FINITE-HORIZON CONTROL CONSISTENCY FOR RANK-DEFICIENT DIFFUSION BRIDGES

Anonymous authors
Paper under double-blind review

ABSTRACT

Control-consistency methods for diffusion bridges recover state-space covectors from Brownian-space controls using instantaneous diffusion geometry, which becomes singular under rank-deficient noise even when the drift propagates forced directions into unforced coordinates. We introduce Finite-Horizon Bridge Control (FBC), replacing singular pointwise recovery under rank deficiency with a finite-horizon controllability lift built from dynamically transported noise directions. We show that the self-consistency identity itself does not require full-rank diffusion and connect the finite-horizon Gramian to Malliavin covariance. We prove exact finite-horizon recovery for controllable time-inhomogeneous affine diffusions with deterministic coefficients and characterize the information scales of dynamically reachable directions over short horizons in the local linearized controllability geometry. For nonlinear systems, we give sufficient executable-consistency conditions and an explicit short-window convergence rate for underdamped Langevin dynamics. FBC thus extends control consistency to rank-deficient and hypoelliptic diffusion bridges by exploiting finite-time controllability.

1 INTRODUCTION

Diffusion bridges describe stochastic dynamics conditioned on terminal information, with applications to rare-event simulation, transition-path sampling, and inference for partially observed diffusions (Baudoin, 2002). Classical approaches construct guided proposals with path-space corrections (Delyon & Hu, 2006; Schauer et al., 2017), while recent neural methods learn conditioned dynamics through score-based, Malliavin-based, variational, or control-based formulations (Heng et al., 2025; Pidstrigach et al., 2025; Yang et al., 2025).

Control Consistency Diffusion Bridges (CCDB) learn Brownian-space controls through an online fixed point in which controlled trajectories generate future training targets (Howard et al., 2026). Their standard construction recovers a state-space covector using instantaneous diffusion geometry. Under rank deficiency, however, this pointwise recovery cannot identify covector components outside the directly forced directions. Instantaneous unobservability can nevertheless coexist with finite-time accessibility: over a finite interval, the drift transports noise-driven directions into coordinates that are not forced directly. In underdamped Langevin dynamics, for example, noise acts on momentum while $dq_t = p_t\,dt$ propagates it into position. This makes finite-time controllability the natural geometry for control recovery.

Motivated by this observation, we introduce *Finite-Horizon Bridge Control* (FBC). FBC retains the executable controller in Brownian space but recovers its state-space covector from noise directions transported over a finite time window. The resulting Gramian captures the dynamically reachable geometry and connects to Malliavin covariance. For full-row-rank diffusions, FBC preserves the standard instantaneous recovery geometry; under rank deficiency, it augments this geometry with dynamically transported noise directions. This perspective separates self-consistency from recovery: rank deficiency obstructs covector recovery, while the control-consistency identity itself remains valid. We therefore study how finite-time controllability restores executable control information beyond instantaneous diffusion geometry.

Our main contributions are:

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.

- We introduce FBC, extending control consistency to rank-deficient diffusions through a finite-horizon recovery map connected to Malliavin covariance and classical controllability.
- We establish exact finite-horizon recovery for controllable affine diffusions with deterministic time-dependent coefficients, characterize short-horizon controllability-layer scaling in the local linearized geometry, and prove nonlinear consistency guarantees, including an explicit underdamped Langevin result.
- We validate FBC across continuous rank loss, rank-deficient nonlinear dynamics, and constrained molecular transition paths.

2 Related Work

Diffusion bridges and control-based learning. Classical bridge methods construct guided proposals with path-law corrections (Delyon & Hu, 2006; Schauer et al., 2017), while recent neural approaches learn conditioned dynamics through score-based, Malliavin-based, variational, or control-based formulations (Heng et al., 2025; Pidstrigach et al., 2025; Yang et al., 2025). CCDB learns Brownian-space controls through self-consistency along controlled trajectories (Howard et al., 2026), with related adjoint constructions in stochastic optimal control and generative-model fine-tuning (Domingo-Enrich et al., 2024; Domingo-Enrich et al., 2025; Havens et al., 2025). FBC retains this learning principle but replaces singular instantaneous covector recovery under rank deficiency with finite-horizon recovery.

Hypoelliptic geometry and transition paths. Degenerate and hypoelliptic bridge methods exploit drift propagation of noise into unforced coordinates, including recent underdamped score-matching formulations (Bierkens et al., 2020; Arnaudon et al., 2022; Li, 2016; Blessing et al., 2025). The finite-time geometry is classical, characterized by Malliavin nondegeneracy (Kusuoka & Stroock, 1985) and controllability (Kalman, 1960); FBC uses this geometry for finite-horizon covector recovery under rank deficiency. Transition-path sampling provides an important nonlinear application setting, with recent methods based on variational Doob transforms, path-integral control, and diffusion path sampling (Du et al., 2024; Holdijk et al., 2023; Seong et al., 2025).

3 Preliminaries

For the class of diffusion bridges considered here, terminal conditioning admits both a path-measure reweighting and a Brownian-space stochastic-control representation. This distinction is central to our setting: the bridge geometry is encoded by a state-space covector, whereas the quantity that is directly executable and learned by control-consistency methods is a Brownian-space control. We first introduce the reference and bridge path measures, then recall the control and Doob-transform representations that underlie finite-horizon recovery. Let $\mathbb{P}$ denote the path law of the reference diffusion

$$dX_t = b(t, X_t)\,dt + G(t, X_t)\,dW_t, \qquad X_0 = x_0, \tag{1}$$

where $X_t \in \mathbb{R}^n$, $W_t \in \mathbb{R}^m$, and $G(t,x) \in \mathbb{R}^{n\times m}$. For a nonnegative terminal weight F, define the corresponding diffusion-bridge law

$$\frac{d\mathbb{Q}^\star}{d\mathbb{P}}(X) = \frac{F(X_T)}{\mathbb{E}_{\mathbb{P}}[F(X_T)]}, \tag{2}$$

with exact endpoint conditioning obtained in the usual singular limit.

When $F = e^{-\Psi}$, Girsanov's theorem represents this bridge as the optimal law of a Brownian-space stochastic-control problem:

$$\begin{aligned} dX_t^v &= [b(t, X_t^v) + G(t, X_t^v)v(t, X_t^v)]\,dt + G(t, X_t^v)\,dW_t, \\ J(v) &= \mathbb{E}\left[\frac{1}{2}\int_0^T \|v(t, X_t^v)\|^2\,dt + \Psi(X_T^v)\right]. \end{aligned} \tag{3}$$

At the optimum of $J(v)$, the controlled path law is $\mathbb{Q}^\star$. To obtain the representation underlying control consistency, define $h(t,x) = \mathbb{E}_{\mathbb{P}}[F(X_T) \mid X_t = x]$. When h is positive and differentiable, Doob's h-transform gives

$$u^\star(t,x) = \nabla_x \log h(t,x), \qquad v^\star(t,x) = G(t,x)^\top u^\star(t,x). \tag{4}$$

Here, $u^\star \in \mathbb{R}^n$ is the state-space covector and $v^\star \in \mathbb{R}^m$ is the executable Brownian-space control. Under rank deficiency, $v^\star$ remains executable even when instantaneous diffusion geometry cannot recover the full covector $u^\star$.

Finally, let $J^b_{t|s} := \partial X_t / \partial X_s$ denote the Jacobian of the reference stochastic flow along a shared Brownian path. Its stochastic variational equation, including state-dependent diffusion terms, is given in eq. (35).

4 Finite-Horizon Bridge Control

4.1 Control Consistency under Rank Deficiency

The key observation is that rank deficiency affects covector recovery, not the self-consistency identity itself.

Theorem 4.1 (Self-consistency)**.** *Under the regularity conditions stated in Appendix B.1, for every* $0 \le s < t < T$*,*

$$u^\star(s,x) = \mathbb{E}_{\mathbb{Q}^\star}\left[\left(J^b_{t|s}\right)^\top u^\star(t, X_t) \,\middle|\, X_s = x \right]. \tag{5}$$

No rank condition on G is required.

Thus, the obstruction lies in covector recovery rather than in self-consistency. If G has full row rank, the state-space covector is recovered pointwise from the executable control:

$$u = \left(GG^\top\right)^{-1} Gv. \tag{6}$$

When $\operatorname{rank} G < n$, however, $GG^\top$ is singular and components in $\ker G^\top$ are invisible to instantaneous diffusion geometry. FBC therefore targets the recovery step, replacing instantaneous covector recovery with finite-horizon recovery from dynamically transported noise directions.

4.2 Finite-Horizon Recovery

To expose directions that become identifiable through the dynamics, we split the drift as $b = a + c$ and define the auxiliary pullback transport

$$b = a + c, \qquad \dot\Phi^a_{t|s} = \nabla_x a(t, X_t)\Phi^a_{t|s}, \qquad \Phi^a_{s|s} = I, \qquad M^a_{s,t} = \left(\Phi^a_{t|s}\right)^{-1}. \tag{7}$$

We choose the auxiliary drift a so that it retains the drift couplings responsible for the leading controllability directions, while the remaining terms are assigned to c. For general nonlinear models, the choice can be guided by the local linearization $(\nabla_x b, G)$, retaining the lowest-order couplings that enlarge the reachable subspace from the directly forced directions; when no clear structure is available, the split remains model dependent. For underdamped dynamics with friction coefficient γ, choosing $a(q,p) = (p, -\gamma p)$ retains the kinetic coupling from momentum to position.

Let $t_k = k\Delta t$, $G_j := G(t_j, X_{t_j})$, and $v_j := v_{\theta^-}(t_j, X_{t_j})$, where θ^- denotes the frozen target-controller parameters. FBC pulls future Brownian directions back to the anchor t_k and aggregates them over a recovery window of length K_w, with ridge parameter $\varepsilon \ge 0$:

$$\begin{aligned} Z_{k,j} &= M^a_{t_k,t_j} G_j, \qquad j = k, \dots, k + K_w - 1, \\ W_k &= \sum_{j=k}^{k+K_w-1} \Delta t\, Z_{k,j} Z_{k,j}^\top, \qquad r_k = \sum_{j=k}^{k+K_w-1} \Delta t\, Z_{k,j} v_j, \\ \widehat{u}_k &= \left(W_k + \varepsilon I\right)^{-1} r_k. \end{aligned} \tag{8}$$

Equation (8) can be viewed as a finite-window ridge problem: transported Brownian directions provide multiple constraints on the anchor covector, allowing drift to reveal directions absent from the instantaneous range of G_k. With $d_w = mK_w \ll n$, the solve costs $O(nd_w^2 + d_w^3)$ operations and $O(nd_w + d_w^2)$ storage; nonlinear a additionally requires pathwise tangent propagation, whereas the kinetic auxiliary used here has constant $\nabla_x a$. Appendix A.3 gives implementation details.

This recovery map has a direct stochastic-geometric counterpart. Let $G_r := G(r, X_r)$ and $M^b_{s,r} := (J^b_{r|s})^{-1}$, and define

$$W^b_{s,t} = \int_s^t M^b_{s,r} G_r G_r^\top \left(M^b_{s,r}\right)^\top dr, \qquad C_{s,t} = \int_s^t J^b_{t|r} G_r G_r^\top \left(J^b_{t|r}\right)^\top dr. \tag{9}$$

Proposition 4.2 (Malliavin–Gramian identity)**.** *The pullback Gramian and Malliavin covariance satisfy*

$$C_{s,t} = J^b_{t|s} W^b_{s,t} \left(J^b_{t|s}\right)^\top, \qquad \operatorname{rank} C_{s,t} = \operatorname{rank} W^b_{s,t}. \tag{10}$$

Hence, finite-window identifiability and Malliavin nondegeneracy encode the same finite-time noise-propagation geometry (Kusuoka & Stroock, 1985; Norris, 1986; Arnaudon & Thalmaier, 2010), grounding finite-window recovery in classical stochastic accessibility.

For affine diffusions $dX_r = (A_r X_r + a_r)\, dr + G_r\, dW_r$ with deterministic time-dependent coefficients, let $\Phi_{r,s}$ be the state-transition matrix and define

$$Z_{s,r} = \Phi^{-1}_{r,s} G_r, \qquad W_{s,t} = \int_s^t Z_{s,r} Z_{s,r}^\top\, dr. \tag{11}$$

Theorem 4.3 (Affine finite-horizon recovery)**.** *Suppose the regularity conditions of Theorem 4.1 hold and the finite-window control moment below is conditionally integrable. If $W_{s,t} \succ 0$, then*

$$u^\star(s,x) = W^{-1}_{s,t} \mathbb{E}_{\mathbb{Q}^\star}\left[\int_s^t Z_{s,r} v^\star(r, X_r)\, dr \,\middle|\, X_s = x\right]. \tag{12}$$

For controllable affine diffusions with deterministic time-dependent coefficients and nonsingular finite-horizon Gramians, executable controls over a finite window are sufficient to recover the anchor covector exactly. Appendix B.2 gives the trajectory-level estimator and the singular-Gramian extension.

4.3 Finite-Horizon Control-Consistency Learning

Finite-horizon recovery provides covector estimates $\widehat{u}_k, \ldots, \widehat{u}_{k+K_c}$ over a short correction horizon of length K_c. Because these estimates live at different times, they must be transported to a common anchor before forming an executable control-consistency target. Define $J^a_\ell := \Phi^a_{t_{\ell+1}|t_\ell}$, $C_\ell := \nabla_x c(t_\ell, X_{t_\ell})$, and $\beta_\ell := (T - t_{\ell+1})/(T - t_\ell)$. Here, J^a_ℓ propagates covectors through the auxiliary dynamics, C_ℓ corrects for the residual drift omitted from that transport, and β_ℓ interpolates according to the remaining horizon.

Starting from $F_{k+K_c} = \widehat{u}_{k+K_c}$, FBC constructs the executable control target by the short backward recursion

$$F_\ell = (J^a_\ell)^\top \left[\beta_\ell F_{\ell+1} + (1-\beta_\ell)\widehat{u}_{\ell+1}\right] + \Delta t\, C_\ell^\top \widehat{u}_\ell, \qquad \widetilde{v}^{\mathrm{FBC}}_k = G_k^\top F_k. \tag{13}$$

Finite-horizon recovery therefore supplies identifiable state-space information, while the backward correction propagates it to the anchor before projection onto the executable Brownian-space control.

We parameterize the controller in residual form as $v_\theta(t,x) = v^g(t,x) + g_\theta(t,x)$, where v^g is the guide control and g_θ is the learned residual. At each outer iteration, we set $\theta^- := \mathrm{sg}(\theta)$, where $\mathrm{sg}(\cdot)$ denotes the stop-gradient operator, and generate detached controlled trajectories. For a set $\mathcal{A}$ of anchor indices, define the relaxed residual target for each $k \in \mathcal{A}$ and the regression objective by

$$\begin{aligned} \widetilde{g}_k = \widetilde{v}^{\mathrm{FBC}}_k - v^g(t_k, X^{\theta^-}_{t_k}), \qquad y_k = (1-\rho) g_{\theta^-}(t_k, X^{\theta^-}_{t_k}) + \rho\, \widetilde{g}_k, \\ \mathcal{L}_{\mathrm{FBC}}(\theta) = \mathbb{E}\left[\frac{1}{|\mathcal{A}|} \sum_{k \in \mathcal{A}} \left\| g_\theta(t_k, X^{\theta^-}_{t_k}) - \mathrm{sg}(y_k) \right\|^2\right]. \end{aligned} \tag{14}$$

Here $\rho \in (0,1]$ controls relaxation toward the newly recovered target. Minimizing $\mathcal{L}_{\mathrm{FBC}}$ therefore amounts to regression onto fixed control-consistency targets; gradients do not pass through trajectory simulation, finite-horizon recovery, or the backward recursion.

5 FINITE-TIME GEOMETRY AND CONSISTENCY

Finite-horizon recovery raises two questions: which state directions become identifiable over a finite window, and whether the recovered information remains consistent after projection to the executable Brownian-space control. We first analyze finite-time identifiability and then nonlinear executable consistency. A short-horizon uniqueness result for the ideal self-consistency fixed point is given in Appendix B.6.

5.1 CONTROLLABILITY GEOMETRY

For a local linearization with drift and diffusion coefficients frozen at the anchor, let A and G denote the resulting drift and diffusion matrices, and define the finite-horizon pullback Gramian

$$W_h^-(A,G) = \int_0^h e^{-A\tau} G G^\top e^{-A^\top \tau}\, d\tau. \tag{15}$$

Its nondegeneracy is characterized exactly by the Kalman controllability condition.

Theorem 5.1 (Finite-time controllability)**.** *For every $h > 0$,*

$$W_h^-(A,G) \succ 0 \quad \Longleftrightarrow \quad \operatorname{rank}\left[G,\, AG,\, \ldots,\, A^{n-1}G\right] = n. \tag{16}$$

Thus, finite-time drift propagation can recover directions absent from the instantaneous range of G (Kalman, 1960). Nondegeneracy alone, however, does not quantify how strongly different directions are identified over short windows. Define the Kalman filtration

$$\mathcal{R}_0 = \{0\}, \qquad \mathcal{R}_j = \operatorname{span}\left\{G,\, AG,\, \ldots,\, A^{j-1}G\right\}, \qquad j \geq 1. \tag{17}$$

The Gramian scale associated with the jth controllability layer $\mathcal{R}_j/\mathcal{R}_{j-1}$ is $\Theta(h^{2j-1})$ as $h \to 0$; the precise statement is given in Appendix B.3. For underdamped Langevin dynamics, the directly forced momentum layer therefore scales as $\Theta(h)$, whereas the position layer generated through one drift coupling scales as $\Theta(h^3)$. Deeper controllability chains yield the hierarchy $h, h^3, h^5, \ldots$; finite-time controllability determines *which* directions are identifiable, while the layer scaling determines how strongly they are identified.

5.2 EXECUTABLE CONSISTENCY

Under rank deficiency, the relevant consistency target is the executable projection $v^\star(s,x) = G_s^\top u^\star(s,x)$, where $G_s := G(s,x)$. Full covector recovery is therefore unnecessary: only components of $u^\star$ visible through $G_s^\top$ affect the executable controller. Fix a preterminal anchor (s,x) and let $X^\star$ denote the Doob bridge conditioned on $X_s^\star = x$. With

$$Z_{s,r}^a = M_{s,r}^a G_r, \qquad G_r := G(r, X_r^\star), \tag{18}$$

define the conditional population Gramian and control moment

$$\begin{aligned} \overline{W}_h^a &= \mathbb{E}_{\mathbb{Q}^\star}\left[\int_s^{s+h} Z_{s,r}^a (Z_{s,r}^a)^\top\, dr \,\middle|\, X_s^\star = x\right], \\ \overline{r}_h^a &= \mathbb{E}_{\mathbb{Q}^\star}\left[\int_s^{s+h} Z_{s,r}^a v^\star(r, X_r^\star)\, dr \,\middle|\, X_s^\star = x\right]. \end{aligned} \tag{19}$$

Define the recovery mismatch and its amplification in executable directions by

$$b_h^a = \overline{r}_h^a - \overline{W}_h^a u^\star(s,x), \qquad \Gamma_s(W) = \left\| W^{\dagger/2} G_s \right\|_{\mathrm{op}}^2. \tag{20}$$

Here $W^\dagger$ denotes the Moore–Penrose pseudoinverse and $W^{\dagger/2} := (W^\dagger)^{1/2}$. The quantity b_h^a is the recovery mismatch, while Γ_s quantifies its amplification by pseudoinversion after projection to the executable control.

