EQUIVARIANCE TRANSPORTS, CLOSURE COMPLETES: GAUGE OBSTRUCTIONS IN LATENT WORLD MODELS

Anonymous authors
Paper under double-blind review

ABSTRACT

Equivariance does not make an exact latent-prediction optimum a predictive state. Two finite constructions prove this: a persistent covariant copy under autonomous self-prediction, and a hidden trajectory gauge despite canonical innovation, hard transport, a proper anchor, and a supplied pose–quotient reconstructor. In a free, observation-aligned Borel class, we prove that deterministic equivariant emission driven by the world's specified joint innovation exists if and only if the invariant code carries the point gauge. This specializes standard factorization and is distinct from full-future-law sufficiency with auxiliary noise. We separate the assumptions for quotient-future and full-future sufficiency. Controlled studies then test how supplied information enters learned dynamics. With supplied physical calibration and known association, calibrated routing with multi-step training attains mean H8 NMSE 0.127500 versus 0.675110 for persistence, winning all 20 frozen-encoder blocks. Readout calibration and one-step retraining fail their tested new-object criteria. A separate matched-update study reverses the advantage of root supervision, and sensor-mounting Pendulum favors system identification over learned dynamics. These contrasts locate calibration's benefits without claiming generic predictive superiority.

1 INTRODUCTION

Latent world models often predict a future representation rather than a future observation. Equivariance suggests that such a representation should separate pose from predictive content, as in equivariant world models and class–pose learning (Park et al., 2022; Marchetti et al., 2023; Ghaemi et al., 2025). We ask whether an exact optimum must retain the observation-future law, and which additional interface closes the gap. Our solution concept is the exact global minimum of the declared forward scalar objective, not stationary points or arbitrary training dynamics.

Three anchors, three boundaries. A *statistical* anchor constrains the embedding law under a declared distributional class (Klindt et al., 2026); adding action excitation still leaves one global orthogonal ambiguity (Zhang et al., 2026a). Even a full-distribution pin constrains only the pooled marginal, not which embeddings share a trajectory. A *token* anchor needs a sufficient predictive horizon (Teoh et al., 2025). A *supervised* class–pose guarantee observes the group element and fixes the latent action (Marchetti et al., 2023). These guarantees leave open the predictive sufficiency of a specified equivariant history code.

The answer is a residual gauge. On a free orbit an equivariant code is $F_{\mathrm{orb}}(p, s) = p\, c(s)$: the physical action transports the orbit coordinate, while the quotient-dependent factor $c(s)$ remains unselected (Marchetti et al., 2023, Proposition 2.1). If that factor is not recoverable from the invariant code, two histories may collide in latent space while inducing different future laws even when all declared losses attain their global lower bounds.

Contributions.

1. We construct two finite exact optima that are not predictive states: a persistent covariant copy and a hidden trajectory gauge that survives canonical innovation, hard transport, a proper anchor, QCI/ICS, and a supplied pose–quotient observation reconstructor (Theorems 3.1 and 3.2).

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.

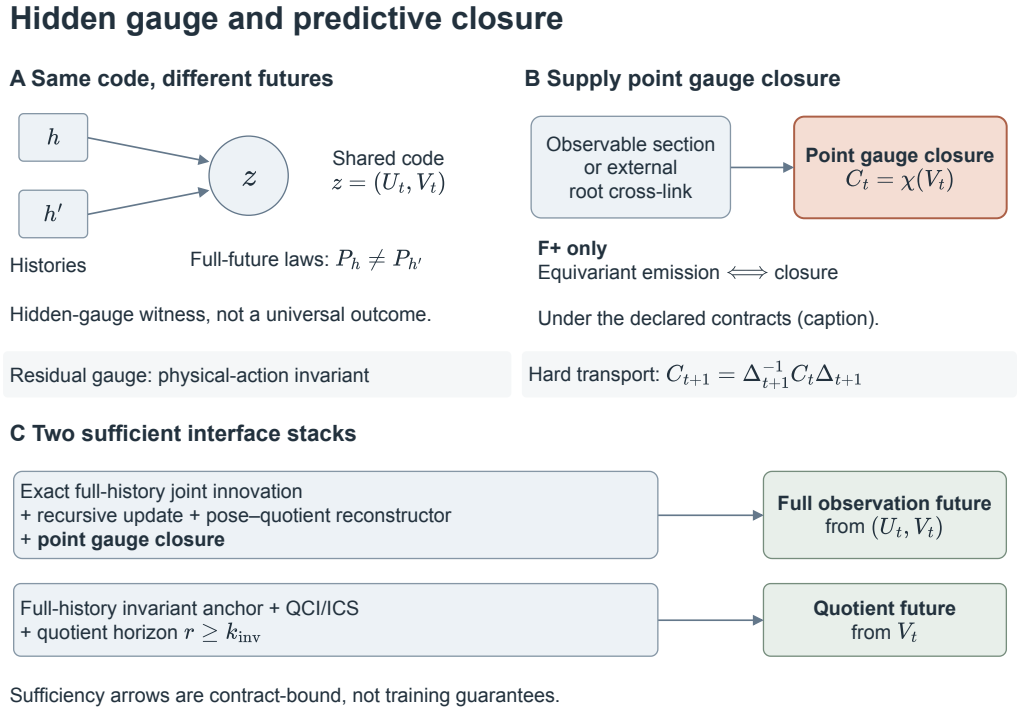


Figure 1: **Hidden gauge and predictive closure.** (A) Equal codes can have distinct future laws; residual non-selection alone need not cause a collision. (B) Sections require the stated Borel and support contracts; cross-links require external calibration per component and a gauge carried in V. Emission equivalence is restricted to F+. (C) Full-future sufficiency uses the exact full-history joint innovation kernel, Borel update, reconstructor and point closure. Quotient-future sufficiency instead uses the full-history anchor, QCI/ICS and $r \geq k_{\mathrm{inv}}$. These are conditional implications, not training guarantees; Table 2 lists their distinct assumptions.

2. We prove an exact emission–gauge equivalence in the declared free deterministic class (Theorem 4.2). The contribution is the specified-innovation boundary, using standard factorization; both emission equivariance and the exact *joint* innovation law are essential.
3. We separate quotient-future from full-future sufficiency (Theorems 5.1 and 5.2) and identify when observable sections or calibrated cross-links supply closure that residual-invariant objectives cannot select (Propositions 6.1 to 6.3).
4. Controlled forecasting and decision studies separate readout calibration, transition routing, and dynamics training. They locate a successful calibrated multi-step recipe and retain adverse comparisons against extra relative supervision and system identification (Section 7).

"Closure completes" means the declared interface stack, not universal identifiability: the joint kernel, recursive update and observation reconstructor remain essential (Figure 1).

2 Setup: Transport, Gauge, and Predictive State

Let G be a locally compact second-countable group acting measurably on standard Borel state and observation spaces X and O. The transition and observation kernels are equivariant, and X/G and O/G are standard Borel. Write

$$H_t = (O_0, \ldots, O_t), \qquad \mathcal{F}_t = \sigma(H_t), \qquad W_t = [O_t] \in O/G.$$

All conditional laws below are fixed Borel versions. All all-time identities hold simultaneously on one countable intersection of full-measure events.

An adapted pose is a Borel $\mathcal{F}_t$-measurable $q_t \in G$; when physical typing is invoked, $q_t(g \cdot h) = g q_t(h)$. Its canonical increment and innovation are

$$\Delta_{t+1} = q_t^{-1} q_{t+1}, \qquad J_{t+1} = (W_{t+1}, \Delta_{t+1}).$$

A supplied pose–quotient observation reconstructor is a Borel map

$$\mathrm{Rec} : (O/G) \times G \to O, \qquad O_t = \mathrm{Rec}(W_t, q_t)$$

almost surely. The supplied reconstructor is a world-side coordinate interface. No latent decoder $D(U_t, V_t) = O_t$ is assumed.

The history encoder is one Borel map

$$F(H_t) = (U_t, V_t) \in G \times Z_{\mathrm{inv}},$$

with physical action $g \cdot (u, v) = (gu, v)$. Its residual right gauge is

$$C_t := q_t^{-1} U_t.$$

Hard right-increment transport $U_{t+1} = U_t \Delta_{t+1}$ gives

$$C_{t+1} = \Delta_{t+1}^{-1} C_t \Delta_{t+1};$$

for a non-abelian group the right gauge conjugates rather than remaining constant.

Definition 2.1 (Point gauge closure and predictive sufficiency)**.** The latent has *point gauge closure*, denoted $\mathrm{GC}_{\mathrm{pt}}$, if one Borel map $\chi : Z_{\mathrm{inv}} \to G$ satisfies $C_t = \chi(V_t)$ almost surely at every governed time. The latent is a predictive state if a Borel kernel R_t satisfies

$$\mathcal{L}(O_{t+1:\infty} \mid H_t) = R_t(U_t, V_t) \quad \text{a.s.}$$

Lemma 2.2 (Known free-orbit rigidity)**.** *For a standard Borel S, the regular left action on $X = G \times S$ and $Z_{\mathrm{orb}} = G$, a Borel map $F_{\mathrm{orb}} : X \to G$ is equivariant if and only if*

$$F_{\mathrm{orb}}(p, s) = p\, c(s)$$

for a Borel $c : S \to G$.

