\rho n/8} \|f\|_{2,\Pi}, \qquad f \in L_0^2(\Pi), \quad n \ge 0.$$

Finally, $d\nu/d\Pi - \bar{h}$ has zero Π mean and $\bar{P}_{\mathrm{L}}^* \bar{h} = \bar{h}$. Therefore

$$\begin{aligned}\|\nu\bar{P}_{\mathrm{L}}^n - \bar{\Pi}\|_{\mathrm{TV}} &= \frac{1}{2}\Pi\left(\left|(\bar{P}_{\mathrm{L}}^*)^n\left(\frac{d\nu}{d\Pi} - \bar{h}\right)\right|\right) \\ &\le \frac{1}{2}\left\|(\bar{P}_{\mathrm{L}}^*)^n\left(\frac{d\nu}{d\Pi} - \bar{h}\right)\right\|_{2,\Pi} \le \frac{C_{\rho,\chi}}{2} e^{-g_\rho n/8}\left\|\frac{d\nu}{d\Pi} - \bar{h}\right\|_{2,\Pi}.\end{aligned}$$

The proof of Theorem C.5, together with the triangle inequality, also gives

$$\left\|\frac{d\nu}{d\Pi} - \bar{h}\right\|_{2,\Pi} \le D_\nu + \|\bar{h} - 1\|_{2,\Pi} \le D_\nu + \frac{2\Delta_{\rho,\mathrm{av}}}{g_\rho}.$$

Combining this with $\|\nu\bar{P}_{\mathrm{L}}^n - \Pi\|_{\mathrm{TV}} \le \|\nu\bar{P}_{\mathrm{L}}^n - \bar{\Pi}\|_{\mathrm{TV}} + \|\bar{\Pi} - \Pi\|_{\mathrm{TV}}$ and Equation (C.109) proves Equation (C.116). The clean-conditional bound follows from Equation (C.26). □

Finally, we remark that the following condition on the reference edges J_j^{F} and J_j^{B} ensures $H_{\rho,p} < \infty$ (Equation (C.100)) without requiring $J_j^{\mathrm{F}} = J_j^{\mathrm{B}}$.

Lemma C.7 (Sufficient conditions for finite movement-density moments)**.** *Assume Equation (C.69). If $J_j^{\mathrm{F}}, J_j^{\mathrm{B}} \ll \pi_j \otimes \pi_{j-1}$ and, for some finite $p > 2$,*

$$\max_{1\le j\le J}\left\{\left\|\frac{dJ_j^{\mathrm{F}}}{d(\pi_j\otimes\pi_{j-1})}\right\|_{L^p(\pi_j\otimes\pi_{j-1})}, \left\|\frac{dJ_j^{\mathrm{B}}}{d(\pi_j\otimes\pi_{j-1})}\right\|_{L^p(\pi_j\otimes\pi_{j-1})}\right\} < \infty, \tag{C.119}$$

then the movement densities in Equation (C.99) exist and $H_{\rho,p} < \infty$.

Proof. Set $b_* := \min\{r_{\min}, v_{\min}\}$. Applying Equation (C.69) to the edge definitions in Equations (B.8) and (B.9) gives

$$\widehat{C}_j(dx,dy) = \sqrt{\widehat{V}_j(x)\widehat{V}_{j-1}(y)}\, J_j^{\mathrm{F}}(dx,dy) \quad\Longrightarrow\quad b_* J_j^{\mathrm{F}} \le \widehat{C}_j \le J_j^{\mathrm{F}}, \qquad j \ge 2,$$

where the same bound holds for $j = 1$ since $\widehat{C}_1(dx,dy) = r(y)J_1^{\mathrm{F}}(dx,dy)$ and $b_* \le r(y) \le 1$. Since the reference edge has marginals

$$J_j^{\mathrm{F}}(dx,\mathcal{X}) = \pi_j(dx), \qquad J_j^{\mathrm{F}}(\mathcal{X},dy) = \pi_{j-1}(dy),$$

each of the two terms $\widehat{C}_j(dx,\mathcal{X})$ and $\widehat{C}_{j+1}(\mathcal{X},dx)$ defining $\mu_j(dx)$ at an interior checkpoint lies between $b_*\pi_j(dx)$ and $\pi_j(dx)$. Thus Equation (C.2) gives

$$\mu_j(dx) = \widehat{C}_j(dx,\mathcal{X}) + \widehat{C}_{j+1}(\mathcal{X},dx) \quad\Longrightarrow\quad 2b_*\pi_j \le \mu_j \le 2\pi_j.$$

At either boundary only one edge contributes, giving $b_*\pi_j \le \mu_j \le \pi_j$. Hence, for every admissible (j,d), Equation (C.4) implies

$$b_*\pi_j \le \mu_j \le 2\pi_j, \qquad \frac{b_*}{\mathcal{Z}}\pi_j(dx) \le \Pi(j,dx,d) = \frac{\mu_j(dx)}{\mathcal{Z}} \le \frac{2}{\mathcal{Z}}\pi_j(dx).$$

The second bound above and the assumption $J_j^{\mathrm{F}}, J_j^{\mathrm{B}} \ll \pi_j \otimes \pi_{j-1}$, which gives $K_j^{\mathrm{F}}(\cdot \mid x) \ll \pi_{j-1}$ and $K_{j-1}^{\mathrm{B}}(\cdot \mid y) \ll \pi_j$ for π_j-almost every x and π_{j-1}-almost every y, respectively (also for the corrected B kernel), imply

$$P_{\mathrm{move}}(z, \cdot) \ll \Pi, \qquad \bar{P}_{\mathrm{move}}(z, \cdot) \ll \Pi \quad \text{for } \Pi\text{-almost every } z,$$

so the two densities in Equation (C.99) exist. For an F move with $2 \le j < J$, multiplying the move probability by the sampling transition gives

$$P_{\mathrm{move}}((j, x, \mathrm{F}), (j-1, dy, \mathrm{F})) = p_j^{\mathrm{F}}(x) Q_j^{\mathrm{F}}(dy \mid x) = \frac{\sqrt{\widehat{V}_j(x)\widehat{V}_{j-1}(y)}}{a_j^{\mathrm{F}}(x) + a_j^{\mathrm{B}}(x)} K_j^{\mathrm{F}}(dy \mid x).$$

Each directional mass a_j^d integrates a weight in $[b_*, 1]$ against a probability kernel, so $a_j^d \ge b_*$. Hence $a_j^{\mathrm{F}} + a_j^{\mathrm{B}} \ge 2b_*$ and the numerator weight is at most 1, giving $P_{\mathrm{move}}((j, x, \mathrm{F}), (j-1, dy, \mathrm{F})) \le K_j^{\mathrm{F}}(dy \mid x)/b_*$. For $z = (j, x, \mathrm{F})$ and $z' = (j-1, y, \mathrm{F})$, taking densities and using $J_j^{\mathrm{F}}(dx, dy) = \pi_j(dx) K_j^{\mathrm{F}}(dy \mid x)$ yields

$$p_{\mathrm{move}}(z, z') \underbrace{\frac{d\Pi(j-1, \cdot, \mathrm{F})}{d\pi_{j-1}}(y)}_{\ge b_*/\mathcal{Z}} \le \frac{1}{b_*} \frac{dK_j^{\mathrm{F}}(\cdot \mid x)}{d\pi_{j-1}}(y) = \frac{1}{b_*} \frac{dJ_j^{\mathrm{F}}}{d(\pi_j \otimes \pi_{j-1})}(x, y).$$