For $\varepsilon_h \geq 0$, define the corresponding population FBC control estimate

$$\overline{v}_{h,\varepsilon_h}^a = G_s^\top \left(\overline{W}_h^a + \varepsilon_h I\right)^{-1} \overline{r}_h^a, \tag{21}$$

using the Moore–Penrose inverse when $\varepsilon_h = 0$. Theorem 5.2 uses conditional population quantities, whereas Algorithm 1 inverts pathwise Gramians before regression; expectation and inversion do not commute in general (Appendix B.7.2).

Theorem 5.2 (Executable consistency)**.** *Under the regularity conditions in Appendix B.4, suppose that* $\mathrm{Range}(G_s) \subseteq \mathrm{Range}(\overline{W}_h^a)$ *for all sufficiently small* h*, and for some* $\alpha, \zeta > 0$*,*

$$\Gamma_s(\overline{W}_h^a) = O(h^{-\alpha}), \qquad \left\| (\overline{W}_h^a)^{\dagger/2} b_h^a \right\| = O(h^{\zeta}). \tag{22}$$

Then

$$\left\| \overline{v}_{h,\varepsilon_h}^a - v^\star(s,x) \right\| = O\left(h^{\zeta - \alpha/2} + \sqrt{\frac{\varepsilon_h}{h^\alpha}} \right). \tag{23}$$

In particular, if $\zeta > \alpha/2$ *and* $\varepsilon_h = o(h^\alpha)$*, then* $\overline{v}_{h,\varepsilon_h}^a \to v^\star(s,x)$ *as* $h \to 0$*.*

The theorem reduces executable consistency to a balance between the normalized recovery mismatch and its amplification in executable directions. Additional exact-tangent bounds are given in Appendix B.4.

For the kinetic auxiliary transport used by FBC, the consistency conditions above can be verified explicitly. For underdamped Langevin dynamics with $G = (0, \sigma I)^\top$ and constant $\sigma \neq 0$, take

$$a(q,p) = (p, -\gamma p), \qquad c(q,p) = (0, -\nabla V(q)), \qquad \nabla c\, G = 0. \tag{24}$$

The last identity follows because the Brownian noise acts only on momentum, so the omitted force curvature does not act directly on the Brownian columns.

Corollary 5.3 (Kinetic consistency)**.** *Under the bounded-curvature and Doob-regularity conditions in Appendix B.5, there exists* $C > 0$*, independent of* h*, such that*

$$\left\| \overline{v}_{h,\varepsilon_h}^a - v^\star(s,x) \right\| \le C \left(h^{1/2} + \sqrt{\frac{\varepsilon_h}{h}} \right). \tag{25}$$

Hence, if $\varepsilon_h = o(h)$*, then* $\overline{v}_{h,\varepsilon_h}^a \to v^\star(s,x)$ *as* $h \to 0$*.*

Indeed, the kinetic scaling from section 5.1 gives $\Gamma_s(\overline{W}_h^a) = O(h^{-1})$, while $\nabla c\, G = 0$ yields a Gramian-normalized residual of order $O(h)$. Thus $\alpha = \zeta = 1$, giving the $O(h^{1/2})$ executable-control rate. Backward-recursion bounds are given in Appendix B.7.4.

6 Experiments

Our experiments ask whether finite-time transport can recover control information lost under instantaneous diffusion geometry. We isolate this mechanism under continuous rank loss and an analytically tractable kinetic OU bridge, then evaluate FBC on strictly rank-deficient nonlinear dynamics, including underdamped Double-Well, Müller–Brown, and constrained alanine. Canonical full-rank benchmarks assess the FBC learning construction in the standard control-consistency setting. Additional experimental details and supplementary analyses are reported in Appendix C.

6.1 Continuous Rank Loss

We continuously degrade the instantaneous diffusion rank to isolate the effect of finite-horizon recovery. Let $G_\eta = \mathrm{diag}(\eta, 1)$ with $\eta \in \{1, 0.3, 0.1, 0.03, 0.01, 0\}$, while keeping the linear drift, terminal observation, model class, and evaluation objective fixed. For every $\eta > 0$, exact pointwise lifting remains within 0.52% of the Doob oracle. At the singular limit $\eta = 0$, however, the pseudoinverse continuation deteriorates to 6.316 ± 0.171 times the oracle objective, whereas finite-horizon recovery attains 1.076 ± 0.003. Figure 1 summarizes the corresponding objective behavior and finite-window Gramian nondegeneracy across the rank-loss sweep.

At exact rank loss, we separate temporal pooling from dynamical transport. An identity window remains rank one and improves the objective ratio to 1.110 ± 0.003, reflecting better use of information already available in the instantaneous subspace. Transporting a single future noise direction raises the lift rank to two and reduces the ratio to 1.071 ± 0.002. The full finite-window construction also has rank two, achieves 1.076 ± 0.003, and gives the strongest terminal-law metrics. The rank change distinguishes temporal pooling from the recovery of the missing state direction by drift transport. The full decomposition is reported in Appendix C.2.3.

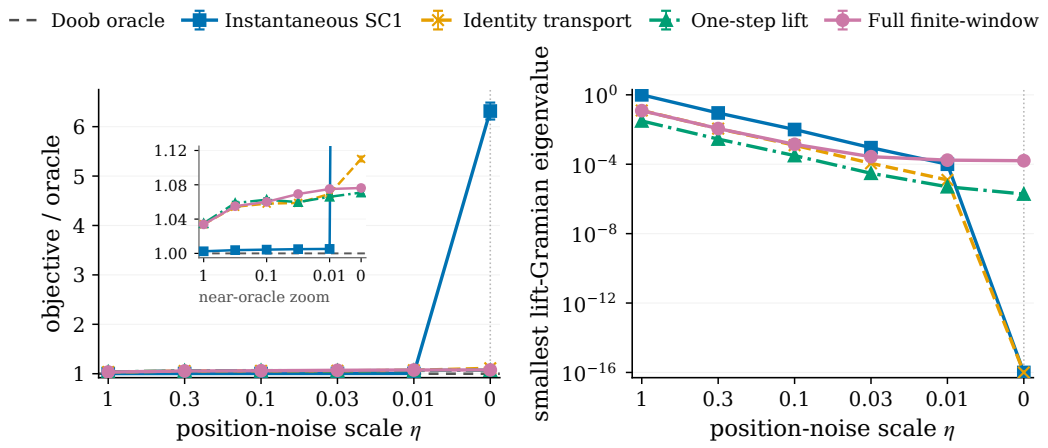
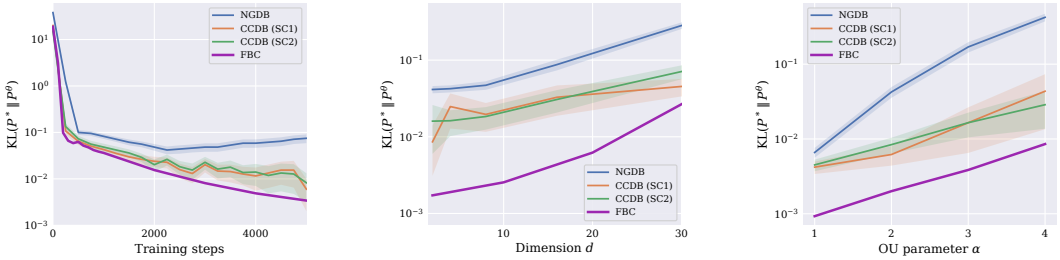


Figure 1: Continuous rank loss in kinetic OU. *Left:* objective relative to the Doob oracle. *Right:* smallest eigenvalue of the lift Gramian, showing finite-window nondegeneracy under rank loss.



(a) Training convergence, $d = 2$, $\alpha = 2$.

(b) Effect of dimension, $\alpha = 2$.

(c) Effect of OU parameter, $d = 2$.

Figure 2: Full-rank Ornstein–Uhlenbeck bridge results across training progress, state dimension, and drift parameter α.

6.2 Ornstein–Uhlenbeck Bridges

Full-rank OU. Ornstein–Uhlenbeck (OU) bridges serve as a canonical full-rank benchmark with an analytically tractable conditioned process (Heng et al., 2025; Yang et al., 2025; Howard et al., 2026). Figure 2 compares FBC with NGDB and CCDB across training progress, state dimension, and drift strength. FBC achieves the lowest observed KL in all three comparisons. Since this regime uses pointwise recovery, these results assess the overall FBC learning construction rather than finite-horizon rank completion; the contribution of the latter is isolated in Section 6.1.

Kinetic OU. We next consider exact rank deficiency while retaining an analytically tractable bridge by augmenting the state with momentum and restricting noise and control to the momentum coordinate:

$$dq_t = p_t\,dt, \qquad dp_t = (-q_t - p_t + v_t)\,dt + dW_t. \tag{26}$$

The instantaneous diffusion has rank one, but the two-dimensional system is controllable through the kinetic coupling. FBC attains an objective of 2.905 ± 0.032, close to the exact Doob value 2.701 ± 0.035, with terminal q-RMSE 0.220 ± 0.003 versus 0.216 ± 0.002. The best instantaneous ridge lift reaches 20.53 ± 1.05, whereas FBC remains close to the analytic bridge under exact rank deficiency. Additional results and analyses are given in Appendix C.2.1.

6.3 Cox–Ingersoll–Ross Bridge

The Cox–Ingersoll–Ross (CIR) process complements the constant-diffusion OU benchmark by evaluating FBC under state-dependent diffusion with a tractable transition law (Cox et al., 1985; Howard et al., 2026). We use

$$dX_t = a(b - X_t)\,dt + \sigma_{\mathrm{CIR}}\sqrt{X_t}\,dW_t$$

Table 1: Quantitative results on the full-rank benchmarks (mean ± std over five random seeds). ● denotes strong performance, ▲ denotes reasonable performance, and ■ denotes failure. The starred SDB value is unreliable because an extreme endpoint control spike causes the numerical time integration used for the KL estimate to underestimate its value.

Method	CIR $\mathrm{KL}(P^\star\|P^\theta)\downarrow$	Double-Well, $v=3.0$ $\mathrm{KL}(P^\star\|P^\theta)\downarrow$	Cell Diffusion, $\mathrm{KL}(P^\theta\|P)\downarrow$		
			Normal	Rare	Multimodal
FBC (Ours)	$\mathbf{0.009 \pm 0.002}$	$\mathbf{0.010 \pm 0.003}$	6.52 ± 0.05	$\mathbf{65.3 \pm 0.2}$	792.4 ± 0.9
CCDB (SC1)	0.023 ± 0.010 ●	0.148 ± 0.073 ●/▲	6.53 ± 0.14 ●	66.1 ± 0.9 ●	792.8 ± 0.9 ●/▲
CCDB (SC2)	0.055 ± 0.009 ●/▲	0.041 ± 0.013 ●	6.53 ± 0.17 ●	66.0 ± 0.5 ●	792.4 ± 1.0 ●/▲
NGDB	0.051 ± 0.012 ●/▲	0.098 ± 0.047 ●/▲	$\mathbf{6.48 \pm 0.15}$ ●	65.5 ± 0.6 ●	$\mathbf{783.4 \pm 0.8}$ ●
SDB	0.104 ± 0.011 ▲	0.229 ± 0.024 ▲	$5.56 \pm 0.39^*$	75.0 ± 5.1 ▲	1845 ± 230 ■
FB	0.074 ± 0.011 ▲	–	6.51 ± 0.14 ●	237 ± 98 ■	4836 ± 954 ■
BEL	0.044 ± 0.004 ●	0.846 ± 0.324 ■	6.93 ± 0.07 ●/▲	94.7 ± 0.5 ▲	869.3 ± 3.2 ▲

and evaluate $\mathrm{KL}(P^\star\|P^\theta)$ against the exact bridge. FBC achieves 0.009 ± 0.002, the lowest KL among the methods reported in Table 1.

6.4 Double-Well Bridges

Overdamped Double-Well. The standard full-rank Double-Well bridge has quartic potential $U(x) = v(x^2-1)^2$ (Pidstrigach et al., 2025; Howard et al., 2026). At $v=3$, FBC achieves $\mathrm{KL}(\mathbb{P}^\star\|\mathbb{P}^\theta) = 0.010 \pm 0.003$, the lowest value among the methods reported in Table 1.

Underdamped Double-Well. We retain the same qualitative two-well transition geometry but change the diffusion geometry by augmenting the state with momentum and restricting noise and control to the momentum coordinate:

$$dq_t = p_t\,dt, \qquad dp_t = [-V'_{\mathrm{DW}}(q_t) - p_t + v_t]\,dt + dW_t, \qquad V_{\mathrm{DW}}(q) = 2(q^2-1)^2. \tag{27}$$

The resulting diffusion is instantaneously rank deficient, while the position remains reachable through the kinetic coupling. Starting from $(q_0, p_0) = (-1, 0)$ and targeting $q_{\mathrm{target}} = 1$ over $T = 2$, FBC improves the stochastic-control objective and terminal metrics relative to NGDB (Table 3). The complete objective decomposition and additional results are reported in Appendix C.4.

6.5 Cell Diffusion

Cell Diffusion complements the analytically tractable benchmarks with a two-dimensional non-gradient differentiation model lacking an exact bridge solution (Wang et al., 2011). Following prior diffusion-bridge studies (Heng et al., 2025; Baker et al., 2025; Yang et al., 2025; Howard et al., 2026), we consider the Normal, Rare, and Multimodal regimes and report $\mathrm{KL}(P^\theta\|P)$. FBC achieves the lowest KL in the Rare regime while remaining competitive on Normal and Multimodal (Table 1).

6.6 Transition Path Sampling

We evaluate FBC on nonlinear transition-path problems using the Müller–Brown and alanine dipeptide benchmarks studied in recent stochastic-control and diffusion-path work (Du et al., 2024; Holdijk et al., 2023; Seong et al., 2025; Howard et al., 2026).

Müller–Brown. On the standard full-rank Müller–Brown benchmark, FBC achieves 25.1 ± 0.1 KL, compared with 27.8 ± 1.7 for CCDB (SC2), 28.6 ± 0.4 for NGDB, and 38.6 ± 0.2 for Doob's Lagrangian (Table 2). We consider an underdamped Müller–Brown system with strict rank deficiency by augmenting the state with momentum and restricting stochastic forcing and control to momentum:

$$dq_t = p_t\,dt, \qquad dp_t = [-\nabla V_{\mathrm{MB}}(q_t) - p_t + v_t]\,dt + dW_t. \tag{28}$$

Here $q, p \in \mathbb{R}^2$ and $T = 2.5$. Under this rank-deficient dynamics, FBC reduces the stochastic-control objective by approximately 12.5% relative to NGDB and simultaneously improves terminal q-RMSE, hit rate, and basin-entry rate (Table 3). Thus, the benefit of finite-horizon recovery persists beyond analytically tractable kinetic models on a nonlinear transition landscape.

Table 2: Transition-path results on Müller–Brown and alanine dipeptide. For underdamped alanine, FBC and NGDB use the matched constrained-underdamped protocol described in Appendix C.6.3. The PIPS and TPS–DPS values are previously reported results from (Seong et al., 2025).

(a) Müller–Brown

Method	KL ↓
FBC (Ours)	$\mathbf{25.1 \pm 0.1}$
CCDB (SC2)	27.8 ± 1.7
NGDB	28.6 ± 0.4
DL	38.6 ± 0.2

(b) Alanine C5→C7ax

Overdamped

Method	ETS (kJ mol^{-1}) ↓
FBC (Ours)	$\mathbf{16.87 \pm 6.56}$
CCDB (SC2)	28.70 ± 5.80
NGDB	26.50 ± 6.80

Underdamped

Method	ETS (kJ mol^{-1}) ↓
FBC (Ours)	$\mathbf{6.07 \pm 0.03}$
NGDB	61.05 ± 3.53
PIPS	28.20 ± 10.90
TPS–DPS	18.40 ± 10.90

Table 3: Results for the strictly rank-deficient underdamped nonlinear benchmarks (mean ± std over five random seeds). Basin uses the task-specific success definition described in Appendix C.4 and Appendix C.5.

Method	Objective ↓	q-RMSE ↓	Hit ↑	Basin ↑
Underdamped Double-Well				
FBC (Ours)	$\mathbf{8.2455 \pm 0.0532}$	$\mathbf{0.2288 \pm 0.0068}$	$\mathbf{0.9189 \pm 0.0071}$	$\mathbf{1.00000 \pm 0.00000}$
NGDB	8.8493 ± 0.1422	0.2614 ± 0.0640	0.8684 ± 0.1089	0.99988 ± 0.00021
Underdamped Müller–Brown				
FBC (Ours)	$\mathbf{9.6200 \pm 1.1659}$	$\mathbf{0.2381 \pm 0.0178}$	$\mathbf{0.9436 \pm 0.0327}$	$\mathbf{0.9404 \pm 0.0505}$
NGDB	10.9918 ± 0.3450	0.2621 ± 0.0354	0.8954 ± 0.0726	0.8746 ± 0.1044

Alanine Dipeptide. On the 22-atom C5→C7ax transition, FBC obtains $16.87 \pm 6.56\,\mathrm{kJ\,mol^{-1}}$ in the overdamped benchmark. In the constrained-underdamped setting, noise and control span 45 tangent-velocity directions in a 90-dimensional state, while the finite-window lift reaches numerical recovery rank 90. FBC obtains $6.07 \pm 0.03\,\mathrm{kJ\,mol^{-1}}$, compared with $61.05 \pm 3.53\,\mathrm{kJ\,mol^{-1}}$ for NGDB. Learned transition paths for Müller–Brown and alanine are shown in Figure 3. Additional results and controllability analyses are given in Appendices C.6.3 and C.7.2.

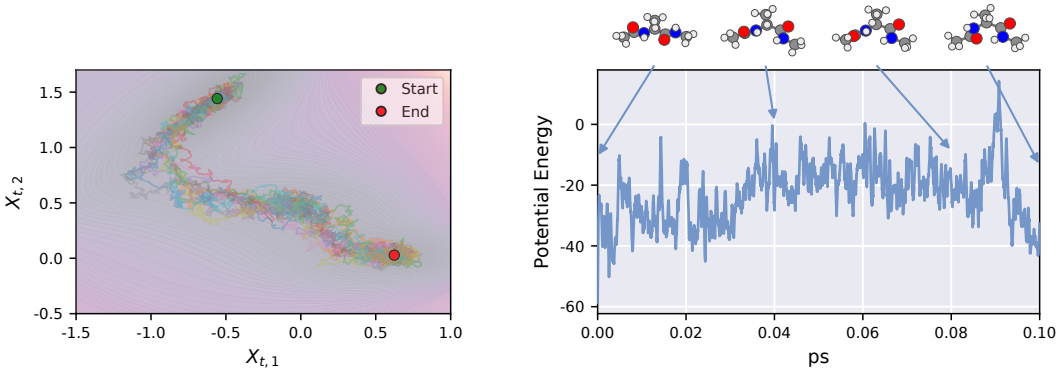


Figure 3: Learned transition paths. Left: FBC trajectories on the Müller–Brown potential. Right: transition-state energy profile along an alanine C5→C7ax trajectory, with conformations shown at four times.

7 Conclusion

Under rank deficiency, FBC replaces singular instantaneous covector recovery with a finite-horizon lift of dynamically transported noise directions. We established rank-free self-consistency, the Malliavin–Gramian connection, exact recovery for controllable affine diffusions with deterministic time-dependent coefficients, small-window controllability-layer scaling, and nonlinear executable consistency, including an explicit short-window executable-control convergence rate for underdamped Langevin dynamics. Experiments spanning continuous rank loss, nonlinear rank-deficient dynamics, and constrained alanine show that FBC recovers control information unavailable to pointwise geometry, supporting finite-time controllability as the relevant recovery geometry.

