WHY PROJECTED HAMILTONIAN FLOW MATCHING CAN LEARN ZERO

Anonymous authors
Paper under double-blind review

ABSTRACT

Generative models have recently incorporated Hamiltonian structure to align the generation process with physical constraints. However, the resulting flows may preserve energy yet fail to learn the required probability transport. We study projected scalar-Hamiltonian velocities $P_H J\nabla h$ and construct smooth positive density paths on compact phase-symmetric shells whose unique population Flow Matching optimum is zero. Orbit averaging identifies the obstruction and characterizes the exact divergence range on round shells. Under a density-ratio assumption, a finite-excess-risk bound gives a quantitative endpoint error floor for approximate fits. We also characterize the closed regression spaces that preserve every teacher's density response: they contain the closure of tangent gradients. A fixed-budget class–path comparison measures how this population capacity translates into neural learning, and a source-filtering diagnostic identifies an exact conditional sampler in the physical-data task. For an explicit two-oscillator construction, the population-optimal flow retains its source law and incurs endpoint total-variation error $1/8$.

1 INTRODUCTION

A constrained generative flow must preserve each sample's energy and redistribute probability within the energy surface. We study the projected scalar-Hamiltonian velocity $v = P_H J\nabla h$, where H is a prescribed invariant and h is a learned scalar. Projection enforces $\nabla H \cdot v = 0$ and selects a particular class of density responses. On phase-symmetric energy shells, this restriction can make population Flow Matching select the zero velocity for a changing, admissible density path. Generation time here parameterizes probability transport and need not coincide with physical time.

A minimal counterexample. On the unit S^3 shell, let $z = q_1^2 + p_1^2$ be the first oscillator's energy share and μ normalized area. The positive path $\rho_t = 1 + t(z - 1/2)$ redistributes this share at fixed total energy. A smooth tangent gradient generates it. Flow Matching over $P_H J\nabla h$ exhibits *regression collapse*: zero is the unique population-optimal velocity for this changing density. The resulting flow retains μ and incurs endpoint TV error $1/8$, or action W_1 error $1/12$. Near-minimizers retain a quantitative endpoint action-W_1 error floor when their own densities have a controlled ratio to ρ_t. Figure 1 summarizes the construction.

Orbit averaging. On compact shells $H = F(I_1, \ldots, I_n)$ with independent oscillator phase symmetries, projected scalar-Hamiltonian fields have zero weighted response to action-dependent tests. On round shells their divergence range is exactly the kernel of an orbit-averaging operator. Projection generally produces a nonzero shell divergence; orbit averaging identifies its range. Orthogonality then gives regression collapse. Generating this response requires a velocity outside the backbone. The same zero optimum persists after adding a convex continuity-residual penalty to the regression objective.

For a legal teacher u, its restricted fit $\Pi_H u$ is response-correct precisely when $D_\rho(u - \Pi_H u) = 0$, where $D_\rho v = -\operatorname{div}_M(\rho v)$. We characterize closed linear regression spaces that preserve this response for every teacher: they must contain all tangent gradients, after closure. Thus both a full gradient class and a gradient-completed backbone suffice. Section 4 compares their learned density transport at a fixed training budget.

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

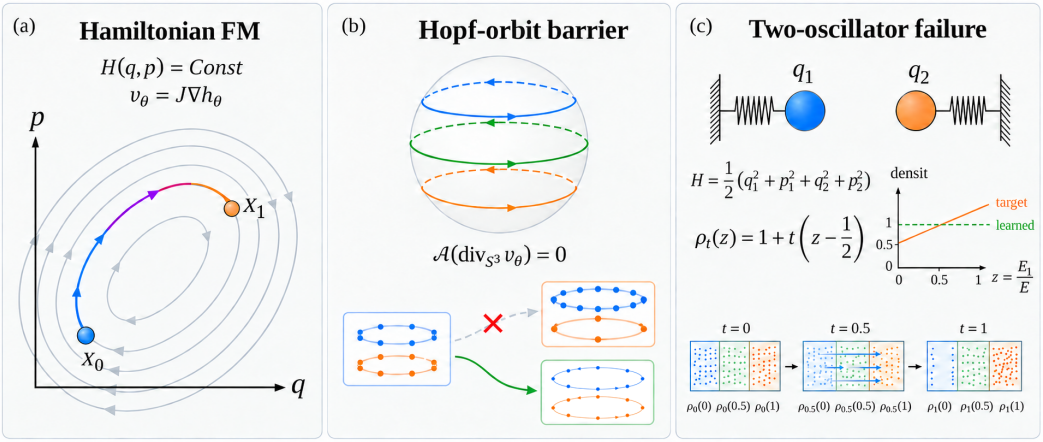


Figure 1: **Exact conservation can coexist with a zero regression optimum.** (a) Hamiltonian phase-space motion motivates the scalar parameterization. The analyzed class is $v_\theta = P_H J\nabla h_\theta$: projection onto the prescribed energy shell is implicit in this overview and enforces conservation of H. (b) On round shells, $\mathcal{A}(\mathrm{div}_{S^3} v_\theta) = 0$ restricts instantaneous density responses; the sphere and orbit drawings are schematic. (c) On the unit S^3, the prescribed energy-share path $\rho_t(z) = 1 + t(z - 1/2)$ changes at fixed total energy, while the restricted population-optimal flow retains its source law. The target endpoint density is $z + 1/2$ and the learned density is one. Generation time parameterizes probability transport. Learned comparisons are in section 4.

Relation to existing transport models. Hamiltonian Neural Networks encode physical structure by learning a scalar Hamiltonian and deriving the dynamics through Hamilton's equations (Greydanus et al., 2019). Constrained variants incorporate explicit mechanical constraints in Cartesian coordinates (Finzi et al., 2020). Hamiltonian Generative Networks extend this perspective to latent dynamics learned from images and introduce a Hamiltonian normalizing flow for density modeling (Toth et al., 2020). Hamiltonian generative flows use augmented constructions (Holderrieth et al., 2024), while FlowPET combines separable Hamiltonian dynamics with measurement-consistency constraints for PET reconstruction (Zhang et al., 2026). Related FM methods incorporate physical and geometric constraints through complementary mechanisms. PCFM applies physics-based corrections to pretrained flows during inference to enforce hard constraints on the final generated solutions (Utkarsh et al., 2025). Gauge FM maps compact convex feasible sets to the unit ball, where reflection or projection enforces feasibility before mapping samples back (Li et al., 2026). Riemannian FM constructs target vector fields adapted to the underlying manifold geometry (Chen and Lipman, 2024). Within this broader effort, we examine the density responses available to the projected scalar class $P_H J\nabla h$. Restrictions of volume-preserving and smooth bijective flows have different mechanisms (Draxler et al., 2024; Cornish et al., 2020). Manifold Flow Matching, Moser transport, and gradient action principles already supply broader density transport classes (Moser, 1965; Rozen et al., 2021; Ben-Hamu et al., 2022; Chen and Lipman, 2024; Khesin and Lee, 2009; Neklyudov et al., 2023). We apply their classical weighted-Poisson construction to the response omitted by restricted regression.

Consequences for learning. The objective controls test the predicted zero optimum after adding exact continuity information. A predeclared two-path, four-class comparison measures learning at a fixed budget. On the analytic sphere path, mean action W_1 falls from 0.0840 for the backbone to 0.0079 for pure gradients; on the geodesic bridge, the generic tangent model has the lowest endpoint error. Both paths remain phase-symmetric; their label distributions, time courses, and learned errors differ. The FPUT diagnostic evaluates source filtering alongside learned transport. The theory characterizes prescribed density paths and population regression spaces; the experiments measure their finite-network realizations.

2 A DENSITY-PATH SEPARATION ON PHASE-SYMMETRIC ENERGY SHELLS

2.1 LEGAL DENSITY PATHS BEYOND THE BACKBONE'S CAPACITY

On phase-symmetric energy shells, conservation-compatible density paths can lie outside the projected scalar-Hamiltonian class. For a constructed tangent teacher, this restriction makes zero the unique population FM-optimal velocity.

Let $x = (q, p) \in \mathbb{R}^{2n}$, $n \geq 2$, with the standard Euclidean metric and $J = \left(\begin{smallmatrix} 0 & I \\ -I & 0 \end{smallmatrix}\right)$. Put $I_i = (q_i^2 + p_i^2)/2$ and $H(x) = F(I_1, \ldots, I_n)$. Assume F is smooth near the compact connected regular level $M = H^{-1}(E)$. Write μ for normalized surface area on M and $P_H = I - \nabla H \nabla H^\top / \|\nabla H\|^2$. Potentials h are smooth in a neighbourhood of M, with unrestricted normal derivatives.

Theorem 1 (Phase-symmetric prescribed-path separation)**.** *Let $a = a(I)|_M$ be smooth, nonzero, and $\int_M a \, d\mu = 0$. Choose $\epsilon \neq 0$ with $|\epsilon| \|a\|_\infty < 1$, and put $\rho_t = 1 + t\epsilon a$, $P_t = \rho_t \mu$ for $t \in [0, 1]$. There is a smooth phase-invariant potential g_t, unique up to constants, such that $u_t = \nabla_M g_t$ generates P_t. Every field $v_t = P_H J \nabla h_t$ has a density response distinct from $\partial_t \rho_t$. For every such field, at each t,*

$$\int_M \rho_t \|v_t - u_t\|^2 \, d\mu = \int_M \rho_t \|u_t\|^2 \, d\mu + \int_M \rho_t \|v_t\|^2 \, d\mu. \tag{1}$$

Thus zero is the unique population-optimal velocity, its teacher-energy-normalized loss is one, and its exact flow stays at μ. Its endpoint error is $\mathrm{TV}(\mu, P_1) = |\epsilon| \int_M |a| \, d\mu / 2 > 0$.

Independent rotations of the (q_i, p_i) planes preserve M, μ, P_H, and J. Averaging an ambient h makes its gradient action-dependent and $J\nabla\bar{h}$ angular. Thus any smooth action-dependent w, ϕ satisfy

$$\int_M w \, \nabla_M \phi \cdot P_H J \nabla h \, d\mu = 0. \tag{2}$$

The phase-invariant zero-mean solution of $-\operatorname{div}_M(\rho_t \nabla_M g_t) = \epsilon a$ supplies the teacher. In (2), $(w, \phi) = (\rho_t, a)$ contradicts the required rate $d\mathbb{E}_{P_t} a / dt = \epsilon \int_M a^2 d\mu$, while (ρ_t, g_t) proves (1). The global proof, a positive regression floor for other legal teachers, and a quartic-shell example appear in subsection A.1.

2.2 AN EXACT DIVERGENCE RANGE ON THE SPHERE

For $H(x) = \|x\|^2/2$, let $S_r^{2n-1} = \{\|x\| = r\}$ and μ_r be its normalized surface measure. Here $P_H = I - xx^\top / r^2$. We compare $\mathcal{V}_{\mathrm{ham}} = \{P_H J \nabla h\}$ and the full tangent class $\mathcal{V}_{\tan} = \{P_H w\}$, with smooth ambient h, w. Both classes conserve H, but the sphere admits a sharper range statement.

Define $Tf = (Jx) \cdot \nabla f$, $Df = x \cdot \nabla f$, and the orbit average

$$\mathcal{A} f(x) = \frac{1}{2\pi} \int_0^{2\pi} f(e^{sJ} x) \, ds.$$

The map e^{sJ} rotates every oscillator phase by the same angle. These periodic orbits are the Hopf circles; functions of the oscillator energy shares are constant along them.

Proposition 2 (Divergence range)**.** *For every smooth ambient scalar h,*

$$\operatorname{div}_{S_r}(P_H J \nabla h) = \frac{1}{r^2} T\big((2n-2)h + Dh\big)\big|_{S_r}, \qquad \{\operatorname{div}_{S_r} v : v \in \mathcal{V}_{\mathrm{ham}}\} = \ker \mathcal{A}. \tag{3}$$

Thus every attainable divergence averages to zero along each Hopf circle, and every smooth function with that property is attainable. In particular, every nonzero Hopf-invariant function lies outside the divergence range. This is an exact range statement on the round shell; Theorem 1 gives an obstruction on the more general shells. The proof is in Appendix A. The orbit-derivative mechanism has a classical contact-geometric counterpart (Bravetti and Tapias, 2015); the displayed formula concerns the unrestricted ambient projected backbone itself.

2.3 Zero optimal velocity with endpoint TV error $1/8$

The range restriction yields an explicit two-oscillator task below. We first state the sphere result under simultaneous phase rotation. The general-shell result in Theorem 1 uses independent phase rotations.

Theorem 3 (Zero optimum for a changing density)**.** *Let a be a nonzero smooth function on S^{2n-1}_r with $Ta = 0$ and $\int a\, d\mu_r = 0$. Choose $\epsilon \neq 0$ with $1 + \epsilon a > 0$, and let b be the mean-zero solution of $-\Delta_{S_r} b = a$. Set*

$$p_t = 1 + t\epsilon a, \qquad P_t = p_t \mu_r, \qquad u_t = \frac{\epsilon \nabla_{S_r} b}{p_t}, \qquad t \in [0, 1]. \tag{4}$$

Then u_t smoothly generates P_t. For every $v_h = P_H J \nabla h$,

$$\mathcal{L}_t(v_h) := \int p_t \|v_h - u_t\|^2 \, d\mu_r = \mathcal{L}_t(0) + \int p_t \|v_h\|^2 \, d\mu_r. \tag{5}$$

Thus zero is the unique minimizing velocity, with teacher-energy-normalized loss one, and the full tangent class attains zero loss. The prescribed path lies outside the backbone class's realizable density paths.

The proof and the resulting family of Hopf-invariant targets are given in Appendix A.

Two oscillators: an explicit density-learning task. On S^3, take $r = 1$,

$$z = q_1^2 + p_1^2, \qquad a = z - \tfrac{1}{2}, \qquad \lambda = 8, \qquad \epsilon = 1.$$

Here z is the first oscillator's energy share. Under μ_1, z is uniform on $[0, 1]$. The eigenfunction identity $-\Delta_{S^3} a = 8a$ gives $b = a/8$. The path and its target are

$$p_t(z) = 1 + t(z - \tfrac{1}{2}), \qquad Q = (z + \tfrac{1}{2})\mu_1.$$

Probability shifts toward larger energy shares while total energy remains fixed. The restricted FM optimum is $v^\star = 0$, whose flow retains μ_1 and incurs

$$\mathrm{TV}(\mu_1, Q) = \frac{1}{2} \int_0^1 |z - \tfrac{1}{2}| \, dz = \frac{1}{8}.$$

The exact sampling map and teacher derivation are in Appendix A. A tangent-gradient field realizes the same teacher; (13) gives its explicit potential.