For the uncorrected F move, whose density relative to Π is $\bar{p}_{\mathrm{move}}$, omitting ρ leaves the numerator unchanged and may change the B mass in the denominator, which remains at least b_*. Dividing both movement measures by the same stationary reference measure Π therefore gives

$$p_{\mathrm{move}}(z, z') \le \frac{\mathcal{Z}}{b_*^2} \frac{dJ_j^{\mathrm{F}}}{d(\pi_j \otimes \pi_{j-1})}(x, y), \qquad \bar{p}_{\mathrm{move}}(z, z') \le \frac{\mathcal{Z}}{b_*^2} \frac{dJ_j^{\mathrm{F}}}{d(\pi_j \otimes \pi_{j-1})}(x, y).$$

The same bound holds at boundaries because the forced transition Q_j^d has denominator $a_j^d \ge b_*$, and on the clean edge because $r \le 1$, with directions reset inward. For B moves, including the correction ρ gives $\pi_{j-1}(dy) K_{j-1}^{\mathrm{B},\rho}(dx \mid y) = J_j^{\mathrm{F}}(dx, dy)$, whereas omitting it gives $\pi_{j-1}(dy) K_{j-1}^{\mathrm{B}}(dx \mid y) = J_j^{\mathrm{B}}(dx, dy)$. The same calculation therefore gives, for $z = (j-1, y, \mathrm{B})$ and $z' = (j, x, \mathrm{B})$ when $j < J$, with $z' = (J, x, \mathrm{F})$ when $j = J$,

$$p_{\mathrm{move}}(z, z') \le \frac{\mathcal{Z}}{b_*^2} \frac{dJ_j^{\mathrm{F}}}{d(\pi_j \otimes \pi_{j-1})}(x, y), \qquad \bar{p}_{\mathrm{move}}(z, z') \le \frac{\mathcal{Z}}{b_*^2} \frac{dJ_j^{\mathrm{B}}}{d(\pi_j \otimes \pi_{j-1})}(x, y).$$

Adding the bounds for p_{move} and $\bar{p}_{\mathrm{move}}$ in each direction yields

$$p_{\mathrm{move}}(z, z') + \bar{p}_{\mathrm{move}}(z, z') \le \frac{\mathcal{Z}}{b_*^2} \begin{cases} 2 \dfrac{dJ_j^{\mathrm{F}}}{d(\pi_j \otimes \pi_{j-1})}(x, y), & \text{F move}, \\ \dfrac{dJ_j^{\mathrm{F}}}{d(\pi_j \otimes \pi_{j-1})}(x, y) + \dfrac{dJ_j^{\mathrm{B}}}{d(\pi_j \otimes \pi_{j-1})}(x, y), & \text{B move}. \end{cases}$$

For each allowed move between checkpoints j and $j-1$, the stationary upper bound gives $\Pi(dz)\Pi(dz') \le 4\pi_j(dx)\pi_{j-1}(dy)/\mathcal{Z}^2$. Using these bounds and $(a+b)^p \le 2^{p-1}(a^p + b^p)$, we integrate over x, y and sum over all edges and both move directions to obtain

$$H_{\rho,p}^p \le C \sum_{j=1}^{J} \sum_{d \in \{\mathrm{F}, \mathrm{B}\}} \left\| \frac{dJ_j^d}{d(\pi_j \otimes \pi_{j-1})} \right\|_{L^p(\pi_j \otimes \pi_{j-1})}^p < \infty$$

for some finite constant C, by Equation (C.119). □

D Finite-Population Synchronized CS-VGB

Running N independent single-chain CS-VGB samplers explores several trajectories in parallel, but the chains generally occupy different checkpoints and directions. In this case, process-verifier values at different noise levels need not be directly comparable, making it difficult to allocate search effort

across chains. We therefore synchronize the checkpoint and direction while retaining distinct particle states, placing all particles at a common noise level where their verifier values can be compared without differences in checkpoint calibration.

We first construct **Multi-chain Synchronized CS-VGB (MCS-CS-VGB)**, which preserves the product clean target $(\pi_0^\star)^{\otimes N}$ for every finite N under the assumptions below. To adapt this synchronized structure to finite-budget search with one-step posterior maps, we share the highest-reward endpoint among Nm_e posterior draws, together with its parent, across particles; each particle then generates its own local candidate pool and selects the highest-weight candidate (Equations (F.4) and (F.5)). This best-endpoint sharing and local selection heuristic combines parallel exploration with recurrent search, but does not inherit the stationary-target guarantee of the exact sampler developed below.

Exact MCS-CS-VGB kernel. For $N \geq 1$, write $\mathbf{x} = (x^{(1)}, \ldots, x^{(N)})$ and $\boldsymbol{\pi}_j := \pi_j^{\otimes N}$. The population state is $(j, \mathbf{x}, d)$, where all particles share j and d. Define the *product edge measure*

$$\widehat{C}_j^{(N)}(d\mathbf{x}, d\mathbf{y}) := \bigotimes_{n=1}^{N} \widehat{C}_j(dx^{(n)}, dy^{(n)}). \tag{D.1}$$

Recalling the single-chain directional masses $a_j^{\mathrm{F}}, a_j^{\mathrm{B}}$ from Equation (B.13), the population directional masses at $1 \leq j < J$ are

$$A_j^{\mathrm{F}}(\mathbf{x}) := \prod_{n=1}^{N} a_j^{\mathrm{F}}(x^{(n)}), \qquad A_j^{\mathrm{B}}(\mathbf{x}) := \prod_{n=1}^{N} a_j^{\mathrm{B}}(x^{(n)}), \tag{D.2}$$

and

$$p_j^{\mathrm{F}} := \frac{A_j^{\mathrm{F}}}{A_j^{\mathrm{F}} + A_j^{\mathrm{B}}}, \qquad p_j^{\mathrm{B}} := \frac{A_j^{\mathrm{B}}}{A_j^{\mathrm{F}} + A_j^{\mathrm{B}}}, \qquad \ell_j := \min\{p_j^{\mathrm{F}}, p_j^{\mathrm{B}}\}.$$

All particles share the action, but conditional on a move each draws from its own transition:

$$Q_j^{\mathrm{F},(N)}(d\mathbf{y} \mid \mathbf{x}) := \prod_{n=1}^{N} Q_j^{\mathrm{F}}(dy^{(n)} \mid x^{(n)}), \qquad Q_j^{\mathrm{B},\rho,(N)}(d\mathbf{y} \mid \mathbf{x}) := \prod_{n=1}^{N} Q_j^{\mathrm{B},\rho}(dy^{(n)} \mid x^{(n)}). \tag{D.3}$$

The action probabilities are those of Table B.1 with the population values $p_j^{\mathrm{F}}, p_j^{\mathrm{B}}, \ell_j$. These shared direction probabilities and the particlewise transitions above are induced by the same product edge measure in Equation (D.1), whose marginals are

$$\widehat{C}_j^{(N)}(d\mathbf{x}, \mathcal{X}^N) = \boldsymbol{\pi}_j(d\mathbf{x}) A_j^{\mathrm{F}}(\mathbf{x}), \qquad \widehat{C}_{j+1}^{(N)}(\mathcal{X}^N, d\mathbf{x}) = \boldsymbol{\pi}_j(d\mathbf{x}) A_j^{\mathrm{B}}(\mathbf{x}). \tag{D.4}$$

Define

$$\begin{aligned} \mu_0^{(N)}(d\mathbf{x}) &:= \widehat{C}_1^{(N)}(\mathcal{X}^N, d\mathbf{x}), & \mu_J^{(N)}(d\mathbf{x}) &:= \widehat{C}_J^{(N)}(d\mathbf{x}, \mathcal{X}^N), \\ \mu_j^{(N)}(d\mathbf{x}) &:= \boldsymbol{\pi}_j(d\mathbf{x}) \left(A_j^{\mathrm{F}}(\mathbf{x}) + A_j^{\mathrm{B}}(\mathbf{x})\right), & & 1 \leq j < J. \end{aligned} \tag{D.5}$$

and

$$\mathcal{Z}_N := \mu_0^{(N)}(\mathcal{X}^N) + 2\sum_{j=1}^{J-1} \mu_j^{(N)}(\mathcal{X}^N) + \mu_J^{(N)}(\mathcal{X}^N) \in (0, \infty). \tag{D.6}$$