This is the second clause of Marchetti et al. (2023, Proposition 2.1); we use it as a known typing result. Equivariance reduces the orbit code to one functional degree of freedom, but does not identify that degree of freedom.

Quotient anchor. Quotient conditional independence, equivalently invariant-channel sufficiency (QCI/ICS), means

$$\mathcal{L}(W_{t+1:\infty} \mid \mathcal{F}_t) = \mathcal{L}(W_{t+1:\infty} \mid W_{0:t}) \quad \text{a.s.}$$

Fix $E := O/G$ and, once for every $\ell \geq 0$, one fixed full-future conditioning scheme

$$Q_\ell : E^\ell \to \mathcal{P}(E^{\mathbb{N}}),$$

where $Q_\ell(h)$ is the conditional law of the future after a length-ℓ quotient history h, reindexed from zero on the common tail space $E^{\mathbb{N}}$. Let $Q_\ell^{(m)}$ be its first-m marginal and let P_ℓ^W be the law of length-ℓ quotient histories. Define k_{inv} as the least $m \geq 0$ for which there is one rectangular family of per-length null sets $(N_\ell)_{\ell \geq 0}$, with $N_\ell \subseteq E^\ell$ Borel and $P_\ell^W(N_\ell) = 0$, chosen once across all lengths, such that

$$Q_a^{(m)}(h) = Q_b^{(m)}(h') \quad \Longrightarrow \quad Q_a(h) = Q_b(h')$$

for every $a, b \geq 0$, $h \in E^a \setminus N_a$, and $h' \in E^b \setminus N_b$. Here $E^0 = \{\varnothing\}$, $P_0^W = \delta_\varnothing$, $Q_\ell^{(0)}$ is the unique law on E^0, and $\inf \varnothing := \infty$; thus both the empty history and horizon zero are included. Any two such Borel schemes disagree only on a Borel P_ℓ^W-null set at each length. Adjoining those disagreement sets to the rectangular family preserves the displayed predicate in both directions, so the choice of version is immaterial and k_{inv} is version-independent.

This is the cross-age reading of the predictive-sufficiency index of Teoh et al. (2025), with explicit tail reindexing and the declared empty-history, horizon-zero, and infinite-tail conventions. The analogous construction gives the full-process cross-age horizon k. Quotienting may increase or decrease this index. We use k_{inv} only in Theorem 5.1 and make no budget-reduction claim.

Autonomous and teacher-forced objectives. The first no-go concerns autonomous self-prediction,

$$L_{\mathrm{sp}} = \sum_t w_t \, \mathbb{E}\, d(\Phi(F(H_t)), F(H_{t+1}))\,,$$

paired with a proper quotient anchor on V_t. The second no-go and the positive results concern a teacher-forced forward-value objective: a proper anchor predicts J_{t+1} (or a quotient block together with Δ_{t+1}), while nonnegative update losses enforce

$$U_{t+1} = U_t \Delta_{t+1}, \qquad V_{t+1} = \Phi_{\mathrm{inv}}(V_t, J_{t+1}).$$

All temporal weights w_t are positive, all component coefficients are positive, the anchor score is integrable and strictly proper, the update losses vanish exactly on equality, and the model class can attain the componentwise lower bounds simultaneously. The weighted Bayes baseline is summable. Under those assumptions an exact global minimum attains the shared full-history Bayes act and both exact updates at every governed time. Autonomous prediction can instead reward a statistic precisely because it ignores the new innovation.

3 Exact optima need not be predictive

3.1 No-go I: a persistent covariant copy

Theorem 3.1 (Persistent covariant copies)**.** *There is a finite equivariant world with free actions and a nonconstant equivariant encoder that jointly minimizes autonomous self-prediction plus a strictly proper quotient anchor, with zero self-prediction loss and Bayes anchor value, while failing predictive sufficiency.*

Construction sketch. Let $G = C_2$ act on a static state (p, s) by $g(p, s) = (gp, s)$, with independent uniform p and full-support s. Observe $O_t = (Y_t, s)$, where the Y_t are conditionally independent copies of p corrupted with probability $1/4$. The equivariant encoder $F(H_t) = (Y_0, s)$ and identity predictor attain zero self-prediction loss and the Bayes quotient-anchor value. Histories $((+1, s), (+1, s))$ and $((+1, s), (-1, s))$ have the same code but next-Y probability $7/10$ and $1/2$, respectively. Appendix A.1 gives the full finite verification. □

The model is not collapsed: it carries a covariant bit and the quotient state. It fails because the objective rewards a statistic that is easy to propagate and never asks the frozen statistic to absorb later evidence.

3.2 No-go II: canonical innovation leaves a trajectory gauge

Theorem 3.2 (A hidden trajectory gauge survives observation reconstruction)**.** *There is a finite equivariant world and a latent satisfying exactly: adapted pose typing, a supplied pose–quotient observation reconstructor, QCI/ICS with $r = k_{\mathrm{inv}} = 1$, a moving pose, a strictly proper bottlenecked canonical anchor at Bayes risk, hard equivariant transport, and zero invariant-update loss. The joint forward-value lower bound is attained, but two positive-probability histories have the same latent and different full observation futures. Point gauge closure is the failed interface; no latent decoder $D(U_t, V_t) = O_t$ is assumed.*

Construction sketch. Use additive $G = C_2$, observations $O_t = (S_t, P_t)$, initial S_0 uniform on $\{a, b\}$ and independent fair P_0, then $S_{t+1} = c$ and $P_{t+1} = P_t + 1$. Set $C(a) = C(c) = 0$, $C(b) = 1$ and $(U_t, V_t) = (P_t + C(S_0), v_\star)$. The quotient future and canonical increment are deterministic, transport is exact, and $\mathrm{Rec}(w, q) = (w, q)$. Initial states $(a, 0)$ and $(b, 1)$ give the same latent but distinct future poses. The invariant code cannot recover the nonconstant $C(S_0)$. Appendix A.2 verifies every hypothesis, including $k_{\mathrm{inv}} = 1$. □

The reconstructor in Theorem 3.2 reads (W_t, q_t), not the latent. Indeed the colliding histories rule out a deterministic latent decoder. The obstruction is that the residual needed to align latent pose with physical pose is not measurable from the invariant code.

4 THE EXACT FREE DETERMINISTIC BOUNDARY

We next ask whether gauge closure is merely one convenient sufficient interface or is forced by the ability to emit the next observation. The answer is exact on the following declared class.

Definition 4.1 (F+ class and equivariant deterministic emission closure)**.** Restrict to a Borel union $O^\circ \subseteq O$ of free observation orbits. Let $\bar{O}^\circ = O^\circ/G$ be standard Borel and suppose $O_t \in O^\circ$ almost surely. Assume an observation-aligned Borel trivialization

$$\Theta : G \times \bar{O}^\circ \to O^\circ, \qquad \Theta(q, w) = \mathrm{Rec}(w, q),$$

with Borel inverse (τ, π) satisfying $\tau(g \cdot o) = g\tau(o)$ and $\pi(g \cdot o) = \pi(o)$. Hence $q_t = \tau(O_t)$ and $W_t = \pi(O_t)$.

Suppose one common Borel kernel K gives the exact *joint* innovation law

$$\mathcal{L}(J_{t+1} \mid \mathcal{F}_t) = K(V_t; \cdot) \quad \text{a.s. at every governed time.}$$

The latent has *equivariant deterministic emission closure*, denoted $\mathrm{EC}_{\mathrm{eq}}$, if one common Borel map

$$\mathcal{E} : (G \times Z_{\mathrm{inv}}) \times (\bar{O}^\circ \times G) \to O^\circ$$

satisfies pointwise $\mathcal{E}((gu, v), j) = g \cdot \mathcal{E}((u, v), j)$ and satisfies

$$O_{t+1} = \mathcal{E}((U_t, V_t), J_{t+1}) \quad \text{a.s. at every governed time.}$$

These contracts define the F+ class; $\mathrm{EC}_{\mathrm{eq}}$ is the property to be characterized, not an assumption of the class.

Theorem 4.2 (Free equivariant emission–gauge equivalence)**.** *For any member of the F+ class in Definition 4.1,*

$$\boxed{\mathrm{EC}_{\mathrm{eq}} \quad \Longleftrightarrow \quad \mathrm{GC}_{\mathrm{pt}}.}$$

After the adapted variables have been supplied, this equivalence consumes neither encoder equivariance, hard latent transport, nor the teacher-forced invariant update.