2.4 Endpoint error under finite excess risk

Exact regression collapse also constrains approximate fits. Let $\mathcal{L}_t(v) = \int \rho_t \|v - u_t\|^2 d\mu_1$ on the preceding S^3 task, where $u_t = \nabla_{S^3} z/(8\rho_t)$.

Proposition 4 (Finite-error consequence under density control)**.** *Let $v_t = P_H J \nabla h_t$ be smooth in space and time on $[0, 1] \times S^3$ and let its flow from μ_1 have density $\widehat{\rho}_t$. Let $C < \infty$ satisfy $\widehat{\rho}_t \leq C\rho_t$ everywhere. Put $\Delta = \int_0^1 [\mathcal{L}_t(v_t) - \mathcal{L}_t(0)] dt$. Then*

$$W_1(z_\# \widehat{P}_1, z_\# P_1) \geq \left(\frac{1}{12} - \sqrt{C\Delta} \right)_+ . \tag{6}$$

The regression identity gives $\Delta = \int\int \rho_t \|v_t\|^2$. Since $\|\nabla_{S^3} z\| \leq 1$, density control bounds the model's change in mean action by $\sqrt{C\Delta}$, whereas the target change is $1/12$; see Appendix B. A sufficient choice is $C = 2\exp(\int_0^1 \|\operatorname{div}_{S^3} v_t\|_\infty dt)$. Applying this bound uniformly along fitted networks requires a uniform bound on the space–time density ratio. The excess risk is a population quantity: for population NMSE, $\Delta = C_u(\mathrm{NMSE} - 1)$, where $C_u = 0.0106027057\ldots$ is the integrated teacher energy. The experiments report empirical losses; using the proposition for a trained network additionally requires certified population risk and density control.

These results concern prescribed density paths and their population FM optima. Time-dependent compositions can reach beyond the instantaneous divergence restriction by following different paths (subsection A.5). The separation requires $n \geq 2$; on S^1 the unrestricted scalar backbone represents every tangent field.

3 WHICH REGRESSION CLASSES PRESERVE DENSITY RESPONSE?

A tangent-gradient branch supplies the missing response. The mixed field is

$$v_{\theta,\phi} = b_\theta + c_\phi, \qquad b_\theta = P_H J \nabla h_\theta, \qquad c_\phi = P_H \nabla g_\phi. \tag{7}$$

Both branches are tangent to the same energy shell, so $\nabla H^\top v_{\theta,\phi} = 0$ in continuous time. The sum defines a conservative probability-transport field. The invariant-law and multiple-constraint background is in subsection D.2.

3.1 A CRITERION FOR RESPONSE-PRESERVING REGRESSION

Fix a compact connected boundaryless regular energy shell M, with induced surface measure σ. Let $\rho_t > 0$ be a smooth normalized density path. At a fixed time, write $\|v\|_\rho^2 = \int_M \rho \|v\|^2 d\sigma$ and define

$$D_\rho v = -\operatorname{div}_M(\rho v), \qquad \langle D_\rho v, \psi \rangle = \int_M \rho v \cdot \nabla_M \psi \, d\sigma.$$

For a tangent backbone b, the missing response is $R_b = \partial_t \rho - D_\rho b$. In particular, let u be a legal teacher and Π_H the weighted orthogonal projection onto $\mathcal{V}_H = \overline{\{P_H J \nabla h\}}$ in $L^2(\rho \, d\sigma; TM)$. The restricted population FM fit satisfies

$$b = \Pi_H u, \qquad R_b = D_\rho(u - \Pi_H u). \tag{8}$$

The required correction is determined by the residual's density response. P_H projects pointwise, while Π_H projects in the density-weighted function space. The closed-space fit belongs to L^2; its response and correction belong to H^{-1} and H^1, respectively. The following criterion uses weak responses. The subsequent minimum-correction result gives a classical generating flow under smoothness assumptions.

Let $\mathcal{H}_\rho = L^2(\rho \, d\sigma; TM)$ and $\mathcal{G}_\rho = \overline{\{\nabla_M \psi : \psi \in C^\infty(M)\}}^{\mathcal{H}_\rho}$.

Proposition 5 (Response-preserving regression)**.** *Let $\mathcal{V} \subset \mathcal{H}_\rho$ be a closed linear subspace and $\Pi_\mathcal{V}$ its orthogonal projection. The following conditions are equivalent:*

$$D_\rho \Pi_\mathcal{V} u = D_\rho u \quad \text{for every } u \in \mathcal{H}_\rho;$$
$$\mathcal{V}^\perp \subseteq \ker D_\rho; \qquad \mathcal{G}_\rho \subseteq \mathcal{V}.$$

Weak integration by parts gives $\ker D_\rho = \mathcal{G}_\rho^\perp$; projection residuals fill $\mathcal{V}^\perp$, proving the equivalence. In particular, both $\mathcal{G}_\rho$ and $\overline{\mathcal{V}_H + \mathcal{G}_\rho}$ preserve every teacher's density response. The equivalence characterizes population least-squares projection onto closed linear spaces.

A sample-based lower witness. For $R_b = D_\rho(u - b)$ and any smooth nonconstant test ψ,

$$\|R_b\|_{\dot{H}^{-1}(\rho)} \geq \frac{|\mathbb{E}_\rho[(u-b) \cdot \nabla_M \psi]|}{\{\mathbb{E}_\rho \|\nabla_M \psi\|^2\}^{1/2}}. \tag{9}$$

Conditional labels V with $\mathbb{E}[V \mid X] = u(X)$ can replace u in the numerator. Each preselected test gives a lower bound on response error, computed from sampled velocity–gradient products. We use $\psi = z$ in the class–path comparison.

3.2 MINIMUM CORRECTION FOR A FIXED BACKBONE

Proposition 6 (Minimum response correction for a fixed backbone)**.** *Under the above assumptions, for any smooth tangent b, the equation*

$$-\operatorname{div}_M(\rho \nabla_M g) = \partial_t \rho - D_\rho b \tag{10}$$

has a unique smooth solution with $\int_M g \, d\sigma = 0$. Its gradient is the unique minimum-norm correction among tangent $L^2(\rho)$ fields satisfying the response constraint weakly:

$$c^\star = \nabla_M g = \operatorname*{arg\,min}_{c:\, D_\rho(b+c) = \partial_t \rho} \frac{1}{2} \|c\|_\rho^2. \tag{11}$$

For smooth time-dependent b_t, the corrected field generates the prescribed density path from its initial law.

Normalization makes the Poisson right side mean zero. Weighted elliptic inversion gives g, and the weighted Hodge decomposition writes every feasible correction as $c = c^\star + k$, where $D_\rho k = 0$. Integration by parts gives

$$\langle c^\star, k\rangle_\rho = 0, \qquad \|c\|_\rho^2 = \|c^\star\|_\rho^2 + \|k\|_\rho^2.$$

The correction has the required response and minimum norm. For a legal teacher,

$$u = b + \underbrace{\nabla_M g}_{\text{minimum response correction}} + k, \qquad D_\rho k = 0. \tag{12}$$

Moreover, $\|\nabla_M g\|_\rho = \|R_b\|_{\dot{H}^{-1}(\rho)}$, linking the missing response to its minimum correction cost; the norm and proof are in subsection D.1. This is the classical weighted-Poisson variational construction (Khesin and Lee, 2009; Neklyudov et al., 2023), applied to a fixed backbone. Its minimum-norm property holds with the backbone fixed.

Explicit conservative completion. For the sphere task the teacher is itself a tangent gradient, with potential

$$g_t^{\mathrm{ex}} = \frac{\log(1 + t(z - 1/2))}{8t} \quad (t > 0), \qquad g_0^{\mathrm{ex}} = (z - 1/2)/8. \tag{13}$$

Both pure gradients and the sum in (7) can therefore fit it. Experiments use unpenalized FM. Appendix C analyzes the population response bias of an optional correction-energy penalty.

4 Class, probability path, and source-access diagnostics

We compare Backbone (B), Gradient (G), Mixed (M), and Tangent (V): $P_H J\nabla h$, $P_H \nabla g$, $P_H(J\nabla h + \nabla g)$, and $P_H w$. Five-seed means and sample SD describe training randomness; conserved energy, teacher fit, and distribution discrepancy are separate quantities. All completion runs use $\gamma = 0$.

4.1 Learning the completed model depends on the path

Both paths connect uniform S^3 area to $(z + 1/2)\mu_1$. The first is the analytic path of Section 2. The second draws endpoints independently and uses the derivative of the shortest geodesic as its conditional FM label. Its marginal is $\rho_t^{\mathrm{geo}} = 1 + \eta(t)(z - 1/2)$, with

$$\eta(t) = \{1 + 2\mathbb{E}\cos[2(1 - t)\Theta]\}/3, \qquad \Theta \sim (2/\pi)\sin^2\theta\, d\theta.$$

Its mean teacher is $\eta'(t)\nabla z/(8\rho_t^{\mathrm{geo}})$; its marginals remain in the same one-dimensional, phase-invariant density family. Both paths retain the exact obstruction (subsection B.3); the geodesic path also has an initial reversal and conditional-label noise. The bridge and budget were fixed before the new trained results were inspected.

All classes use three width-66 layers, matched minibatches within each path, and 1,000 AdamW updates. Parameter counts span 10,297–10,498, within 2%; this matches model size and update count. Wall-clock costs are reported separately. We reuse the ten retained analytic B/M checkpoints and train only the thirty missing cells. All forty are evaluated by the same current code; the imported cells were trained in an earlier runtime. Checkpoints are the raw weights at the prescribed final update. Evaluation uses 4,096 sources, 64-step RK4, and 64 shared SW1 directions. The primary endpoint metric in Table 1 is one-dimensional action W_1 against the exact target CDF. The weak diagnostic measures the normalized RMS error of the action rate $\mathbb{E}_{\rho_t}[v \cdot \nabla z]$ against $\eta'(t)/12$. It measures the response in the action direction. Full protocol, costs, and all seed results are in subsection E.2. For the geodesic runs, the mean CPU time for a 4,096-state rollout is 0.352 s for M and 0.150 s for V (Table 7); these are descriptive timings on one CPU environment.

Learning across the two paths. On the analytic path, G, M, and V recover the action redistribution, while B stays near the zero-flow error. On the geodesic bridge, G improves on B in all five seeds but recovers only part of the transport; M gains little, and V has the lowest endpoint error in all five paired seeds. Both mean teachers belong to the full gradient class. Their finite-network fits

Table 1: **Frozen class–path comparison, five seeds.** Mean ± sample SD. All errors are lower-is-better. Mean-teacher NMSE removes conditional-label variance; geodesic conditional-label NMSE is reported separately in the supplement. Analytic B/M are imported; all other cells are new at the same 1,000-update budget.

Path	Class	Action W_1	Full-state SW1	Mean-teacher NMSE
Analytic	B	0.0840 ± 0.0024	0.0248 ± 0.0018	$1.0000 \pm 4.6\times10^{-5}$
	G	0.0079 ± 0.0018	0.0122 ± 0.0018	0.0174 ± 0.0008
	M	0.0076 ± 0.0019	0.0122 ± 0.0018	0.0164 ± 0.0012
	V	0.0079 ± 0.0018	0.0122 ± 0.0018	0.0174 ± 0.0006
Geodesic	B	0.0840 ± 0.0024	0.0250 ± 0.0020	1.0009 ± 0.0010
	G	0.0652 ± 0.0103	0.0210 ± 0.0015	0.7844 ± 0.0879
	M	0.0817 ± 0.0026	0.0246 ± 0.0021	0.9707 ± 0.0196
	V	0.0107 ± 0.0013	0.0131 ± 0.0015	0.3341 ± 0.0099

differ under the two path and label distributions; label variance, parameterization, initialization, and optimization are possible contributors.

The primary B-versus-G separation criterion holds on both paths, but its paired mean action-error reduction is 0.07617 for the analytic path and 0.01882 for the geodesic bridge. All forty integration gates pass.

Endpoint and path response. V's geodesic weak-response error is 0.526, accompanying its low endpoint error. The bridge's irreducible conditional-label variance is 99.42% of label energy, so conditional-label NMSE near one is compatible with substantial learning of the mean teacher. The comparison measures both the class restriction and the path-dependent accuracy of the finite-network fits.

4.2 Objective Controls and a Nonquadratic Shell

Objective control. An objective-based explanation predicts that supplying exact continuity information should recover the required transport. The class restriction instead predicts that zero remains population-optimal under the augmented objective. The FM/RFM and sphere-adapted FDM runs test this distinction on the same analytic sphere path, teacher, and backbone. RFM's induced metric gives the same squared loss here, so it reuses FM. FDM adds the known intrinsic continuity residual; this is an analytic-marginal adaptation of Huang et al. (2025), using the known marginal in place of their conditional estimator. The scan at weights $0.1, 1, 10$ leaves backbone NMSE near one and action W_1 near its prior value. Proposition 7 explains this absolute-residual floor by phase averaging and convexity. The mixed FDM positive control checks that the augmented objective can learn once the class contains the missing direction; the full scan and rollout figure are in subsection E.1.

Nonquadratic shell. The nonquadratic shell $H = I_1 + 2I_2 + I_1^2 + I_2^2 = 1$ supplies a second geometry. With $c = \mathbb{E}_\mu I_1$, a binary sensor gives posterior paths $\rho_{t,y} = 1 + ty(I_1 - c)$, $y \in \{-1, 1\}$, learned from exact quantile samples. Table 2 reports the 300-update comparison. These are short mechanism experiments: the generic tangent and mixed models attain comparable errors. Protocols and conditional response plots are in subsection E.3.

Table 2: **Nonquadratic conditional generation.** Five seeds, 300 updates; mean ± sample SD. All errors are lower-is-better. W_1 averages the two conditions; event error takes the worse absolute error in $P(I_1 > c \mid y)$. Mean absolute energy errors are below 2.2×10^{-7}.