Theorem D.1 (Stationarity of MCS-CS-VGB)**.** *Under the assumptions of Theorem C.1, the MCS-CS-VGB kernel has stationary law*

$$\Pi_N(0, d\mathbf{x}, \mathrm{B}) := \frac{\mu_0^{(N)}(d\mathbf{x})}{\mathcal{Z}_N}, \quad \Pi_N(J, d\mathbf{x}, \mathrm{F}) := \frac{\mu_J^{(N)}(d\mathbf{x})}{\mathcal{Z}_N}, \quad \Pi_N(j, d\mathbf{x}, d) := \frac{\mu_j^{(N)}(d\mathbf{x})}{\mathcal{Z}_N}, \tag{D.7}$$

where $1 \leq j < J$ *and* $d \in \{\mathrm{F}, \mathrm{B}\}$. *Its clean conditional is*

$$\Pi_N(d\mathbf{x} \mid j = 0) = \left(\pi_0^\star\right)^{\otimes N}(d\mathbf{x}). \tag{D.8}$$

Proof. Write P_N for the MCS-CS-VGB kernel. We verify $\Pi_N P_N = \Pi_N$ by summing the one-update measures entering each state. At an interior state $(j, \mathbf{x}, \mathrm{F})$, an F move arrives from $j+1$. Since $Q_{j+1}^{\mathrm{F},(N)}$ is the conditional transition of the product edge measure, Equations (D.4) and (D.5) gives

$$\int_{\mathcal{X}^N} \Pi_N(j+1, d\mathbf{x}', \mathrm{F}) p_{j+1}^{\mathrm{F}}(\mathbf{x}') Q_{j+1}^{\mathrm{F},(N)}(d\mathbf{x} \mid \mathbf{x}') = \frac{1}{\mathcal{Z}_N} \int_{\mathcal{X}^N} \widehat{C}_{j+1}^{(N)}(d\mathbf{x}', d\mathbf{x}) = \frac{\mu_j^{(N)}(d\mathbf{x}) p_j^{\mathrm{B}}(\mathbf{x})}{\mathcal{Z}_N}. \tag{D.9}$$

Here $p_J^{\mathrm{F}} = 1$ for the forced move at the noise boundary. A B-to-F flip at the same $(j, \mathbf{x})$ contributes $\Pi_N(j, d\mathbf{x}, \mathrm{B})[p_j^{\mathrm{F}}(\mathbf{x}) - \chi \ell_j(\mathbf{x})]$, and an F stay contributes $\Pi_N(j, d\mathbf{x}, \mathrm{F}) \chi \ell_j(\mathbf{x})$. Thus

$$(\Pi_N P_N)(j, d\mathbf{x}, \mathrm{F}) = \frac{\mu_j^{(N)}(d\mathbf{x})}{\mathcal{Z}_N} \left(p_j^{\mathrm{B}} + p_j^{\mathrm{F}} - \chi \ell_j + \chi \ell_j \right) = \Pi_N(j, d\mathbf{x}, \mathrm{F}).$$

The B case is symmetric. At the boundaries, the incoming measures are $\widehat{C}_1^{(N)}(\mathcal{X}^N, d\mathbf{x})/\mathcal{Z}_N = \Pi_N(0, d\mathbf{x}, \mathrm{B})$ and $\widehat{C}_J^{(N)}(d\mathbf{x}, \mathcal{X}^N)/\mathcal{Z}_N = \Pi_N(J, d\mathbf{x}, \mathrm{F})$, respectively, proving stationarity. Finally,

$$\mu_0^{(N)}(d\mathbf{x}) = \bigotimes_{n=1}^{N} \left(r(x^{(n)}) \pi_0(dx^{(n)}) \right) = Z^N \left(\pi_0^\star \right)^{\otimes N}(d\mathbf{x}),$$

which proves Equation (D.8) after conditioning on $j = 0$. □

Finite-candidate implementation of MCS-CS-VGB. Use the local pools and coordinatewise reverse insertion of Appendix B.3. For fixed $\alpha > 0$, denote the depth-weighted candidate contributions in Equation (B.18) for particle n by $e_{j,n,i}^d$, with $d \in \{\mathrm{F}, \mathrm{B}\}$. Define

$$\begin{aligned} \widehat{A}_j^{\mathrm{F}} &:= \prod_{n=1}^{N} \left(\frac{1}{L_F} \sum_{i=1}^{L_F} e_{j,n,i}^{\mathrm{F}} \right), \quad E_{j,\mathbf{i}}^{\mathrm{F}} := \prod_{n=1}^{N} e_{j,n,i_n}^{\mathrm{F}}, \\ \widehat{A}_j^{\mathrm{B}} &:= \prod_{n=1}^{N} \left(\frac{1}{L_B} \sum_{k=1}^{L_B} e_{j,n,k}^{\mathrm{B}} \right), \quad E_{j,\mathbf{k}}^{\mathrm{B}} := \prod_{n=1}^{N} e_{j,n,k_n}^{\mathrm{B}}. \end{aligned} \tag{D.10}$$

Without lazy pool evaluation, substitute $\widehat{A}_j^{\mathrm{F}}, \widehat{A}_j^{\mathrm{B}}$ into the synchronized action rule. On a move in direction $d \in \{\mathrm{F}, \mathrm{B}\}$, independently select each particle's candidate with probability

$$\mathbb{P}(I_n = i \mid d \text{ move}) = \frac{e_{j,n,i}^d}{\sum_{h=1}^{L_d} e_{j,n,h}^d}, \qquad \mathbb{P}(\mathbf{I} = \mathbf{i} \mid d \text{ move}) = \prod_{n=1}^{N} \frac{e_{j,n,i_n}^d}{\sum_{h=1}^{L_d} e_{j,n,h}^d} = \frac{E_{j,\mathbf{i}}^d}{L_d^N \widehat{A}_j^d}, \tag{D.11}$$

and reverse-insert the previous states into independent uniform opposite-pool slots.

With *lazy pool evaluation*, an F update uses the full F pools and B mark $\mathbf{k}$; a B update uses the full B pools and F mark $\mathbf{i}$. In direction F, draw $I_n \sim \mathrm{Unif}([L_F])$ independently and write $\mathbf{I} = (I_1, \ldots, I_N)$; in direction B, draw $K_n \sim \mathrm{Unif}([L_B])$ independently and write $\mathbf{K} = (K_1, \ldots, K_N)$. Conditional on these indices, the probabilities of (move, stay, flip) are

$$\begin{aligned} \mathrm{F}: &\quad \frac{\left(\widehat{A}_j^{\mathrm{F}},\ \chi \min\{E_{j,\mathbf{I}}^{\mathrm{F}}, E_{j,\mathbf{k}}^{\mathrm{B}}\},\ E_{j,\mathbf{k}}^{\mathrm{B}} - \chi \min\{E_{j,\mathbf{I}}^{\mathrm{F}}, E_{j,\mathbf{k}}^{\mathrm{B}}\} \right)}{\widehat{A}_j^{\mathrm{F}} + E_{j,\mathbf{k}}^{\mathrm{B}}}, \\ \mathrm{B}: &\quad \frac{\left(\widehat{A}_j^{\mathrm{B}},\ \chi \min\{E_{j,\mathbf{K}}^{\mathrm{B}}, E_{j,\mathbf{i}}^{\mathrm{F}}\},\ E_{j,\mathbf{i}}^{\mathrm{F}} - \chi \min\{E_{j,\mathbf{K}}^{\mathrm{B}}, E_{j,\mathbf{i}}^{\mathrm{F}}\} \right)}{\widehat{A}_j^{\mathrm{B}} + E_{j,\mathbf{i}}^{\mathrm{F}}}. \end{aligned} \tag{D.12}$$

On a move, discard the uniform indices and redraw by Equation (D.11), and mark the reverse-inserted slots at an interior destination. On an F-to-B flip, mark the preselected indices $\mathbf{I}$ in the retained F pools and complete only missing B entries; the B-to-F case uses the preselected indices $\mathbf{K}$ symmetrically. A stay changes nothing; generated entries are retained across flips and stays. Boundary moves are forced without marks.