Proof idea. Assume $\mathrm{EC}_{\mathrm{eq}}$ and write $j = (w', \delta)$. Define

$$r(u, v, j) := \tau(\mathcal{E}((u, v), j))^{-1}\, u\delta.$$

Equivariance of $\mathcal{E}$ and τ makes r invariant in u, hence $r(u, v, j) = \bar{r}(v, j)$. On a governed transition,

$$\rho(V_t, J_{t+1}) := \Delta_{t+1}\bar{r}(V_t, J_{t+1})\Delta_{t+1}^{-1} = C_t.$$

Push the *joint* kernel forward: $D(v) = \rho(v, \cdot)_{\#}K(v; \cdot) \in \mathcal{P}(G)$. Because C_t is $\mathcal{F}_t$-measurable, conditional expectation on a countable measure-determining algebra yields $D(V_t) = \delta_{C_t}$ almost surely. The Dirac embedding of a standard Borel group has Borel image and Borel inverse, so $C_t = \chi(V_t)$ for one total Borel χ.

Conversely, from $C_t = \chi(V_t)$ define

$$\mathcal{E}_{\mathrm{GC}}((u, v), (w', \delta)) := \mathrm{Rec}\big(w', u\chi(v)^{-1}\delta\big)\,.$$

It is Borel and equivariant, and on governed transitions its pose argument is $U_t C_t^{-1}\Delta_{t+1} = q_{t+1}$, so it emits O_{t+1}. The complete measurable argument is in Appendix B. □

The exact constructions are in Appendix B. Separate innovation marginals are not a technical approximation to the theorem: their coupling is the information that collapses the conditional law of the recovered gauge to a Dirac mass.

Fixed innovation versus auxiliary noise. The proof specializes regular-action factorization and conditional degeneracy, rather than introducing a general stochastic-representation principle (Bloem-Reddy & Teh, 2020). The distinction from auxiliary-noise representations is the *specified coupling* to J_{t+1}, not equality in law versus almost-sure equality. Appendix G gives a four-state process whose latent determines the entire future law and admits an equivariant auxiliary-noise representation, yet no deterministic emitter from that latent and the actual innovation succeeds with probability above $1/2$. Point closure is therefore not necessary for every predictive state.

Table 1: Both contracts in Theorem 4.2 are load-bearing. The counterexamples are finite C_2 processes, exact at every time, and retain a fresh random future.

Removed contract	What still holds	Why point GC fails
Emission equivariance	one shared exact joint kernel; one common deterministic emission	the emitter ignores how the latent pose transforms
Exact joint coupling	equivariant deterministic emission; both innovation marginals are exact	their correlation still reveals the hidden gauge

Table 2: The two positive results are not one theorem. A check mark means that the proof consumes the item.

Interface	Quotient future	Full future
Full-history anchor identity	✓	one-step J kernel
QCI/ICS and $r \geq k_{\mathrm{inv}}$	✓	—
Teacher-forced recursive update	—	✓
Pose–quotient reconstructor	—	✓
Point gauge closure	—	✓
Pose-equivariance identity itself	—	—

5 What closure buys

The exact boundary above concerns one-step equivariant emission. Two longer-horizon conclusions must remain separate.

Interfaces. Fix $r \geq 1$. The invariant anchor identity is the full-history factorization

$$\mathcal{L}((W_{t+1:t+r}, \Delta_{t+1}) \mid \mathcal{F}_t) = A_r(V_t; \cdot) \quad \text{a.s. at every governed time}$$

for one common Borel kernel A_r realized by a V_t-only anchor under an integrable strictly proper score. This is stronger than fitting only the conditional law given V_t. Let the exact teacher-forced updates and the supplied reconstructor identity be as in Section 2.

Theorem 5.1 (Anchor-only quotient sufficiency)**.** *Under the displayed full-history anchor identity, if QCI/ICS holds and $r \geq k_{\mathrm{inv}}$, then the invariant code V_t alone determines the quotient-future law. The result uses neither the recursive update, the pose–quotient reconstructor, nor gauge closure.*

Theorem 5.2 (Gauge-closed full-future sufficiency)**.** *Suppose one common Borel kernel satisfies*

$$\mathcal{L}(J_{t+1} \mid \mathcal{F}_t) = K(V_t; \cdot) \quad \textit{a.s. at every governed time}.$$

If the teacher-forced latent update is deterministic and Borel, the supplied reconstructor identity holds, and $\mathrm{GC}_{\mathrm{pt}}$ *holds, then a Borel kernel R_t satisfies*

$$\mathcal{L}(O_{t+1:\infty} \mid H_t) = R_t(U_t, V_t) \quad \textit{a.s.}$$

This result uses neither QCI/ICS nor $r \geq k_{\mathrm{inv}}$.

Proof roadmap. For Theorem 5.1, equality of V gives equality of the anchored next-r law. QCI/ICS identifies that law with the quotient-history law, and $r \geq k_{\mathrm{inv}}$ upgrades equality of the block to equality of the reindexed quotient future. Borel factorization through the standard-Borel space of next-r laws, followed by composition with the anchor, yields the desired kernel without an additional standard-Borel assumption on Z_{inv}.

For Theorem 5.2, the one-step kernel and exact updates give a Borel transition on (U, V, J). At the next time point, closure recovers $q_{t+1} = U_{t+1}\chi(V_{t+1})^{-1}$; the supplied reconstructor then emits O_{t+1}. Finite-cylinder induction along the full-history filtration and Ionescu–Tulcea iteration produce the complete conditional future law. Full proofs, including the measurable factorization step, are in Appendix C. □

Suppose the V_t-only anchor realizes the full-history conditional Bayes act together with both zero update losses, and the weighted Bayes baseline is summable. Positive objective and temporal weights then force every exact global minimum to attain that act and both updates. To apply Theorem 5.2 to these minima, the model class must also supply closure; Theorem 3.2 shows why this restriction is substantive.

Proposition 5.3 (Non-abelian gauge law)**.** *With $C_t = q_t^{-1}U_t$, hard right-increment transport gives $C_{t+1} = \Delta_{t+1}^{-1} C_t \Delta_{t+1}$; equivalently the left factor $U_t q_t^{-1}$ is trajectory invariant. Thus the right gauge generally conjugates rather than remaining constant.*

6 What can produce closure

The next results delimit what the objective can select and give two sources of closure: observable sections and external root cross-links.

Proposition 6.1 (The objective cannot choose the gauge)**.** *Let $\mathfrak{M}$ be a declared model class. Let $E_{\mathrm{res}} \subseteq \mathfrak{M} \times \mathfrak{M}$ be a symmetric edge relation encoding residual transformations admitted by the comparison and preserving predictive adequacy, and let $\sim_{\mathrm{res}}$ be its finite-chain closure. If every objective component is invariant on every E_{res} edge, the objective is constant on each $\sim_{\mathrm{res}}$ class and the exact-global-minimizer set is saturated under that relation. Hence, if a closed and a non-closed model share a class, the objective cannot prefer one to the other.*

Physical group action and residual-gauge action are different symmetries. A term may preserve physical equivariance while deliberately breaking residual-gauge symmetry.

Proposition 6.2 (An observable section produces closure)**.** *Let $O^\circ \subseteq O$ be a Borel union of free observation orbits, let $\bar{O}^\circ = O^\circ/G$ be standard Borel, and suppose its quotient map admits a Borel section. Suppose $O_t \in O^\circ$ almost surely at every governed time. Suppose the pose–quotient reconstructor is equivariant in its pose argument and suppose its identity-point map $s_{\mathrm{OR}}(w) = \mathrm{Rec}(w, e)$ is itself a Borel section on $\bar{O}^\circ$. Its transporter is then Borel. The architectural choice*

$$U_t = \tau_{\mathrm{OR}}(O_t)$$

gives $C_t = e$ and hence point gauge closure. More generally, a different Borel section also yields point closure when the encoder uses its transporter, hard transport holds, and V_t supplies the shared exact joint innovation kernel.

Proposition 6.3 (Synchronization propagates but does not create)**.** *Flat relative-frame equations determine frames only up to one global factor per connected component. A spanning tree already propagates the relative frames; consistent loops add no absolute information. One physical–latent root cross-link per component removes the remaining factor. Under an architecture that carries the resulting gauge measurably in V, this gives point closure, but the cross-link is external calibration rather than information created by synchronization.*

The synchronization count is classical (Gao et al., 2021; Singer & Wu, 2012). In the common-left-factor convention used here, if f and $\tilde{f}$ satisfy the same relative edge equations, then $\tilde{f}_i f_i^{-1} = \tilde{f}_j f_j^{-1}$ on every edge, so the left ratio is constant on each connected component. The complete proofs of Propositions 6.1 to 6.3 are in Appendix D.

7 From exact obstructions to learned prediction

The experiments follow three questions: do the exact obstructions occur, where does supplied alignment help a predictor, and does that benefit survive stronger baselines? The theory specifies interfaces; these studies test their practical motivation, not attainment of the exact hypotheses. Fixed protocols retain both favorable and adverse outcomes.

Exact landscape and mechanism controls. A pre-specified population-loss audit finds four zero-loss persistent copies among 16 local equivariant triples on one quotient fibre and one transition, each with next-law TV gap $1/5$. Smooth controls demonstrate quantized collisions and direct calibration under an algebraically flat baseline, with all fixed-seed outcomes and misses in Appendix E and fig. 2. They support mechanisms rather than generic convergence claims.