Model	Teacher NMSE	Action W_1	Event error
Backbone (B)	$1.0001 \pm 5.4\times10^{-5}$	$0.0333 \pm 1.7\times10^{-6}$	0.0876 ± 0.0086
Mixed (M)	0.0399 ± 0.0449	0.0053 ± 0.0018	0.0130 ± 0.0067
Tangent (V)	0.0235 ± 0.0083	0.0045 ± 0.0019	0.0105 ± 0.0080
Prior	–	$0.0333 \pm 2.0\times10^{-6}$	0.0876 ± 0.0086
Exact-flow oracle	–	0.0036 ± 0.0022	0.0094 ± 0.0090

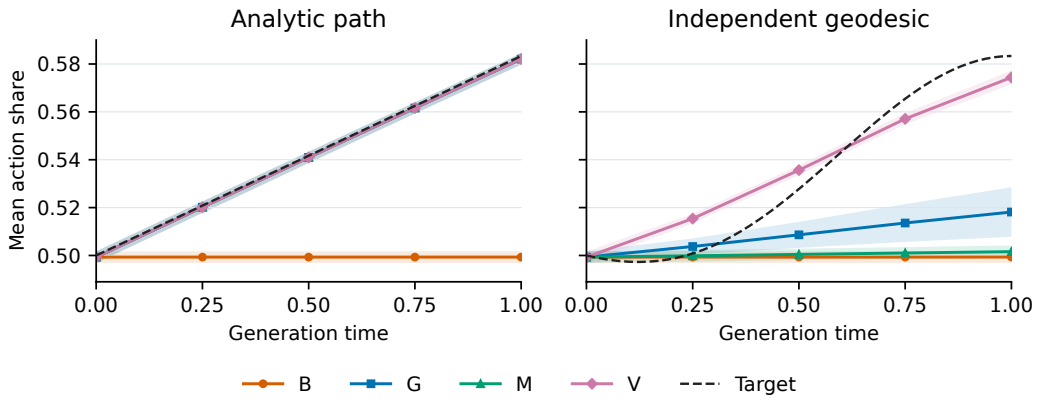


Figure 2: **The probability path changes the learned transport at a fixed training budget.** B denotes Backbone, G Gradient, M Mixed, and V Tangent. B stays near its source mean. G and M track the analytic target and lag behind the geodesic target. V achieves the closest geodesic endpoint mean and has a visible intermediate-path discrepancy. Curves and shading show five-seed means and sample SD after the prescribed training budget; the dashed target is the exact marginal mean.

4.3 FPUT-32: CONDITIONAL SAMPLING FROM THE SOURCE MIXTURE

The FPUT task draws source states from an approximately equal mixture of three condition-specific pools of a published trajectory (Marchetti, 2025). The observable condition is the interval of the harmonic-mode fraction $r(x)$, with thresholds 0.05 and 0.8. For disjoint labels c and source $P_0 = (Q_0 + Q_1 + Q_2)/3$, conditioning the source gives

$$P_0(\cdot \mid \mathrm{bin}(r(X)) = c) = Q_c.$$

Filtering therefore returns the exact stipulated conditional law with acceptance probability $1/3$. This is an endpoint sampler under the same source access granted to the neural models. Filtering accepts the source states in the requested condition; the neural models learn a flow from the full source using bridge velocity labels.

Table 3: **FPUT endpoint diagnostic under the granted source access.** Neural entries are five-seed means after a 15,000-update budget with validation-selected checkpoints. Filtering is one deterministic evaluation of the fixed source cache. All entries use the frozen 64 SW1 directions.

Sampler	Condition hit (%)	Mode W_1	Full-state SW1
Backbone	46.99	0.29694	0.20149
Gradient	94.21	0.03477	0.17323
Mixed	97.94	0.02207	0.16493
Source filtering	100.00	0.001665	0.021443

We partitioned the 1,536 saved source states into 513, 513, and 510 outputs and compared them with 512 targets per condition (Table 3). Target coordinates were used only for evaluation. Filtering each original query's 512-state source batch separately gives 171, 171, and 170 outputs and mean SW1 0.023535, giving the same ordering under separate-query evaluation. The neural comparison measures learned transport at a fixed budget; source filtering provides the more accurate endpoint sampler in this task. The data come from one orbit with time-block train/test splits. Appendix F, subsection E.4 report the neural runs, uncertainties, distribution-shift diagnostics, and protocol.

5 DISCUSSION AND CONCLUSION

On phase-symmetric energy shells, projected scalar-Hamiltonian Flow Matching can select zero as its population optimum for a changing density path. Orbit averaging characterizes the available

responses, and density control extends the endpoint error floor to approximate fits. The response-preservation criterion identifies tangent gradients as sufficient for population density transport.

The learned comparison shows how this capacity interacts with the probability path: gradient models fit the analytic teacher accurately, while the generic tangent model achieves the best geodesic endpoint accuracy at the fixed budget. The FPUT diagnostic shows that source conditioning itself supplies the target sampling law. These results identify density response as the quantity to examine when designing a conservative generative model.

AI Use Statement

Generative AI tools assisted with polishing the paper only. The authors take responsibility for the paper's text, claims, and code.

Reproducibility Statement

Appendices specify the assumptions, proofs, data, teachers, estimands, and training protocols. The accompanying records retain per-seed metrics, checkpoints, source-range checksums for the published FPUT subset, and figure-generation inputs. Reproduction entry points are given with each experiment. The class–path comparison retains a pre-execution protocol hash, a start ledger, raw final checkpoints, and all outcomes. Verification scripts check the saved metrics, source subsets, and artifact hashes.

References

H. Ben-Hamu, S. Cohen, J. Bose, B. Amos, M. Nickel, A. Grover, R. T. Q. Chen, and Y. Lipman. Matching normalizing flows and probability paths on manifolds. In *Proc. 39th Int. Conf. on Machine Learning (ICML)*, PMLR 162, 2022. arXiv:2207.04711.

A. Bravetti and D. Tapias. Liouville's theorem and the canonical measure for nonconservative systems from contact geometry. *J. Phys. A: Math. Theor.*, 48:245001, 2015. arXiv:1412.0026.

R. T. Q. Chen and Y. Lipman. Flow matching on general geometries. In *ICLR*, 2024. arXiv:2302.03660.

R. Cornish, A. Caterini, G. Deligiannidis, and A. Doucet. Relaxing bijectivity constraints with continuously indexed normalising flows. In *Proc. 37th Int. Conf. on Machine Learning (ICML)*, PMLR 119, 2020. arXiv:1909.13833.

F. Draxler, S. Wahl, C. Schnörr, and U. Köthe. On the universality of volume-preserving and coupling-based normalizing flows. In *Proc. 41st Int. Conf. on Machine Learning (ICML)*, PMLR 235, 2024. arXiv:2402.06578.

M. Finzi, K. A. Wang, and A. G. Wilson. Simplifying Hamiltonian and Lagrangian neural networks via explicit constraints. In *NeurIPS*, 2020. arXiv:2010.13581.

S. Greydanus, M. Dzamba, and J. Yosinski. Hamiltonian neural networks. In *NeurIPS*, 2019. arXiv:1906.01563.

P. Holderrieth, Y. Xu, and T. Jaakkola. Hamiltonian score matching and generative flows. In *NeurIPS*, 2024. arXiv:2410.20470.

Y. Huang, T. Transue, S.-H. Wang, W. M. Feldman, H. Zhang, and B. Wang. Improving flow matching by aligning flow divergence. In *Proc. 42nd Int. Conf. on Machine Learning (ICML)*, PMLR 267, pages 25813–25834, 2025. https://proceedings.mlr.press/v267/huang25ag.html.

B. Khesin and P. Lee. A nonholonomic Moser theorem and optimal transport. *J. Symplectic Geom.*, 7(4):381–414, 2009. arXiv:0802.1551.

Z. Li, K. Alim, and N. Azizan. HardFlow: Hard-constrained sampling for flow-matching models via trajectory optimisation. arXiv:2511.08425, 2025.

X. Li, E. Liang, and M. Chen. Gauge flow matching: Efficient constrained generative modeling over general convex set and beyond. In *ICLR*, 2026.

Y. Lipman, R. T. Q. Chen, H. Ben-Hamu, M. Nickel, and M. Le. Flow matching for generative modeling. In *ICLR*, 2023.

G. Marchetti. Dataset of Entire Trajectories of Fermi–Pasta–Ulam–Tsingou Model β ($N = 32$, $k = 1$, $A = 10$, $\beta \in [1.6, 3]$). Zenodo, 2025. https://doi.org/10.5281/zenodo.15873646.

J. Moser. On the volume elements on a manifold. *Trans. Amer. Math. Soc.*, 120(2):286–294, 1965.

K. Neklyudov, R. Brekelmans, D. Severo, and A. Makhzani. Action matching: Learning stochastic dynamics from samples. In *Proc. 40th Int. Conf. on Machine Learning (ICML)*, PMLR 202, pages 25858–25889, 2023.

N. Rozen, A. Grover, M. Nickel, and Y. Lipman. Moser flow: Divergence-based generative modeling on manifolds. In *NeurIPS*, 2021. arXiv:2108.08052.

P. Toth, D. J. Rezende, A. Jaegle, S. Racanière, A. Botev, and I. Higgins. Hamiltonian generative networks. arXiv:1909.13789, version 2, 2020.

U. Utkarsh, P. Cai, A. Edelman, R. Gomez-Bombarelli, and C. V. Rackauckas. Physics-constrained flow matching: Sampling generative models with hard constraints. In *NeurIPS*, 2025. arXiv:2506.04171.

Z. Zhang, H. Tang, Y. Hu, Z. Hu, and J. Qin. FlowPET: Physics-informed symplectic flow matching for low-count PET reconstruction. In *ICML*, 2026. arXiv:2607.11104.

APPENDIX GUIDE

Appendix A proves the separation and its sphere specialization; Appendix B gives the finite-error proof, convex residual floor, and exact geodesic marginals. Appendix D, Appendix C give the classical minimum correction and optional regularization analysis. Appendix E, Appendix G specify the experimental protocols and additional metrics. Appendix F reports source filtering and neural transport on FPUT data. Appendix H gives all five paired physical controls. Appendix I relates the projected class to other transport constructions.

H is the prescribed energy, h the learned Hamiltonian potential, and g the correction potential. P_H acts pointwise; Π_H projects in the density-weighted velocity space. P_t denotes a probability law, while p_t, ρ_t are densities relative to their stated measure. μ, μ_r are normalized shell area; σ in Section 3 is induced area. The scalar b in the sphere Poisson construction is distinct from the backbone velocity b in Section 3. The density operator D_ρ is distinct from the radial derivative $D = x \cdot \nabla$. Generation time parameterizes the probability path.

A PHASE SYMMETRY, THE PROJECTED BACKBONE, AND THE SPHERE SPECIALIZATION

A.1 PROOF OF THE PHASE-SYMMETRIC SEPARATION

Let R_α, $\alpha \in \mathbb{T}^n$, act by independent orthogonal rotations of the (q_i, p_i) planes. Since $H = F(I)$, the action preserves M, its induced metric, and normalized surface area μ. The standard J commutes with R_α and $P_H(R_\alpha x) = R_\alpha P_H(x) R_\alpha^\top$. For a smooth ambient potential h, put $\bar{h}(x) = \int_{\mathbb{T}^n} h(R_\alpha x)\, d\alpha$. Equivariance of the gradient, J, and P_H implies that the group average of $P_H J \nabla h$ is $P_H J \nabla \bar{h}$. The smooth invariant $\bar{h}$ is a function of the actions locally on the principal stratum $I_i > 0$. The gradients of $\bar{h}$, H, and any invariant ϕ are radial in each oscillator plane, whereas $J\nabla\bar{h}$ is angular. Hence $P_H J \nabla \bar{h} = J \nabla \bar{h}$ and $\nabla_M \phi \cdot P_H J \nabla \bar{h} = 0$ there. Continuity extends this equality across $I_i = 0$. Averaging the integral of the invariant weight $w \nabla_M \phi$ against $P_H J \nabla h$ proves (2).

The zero-mean right side ϵa and uniformly positive ρ_t make $L_t = -\operatorname{div}_M(\rho_t \nabla_M)$ coercive on mean-zero $H^1(M)$. Poincaré's inequality, elliptic regularity, and smooth parameter dependence give the unique mean-zero smooth solution g_t of $L_t g_t = \epsilon a$. Each R_α preserves this equation and the mean-zero gauge, so uniqueness makes g_t invariant. The smooth field $u_t = \nabla_M g_t$ obeys $\partial_t \rho_t + \operatorname{div}_M(\rho_t u_t) = 0$. Its complete flow on compact M pushes μ to $\rho_t \mu$. Equation (2) with $(w, \phi) = (\rho_t, g_t)$ gives $\int_M \rho_t u_t \cdot v_h \, d\mu = 0$ for every backbone field. Expansion proves (1). Since a is nonzero, u_t is nonzero; strict positivity of ρ_t makes zero the unique minimizing velocity, although distinct potentials can represent it. The integrated optimum is the zero flow.

If a backbone field generated the prescribed path, testing its continuity equation with a would give

$$\epsilon \int_M a^2 \, d\mu = \frac{d}{dt} \int_M a \rho_t \, d\mu = \int_M \rho_t \nabla_M a \cdot P_H J \nabla h_t \, d\mu = 0,$$

a contradiction. The argument uses an averaged action response, allowing individual backbone trajectories to change their actions. For any other legal teacher $\widetilde{u}_t$, the same test and Cauchy–Schwarz yield, for every v_h,

$$\|\widetilde{u}_t - v_h\|_{\rho_t}^2 \geq \frac{\epsilon^2 (\int_M a^2 \, d\mu)^2}{\int_M \rho_t \|\nabla_M a\|^2 \, d\mu} > 0.$$

The denominator is positive because a nonconstant smooth a on connected M has nonzero gradient. This gives a positive regression floor for the alternative teacher; its optimal velocity and endpoint depend on that teacher.

A nonquadratic shell. Take $n = 2$, $H = I_1 + 2I_2 + I_1^2 + I_2^2$, and $E = 1$. Since H increases along each nonzero radial ray in $\mathbb{R}^4$, the level is smooth, compact, and connected. On it, $0 \leq z = I_1 \leq (\sqrt{5} - 1)/2$ and $I_2(z) = \sqrt{2 - z - z^2} - 1$. Put $a = z - c$, where $c = \int_M z \, d\mu$, and choose

$\epsilon = 1$. The phase-integrated surface-area weight in z is proportional to

$$w(z) = \frac{\sqrt{2z(1+2z)^2 + 2I_2(z)(2+2I_2(z))^2}}{2+2I_2(z)}.$$

This follows by restricting the Euclidean metric $\sum_i \{dI_i^2/(2I_i) + 2I_i\, d\theta_i^2\}$ to the level curve $I_2 = I_2(z)$; endpoint coordinate singularities are integrable. One-dimensional quadrature gives $c \simeq 0.323041$ and $\mathrm{TV}(\mu, (1+a)\mu) \simeq 0.079527$. Both values are obtained by deterministic quadrature.