REFERENCES

Alexis Arnaudon, Frank van der Meulen, Moritz Schauer, and Stefan Sommer. Diffusion bridges for stochastic hamiltonian systems and shape evolutions. *SIAM Journal on Imaging Sciences*, 15(1): 293–323, 2022.

Marc Arnaudon and Anton Thalmaier. The differentiation of hypoelliptic diffusion semigroups. *Illinois Journal of Mathematics*, 54(4):1285–1311, 2010.

Elizabeth Louise Baker, Moritz Schauer, and Stefan Sommer. Score matching for bridges without learning time-reversals. In Yingzhen Li, Stephan Mandt, Shipra Agrawal, and Emtiyaz Khan (eds.), *Proceedings of The 28th International Conference on Artificial Intelligence and Statistics*, volume 258 of *Proceedings of Machine Learning Research*, pp. 775–783. PMLR, 03–05 May 2025.

Fabrice Baudoin. Conditioned stochastic differential equations: theory, examples and application to finance. *Stochastic Processes and their Applications*, 100(1-2):109–145, 2002.

Joris Bierkens, Frank Van Der Meulen, and Moritz Schauer. Simulation of elliptic and hypo-elliptic conditional diffusions. *Advances in Applied Probability*, 52(1):173–212, 2020.

Denis Blessing, Julius Berner, Lorenz Richter, and Gerhard Neumann. Underdamped diffusion bridges with applications to sampling. In *International Conference on Learning Representations*, volume 2025, pp. 2970–3002, 2025.

John C Cox, Jr Ingersoll, Jonathan E, and Stephen A Ross. A theory of the term structure of interest rates. *Econometrica*, 53(2):385–407, March 1985.

Bernard Delyon and Ying Hu. Simulation of conditioned diffusion and application to parameter estimation. *Stochastic Processes and their Applications*, 116(11):1660–1675, 2006.

Carles Domingo-Enrich, Jiequn Han, Brandon Amos, Joan Bruna, and Ricky T Chen. Stochastic optimal control matching. *Advances in Neural Information Processing Systems*, 37:112459–112504, 2024.

Carles Domingo-Enrich, Michal Drozdzal, Brian Karrer, and Ricky T. Q. Chen. Adjoint matching: Fine-tuning flow and diffusion generative models with memoryless stochastic optimal control. In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net, 2025.

Yuanqi Du, Michael Plainer, Rob Brekelmans, Chenru Duan, Frank Noe, Carla P Gomes, Alan Aspuru-Guzik, and Kirill Neklyudov. Doob's lagrangian: A sample-efficient variational approach to transition path sampling. *Advances in Neural Information Processing Systems*, 37:65791–65822, 2024.

Aaron J Havens, Benjamin Kurt Miller, Bing Yan, Carles Domingo-Enrich, Anuroop Sriram, Daniel S. Levine, Brandon M Wood, Bin Hu, Brandon Amos, Brian Karrer, Xiang Fu, Guan-Horng Liu, and Ricky T. Q. Chen. Adjoint sampling: Highly scalable diffusion samplers via adjoint matching. In Aarti Singh, Maryam Fazel, Daniel Hsu, Simon Lacoste-Julien, Felix Berkenkamp, Tegan Maharaj, Kiri Wagstaff, and Jerry Zhu (eds.), *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 22204–22237. PMLR, 13–19 Jul 2025.

Jeremy Heng, Valentin De Bortoli, Arnaud Doucet, and James Thornton. Simulating diffusion bridges with score matching. *Biometrika*, 112(4):asaf048, 2025.

Lars Holdijk, Yuanqi Du, Ferry Hooft, Priyank Jaini, Berend Ensing, and Max Welling. Stochastic optimal control for collective variable free sampling of molecular transition paths. In *Advances in Neural Information Processing Systems*, volume 36, pp. 79540–79556. Curran Associates, Inc., 2023. doi: 10.52202/075280-3481.

Samuel Howard, Nikolas Nüsken, and Jakiw Pidstrigach. Control consistency losses for diffusion bridges. In *Proceedings of the 43rd International Conference on Machine Learning*, volume 306 of *Proceedings of Machine Learning Research*. PMLR, 2026.

Rudolf Emil Kalman. Contributions to the theory of optimal control. *Bol. soc. mat. mexicana*, 5: 102–119, 1960.

Shigeo Kusuoka and Daniel W. Stroock. Applications of the malliavin calculus. ii. *Journal of the Faculty of Science, the University of Tokyo. Sect. 1 A, Mathematics*, 32:1–76, 1985.

Xue-Mei Li. On hypoelliptic bridge. *Electronic Communications in Probability*, 21(24):1–12, 2016. doi: 10.1214/16-ECP4646.

James Norris. Simplified malliavin calculus. In Jacques Azéma and Marc Yor (eds.), *Séminaire de Probabilités XX 1984/85*, pp. 101–130. Springer Berlin Heidelberg, Berlin, Heidelberg, 1986. ISBN 978-3-540-39860-8.

Jakiw Pidstrigach, Elizabeth Louise Baker, Carles Domingo-Enrich, George Deligiannidis, and Nikolas Nüsken. Conditioning diffusions using malliavin calculus. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 49292–49315. PMLR, 13–19 Jul 2025.

Moritz Schauer, Frank Van Der Meulen, and Harry Van Zanten. Guided proposals for simulating multi-dimensional diffusion bridges. *Bernoulli*, 23(4A):2917–2950, 2017.

Kiyoung Seong, Seonghyun Park, Seonghwan Kim, Woo Youn Kim, and Sungsoo Ahn. Transition path sampling with improved off-policy training of diffusion path samplers. In Y. Yue, A. Garg, N. Peng, F. Sha, and R. Yu (eds.), *International Conference on Learning Representations*, volume 2025, pp. 93040–93062, 2025.

Jin Wang, Kun Zhang, Li Xu, and Erkang Wang. Quantifying the waddington landscape and biological paths for development and differentiation. *Proceedings of the National Academy of Sciences*, 108 (20):8257–8262, 2011.

Gefan Yang, Frank Van Der Meulen, and Stefan Sommer. Neural guided diffusion bridges. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 71210–71230. PMLR, 13–19 Jul 2025.

AI USE STATEMENT

Generative AI tools were used for assistance with manuscript editing, literature organization, and implementation debugging. All theoretical developments, experiments, and final claims were independently verified by the authors.

A METHOD AND IMPLEMENTATION DETAILS

This appendix collects notation and implementation details that are omitted from the main text. Section A.1 summarizes the discrete FBC quantities, section A.2 describes the full-row-rank and state-dependent-diffusion cases, section A.3 gives the finite-horizon linear algebra and numerical stabilization details, and section A.4 summarizes the complete training procedure.

A.1 NOTATION AND DISCRETE CONVENTIONS

We use the time grid $t_k = k\Delta t$. A subscript k denotes evaluation along the detached controlled trajectory at (t_k, X_{t_k}). Table 4 summarizes the notation used throughout the discrete FBC construction.

Table 4: Notation used in the FBC implementation.

Symbol	Meaning
$X_t \in \mathbb{R}^n$, $W_t \in \mathbb{R}^m$	State and Brownian motion
$G_k \in \mathbb{R}^{n\times m}$	Diffusion/control matrix at t_k
$v_k \in \mathbb{R}^m$	Executable Brownian-space control
$\Phi^a_{t\|s}$, $M^a_{s,t}$	Auxiliary tangent flow and pullback transport
$Z_{k,j} = M^a_{t_k,t_j} G_j$	Future Brownian direction transported to t_k
W_k, r_k	Finite-window Gramian and control moment
$\widehat{u}_k$	Recovered finite-horizon covector
F_k	Backward-corrected covector target
K_w, K_c	Recovery-window and correction-horizon lengths
ε	Ridge regularization of the Gramian solve
ρ	Relaxation toward the new FBC target
$\mathrm{sg}(\cdot)$	Stop-gradient operator

A.2 FULL-RANK RECOVERY AND STATE-DEPENDENT DIFFUSION

When G_k has full row rank, the state covector is recovered pointwise as

$$\widehat{u}_k = (G_k G_k^\top)^{-1} G_k v_k. \tag{29}$$

Finite-window rank completion is therefore unnecessary in the full-row-rank case.

For state-dependent diffusion, however, the consistency recursion must retain the dependence of the one-step tangent on the realized Brownian increment. In the full-row-rank state-dependent setting, we use the generalized control-consistency construction with a zero-drift auxiliary dynamics. Suppose a detached controlled trajectory is simulated using Brownian increment $\Delta W_k^{\mathbb{Q}} \sim \mathcal{N}(0, \Delta t\, I)$ over $[t_k, t_{k+1}]$. Here $\Delta W_k^{\mathbb{Q}}$ denotes the realized noise increment used by that controlled rollout. The corresponding auxiliary Euler tangent is

$$J_k^a = \partial_x \left[x + G(t_k, x)\Delta W_k^{\mathbb{Q}} \right]\Big|_{x=X_{t_k}}. \tag{30}$$

Let

$$H_k = \beta_k F_{k+1} + (1-\beta_k)\widehat{u}_{k+1}. \tag{31}$$

The full-row-rank consistency recursion is then

$$F_k = (J_k^a)^\top H_k + \Delta t\, \nabla_x b(t_k, X_{t_k})^\top \widehat{u}_k. \tag{32}$$

Let K_{rec} denote the last recursion index at which the recovered covector is defined. We initialize $F_{K_{\mathrm{rec}}} = \widehat{u}_{K_{\mathrm{rec}}}$ and apply eq. (32) backward for $k = K_{\mathrm{rec}} - 1, \ldots, 0$, or only as far as required by the

selected training anchors. For hard endpoint bridges, $t_{K_{\mathrm{rec}}} < T$. Here J_k^a is the zero-drift auxiliary tangent used to construct the training target, rather than the reference-flow Jacobian J_k^b appearing in Theorem 4.1. The drift omitted from the auxiliary tangent therefore reappears through the running correction in eq. (32).

Equations (30) and (32) define the state-dependent full-row-rank construction used for CIR. When G is constant, $J_k^a = I$, and the construction reduces to the corresponding full-horizon control-consistency recursion used in the constant-diffusion full-row-rank experiments.

For comparison, the exact reference stochastic-flow Jacobian used in the theory satisfies the state-dependent variational equation in eq. (35). The distinction is intentional: J^b characterizes the exact self-consistency identity, whereas J^a is the auxiliary tangent used to construct the executable training target.

A.3 Finite-Horizon Linear Algebra and Numerical Choices

Window-space formulation. The finite-horizon recovery in eq. (8) can be evaluated without explicitly factorizing the $n \times n$ Gramian. For an anchor t_k, define

$$\begin{aligned} U_k &= \left[\sqrt{\Delta t}\, Z_{k,k}, \ldots, \sqrt{\Delta t}\, Z_{k,k+K_w-1}\right] \in \mathbb{R}^{n\times d_w}, \\ \mathbf{v}_k^{\mathrm{win}} &= \begin{bmatrix} \sqrt{\Delta t}\, v_k \\ \vdots \\ \sqrt{\Delta t}\, v_{k+K_w-1} \end{bmatrix} \in \mathbb{R}^{d_w}, \qquad d_w = mK_w. \end{aligned} \tag{33}$$

Then $W_k = U_k U_k^\top$ and $r_k = U_k \mathbf{v}_k^{\mathrm{win}}$, and for $\varepsilon > 0$ the recovery admits the equivalent form

$$\widehat{u}_k = (W_k + \varepsilon I)^{-1} r_k = U_k \left(U_k^\top U_k + \varepsilon I\right)^{-1} \mathbf{v}_k^{\mathrm{win}}. \tag{34}$$

Thus, when $d_w = mK_w \ll n$, the recovery can be computed through a d_w-dimensional linear system rather than an n-dimensional one. Once U_k is available, forming $U_k^\top U_k$ costs $O(nd_w^2)$ operations, the dense solve costs $O(d_w^3)$, and the required storage is $O(nd_w + d_w^2)$. The cost of constructing U_k depends on the tangent implementation: dense Jacobians require matrix–block products, whereas Jacobian–vector or Jacobian–block products avoid materializing the full $n \times n$ tangent.

Numerical stability. For $\varepsilon = 0$, recovery is interpreted on the numerically supported subspace using the Moore–Penrose inverse. We monitor the numerical rank of the finite-window Gramian together with its spectrum and the regularized condition number

$$\kappa_\varepsilon(W_k) = \frac{\lambda_{\max}(W_k) + \varepsilon}{\lambda_{\min}(W_k) + \varepsilon}.$$

These quantities distinguish two separate failure modes: insufficient finite-window rank indicates that the selected transported directions have not exposed the required controllability subspace, whereas a large $\kappa_\varepsilon(W_k)$ indicates weakly identified directions within an already recovered span. The numerical-rank convention used in the experiments is given in eq. (99).

Recovery, correction, and regularization. The recovery-window length K_w and correction horizon K_c play different roles. The physical recovery duration $K_w \Delta t$ determines which transported Brownian directions enter the Gramian, whereas $H_c = K_c \Delta t$ determines how far the recovered covectors are propagated by the backward recursion in eq. (13). Increasing K_w can enlarge the identifiable span only when the additional transported directions expose a new controllability layer; otherwise it only adds information within an already observed span. The ridge parameter ε acts after this geometric construction and stabilizes weak finite-window directions without changing $\mathrm{Range}(U_k)$. Its asymptotic role is analyzed in Appendix B.4. Finally, the relaxation parameter $\rho \in (0, 1]$ controls the outer update toward the detached FBC target defined in section 4.3. Thus, K_w determines recovery geometry, K_c determines local backward propagation, ε controls inversion stability, and ρ controls the outer fixed-point update.

Choice of auxiliary transport. The auxiliary split $b = a + c$ is chosen so that a retains the drift couplings responsible for the leading controllability layers, while c contains the remaining drift contribution handled by the local correction. For the underdamped systems considered in this work, we use

$$a(q, p) = (p, -\gamma p), \qquad c(q, p) = (0, -\nabla V(q)).$$

This split preserves the kinetic momentum-to-position coupling in the auxiliary transport. Moreover, $\nabla_c G = 0$ because the Brownian columns lie entirely in momentum, giving the structural cancellation used in the kinetic consistency analysis of Appendix B.5.

A.4 Training Algorithm

Algorithm 1 summarizes the complete training procedure. In the full-row-rank setting, state covectors are recovered pointwise and propagated using the consistency construction of Appendix A.2. Under rank deficiency, each required covector is instead reconstructed from a finite window of dynamically transported Brownian directions and then propagated over the short correction horizon using eq. (13). Both constructions ultimately produce a detached Brownian-space target for regression of the residual controller.

Algorithm 1 Finite-Horizon Bridge Control training

Require: guide v^g, residual controller g_θ, time grid $\{t_k\}_{k=0}^K$, training anchors $\mathcal{A}$, relaxation ρ; for rank-deficient recovery, window length K_w, correction horizon K_c, ridge parameter ε, and anchors satisfying $k + K_c + K_w - 1 \le K$, with strict inequality for hard endpoint bridges

for each outer iteration **do**
  Freeze the target controller: $\theta^- \leftarrow \mathrm{sg}(\theta)$
  Simulate detached controlled trajectories under $v_{\theta^-} = v^g + g_{\theta^-}$
  Evaluate detached controls $v_j = v_{\theta^-}(t_j, X_{t_j}^{\theta^-})$ at the required time points
  if the modeled diffusion is full row rank on the relevant state region **then**
    Recover the required covectors pointwise: $\widehat{u}_j = (G_j G_j^\top)^{-1} G_j v_j$
    Apply the full-row-rank consistency recursion of Appendix A.2 to obtain F_k for $k \in \mathcal{A}$
  else
    for each anchor $k \in \mathcal{A}$ **do**
      for $\ell = k, \dots, k + K_c$ **do**
        Transport the recovery window to t_ℓ: $Z_{\ell,j} = M^a_{t_\ell, t_j} G_j$, $j = \ell, \dots, \ell + K_w - 1$
        Form $W_\ell = \sum_{j=\ell}^{\ell+K_w-1} \Delta t\, Z_{\ell,j} Z_{\ell,j}^\top$ and $r_\ell = \sum_{j=\ell}^{\ell+K_w-1} \Delta t\, Z_{\ell,j} v_j$
        Recover $\widehat{u}_\ell = (W_\ell + \varepsilon I)^{-1} r_\ell$
      end for
      Set $F_{k+K_c} = \widehat{u}_{k+K_c}$ and apply eq. (13) for $\ell = k + K_c - 1, \dots, k$ to obtain F_k
    end for
  end if
  for each anchor $k \in \mathcal{A}$ **do**
    Project to the executable target: $\widetilde{v}_k^{\mathrm{FBC}} = G_k^\top F_k$
    Form the residual target: $\widetilde{g}_k = \widetilde{v}_k^{\mathrm{FBC}} - v^g(t_k, X_{t_k}^{\theta^-})$
    Relax toward the detached target: $y_k = (1 - \rho) g_{\theta^-}(t_k, X_{t_k}^{\theta^-}) + \rho\, \widetilde{g}_k$
  end for
  Update θ by minimizing $\mathcal{L}_{\mathrm{FBC}}(\theta)$ in eq. (14)
end for

All trajectories and quantities used to construct F_k and y_k are detached from the optimization graph. Each outer iteration therefore consists of three conceptually separate stages: rollout, target construction, and regression. The network update does not differentiate through the SDE solver, finite-horizon recovery, or the consistency recursion.

B Theory and Proofs

This appendix provides the proofs and auxiliary results used in sections 3 to 5. We first establish the rank-free self-consistency identity, then develop the finite-horizon Malliavin and controllability

geometry, affine recovery, nonlinear executable consistency, the underdamped Langevin specialization, and stability results for the implemented estimator.

B.1 Stochastic Flow and Self-Consistency

Regularity. Fix an arbitrary preterminal horizon $\tau < T$. On $[0, \tau]$, assume that the reference SDE admits a C^1 stochastic flow in the initial condition, that

$$h(t, x) = \mathbb{E}_{\mathbb{P}}[F(X_T) \mid X_t = x]$$

is positive and continuously differentiable on the region under consideration, and that the relevant flow derivatives and gradients are integrable so that differentiation may be interchanged with conditional expectation. These assumptions are required only on preterminal intervals; in particular, the argument does not require regularity at the terminal time itself. They allow rectangular and state-dependent diffusion matrices.

Reference-flow Jacobian. For $0 \le s \le r \le \tau$, write

$$J^b_{r|s} = \frac{\partial X^{s,x}_r}{\partial x}$$

for the Jacobian of the *reference* stochastic flow initialized at $X_s = x$. If $G_i(t, x)$ denotes the ith column of $G(t, x)$, then for state-dependent diffusion the Jacobian satisfies

$$dJ^b_{r|s} = \nabla_x b(r, X_r) J^b_{r|s}\, dr + \sum_{i=1}^{m} \nabla_x G_i(r, X_r) J^b_{r|s}\, dW_{r,i}, \qquad J^b_{s|s} = I. \tag{35}$$

Thus, $J^b_{r|s}$ is the derivative of the full reference stochastic flow rather than merely the deterministic drift flow. When G is independent of x, the stochastic term in eq. (35) vanishes.