Table 3: Calibrated routing with multi-step training beats persistence; native multi-step training does not on average. Mean H8 NMSE over the same 20 frozen ROOT-encoder blocks; lower is better. Four new heads per block share initial weights, raw sequences, supplied one-record four-harmonic calibration and 1,000 updates. TF and AR denote teacher-forced and multi-step autoregressive *training*; all forecasts are autonomous. Matched updates are not equal FLOPs or training observation access. Parent encoders were already root-supervised.

Evaluation interface	TF training	AR training
Native carrier	1.221177	0.707873
Calibrated routing	0.964155	**0.127500**
Association placebo	1.320647	0.901203
Routed, permuted actions	1.358783	1.005090
Matched persistence (same forecasts)	0.675110	0.675110

Alignment at the readout versus inside the transition. In a noisy $\mathrm{SO}(2)$ point-cloud world, root supervision (ROOT) beats relative-only supervision (REL) on new trajectories of training objects in all 20 paired blocks; shuffled root association is a control. ROOT has extra physical information, not an equal-supervision advantage. Freezing all 60 models, a new-object readout-only calibration test still loses to matched persistence in $59/60$ arm/block comparisons (Appendix F). Routing the carrier through the supplied calibration frame before transition improves the frozen ROOT models in fresh-object and bounded shape/noise tests (Appendix H.1). However, a subsequent one-step physical-objective retraining recipe has H8 NMSE 1.156480 against persistence 0.683852, winning only $2/20$ blocks. Readout repair, routing, and transition training are distinct interventions.

A controlled multi-step intervention. We retain all 20 ROOT encoders and train four matched heads on 320 new objects, testing once on 320 disjoint objects and 15,360 H8 trajectories. Calibrated routing with AR training attains mean H8 NMSE 0.127500, versus 0.675110 for persistence and 0.707873 for native AR. It beats routed TF, native AR, persistence, association-placebo AR and action-permuted routed AR in all 20 blocks (Table 3 and appendix H.2). The routing-by-training interaction favors AR in $18/20$ blocks, not every block; native AR remains worse than persistence on average. This supports the joint recipe under supplied physical calibration and known association. A changed descriptor convention prevents attributing the earlier one-step failure to horizon alone.

Root supervision is not uniformly beneficial. A separate 20-block experiment compares REL, ROOT and extra-pair relative supervision (EXTRA_REL), with 60 parents and 120 downstream heads. The primary calibrated-AR contrast is adverse: ROOT minus EXTRA_REL H8 NMSE is $+0.140583$ (marginal approximate 95% t interval $[0.094076, 0.187091]$), with ROOT worse in all 20 blocks. Under the original training recipe, ROOT instead wins all 20 blocks on both same-object and held-out panels; its held-out NMSE is still 1.888265. A physically supervised six-feature system-identification (SID) baseline initialized from observed moments attains 0.000378, below all three learned AR heads. Regimes differ in information and objective; these comparisons do not isolate a unique failure mechanism (Appendix H.3).

Decision value on a named control task. On sensor-mounting Pendulum with a fixed H16, 27-action-sequence library, 12 independent training/test blocks compare current versus history inputs, with versus without calibration, at 0° and 5° angular noise. History lowers cost, yet calibration retains value under history for both SID and learned dynamics: all eight primary Bonferroni-adjusted 95% t intervals are positive. Distribution-free Hoeffding intervals all include zero. Learned dynamics has higher mean cost than SID in all eight cells (differences 0.150–0.251); both use physical labels, and SID has a stronger dynamics prior. Numerical reference gaps and the lack of full source certification remain explicit (Appendix H.4).

A bounded consequence of missing information. For a continuous controlled-circle law with Haar initial angle, a fresh rotation-invariant object prior making the initial cloud independent of that angle, and a known fixed action sequence, the unrestricted four-harmonic Bayes risk has an exact conditional-mean formula. Certified integration on a fixed 16-sequence library gives H8 risk

enclosures within $[0.73704901, 0.98708800]$. This is an exact-real comparison law, not a statistical interval or a floor for the calibrated learned studies (Appendix H.5).

8 Related work

Predictive states and latent prediction. Predictive-state and observable-operator models represent future laws (Jaeger, 2000; Littman et al., 2001; Singh et al., 2004). Latent prediction instead learns an embedding, whose identification depends on the objective and model (Guo et al., 2020; Ni et al., 2024; Bagatella et al., 2026; Teoh et al., 2025; Klindt et al., 2026; Zhang et al., 2026a).

Shortcut solutions and identification. Sobal et al. (2022) give a persistent Gaussian noise copy with an identity predictor for VICReg and infinite-negative InfoNCE. Our finite equivariant witness additionally attains a proper quotient anchor and gives an exact predictive-law collision. Shahverdi et al. (2026) prove layerwise equivariance from weak identifiability and end-to-end equivariance under their model/adjunction assumptions; latent actions are existential. Mikulasch & Zenke (2026) establish affine encoder identification with nondegenerate Gaussian predictive laws, injective generation, invertibility on the data manifold, perfect distribution matching, and full-row-rank predictor Jacobians; Remark 1 qualifies extensions to internal variables. We instead test future-law sufficiency of a specified history code and characterize its gauge under the declared emitter contract.

Equivariance and identifiability. Structural identification uses observed group elements, paired action data, or task structure (Marchetti et al., 2023; Park et al., 2022; Ghaemi et al., 2025; Schmidt et al., 2025; Liu et al., 2023). We use the free-orbit rigidity of Marchetti et al. (2023) as known. Our additional question is when equivariant emission forces the residual gauge into the invariant code.

Canonicalization and symmetry breaking. Canonicalization and probabilistic symmetry breaking select representatives or representative distributions (Kaba et al., 2023; Ma et al., 2024; Lawrence et al., 2025; Urbano et al., 2026). These are producers, not universal necessity statements. Proposition 6.2 gives the measurable temporal section contract; Proposition 6.3 distinguishes relative-frame propagation from external absolute calibration (Gao et al., 2021; Singer & Wu, 2012).

Structured world models in practice. Multi-step dynamics training is established practice, including latent overshooting (Hafner et al., 2019); our intervention diagnoses its interaction with a specified calibration interface, not a new training principle. Recent systems separate ego-motion, world effects, or group flows (Wang et al., 2026a; Zhang et al., 2026b; Lillemark et al., 2026; Wang et al., 2026b). These motivate, but do not settle, our interface question. Our necessity result concerns the F+ emitter contract, not arbitrary learned predictors.

9 Limitations and conclusion

The theory concerns exact global values under declared Borel interfaces, not high-dimensional optimization or general stationary points. Theorem 4.2 requires a free observation-aligned trivialization, equivariant deterministic emission, and an exact joint innovation kernel. It characterizes emission driven by that innovation, not arbitrary full-future-law representations; point closure is not asserted to be the weakest resolver. Non-free actions are outside scope.

The empirical studies use synthetic generators, privileged physical labels, and supplied association or sensor models. Multi-step routing succeeds in its declared setting, but adverse ROOT and SID comparisons preclude generic superiority. Neither operational posterior certification nor fresh population confirmation follows from the conditional numerical references; some bounds remain unresolved (Appendix H). Real-world transfer and attainment of the exact interfaces are unestablished. Under their fixed budgets, neither quantized-law A1 implementation attained the prespecified excess-enclosure width of $1/500$ on any positive-evidence case with $H > 0$ (Appendix H.7). Our exact results identify the missing gauge interface and separate the assumptions needed for quotient and full-future prediction. The experiments show where supplied calibration helps and where stronger baselines still win. Equivariance supplies transport without selecting a predictive representative; information placement and learning accurate dynamics remain separate tasks.

REFERENCES

Marco Bagatella, Matteo Pirotta, Ahmed Touati, Alessandro Lazaric, and Andrea Tirinzoni. TD-JEPA: Latent-predictive representations for zero-shot reinforcement learning. In *International Conference on Learning Representations (ICLR)*, 2026. Oral. arXiv:2510.00739.

Benjamin Bloem-Reddy and Yee Whye Teh. Probabilistic symmetries and invariant neural networks. *Journal of Machine Learning Research*, 21(90):1–61, 2020. URL https://jmlr.org/papers/v21/19-322.html.

Tingran Gao, Jacek Brodzki, and Sayan Mukherjee. The geometry of synchronization problems and learning group actions. *Discrete & Computational Geometry*, 65(1):150–211, 2021. doi: 10.1007/s00454-019-00100-2. First published online 28 May 2019; preprint arXiv:1610.09051.

Hafez Ghaemi, Eilif B. Muller, and Shahab Bakhtiari. seq-JEPA: Autoregressive predictive learning of invariant-equivariant world models. In *Advances in Neural Information Processing Systems*, volume 38, 2025. URL https://proceedings.neurips.cc/paper_files/paper/2025/hash/2f63d2963526bdd9ff1b8bcc2dc9905a-Abstract-Conference.html.

Daniel Guo, Bernardo Avila Pires, Bilal Piot, Jean-Bastien Grill, Florent Altché, Rémi Munos, and Mohammad Gheshlaghi Azar. Bootstrap latent-predictive representations for multitask reinforcement learning. In *Proceedings of the 37th International Conference on Machine Learning (ICML)*, volume 119 of *Proceedings of Machine Learning Research*, 2020.