In the following sphere-specific subsections, densities are relative to normalized surface measure μ_r on S_r^{2n-1}, unless specified otherwise. The canonical J satisfies $J^\top = -J$ and $J^2 = -I$. Scalar potentials are smooth on an annulus containing the sphere. In particular, their normal derivatives are part of the model class.

A.2 Proof of the divergence characterization

Write $R = \|x\|$, $D = x \cdot \nabla$, and $T = (Jx) \cdot \nabla$ on the annulus. Since $x \cdot J\nabla h = -Th$,

$$v_h = J\nabla h + \frac{Th}{R^2} x.$$

The symmetry of the Hessian gives $\mathrm{div}(J\nabla h) = 0$. The radial and rotation generators commute, $[D, T] = 0$, and $D(R^{-2}) = -2R^{-2}$. Therefore

$$\mathrm{div}\, v_h = \frac{D(Th) + (2n-2)Th}{R^2} = \frac{T(Dh + (2n-2)h)}{R^2}.$$

This extension is tangent to every concentric sphere. If $\nu = x/R$, differentiating $v_h \cdot \nu = 0$ in the radial direction gives $\nu^\top Dv_h\, \nu = 0$. The surface divergence $\mathrm{tr}((I - \nu\nu^\top)Dv_h)$ thus equals the ambient divergence on S_r, proving the first identity in Proposition 2.

The action e^{sJ} is orthogonal, has period 2π, and preserves μ_r. Hence $\mathcal{A}(Tf) = 0$ and $\int f\, Tg\, d\mu_r = -\int (Tf) g\, d\mu_r$. For the converse, take $f \in C^\infty(S_r)$ with $\mathcal{A}f = 0$ and define

$$b(x) = \frac{1}{2\pi} \int_0^{2\pi} s\, f(e^{sJ}x)\, ds.$$

Differentiating along T and integrating by parts gives $Tb = f - \mathcal{A}f = f$. Extend b radially to the annulus and set $h = (R^2 - r^2)b/2$. On S_r, $\nabla h = bx$, so $v_h = bJx$ and $\mathrm{div}_{S_r} v_h = Tb = f$. This proves the exact range statement. In particular, for every smooth A with $TA = 0$,

$$\int \nabla_{S_r} A \cdot v_h\, d\mu_r = -\int A\ \mathrm{div}_{S_r} v_h\, d\mu_r = 0. \tag{14}$$

A.3 Proof of the population separation

Under Theorem 3, compactness and $1 + \epsilon a > 0$ make $p_t = (1-t) + t(1+\epsilon a)$ uniformly positive. The function a has zero mean, so $p_t \mu_r$ is a probability measure. The mean-zero Poisson solution $b = (-\Delta_{S_r})^{-1} a$ is smooth. Every Hopf rotation is an isometry preserving a and μ_r; composing b with such a rotation solves the same normalized Poisson equation. Uniqueness therefore gives $Tb = 0$. Direct substitution gives

$$\partial_t p_t + \mathrm{div}_{S_r}(p_t u_t) = \epsilon a + \epsilon \Delta_{S_r} b = 0.$$

The smooth field u on a compact manifold has a complete evolution over $[0, 1]$. The change-of-variables identity, or uniqueness of this smooth continuity equation, therefore gives $(\varphi_{t,s})_{\#} P_s = P_t$.

Using (14) with $A = b$ gives

$$\int p_t\, u_t \cdot v_h\, d\mu_r = \epsilon \int \nabla_{S_r} b \cdot v_h\, d\mu_r = 0.$$

Expanding the square proves (5). Strict positivity of p_t makes zero the unique minimizing velocity; distinct scalar potentials may represent that same field. Because a is nonzero and mean-zero, b

is nonconstant and $\int p_t \|u_t\|^2 \, d\mu_r > 0$, so the normalized loss is defined. For uniform training time, the excess integrated risk is $\int_0^1 \int p_t \|v_t\|^2 \, d\mu_r \, dt$. A continuous minimizing velocity vanishes everywhere on the sphere for every time. Its exact flow is the identity on the base support.

Every backbone generator violates the continuity equation required by the prescribed path. Indeed, along that path,

$$\frac{d}{dt} \int a \, p_t \, d\mu_r = \epsilon \int a^2 \, d\mu_r \neq 0.$$

A generating field would give this derivative as $\int p_t v_h \cdot \nabla a \, d\mu_r$. Here $p_t \nabla a = \nabla(a + t\epsilon a^2/2)$, and the potential in parentheses is T-invariant. Equation (14) makes the integral zero. This proves the final claim of Theorem 3 independently of the selected teacher.

For any smooth positive normalized $q \neq 1$ with $Tq = 0$, take $a = q - 1$ and $\epsilon = 1$. The path satisfies $p_t \geq \min(1, \min q) > 0$. Thus the theorem covers every such target, with endpoint TV error $\frac{1}{2} \int |q - 1| \, d\mu_r$. The eigenfunction choice below makes its Poisson solution explicit.

A.4 Closed-form task and endpoint discrepancy

On S_r^3, write $z = (q_1^2 + p_1^2)/r^2$ and $a = z - 1/2$. The homogeneous polynomial $q_1^2 + p_1^2 - \|x\|^2/2$ is harmonic of degree two in $\mathbb{R}^4$. Its restriction divided by r^2 is therefore a spherical eigenfunction with $\lambda = 8/r^2$, so $b = a/\lambda$. It is invariant under simultaneous phase rotation. Explicitly,

$$\nabla_{S_r} a = \frac{2}{r^2}\big((1-z)q_1, -zq_2, (1-z)p_1, -zp_2\big), \qquad \|\nabla_{S_r} a\|^2 = \frac{4}{r^2} z(1-z).$$

Normalized independent two-dimensional Gaussian blocks give uniform surface measure. Their squared radii are independent equal-scale gamma variables of shape one; their ratio z is therefore uniform on $[0, 1]$, independently of the two uniform phases. For $0 < |\epsilon| < 2$ the target is positive and

$$\mathrm{TV}(\mu_r, p_1 \mu_r) = \frac{|\epsilon|}{2} \int_0^1 |z - \tfrac{1}{2}| \, dz = \frac{|\epsilon|}{8}, \qquad \mathcal{L}_0(0) = \frac{\epsilon^2 r^2}{96}.$$

The target z-CDF at time t is $G_t(z) = z + (t\epsilon/2)(z^2 - z)$, with derivative $p_t(z) > 0$. Keeping the phases fixed and setting $z_t = G_t^{-1}(z_0)$ gives

$$\dot{z}_t = \frac{\epsilon}{2} \frac{z_t(1 - z_t)}{p_t(z_t)} = \nabla_{S_r} z \cdot u_t.$$

The Cartesian formula for u_t is smooth through $z = 0, 1$, so the coordinate description extends to a smooth flow on the entire sphere. For any $n \geq 2$, the same construction uses $a = (q_1^2 + p_1^2)/r^2 - 1/n$, with eigenvalue $4n/r^2$. The exact $1/8$ value in Table 4 refers specifically to $n = 2$ and $\epsilon = 1$.

Finally, $\log(1 + t\epsilon a)/(\lambda t) = \epsilon a/\lambda - t\epsilon^2 a^2/(2\lambda) + O(t^2)$. This supplies the smooth extension of (19) through $t = 0$, and differentiation gives $\nabla_{S_r} g_t = u_t$.

Table 4: **Exact analytic comparison on the same positive path.** Entries are population values for Theorem 3 with $a = z - 1/2$ and $\epsilon = 1$, using arbitrary smooth function classes, the same noiseless teacher, and exact flows. The gradient-energy penalty is absent ($\gamma = 0$). The endpoint columns use each class's population FM-optimal velocity. $\mathrm{TV}(\mu_1, p_1\mu_1) = \frac{1}{2}\int_0^1 |z - \frac{1}{2}| \, dz = 1/8$. These are derived values; Figure 3, Table 5 report the matched finite-network study.

Velocity class	energy error	optimal normalized FM loss	endpoint TV error
$P_H J \nabla h$	0	1	$1/8$
$P_H \nabla g$	0	0	0
$P_H w$	0	0	0
$P_H J \nabla h + P_H \nabla g$	0	0	0

A.5 The boundary between a velocity class and its compositions

The divergence restriction is instantaneous. To see why it gives a path/objective result and leaves endpoint controllability separate, consider the unit S^3. For $T = Jx$, take

$$V = \nabla_S q_1 = P_H J \nabla p_1, \qquad W = p_1 T = P_H J \nabla\big[(\|x\|^2 - 1)p_1/2\big].$$

Both are backbone fields. Their divergences are $-3q_1$ and $-q_1$, respectively. The identity $\operatorname{div}[V,W] = V(\operatorname{div} W) - W(\operatorname{div} V)$ gives

$$\operatorname{div}_S[V,W] = -1 + q_1^2 + 3p_1^2, \qquad \mathcal{A}(\operatorname{div}_S[V,W]) = -1 + 2z = 2a.$$

The bracket escapes the instantaneous range. This calculation characterizes the first additional direction available to compositions. Accordingly, Theorem 3 is a prescribed-path theorem, distinct from endpoint-orbit controllability.

On S^1, each Hopf orbit is the entire circle and every tangent field is bT; the unrestricted normal derivative construction above realizes every such field. The density-learning separation therefore requires $n \geq 2$. The round-shell geometry is part of the hypothesis, and the theorem supplies an exact test case for the architecture. The physical studies report protocol-specific behavior on other geometries. Smooth radial cutoffs supported on an annulus extend the displayed fields to complete ambient fields while preserving their values on the sphere.

B Finite error, residual penalties, and the geodesic bridge

B.1 Proof of the finite-error consequence

The orthogonality identity in Theorem 3 gives

$$\Delta = \int_0^1 \int_{S^3} \rho_t \|v_t\|^2 \, d\mu_1 \, dt.$$

The model's own continuity equation and $\|\nabla_{S^3} z\|^2 = 4z(1-z) \leq 1$ imply

$$\begin{aligned} \left| \mathbb{E}_{\widehat{P}_1} z - \mathbb{E}_{\mu_1} z \right| &= \left| \int_0^1 \int \widehat{\rho}_t v_t \cdot \nabla_{S^3} z \, d\mu_1 \, dt \right| \\ &\leq \left(\int_0^1 \int \widehat{\rho}_t \|v_t\|^2 \right)^{1/2} \left(\int_0^1 \int \widehat{\rho}_t \|\nabla_{S^3} z\|^2 \right)^{1/2} \leq \sqrt{C\Delta}. \end{aligned}$$

The desired mean change is $\int_0^1 z(z-1/2)\, dz = 1/12$. The Kantorovich lower bound by the difference of means and the triangle inequality give (6). For the stated sufficient density control, the characteristic equation is $\frac{d}{dt} \log \widehat{\rho}_t(X_t) = -\operatorname{div}_{S^3} v_t(X_t)$. The initial density is one, so $\widehat{\rho}_t \leq e^A$, where $A = \int_0^1 \| \operatorname{div}_{S^3} v_t \|_\infty dt$. Since $\rho_t \geq 1/2$, $C = 2e^A$ suffices. Smoothness on compact space–time makes A finite for each field. Uniform application across fitted models requires a common bound on A. The teacher energy is the deterministic integral

$$C_u = \int_0^1 \int_0^1 \frac{z(1-z)}{16[1 + t(z-1/2)]} \, dz \, dt = 0.0106027057249\ldots.$$

This establishes the population NMSE conversion. For example, the assumptions $C \leq 2$ and population NMSE ≤ 1.01 imply an action W_1 lower bound of about 0.0688. This numerical illustration assumes certified density control and population NMSE; the trained networks are evaluated using empirical losses.

B.2 Convex continuity penalties on the phase-symmetric path

Use the setting and phase-invariant gradient teacher of Theorem 1. At a fixed time put $r_v = (\partial_t \rho + \operatorname{div}_M(\rho v))/\rho$.

Proposition 7 (Convex residual floor)**.** *Let $v = P_H J \nabla h$ be smooth and let $\Phi : \mathbb{R} \to \mathbb{R}$ be convex. Then*

$$\int_M \rho \Phi(r_v) \, d\mu \geq \int_M \rho \Phi(\partial_t \rho / \rho) \, d\mu.$$

For every $\lambda \geq 0$, zero is the unique minimizing velocity of $\|v - u\|_\rho^2 + \lambda \int_M \rho \Phi(r_v) \, d\mu$ over the smooth projected scalar-Hamiltonian class.

Let $\mathcal{Q}$ average over the independent phase rotations. These isometries preserve H, ρ, μ and commute with the divergence operator. The averaged potential is action-dependent, so $\mathcal{Q}v$ is an angular field with action-dependent coefficients. Each phase generator is divergence-free, and its coefficients and ρ are invariant along that generator. Hence $\mathcal{Q}\,\mathrm{div}_M(\rho v) = \mathrm{div}_M(\rho \mathcal{Q}v) = 0$, so $\mathcal{Q}r_v = \partial_t \rho/\rho$. Jensen's inequality and integration against the invariant measure $\rho\mu$ prove the first conclusion. Equation (1) gives an objective gap of at least $\|v\|_\rho^2$, proving uniqueness. This covers the absolute-residual penalty used in the sphere-adapted FDM experiment under the independent phase invariance assumed here.

B.3 Exact marginals of the independent-endpoint geodesic bridge

Let $X_0 \sim \mu_1$ and $X_1 \sim (1+a)\mu_1$ independently on S^3, where $a = z - 1/2$. Apart from the probability-zero antipodal event, write $X_1 = \cos\Theta\, X_0 + \sin\Theta\, U$ with $\|U\| = 1$ and $U \perp X_0$. Define

$$X_t = \cos(t\Theta)X_0 + \sin(t\Theta)U, \qquad V_t = \Theta[-\sin(t\Theta)X_0 + \cos(t\Theta)U].$$

In the following expectation, Θ has density $2\sin^2\theta/\pi$ on $[0,\pi]$. Set

$$\eta(t) = \frac{1 + 2\mathbb{E}\cos(2(1-t)\Theta)}{3}, \qquad \eta'(t) = \frac{4}{3}\mathbb{E}[\Theta \sin(2(1-t)\Theta)]. \tag{15}$$

Proposition 8 (Geodesic marginal and mean teacher)**.** *For this independent-endpoint bridge,*

$$\rho_t^{\mathrm{geo}} = 1 + \eta(t)a, \qquad \mathbb{E}[V_t \mid X_t = x] = \frac{\eta'(t)\nabla_{S^3} z(x)}{8\rho_t^{\mathrm{geo}}(x)}.$$

The coefficients satisfy $\eta(0) = 0$, $\eta(1) = 1$, $-1/3 \le \eta(t) \le 1$, and $\rho_t^{\mathrm{geo}} \ge 1/2$. The unique population FM-optimal projected scalar-Hamiltonian velocity is zero. Moreover, $\mathbb{E}\|V_t\|^2 = \pi^2/3 - 1/2$ for every t.