Proof of Theorem 4.1. Fix $0 \le s < t \le \tau$ and let $\mathbb{P}^{s,x}$ denote the reference path law restarted from $X_s = x$. By the Markov property,

$$h(s, x) = \mathbb{E}^{s,x}_{\mathbb{P}}[h(t, X_t)]\,. \tag{36}$$

Differentiating eq. (36) with respect to the initial state and using the stochastic-flow chain rule gives

$$\nabla_x h(s, x) = \mathbb{E}^{s,x}_{\mathbb{P}}\left[\left(J^b_{t|s}\right)^\top \nabla_x h(t, X_t)\right]. \tag{37}$$

Since $\nabla_x h = h\, u^\star$, division by $h(s, x) > 0$ yields

$$u^\star(s, x) = \frac{1}{h(s, x)} \mathbb{E}^{s,x}_{\mathbb{P}}\left[h(t, X_t)\left(J^b_{t|s}\right)^\top u^\star(t, X_t)\right]. \tag{38}$$

Let $\mathbb{Q}^{\star,s,x}$ denote the terminally tilted path law conditional on $X_s = x$. On $\mathcal{F}_t$ its Radon–Nikodym derivative with respect to $\mathbb{P}^{s,x}$ is

$$\left.\frac{d\mathbb{Q}^{\star,s,x}}{d\mathbb{P}^{s,x}}\right|_{\mathcal{F}_t} = \frac{h(t, X_t)}{h(s, x)}. \tag{39}$$

Applying this change of measure to eq. (38) gives

$$u^\star(s, x) = \mathbb{E}_{\mathbb{Q}^\star}\left[\left(J^b_{t|s}\right)^\top u^\star(t, X_t) \,\middle|\, X_s = x\right],$$

which proves Theorem 4.1. Here $J^b_{t|s}$ remains the Jacobian of the reference stochastic flow: the change of measure reweights the same preterminal path functional and does not replace $J^b_{t|s}$ by the Jacobian of the controlled drift. Since $\tau < T$ was arbitrary, the identity holds for every $0 \le s < t < T$. The proof is rank-free: it uses stochastic-flow differentiation and change of measure without inverting G, $GG^\top$, or a finite-time covariance matrix. Rank deficiency enters only in the subsequent recovery of the state covector from $v^\star = G^\top u^\star$.

Hard endpoint conditioning. Let $y \in \mathbb{R}^n$ and suppose that the transition density $p(t, x; T, y)$ is positive and continuously differentiable in x on the preterminal region under consideration, with the same stochastic-flow and integrability assumptions as above. Setting $h_y(t, x) = p(t, x; T, y)$, the bridge law restricted to $\mathcal{F}_t$, $t < T$, is the corresponding Doob transform, with

$$\left.\frac{d\mathbb{Q}^{y,s,x}}{d\mathbb{P}^{s,x}}\right|_{\mathcal{F}_t} = \frac{p(t, X_t; T, y)}{p(s, x; T, y)}. \tag{40}$$

Applying the preceding argument with h_y therefore gives

$$\nabla_x \log p(s, x; T, y) = \mathbb{E}_{\mathbb{Q}^y}\left[\left(J^b_{t|s}\right)^\top \nabla_x \log p(t, X_t; T, y) \mid X_s = x\right], \qquad 0 \le s < t < T. \tag{41}$$

Thus Theorem 4.1 extends to hard endpoint bridges on every preterminal region satisfying the stated regularity conditions.

B.2 Malliavin–Gramian Geometry and Affine Recovery

We first relate the pullback geometry used by FBC to the usual Malliavin covariance and then establish exact finite-window recovery for affine diffusions with deterministic time-dependent coefficients. Whenever the inverse derivative flow $M^b_{s,r} = (J^b_{r|s})^{-1}$ is used below, we work on the stochastic-flow domain where $J^b_{r|s}$ is invertible.

Proof of Proposition 4.2. By the derivative-flow composition rule,

$$J^b_{t|r} = J^b_{t|s}\left(J^b_{r|s}\right)^{-1} = J^b_{t|s} M^b_{s,r}.$$

Substituting this identity into the Malliavin covariance gives

$$C_{s,t} = J^b_{t|s}\left[\int_s^t M^b_{s,r} G_r G_r^\top \left(M^b_{s,r}\right)^\top dr\right]\left(J^b_{t|s}\right)^\top = J^b_{t|s} W^b_{s,t}\left(J^b_{t|s}\right)^\top.$$

Since $J^b_{t|s}$ is invertible on the stochastic-flow domain, congruence preserves rank, yielding

$$\operatorname{rank} C_{s,t} = \operatorname{rank} W^b_{s,t}.$$

Finite-window least-squares interpretation. For a fixed anchor k, consider

$$\mathcal{J}_k(u) = \sum_{j=k}^{k+K_w-1} \Delta t \, \left\| Z_{k,j}^\top u - v_j \right\|^2 + \varepsilon \|u\|^2. \tag{42}$$

Differentiating $\mathcal{J}_k$ with respect to u gives the normal equation

$$(W_k + \varepsilon I)u = r_k. \tag{43}$$

Hence, for $\varepsilon > 0$, the unique minimizer is

$$\widehat{u}_k = (W_k + \varepsilon I)^{-1} r_k,$$

while for $\varepsilon = 0$ the same expression holds whenever $W_k \succ 0$.

Finite-window recovery fits a single anchor covector to executable controls through the dynamically transported directions $Z_{k,j}$ over the recovery window. This provides an algebraic least-squares construction at the trajectory level. Exact affine recovery emerges after conditional expectation, while the nonlinear recovery error is characterized separately in Appendix B.4.

Affine setting. Consider the affine reference diffusion

$$dX_r = (A_r X_r + a_r)\, dr + G_r\, dW_r, \tag{44}$$

where A_r, a_r, and G_r are deterministic functions of time. Let $\Phi_{r,s}$ denote the state-transition matrix,

$$\frac{d}{dr}\Phi_{r,s} = A_r \Phi_{r,s}, \qquad \Phi_{s,s} = I.$$

Then the reference-flow Jacobian is deterministic, $J^b_{r|s} = \Phi_{r,s}$. Define

$$Z_{s,r} = \Phi^{-1}_{r,s} G_r, \qquad W_{s,t} = \int_s^t Z_{s,r} Z_{s,r}^\top \, dr.$$

Proof of Theorem 4.3. For any $r \in (s,t)$, Theorem 4.1 and the deterministic Jacobian give

$$u^\star(s,x) = \Phi_{r,s}^\top \mathbb{E}_{\mathbb{Q}^\star}[\,u^\star(r,X_r) \mid X_s = x\,],$$

hence

$$\mathbb{E}_{\mathbb{Q}^\star}[\,u^\star(r,X_r) \mid X_s = x] = \Phi_{r,s}^{-\top} u^\star(s,x). \tag{45}$$

Using $v^\star(r,X_r) = G_r^\top u^\star(r,X_r)$ and $Z_{s,r} = \Phi_{r,s}^{-1} G_r$, we obtain

$$\mathbb{E}_{\mathbb{Q}^\star}[\,Z_{s,r} v^\star(r,X_r) \mid X_s = x] = Z_{s,r} G_r^\top \Phi_{r,s}^{-\top} u^\star(s,x) \tag{46}$$

$$= Z_{s,r} Z_{s,r}^\top u^\star(s,x). \tag{47}$$

Assuming the integrability required for conditional Fubini, integration over $r \in [s,t]$ gives

$$\mathbb{E}_{\mathbb{Q}^\star}\left[\int_s^t Z_{s,r} v^\star(r,X_r)\,dr \,\middle|\, X_s = x\right] = W_{s,t} u^\star(s,x). \tag{48}$$

If $W_{s,t} \succ 0$, multiplication by $W_{s,t}^{-1}$ proves Theorem 4.3.

Trajectory-level estimator. In the nonsingular affine setting, define

$$\widehat{u}_{s,t}^{\mathrm{traj}} = W_{s,t}^{-1} \int_s^t Z_{s,r} v^\star(r,X_r)\,dr. \tag{49}$$

Because $W_{s,t}$ is deterministic, eq. (48) implies

$$\mathbb{E}_{\mathbb{Q}^\star}\left[\,\widehat{u}_{s,t}^{\mathrm{traj}} \,\middle|\, X_s = x\right] = u^\star(s,x). \tag{50}$$

Thus, when the exact executable control is used, the trajectory-level finite-window estimator is conditionally unbiased. Theorem 4.3 is therefore an exact population recovery identity; an individual continuation generally retains sampling variance and need not recover $u^\star(s,x)$ pathwise.

With ridge regularization,

$$\widehat{u}_{s,t}^{\varepsilon} = (W_{s,t} + \varepsilon I)^{-1} \int_s^t Z_{s,r} v^\star(r,X_r)\,dr,$$

and eq. (48) gives

$$\mathbb{E}\big[\,\widehat{u}_{s,t}^{\varepsilon} \,\big|\, X_s = x\big] - u^\star(s,x) = -\varepsilon (W_{s,t} + \varepsilon I)^{-1} u^\star(s,x). \tag{51}$$

For fixed $W_{s,t} \succ 0$, the bias therefore vanishes as $\varepsilon \downarrow 0$.

Singular Gramian and executable recovery. Full covector recovery is unnecessary for bridge execution, and the affine normal equation eq. (48) remains valid when $W_{s,t}$ is singular.

Corollary B.1 (Singular affine recovery)**.** *Suppose* $W_{s,t} \succeq 0$ *and* $\mathrm{Range}(G_s) \subseteq \mathrm{Range}(W_{s,t})$. *Then the executable control is exactly identified by*

$$v^\star(s,x) = G_s^\top W_{s,t}^\dagger \mathbb{E}_{\mathbb{Q}^\star}\left[\int_s^t Z_{s,r} v^\star(r,X_r)\,dr \,\middle|\, X_s = x\right]. \tag{52}$$

Moreover,

$$\widehat{v}_{s,t}^{\mathrm{traj}} = G_s^\top W_{s,t}^\dagger \int_s^t Z_{s,r} v^\star(r,X_r)\,dr \tag{53}$$

is conditionally unbiased for $v^\star(s,x)$ *whenever it is integrable.*

Proof. By eq. (48), the right-hand side of eq. (52) equals $G_s^\top W_{s,t}^\dagger W_{s,t} u^\star(s,x)$. Since $W_{s,t}^\dagger W_{s,t}$ is the orthogonal projector onto $\mathrm{Range}(W_{s,t})$ and $\mathrm{Range}(G_s) \subseteq \mathrm{Range}(W_{s,t})$, we have $G_s^\top W_{s,t}^\dagger W_{s,t} = G_s^\top$, proving the first claim. Taking the conditional expectation of eq. (53) gives the second. □

B.3 Controllability and Small-Window Scaling

This section characterizes both *whether* the finite-horizon Gramian recovers missing state directions and *how strongly* the resulting controllability layers are identified as the recovery window shrinks.

For constant matrices $A \in \mathbb{R}^{n\times n}$ and $G \in \mathbb{R}^{n\times m}$, recall the pullback controllability Gramian

$$W_h^-(A,G) = \int_0^h e^{-A\tau} G G^\top e^{-A^\top \tau}\, d\tau. \tag{54}$$

Proof of Theorem 5.1. For any $x \in \mathbb{R}^n$, $x^\top W_h^-(A,G)x = \int_0^h \|G^\top e^{-A^\top \tau}x\|^2\, d\tau$. Hence x lies in the null space of $W_h^-(A,G)$ if and only if $G^\top e^{-A^\top \tau}x = 0$ for all $\tau \in [0,h]$. Analyticity and differentiation at $\tau = 0$, together with Cayley–Hamilton, show that this is equivalent to $G^\top (A^\top)^k x = 0$ for $k = 0,\dots,n-1$, or equivalently $x \perp \mathrm{Range}[G, AG, \dots, A^{n-1}G]$. Thus $W_h^-(A,G) \succ 0$ exactly when the Kalman controllability matrix has full row rank.

Finite-time controllability determines which directions are eventually identifiable, but it does not quantify the strong anisotropy that appears as $h \downarrow 0$. We next make this structure explicit.

Theorem B.2 (Small-window controllability scaling)**.** *Assume (A,G) is controllable and define the Kalman filtration*

$$\mathcal{R}_0 = \{0\}, \qquad \mathcal{R}_j = \mathrm{span}\{G, AG, \dots, A^{j-1}G\}. \tag{55}$$

Let r be the smallest index such that $\mathcal{R}_r = \mathbb{R}^n$, and define the orthogonal controllability layers

$$\mathcal{V}_j = \mathcal{R}_j \cap \mathcal{R}_{j-1}^{\perp}, \qquad j = 1,\dots,r,$$

with orthogonal projector P_j onto $\mathcal{V}_j$. Let $d_j = \dim \mathcal{V}_j$, $q_j = \sum_{\ell=1}^{j} d_\ell$, and $q_0 = 0$. Introduce the anisotropic dilation

$$D_h = \sum_{j=1}^{r} h^{j-\frac{1}{2}} P_j. \tag{56}$$

Then there exist constants $c, C, h_0 > 0$ such that, for every $0 < h \le h_0$,

$$cI \preceq D_h^{-1} W_h^-(A,G) D_h^{-1} \preceq CI. \tag{57}$$

Consequently, for every $x \in \mathcal{V}_j$,

$$c\, h^{2j-1}\|x\|^2 \le x^\top W_h^-(A,G)x \le C\, h^{2j-1}\|x\|^2. \tag{58}$$

Moreover, if the eigenvalues are ordered as $\lambda_1(W_h^-) \ge \cdots \ge \lambda_n(W_h^-) > 0$, then for every $j = 1,\dots,r$ and every $q_{j-1} < i \le q_j$,

$$\lambda_i(W_h^-) = \Theta\big(h^{2j-1}\big) \qquad (h \downarrow 0). \tag{59}$$

In particular,

$$\lambda_{\max}(W_h^-) = \Theta(h), \qquad \lambda_{\min}(W_h^-) = \Theta\big(h^{2r-1}\big). \tag{60}$$

Proof. Rescale time by $\tau = hs$ and define

$$B_h(s) = h^{1/2} D_h^{-1} e^{-Ahs} G, \qquad 0 \le s \le 1. \tag{61}$$

Then

$$D_h^{-1} W_h^-(A,G) D_h^{-1} = \int_0^1 B_h(s) B_h(s)^\top\, ds. \tag{62}$$

Because $\mathrm{Range}(A^k G) \subseteq \mathcal{R}_{k+1}$, orthogonality of the filtration gives $P_j A^k G = 0$ whenever $k < j-1$. Expanding e^{-Ahs} therefore yields, uniformly for $s \in [0,1]$,

$$P_j B_h(s) = \frac{(-s)^{j-1}}{(j-1)!} P_j A^{j-1} G + O(h). \tag{63}$$

Hence B_h converges uniformly to

$$B_0(s) = \sum_{j=1}^{r} \frac{(-s)^{j-1}}{(j-1)!} P_j A^{j-1} G. \tag{64}$$

Set

$$Q_0 = \int_0^1 B_0(s) B_0(s)^\top \, ds.$$

We claim that $Q_0 \succ 0$. If $x^\top Q_0 x = 0$, continuity implies $B_0(s)^\top x = 0$ for every $s \in [0,1]$. Writing $x = \sum_{j=1}^{r} x_j$ with $x_j \in \mathcal{V}_j$ and equating coefficients of the resulting vector-valued polynomial gives

$$x_j^\top P_j A^{j-1} G = 0, \qquad j = 1, \dots, r.$$

Since $\mathcal{R}_j = \mathcal{R}_{j-1} + \mathrm{Range}(A^{j-1}G)$, projection onto $\mathcal{V}_j$ gives

$$\mathrm{Range}(P_j A^{j-1} G) = \mathcal{V}_j.$$

Therefore $x_j = 0$ for every j, so $x = 0$ and $Q_0 \succ 0$.

Uniform convergence of B_h implies that

$$Q_h := D_h^{-1} W_h^-(A, G) D_h^{-1}$$

and Q_h^{-1} remain uniformly bounded for sufficiently small h, proving eq. (57). Equivalently,

$$c D_h^2 \preceq W_h^-(A, G) \preceq C D_h^2.$$

For $x \in \mathcal{V}_j$, $D_h x = h^{j-1/2} x$, which gives eq. (58). The eigenvalues of D_h^2 are h^{2j-1} with multiplicity d_j; monotonicity of ordered eigenvalues under the Loewner order therefore yields eq. (59) and eq. (60). □

Thus, controllability guarantees identifiability for every $h > 0$, while the scaling h^{2j-1} quantifies the increasingly weak short-window information carried by deeper controllability layers.

Corollary B.3 (First-layer executable amplification)**.** *Under the assumptions of Theorem B.2,*

$$\left\| \left(W_h^-(A, G)\right)^{-1/2} G \right\|_{\mathrm{op}}^2 = O(h^{-1}) \qquad (h \downarrow 0). \tag{65}$$

Proof. Since $\mathrm{Range}(G) = \mathcal{V}_1$, we have $D_h^{-1} G = h^{-1/2} G$. Writing $W_h^- = D_h Q_h D_h$, with Q_h^{-1} uniformly bounded by Theorem B.2, gives $G^\top (W_h^-)^{-1} G = h^{-1} G^\top Q_h^{-1} G$. Taking the operator norm proves eq. (65). □

This distinction is important for executable consistency. Full-covector inversion can be controlled by the deepest scale h^{2r-1}, whereas the current executable control $G^\top u$ probes only the directly forced layer and therefore incurs the milder h^{-1} inverse amplification.

Kinetic geometry. Consider the $2d$-dimensional kinetic auxiliary dynamics

$$A_{\mathrm{kin}} = \begin{pmatrix} 0 & I_d \\ 0 & -\gamma I_d \end{pmatrix}, \qquad G_{\mathrm{kin}} = \begin{pmatrix} 0 \\ \sigma I_d \end{pmatrix}, \qquad \sigma \neq 0. \tag{66}$$

Corollary B.4 (Kinetic controllability)**.** *The pair $(A_{\mathrm{kin}}, G_{\mathrm{kin}})$ is controllable. Consequently,*

$$W_h^-(A_{\mathrm{kin}}, G_{\mathrm{kin}}) \succ 0 \qquad \textit{for every } h > 0,$$

although $G_{\mathrm{kin}} G_{\mathrm{kin}}^\top$ vanishes on the position coordinates.

Proof. The controllability matrix satisfies $[G_{\mathrm{kin}}, A_{\mathrm{kin}} G_{\mathrm{kin}}] = \begin{pmatrix} 0 & \sigma I_d \\ \sigma I_d & -\gamma \sigma I_d \end{pmatrix}$, which has rank $2d$ for $\sigma \neq 0$. Theorem 5.1 therefore gives $W_h^-(A_{\mathrm{kin}}, G_{\mathrm{kin}}) \succ 0$ for every $h > 0$. □

The kinetic pair has exactly two controllability layers: momentum is directly forced, whereas position is exposed after one drift coupling. Theorem B.2 therefore predicts eigenvalue scales h and h^3. For this system the leading constants can also be computed explicitly.

Corollary B.5 (Kinetic Gramian scaling)**.** *For the pair in eq.* (66)*, as* $h \downarrow 0$*,*

$$\lambda_{\max}\big(W_h^-\big) = \sigma^2 h\,[1 + O(h)], \qquad \lambda_{\min}\big(W_h^-\big) = \frac{\sigma^2}{12} h^3 [1 + O(h)]. \tag{67}$$

Each of the two eigenvalues has multiplicity d*.*

Proof. Because the coordinate pairs decouple, it suffices to consider one position–momentum pair. As $\tau \downarrow 0$,

$$e^{-A_{\mathrm{kin}}\tau} G_{\mathrm{kin}} = \sigma \begin{pmatrix} -\tau + O(\tau^2) \\ 1 + \gamma\tau + O(\tau^2) \end{pmatrix}.$$

Substitution into eq. (54) gives

$$W_h^- = \sigma^2 \begin{pmatrix} \frac{h^3}{3} + O(h^4) & -\frac{h^2}{2} + O(h^3) \\ -\frac{h^2}{2} + O(h^3) & h + O(h^2) \end{pmatrix}.$$

Hence $\mathrm{tr}(W_h^-) = \sigma^2 h[1 + O(h)]$ and $\det(W_h^-) = \frac{\sigma^4}{12} h^4[1 + O(h)]$. The larger eigenvalue is therefore $\sigma^2 h[1 + O(h)]$, while the smaller is the determinant divided by the larger eigenvalue, giving $\frac{\sigma^2}{12} h^3[1 + O(h)]$. The same two eigenvalues occur independently in each of the d coordinate pairs. □

Thus, the directly forced momentum layer is observed at order h, whereas the hidden position layer is recovered at order h^3. This is the two-layer geometry used in the underdamped consistency analysis below.