Danijar Hafner, Timothy Lillicrap, Ian Fischer, Ruben Villegas, David Ha, Honglak Lee, and James Davidson. Learning latent dynamics for planning from pixels. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 2555–2565. PMLR, 2019. URL https://proceedings.mlr.press/v97/hafner19a.html.

Herbert Jaeger. Observable operator models for discrete stochastic time series. *Neural Computation*, 12(6):1371–1398, 2000. doi: 10.1162/089976600300015411.

Sékou-Oumar Kaba, Arnab Kumar Mondal, Yan Zhang, Yoshua Bengio, and Siamak Ravanbakhsh. Equivariance with learned canonicalization functions. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 15546–15566. PMLR, 2023. URL https://proceedings.mlr.press/v202/kaba23a.html.

David Klindt, Yann LeCun, and Randall Balestriero. When does LeJEPA learn a world model? *arXiv preprint arXiv:2605.26379*, 2026.

Hannah Lawrence, Vasco Portilheiro, Yan Zhang, and Sékou-Oumar Kaba. Improving equivariant networks with probabilistic symmetry breaking. In *International Conference on Learning Representations (ICLR)*, 2025. arXiv:2503.21985.

Hansen Jin Lillemark, Benhao Huang, Fangneng Zhan, Yilun Du, and Thomas Anderson Keller. Flow equivariant world models: Memory for partially observed dynamic environments. *arXiv preprint arXiv:2601.01075*, 2026.

Michael L. Littman, Richard S. Sutton, and Satinder Singh. Predictive representations of state. In *Advances in Neural Information Processing Systems 14*, pp. 1555–1561, 2001.

Yu-Ren Liu, Biwei Huang, Zhengmao Zhu, Honglong Tian, Mingming Gong, Yang Yu, and Kun Zhang. Learning world models with identifiable factorization. In *Advances in Neural Information Processing Systems*, volume 36, 2023. URL https://proceedings.neurips.cc/paper_files/paper/2023/hash/65496a4902252d301cdf219339bfbf9e-Abstract-Conference.html.

George Ma, Yifei Wang, Derek Lim, Stefanie Jegelka, and Yisen Wang. A canonicalization perspective on invariant and equivariant learning. In *Advances in Neural Information Processing Systems*, 2024. arXiv:2405.18378.

Giovanni Luca Marchetti, Gustaf Tegnér, Anastasiia Varava, and Danica Kragic. Equivariant representation learning via class-pose decomposition. In Francisco Ruiz, Jennifer Dy, and Jan-Willem van de Meent (eds.), *Proceedings of the 26th International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 206 of *Proceedings of Machine Learning Research*, pp. 4745–4756. PMLR, 25–27 Apr 2023. URL https://proceedings.mlr.press/v206/marchetti23b.html.

Fabian A Mikulasch and Friedemann Zenke. Understanding self-supervised learning via latent distribution matching, 2026. URL https://arxiv.org/abs/2605.03517v4. Version 4.

Tianwei Ni, Benjamin Eysenbach, Erfan Seyedsalehi, Michel Ma, Clement Gehring, Aditya Mahajan, and Pierre-Luc Bacon. Bridging state and history representations: Understanding self-predictive RL. In *International Conference on Learning Representations (ICLR)*, 2024. arXiv:2401.08898.

Jung Yeon Park, Ondrej Biza, Linfeng Zhao, Jan-Willem Van De Meent, and Robin Walters. Learning symmetric embeddings for equivariant world models. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 17372–17389. PMLR, 2022. URL https://proceedings.mlr.press/v162/park22a.html.

Tobias Schmidt, Steffen Schneider, and Matthias Bethge. Equivariance by contrast: Identifiable equivariant embeddings from unlabeled finite group actions. In *Advances in Neural Information Processing Systems*, 2025. arXiv:2510.21706.

Vahid Shahverdi, Giovanni Luca Marchetti, Georg Bökman, and Kathlén Kohn. Identifiable equivariant networks are layerwise equivariant, 2026. URL https://arxiv.org/abs/2601.21645v2. Version 2.

Amit Singer and Hau-Tieng Wu. Vector diffusion maps and the connection Laplacian. *Communications on Pure and Applied Mathematics*, 65(8):1067–1144, 2012. doi: 10.1002/cpa.21395.

Satinder Singh, Michael R. James, and Matthew R. Rudary. Predictive state representations: A new theory for modeling dynamical systems. In *Uncertainty in Artificial Intelligence (UAI)*, pp. 512–519, 2004.

Vlad Sobal, Jyothir S V, Siddhartha Jalagam, Nicolas Carion, Kyunghyun Cho, and Yann LeCun. Joint embedding predictive architectures focus on slow features, 2022. URL https://arxiv.org/abs/2211.10831v1. Version 1.

Jayden Teoh, Manan Tomar, Kwangjun Ahn, Edward S. Hu, Tim Pearce, Pratyusha Sharma, Akshay Krishnamurthy, Riashat Islam, Alex Lamb, and John Langford. Next-latent prediction transformers learn compact world models. *arXiv preprint arXiv:2511.05963*, 2025.

Alonso Urbano, David W. Romero, Max Zimmer, and Sebastian Pokutta. RECON: Robust symmetry discovery via explicit canonical orientation normalization. In *International Conference on Learning Representations (ICLR)*, 2026. arXiv:2505.13289.

Zhidong Wang, Jingsong Liang, Zirui Li, Zhan Chen, Han Yu, and Chen Lv. Ego-dynamics-augmented world model for autonomous driving with zero-shot cross-chassis adaptation. *arXiv preprint arXiv:2607.13410*, 2026a.

Zijie Wang, Wei Zhang, Weiming Zhang, Fanqi Zhang, Xiao Tan, Yipeng Qin, and Guanbin Li. World models as group actions. *arXiv preprint arXiv:2605.24578*, 2026b.

Xiangteng Zhang, Yang Guan, Bo Zhang, Hongyang Li, Ya-Qin Zhang, and Shengbo Eben Li. On the identifiability of controlled world models. *arXiv preprint arXiv:2607.22430*, 2026a.

Yi-Ge Zhang, Tianqi Du, Qi Zhang, and Yisen Wang. DWM: Separating world effects from actions in latent world models. *arXiv preprint arXiv:2607.18715*, 2026b.

Part I: Proofs, audits, and reproducibility

A Exact finite no-go audits

A.1 Complete verification of No-go I

We verify the finite construction used in Theorem 3.1. Write $C_2 = \{+1, -1\}$, let $S = \{0, 1\}$, and take $X = C_2 \times S$ with the free action $g \cdot (p, s) = (gp, s)$. The state is static, with initial law $\mathrm{Unif}(C_2) \otimes \nu$, where ν has full support on S. Conditionally on (p, s), the observations $O_t = (Y_t, s)$ are independent and satisfy

$$\mathbb{P}(Y_t = p \mid p, s) = \frac{3}{4}, \qquad \mathbb{P}(Y_t = -p \mid p, s) = \frac{1}{4}.$$

The observation channel is equivariant under $g \cdot (y, s) = (gy, s)$, and all spaces are finite standard Borel spaces.

Define the encoder and autonomous predictor by

$$F^\flat((y_0, s), \ldots, (y_t, s)) = (y_0, s), \qquad \Phi^\flat = \mathrm{id}\,.$$

The orbit component carries the regular free C_2 action, and $F^\flat(g \cdot H_t) = g \cdot F^\flat(H_t)$. Since appending a new observation leaves the first observation unchanged,

$$\Phi^\flat(F^\flat(H_t)) = F^\flat(H_{t+1})$$

pointwise. Hence the nonnegative autonomous self-prediction loss is zero. The quotient observation is exactly s, so an anchor reading the invariant component predicts every future quotient block at its conditional Bayes law. For any integrable strictly proper score, this attains the model-independent conditional Bayes value. The two terms are attained simultaneously, so the displayed triple is a joint global-value minimizer.

It remains to verify the predictive collision. Under the uniform pose prior, two repeated positive observations give posterior odds $9 : 1$ for $p = +1$:

$$\mathbb{P}(p = +1 \mid Y_0 = Y_1 = +1) = \frac{9}{10}.$$

Therefore

$$\mathbb{P}(Y_2 = +1 \mid Y_0 = Y_1 = +1) = \frac{9}{10}\frac{3}{4} + \frac{1}{10}\frac{1}{4} = \frac{7}{10}.$$

After one positive and one negative observation, the posterior on p is uniform, and thus

$$\mathbb{P}(Y_2 = +1 \mid Y_0 = +1, Y_1 = -1) = \frac{1}{2}.$$

For either value of s, both histories have positive probability and the encoder maps them to the same code $(+1, s)$. Their next-observation laws differ, so the code is not even one-step sufficient and hence is not sufficient for the complete observation future.