To verify the marginal, first take both endpoints uniform. Rotational invariance makes X_t uniform and independent of Θ. Given $X_t = x, \Theta = \theta$, the endpoint X_1 is uniform on the latitude at distance $(1-t)\theta$ from x. Write $a(x) = x^\top A x$, where $A = \mathrm{diag}(1/2, -1/2, 1/2, -1/2)$. For a uniform tangent unit vector W at x, $\mathbb{E}[W^\top A W] = (\mathrm{tr}\, A - a(x))/3 = -a(x)/3$ and $\mathbb{E}W = 0$. Thus the conditional latitude mean of $a(X_1)$ is $[\cos^2((1-t)\theta) - \sin^2((1-t)\theta)/3]a(x)$. Tilting the uniform endpoint pair by $1 + a(X_1)$ proves (15) and the marginal density. Direct integration gives the endpoint coefficients and the displayed range gives positivity.

The bridge and tilted law are equivariant under $O(2) \times O(2)$ acting on the two oscillator planes. Reflections rule out angular components of the conditional mean; its remaining tangent component is in the ∇z direction. The marginal continuity equation determines its flux up to a constant. Smoothness at $z = 0, 1$ gives zero boundary flux, and $-\Delta_{S^3}(z - 1/2) = 8(z - 1/2)$ gives the displayed teacher. Equivalently, its action flux is $\eta'(t)z(1-z)/2$. The weighted product of this teacher with a backbone field integrates to zero by (14). The conditional-mean regression identity then makes zero the unique population-optimal velocity even for the sampled geodesic labels. Since X_0 is uniform independently of X_1, the distance distribution remains $2\sin^2\theta/\pi$ under the target tilt, and $\|V_t\|^2 = \Theta^2$ gives the stated label energy.

The coefficient initially decreases: $\eta'(0) = -1/2$, while $\eta(1/2) = 1/3$ and $\eta'(1) = 0$. The initial reversal gives a nonmonotone traversal of the density family. Quadrature gives integrated mean-teacher energy $0.0162288618\ldots$ and conditional label energy $2.7898681337\ldots$; their difference is irreducible regression noise. About 99.42% of the latter is noise. These deterministic values were computed before training. The finite-error proof also applies to this path with its own excess marginal risk, because the same orthogonality, endpoint mean shift, and density lower bound hold. Both paths belong to the same phase-symmetric density family on the fixed sphere.

C Population analysis of correction regularization

The experiments use $\gamma = 0$. This section analyzes the population objective for $\gamma \ge 0$.

C.1 Joint learning with a gradient-energy penalty

An ideal design minimizes correction energy subject to the required density response:

$$\min_{\theta,\phi} \int_0^1 \|c_\phi(t,\cdot)\|_{\rho_t}^2 \, dt \quad \text{subject to} \quad D_{\rho_t}(b_\theta + c_\phi) = \partial_t \rho_t. \tag{16}$$

This constrained criterion motivates a penalized regression objective. Given FM labels V_t at $X_t \sim P_t$, consider

$$\boxed{\mathcal{L}_\gamma(\theta,\phi) = \mathbb{E}_{t,X_t,V_t}\left[\|b_\theta(t,X_t) + c_\phi(t,X_t) - V_t\|^2 + \gamma\|c_\phi(t,X_t)\|^2\right], \quad \gamma \geq 0.} \tag{17}$$

Both potentials are optimized jointly. The penalty is the squared norm of the tangent correction velocity. The gradient output can be initialized at zero. Automatic differentiation evaluates both branches using the same projector during training and sampling, and sampled labels suffice to evaluate the objective.

With finite second moments, the label mean $u_t(x) = \mathbb{E}[V_t \mid t, X_t = x]$ replaces V_t in (17) up to a parameter-independent noise term. We assume this mean is a tangent teacher generating P_t. Fitting its velocity is stronger than matching its density response. Thus (17) penalizes velocity error, while (16) constrains the density response. Same-path teacher changes can alter a restricted FM fit (subsection D.4); the chosen teacher remains part of the protocol.

C.2 Response fidelity and the cost of regularization

Population response identity. At a fixed time, suppose a population stationary pair (b, c) permits arbitrary smooth gradient variations of c, with b held fixed during that variation. For every smooth ψ,

$$\langle b + c - u, \nabla_M \psi\rangle_\rho + \gamma\langle c, \nabla_M \psi\rangle_\rho = 0.$$

Consequently,

$$D_\rho(b + c - u) = -\gamma D_\rho c, \qquad \partial_t \rho - D_\rho(b + c) = \gamma D_\rho c. \tag{18}$$

At $\gamma = 0$, the completed population fit preserves the teacher's density response. Generating the path additionally requires a sufficiently regular velocity and uniqueness of the continuity equation. For $\gamma > 0$, the penalty generally leaves response error when the gradient carries nonzero response. The identity applies to stationary pairs admitting the stated gradient variations (subsection D.3).

Response bias in the two-oscillator example. For the two-oscillator example, $a = z - 1/2$ and $p_t = 1 + ta$. Its teacher is the gradient of

$$g_t^{\mathrm{ex}} = \frac{\log(1+ta)}{8t} \quad (t > 0), \qquad g_0^{\mathrm{ex}} = \frac{a}{8}, \qquad u_t = \nabla_M g_t^{\mathrm{ex}}. \tag{19}$$

The extension through $t = 0$ is smooth. The unpenalized analytic class comparison is in Table 4 (Appendix A). Since $u_t \perp \mathcal{V}_H$ in $L^2(\rho_t)$, the complete function classes have, for $\gamma > 0$, the unique optimal velocity pair

$$b_\gamma^\star = 0, \qquad c_\gamma^\star = \frac{u_t}{1+\gamma}, \qquad \partial_t \rho_t - D_{\rho_t}(b_\gamma^\star + c_\gamma^\star) = \frac{\gamma}{1+\gamma}\partial_t \rho_t. \tag{20}$$

The proof is in subsection D.3; at $\gamma = 0$, $(0, u_t)$ is an optimal pair, and every admissible pair whose sum is u_t has the same risk. The required response remains in the gradient branch. The displayed identity evaluates response error on the prescribed path; endpoint error uses the generated law.

Experimental setting. The learned experiments use unpenalized completion ($\gamma = 0$) and measure regression, density response, endpoint discrepancy, and conservation. The divergence-objective weights in subsection 4.2 are distinct from γ.

D Density response, minimum correction, and regression

D.1 Proof of minimum response correction

Fix the compact connected boundaryless leaf in Proposition 6. Smooth strict positivity of ρ implies uniform positivity. Normalization and the divergence theorem give $\int_M R_b \, d\sigma = 0$. The weighted elliptic operator $L_\rho = -\operatorname{div}_M(\rho \nabla_M)$ is coercive on mean-zero $H^1(M)$ by Poincaré's inequality. Lax–Milgram and elliptic regularity give a unique smooth mean-zero solution g of $L_\rho g = R_b$. If c is any feasible tangent $L^2(\rho)$ correction, then $k = c - \nabla_M g$ satisfies $D_\rho k = 0$ weakly. Taking g as test function gives $\langle k, \nabla_M g \rangle_\rho = 0$, hence

$$\|c\|_\rho^2 = \|\nabla_M g\|_\rho^2 + \|k\|_\rho^2.$$

This proves uniqueness of the minimum correction as a velocity, independently of the constant gauge for its potential. The smooth family L_{ρ_t} has smooth normalized solutions for smooth data. Thus $b_t + \nabla_M g_t$ is a smooth time-dependent tangent field with complete evolution on compact M; uniqueness of the continuity equation gives its pushforward identity for the prescribed density path.

For a zero-mean response R, define the homogeneous weighted negative Sobolev norm by

$$\|R\|_{\dot{H}^{-1}(\rho)} = \sup_{\psi \in C^\infty(M),\, \|\nabla_M \psi\|_\rho \le 1} |\langle R, \psi \rangle|.$$

Constants are immaterial because R has zero mean. If $R = L_\rho g$, integration by parts and Cauchy–Schwarz bound this supremum by $\|\nabla_M g\|_\rho$. For nonzero R, the test $\psi = g/\|\nabla_M g\|_\rho$ attains the bound; for $R = 0$ both sides vanish. This proves the norm identity in the main text. For $R_b = D_\rho(u - b)$, the same construction yields (12). The orthogonal pair is the correction and the zero-response residual. Inner products involving the backbone may be nonzero.

Closed-space residuals. At a fixed smooth positive density on compact connected M, the residual $r = u - \Pi_H u$ belongs to $L^2(\rho; TM)$ and $D_\rho r$ is a mean-zero H^{-1} functional. Lax–Milgram gives a unique mean-zero weak solution $g \in H^1(M)$ of $-\operatorname{div}_M(\rho \nabla_M g) = D_\rho r$. For $k = r - \nabla_M g$, weak integration by parts gives $D_\rho k = 0$ and $\langle k, \nabla_M g \rangle_\rho = 0$, proving (12). The same minimum-norm argument applies in this weak setting as well. Smooth response data yield a smooth potential by elliptic regularity; only then do we use a smooth ambient extension to write $\nabla_M g = P_H \nabla \widetilde{g}$. Classical flow realization requires the smoothness hypotheses above; the weak decomposition gives orthogonality between the gradient correction and k.

D.2 Tangency, completeness, and multiple constraints

For a regular invariant map $F = (F_1, \dots, F_k)$, let $N = [\nabla F_1 \ \cdots \ \nabla F_k]$. The pointwise projector is

$$P_F = I - N(N^\top N)^{-1} N^\top. \tag{21}$$

Along a tangent flow, the chain rule gives $dF(X_t)/dt = 0$. Thus $F_\# P_t = F_\# P_0$: the base invariant law must already agree with the target's. This elementary compatibility requirement is distinct from the backbone restriction on a fixed compatible shell. In this subsection F denotes a vector-valued invariant map; Section 2 uses $F(I)$ for a scalar action function.

A smooth time-dependent tangent field on a compact regular shell is complete on the finite generation interval. The weighted Poisson construction therefore supplies the classical path realization under the smooth hypotheses of Proposition 6. The closed-space criterion gives a weak response identity. Numerical integration approximates the continuous-time flow. The experiments report conservation on raw sampled states, with their stated integrators.

For a smooth tangent teacher, fitting the full gradient class gives $-\operatorname{div}_M(\rho \nabla_M g) = D_\rho u$ and hence the same density response. The fitted field may differ pointwise from the teacher by a zero-response component. In the general sphere construction, $u^\flat = \epsilon \, db/p_t$ and $d(u^\flat) = -t\epsilon^2 \, da \wedge db/p_t^2$ can be nonzero. For the eigenfunction example $b = a/\lambda$, the teacher itself is a gradient.

D.3 Regularized population fitting and the sphere example

Assume finite second moments and let $u = \mathbb{E}[V \mid t, X]$. Conditional expectation gives

$$\mathbb{E}\|b + c - V\|^2 = \mathbb{E}\|b + c - u\|^2 + \mathbb{E}\|V - u\|^2.$$

The second term is independent of both learned potentials. At a fixed time and density, write $F_\gamma(b,c) = \|b+c-u\|_\rho^2 + \gamma\|c\|_\rho^2$. For a stationary pair that admits all smooth gradient variations, differentiation at $c + s\nabla_M\psi$ gives

$$\tfrac{1}{2}\frac{d}{ds}F_\gamma(b, c+s\nabla_M\psi)\Big|_{s=0} = \langle b+c-u+\gamma c, \nabla_M\psi\rangle_\rho = 0.$$

The weak definition of D_ρ proves (18). For an integrated-in-time function-space objective, arbitrary smooth time-localized variations yield this identity for almost every time; smooth fitted fields extend it to every time. The conclusion applies to a stationary pair admitting the stated function-space variations. Applying it to a neural stationary point requires those variations to be available in its parameterization.

For the sphere task in Equation 19, the smooth teacher is a gradient and Theorem 3 gives $\langle u, b\rangle_\rho = 0$ for all $b \in \mathcal{V}_H$, including its weighted closure. Let $c_0 = u/(1+\gamma)$. For any admissible b and gradient $c = c_0 + d$, direct expansion gives

$$F_\gamma(b,c) - F_\gamma(0,c_0) = \|b+d\|_\rho^2 + \gamma\|d\|_\rho^2.$$

The cross terms in d cancel because $c_0 - u = -\gamma c_0$, and the remaining cross term in b vanishes by orthogonality. For $\gamma > 0$ equality requires $d = 0$ and $b = 0$, proving the unique optimal velocity pair in (20). At $\gamma = 0$, $(0, u)$ is optimal but any admissible pair with $b + c = u$ has the same risk. Distinct potentials may always represent the same velocity. Applying D_ρ gives the stated response error, evaluated on the prescribed density. Endpoint error uses the learned flow's own marginals.

More generally, for a fixed backbone the unpenalized gradient fit is the minimum response correction $c^\star$. Orthogonal projection onto the full gradient class gives the penalized fit $c^\star/(1+\gamma)$. When the backbone is also optimized this scaling requires additional structure; the sphere example supplies the orthogonality needed for the explicit joint solution. Correction energy measures the population velocity penalty, while endpoint accuracy is evaluated under the generated distribution.

D.4 Path realizability and teacher-dependent regression

At a fixed density, write $s = \partial_t\rho$. The condition $s \in D_\rho(\mathcal{V}_H)$ states that a velocity in the class has the required response. By contrast, for a chosen legal teacher u, its restricted population FM fit is response-correct exactly when $D_\rho(u - \Pi_H u) = 0$. These are different conditions. They concern fixed-time weak responses; generating a path additionally requires a sufficiently regular time-dependent selection. If $D_\rho z = 0$, then $u + z$ generates the same response as u, but linearity of the projection gives

$$D_\rho\Pi_H(u+z) - D_\rho\Pi_H u = D_\rho\Pi_H z.$$

Thus a same-path teacher change can alter the fitted response unless this last term vanishes. The sampled-label decomposition and a separate normal-velocity gauge example are in subsection D.5.