B.4 Executable-Control Consistency

The finite-horizon lift need not recover every component of the state-space covector accurately. Bridge execution only depends on the projection $G_s^\top u$, so the relevant inverse problem is whether the finite-window statistics identify and stably recover this projection. We first isolate the required linear-algebraic statement.

Executable identifiability. Let $W \succeq 0$ and $r = Wu^\dagger$. Since the solution set of $Wu = r$ is $u^\dagger + \ker W$, the executable projection $G_s^\top u$ is uniquely determined if and only if $\ker W \subseteq \ker G_s^\top$, equivalently, by symmetry of W, $\mathrm{Range}(G_s) \subseteq \mathrm{Range}(W)$. Under this condition, $G_s^\top u^\dagger = G_s^\top W^\dagger r$. Thus, singularity of W is harmless for bridge execution whenever its unidentified directions are invisible to $G_s^\top$.

Control-relevant amplification. Let $P_W = W^\dagger W$ denote the orthogonal projector onto $\mathrm{Range}(W)$, and define

$$\Gamma_s(W) = \left\| W^{\dagger/2} G_s \right\|_{\mathrm{op}}^2. \tag{68}$$

Unlike the global pseudoinverse norm $\|W^\dagger\|_{\mathrm{op}}$, $\Gamma_s(W)$ measures inverse amplification only in state directions that contribute to the executable control.

Proposition B.6 (Executable stability)**.** *Let* $W \succeq 0$ *satisfy* $\mathrm{Range}(G_s) \subseteq \mathrm{Range}(W)$ *and suppose*

$$r = Wu^\dagger + e.$$

For the unregularized estimator

$$\widehat{v}_0 = G_s^\top W^\dagger r,$$

we have

$$\left\| \widehat{v}_0 - G_s^\top u^\dagger \right\| \le \sqrt{\Gamma_s(W)} \left\| W^{\dagger/2} e \right\|. \tag{69}$$

For $\varepsilon > 0$*, define*

$$\widehat{v}_\varepsilon = G_s^\top (W + \varepsilon I)^{-1} r.$$

Then

$$\left\|\widehat{v}_\varepsilon - G_s^\top u^\dagger\right\| \le \sqrt{\Gamma_s(W)}\left(\left\|W^{\dagger/2}e\right\| + \frac{\sqrt{\varepsilon}}{2}\left\|P_W u^\dagger\right\|\right). \tag{70}$$

Proof. The range condition gives $P_W G_s = G_s$ and $G_s^\top(I - P_W) = 0$. For $\varepsilon = 0$, $\widehat{v}_0 - G_s^\top u^\dagger = (W^{\dagger/2}G_s)^\top W^{\dagger/2}e$, so eq. (69) follows by Cauchy–Schwarz. For $\varepsilon > 0$, $\widehat{v}_\varepsilon - G_s^\top u^\dagger = G_s^\top(W + \varepsilon I)^{-1}e - \varepsilon G_s^\top(W+\varepsilon I)^{-1}P_W u^\dagger$. On $\mathrm{Range}(W)$, the spectral bounds $\sqrt{\lambda}/(\lambda+\varepsilon) \le 1/\sqrt{\lambda}$ and $\varepsilon\sqrt{\lambda}/(\lambda+\varepsilon) \le \sqrt{\varepsilon}/2$ for $\lambda > 0$, followed by Cauchy–Schwarz, give eq. (70). □

Proof of Theorem 5.2. Set $W = \overline{W}_h^a$, $r = \overline{r}_h^a$, $u^\dagger = u^\star(s,x)$, and $e = b_h^a$, so that $r = Wu^\dagger + e$. Applying Proposition B.6 and the range condition gives

$$\left\|\overline{v}_{h,\varepsilon_h}^a - v^\star(s,x)\right\| \le \sqrt{\Gamma_s(\overline{W}_h^a)}\left(\left\|(\overline{W}_h^a)^{\dagger/2}b_h^a\right\| + \frac{\sqrt{\varepsilon_h}}{2}\|u^\star(s,x)\|\right), \tag{71}$$

where we used $\|P_h u^\star(s,x)\| \le \|u^\star(s,x)\|$ for the orthogonal projector $P_h = (\overline{W}_h^a)^\dagger \overline{W}_h^a$. Under the assumed bounds $\Gamma_s(\overline{W}_h^a) = O(h^{-\alpha})$ and $\|(\overline{W}_h^a)^{\dagger/2}b_h^a\| = O(h^\zeta)$,

$$\left\|\overline{v}_{h,\varepsilon_h}^a - v^\star(s,x)\right\| = O\left(h^{\zeta-\alpha/2} + \sqrt{\frac{\varepsilon_h}{h^\alpha}}\right).$$

Both terms vanish when $\zeta > \alpha/2$ and $\varepsilon_h = o(h^\alpha)$, proving Theorem 5.2.

Exact reference-tangent residual. For additive diffusion with constant G and $\Sigma_G = GG^\top$, let the finite-window transport use the exact reference tangent. Along the Doob bridge conditioned on $X_s^\star = x$, define $J_r = J_{r|s}^b$, $M_r = J_r^{-1}$, $H_r = M_r \Sigma_G M_r^\top$, and $Y_r = J_r^\top u^\star(r, X_r^\star)$. Let b_h^b denote the corresponding population recovery residual obtained from the exact reference transport.

Corollary B.7 (Exact-tangent residual rate)**.** *Assume the regularity of Theorem 4.1, sufficient square integrability on $[s, s+h]$, and suppose that on the window*

$$\|\nabla_x b(r, X_r^\star)\|_{\mathrm{op}} \le B, \qquad \|\nabla_x u^\star(r, X_r^\star)G\|_F \le S. \tag{72}$$

Then

$$\|b_h^b\| \le \frac{8}{5}B\|\Sigma_G\|_F S e^{3Bh}h^{5/2}. \tag{73}$$

Proof. By the flow composition rule and Theorem 4.1, $(Y_r)_{r\ge s}$ is a $\mathbb{Q}^\star$-martingale with $Y_s = u^\star(s,x)$. Since $u^\star(r, X_r^\star) = M_r^\top Y_r$ and $Z_{s,r}^b = M_r G$, we have $Z_{s,r}^b v^\star(r, X_r^\star) = H_r Y_r$. Writing $\overline{H}_r = \mathbb{E}_{\mathbb{Q}^\star}[H_r \mid X_s^\star = x]$ and using $\mathbb{E}[Y_r - Y_s \mid X_s^\star = x] = 0$ therefore gives

$$b_h^b = \int_s^{s+h} \mathbb{E}_{\mathbb{Q}^\star}\left[(H_r - \overline{H}_r)(Y_r - Y_s)\,\middle|\, X_s^\star = x\right] dr. \tag{74}$$

The bounds $\|J_r\|_{\mathrm{op}}, \|M_r\|_{\mathrm{op}} \le e^{B(r-s)}$ imply

$$\left(\mathbb{E}\left[\|H_r - \overline{H}_r\|_F^2 \,\middle|\, X_s^\star = x\right]\right)^{1/2} \le 4B\|\Sigma_G\|_F e^{2Bh}(r-s). \tag{75}$$

Moreover, Itô's formula for the Doob covector gives $dY_r = J_r^\top \nabla_x u^\star(r, X_r^\star) G\, dW_r^\star$, so conditional Itô isometry yields

$$\left(\mathbb{E}\left[\|Y_r - Y_s\|^2 \,\middle|\, X_s^\star = x\right]\right)^{1/2} \le e^{Bh}S(r-s)^{1/2}. \tag{76}$$

Conditional Cauchy–Schwarz and integration of $(r-s)^{3/2}$ over $[s, s+h]$ give eq. (73). □

B.5 KINETIC LANGEVIN SPECIALIZATION

We now verify the executable-consistency conditions for the kinetic auxiliary transport used in the underdamped experiments. Consider the reference diffusion

$$dq_t = p_t\,dt, \qquad dp_t = [-\nabla V(q_t) - \gamma p_t]\,dt + \sigma\,dW_t, \tag{77}$$

with $\sigma \neq 0$, and split the drift as

$$a(q,p) = (p, -\gamma p), \qquad c(q,p) = (0, -\nabla V(q)). \tag{78}$$

Writing

$$A := \nabla a = \begin{pmatrix} 0 & I \\ 0 & -\gamma I \end{pmatrix}, \qquad C_r := \nabla c(X_r^\star) = \begin{pmatrix} 0 & 0 \\ -\nabla^2 V(q_r^\star) & 0 \end{pmatrix}, \qquad G = \begin{pmatrix} 0 \\ \sigma I \end{pmatrix}, \tag{79}$$

we have the structural cancellation

$$C_r G = 0. \tag{80}$$

Thus the force-curvature term omitted from the auxiliary tangent does not act directly on the instantaneous Brownian columns. The stronger second-order matching used below follows from this cancellation together with the short-time kinetic transport generated by A.

Second-order matching of transported Brownian directions. Fix a preterminal anchor (s,x) and a realization of the Doob bridge $X^\star = (q^\star, p^\star)$ starting from $X_s^\star = x$. Let $M_{s,r}^b$ denote the inverse tangent of the *reference* drift, evaluated along this bridge realization, so that

$$\partial_r M_{s,r}^b = -M_{s,r}^b(A + C_r), \qquad M_{s,s}^b = I.$$

The kinetic auxiliary inverse tangent is deterministic and given by

$$M_{s,r}^a = e^{-A(r-s)}.$$

Lemma B.8 (Noise-column matching)**.** *Assume that for some $h_0 > 0$,*

$$\sup_{r\in[s,s+h_0]} \|\nabla^2 V(q_r^\star)\|_{\mathrm{op}} \le L_2, \qquad \sup_{r\in[s,s+h_0]} \|M_{s,r}^b\|_{\mathrm{op}} \le K_M. \tag{81}$$

Then there exists a constant C_Z, depending only on $K_M, L_2, |\sigma|, |\gamma|$, and h_0, such that

$$\left\| \left(M_{s,s+\tau}^b - M_{s,s+\tau}^a\right) G \right\|_{\mathrm{op}} \le C_Z \tau^2, \qquad 0 \le \tau \le h_0. \tag{82}$$

Proof. Variation of constants gives $M_{s,r}^b - M_{s,r}^a = -\int_s^r M_{s,u}^b C_u e^{-A(r-u)}\,du$. Since $C_u G = 0$, we have $C_u e^{-A\theta} G = C_u(e^{-A\theta} - I)G$. Let $K_A = \sup_{0<\theta\le h_0} \|(e^{-A\theta} - I)G\|_{\mathrm{op}}/\theta < \infty$. The curvature bound then gives $\|C_u e^{-A\theta} G\|_{\mathrm{op}} \le L_2 K_A \theta$. Using $\|M_{s,u}^b\|_{\mathrm{op}} \le K_M$ and integrating with $\theta = r - u$ yields $\|(M_{s,s+\tau}^b - M_{s,s+\tau}^a)G\|_{\mathrm{op}} \le \frac{1}{2} K_M L_2 K_A \tau^2$, proving the claim. □

The second-order match is specific to the kinetic structure: $C_r G = 0$ removes the zeroth-order interaction, so the omitted force curvature acts only after kinetic transport has generated an $O(\tau)$ position component.

Auxiliary population residual. Define

$$Z_{s,r}^b = M_{s,r}^b G, \qquad Z_{s,r}^a = M_{s,r}^a G,$$

and, for $q \in \{a, b\}$,

$$\overline{W}_h^q = \mathbb{E}_{\mathbb{Q}^\star}\left[\int_s^{s+h} Z_{s,r}^q (Z_{s,r}^q)^\top\,dr \,\middle|\, X_s^\star = x \right].$$

The corresponding population residual is

$$b_h^q = \mathbb{E}_{\mathbb{Q}^\star}\left[\int_s^{s+h} Z_{s,r}^q v^\star(r, X_r^\star)\,dr \,\middle|\, X_s^\star = x \right] - \overline{W}_h^q u^\star(s,x). \tag{83}$$

Proposition B.9 (Auxiliary residual)**.** *Assume the conditions of Lemma B.8. Suppose additionally that $v^\star(r, X_r^\star)$ and the exact and auxiliary transported Brownian columns are uniformly bounded on the localized preterminal window, and that the fixed anchor covector $u^\star(s, x)$ is finite. Then*

$$\|b_h^a - b_h^b\| = O(h^3). \tag{84}$$

If, in addition, the assumptions of Corollary B.7 hold, then

$$\|b_h^a\| = O(h^{5/2}). \tag{85}$$

Proof. Lemma B.8 gives $\|Z_{s,s+\tau}^a - Z_{s,s+\tau}^b\|_{\mathrm{op}} = O(\tau^2)$. Under the stated boundedness assumptions, both the control-moment difference and $\|\overline{W}_h^a - \overline{W}_h^b\|_{\mathrm{op}}$ are therefore $O(\int_0^h \tau^2\, d\tau) = O(h^3)$, which gives $\|b_h^a - b_h^b\| = O(h^3)$. Corollary B.7 gives $\|b_h^b\| = O(h^{5/2})$; since $h^3 = o(h^{5/2})$, the triangle inequality yields $\|b_h^a\| = O(h^{5/2})$. □

Thus, replacing the nonlinear reference tangent by the kinetic auxiliary transport does not change the leading $h^{5/2}$ order of the population recovery residual.

Proof of Corollary 5.3. Corollary B.4 gives $\overline{W}_h^a \succ 0$ for every $h > 0$, so the range condition in Theorem 5.2 is automatic. Corollary B.3 gives $\Gamma_s(\overline{W}_h^a) = O(h^{-1})$, while Corollary B.5 gives $\lambda_{\min}(\overline{W}_h^a) = \Omega(h^3)$. Together with Proposition B.9, this yields $\|(\overline{W}_h^a)^{-1/2} b_h^a\| = O(h)$, so Theorem 5.2 applies with $\alpha = \zeta = 1$ and gives $\|\overline{v}_{h,\varepsilon_h}^a - v^\star(s,x)\| = O(h^{1/2} + \sqrt{\varepsilon_h/h})$. Hence $\varepsilon_h = o(h)$ implies $\overline{v}_{h,\varepsilon_h}^a \to v^\star(s,x)$ as $h \downarrow 0$, proving Corollary 5.3.

B.6 Short-Horizon Fixed-Point Uniqueness

The self-consistency identity in Theorem 4.1 is a necessary property of the Doob covector. We now show that, for constant additive diffusion, the corresponding *exact population terminal-consistency map* is contractive on sufficiently short time slabs. The result is local in both time and function class and establishes uniqueness of the Doob covector among fields in a prescribed invariant bounded-Lipschitz class. The finite-window, ridge-regularized, discrete FBC update is analyzed separately in Appendix B.7.

Terminal-consistency map. Let $G \in \mathbb{R}^{n\times m}$ be constant and set $\Sigma_G = GG^\top$. Given a candidate state-covector field u, define

$$dX_t^u = [b(t, X_t^u) + \Sigma_G u(t, X_t^u)]\, dt + G\, dW_t. \tag{86}$$

Fix a terminal time $\tau \le T$ and a terminal covector field $g : \mathbb{R}^n \to \mathbb{R}^n$. For hard endpoint conditioning, we take $\tau < T$ so that the terminal covector at τ is preterminal and regular. Let $J^b[X^u]_{r|s}$ denote the derivative flow generated by the *reference drift* b, evaluated along the coordinate path X^u:

$$\frac{d}{dr} J^b[X^u]_{r|s} = \nabla_x b(r, X_r^u) J^b[X^u]_{r|s}, \qquad J^b[X^u]_{s|s} = I. \tag{87}$$

This is intentionally not the Jacobian of the candidate-controlled drift $b + \Sigma_G u$.

Define

$$(\mathcal{T}_{\tau,g} u)(s,x) = \mathbb{E}\Big[\big(J^b[X^u]_{\tau|s}\big)^\top g(X_\tau^u) \,\Big|\, X_s^u = x \Big], \qquad s \in [\tau - H, \tau]. \tag{88}$$

Regularity class. Let $\mathcal{U}$ be a closed subset of the Banach space of bounded continuous vector fields on $[\tau - H, \tau] \times \mathbb{R}^n$, equipped with the supremum norm. Assume that:

1. $b(t,\cdot)$ is L_b-Lipschitz and $\|\nabla_x b(t,x)\|_{\mathrm{op}} \le B_b$;
2. $\nabla_x b(t,\cdot)$ is L_J-Lipschitz;
3. every $u \in \mathcal{U}$ is L_u-Lipschitz in the state variable, with the same constant L_u;
4. the terminal field satisfies $\|g(x)\| \le M_g$ and $\|g(x) - g(y)\| \le L_g \|x - y\|$;
5. $\mathcal{T}_{\tau,g}$ maps $\mathcal{U}$ into itself.

Set

$$a_G = \|\Sigma_G\|_{\mathrm{op}}, \qquad \kappa = L_b + a_G L_u, \tag{89}$$

and

$$q_H = a_G e^{(\kappa+B_b)H}\left(L_g H + M_g L_J H^2\right). \tag{90}$$

Theorem B.10 (Short-horizon fixed-point uniqueness)**.** *Under the assumptions above,*

$$\|\mathcal{T}_{\tau,g}u - \mathcal{T}_{\tau,g}w\|_\infty \le q_H \|u-w\|_\infty, \qquad u, w \in \mathcal{U}. \tag{91}$$

Since $q_H = O(H)$ as $H \downarrow 0$, the map $\mathcal{T}_{\tau,g}$ is a contraction whenever H is sufficiently small that $q_H < 1$, and therefore has a unique fixed point in $\mathcal{U}$.

If $g = u^\star(\tau,\cdot)$ and the restriction of the Doob covector to $[\tau - H, \tau]$ belongs to $\mathcal{U}$, then this unique fixed point is $u^\star$ on that slab.

Proof. Fix $s \in [\tau - H, \tau]$, $x \in \mathbb{R}^n$, and couple X^u and X^w with the same Brownian path. Writing $d = \|u - w\|_\infty$, additivity of the diffusion and the Lipschitz assumptions give, by Gronwall, $\sup_{q\in[s,r]} \|X_q^u - X_q^w\| \le a_G(r-s)e^{\kappa(r-s)}d$. Moreover, $\|J^b[X^u]_{r|s}\|_{\mathrm{op}}, \|J^b[X^w]_{r|s}\|_{\mathrm{op}} \le e^{B_b(r-s)}$, and variation of constants together with the L_J-Lipschitz property of ∇b yields $\|J^b[X^u]_{r|s} - J^b[X^w]_{r|s}\|_{\mathrm{op}} \le a_G L_J (r-s)^2 e^{(\kappa+B_b)(r-s)}d$. Therefore, with $\Delta = \tau - s \le H$,

$$\|(\mathcal{T}_{\tau,g}u)(s,x) - (\mathcal{T}_{\tau,g}w)(s,x)\| \le a_G e^{(\kappa+B_b)H}\left(L_g H + M_g L_J H^2\right) d = q_H d.$$

Taking the supremum proves eq. (91). When $g = u^\star(\tau,\cdot)$ and $\tau < T$, Theorem 4.1 gives $\mathcal{T}_{\tau,g}u^\star = u^\star$. If $\tau = T$ with a smooth positive terminal weight F, the same identity follows by differentiating $h(s,x) = \mathbb{E}_{\mathbb{P}}^{s,x}[F(X_T)]$, with $g = \nabla \log F$. Since $\mathcal{U}$ is complete, invariance and $q_H < 1$ imply uniqueness by Banach's fixed-point theorem. □

The invariant class $\mathcal{U}$ may contain non-gradient fields. Gradient structure identifies the Doob covector from the terminal tilt; once the terminal field is fixed, contraction yields uniqueness within the invariant class.