Write $\mathsf{G}_j^{(N)} := \prod_n \mathsf{G}_j$. Let $\mathcal{Z}_{N,\alpha}$ be the normalizer in Equation (D.6) with each edge $\widehat{C}_j^{(N)}$ replaced by $\alpha^{N(j-1)} \widehat{C}_j^{(N)}$.

Theorem D.2 (Stationary clean target with finite local pools)**.** *Without lazy evaluation, the interior stationary law is*

$$\Pi_{\mathrm{pool},N}(j, d\mathbf{x}, d\mathbf{u}, d) = \frac{\boldsymbol{\pi}_j(d\mathbf{x})\mathsf{G}_j^{(N)}(d\mathbf{u} \mid \mathbf{x})}{\mathcal{Z}_{N,\alpha}}(\widehat{A}_j^{\mathrm{F}} + \widehat{A}_j^{\mathrm{B}}). \tag{D.13}$$

With lazy evaluation, the F law is

$$\Pi_{\mathrm{mark},N}(j, d\mathbf{x}, d\mathbf{u}, \mathrm{F}, \mathbf{k}) = \frac{\boldsymbol{\pi}_j(d\mathbf{x})\mathsf{G}_j^{(N)}(d\mathbf{u} \mid \mathbf{x})}{\mathcal{Z}_{N,\alpha}} \frac{\widehat{A}_j^{\mathrm{F}} + E_{j,\mathbf{k}}^{\mathrm{B}}}{L_B^N}, \tag{D.14}$$

and the B law is symmetric. At a boundary, both laws use only the available directional mass. Summing over the mark recovers $\Pi_{\mathrm{pool},N}$*, and*

$$\Pi_{\mathrm{mark},N}(d\mathbf{x} \mid j = 0) = \Pi_{\mathrm{pool},N}(d\mathbf{x} \mid j = 0) = (\pi_0^\star)^{\otimes N}(d\mathbf{x}). \tag{D.15}$$

Proof. The reverse-insertion identity in Equation (C.13) applies independently to each particle, so taking products gives the corresponding identity for the population. With the factorized selection rule in Equation (D.11) and independent uniform insertion slots, the proof of Theorem C.2 therefore applies with

$$(L_d, e_{j,i}^d, \widehat{a}_j^d) \longmapsto (L_d^N, E_{j,\mathbf{i}}^d, \widehat{A}_j^d).$$

The uniform tuple draws and conditional action probabilities in Equation (D.12) give the same flip–stay cancellation. Thus both stated laws are stationary, and summing over the mark recovers the unmarked law. At the clean boundary, the edge weight is one and integrating out the pools gives a measure proportional to $\bigotimes_{n=1}^N [r(x^{(n)})\pi_0(dx^{(n)})]$, whose normalization is $(\pi_0^\star)^{\otimes N}$.

□

E Additional Related Works

Guidance and particle sampling. Like CS-VGB, guidance and candidate-selection methods use intermediate information to favor desirable outputs, either by modifying denoising predictions (Dhariwal & Nichol, 2021; Ho & Salimans, 2022; Song et al., 2023) or selecting candidates using soft values (Li et al., 2025). TDS and Feynman–Kac steering address the related target-sampling problem through particle weighting and resampling (Wu et al., 2023a; Singhal et al., 2025). TDS retains asymptotic exactness through importance corrections as the particle count increases, whereas our corrected construction preserves the stationary clean target at fixed finite pool sizes.

Search and iterative refinement. Our use of additional compute to revisit generation decisions connects CS-VGB to verifier-based search (Ma et al., 2025; Zhang et al., 2026) and the noising–denoising refinement of RERD and ABCD (Uehara et al., 2025; Lee et al., 2025a). Distributional guarantees also appear in DTS, which improves tree-based values as rollouts accumulate (Jain et al., 2025), and IterRef, which applies multiple-try Metropolis to intermediate distributions in discrete diffusion (Lee et al., 2025b). Our focus is instead on preserving the continuous-state clean target despite process-verifier error, without requiring those estimates to become exact.

Verifier error and backtracking. The closest conceptual predecessors of CS-VGB are AR-VGB and MDM-VGB, which separate imperfect process verifiers from the terminal target in discrete generation (Rohatgi et al., 2026; Jeon et al., 2026). CS-VGB extends this principle to continuous states by pairing denoising and renoising transitions. A related approach, DOIT, uses Monte Carlo estimates of a Doob h-correction to steer pretrained diffusion models toward a reward-tilted target and analyzes the resulting approximation error (Zhu et al., 2026). CS-VGB instead uses process verifiers within a bidirectional stationary construction, where verifier error affects convergence rates but not the clean target under correct transition pairing. We separately quantify the bias caused by uncorrected transition mismatch.

Stochastic transitions and posterior maps. We instantiate CS-VGB with diffusion models and one-step posterior maps. For diffusion models, reverse denoising steps and forward noising transitions

provide the checkpoint kernels (Ho et al., 2020; Song et al., 2021), whereas Meta Flow Maps and Posterior Diamond Maps generate clean posterior samples that we combine with Gaussian bridges to reach intermediate checkpoints (Potaptchik et al., 2026; Holderrieth et al., 2026). Related approaches include Noise Conditional Variational Score Distillation (NCVSD), which distills diffusion models into stochastic denoisers that approximate clean posterior samples across noise levels (Peng et al., 2025), and Strong Stochastic Flow Maps (SSFM), which learn the pathwise solution map of an additive-noise SDE and thus approximate its conditional transitions (McCallum et al., 2026). These methods provide transition samplers that can be used within CS-VGB, which determines how to traverse checkpoint kernels in both directions while preserving the stationary clean target under correct transition pairing. In particular, as conditional samplers of clean states, NCVSD and SSFM are also compatible with CS-VGB through the Gaussian-bridge construction (Proposition B.1).

F Experiment Settings and Supplementary Results

Settings. All experiments were implemented in Python using PyTorch (Paszke et al., 2019) and conducted on Ubuntu Linux. We used NVIDIA H100 GPUs with 80 GB of memory, NVIDIA H200 GPUs with 141 GB of memory, and NVIDIA RTX PRO 6000 Blackwell GPUs with 96 GB of memory.

F.1 Shared Search Protocol

Unlike the analytic sampling experiments, practical search sets $\rho_j = 1$ and proposal importance weights $dK/dG = 1$. It archives clean outputs in $\mathcal{S}$ and returns $x^\star \in \arg\max_{x\in\mathcal{S}} r(x)$, with task-specific endpoint archiving as described below. These choices prioritize reward rather than exact target sampling.