D.5 Projected regression and marginal-path compatibility

Let V_t^* be a sampled target and $u_t(x) = \mathbb{E}[V_t^* \mid t, X_t = x]$ its marginal mean. With finite second moments and full-rank $DF(X_t)$, the projected loss is

$$\mathcal{L}_{\tan}(\theta) = \mathbb{E}\|P_F w_\theta - V_t^*\|^2 \tag{22}$$

$$= \underbrace{\mathbb{E}\|P_F(w_\theta - u_t)\|^2}_{\text{tangent approximation}} + \underbrace{\mathbb{E}\|(I-P_F)u_t\|^2}_{\text{normal marginal velocity}} + \underbrace{\mathbb{E}\|V_t^* - u_t\|^2}_{\text{conditional-target noise}}. \tag{23}$$

Let $\zeta = V_t^* - u_t(X_t)$, so $\mathbb{E}[\zeta \mid t, X_t] = 0$. The conditional-mean identity gives

$$\mathbb{E}\|P_F w_\theta - V_t^*\|^2 = \mathbb{E}\|P_F w_\theta - u_t\|^2 + \mathbb{E}\|\zeta\|^2.$$

Since P_F is orthogonal, $P_F(w_\theta - u_t)$ and $(I - P_F)u_t$ are perpendicular. Splitting the first term proves Equation 23. The same argument gives

$$\mathbb{E}\|(I-P_F)V_t^*\|^2 = \mathbb{E}\|(I-P_F)u_t\|^2 + \mathbb{E}\|(I-P_F)\zeta\|^2.$$

Thus the normal component of a sampled training target contains both marginal-velocity mismatch and conditional noise. This is the usual conditional-to-marginal flow-matching identity (Lipman et al., 2023; Chen and Lipman, 2024), followed by orthogonal projection.

A constant invariant law with a normal generating velocity. Take $P_t = \mathcal{N}(0, I_2)$ for every t, $F(x_1, x_2) = x_1$, and $u(x) = (-x_2, x_1)$. Its flow is rotation, which preserves P_t, so $F_{\#}P_t$ is constant. Its normal component is $(-x_2, 0)$, with squared $L^2(P_t)$ norm one. The zero velocity also generates this stationary path and is tangent to F. Hence a normal component of a chosen velocity diagnoses that velocity, while the marginals may admit a tangent alternative.

E Data, training, and evaluation protocols

Velocity classes and metrics. We write B (Backbone) for $P_H J \nabla h$, G (Gradient) for $P_H \nabla g$, M (Mixed) for their sum, and V (Tangent) for $P_H w$. Multiple constraints use P_F from (21). MSE is $\mathbb{E}\|v - u\|^2$ and marginal-teacher NMSE divides it by $\mathbb{E}\|u\|^2$. Geodesic-bridge and FPUT conditional-label NMSE instead include conditional target variance, as explained below and in subsection D.5. Neural comparisons report five training seeds and sample SD unless stated otherwise. This variation measures training randomness conditional on the fixed data source. FM and prescribed-path RFM share one sphere run.

Action W_1 is one-dimensional: z on the sphere and I_1 on the nonquadratic shell. FPUT-32 reports mode-fraction W_1 and full-state sliced W_1 (SW1) in unstandardized coordinates. The two metrics measure discrepancies in different spaces. In the supplementary quartic and physical diagnostics, MMD^2 is the nonnegative biased RBF estimate on 1,024 states; target-only standardization and median calibration use an independent target sample, with kernel scales $0.5, 1, 2$ times the median. Independent target references can give nonzero oracle error at finite sample size.

Conservation summaries remain protocol-specific: mean absolute error for sphere/quartic, mean relative drift for physical controls, and maximum relative error over evaluated trajectories and times for FPUT-32. Each table identifies its conservation estimator.

E.1 Fixed-teacher sphere: objective and protocol

This is the protocol for subsection 4.2. This study contains 30 trained cells: FM with the Hamiltonian backbone, three divergence weights with that backbone, FM with the mixed class, and a mixed-class divergence positive control, each with five seeds. FM and prescribed-path RFM are one shared set of runs.

Write $a = z - 1/2$, $p_t = 1 + ta$, and let μ be normalized sphere area. The residual is

$$r_\theta = \mathrm{div}_{S^3} v_\theta + \langle \nabla_{S^3} \log p_t, v_\theta \rangle + \frac{a}{p_t}.$$

Because $\mathrm{div}_{S^3}(p_t u_t) = -a$, this is the weighted divergence of the velocity error, divided by p_t. The implemented objective is

$$L_\lambda = \frac{\mathbb{E}_{t,p_t}\|v_\theta - u_t\|^2}{C_u} + \lambda \frac{\mathbb{E}_{t,p_t}|r_\theta|}{1/4}, \qquad C_u = \mathbb{E}_{t,p_t}\|u_t\|^2 \approx 0.0106027057.$$

The constants are independent of the model: C_u is computed by 128-point quadrature in each of time and action, and $1/4$ is the exact zero-velocity response baseline. FM has $\lambda = 0$; the sphere-adapted FDM scan uses $0.1, 1, 10$. The comparison table reports the middle weight of the fixed scan. This is an analytic-marginal, intrinsic adaptation of the divergence objective in Huang et al. (2025), with the known marginal replacing the conditional-score estimator. The four-dimensional trace is evaluated exactly by autodifferentiation.

Training uses batch 128, a three-layer width-66 MLP, 16 Fourier time features, AdamW learning rate 0.002 with cosine decay, weight decay 10^{-4}, gradient clipping at one, and 1,000 updates, followed by evaluation of the raw final weights. Initial h and minibatches match within seeds; the extra g output is initialized to zero. All seeds and weights are retained.

The full weight scan and mixed-class positive control are in Table 9.

The directory `02_sphere` contains per-seed histories, checkpoints, endpoint arrays, intermediate action means, and CSVs. Additional columns give absolute MSE, full-state SW1, and the moment-response witness. Run `reproduce_sphere_experiment.py` with `--full` at the project

root; `--check-only` checks the implementation. Each run creates a fresh output directory. The source imports the adjacent sphere and model utilities. An independent orthonormal tangent-frame calculation reproduces the regression norm and the directly differentiated flux response on saved models, including parameter gradients. A finite-difference check of the divergence-loss parameter derivative has error below 2×10^{-12}; 64/128-step RK4 refinement on subsets differs by at most 1.1×10^{-6} in endpoint norm. These checks verify the implementation of the specified teacher, loss, derivatives, and integration scheme.

Table 5: **Fixed sphere path: training objective versus velocity class.** Five seeds, 1,000 updates; mean $\pm$ sample SD. All columns are lower-is-better. FM and RFM share a single set of runs because their induced-metric objectives coincide here. FDM denotes the sphere-adapted, analytic-marginal objective with $\lambda = 1$. The oracle's finite-sample action W_1 is 0.0031 ± 0.0003.

Training objective / velocity class	Teacher NMSE	Action W_1	Mean $\lvert H - E \rvert$
FM=RFM / Backbone (B)	$1.0000 \pm 4.6\times10^{-5}$	0.0840 ± 0.0024	$1.1\times10^{-7} \pm 9.9\times10^{-10}$
Sphere FDM / Backbone (B)	$1.0000 \pm 8.6\times10^{-5}$	0.0840 ± 0.0024	$1.1\times10^{-7} \pm 1.5\times10^{-9}$
FM / Mixed (M)	0.0164 ± 0.0012	0.0076 ± 0.0019	$1.1\times10^{-7} \pm 9.6\times10^{-10}$

E.2 Frozen class–path comparison

The directory `09_runs/class_path_20260924_frozen` contains the pre-execution protocol, source, smoke results, SHA-256 freeze manifest, start marker, training ledger, raw final checkpoints, per-cell arrays, and summary. The protocol fixes all hyperparameters and paths before the new training runs. Its SHA-256 hash is

`e8690be3cd4a37f1fec91d00e49df66dc50e94f648210764eaf6f408e5d1e8d9`.

Width is 66 in each of three SiLU hidden layers, with the retained 16-dimensional Fourier time embedding. B/G/M/V have 10,297/10,297/10,364/10,498 parameters. Initial B and G potentials agree; M copies the B trunk and Hamiltonian head and initializes its gradient head to zero; V copies the same hidden trunk with a seeded vector head. Each new cell uses 1,000 updates, batch size 128, AdamW learning rate 0.002, weight decay 10^{-4}, cosine decay, and global gradient clipping at one. Training runs for the full budget and evaluation uses the raw final weights. The ten existing analytic B/M final checkpoints are imported; analytic G/V and all geodesic cells give thirty new training starts. The cost table reports new training times in the current environment. All forty evaluations run in the current CPU environment.

Seeds are $s = 0, \ldots, 4$. Initialization uses $20000 + s$, training sources and times $31000 + s$, and independent geodesic training endpoints $71000 + s$. The analytic squared loss is divided by its label energy C_u; the geodesic loss is divided by $\pi^2/3 - 1/2$. Each path therefore uses its own label-energy normalization; their signal-to-noise ratios differ. Mean-teacher and conditional-label NMSE are reported separately, each divided by the corresponding evaluation energy. The exact mean is used for evaluation; geodesic training uses the sampled conditional labels.

Evaluation sources/times use $41000 + s$ and geodesic evaluation endpoints $91000 + s$, disjoint from the training endpoint stream. Rollout sources use $51000 + s$; independent references use $61000 + s$; SW1 directions use $81000 + s$. Sample sizes are 4,096, with 64 SW1 directions. Evaluation uses the raw output of 64-step RK4. Saved times are $0, 1/4, 1/2, 3/4, 1$. Action W_1 integrates the empirical CDF difference against $G_t(z) = z + \eta(t)(z^2 - z)/2$ on a 4,097-point uniform grid. SW1 compares equally sized sorted projections in original coordinates. Path integrals use the fixed five-node trapezoid rule to estimate the continuous-time integrals.

For the weak diagnostic, nine equally spaced times use 4,096 fresh exact marginal samples each, with seeds $101000 + 100s + j$. For $\psi = z$, record

$$e_j = \widehat{\mathbb{E}}_{\rho_{t_j}}[v \cdot \nabla z] - \eta'(t_j)/12, \qquad W_{\mathrm{resp}} = \frac{\{\mathrm{trapz}_j\, e_j^2\}^{1/2}}{\{\mathrm{trapz}_j(\eta'(t_j)/12)^2\}^{1/2}}.$$

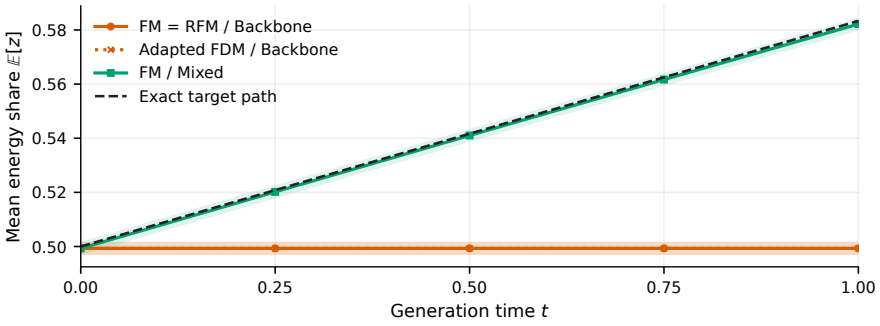


Figure 3: Retained analytic-path objective controls: mean action over five learned rollouts (mean and sample SD). The target has $\mathbb{E}z = 1/2 + t/12$. These curves precede the frozen class–path extension and are retained unchanged.

For the analytic path, $\eta(t) = t$. Arrays also retain paired teacher-residual means and pointwise Monte Carlo standard errors. This diagnostic evaluates the residual in the action direction selected by $\psi = z$.

The integration gate compares 64- and 128-step endpoints for the first 256 sources; maximum position discrepancy must be at most 10^{-4}. Every cell retains its original solver, checkpoint, and gate outcome. Conservation reports mean endpoint and maximum saved-frame absolute $|H - 1/2|$. One warmed 4,096-sample rollout and each new training run are timed on the CPU with four PyTorch threads; these are descriptive costs for the current environment.

The primary contrast is paired B versus G endpoint action W_1 for each path. The predeclared descriptive separation criterion requires positive mean B-minus-G error and at least four of five G wins. M/G comparisons are descriptive; pure gradients already satisfy the population response-preservation criterion.

Execution outcome. Exactly 30 new training starts, 10 imports, and 40 evaluations completed in a single formal execution. All integration gates passed; the largest refinement difference was 1.10×10^{-6} and the largest saved-frame absolute energy error was 5.96×10^{-7}. The exact endpoint sampler has action W_1 0.0031 ± 0.0003 and SW1 0.0121 ± 0.0018. The new CPU runtime used PyTorch 2.13.0, Python 3.13.2, and four threads. V combines low geodesic endpoint error with nonzero weak-response and integrated-path errors.

E.3 Nonquadratic conditional generation

Write $z = I_1$ and $z_{\max} = (\sqrt{5} - 1)/2$. On $H = 1$, $I_2(z) = \sqrt{2 - z - z^2} - 1$. We tabulate the normalized surface-area pushforward $f_0(z)$ using the weight in Appendix A and set

$$c = \int_0^{z_{\max}} z f_0(z)\, dz, \qquad A(z) = \int_0^z (s - c) f_0(s)\, ds, \qquad F_{t,y}(z) = F_0(z) + tyA(z). \tag{24}$$

The conditional task draws $y = \pm 1$ with equal probability. Source $u \sim \mathrm{Unif}(0, 1)$ is converted to $z_t = F_{t,y}^{-1}(u)$. Two uniform phase angles and $q_i = \sqrt{2I_i}\cos\theta_i$, $p_i = \sqrt{2I_i}\sin\theta_i$ give a shell state. The teacher is

$$\dot{z} = -\frac{yA(z)}{[1 + ty(z - c)]f_0(z)}, \qquad u_{t,y}(x) = \frac{\dot{z}}{\|P_H \nabla I_1(x)\|^2} P_H \nabla I_1(x). \tag{25}$$

The same tabulation directly evaluates the teacher. The Python implementation uses 65,537 quadrature points, 27 rounds of inverse-CDF bisection, and explicit projection onto $\ker dH$.