For hard endpoint bridges, the result is applied only at preterminal $\tau < T$ where the terminal covector satisfies the stated regularity assumptions.

B.7 Finite-Sample and Discrete-Target Stability

The preceding results characterize the ideal finite-horizon recovery at the population level. We now connect this analysis to the discrete FBC target used in Algorithm 1. We separate five sources of approximation: finite-window observation error, finite-continuation sampling, auxiliary-transport error, time discretization, and propagation through the short backward correction. We finally record the effect of the relaxation parameter used in the outer fixed-point update.

B.7.1 Pathwise Finite-Window Recovery

We begin with a deterministic perturbation bound for one finite recovery window after the empirical Gramian has become positive definite; singular windows are covered by the executable-space pseudoinverse analysis of Appendix B.4. Fix an anchor t_k and suppose

$$v_j = Z_{k,j}^\top u_k^\dagger + \xi_j, \qquad j = k, \dots, k + K_w - 1,$$

for an anchor covector $u_k^\dagger$, and define $e_k = \sum_{j=k}^{k+K_w-1} \Delta t\, Z_{k,j}\xi_j$.

Proposition B.11 (Pathwise window recovery)**.** *Suppose $W_k \succeq \mu_k I$ for some $\mu_k > 0$. Then the ridge estimator $\widehat{u}_k = (W_k + \varepsilon I)^{-1} r_k$ satisfies*

$$\widehat{u}_k - u_k^\dagger = (W_k + \varepsilon I)^{-1}(e_k - \varepsilon u_k^\dagger),$$

and consequently

$$\|\widehat{u}_k - u_k^\dagger\| \le \frac{\|e_k\| + \varepsilon\|u_k^\dagger\|}{\mu_k + \varepsilon}. \tag{92}$$

Proof. The observation model gives $r_k = W_k u_k^\dagger + e_k$. Substitution into $\widehat{u}_k = (W_k + \varepsilon I)^{-1} r_k$ yields the stated identity, while $W_k \succeq \mu_k I$ implies $\|(W_k + \varepsilon I)^{-1}\|_{\mathrm{op}} \le (\mu_k + \varepsilon)^{-1}$ and hence eq. (92). □

Thus the local recovery error is controlled by the accumulated window mismatch e_k and ridge bias, amplified by the conditioning of the realized finite-window Gramian.

B.7.2 Finite-Continuation Sampling

For nonlinear systems, the finite-window Gramian is generally path dependent, so inversion and conditional expectation do not commute:

$$\mathbb{E}\big[(W + \varepsilon I)^{-1} r \mid X_s = x\big] \neq \left(\mathbb{E}[W \mid X_s = x] + \varepsilon I\right)^{-1} \mathbb{E}[r \mid X_s = x]$$

in general. FBC therefore forms continuation-level pathwise lifts before regression rather than pooling random Gramians prior to inversion. To record the effect of repeated continuations, let $\widehat{u}_1, \dots, \widehat{u}_N$ be conditionally i.i.d. pathwise lifts at a fixed anchor, define $\overline{u}_N = N^{-1} \sum_{i=1}^N \widehat{u}_i$, and write $b_\varepsilon = \mathbb{E}[\widehat{u}_i - u^\star(s,x) \mid X_s = x]$ and $\sigma_\varepsilon^2 = \mathbb{E}[\|\widehat{u}_i - u^\star(s,x) - b_\varepsilon\|^2 \mid X_s = x]$. Then

$$\mathbb{E}\big[\, \|\overline{u}_N - u^\star(s,x)\|^2 \,\big|\, X_s = x\big] = \|b_\varepsilon\|^2 + \frac{\sigma_\varepsilon^2}{N}.$$

The same decomposition holds after projection by $G_s^\top$. Thus additional independent continuations reduce centered Monte Carlo variance at the usual N^{-1} rate but do not remove population or ridge bias.

B.7.3 Auxiliary Transport and Temporal Discretization

We record two standard perturbation estimates that connect the continuous finite-window geometry to the transport and discretization used by FBC. For a fixed trajectory on $[s,t]$, let $Z_r^q = M_{s,r}^q G_r$ and $W_{s,t}^q = \int_s^t Z_r^q (Z_r^q)^\top \, dr$, $q \in \{a, b\}$. If, with $h = t - s$,

$$\sup_{r \in [s,t]} \|M_{s,r}^a - M_{s,r}^b\|_{\mathrm{op}} \le \delta, \qquad \sup_{r \in [s,t]} \|M_{s,r}^b\|_{\mathrm{op}} \le K, \qquad \sup_{r \in [s,t]} \|G_r\|_{\mathrm{op}} \le g,$$

then the product-difference bound gives

$$\|W_{s,t}^a - W_{s,t}^b\|_{\mathrm{op}} \le h g^2 \delta (2K + \delta). \tag{93}$$

Hence, if $W_{s,t}^b \succeq \mu I$ and the right-hand side of eq. (93) is smaller than μ, the auxiliary Gramian remains positive definite. For temporal discretization, let A and G be constant, $K_w \Delta t = h$, and

$$W_h^{\Delta t} = \sum_{j=0}^{K_w - 1} \Delta t \, e^{-Aj\Delta t} G G^\top e^{-A^\top j \Delta t}.$$

With $F(r) = e^{-Ar} G G^\top e^{-A^\top r}$ and $L_F = \sup_{0 \le r \le h} \|F'(r)\|_{\mathrm{op}}$, the left-endpoint quadrature error satisfies

$$\|W_h^{\Delta t} - W_h^-(A, G)\|_{\mathrm{op}} \le \frac{h L_F}{2} \Delta t. \tag{94}$$

Thus $W_h^{\Delta t} \to W_h^-(A,G)$ at first order for fixed h, and any strictly positive finite-window spectral gap is preserved for sufficiently small Δt. Perturbations of the transported directions contribute an additional product-difference error of the same form, so uniformly first-order transport accuracy preserves the overall $O(\Delta t)$ Gramian approximation.

B.7.4 Backward-Correction Stability

We now propagate local finite-window recovery errors through the short backward correction in eq. (13). In the additive-diffusion setting, write $b = a + c$ and $A_\ell = \nabla a(t_\ell, X_{t_\ell})$, $C_\ell = \nabla c(t_\ell, X_{t_\ell})$.

Local correction accuracy. Suppose the reference and auxiliary one-step tangents satisfy $J_\ell^b = I + \Delta t(A_\ell + C_\ell) + R_\ell^b$ and $J_\ell^a = I + \Delta t A_\ell + R_\ell^a$, with $\|R_\ell^a\|_{\mathrm{op}} \le K_a \Delta t^2$ and $\|R_\ell^b\|_{\mathrm{op}} \le K_b \Delta t^2$. If $\|H_\ell\| \le M_H$, $\|C_\ell\|_{\mathrm{op}} \le C_0$, and $\|H_\ell - u_\ell\| \le L_H \Delta t$, then

$$\left\|(J_\ell^a)^\top H_\ell + \Delta t\, C_\ell^\top u_\ell - (J_\ell^b)^\top H_\ell\right\| \le \left[C_0 L_H + (K_a + K_b) M_H\right] \Delta t^2.$$

Indeed, the difference equals $\Delta t\, C_\ell^\top (u_\ell - H_\ell) + (R_\ell^a - R_\ell^b)^\top H_\ell$. Thus the explicit $C_\ell^\top u_\ell$ term cancels the first-order drift component omitted from the auxiliary tangent whenever $H_\ell - u_\ell = O(\Delta t)$. For the underdamped split in eq. (78), bounded $\nabla^2 V$ gives the same $O(\Delta t^2)$ local correction error; this remains true for the exponential auxiliary tangent $J_\ell^a = e^{A\Delta t}$ because $e^{A\Delta t} = I + \Delta t A + O(\Delta t^2)$.

Propagation through the correction horizon. Fix an anchor k and let $K = k + K_c$. Let ideal lifted covectors $u_\ell^\dagger$ and reference tangents define

$$F_K^\dagger = u_K^\dagger, \qquad F_\ell^\dagger = (J_\ell^b)^\top \left[\beta_\ell F_{\ell+1}^\dagger + (1-\beta_\ell) u_{\ell+1}^\dagger\right],$$

while the implemented recursion is

$$\widehat{F}_K = \widehat{u}_K, \qquad \widehat{F}_\ell = (J_\ell^a)^\top \left[\beta_\ell \widehat{F}_{\ell+1} + (1-\beta_\ell)\widehat{u}_{\ell+1}\right] + \Delta t\, C_\ell^\top \widehat{u}_\ell.$$

Define $E_\ell^u = \|\widehat{u}_\ell - u_\ell^\dagger\|$, $L_\ell = \|J_\ell^a\|_{\mathrm{op}}$, $a_\ell = \beta_\ell L_\ell$, $H_\ell^\dagger = \beta_\ell F_{\ell+1}^\dagger + (1-\beta_\ell) u_{\ell+1}^\dagger$, and

$$d_\ell^{\mathrm{corr}} = (J_\ell^a)^\top H_\ell^\dagger + \Delta t\, C_\ell^\top u_\ell^\dagger - (J_\ell^b)^\top H_\ell^\dagger, \qquad D_\ell = \|d_\ell^{\mathrm{corr}}\|.$$

For $r \ge k$, set $\Pi_{k,r} = \prod_{\ell=k}^{r-1} a_\ell$ and $\Pi_{k,k} = 1$.

Theorem B.12 (Discrete FBC target error)**.** *The correction error at the anchor satisfies*

$$\|\widehat{F}_k - F_k^\dagger\| \le \Pi_{k,K} E_K^u + \sum_{r=k}^{K-1} \Pi_{k,r} \left[L_r(1-\beta_r) E_{r+1}^u + \Delta t\, \|C_r\|_{\mathrm{op}} E_r^u + D_r\right]. \tag{95}$$

Consequently,

$$\left\|G_k^\top \widehat{F}_k - G_k^\top F_k^\dagger\right\| \le \|G_k\|_{\mathrm{op}} \times \textit{the right-hand side of eq. (95)}. \tag{96}$$

Proof. Subtracting the ideal and implemented recursions gives

$$\|\widehat{F}_\ell - F_\ell^\dagger\| \le a_\ell \|\widehat{F}_{\ell+1} - F_{\ell+1}^\dagger\| + L_\ell(1-\beta_\ell) E_{\ell+1}^u + \Delta t\, \|C_\ell\|_{\mathrm{op}} E_\ell^u + D_\ell.$$

Since $\|\widehat{F}_K - F_K^\dagger\| = E_K^u$, iteration proves eq. (95), and projection by $G_k^\top$ gives eq. (96). □

If $K_c \Delta t \le H_c$, $L_\ell \le 1 + L\Delta t$, $\|C_\ell\|_{\mathrm{op}} \le C$, $\sum_{\ell=k}^{K-1}(1-\beta_\ell) \le B_\beta$, $\max_{k \le \ell \le K} E_\ell^u \le B_u$, and $D_\ell \le D\Delta t^2$, then $\Pi_{k,r} \le e^{LH_c}$ and Theorem B.12 gives

$$\left\|G_k^\top \widehat{F}_k - G_k^\top F_k^\dagger\right\| \le \|G_k\|_{\mathrm{op}} e^{LH_c} \left[(1 + B_\beta + CH_c) B_u + D H_c \Delta t\right].$$

Hence the backward correction propagates local lift errors by a bounded factor and accumulates $O(\Delta t^2)$ local correction errors to $O(\Delta t)$ over a fixed physical correction horizon.

Kinetic scaled bound. For the underdamped split, let $S_h = \mathrm{diag}(hI, I)$, where h is the recovery-window duration, and use $J_\ell^a = e^{A\Delta t}$. Suppose $\Delta t \le h$, $K_c \Delta t \le c_h h$, $\|\nabla^2 V(q_\ell)\|_{\mathrm{op}} \le L_2$, $\sum_{\ell=k}^{K-1}(1-\beta_\ell) \le B_\beta$,

$$\max_{k \le \ell \le K} \|S_h(\widehat{u}_\ell - u_\ell^\dagger)\| \le \mathcal{E}_h, \qquad \|S_h d_\ell^{\mathrm{corr}}\| \le D\Delta t^2.$$

Then there exists a constant C_{rec}, independent of sufficiently small h and Δt, such that

$$\left\|G^\top(\widehat{F}_k - F_k^\dagger)\right\| \le |\sigma| C_{\mathrm{rec}} \left(\mathcal{E}_h + Dh\Delta t\right). \tag{97}$$

Indeed, $\|S_h (J_\ell^a)^\top S_h^{-1}\|_{\mathrm{op}} \le 1 + C_A \Delta t/h$ and $\|S_h C_\ell^\top S_h^{-1}\|_{\mathrm{op}} \le hL_2$, so products of the scaled tangent factors remain uniformly bounded over $K_c \Delta t = O(h)$. Applying the scaled version of the recurrence in the proof of Theorem B.12 gives an $O(\mathcal{E}_h + Dh\Delta t)$ bound in the S_h norm, and $\|G^\top z\| \le |\sigma| \|S_h z\|$ yields eq. (97).

B.7.5 Relaxation and Inexact Outer Updates

We finally isolate the effect of the relaxation parameter in the outer FBC update. Let $\mathcal{X}$ be a Banach space, let $\mathcal{T} : \mathcal{X} \to \mathcal{X}$ denote an exact consistency operator, and for $0 < \rho \leq 1$ define $\mathcal{T}_\rho = (1-\rho)I + \rho\mathcal{T}$.

Proposition B.13 (Relaxation and inexact outer updates)**.** *For every $0 < \rho \leq 1$, $\mathrm{Fix}(\mathcal{T}_\rho) = \mathrm{Fix}(\mathcal{T})$. If $\mathcal{T}$ is q-contractive with $0 \leq q < 1$ and fixed point $z^\star$, then $\mathcal{T}_\rho$ is q_ρ-contractive with $q_\rho = 1 - \rho + \rho q < 1$. Moreover, suppose approximate target maps $\widehat{\mathcal{T}}_n$ satisfy $\sup_{z\in\mathcal{X}} \|\widehat{\mathcal{T}}_n z - \mathcal{T} z\|_{\mathcal{X}} \leq \eta_n$, and consider*

$$z_{n+1} = (1-\rho)z_n + \rho\widehat{\mathcal{T}}_n z_n.$$

Then

$$\|z_N - z^\star\|_{\mathcal{X}} \leq q_\rho^N \|z_0 - z^\star\|_{\mathcal{X}} + \rho \sum_{n=0}^{N-1} q_\rho^{N-1-n}\eta_n. \tag{98}$$

Consequently, if $\eta_n \leq \overline{\eta}$ uniformly, then $\limsup_{N\to\infty} \|z_N - z^\star\|_{\mathcal{X}} \leq \overline{\eta}/(1-q)$, while $\eta_n \to 0$ implies $z_N \to z^\star$.

Proof. Since $\mathcal{T}_\rho z - z = \rho(\mathcal{T} z - z)$, the two operators have the same fixed points. Contractivity follows from $\|\mathcal{T}_\rho z - \mathcal{T}_\rho w\|_{\mathcal{X}} \leq (1-\rho+\rho q)\|z-w\|_{\mathcal{X}}$. For the inexact iteration,

$$\|z_{n+1} - z^\star\|_{\mathcal{X}} \leq q_\rho \|z_n - z^\star\|_{\mathcal{X}} + \rho\eta_n.$$

Iterating gives eq. (98). If $\eta_n \leq \overline{\eta}$, summing the geometric series and using $1 - q_\rho = \rho(1-q)$ gives the stated asymptotic bound; if $\eta_n \to 0$, the geometrically weighted error term vanishes. □

Thus relaxation preserves the exact population fixed point while controlling the size of each inexact outer update. In FBC, finite-window recovery, continuation sampling, auxiliary transport, temporal discretization, backward correction, function approximation, and regression optimization can be viewed collectively as producing the approximation error η_n.

Summary of the implementation error chain. Proposition B.11 controls pathwise finite-window recovery; finite-continuation averaging separates population bias from Monte Carlo variance; the perturbation estimates in Appendix B.7.3 control auxiliary-transport and temporal-discretization errors; Theorem B.12 propagates local lift errors through the backward correction; and Proposition B.13 describes their propagation through the relaxed outer update. Together, these results connect finite-horizon population recovery to the detached regression targets used in Algorithm 1.

C Experimental Details and Additional Results

This appendix provides the benchmark definitions, evaluation details, and supplementary analyses for the experiments in section 6. We first describe the common evaluation conventions, then follow the scientific progression of the main text: continuous rank loss and kinetic OU, canonical full-rank benchmarks, strictly rank-deficient nonlinear dynamics, and constrained molecular transition paths. Additional controllability and stability analyses are provided at the end.

C.1 Common Evaluation Settings and Metrics

Replicates and evaluation. Repeated results from our experiments are reported as the mean ± std over five random seeds.

Rank-deficient NGDB baselines. For the underdamped Double-Well and Müller–Brown tasks, NGDB retains its native linear kinetic auxiliary and cFH backward filter, with executable control acting only through the momentum coordinates. In constrained alanine, NGDB is run on the same constrained dynamics with control restricted to the tangent-velocity subspace and uses the cFH guide described in Appendix C.6. No pseudoinverse or state-space covector reconstruction is introduced into NGDB; the baseline remains defined in its executable Brownian control space. The ridge- and pseudoinverse-based SC1/SC2 variants are used only as diagnostic adaptations in the kinetic-OU

experiment, where they probe the effect of numerically regularizing the singular instantaneous recovery map. Such regularization does not enlarge the state subspace identifiable from instantaneous diffusion geometry, so we do not treat these variants as standard baselines for general nonlinear rank-deficient systems. For the nonlinear rank-deficient tasks, we instead compare against NGDB, whose control is defined directly in executable Brownian space without requiring instantaneous state-space covector recovery.

Stochastic-control and terminal metrics. We evaluate controlled bridge models using the stochastic-control objective

$$J(v) = \mathbb{E}\left[\frac{1}{2}\int_0^T \|v(t, X_t^v)\|^2\, dt + \Psi(X_T^v)\right].$$

For the continuous rank-loss experiments, we report the relative objective

$$J_{\mathrm{rel}} = \frac{J(v)}{J(v^\star)},$$

so that $J_{\mathrm{rel}} = 1$ corresponds to the analytic oracle. For benchmarks with target position q_{tar}, terminal accuracy is measured by

$$\mathrm{RMSE}_q = \left(\mathbb{E}\|q_T - q_{\mathrm{tar}}\|^2\right)^{1/2}.$$

For kinetic OU and the exact-rank-loss ablation, we also report the pathwise control error

$$\mathrm{RMSE}_v = \left(\frac{1}{T}\mathbb{E}\int_0^T \|v_\theta(t, X_t) - v^\star(t, X_t)\|^2\, dt\right)^{1/2}.$$

Hit and basin metrics are benchmark specific and are defined in the corresponding subsections.