A.2 Complete verification of No-go II

We audit every premise of Theorem 3.2. All operations in $G = C_2 = \{0, 1\}$ are additive. Let $S = \{a, b, c\}$ and $X = O = S \times C_2$, with free action $g \cdot (s, p) = (s, g + p)$ and the identity observation channel. Initially, S_0 is uniform on $\{a, b\}$ and P_0 is uniform on C_2, independently. The deterministic transition is

$$S_{t+1} = c, \qquad P_{t+1} = P_t + 1.$$

Set

$$W_t = S_t, \qquad q_t(H_t) = P_t, \qquad C(a) = 0, \qquad C(b) = 1,$$

extend C arbitrarily off the governed support, and define

$$U_t = P_t + C(S_0), \qquad V_t = v_\star.$$

1. **Regularity and equivariance.** Every state, observation, quotient, and latent space is finite discrete. The world transition and identity observation channel are equivariant, the action on $X = O$ is free, and the quotient is S with an evident Borel section.
2. **Adapted pose and encoder equivariance.** The pose $q_t = P_t$ is measurable from the current observation and satisfies $q_t(g \cdot H_t) = g + q_t(H_t)$. Since S_0 is invariant under the physical action,
$$U_t(g \cdot H_t) = g + P_t + C(S_0) = g + U_t(H_t),$$
while V_t is invariant.
3. **Canonical innovation and hard transport.** The nontrivial increment is $\Delta_{t+1} = 1$ at every time, and
$$U_{t+1} = P_t + 1 + C(S_0) = U_t + \Delta_{t+1}.$$
4. **Observation reconstruction.** The Borel map $\mathrm{Rec}(w, q) = (w, q)$ gives $O_t = \mathrm{Rec}(W_t, q_t)$ pointwise. This is the declared pose–quotient reconstruction interface; the theorem does not assume a separate latent decoder $D(U_t, V_t) = O_t$.
5. **Invariant update.** The constant map $\Phi_{\mathrm{inv}}(v_\star, j) = v_\star$ gives the invariant update exactly and therefore has zero update loss.
6. **Canonical anchor and exact joint kernel.** At each governed time, $J_{t+1} = (W_{t+1}, \Delta_{t+1}) = (c, 1)$ deterministically. Hence the one-point invariant code supports the shared exact joint kernel
$$K(v_\star) = \delta_{(c,1)}.$$
A finite-outcome Brier score is strictly proper, and this point mass is its unique conditional Bayes act.
7. **QCI/ICS and the anchor horizon.** The quotient process has $W_0 \in \{a, b\}$ and $W_t = c$ for $t \geq 1$. Its future is conditionally independent of the pose history given the quotient history. Under the declared cross-age convention including the empty history, the empty history has a nondegenerate next quotient observation, while every nonempty history has next observation c and the same constant future. Thus next-one equality determines the complete quotient future, while horizon zero fails by comparing an empty and a nonempty history. Therefore $k_{\mathrm{inv}} = 1$, and the construction takes $r = 1$.
8. **Joint lower bound.** The same model attains the anchor Bayes value and makes both nonnegative update losses zero. With strictly positive coefficients, this is the joint model-independent lower bound of the forward-value objective.
9. **The unique failed interface.** The right gauge is
$$q_t^{-1} U_t = C(S_0).$$
It takes both values with positive probability while $V_t = v_\star$ is constant. No Borel map χ can satisfy $q_t^{-1} U_t = \chi(V_t)$ on both events. Thus point gauge closure is the failed interface.

Finally, at time zero consider the two histories
$$h = ((a, 0)), \qquad h' = ((b, 1)).$$
Each has probability $1/4$, and both have latent code $(0, v_\star)$. Their next-observation laws are nevertheless
$$\mathcal{L}(O_1 \mid H_0 = h) = \delta_{(c,1)}, \qquad \mathcal{L}(O_1 \mid H_0 = h') = \delta_{(c,0)}.$$
The common latent value therefore cannot index both conditional laws. The latent is not one-step sufficient and hence is not full-future sufficient.

B Proof and sharpness of the exact free boundary

B.1 Full proof of the free emission–gauge equivalence

Proof of Theorem 4.2. Take one common equivariant deterministic emission witness $\mathcal{E}$. For $j = (w', \delta)$ define
$$r(u, v, j) := \tau(\mathcal{E}((u, v), j))^{-1}\, u\delta.$$

Emission equivariance and transporter equivariance imply, pointwise,

$$\begin{aligned} r(gu, v, j) &= \left[g\tau(\mathcal{E}((u, v), j))\right]^{-1} gu\delta \\ &= r(u, v, j). \end{aligned}$$

Consequently $r(u, v, j) = \bar{r}(v, j)$ for the Borel map $\bar{r}(v, j) := r(e, v, j)$. On a governed transition, emission closure gives

$$\begin{aligned} \bar{r}(V_t, J_{t+1}) &= q_{t+1}^{-1} U_t \Delta_{t+1} \\ &= (q_t \Delta_{t+1})^{-1} (q_t C_t) \Delta_{t+1} \\ &= \Delta_{t+1}^{-1} C_t \Delta_{t+1}. \end{aligned}$$

Thus the Borel statistic

$$\rho(v, (w', \delta)) := \delta\, \bar{r}(v, (w', \delta))\, \delta^{-1}$$

satisfies $\rho(V_t, J_{t+1}) = C_t$ almost surely.

Push the exact joint innovation kernel forward:

$$D(v) := \rho(v, \cdot)_{\#} K(v; \cdot) \in \mathcal{P}(G).$$

For every Borel $B \subseteq G$,

$$D(v)(B) = \int \mathbf{1}_B(\rho(v, j)) K(v; dj)$$

is Borel in v. Choose a countable measure-determining algebra $\mathcal{A} \subseteq \mathcal{B}(G)$. For every $A \in \mathcal{A}$, the exact conditional-law identity gives

$$\begin{aligned} D(V_t)(A) &= \mathbb{E}[\mathbf{1}_A(\rho(V_t, J_{t+1})) \mid \mathcal{F}_t] \\ &= \mathbf{1}_A(C_t) \qquad \text{a.s.} \end{aligned}$$

The second equality uses both $\rho(V_t, J_{t+1}) = C_t$ and $\mathcal{F}_t$-measurability of C_t. Intersecting the full-measure events over the countable algebra and the countable governed time set yields

$$D(V_t) = \delta_{C_t} \qquad \text{a.s. at every governed time.}$$

The Dirac embedding

$$d_G : G \to \mathcal{P}(G), \qquad d_G(c) = \delta_c,$$

is Borel and injective. Lusin–Souslin makes its image Borel and its inverse Borel on that image. Extend the inverse by e off the image:

$$\chi(v) := \begin{cases} d_G^{-1}(D(v)), & D(v) \in d_G(G), \\ e, & D(v) \notin d_G(G). \end{cases}$$

Then χ is one total Borel map and $C_t = \chi(V_t)$ almost surely at every governed time. This proves the forward implication.

Conversely, let χ witness point gauge closure and define

$$\mathcal{E}_{\mathrm{GC}}((u, v), (w', \delta)) := \mathrm{Rec}\big(w', u\chi(v)^{-1}\delta\big)\,.$$

This map is Borel, and equivariance of Rec gives

$$\mathcal{E}_{\mathrm{GC}}((gu, v), j) = g \cdot \mathcal{E}_{\mathrm{GC}}((u, v), j).$$

On a governed transition,

$$U_t \chi(V_t)^{-1} \Delta_{t+1} = U_t C_t^{-1} \Delta_{t+1} = q_t \Delta_{t+1} = q_{t+1}.$$

Since the first coordinate of J_{t+1} is W_{t+1},

$$\mathcal{E}_{\mathrm{GC}}((U_t, V_t), J_{t+1}) = \mathrm{Rec}(W_{t+1}, q_{t+1}) = O_{t+1} \quad \text{a.s.}$$

This proves the reverse implication. □

B.2 Emission equivariance is load-bearing

Example B.1. Write $G = \bar{O}^\circ = \mathbb{Z}/2\mathbb{Z}$ additively and $O^\circ = G \times \bar{O}^\circ$, with $g \cdot (q, w) = (g + q, w)$. On a product space carrying independent fair bits $C, R_1, R_2, \ldots$, set

$$q_t = 0, \qquad W_0 = C, \qquad W_t = R_t \quad (t \geq 1),$$

let $O_t = (q_t, W_t)$ and $\mathcal{F}_t = \sigma(C, R_1, \ldots, R_t)$, and put

$$U_t = C, \qquad V_t = \star, \qquad \Delta_{t+1} = 0.$$

The shared exact joint kernel is

$$K(\star) = \tfrac{1}{2}\delta_{(0,0)} + \tfrac{1}{2}\delta_{(1,0)}.$$

The common deterministic emission

$$\mathcal{E}_{\mathrm{bad}}((u, \star), (w', \delta)) = (0, w')$$

reconstructs O_{t+1} at every time, while $C_t = q_t^{-1} U_t = C$ is nonconstant and cannot be a function of $V_t = \star$. The map $\mathcal{E}_{\mathrm{bad}}$ fails emission equivariance at the nonidentity group element. Thus deterministic emission without the equivariance clause does not imply point gauge closure. The fresh bit R_{t+1} keeps the future nontrivial at every time.