Gradient guidance for posterior bridges. For ImageNet and CelebA, we use normalized proposal-mean shifts before candidate selection, rather than the covariance-scaled bridge guidance in Equation (B.36). Following the Monte Carlo reward-gradient estimation used in MFM and PDM (Potaptchik et al., 2026; Holderrieth et al., 2026), we differentiate the log mean reward of posterior samples. Write $t_j = t_{\kappa_j}$ and use m_j, a_j, σ_j from Equation (B.30). Let $\zeta_\ell \overset{\text{iid}}{\sim} \mathcal{N}(0, I)$, $\ell = 1, \dots, m_g$, be the noise inputs to the posterior sampler. Define

$$Z_\ell(x) = \Phi^\theta_{t_j,0}(x, \zeta_\ell), \qquad h_j(x) = \nabla_x \log\left[\frac{1}{m_g}\sum_{\ell=1}^{m_g} r(Z_\ell(x))\right], \qquad \bar{Z} = \frac{1}{m_g}\sum_{\ell=1}^{m_g} Z_\ell(x). \tag{F.1}$$

The gradient is taken through the posterior sampler; at $j = 0$, $h_0 = \nabla \log r$. The implemented guided proposals are

$$G^{\mathrm{F}}_j(dy \mid x, Z) = \mathcal{N}\left(dy; m_j(x, Z) + g_F(t_j - t_{j-1})\left\|\frac{\bar{Z} - x}{t_j}\right\|_2 \frac{h_j(x)}{\|h_j(x)\|_2}, \frac{t^2_{j-1}\sigma^2_j}{t^2_j} I\right), \tag{F.2}$$

$$G^{\mathrm{B}}_j(dy \mid x) = \mathcal{N}\left(dy; a_{j+1}x + g_B\sigma_{j+1}\frac{h_j(x)}{\mathrm{RMS}(h_j(x))}, \sigma^2_{j+1} I\right), \tag{F.3}$$

These kernels are conditional on the guidance draws, with Z a clean posterior draw and $\mathrm{RMS}(h) = \|h\|_2/\sqrt{d}$ for state dimension d. Both guidance shifts are defined to be zero when $h_j(x) = 0$. Averaging over the posterior draws gives the proposal kernels conditional on x alone.

Best-endpoint sharing and local candidate selection. For unguided search with one-step posterior maps, we follow MFM-Search (Potaptchik et al., 2026) in sharing the highest-reward clean endpoint across particles. From each parent $x^{(n)}_j$, draw m_e posterior endpoints $Z_{n,\ell}$ and select

$$(n^\star, \ell^\star) \in \arg\max_{1\le n\le N,\, 1\le \ell\le m_e} r(Z_{n,\ell}), \qquad (x^\star, Z^\star) := (x^{(n^\star)}_j, Z_{n^\star,\ell^\star}). \tag{F.4}$$

Each particle then forms its own L_F bridge candidates from this shared parent–endpoint pair and selects the highest-weight local child:

$$Y_{n,i} = m_j(x^\star, Z^\star) + \frac{t_{j-1}\sigma_j}{t_j}\xi_{n,i}, \qquad \xi_{n,i} \overset{\text{iid}}{\sim} \mathcal{N}(0, I), \qquad i^\star_n \in \arg\max_{1\le i\le L_F} e^{\mathrm{F}}_{j,n,i}. \tag{F.5}$$

Here $e^{\mathrm{F}}_{j,n,i}$ is the depth-weighted candidate contribution in Equation (B.18), with the search corrections specified above. After moving to $Y_{n,i^\star_n}$, reverse insertion uses the shared parent $x^\star$, not the particle's previous parent, and the reverse-insertion mark is updated accordingly at interior destinations.

Renoising warp. Finite-budget search favors frequent clean outputs, motivating a smaller depth weight α. However, this can cause the chain to switch back to denoising after only shallow renoising, spending NFE on repeated local refinements near the clean endpoint. To encourage broader exploration within the same budget, we optionally force renoising to depth $j = \lceil \beta J \rceil$ after archiving a clean sample; $\beta = 0$ disables this warp. For $\beta > 0$,

$$x_{\mathrm{w}} \sim K^{\mathrm{B}}_{0\to j}(\cdot \mid x), \qquad K^{\mathrm{B}}_{0\to j} := K^{\mathrm{B}}_0 \cdots K^{\mathrm{B}}_{j-1}.$$

We reset pools and marks and resume from $(j, x_{\mathrm{w}}, \mathrm{F})$. This direct warp is a search-only heuristic, distinct from ordinary backtracking.

Skipping stays. Since stays leave the state, pools, and mark unchanged, we sample the next non-stay action directly:

$$\mathbb{P}(\mathsf{move} \mid \text{not stay}) = \frac{p_{\mathrm{move}}}{1 - p_{\mathrm{stay}}}, \qquad \mathbb{P}(\mathsf{flip} \mid \text{not stay}) = \frac{p_{\mathrm{flip}}}{1 - p_{\mathrm{stay}}}.$$

Candidate and mark selection are unchanged. Skipping stays alone preserves the chain's stationary clean conditional because clean-boundary states have no stays.

F.2 Target Recovery in Two Dimensions

We use a 24-component Gaussian mixture as the reference and define the reward by

$$\pi_0(x) = \sum_{k=1}^{24} w_k \mathcal{N}(x; \mu_k, \Sigma_k), \qquad r(x) = 10^{-3} + \sum_{\ell=1}^{6} a_\ell \exp\!\left[-\tfrac{1}{2}(x - c_\ell)^\top S_\ell^{-1}(x - c_\ell)\right].$$

Here w_k, μ_k, and Σ_k are the mixture weights, means, and covariance matrices, while a_ℓ, c_ℓ, and S_ℓ specify the strength, center, and spatial spread of each reward term. We use 32 intervals uniformly spaced in $-\log(1-t)$, with $t_j = 1 - 100^{-j/32}$ for $j = 0, \ldots, 32$. Renoising is Gaussian and denoising uses its exact conditional reverse, so $\rho_j = 1$. The initial state follows the reference marginal at $t_{32} = 0.99$.

Verifier error. The ideal verifier $V_j(x)$ is evaluated analytically using Gaussian-mixture posterior expectations of the terminal reward. To model a spatially and temporally coherent error rather than fresh evaluation noise, we use the fixed random Fourier field

$$\eta(x,t) = \sqrt{\frac{2}{M}} \sum_{m=1}^{M} c_m \cos(\omega_m^\top x + \nu_m t + \phi_m), \qquad s_\eta^2 = \frac{1}{M}\sum_{m=1}^{M} c_m^2, \tag{F.6}$$

with $M = 16$, $\omega_m \sim \mathcal{N}(0, I_2)$, $\nu_m, c_m \sim \mathcal{N}(0,1)$, and $\phi_m \sim \mathrm{Unif}[0, 2\pi)$. The field is fixed across methods and repetitions; s_η^2 is its variance over random phases conditional on the amplitudes.

The perturbed verifier is defined by

$$\log \widehat{V}_j(x) = \log V_j(x) + \lambda t\, \eta(x,t) - \tfrac{1}{2}(\lambda t)^2 s_\eta^2, \qquad t = t_{\kappa_j}. \tag{F.7}$$

We use $\lambda \in \{0, 3, 10, 15, 20\}$, and both perturbation terms vanish at $t = 0$, preserving the exact terminal reward.

Baseline protocols. Noisy guidance uses 16,384 Euler–Maruyama reverse-SDE steps, the analytic reference score, and unit-strength $\nabla \log \widehat{V}_t$ guidance. SMC uses 512 particles, 32 exact denoising intervals, incremental weights $\widehat{V}_{j-1}/\widehat{V}_j$, and resampling at every interval. RERD uses 512 samples, forward candidate pools of four, and 65,536 refinement updates with uniform renoising depths.

F.3 IMAGENET PROMPT-GUIDED GENERATION

Generator and evaluation prompts. We use the pretrained class-conditional MFM checkpoint (Potaptchik et al., 2026)[1] for ImageNet 256×256 (Deng et al., 2009), with Stability AI's fine-tuned VAE decoder[2], classifier-free guidance scale 2.5, and one model step per posterior draw. Class guidance is shared by all methods.