KL metrics. For tractable bridge benchmarks such as OU, CIR, and the one-dimensional overdamped Double-Well, we follow the benchmark-specific evaluation used in the corresponding prior work and report $\mathrm{KL}(P^\star \| P^\theta)$ against the reference bridge. For Cell Diffusion, where no analytic bridge is available, we report $\mathrm{KL}(P^\theta \| P)$ relative to the unconditioned reference process, following the established benchmark convention. Because these KL quantities have different orientations and reference measures, values should be compared only within the same benchmark family. For the kinetic-OU rank-loss diagnostics, "Gaussian KL" denotes the KL from the moment-matched Gaussian of the learned terminal samples to the analytic Doob terminal Gaussian.

Molecular transition-path metric. For alanine dipeptide, we report transition-state energy (ETS) in $\mathrm{kJ\,mol^{-1}}$.

Numerical rank. For finite-window controllability analysis, we define the numerical rank of a positive-semidefinite Gramian W in the fixed tangent-coordinate representation by

$$\mathrm{rank}_{\mathrm{num}}(W) = \#\left\{i : \lambda_i(W) > \max\left(10^{-18}, 10^{-10}\lambda_{\max}(W)\right)\right\}. \tag{99}$$

Here 10^{-18} is an absolute numerical floor, whereas $10^{-10}\lambda_{\max}(W)$ is the relative spectral threshold. We use this numerical rank solely for the reported finite-window analysis. Whenever layer scaling is reported, the scaling exponents are obtained by linear fits in log–log coordinates using the corresponding Gramian eigenvalue or layer statistic and the physical window duration h.

Table 5 summarizes the principal evaluation settings used in this work.

For the rank-deficient kinetic OU, underdamped Double-Well, and underdamped Müller–Brown experiments, FBC uses $K_w = K_c = 8$; constrained-underdamped alanine uses $K_w = K_c = 4$ as specified in Appendix C.6.3.

C.2 Continuous Rank Loss and Kinetic OU

Kinetic OU provides an analytically tractable setting in which instantaneous diffusion rank can be removed while finite-time controllability is retained. We first evaluate FBC at strict rank deficiency against the exact Doob controller and instantaneous control-consistency baselines. We then vary the instantaneous diffusion rank continuously and isolate the respective contributions of temporal pooling and dynamical transport.

Table 5: Principal evaluation settings for experiments in this work.

Benchmark	State dim. / rank(G)	Seeds	Numerical steps
Continuous rank-loss OU	$2/1$ at $\eta = 0$	5	256
Kinetic OU	$2/1$	5	256
Underdamped Double-Well	$2/1$	5	256
Underdamped Müller–Brown	$4/2$	5	256
CIR	$1/1$	5	500
Overdamped Double-Well	$1/1$	5	500
Cell Diffusion	$2/2$	5	400
Full-rank OU sweeps	d/d (dimension sweep)	5	500
Constrained underdamped alanine	$90/45$	5	30,000

C.2.1 KINETIC OU BRIDGE

We consider the rank-deficient kinetic OU bridge

$$dq_t = p_t\,dt, \qquad dp_t = (-q_t - p_t + v_t)\,dt + dW_t,$$

or, with $X_t = (q_t, p_t)^\top$,

$$dX_t = (AX_t + Gv_t)\,dt + G\,dW_t, \qquad A = \begin{pmatrix} 0 & 1 \\ -1 & -1 \end{pmatrix}, \qquad G = \begin{pmatrix} 0 \\ 1 \end{pmatrix}.$$

We use $(q_0, p_0) = (-1, 0)$ and $T = 2$. The terminal observation acts only on position, with $H = (1, 0)$, target $y = 1$, and observation variance $R = 0.04$. For $\tau = T - t$, define $\Phi_\tau = e^{A\tau}$ and

$$Q_\tau = \int_0^\tau e^{As} G G^\top e^{A^\top s}\,ds.$$

The exact continuous-time Doob controller is

$$v^\star(t, x) = G^\top \Phi_\tau^\top H^\top \left(R + HQ_\tau H^\top\right)^{-1} (y - H\Phi_\tau x)\,. \tag{100}$$

We report the stochastic-control objective, terminal position RMSE, control RMSE RMSE_v relative to eq. (100), and the terminal hit rate

$$\mathrm{Hit} = \mathbb{P}\Big(|q_T - y| \le 2\sqrt{R}\Big)\,,$$

which corresponds to threshold 0.4 for $R = 0.04$.

The exact Doob controller obtains $J(v^\star) = 2.701 \pm 0.035$ and $\mathrm{RMSE}_q = 0.216 \pm 0.002$. FBC obtains $J(v_{\mathrm{FBC}}) = 2.905 \pm 0.032$ and $\mathrm{RMSE}_q = 0.220 \pm 0.003$, with control RMSE 0.455 ± 0.003 and hit rate 0.935 ± 0.004.

To determine whether numerical regularization alone can repair instantaneous recovery, we additionally evaluate the SC1 and SC2 instantaneous control-consistency constructions using ridge inverses and the Moore–Penrose pseudoinverse. We retain the SC1/SC2 terminology of the corresponding CCDB baselines. The complete comparison is reported in Table 6.

The best instantaneous ridge result is 20.531 ± 1.045, more than seven times the FBC objective. Varying the ridge strength changes numerical conditioning while leaving the missing position direction unobserved; finite-time transport instead enlarges the identifiable subspace.

Source of the residual oracle gap. Figure 4 decomposes the stochastic-control objective into control energy and terminal cost. The exact Doob objective decomposes as $2.117 + 0.584 = 2.701$, whereas the FBC objective decomposes as $2.298 + 0.607 = 2.905$. Relative to the exact Doob controller, FBC has approximately $1.085\times$ the control energy, $1.039\times$ the terminal cost, and $1.019\times$ the terminal q-RMSE. Thus, the remaining objective gap is dominated by control energy, while the terminal cost and terminal q-RMSE remain close to the Doob values.

Table 6: Kinetic OU at strict rank deficiency (mean ± std over five random seeds). The final column gives simultaneous 95% Bonferroni confidence intervals for $\Delta J = J(\text{baseline}) - J(\text{FBC})$.

Method	Objective ↓	q-RMSE ↓	Control RMSE ↓	Hit ↑	FWER CI for ΔJ
FBC (Ours)	$\mathbf{2.905 \pm 0.032}$	$\mathbf{0.220 \pm 0.003}$	$\mathbf{0.455 \pm 0.003}$	$\mathbf{0.935 \pm 0.004}$	–
SC1 + ridge ($\varepsilon = 0.1$)	21.368 ± 0.896	1.076 ± 0.022	4.353 ± 0.102	0.169 ± 0.015	$[16.39, 20.53]$
SC1 + ridge ($\varepsilon = 0.03$)	21.326 ± 1.240	1.077 ± 0.028	4.347 ± 0.137	0.166 ± 0.011	$[15.55, 21.29]$
SC1 + ridge ($\varepsilon = 0.01$)	20.531 ± 1.045	1.058 ± 0.022	4.254 ± 0.122	0.172 ± 0.012	$[15.19, 20.07]$
SC1 + pseudoinverse	24.220 ± 0.565	1.150 ± 0.012	4.673 ± 0.059	0.132 ± 0.004	$[19.99, 22.64]$
SC2 + ridge ($\varepsilon = 0.1$)	218.892 ± 12.224	3.462 ± 0.103	14.780 ± 0.416	0.002 ± 0.001	$[187.20, 244.77]$
SC2 + ridge ($\varepsilon = 0.03$)	212.125 ± 13.566	3.411 ± 0.130	14.546 ± 0.461	0.001 ± 0.001	$[177.32, 241.12]$
SC2 + ridge ($\varepsilon = 0.01$)	191.492 ± 11.629	3.214 ± 0.091	13.813 ± 0.427	0.002 ± 0.001	$[161.21, 215.97]$
SC2 + pseudoinverse	132.144 ± 11.764	2.563 ± 0.094	11.438 ± 0.525	0.008 ± 0.002	$[101.55, 156.93]$

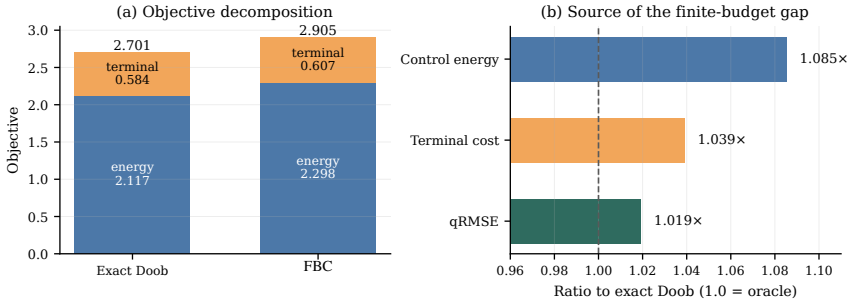


Figure 4: Objective decomposition for kinetic OU at exact rank deficiency. Left: control-energy and terminal-cost contributions for the exact Doob controller and FBC. Right: FBC-to-Doob ratios for control energy, terminal cost, and terminal q-RMSE.

C.2.2 Continuous Rank Loss

To isolate the transition from full-rank to rank-deficient diffusion geometry, we keep the drift, terminal observation, model class, and evaluation objective of kinetic OU fixed and replace the diffusion by

$$G_\eta = \begin{pmatrix} \eta & 0 \\ 0 & 1 \end{pmatrix}, \qquad \eta \in \{1, 0.3, 0.1, 0.03, 0.01, 0\}. \tag{101}$$

For every $\eta > 0$, G_η has full row rank and pointwise recovery is well posed. At $\eta = 0$, position noise disappears and the instantaneous diffusion rank drops from two to one, while the drift matrix A remains controllable with the remaining momentum direction.

We compare three recovery constructions at each value of η: *Pointwise* uses instantaneous covector recovery from the current diffusion geometry; *Finite window* applies finite-horizon recovery at every value of η; and *FBC* uses pointwise recovery while G_η has full row rank and finite-horizon recovery at the singular point $\eta = 0$.

Objectives in Table 7 are normalized by the exact Doob objective under the same value of η.

Table 7: Continuous rank-loss sweep (mean ± std over five random seeds). Objective values are normalized by the corresponding exact Doob objective.

η	FBC	Finite window	Pointwise	$\lambda_{\min}(W)$	FBC recovery
1	$\mathbf{1.002 \pm 0.001}$	1.034 ± 0.001	1.002 ± 0.001	1.22×10^{-1}	Pointwise
0.3	$\mathbf{1.004 \pm 0.001}$	1.055 ± 0.002	1.004 ± 0.001	1.14×10^{-2}	Pointwise
0.1	$\mathbf{1.004 \pm 0.003}$	1.060 ± 0.002	1.004 ± 0.003	1.41×10^{-3}	Pointwise
0.03	$\mathbf{1.005 \pm 0.003}$	1.069 ± 0.002	1.005 ± 0.003	2.72×10^{-4}	Pointwise
0.01	$\mathbf{1.005 \pm 0.004}$	1.075 ± 0.003	1.005 ± 0.004	1.72×10^{-4}	Pointwise
0	$\mathbf{1.076 \pm 0.003}$	$\mathbf{1.076 \pm 0.003}$	6.316 ± 0.171	1.59×10^{-4}	Finite window

For every $\eta > 0$, exact pointwise recovery remains within 0.52% of the oracle objective despite the progressively smaller instantaneous position-noise scale. The qualitative failure occurs at the singular point itself: when $\eta = 0$, the pointwise continuation jumps to 6.316 ± 0.171, whereas finite-horizon recovery remains at 1.076 ± 0.003.

At $\eta = 0$, the remaining forced direction is $g = (0, 1)^\top$, while its drift image is $Ag = (1, -1)^\top$. Hence $[g, Ag]$ has rank two, and Theorem 5.1 implies $W_h^-(A, G_0) \succ 0$ for every $h > 0$. Thus the experiment separates poor instantaneous conditioning from exact loss of instantaneous identifiability: the former does not by itself destroy pointwise recovery, whereas the latter requires dynamical transport.

C.2.3 Temporal Pooling versus Dynamical Transport

At exact rank loss, we further separate the effect of pooling controls across future times from the effect of dynamically transporting the corresponding Brownian directions. All variants use the same rank-deficient system with $\eta = 0$. The *instantaneous pseudoinverse* uses only the current rank-one diffusion direction; the *identity window* pools future controls without transporting their directions and therefore remains rank one; the *one-step transported lift* introduces one drift-transported future direction; and the *full finite window* uses the complete FBC transported-window construction. The resulting ablation is summarized in Table 8.

Table 8: Recovery ablation at exact rank loss ($\eta = 0$) (mean $\pm$ std over five random seeds). The identity window adds temporal observations without changing the recovery rank, whereas transported windows expose the missing state direction.

Method	Obj./oracle ↓	q-RMSE ↓	Control RMSE ↓	Terminal Gaussian KL ↓	Lift rank
Instantaneous pseudoinverse	6.316 ± 0.171	1.020 ± 0.012	2.700 ± 0.046	12.387 ± 0.441	1
Identity window	1.110 ± 0.003	0.293 ± 0.003	0.394 ± 0.005	0.228 ± 0.015	1
One-step transported lift	$\mathbf{1.071 \pm 0.002}$	0.264 ± 0.002	$\mathbf{0.315 \pm 0.004}$	0.107 ± 0.006	2
Full finite window	1.076 ± 0.003	$\mathbf{0.222 \pm 0.002}$	0.321 ± 0.003	$\mathbf{0.051 \pm 0.004}$	2

The identity window improves substantially over the instantaneous pseudoinverse while remaining rank one, reflecting the benefit of temporal pooling within the directly forced subspace. Both transported constructions raise the lift rank from one to two and expose the missing position direction through drift transport.

The one-step transported lift attains the lowest scalar control objective and control RMSE among the transported variants, whereas the full finite window gives the lowest terminal q-RMSE and Gaussian terminal KL. The ablation therefore separates two effects: temporal pooling improves estimation within the instantaneously observable subspace, while dynamical transport changes the identifiable state-space geometry itself. The transported variants exhibit different tradeoffs between control and terminal metrics as the recovery window changes.

C.3 Full-Rank Benchmarks

The benchmarks in this section evaluate the FBC learning construction in full-rank settings using the full-row-rank recovery procedure described in Appendix A.2, followed by the corresponding detached control-consistency update. The benchmark-specific KL conventions follow Appendix C.1.

C.3.1 Ornstein–Uhlenbeck Bridges

We use the standard isotropic Ornstein–Uhlenbeck bridge

$$dX_t = -\alpha X_t \, dt + \sigma \, dW_t.$$

For the canonical two-dimensional configuration, we set $\alpha = 2$, $\sigma = 0.1$, $T = 1$, $x_0 = (1, 0)$, and $x_T = (0, 1)$. Because the transition law is Gaussian, the exact conditioned process is available analytically and performance is evaluated using $\mathrm{KL}(P^\star \| P^\theta)$.

Figure 2 in the main text evaluates optimization progress in the canonical two-dimensional problem, scaling with the state dimension, and sensitivity to the drift parameter α. Since $G = \sigma I$ has full row

rank, the pointwise state covector is identifiable at every time; this experiment evaluates the FBC learning construction in a canonical full-rank setting.

C.3.2 Cox–Ingersoll–Ross Model

The one-dimensional Cox–Ingersoll–Ross process is

$$dX_t = a(b - X_t)\,dt + \sigma_{\mathrm{CIR}}\sqrt{X_t}\,dW_t,$$

with $a = b = \sigma_{\mathrm{CIR}} = 1$, $T = 1$, and $x_0 = x_T = 2$, following the bridge configuration used in prior control-consistency experiments (Cox et al., 1985; Howard et al., 2026). The CIR transition density is available in closed form through the noncentral chi-squared representation, equivalently through the modified-Bessel transition density, so the conditioned process remains tractable and performance is evaluated using $\mathrm{KL}(P^\star \| P^\theta)$.

CIR differs from the constant-diffusion benchmarks because its diffusion coefficient depends on the state. Accordingly, the state-dependent one-step tangent construction described in Appendix A.2 is retained when constructing the control-consistency target. The corresponding quantitative results are reported in Table 1 of the main text.

C.3.3 Overdamped Double-Well

The standard full-rank metastable benchmark is the one-dimensional overdamped Langevin diffusion

$$dX_t = -U'(X_t)\,dt + \sigma\,dW_t, \qquad U(x) = v(x^2 - 1)^2.$$

We use $v = 3$, $\sigma = 1$, $T = 1$, $x_0 = 1$, and $x_T = -1$, following the benchmark used in previous neural diffusion-bridge studies (Pidstrigach et al., 2025; Howard et al., 2026). Because the problem is one dimensional, the backward Kolmogorov equation can be solved numerically to obtain a reference conditioned process, and performance is evaluated using $\mathrm{KL}(P^\star \| P^\theta)$. Evaluation uses a 500-step numerical grid.

The corresponding quantitative results are reported in Table 1 of the main text. The underdamped Double-Well experiment in Appendix C.4 retains the same qualitative two-well transition geometry but introduces momentum and momentum-only actuation; the two benchmarks therefore test distinct diffusion geometries.

C.3.4 Cell Diffusion

We use the two-dimensional non-gradient differentiation model adopted in prior diffusion-bridge studies,

$$dX_t = b_{\mathrm{cell}}(X_t)\,dt + \sigma\,dW_t, \qquad \sigma = 0.1,$$

where, for $x = (x_1, x_2)$ and $\delta = 2^{-4}$,

$$b_{\mathrm{cell},1}(x) = \frac{x_1^4}{\delta + x_1^4} + \frac{\delta}{\delta + x_2^4} - x_1,$$

$$b_{\mathrm{cell},2}(x) = \frac{x_2^4}{\delta + x_2^4} + \frac{\delta}{\delta + x_1^4} - x_2.$$

We use the established Normal, Rare, and Multimodal tasks. The Normal and Rare tasks start from $x_0 = (0.1, -0.1)$ and target $x_T = (2.0, -0.1)$ and $x_T = (1.0, -0.1)$, respectively, with $T = 4$. The Multimodal task starts and ends at $(-1, -1)$ with $T = 5$.

No analytic bridge is available for this benchmark. Following the established benchmark convention, we report $\mathrm{KL}(P^\theta \| P)$ relative to the unconditioned reference process rather than $\mathrm{KL}(P^\star \| P^\theta)$. The corresponding results for the Normal, Rare, and Multimodal regimes are reported in Table 1 of the main text.

C.3.5 Full-Rank Müller–Brown

The Müller–Brown potential provides a standard two-dimensional transition-path benchmark. The full-rank experiment follows the established overdamped Müller–Brown setting used in the comparison source (Du et al., 2024; Howard et al., 2026). Unlike the underdamped experiment in

Appendix C.5, stochastic forcing acts directly in the full configuration space, so the instantaneous diffusion geometry used for recovery is full rank.

Performance is reported using the KL metric of the established full-rank benchmark, with quantitative results given in Table 2 of the main text.

The underdamped Müller–Brown experiment in Appendix C.5 instead augments the state with momentum and restricts stochastic forcing and control to momentum, producing a strictly rank-deficient diffusion. It is therefore evaluated using stochastic-control and terminal-transition metrics rather than the full-rank KL metric.