B.3 Separate innovation marginals do not suffice

Example B.2. On a product space carrying independent fair bits $Q, C, R_1, R_2, \ldots$, retain the same group, observation space, and action, and set

$$q_0 = Q, \qquad W_0 = C,$$

then recurse by

$$W_{t+1} = R_{t+1}, \qquad \Delta_{t+1} = R_{t+1} + C, \qquad q_{t+1} = q_t + \Delta_{t+1}.$$

Let $O_t = (q_t, W_t)$, $\mathcal{F}_t = \sigma(O_0, \ldots, O_t)$, and put

$$U_t = q_t + C, \qquad V_t = \star.$$

The common map

$$\mathcal{E}((u, \star), (w', \delta)) = (u + w', w')$$

is equivariant and emits the correct next observation because $U_t + W_{t+1} = q_t + C + R_{t+1} = q_{t+1}$. Both coordinate marginals of $J_{t+1} = (W_{t+1}, \Delta_{t+1})$ are fair given either value of C, but their joint laws are

$$C = 0: \quad \tfrac{1}{2}\delta_{(0,0)} + \tfrac{1}{2}\delta_{(1,1)},$$
$$C = 1: \quad \tfrac{1}{2}\delta_{(0,1)} + \tfrac{1}{2}\delta_{(1,0)}.$$

Hence there is no shared exact joint kernel $K(\star)$, while $C_t = q_t^{-1} U_t = C$ remains nonconstant and point gauge closure fails. The coupling, not merely the two marginals, is load-bearing.

C Proofs of the positive results

C.1 A measurable factorization lemma

Lemma C.1 (Borel factorization through a standard-Borel code)**.** *Let X, Y, Z be standard Borel spaces with $Z \neq \varnothing$, and let $f : X \to Y$ and $g : X \to Z$ be Borel. If*

$$f(x) = f(x') \quad \Longrightarrow \quad g(x) = g(x'),$$

then there is a total Borel map $\psi : Y \to Z$ such that $g = \psi \circ f$.

Proof. Choose a Borel embedding $\iota : Z \to \{0,1\}^{\mathbb{N}}$ and write $B_n = \{z : \iota(z)_n = 1\}$. For every n, the sets

$$A_n = f(g^{-1}(B_n)), \qquad A'_n = f(g^{-1}(Z \setminus B_n))$$

are disjoint analytic subsets of Y. Lusin separation gives a Borel $D_n \subseteq Y$ with $A_n \subseteq D_n$ and $A'_n \subseteq Y \setminus D_n$. The Borel map

$$u(y) = (\mathbf{1}_{D_n}(y))_{n \geq 0}$$

satisfies $u(f(x)) = \iota(g(x))$. By Lusin–Souslin, $\iota(Z)$ is Borel and ι^{-1} is Borel on that image. Extending the inverse by a fixed $z_0 \in Z$ off the image gives the required total Borel map ψ. □

C.2 Proof of anchor-only quotient sufficiency

Proof of Theorem 5.1. Put $\bar{O} = O/G$ and

$$\mathcal{H}^* = \bigsqcup_{s\geq 0} O^{s+1}.$$

The countable disjoint union $\mathcal{H}^*$ is standard Borel. By taking countable unions of the exceptional sets in the anchor identity, QCI/ICS, and the selected quotient-future conditioning scheme, together with the pullbacks of the fixed rectangular null family defining k_{inv}, work on a Borel governed set $\mathcal{H}_0 \subseteq \mathcal{H}^*$ whose intersection with each history-length component is conull.

For a governed history $h \in O^{s+1}$, let

$$\overline{Q}(h) \in \mathcal{P}(\bar{O}^{\mathbb{N}})$$

be a Borel version of the reindexed quotient-future conditional law given the full history. QCI/ICS identifies it with the kernel conditioned on the quotient history. Marginalizing the exact anchor gives a Borel kernel $\overline{A}_r$ such that

$$\overline{Q}^{(r)}(h) = \overline{A}_r(V(h)) \qquad (h \in \mathcal{H}_0).$$

Define the Borel map

$$f : \mathcal{H}_0 \to \mathcal{P}(\bar{O}^r), \qquad f(h) = \overline{A}_r(V(h)).$$

If $f(h) = f(h')$, their first-k_{inv} marginals agree because $r \geq k_{\mathrm{inv}}$. The fixed cross-age condition therefore gives

$$\overline{Q}(h) = \overline{Q}(h').$$

Both $\mathcal{P}(\bar{O}^r)$ and $\mathcal{P}(\bar{O}^{\mathbb{N}})$ are standard Borel, and the latter is nonempty. Applying Lemma C.1 to f and $\overline{Q}$ yields a total Borel map

$$\psi : \mathcal{P}(\bar{O}^r) \to \mathcal{P}(\bar{O}^{\mathbb{N}})$$

with $\overline{Q} = \psi \circ f$. Thus $\Psi := \psi \circ \overline{A}_r$ is one total Borel map on Z_{inv} satisfying $\overline{Q}(h) = \Psi(V(h))$ on $\mathcal{H}_0$. This uses no standard-Borel hypothesis on Z_{inv}, teacher-forced update, observation reconstruction, or gauge closure. □

C.3 Proof of gauge-closed full-future sufficiency

Proof of Theorem 5.2. Set

$$\bar{O} = O/G, \qquad \mathsf{Z} = G \times Z_{\mathrm{inv}}, \qquad \mathsf{J} = (O/G) \times G, \qquad Z_s = (U_s, V_s).$$

All equalities below hold on one common probability-one event, obtained by intersecting over the countable governed time set.

Step 1: exact global values force the learned interfaces. This step is needed only for the global-minimizer clause. For a governed time s, set

$$Y_s = (W_{s+1:s+r}, \Delta_{s+1}), \qquad \Pi_s = \mathcal{L}(Y_s \mid \mathcal{F}_s), \qquad \widehat{\Pi}_s = A_r(V_s).$$

Because V_s is $\mathcal{F}_s$-measurable, $\widehat{\Pi}_s$ is an admissible forecast. Conditional strict propriety gives

$$\mathbb{E}[S(\widehat{\Pi}_s, Y_s) \mid \mathcal{F}_s] \geq \mathbb{E}[S(\Pi_s, Y_s) \mid \mathcal{F}_s] \quad \text{a.s.},$$

with equality exactly when $\widehat{\Pi}_s = \Pi_s$ almost surely. After integration and governed-time aggregation, the finite conditional Bayes value is a model-independent lower bound for the anchor term. Subtracting this baseline from the forward-value objective leaves

$$\lambda_{\mathrm{can}}(\mathcal{J}_{\mathrm{can}} - \mathcal{B}_{\mathrm{can}}) + \lambda_{\mathrm{orb}}\mathcal{J}_{\mathrm{orb}} + \lambda_{\mathrm{inv}}\mathcal{J}_{\mathrm{inv}} \geq 0.$$

Simultaneous realizability attains the baseline and zero in both update terms. At any global minimizer all three nonnegative excesses therefore vanish. Strict positivity of the coefficients and of every governed-time weight promotes the aggregate equality to every governed time. Strict propriety gives the exact anchor identity; the zero sets of the update losses give the two exact teacher-forced updates.

Step 2: use the Borel one-step innovation kernel. Use the kernel K supplied by Theorem 5.2. In the global-minimizer application, obtain this kernel by marginalizing A_r as follows. Let

$$\rho : \bar{O}^r \times G \to \mathsf{J}, \qquad \rho((w_1, \dots, w_r), \delta) = (w_1, \delta),$$

and define $K(v; B) = A_r(v; \rho^{-1}(B))$. This is a Borel Markov kernel, and the one-step marginal of the anchor identity gives

$$\mathcal{L}(J_{s+1} \mid \mathcal{F}_s) = K(V_s; \cdot) \quad \text{a.s.}$$

Define the deterministic Borel latent transition

$$\mathsf{T}((u, v), (w, \delta)) = \big(u\delta, \Phi_{\mathrm{inv}}(v, (w, \delta))\big).$$

The update interface states $Z_{s+1} = \mathsf{T}(Z_s, J_{s+1})$ almost surely. Hence the conditional law of (J_{s+1}, Z_{s+1}) depends on the past only through Z_s.

Step 3: reconstruct the next full observation. For $z = (u, v)$ and $j = (w, \delta)$ write $(u^+, v^+) = \mathsf{T}(z, j)$ and set

$$q^+(z, j) = u^+ \chi(v^+)^{-1}.$$

Gauge closure at time $s + 1$ gives $q_{s+1} = q^+(Z_s, J_{s+1})$. The Borel emission

$$\mathsf{E}(z, (w, \delta)) = \mathrm{Rec}(w, q^+(z, (w, \delta)))$$

therefore satisfies $O_{s+1} = \mathsf{E}(Z_s, J_{s+1})$ almost surely.