We construct 100 fixed evaluation prompts by pairing ImageNet classes with additional conditions: 21 actions, 20 object interactions, 20 visual states, 20 contexts, and 19 spatial relations. The generator is conditioned on the class label, while the full prompt is used to compute the reward and evaluate prompt alignment. Figure F.1 provides ten example prompts.

ImageNet: example evaluation prompts

0. **volcano (980)**
   `majestic volcano erupting with lava flowing down its sides`
1. **space shuttle (812)**
   `space shuttle launching into space, fiery fumes`
2. **hammerhead (4)**
   `a hammerhead nudging a floating buoy with its snout`
3. **stingray (6)**
   `a stingray partially buried in pale sand`
4. **hen (8)**
   `a hen in a sunlit farmyard`
5. **brambling (10)**
   `a brambling perched on a frosted branch`
6. **house finch (12)**
   `a house finch pulling a seed from a sunflower head`
7. **indigo bunting (14)**
   `an indigo bunting with rain-soaked plumage`
8. **bulbul (16)**
   `a bulbul taking flight`
9. **magpie (18)**
   `a magpie in a rain-soaked city park`

Figure F.1: Example prompts for ImageNet generation.

Reward and independent evaluation. We use HPSv2 (Wu et al., 2023b) with v2.1 weights. The reported score is $s(x, c) = 100\langle f_I(x), f_T(c)\rangle$, where $f_I(x)$ and $f_T(c)$ are the unit-normalized embeddings of image x and prompt c produced by the HPSv2 image and text encoders, respectively. For search, we use the calibrated log reward

$$\log r(x; c) = \frac{s(x, c) - \mu_c}{T\sigma_c}, \qquad T = 0.5,$$

where μ_c and σ_c are the mean and standard deviation of HPSv2 scores for prompt c, computed from 64 images generated by the base MFM under the corresponding class label without reward guidance or search. The same normalization statistics are used for all methods. We remark that the plotted HPSv2 values are raw scores, not standardized log rewards. We evaluate prompt alignment with VQAScore (Lin et al., 2024), using the 3B-parameter instruction-tuned Qwen2.5-VL model (Bai et al., 2025) as the evaluation backbone[3] (Figure 5). For the question *Does this figure show {prompt}?*

[1] https://github.com/adh1s/mfm
[2] https://huggingface.co/stabilityai/sd-vae-ft-mse
[3] https://huggingface.co/Qwen/Qwen2.5-VL-3B-Instruct

Please answer Yes or No., VQAScore reports the full-vocabulary next-token probability of *Yes*. CLIP-IQA (Wang et al., 2023) is shown for both ImageNet and CelebA in Figure F.3. Neither metric is used for optimization.

Guided CS-VGB. CS-VGB uses the posterior bridges in Appendix B.4, with $t_j = 0.99j/J$ and an initial unguided transition from $t = 1$ to 0.99. We set $\alpha = 0.3$, $\chi = 1$, and one posterior sample per value estimate, with lazy evaluation, skipped stays, and reverse insertion. We set $\beta = 0.2$. If the update budget is exhausted before the chain reaches the clean level, a forward-only fallback denoises the current state to the clean level to complete the output. Edge and proposal corrections are omitted as specified in Appendix F.1.

Guided runs use one chain; F guides forward proposals only, while F+B also guides backward proposals. Table F.1 lists F+B settings: posterior draws m_g, guidance scales g_F, g_B. The output is the best clean visit, without archiving extra verifier endpoints.

Table F.1: **ImageNet guided CS-VGB configurations.** B is the requested NFE budget; J counts checkpoint intervals.

B	J	L_F	L_B	m_g	g_F	g_B
2^{12}	64	1	1	4	1.75	1.5
2^{13}	64	4	4	8	2	1.5
10,000	64	4	4	8	2	1.5
19,609	128	8	8	16	3	2
42,362	129	16	8	16	3.5	2
2^{16}	129	16	8	16	3.5	2

Unguided CS-VGB search. Unguided runs use $L_F = L_B = 4$ and a budget-dependent population size N, drawing m_e clean endpoint candidates per particle for each denoising proposal (Table F.2). Following MFM-Search (Potaptchik et al., 2026), we select the highest-reward clean endpoint across particles. We then select the highest-weight child among the L_F local candidates for each particle and reverse-insert its parent. Following the search protocol of Posterior Diamond Maps (Holderrieth et al., 2026), we also archive clean endpoints generated for verifier evaluation and return the highest-reward image in the archive.

Table F.2: **ImageNet unguided CS-VGB configurations.** Columns give requested NFE budget B, intervals J, particles N, and clean endpoint candidates per particle m_e.

B	J	N	m_e	B	J	N	m_e
1,433	16	4	8	10,623	32	8	24
1,854	16	4	16	12,637	32	8	24
1,902	16	4	16	12,767	32	8	24
1,953	16	4	16	16,827	32	8	24
2,037	16	4	16	21,716	64	8	24
3,975	16	4	32	25,859	64	8	24
6,230	16	8	24	34,333	64	8	24
7,956	16	8	32	64,954	64	8	24
10,393	16	8	24	131,029	64	16	52

F.3.1 Qualitative comparisons on ImageNet

At approximately 65k NFE, we compare outputs for five of the first 20 evaluation prompts.

Table F.3: **ImageNet prompt-aligned outputs near 65k NFE.** Generated images with HPSv2 scores (reported on the $\times 100$ scale) and actual NFE. Bold marks the largest HPSv2 score in each row.

Prompt	BoN	MFM Search	CS-VGB Search	MFM-G	CS-VGB-G
ID 004 a hen in a sunlit farmyard.					
	30.25 NFE 65,536	35.81 NFE 65,536	36.61 NFE 64,768	46.96 NFE 64,990	**55.29** NFE 65,548
ID 005 a brambling perched on a frosted branch.					
	32.26 NFE 65,536	31.56 NFE 65,536	38.72 NFE 64,712	48.05 NFE 64,990	**51.73** NFE 65,536
ID 009 a magpie in a rain-soaked city park.					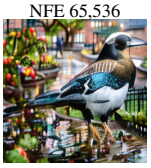
	30.97 NFE 65,536	35.86 NFE 65,536	39.32 NFE 64,904	47.33 NFE 64,990	**56.08** NFE 65,567
ID 014 a spotted salamander on a damp forest floor.	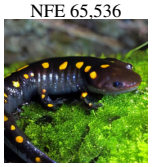	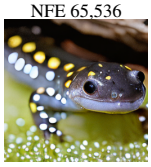			
	31.61 NFE 65,536	31.96 NFE 65,536	36.58 NFE 64,864	45.95 NFE 64,990	**51.71** NFE 65,573
ID 016 a tailed frog leaping from the water.	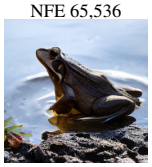				
	29.55 NFE 65,536	34.16 NFE 65,536	37.49 NFE 64,736	49.16 NFE 64,990	**51.96** NFE 65,575

CS-VGB-G depicts a wet urban setting for ID 009 and an airborne frog with splashes for ID 016, with HPSv2 gains of 2.80–8.75 over MFM-G across these cases. Higher scores do not verify every detail; aggregate reward and alignment results appear in Figure 5.

F.4 CelebA Facial-Attribute Generation

Generator and conditions. We use the pretrained PDM checkpoint (Holderrieth et al., 2026)[4] for CelebA 64×64 (Liu et al., 2015) and one model step per posterior draw or transition. PDM-G additionally uses the flow-map checkpoint from the same release for its drift.