Training and evaluation. Five seeds use 300 AdamW updates, width 66, batch 512, learning rate 0.002 with cosine decay, weight decay 10^{-4}, and clipping at one. We evaluate raw final weights. The mixed model shares a trunk, copies the backbone initialization, and initializes the extra gradient head to zero. Backbone/vector checkpoints are retained matched controls. Each condition uses

Table 6: **Frozen comparison: response, conditional labels, and path discrepancy.** Five-seed mean $\pm$ sample SD. W_{resp} is the fixed action-response diagnostic. Path integrals use five saved times. Analytic labels are their own marginal mean.

Path	Class	W_{resp}	Label NMSE	$\int W_1^z\, dt$	$\int \mathrm{SW1}\, dt$
Analytic	B	$1.0000 \pm 8.6\times10^{-6}$	$1.0000 \pm 4.6\times10^{-5}$	0.0427 ± 0.0021	0.0170 ± 0.0016
	G	0.0093 ± 0.0045	0.0174 ± 0.0008	0.0043 ± 0.0009	0.0121 ± 0.0017
	M	0.0119 ± 0.0045	0.0164 ± 0.0012	0.0042 ± 0.0009	0.0121 ± 0.0017
	V	0.0075 ± 0.0021	0.0174 ± 0.0006	0.0043 ± 0.0009	0.0121 ± 0.0017
Geodesic	B	$1.0000 \pm 4.0\times10^{-5}$	$1.0000 \pm 2.4\times10^{-5}$	0.0354 ± 0.0016	0.0164 ± 0.0017
	G	0.8591 ± 0.0699	0.9988 ± 0.0004	0.0274 ± 0.0041	0.0150 ± 0.0014
	M	0.9825 ± 0.0117	0.9999 ± 0.0001	0.0344 ± 0.0015	0.0163 ± 0.0018
	V	0.5260 ± 0.0129	0.9962 ± 0.0009	0.0096 ± 0.0006	0.0125 ± 0.0015

Table 7: **CPU costs and conservation.** Five-seed means; full SD and individual values are in the released CSV/JSON. Training times exclude evaluation; imported times are omitted. Sampling is one warmed 4,096-state, 64-step rollout per cell. Mean absolute endpoint energy errors are in units of 10^{-7}.

Path	Class	Train (s)	Sample (s)	Mean $\lvert H - 1/2\rvert$ (10^{-7})
Analytic	B	imported	0.250	1.044
	G	0.583	0.249	1.044
	M	imported	0.355	1.052
	V	0.418	0.146	1.048
Geodesic	B	0.602	0.256	1.052
	G	0.605	0.248	1.051
	M	0.805	0.352	1.050
	V	0.445	0.150	1.044

2,048 endpoint source states and 32-step RK4. Action W_1 and regression metrics average the two conditions within each seed; event error takes the larger absolute error in $P(I_1 > c \mid y)$. Additional metrics are in Table 10.

Reproduction. The tabulation and exact teacher are implemented in `06_shared_core/quartic_shell.py`. Run the mixed comparison with `03_quartic_and_physical/train_quartic_and_sphere_baselines.py`, using `--section e3`. Outputs include per-seed metrics and conditional endpoint CDFs. Run training in a fresh output directory: the implementation skips cells whose output files already exist.

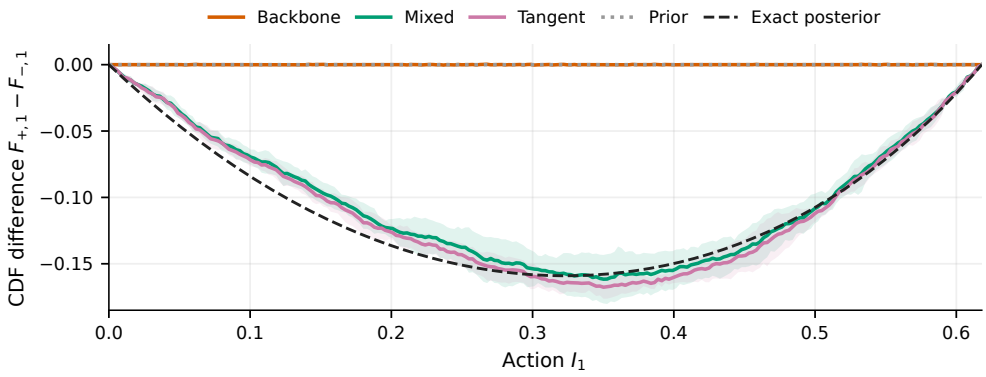


Figure 4: **Conditional response on a nonquadratic shell.** At the endpoint, the difference of action CDFs for $y = +1$ and $y = -1$ follows the exact posterior response for Mixed and Tangent models; Backbone and prior remain near zero. Lines and bands show five-seed mean and sample SD (300 updates).

E.4 Published FPUT-32 states: conditional task and protocol

Source and preprocessing. We use `dataMatrix_3.0.npy` from Zenodo record 15873646 (Marchetti, 2025), released under CC BY 4.0. The file contains four million 64-dimensional states of a fixed-end FPUT-β chain, $N = 32$, $\beta = 3$, initially excited in mode one with amplitude $A = 10$ and physical integration step 0.05. The local data consist of four 4,096-row ranges beginning at indices 0, 100000, 1000000, and 3000000. Verification covers these downloaded ranges and their checksums. The first 32 columns are positions and the last 32 are momenta, checked against the Hamiltonian equations. The full energy is

$$H(q,p) = \frac{1}{2}\sum_{i=1}^{32} p_i^2 + \sum_{i=0}^{32}\left[\frac{1}{2}(q_{i+1}-q_i)^2 + \frac{3}{4}(q_{i+1}-q_i)^4\right], \qquad q_0 = q_{33} = 0.$$

For each range, local rows $[0, 2457)$ are used for training, $[2662, 3276)$ for validation, and $[3481, 4096)$ for testing. These time blocks retain serial dependence and phase distribution shift. The low/middle/high condition groups respectively contain $2457/4914/2457$ training states, $614/1228/614$ validation states, and $615/1230/615$ test states. These are conditional empirical occupation samples from one orbit, with the dependence structure inherited from that trajectory.

Observation and bridge. Apply the orthonormal discrete sine transform to obtain modal positions Q_k and momenta P_k, with $\omega_k = 2\sin(k\pi/66)$ and $E_k^{\mathrm{harm}} = (P_k^2 + \omega_k^2 Q_k^2)/2$. The observed label is the interval of $r = E_1^{\mathrm{harm}}/\sum_k E_k^{\mathrm{harm}}$. All states are retracted to $E = H(x_{\mathrm{first}}) \approx 0.480767040670$. Writing $H(sy) = A(y)s^2 + B(y)s^4$, the positive radial scale satisfies

$$s^2 = \frac{2E}{A + \sqrt{A^2 + 4BE}}, \qquad R_E(y) = sy.$$

The raw relative energy drift is at most 6.66×10^{-4} and the maximum relative state correction is 3.28×10^{-4}. We form 4,096 training pairs and 512 validation and test pairs per condition, using an equal mixture of the three source groups and size-32 Euclidean optimal-assignment minibatches. The target endpoint is sampled within the requested group. For $y_t = (1-t)x_0 + tx_1$, $x_t = s_t y_t$, and $d = x_1 - x_0$,

$$\dot{x}_t = s_t d + \dot{s}_t y_t, \qquad \dot{s}_t = -s_t \frac{\nabla H(x_t)\cdot d}{\nabla H(x_t)\cdot y_t}.$$

Bridge energy, tangency, and derivative finite-difference checks pass; the derivative relative error across sampled conditions is below 10^{-9}. This conditional target depends on the endpoint pair. Its regression NMSE includes conditional target variance. The sphere's marginal-teacher NMSE evaluates the known conditional mean directly.

Training and evaluation. Five seeds vary initialization and minibatch/time draws on fixed data. Each model uses Adam with learning rate 10^{-3}, batch 256, two width-96 tanh layers, and 15,000 updates. Validation loss is checked every 250 updates. At budgets 3,000, 7,500, and 15,000, evaluate the checkpoint with the best validation loss available at that budget. All models see matched minibatches within a seed. The wide-backbone diagnostic uses width 192 and seeds 0–2, with the same optimizer and update budget. Generation uses 512 common test source samples per condition and raw RK4 with step 0.01. At 21 equally spaced generation times, compare with the held-out bridge marginals. Full-state SW1 uses 64 fixed Euclidean unit directions in the original coordinates; integrate by the trapezoidal rule. Mode-fraction W_1 is one-dimensional. SW1, mode-fraction W_1, and hit rates average conditions within each seed. Conditional-target NMSE instead pools squared errors and squared target norms over all conditions before taking their ratio. Energy drift is the maximum over evaluated trajectories and times. We then report the five-seed mean and sample SD for the distribution and regression metrics. Seed variation measures training randomness on the fixed data.

Additional robustness and distribution-shift diagnostics are in subsection G.3.

F FPUT SOURCE FILTERING AND NEURAL TRANSPORT

The B/G/M comparisons measure neural transport under the chosen bridge and training budget. The source-filtering diagnostic in subsection 4.3 evaluates the conditional endpoint sampler provided by the source mixture. It uses the original neural checkpoints and their original selection rule.

Filtering replay and access. The diagnostic reads `04_fput32/extended/frozen.npz` and computes condition labels from the source coordinates only. The fixed 64 directions and target arrays are used only to evaluate outputs. Filtering acceptance is exactly $1/3$ for the stipulated balanced mixture; finite source batches differ by integer rounding. Average acceptance cost is three independent source draws per accepted output. The pooled-cache audit returns 513/513/510 outputs; the separate-query audit returns 171/171/170. Both evaluate accepted subsets of the fixed source cache.

The retained cache, source checksums, and result verification are recorded in `04_fput32/extended/frozen.npz`, `04_fput32/extended/frozen_manifest.json`, and `04_fput32/extended/verification.json`. The extended report and per-seed records preserve the filtering inputs and saved neural endpoint metrics. The fifteen retained neural SW1 values agree with direct rollout-array replay to 2.8×10^{-17}, verifying the reported arithmetic from the saved arrays.

F.1 NEURAL COMPARISONS

Data and task. Four windows from the published fixed-end FPUT trajectory ($N = 32$, $\beta = 3$; Marchetti, 2025) supply 16,384 states $x = (q, p) \in \mathbb{R}^{64}$. We condition on three intervals of the harmonic-mode fraction $r(x) = E_1^{\mathrm{harm}} / \sum_{k=1}^{32} E_k^{\mathrm{harm}}$: $r < 0.05$, $0.05 \le r < 0.8$, and $r \ge 0.8$. Given the interval label, a common source mixing the three groups equally is transported to the corresponding empirical conditional distribution. Generation time parameterizes this conditional transport. The observable r measures harmonic modal energy shares; the full nonlinear energy also contains quartic interactions.

Disjoint time blocks with gaps define training, validation, and test sets. A radial correction of at most 3.28×10^{-4} in relative state norm places all states on $H = E \approx 0.480767$. The bridge $x_t = R_E((1-t)x_0 + tx_1)$ uses the exact radial retraction R_E; its derivative is the conditional FM target. We compare backbone, tangent-gradient, and mixed models with paired minibatches, two width-96 hidden layers, and 15,000 updates. The mixed model adds one scalar head (16,130 versus 16,033 parameters). Validation regression loss selects checkpoints. Full protocols are in subsection E.4.

Conditional accuracy. Completion raises condition satisfaction from 47.0% to 97.9% and reduces endpoint full-state sliced W_1 (SW1) by 18.1% in the seed means, improving all five paired seeds (Table 8). Pure gradients also improve condition control. The mixed model's full-state SW1 remains

above the target–target resampling reference (SW1 ≈ 0.0162), showing residual distribution error within the requested intervals.

Table 8: **Conditional generation from published FPUT-32 states.** 15,000 updates, five training seeds; mean $\pm$ sample SD after averaging conditions. Hit rate is higher-is-better; both distances are lower-is-better. Maximum relative energy error across all evaluated budgets and models is 1.25×10^{-5}. The seeds vary training randomness on a fixed single-orbit data split.

Velocity class	Hit rate (%) ↑	Mode W_1 ↓	SW1 ↓
Backbone (B)	47.0 ± 2.5	0.2969 ± 0.0091	0.2015 ± 0.0064
Gradient (G)	94.2 ± 3.2	0.0348 ± 0.0102	0.1732 ± 0.0041
Mixed (M)	97.9 ± 1.5	0.0221 ± 0.0012	0.1649 ± 0.0123

Budget and capacity. Across 3,000–15,000 updates, backbone endpoint SW1 improves from 0.2097 to 0.2015 (Figure 5). A width-192 backbone (50,497 parameters) reaches 0.1939–0.2119 in three seeds, still above paired mixed runs. The path-integrated improvement is 3.0% in the seed means at 15,000 updates, with wins in three of five seeds. These measurements describe the endpoint and path accuracy of finite networks trained on a single orbit. The time-block distribution shift and bridge protocol are specified in subsection E.4.

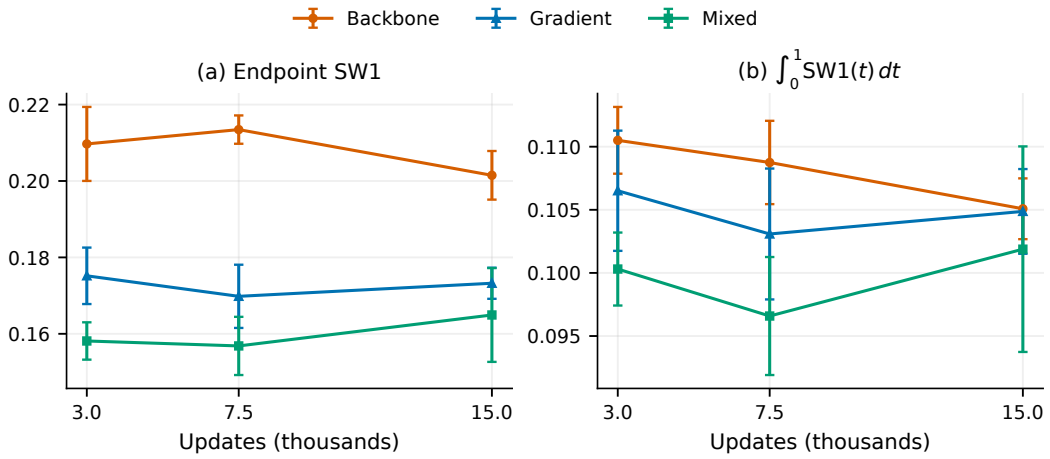


Figure 5: **FPUT-32 conditional generation across training budgets.** Five-seed mean $\pm$ sample SD. The endpoint advantage persists; the integrated-path advantage weakens. SW1 averages 64 fixed unit-direction projections in the original coordinates. The horizontal axis is the training budget; the path integral uses generation time $t \in [0, 1]$.