C.4 Underdamped Double-Well

The underdamped Double-Well experiment tests FBC on a nonlinear metastable system with strict instantaneous rank deficiency. We use the quartic potential $V_{\mathrm{DW}}(q) = \kappa(q^2-1)^2$ with $\kappa = 2$ and controlled dynamics

$$dq_t = p_t\,dt, \qquad dp_t = [-V'_{\mathrm{DW}}(q_t) - \gamma p_t + \sigma v_t]\,dt + \sigma\,dW_t.$$

We set $T = 2$, $(q_0, p_0) = (-1, 0)$, $q_{\mathrm{target}} = 1$, $\gamma = \sigma = 1$, and $R_q = 0.04$. Brownian forcing and executable control act only through momentum, corresponding to diffusion matrix $G = (0,1)^\top$. The instantaneous diffusion rank is therefore one in a two-dimensional state space, while the kinetic coupling $dq_t = p_t\,dt$ transports the forced momentum direction into position over finite time. The FBC controller is a three-hidden-layer width-128 SiLU MLP with 33,665 parameters; the matched NGDB baseline uses the same parameter count.

Evaluation. We use 256 numerical time steps. We report the stochastic-control objective together with its control-energy and terminal-cost components and terminal position RMSE. The terminal hit rate is

$$\mathrm{Hit} = \mathbb{P}\Big(|q_T - q_{\mathrm{target}}| \le 2\sqrt{R_q}\Big)\,,$$

corresponding to threshold 0.4 for $R_q = 0.04$. We also report the basin statistic $\mathrm{Basin} = \mathbb{P}(q_T \ge 0)$, which records whether the final state lies in the target-side well. The complete objective decomposition and terminal metrics are reported in Table 9.

Table 9: Underdamped Double-Well results. Lower is better for Objective, Energy, Terminal, and q-RMSE; higher is better for Hit and Basin.

Method	Objective ↓	Energy ↓	Terminal ↓	q-RMSE ↓	Hit ↑	Basin ↑
NGDB	8.8493 ± 0.1422	7.9539 ± 0.3472	0.8954 ± 0.4216	0.2614 ± 0.0640	0.8684 ± 0.1089	0.99988 ± 0.00021
FBC (Ours)	$\mathbf{8.2455 \pm 0.0532}$	$\mathbf{7.5905 \pm 0.0792}$	$\mathbf{0.6551 \pm 0.0385}$	$\mathbf{0.2288 \pm 0.0068}$	$\mathbf{0.9189 \pm 0.0071}$	$\mathbf{1.00000 \pm 0.00000}$

FBC lowers both the control-energy and terminal-cost components. The terminal metrics show the same trend: FBC reduces terminal position RMSE and increases the hit rate. Both methods almost always finish on the target side of the barrier, however, so the basin statistic is nearly saturated and is less discriminative than terminal RMSE and hit rate in this configuration.

Relation to finite-time geometry. Instantaneously, the diffusion exposes only the momentum covector and pointwise recovery cannot identify the position component. Over a finite window, the kinetic coupling transports this forced direction into position. The experiment therefore provides a nonlinear realization of the two-layer controllability geometry analyzed in Appendix B.3, where the directly forced momentum layer appears at order h and the position layer generated through one drift coupling appears at order h^3.

C.5 Underdamped Müller–Brown

The underdamped Müller–Brown experiment extends the nonlinear rank-deficient evaluation from a one-dimensional metastable coordinate to a two-dimensional transition landscape. We use the scaled Müller–Brown potential

$$V_{\mathrm{MB}}(x,y) = s_V \sum_{i=1}^{4} A_i \exp(\phi_i(x,y))\,, \qquad \phi_i(x,y) = a_i(x-\bar{x}_i)^2 + b_i(x-\bar{x}_i)(y-\bar{y}_i) + c_i(y-\bar{y}_i)^2,$$

with scale factor $s_V = 0.02$. The potential parameters are listed in Table 10.

Table 10: Parameters of the Müller–Brown potential.

i	1	2	3	4
A_i	-200	-100	-170	15
a_i	-1	-1	-6.5	0.7
b_i	0	0	11	0.6
c_i	-10	-10	-6.5	0.7
$\bar{x}_i$	1	0	-0.5	-1
$\bar{y}_i$	0	0.5	1.5	1

The controlled state $X_t = (q_t, p_t) \in \mathbb{R}^4$ evolves according to

$$dq_t = p_t\,dt, \qquad dp_t = [-\nabla V_{\mathrm{MB}}(q_t) - \gamma p_t + \sigma v_t]\,dt + \sigma\,dW_t.$$

We set $\gamma = \sigma = 1$ and $T = 2.5$, with $q_0 = (-0.558, 1.442)$, $p_0 = (0, 0)$, and $q_{\mathrm{target}} = (0.623, 0.028)$. The terminal penalty is $\Psi(q_T) = \|q_T - q_{\mathrm{target}}\|^2/(2R_q)$ with $R_q = 0.04$. Brownian forcing and executable control act only through the two momentum coordinates, corresponding to $G = (0_{2\times 2}, I_2)^\top$. Thus, the instantaneous diffusion rank is two in a four-dimensional state space, while the kinetic coupling $dq_t = p_t\,dt$ transports the forced momentum directions into position over finite time.

Evaluation. We use 256 numerical time steps. In addition to the stochastic-control objective and terminal position RMSE, we report two transition-success metrics. The terminal hit rate is $\mathrm{Hit} = \mathbb{P}(\|q_T - q_{\mathrm{target}}\| \le 2\sqrt{R_q})$, corresponding to radius 0.4 for $R_q = 0.04$. The basin-entry rate is $\mathrm{Basin} = \mathbb{P}(\min_{0\le t\le T} \|q_t - q_{\mathrm{target}}\| \le 0.35)$. Hit measures final proximity to the target, whereas Basin records whether the trajectory enters the target neighborhood at any point along the transition. The corresponding quantitative results are reported in Table 3 of the main text.

Relation to finite-time recovery. The experiment has the same two-layer kinetic structure as the underdamped Double-Well, but with two position and two momentum coordinates and a strongly nonlinear two-dimensional potential. Instantaneously, executable controls expose only the two momentum covector components. Over a finite window, kinetic transport propagates these directions into position and provides the rank-completing geometry used by FBC. This experiment therefore tests whether the finite-time recovery mechanism persists beyond affine dynamics and the one-dimensional metastable setting.

C.6 Alanine Dipeptide

We evaluate FBC on the alanine dipeptide C5→C7ax transition in both the established overdamped setting and a constrained-underdamped rank-deficient setting.

C.6.1 Benchmark Provenance and Overdamped Setting

We consider the 22-atom alanine dipeptide C5→C7ax transition used in prior molecular transition-path studies (Holdijk et al., 2023; Seong et al., 2025; Howard et al., 2026). The molecular system follows the setup of Howard et al. (2026), using the AMBER14/ff14SB force field through OpenMM and DMFF for differentiable force evaluation. The C5 and C7ax endpoint conformations follow Seong et al. (2025) and are aligned using the same Kabsch-alignment procedure.

Performance is measured by transition-state energy (ETS), reported in $\mathrm{kJ\,mol^{-1}}$, using the established molecular transition-path evaluator. The corresponding quantitative results are reported in Table 2 of the main text.

This overdamped experiment evaluates FBC in the established alanine transition-path setting. The constrained-underdamped experiment below instead modifies the dynamics to expose strict instantaneous rank deficiency and is analyzed separately.

C.6.2 CONSTRAINED-UNDERDAMPED GEOMETRY

To obtain strict instantaneous rank deficiency in the molecular setting, we use constrained underdamped dynamics with Cartesian position and velocity states. AllBonds constraints are enforced, and Brownian forcing and executable control act only in the tangent-velocity subspace. Alanine dipeptide has 22 atoms and hence 66 Cartesian coordinates; after imposing 21 independent bond constraints, the tangent configuration space has dimension 45. The constrained position–velocity state therefore has dimension 90, while the instantaneous effective diffusion rank is only 45.

The missing tangent-position directions become identifiable over a finite window through the kinetic coupling between velocity and position. Under the numerical-rank criterion in eq. (99), a one-time-point window contains the 45 instantaneous Brownian directions and retains numerical rank 45, whereas every tested window with $K_w \geq 2$ reaches numerical rank 90. Thus, finite-time transport raises the recovery geometry from numerical rank 45 to 90 in the constrained tangent state. Detailed rank and short-window scaling analyses are reported in Appendix C.7.2.

C.6.3 CONSTRAINED-UNDERDAMPED COMPARISON

We compare FBC with an NGDB adaptation on the constrained-underdamped C5→C7ax system. FBC and NGDB use the same constrained dynamics, endpoint task, integration grid, and evaluation protocol. Evaluation uses a $1\,\mathrm{ps}$ horizon with 30,000 integration steps. Terminal RMSE is computed from the full 66 Cartesian position coordinates of the terminal configuration relative to the target C7ax conformation and is reported in nm. For FBC, the recovery and correction horizons are fixed at $K_w = K_c = 4$. The FBC residual controller is a width-128 MLP trained for 250 updates at learning rate 10^{-4}, with relaxation $\rho = 0.25$ and eight training anchors. The matched NGDB baseline uses a three-layer width-128 `tanh` MLPSmall, batch size 16, learning rate 10^{-4}, the cFH guide, and the final 800-update iterate. For the five-seed comparison, we use fixed training budgets of 250 updates for FBC and 800 updates for NGDB, and report the final iterate from each run. The corresponding comparison is reported in Table 11.

Table 11: Constrained-underdamped alanine comparison (mean ± std over five random seeds).

Method	ETS (kJ mol^{-1}) ↓	Terminal Cartesian RMSE (nm) ↓
FBC (Ours)	$\mathbf{6.0692 \pm 0.0329}$	$\mathbf{0.029273 \pm 0.000001}$
NGDB	61.0522 ± 3.5259	0.031921 ± 0.000111

FBC lowers both ETS and terminal RMSE. The NGDB-minus-FBC difference in mean ETS is $54.9830\,\mathrm{kJ\,mol}^{-1}$, with a Welch 95% confidence interval of $[50.6051, 59.3609]\,\mathrm{kJ\,mol}^{-1}$.

Table 12: Training configurations for the constrained-underdamped alanine comparison. Wall times were measured on NVIDIA GeForce RTX 4080 hardware and are reported as mean ± std over five random seeds.

Method	Training configuration	Parameters	Train time (s)
FBC (Ours)	250 updates, lr 10^{-4}	50,626	613.4 ± 27.9
NGDB	800 updates, lr 10^{-4}	58,690	2493.0 ± 15.6

Table 12 reports the training configurations, parameter counts, and observed wall-clock times for the constrained-underdamped comparison.

C.7 ADDITIONAL CONTROLLABILITY AND STABILITY ANALYSES

The experiments above evaluate the learned bridge controllers. This section isolates the finite-time geometry itself. We first verify the general h^{2j-1} controllability-layer law on a four-layer linear chain, then examine rank completion and layer scaling in the 90-dimensional constrained alanine tangent state.

C.7.1 Four-Layer Controllability Chain

To verify the small-window layer law beyond the two-layer kinetic setting, we consider the four-dimensional integrator chain

$$\begin{aligned} dx_t^{(0)} &= v_t\,dt + dW_t, \\ dx_t^{(1)} &= x_t^{(0)}\,dt, \\ dx_t^{(2)} &= x_t^{(1)}\,dt, \\ dx_t^{(3)} &= x_t^{(2)}\,dt. \end{aligned}$$

Noise and executable control act only on $x^{(0)}$, so the instantaneous control rank is one. Successive drift couplings expose one additional state direction at each controllability layer, and the Kalman filtration therefore grows from dimension 1 to 2, 3, and finally 4.

For this nilpotent chain, the pullback Gramian of section 5.1 has entries

$$[W_h^-]_{ij} = (-1)^{i+j}\frac{h^{i+j+1}}{i!\,j!\,(i+j+1)}, \qquad i,j \in \{0,1,2,3\}.$$

The alternating signs arise from the pullback factor $e^{-A\tau}$. Writing $D = \mathrm{diag}(1,-1,1,-1)$, this Gramian is related by $W_h^- = D\widetilde{W}_h D$ to the corresponding all-positive moment matrix $\widetilde{W}_h$, and therefore the two matrices have identical eigenvalues. The controllability-layer result of Theorem B.2 predicts

$$\lambda_1(W_h^-) \asymp h, \qquad \lambda_2(W_h^-) \asymp h^3, \qquad \lambda_3(W_h^-) \asymp h^5, \qquad \lambda_4(W_h^-) \asymp h^7.$$

We estimate the scaling exponents by linear fits in log–log coordinates over $h \in \{0.005, 0.008, 0.012, 0.018, 0.027, 0.04, 0.06, 0.09\}$, using both the exact continuous Gramian and a 64-substep Euler approximation. The fitted exponents are reported in Table 13 and visualized in Figure 5.

Table 13: Short-window eigenvalue-scaling exponents for the four-layer integrator chain. Both the exact Gramian and its Euler–64 approximation recover the theoretical $1, 3, 5, 7$ controllability hierarchy.

Estimator	λ_1	λ_2	λ_3	λ_4	$R^2_{\min}$
Theory	1	3	5	7	–
Exact Gramian fit	1.00057	3.00000	5.00000	6.99943	0.99999983
Euler–64 fit	1.00055	2.99998	4.99998	6.99948	0.99999984

The fitted exponents are numerically indistinguishable from the predicted odd-integer sequence. In particular, the deepest controllability direction is identifiable for every positive horizon but enters the Gramian only at order h^7, illustrating the distinction between finite-time identifiability and short-window conditioning. The Euler approximation preserves the same hierarchy: as additional transported time points are included, the discrete controllability span successively exposes all four state directions. Thus, the layer structure underlying FBC survives the temporal discretization used to construct the finite-window recovery operator.

C.7.2 Constrained Alanine Controllability

We next examine the finite-time controllability geometry of the constrained-underdamped alanine system introduced in Appendix C.6.2. The tangent state has dimension 90, comprising 45 tangent-position and 45 tangent-velocity directions. Brownian forcing and executable control act only through the tangent-velocity subspace, so the instantaneous effective diffusion rank is 45. Numerical ranks below use $\mathrm{rank}_{\mathrm{num}}(W)$ from eq. (99), with the resulting rank and scaling summary reported in Table 14.

A one-time-point window contains the 45 instantaneous Brownian directions and spans only the directly forced tangent-velocity subspace. Every tested window with $K_w \geq 2$ reaches numerical rank 90, as transported directions from later time points expose the complementary tangent-position

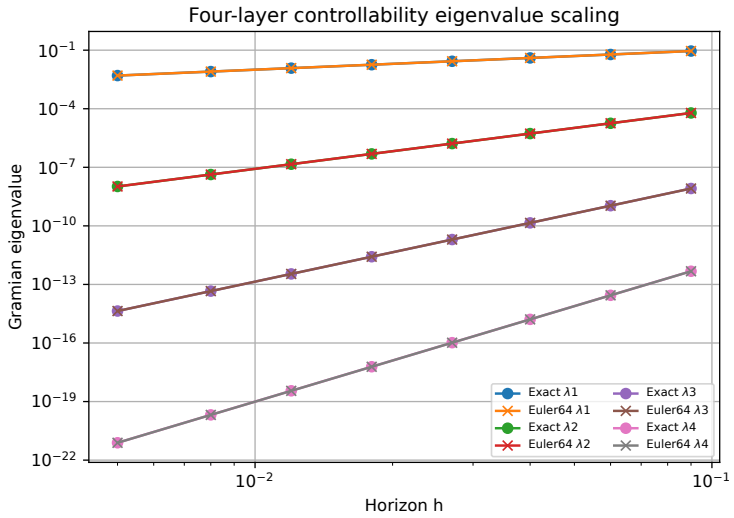


Figure 5: Short-window controllability scaling for the four-layer integrator chain. The exact finite-horizon Gramian and its Euler–64 discretization recover the predicted h, h^3, h^5, and h^7 eigenvalue hierarchy.

Table 14: Finite-window controllability results for constrained-underdamped alanine.

Quantity		Result
Tangent-state dimension		90
Instantaneous effective diffusion rank		45
One-time-point window rank		45
Finite-window rank for every tested $K_w \geq 2$		90
Direct velocity-layer scaling	$h^{1.00538}$	($R^2 = 0.999998$)
Hidden position-layer scaling	$h^{3.12836}$	($R^2 = 0.999477$)

subspace. Finite-window transport therefore reaches full numerical rank in the constrained tangent state despite instantaneous rank deficiency.

For the layer-scaling analysis, let $0 \leq \lambda_1(W) \leq \cdots \leq \lambda_{90}(W)$ denote the tangent-Gramian eigenvalues for each evaluated path-anchor pair. We define the hidden-position and direct-velocity layer scales as the geometric means of the lower and upper 45 eigenvalues, respectively,

$$S_{\mathrm{pos}}(W) = \exp\left(\frac{1}{45}\sum_{i=1}^{45}\log\lambda_i(W)\right), \qquad S_{\mathrm{vel}}(W) = \exp\left(\frac{1}{45}\sum_{i=46}^{90}\log\lambda_i(W)\right).$$

The plotted layer statistics are the medians of $S_{\mathrm{pos}}(W)$ and $S_{\mathrm{vel}}(W)$ over the path-anchor samples, and the scaling exponents are obtained by linear fits using $K_w \in \{2, 4, 8, 16\}$ on a 0.001 ps storage grid, corresponding to $h \in \{0.002, 0.004, 0.008, 0.016\}$ ps.

The short-window layer strengths also agree with the kinetic prediction. The directly forced tangent-velocity layer has fitted scaling exponent 1.00538, close to the theoretical h behavior, whereas the dynamically exposed tangent-position layer has exponent 3.12836, close to the theoretical h^3 behavior. The corresponding fits have $R^2 = 0.999998$ and $R^2 = 0.999477$, respectively. Figure 6 shows these fitted layer statistics together with the theoretical h and h^3 reference slopes.

Rank versus conditioning. Finite-window rank completion establishes identifiability but does not imply uniform conditioning. The four-layer chain in Appendix C.7.1 makes this distinction explicit: all four directions are identifiable for every positive horizon, yet the deepest controllability layer enters only at order h^7. The same phenomenon appears in constrained alanine, where the directly

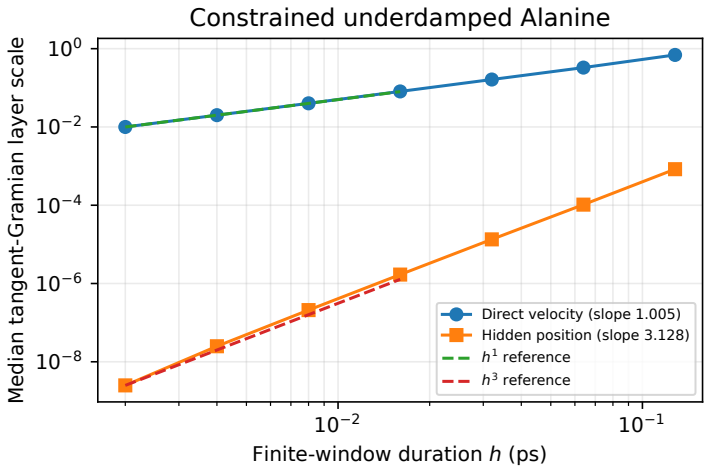


Figure 6: Short-window controllability scaling in constrained-underdamped alanine. The directly forced tangent-velocity layer follows the predicted h scaling, whereas the dynamically exposed tangent-position layer follows approximately h^3. Lines show the fitted layer statistics together with the theoretical h and h^3 reference slopes.

forced velocity layer scales approximately as h, while the position layer exposed through kinetic transport scales approximately as h^3.

Accordingly, the recovery window and ridge regularization play distinct roles. The window determines which dynamically transported directions become identifiable, whereas the ridge parameter ε stabilizes inversion of the resulting anisotropic Gramian without creating directions outside its transported span. These roles correspond to the finite-time geometry analyzed in Appendix B.3 and the executable-consistency bounds in Appendix B.4.