Attribute conditions. We construct 100 fixed attribute conditions, each requesting presence or absence of ten attributes: Smiling, Eyeglasses, Bangs, Wearing Hat, Mouth Slightly Open, No Beard, Bushy Eyebrows, Wearing Lipstick, Wearing Earrings, and one of Black Hair, Brown Hair, or Blond Hair. Negative hair-color targets exclude that color. Figure F.2 shows ten examples.

[4] https://github.com/PeterHolderrieth/diamond_maps

CelebA: example attribute conditions

```
a close-up portrait photo of a person who

0. is not smiling, is not wearing eyeglasses, has bangs, is not
   wearing a hat, has their mouth slightly open, has blond hair, has
   a beard, has bushy eyebrows, is wearing lipstick, and is not
   wearing earrings

1. is smiling, is wearing eyeglasses, does not have bangs, is wearing
   a hat, has their mouth closed, does not have black hair, does not
   have a beard, does not have bushy eyebrows, is not wearing
   lipstick, and is wearing earrings

2. is not smiling, is wearing eyeglasses, has bangs, is not wearing a
   hat, has their mouth closed, does not have brown hair, has a
   beard, has bushy eyebrows, is wearing lipstick, and is wearing
   earrings

3. is smiling, is not wearing eyeglasses, does not have bangs, is
   wearing a hat, has their mouth slightly open, has black hair, does
   not have a beard, does not have bushy eyebrows, is not wearing
   lipstick, and is not wearing earrings

4. is smiling, is wearing eyeglasses, does not have bangs, is wearing
   a hat, has their mouth slightly open, has brown hair, has a beard,
   does not have bushy eyebrows, is wearing lipstick, and is wearing
   earrings

5. is not smiling, is not wearing eyeglasses, has bangs, is wearing a
   hat, has their mouth closed, does not have blond hair, does not
   have a beard, has bushy eyebrows, is not wearing lipstick, and is
   not wearing earrings

6. is smiling, is not wearing eyeglasses, has bangs, is not wearing a
   hat, has their mouth closed, does not have brown hair, has a
   beard, does not have bushy eyebrows, is not wearing lipstick, and
   is not wearing earrings

7. is not smiling, is wearing eyeglasses, does not have bangs, is not
   wearing a hat, has their mouth slightly open, has blond hair, does
   not have a beard, has bushy eyebrows, is wearing lipstick, and is
   wearing earrings

8. is smiling, is not wearing eyeglasses, has bangs, is not wearing a
   hat, has their mouth slightly open, has brown hair, does not have
   a beard, has bushy eyebrows, is not wearing lipstick, and is
   wearing earrings

9. is not smiling, is wearing eyeglasses, does not have bangs, is
   wearing a hat, has their mouth closed, does not have black hair,
   has a beard, does not have bushy eyebrows, is wearing lipstick,
   and is not wearing earrings
```

Figure F.2: Example attribute conditions for CelebA generation.

Reward and evaluation metrics. We construct the attribute reward using the CelebA-trained FaRL attribute classifier (Zheng et al., 2022) and a MobileNet-based RetinaFace detector from FACER[5], upscaling images to 256×256 before evaluation. Let $p_a(x)$ denote the classifier probability of the requested presence or absence of attribute a, clipped to $[10^{-6}, 1 - 10^{-6}]$. With reward temperature

[5] https://github.com/FacePerceiver/facer

one, the optimization reward and the reported *attribute geometric mean* are

$$\log r(x) = \sum_{a=1}^{10} \log p_a(x), \qquad \mathrm{GM}(x) = \exp\left(\frac{1}{10}\sum_{a=1}^{10} \log p_a(x)\right).$$

If no face is detected, optimization assigns $p_a = 10^{-4}$ with zero image gradient, while evaluation assigns GM zero.

Held-out metrics are CLIP-IQA (Wang et al., 2023) and VQAScore (Lin et al., 2024). VQAScore uses the same evaluation procedure and Qwen2.5-VL 3B backbone as in the ImageNet experiments.

Guided CS-VGB. CS-VGB uses the posterior bridges in Appendix B.4, a uniform grid $t_j = j/J$, $\alpha = 0.3$, $\chi = 1$, lazy evaluation, skipped stays, and one posterior draw per value estimate. Edge and proposal corrections are omitted. We set $\beta = 0.2$ and use the same forward-only fallback as in the ImageNet experiments when the update budget ends away from the clean level. Guided variants use one chain and $L_F = L_B = 2$, with $g_F = 0.1$ and $g_B = 0$ (F) or 0.1 (F+B).

Unguided CS-VGB search. Unguided runs use the same posterior bridges, depth weight, persistence, lazy evaluation, renoising, fallback, and archive policy described above, without gradient guidance. Budget-dependent settings appear in Table F.4. Each proposal draws m_e endpoints per particle and builds bridges from the highest-reward endpoint's parent.

Table F.4: **CS-VGB settings for CelebA.** Guided columns apply to both F-only and F+B variants. Here, m_g denotes the number of posterior draws for guidance, J the number of checkpoint intervals, N the number of particles, $L = L_F = L_B$ the candidate-pool size, and m_e the number of endpoint draws per particle.

	Guided ($N = 1$, $L = 2$)		Unguided			
Target NFE	J	m_g	J	N	L	m_e
2^6	4	1	2	2	1	4
2^7	4	1	4	2	4	4
2^8	4	1	4	4	1	8
2^9	8	1	8	4	4	8
2^{10}	16	2	16	8	1	4
2^{11}	16	2	8	8	4	16
2^{12}	32	4	16	8	4	8
2^{13}	32	4	16	8	4	32

F.4.1 Qualitative comparisons on CelebA

At around 4,096-NFE target, we compare outputs of baselines for five attribute conditions.

Table F.5: **CelebA attribute-conditioned outputs at a 4,096-NFE target budget.** +/− denote required presence/absence; beard is expressed directly. Bold marks the largest GM in each row.

Condition	BoN	PDM-SMC	CS-VGB Search	PDM-G	CS-VGB-G
ID 011 − smiling, + eyeglasses, − bangs, + hat, − mouth open, + black hair, Beard, + bushy brows, − lipstick, + earrings.	 0.2688 NFE 4,096	 0.7476 NFE 4,068	 0.7634 NFE 4,152	 0.9453 NFE 4,068	 **0.9805** NFE 4,101
ID 013 − smiling, + eyeglasses, + bangs, − hat, + mouth open, − blond hair, No beard, + bushy brows, + lipstick, − earrings.	 0.4232 NFE 4,096	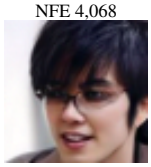 0.7046 NFE 4,068	 0.9511 NFE 4,184	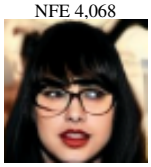 0.9668 NFE 4,068	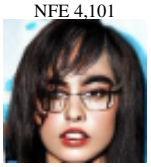 **0.9923** NFE 4,103
ID 001 + smiling, + eyeglasses, − bangs, + hat, − mouth open, − black hair, No beard, − bushy brows, − lipstick, + earrings.	 0.5476 NFE 4,096	 0.9456 NFE 4,068	 0.9775 NFE 4,104	 0.9678 NFE 4,068	 **0.9900** NFE 4,101
ID 003 + smiling, − eyeglasses, − bangs, + hat, + mouth open, + black hair, No beard, − bushy brows, − lipstick, − earrings.	 0.6736 NFE 4,096	 0.7743 NFE 4,068	 0.9759 NFE 4,184	 0.9670 NFE 4,068	 **0.9889** NFE 4,107
ID 006 + smiling, − eyeglasses, + bangs, − hat, − mouth open, − brown hair, Beard, − bushy brows, − lipstick, − earrings.	 0.7856 NFE 4,096	 0.9525 NFE 4,068	 0.9921 NFE 4,168	 0.9782 NFE 4,068	 **0.9969** NFE 4,106