G Additional metrics and robustness

G.1 Sphere objective scan

The same fixed teacher and training protocol are used for all rows below. The weight scan keeps the backbone class fixed. The mixed-class FDM row checks the objective in a class containing the required response.

Table 9: Fixed-teacher controls, mean ± sample SD. Response is normalized weighted-divergence L1; zero velocity has population value one. The last row is an additional implementation control.

Objective / class	NMSE	Response L1	Action W_1	KS
FM / Backbone (B)	$1.0000 \pm 4.6\times10^{-5}$	1.0049 ± 0.0121	0.0840 ± 0.0024	0.1290 ± 0.0053
FDM $\lambda = 0.1$ / Backbone (B)	$1.0000 \pm 4.4\times10^{-5}$	1.0049 ± 0.0121	0.0840 ± 0.0024	0.1290 ± 0.0053
FDM $\lambda = 1$ / Backbone (B)	$1.0000 \pm 8.6\times10^{-5}$	1.0049 ± 0.0121	0.0840 ± 0.0024	0.1289 ± 0.0053
FDM $\lambda = 10$ / Backbone (B)	1.0001 ± 0.0002	1.0049 ± 0.0121	0.0840 ± 0.0024	0.1289 ± 0.0054
FM / Mixed (M)	0.0164 ± 0.0012	0.1820 ± 0.0080	0.0076 ± 0.0019	0.0169 ± 0.0034
FDM $\lambda = 1$ / Mixed (M)	0.0092 ± 0.0027	0.1082 ± 0.0155	0.0048 ± 0.0017	0.0120 ± 0.0024

G.2 Nonquadratic conditional metrics

Table 10: Conditional quartic: regression and distribution metrics (mean ± sample SD). W_1 compares I_1; metrics average conditions within seed.

Model	MSE	NMSE	W_1	MMD^2	KS
Backbone (B)	$0.0045 \pm 2.4\times10^{-5}$	$1.0001 \pm 5.4\times10^{-5}$	$0.0333 \pm 1.7\times10^{-6}$	0.0010 ± 0.0002	0.0857 ± 0.0015
Mixed (M)	0.0002 ± 0.0002	0.0399 ± 0.0449	0.0053 ± 0.0018	0.0007 ± 0.0002	0.0224 ± 0.0053
Tangent (V)	$0.0001 \pm 3.7\times10^{-5}$	0.0235 ± 0.0083	0.0045 ± 0.0019	0.0007 ± 0.0002	0.0201 ± 0.0066
Prior	–	–	$0.0333 \pm 2.0\times10^{-6}$	0.0010 ± 0.0002	0.0857 ± 0.0016
Exact-flow oracle	–	–	0.0036 ± 0.0022	0.0007 ± 0.0002	0.0169 ± 0.0078

G.3 FPUT-32 budget, capacity, and distribution shift

Detailed robustness results. At 15,000 updates, backbone/gradient/mixed conditional-target NMSE is respectively 0.3600 ± 0.0048, 0.3944 ± 0.0055, and 0.3136 ± 0.0082. Integrated SW1 is 0.1051 ± 0.0024, 0.1049 ± 0.0034, and 0.1019 ± 0.0081. Mixed endpoint SW1 is lower in all five paired seeds; its integrated SW1 is lower in only three. The wide backbone has 50,497 parameters and endpoint SW1 0.2119, 0.1939, and 0.2050, versus paired mixed values 0.1667, 0.1754, and 0.1438. These measurements compare the specified finite networks at the stated budgets.

Distribution shift. Train-target versus test-target SW1 is approximately 0.0692, 0.0772, and 0.7586 for the three conditions, quantifying the empirical shift between time blocks. A separate generation check uses the training distribution. Backbone/gradient/mixed SW1 averages $0.1294/0.0380/0.0508$, with condition hit rates $47.6\%/93.8\%/99.2\%$. This check measures fit to the training distribution. The FPUT comparison is an empirical conditional-transport study on the chosen bridge; the analytic separation theorem concerns smooth phase-symmetric densities. Here tangent-gradient completion improves the measured endpoint transport, with the ranking between pure-gradient and mixed models depending on the metric and data split.

H When does the Hamiltonian backbone already suffice?

Matched physical-system controls compare completion with the backbone on Hénon–Heiles (HH), a central-force system preserving H and angular momentum L, and periodic FPUT chains with $N = 4, 16$. The HH dynamics teacher is $0.2J\nabla H$ and is exactly representable by the backbone. For the other tasks, the teacher is designed as $0.15J\nabla H + P\nabla g$. Here P projects onto the specified invariant leaf, including both constraints for the central-force case. Each target is the pushforward of the specified source by its designed teacher. The published fixed-end FPUT-32 task instead uses observed trajectory states.

The retained runs share training states, initialization, a 300-update budget, 2,048 endpoint samples, and 64-step RK4 across each backbone/mixed pair. Table 11 reports the complete comparison; Table 12 records auxiliary distribution diagnostics. On HH dynamics, both models have essentially identical performance, as expected for a representable teacher. The designed HH transport and FPUT-4 tasks show larger benefits from completion, while central-force gains are smaller. FPUT-16 remains unresolved: the observable error falls from 1.5725 to 1.0589 target standard deviations, but both models are far from the oracle. The controls measure task-dependent gains at 300 updates.

The HH dynamics teacher belongs to the backbone class; the designed transport teachers include gradient components represented by the completed class.

H.1 Systems, sources, and teachers

For HH, $H = \frac{1}{2}\|p\|^2 + \frac{1}{2}(q_1^2 + q_2^2) + q_1^2 q_2 - q_2^3/3$ at $E = 0.125$. The source is the Liouville probability on the compact central-well component: draw q uniformly in the accessible central well, a uniform momentum angle, and $\|p\| = \sqrt{2(E - V(q))}$. Exterior noncompact components are excluded. For central force, $H = \frac{1}{2}\|p\|^2 - 1/r + 0.15/r^2$, $r = \|q\|$, and $L = q_1 p_2 - q_2 p_1$, with $E = -0.3$, $L = 1$. Draw r uniformly between the turning points, a uniform position angle, and a uniform radial-momentum sign. Set $p_{\tan} = L/r$ and $p_r = \pm\sqrt{2(E + 1/r - 0.15/r^2) - L^2/r^2}$. These draws define the source law for the controlled transport task. For periodic FPUT, $q_{N+1} = q_1$ and

$$H = \frac{1}{2}\sum_i p_i^2 + \sum_i \left\{ \frac{1}{2}(q_{i+1} - q_i)^2 + \frac{1}{4}(q_{i+1} - q_i)^4 \right\}, \qquad E = 0.1N, \quad N = 4, 16.$$

The source centers Gaussian q, p and radially rescales to the energy shell; $\sum_i q_i = \sum_i p_i = 0$ are also projected constraints. This procedure induces a specific, generally nonuniform law on the shell. Unlike the fixed-end FPUT-32 task, these are designed teacher-pushforward controls. The transport potentials are $g = -0.15(q_1^2 + p_1^2)$ for HH, $g = -0.075\|q\|^2$ for central force, and $g = -0.15[(q_1 - q_2)^2 + (p_1 - p_2)^2]$ for periodic FPUT.

H.2 Training, sampling, and diagnostics

Models use width 66, batch 256, 300 AdamW updates, learning rate 0.002 with cosine decay, weight decay 10^{-4}, clipping at one, and raw final weights. Each task has 16,384 training and 4,096 held-out states. Mixed heads share a trunk; the extra gradient head is initialized to zero. Teacher paths use 64-step RK4 for training and 128 steps for independent endpoint references; generated endpoints use 64 steps and 2,048 source states. Central projection solves the two-constraint Gram system. Periodic FPUT first removes the position/momentum mean directions and then projects off the energy normal. Endpoint errors are measured on the raw integrator output.

Observable W_1 uses $q_1^2 + p_1^2$ (HH), $\|q\|$ (central), and $[(q_1 - q_2)^2 + (p_1 - p_2)^2]/2$ (FPUT), divided by the population-form SD of an independent target sample. Across training seeds, error bars instead use sample SD. Energy drift is $\mathbb{E}|H(X_1) - H(X_0)|/|E|$, measured before any correction.

Surface-entropy change is estimated along each model's own flow as $\widehat{\Delta S} = \int_0^1 \widehat{\mathbb{E}}\, \mathrm{div}_M\, v_t(X_t)\, dt$, using 128 source states, four projected Hutchinson probes, and nine trapezoidal nodes. The entropy-error column is the absolute model–teacher difference within each seed, then averaged across seeds. It includes Monte Carlo and discretization errors. The teacher mean $\pm$ sample SD is respectively -0.0047 ± 0.0060, -0.0435 ± 0.0334, -0.1997 ± 0.0213, -0.4906 ± 0.0226, and -0.6198 ± 0.0483 in the table's task order. FPUT-16's mixed entropy error exceeds the backbone value.

Per-seed records and source code are in `03_quartic_and_physical`. The entry point is `train_physical_system_controls.py` in that directory. Use a fresh output directory for reproduction. The CSVs include unrounded values and prior/oracle references.

Table 11: **Physical-system transport under matched short budgets.** Five seeds, 300 updates; mean $\pm$ sample SD. W_1 is divided by the target observable's SD. Energy drift is relative to $|E|$. Entropy error is the absolute difference between model and teacher estimates of surface-entropy change, computed within each seed before aggregation; it includes Monte Carlo error. Lower is better for all columns.

Task	NMSE	Observable W_1/SD	Energy drift	Entropy error
Backbone (B)				
HH dynamics	0.0977 ± 0.0018	0.0457 ± 0.0218	$2.0\times10^{-7} \pm 5.6\times10^{-9}$	0.0089 ± 0.0101
HH transport	0.4913 ± 0.0020	0.3665 ± 0.0519	$2.0\times10^{-7} \pm 3.9\times10^{-9}$	0.0488 ± 0.0314
Central H, L	0.0236 ± 0.0005	0.0599 ± 0.0209	$3.9\times10^{-7} \pm 1.1\times10^{-8}$	0.0408 ± 0.0137
FPUT $N = 4$	0.2661 ± 0.0023	0.9175 ± 0.0570	$2.0\times10^{-7} \pm 4.4\times10^{-9}$	0.2943 ± 0.0231
FPUT $N = 16$	0.1140 ± 0.0009	1.5725 ± 0.0885	$1.3\times10^{-7} \pm 2.3\times10^{-9}$	0.4619 ± 0.0209
Mixed (M)				
HH dynamics	0.0977 ± 0.0018	0.0459 ± 0.0219	$2.0\times10^{-7} \pm 2.3\times10^{-9}$	0.0089 ± 0.0102
HH transport	0.0738 ± 0.0260	0.0517 ± 0.0263	$2.1\times10^{-7} \pm 2.9\times10^{-9}$	0.0145 ± 0.0110
Central H, L	0.0125 ± 0.0002	0.0377 ± 0.0177	$3.9\times10^{-7} \pm 4.7\times10^{-9}$	0.0214 ± 0.0214
FPUT $N = 4$	0.0239 ± 0.0117	0.1192 ± 0.0206	$2.0\times10^{-7} \pm 7.3\times10^{-10}$	0.0285 ± 0.0359
FPUT $N = 16$	0.0946 ± 0.0024	1.0589 ± 0.0825	$1.3\times10^{-7} \pm 5.5\times10^{-10}$	0.4886 ± 0.0561

Table 12: Auxiliary physical-control metrics (mean $\pm$ sample SD). B: Backbone; M: Mixed. MMD^2 is the biased full-state estimate; KS uses the task observable. These metrics complement the teacher-fit, observable-error, and conservation measurements.

Task	Class	MSE	MMD^2	KS
HH dynamics	B	$0.0009 \pm 1.9\times10^{-5}$	0.0010 ± 0.0006	0.0315 ± 0.0102
HH dynamics	M	$0.0009 \pm 1.9\times10^{-5}$	0.0010 ± 0.0006	0.0314 ± 0.0102
HH transport	B	$0.0049 \pm 3.7\times10^{-5}$	0.0026 ± 0.0004	0.1531 ± 0.0195
HH transport	M	0.0007 ± 0.0003	0.0010 ± 0.0005	0.0365 ± 0.0127
Central H, L	B	$0.0015 \pm 2.9\times10^{-5}$	0.0010 ± 0.0005	0.0354 ± 0.0052
Central H, L	M	$0.0008 \pm 1.1\times10^{-5}$	0.0010 ± 0.0005	0.0256 ± 0.0080
FPUT $N = 4$	B	0.0197 ± 0.0002	0.0048 ± 0.0006	0.3283 ± 0.0087
FPUT $N = 4$	M	0.0018 ± 0.0009	0.0009 ± 0.0003	0.0619 ± 0.0103
FPUT $N = 16$	B	0.0297 ± 0.0002	$0.0011 \pm 9.1\times10^{-5}$	0.3875 ± 0.0151
FPUT $N = 16$	M	0.0246 ± 0.0006	$0.0010 \pm 8.3\times10^{-5}$	0.3095 ± 0.0158

I Relation to neighboring transport classes

The restricted class is essential. With constant canonical J, $J\nabla h$ is divergence-free and its complete flow preserves ambient volume. Classical equimeasurability restrictions therefore apply to that class (Draxler et al., 2024). Projection generally produces a nonzero shell divergence. For $P_H J\nabla h$, the restriction studied here is the phase-averaged density response.

Manifold transport and minimum correction. Riemannian flow matching learns tangent velocities on a manifold (Chen and Lipman, 2024), while Moser-type constructions and action principles supply transports of smooth positive densities (Moser, 1965; Rozen et al., 2021; Khesin and Lee, 2009; Neklyudov et al., 2023). Our use of the weighted Poisson equation isolates the minimum response correction for a fixed backbone. The result characterizes density-response capacity and the minimum correction norm in the stated function space.

Hard constraints and the chosen path. PCFM and HardFlow impose physical constraints on generation (Utkarsh et al., 2025; Li et al., 2025). Here the invariant is fixed throughout. Our separation concerns prescribed density paths. Composed fields have additional endpoint-transport directions, as illustrated in subsection A.5. The physical controls measure distribution transport under their specified teachers.