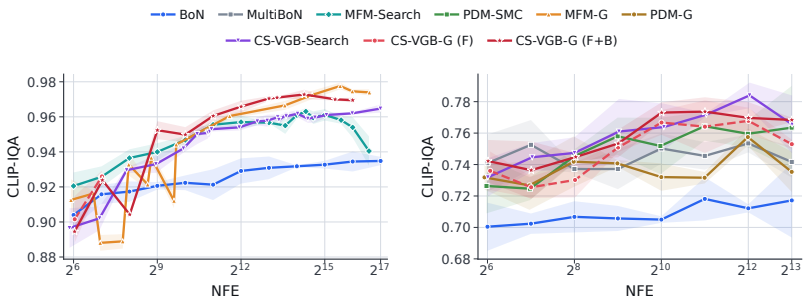

(a) ImageNet with MFM (b) CelebA with PDM

Figure F.3: **Independent image-quality evaluation.** CLIP-IQA of images selected using HPSv2 on ImageNet and attribute GM on CelebA. Curves show mean ±1 SD across three seeds, each evaluated on 100 conditions, against actual mean NFE. CLIP-IQA is not used for reward optimization.

F.5 Multi-Robot Planning with Diffusion Models

Tasks and reference models. We use MMD's pretrained diffusion models (Shaoul et al., 2025)[6]. Highways and Circle involve 16 robots; Conveyor Moderate and Hard involve four. Circle has a central obstacle of radius 0.20, and the Conveyor tasks differ in blocking geometry (Figure F.4). Joint success requires the entire multi-robot trajectory to be feasible. Each task has 64 fixed instances generated by perturbing starts, goals, and obstacles.

CS-VGB updates and guidance. CS-VGB uses one chain, checkpoint grid $[6, 2, 0]$, and guided DDPM rollouts between checkpoints. We set $L_F = L_B = 2$, $\alpha = 0.3$, and $\chi = 1$, with reverse insertion, lazy pool evaluation, and skipped stays. We use the search protocol in Appendix F.1, archive evaluated clean trajectories, and stop upon feasibility or the NFE/update limit (12,288 updates). In this case, renoising warp and forward-only fallback are disabled.

Evaluation and computational cost. Guided NFE caps are $2^9, 2^{10}, 2^{11}, 2^{12}, 2^{13}, 12{,}288, 2^{14}$. We plot actual mean model NFE per joint trial, including early stopping, without dividing by robot count. Table F.6 reports overall and task-wise success rates alongside actual mean NFE at the largest cap. CS-VGB-G achieves the highest overall success with the fewest model evaluations on average; its advantage is concentrated in Highways and Circle, while all methods are near ceiling on the Conveyor tasks.

Table F.6: **Planning results at the largest evaluated setting.** All methods use NFE cap 2^{14}. Success entries are percentages.

Method	Overall ± SD	Mean NFE	Highways	Circle	Conv. Mod.	Conv. Hard
BoN-G	67.58 ± 1.56	6,355.5	64.69	5.62	100.00	100.00
DFS-G	60.08 ± 1.25	8,029.7	40.00	0.31	100.00	100.00
ABCD-G	71.17 ± 1.52	5,765.2	73.12	11.56	100.00	100.00
CS-VGB-G	84.84 ± 1.98	5,044.8	80.00	60.00	99.69	99.69

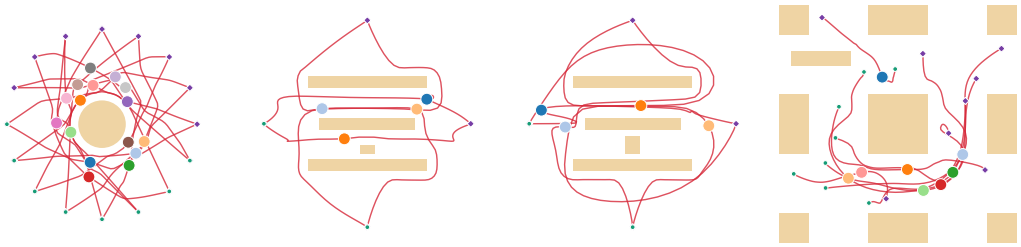

(a) Circle (b) Conveyor Moderate (c) Conveyor Hard (d) Highways

Figure F.4: **Successful multi-robot trajectories from CS-VGB-G.** Examples in Circle, Conveyor Moderate, Conveyor Hard, and Highways. Red curves show trajectories, green dots mark starts, purple diamonds mark goals, and shaded regions indicate obstacles. The Highways example uses eight robots, whereas the quantitative Highways experiments use 16.

[6] https://github.com/yoraish/mmd

F.6 Additional Ablations on Search Dynamics

We study lazy pool evaluation and directional persistence on 100 CelebA conditions (examples in Figure F.2). To measure evaluation cost, we compare lazy ON and OFF at $\chi = 1$ with the same number of updates within each search variant: 313 for guided CS-VGB and 45 for unguided search. Unlike a fixed-NFE comparison, this protocol measures how much computation each variant consumes for a fixed number of search updates.

Table F.7: **Lazy pool evaluation at matched update counts on CelebA.** NFE averages all 100 runs.

		Mean NFE			Attribute GM	
Variant	Updates	ON	OFF	Saved	ON	OFF
Guided	313	4,103.89	5,955.02	31.1%	0.98138	0.98182
Unguided	45	4,174.24	5,738.72	27.3%	0.83989	0.83953

Lazy pool evaluation. At matched update counts, lazy evaluation reduces mean NFE by 31.1% with guidance (Table F.7). Guided search retains nearly the same mean GM, while mean runtime decreases from 73.27 to 48.63 seconds, a 33.6% reduction. A similar computational saving is observed in unguided search, where lazy evaluation reduces mean NFE by 27.3% with similar mean GM among runs that reach a clean output. These results illustrate how deferring opposite-pool evaluations reduces the computational cost of continued search.

Directional persistence. We also vary directional persistence with lazy ON and an initial target budget of 4,096 NFE (Table F.8). Direction reversals and clean visits are measured at this initial budget over all 100 runs, including zero visits for runs without a clean output. Reward averages use all 100 conditions after supplementing missing outputs at higher NFE caps and are therefore not compute-matched across χ.

Table F.8: **Directional persistence on CelebA with lazy ON.** G/U denote guided/unguided search. Panel (a) reports supplemented attribute GM; panel (b) reports mean counts per run at the initial 4,096-NFE target, comparing $\chi = 0$ with $\chi = 1$.

(a) Attribute GM

χ	G	U
0	0.9787	0.8214
0.3	0.9803	0.8296
0.5	0.9816	0.8233
0.7	0.9804	0.8282
1	0.9804	0.8264

(b) Search dynamics: $\chi = 0 \rightarrow \chi = 1$

Variant	F-to-B flips	B-to-F flips	Clean visits
G	$45.81 \rightarrow 8.85$	$55.19 \rightarrow 20.53$	$12.47 \rightarrow 15.99$
U	$7.78 \rightarrow 4.35$	$9.10 \rightarrow 6.17$	$1.81 \rightarrow 2.21$

Increasing χ from 0 to 1 reduces flips in both directions and increases clean visits for both variants. This pattern is qualitatively consistent with the role of persistence in Proposition 2.1, which establishes $\Theta(J)$ expected updates to the first clean visit under its stated assumptions. Reward changes are small and nonmonotonic, with guided GM varying by 0.0029 across the sweep. Persistence therefore primarily changes checkpoint traversal in these experiments, rather than uniformly improving reward.