Step 4: factor the complete future recursively. Define the one-step kernel

$$\mathsf{M}(z; C) = \int_{\mathsf{J}} \mathbf{1}_C\big(j, \mathsf{T}(z, j), \mathsf{E}(z, j)\big) K(v; dj).$$

Starting from z, Ionescu–Tulcea iteration gives a Borel probability kernel $\mathbf{P}_z$ on

$$(J_1, Z_1, O_1, J_2, Z_2, O_2, \dots).$$

The one-step conditional identity establishes the first finite-cylinder case. Conditioning at time $t + n$, applying the same one-step identity, and then using the tower property gives the $(n + 1)$-step case. A monotone-class argument extends the result from product functions to all bounded Borel functions.

Projecting $\mathbf{P}_z$ to the observation coordinates gives a Borel kernel $R_t(z; \cdot)$ on $O^{\mathbb{N}}$. Since a countable cylinder generator determines probability laws on a standard-Borel product space,

$$\mathcal{L}(O_{t+1:\infty} \mid H_t) = R_t(U_t, V_t) \quad \text{a.s.}$$

The argument uses neither QCI/ICS nor $r \geq k_{\mathrm{inv}}$. □

C.4 The non-abelian gauge law

Proof of Proposition 5.3. By definition, $q_{t+1} = q_t \Delta_{t+1}$. Hard right-increment transport therefore gives

$$\begin{aligned} C_{t+1} &= q_{t+1}^{-1} U_{t+1} \\ &= (q_t \Delta_{t+1})^{-1} (U_t \Delta_{t+1}) \\ &= \Delta_{t+1}^{-1} C_t \Delta_{t+1}. \end{aligned}$$

For the left factor $L_t = U_t q_t^{-1}$,

$$L_{t+1} = U_t \Delta_{t+1} (q_t \Delta_{t+1})^{-1} = U_t q_t^{-1} = L_t.$$

If point gauge closure holds, substituting $C_t = \chi(V_t)$ and $C_{t+1} = \chi(V_{t+1})$ gives

$$\chi(V_{t+1}) = \Delta_{t+1}^{-1} \chi(V_t) \Delta_{t+1}.$$

□

D PROOFS FOR CLOSURE PRODUCTION

D.1 INVARIANT-OBJECTIVE FLATNESS

Proof of Proposition 6.1. Let $\mathcal{J}_j$ denote an objective component. By hypothesis, for every admitted adequate edge (m, am),

$$\mathcal{J}_j(am) = \mathcal{J}_j(m) \qquad \text{for every component } j.$$

If m and m' are related by the intrinsic residual relation, a finite chain of such edges, in either orientation, joins them. Componentwise invariance and finite induction give

$$\mathcal{J}_j(m) = \mathcal{J}_j(m') \qquad \text{for every } j.$$

The aggregate objective is therefore constant on every residual class. If m is a global minimizer, every residual-equivalent m' is another global minimizer. In particular, if one class contains a gauge-closed minimizer and a non-gauge-closed minimizer, the unchanged invariant objective retains both and cannot prefer the gauge-closed one. Distinguishing the two requires an architectural restriction, a term that is not invariant on an admitted residual edge, or external calibration. □

D.2 AN OBSERVABLE SECTION PRODUCES CLOSURE

Proof of Proposition 6.2. By the support hypothesis and the countable time convention, all governed observations lie in O° on one common probability-one event. Let $s : \bar{O}^\circ \to O^\circ$ be a Borel section and define

$$\Theta_s : G \times \bar{O}^\circ \to O^\circ, \qquad \Theta_s(g, w) = g \cdot s(w).$$

Joint Borelness of the action makes Θ_s Borel. It is surjective because $s(w)$ lies in orbit w, and it is injective because the action is free. Lusin–Souslin makes its inverse Borel. Hence

$$\tau_s := \mathrm{pr}_G \circ \Theta_s^{-1} : O^\circ \to G$$

is Borel and satisfies $\tau_s(g \cdot o) = g\tau_s(o)$. Conversely, an equivariant G-valued point transporter implies freeness by right cancellation. Extending τ_s by e on $O \setminus O^\circ$ makes the architectural map total; all transporter identities below are on the governed support.

For the aligned section $s_{\mathrm{OR}}(w) = \mathrm{Rec}(w, e)$, equivariance of the pose–quotient reconstructor gives $\mathrm{Rec}(w, q) = q \cdot s_{\mathrm{OR}}(w)$ and therefore

$$\tau_{\mathrm{OR}}(\mathrm{Rec}(w, q)) = q.$$

The architectural choice $U_t = \tau_{\mathrm{OR}}(O_t)$ gives $q_t^{-1}U_t = e$, which is point gauge closure with $\chi \equiv e$.

For a possibly misaligned Borel section s, define $a(w) = \tau_s(s_{\mathrm{OR}}(w))$. The map $w \mapsto [s(w)]$ into O/G is Borel and injective, so Lusin–Souslin identifies $\bar{O}^\circ$ with a Borel subset of O/G. For the pushforward below, extend a by e outside this subset. On the governed support,

$$U_t = \tau_s(O_t) = q_t a(W_t).$$

Hard transport and $q_{t+1} = q_t\Delta_{t+1}$ imply

$$a(W_t) = \Delta_{t+1} a(W_{t+1}) \Delta_{t+1}^{-1}.$$

Set $\varkappa(w, \delta) = \delta a(w)\delta^{-1}$. If $K(V_t; \cdot)$ is the shared exact joint law of (W_{t+1}, Δ_{t+1}) given $\mathcal{F}_t$, conditional expectation on a countable measure-determining algebra, using $\mathcal{F}_t$-measurability of $a(W_t)$, gives

$$\varkappa_{\#} K(V_t; \cdot) = \delta_{a(W_t)} \quad \text{a.s.}$$

The Dirac embedding of G has Borel image and Borel inverse. Applying that inverse to the Borel map $v \mapsto \varkappa_{\#} K(v; \cdot)$, and extending by e off the Dirac image, gives one total Borel χ with

$$q_t^{-1}U_t = a(W_t) = \chi(V_t) \quad \text{a.s.}$$

Thus the section produces point gauge closure. The exact joint kernel is load-bearing because $\varkappa$ depends on the coupling of w and δ. □

D.3 SYNCHRONIZATION PROPAGATES BUT DOES NOT CREATE A GAUGE

Proof of Proposition 6.3. Fix a connected component $\mathcal{C}$ of the synchronization graph and a root i_0. Choose a spanning tree and normalize the root relative frame to $\tau_{i_0} = e$. The exact edge labels recursively determine a unique relative frame τ_i at every other tree vertex. Flatness makes this value path independent. Every exact family of absolute initial frames then has the form

$$x_i = B_{\mathcal{C}}\tau_i, \qquad i \in \mathcal{C},$$

where $B_{\mathcal{C}} \in G$ is arbitrary. Conversely, each choice of $B_{\mathcal{C}}$ gives another exact solution because the common left factor cancels from every relative constraint. Consistent loops test flatness but cannot shrink this residual factor.

Ordered temporal increments give

$$q_{i,t} = x_i\Delta_{i,1}\cdots\Delta_{i,t} = B_{\mathcal{C}}\tau_i r_{i,t}, \qquad r_{i,t} := \Delta_{i,1}\cdots\Delta_{i,t}.$$

If the physical and latent systems have the same relative labels, write their component factors as $B_{\mathcal{C}}$ and $A_{\mathcal{C}}$. Their right gauge is

$$C_{i,t} = r_{i,t}^{-1}\tau_i^{-1}D_{\mathcal{C}}\tau_i r_{i,t}, \qquad D_{\mathcal{C}} := B_{\mathcal{C}}^{-1}A_{\mathcal{C}}.$$

No relative edge or increment determines $D_{\mathcal{C}}$. One physical–latent root cross-link fixes it and hence fixes the point gauge on that component, but that cross-link is external calibration. A disconnected graph requires one such cross-link per component. □

E CONTROLLED OPTIMIZATION AUDIT: PROTOCOL AND OUTPUTS

Outcome-blind freeze. Before executing any outcome command, we fixed the hypotheses, models, exact population laws, seeds $0, \ldots, 99$, optimizer, schedules, stopping rules, thresholds, failure handling, and conjunctive promotion rule. The ordered pre-specification contract has SHA-256 identities

protocol: `fca935abc304e6046ff8f25bcfc3f053dad839f34925f0ab93b76d0aa8572b26`,
environment erratum: `c590ae85b69e9953ee05580435a19c7a459a7bfd4b17ee3f98def4b44315e0f3`,
ordered manifest: `9378b6b9083137a9db84a1dbdbe76544ab51b39d6d154468eb686c70167fa635`.

This is a local outcome-blind pre-specification, not a claim of registration in a public repository.

P-TAB. This lane fixes one quotient fibre and the single transition $0 \to 1$. Its finite local class consists of all equivariant $f_0 : C_2 \to C_2$ (2 choices), $f_1 : C_2^2 \to C_2$ (4 choices), and equivariant predictors $\phi : C_2 \to C_2$ (2 choices). The script enumerates all 16 models with exact rational population weights. Write $f_0(y) = a_0 y$, $\phi(u) = a_\phi u$, and $f_1(y_0, y_1) = y_0 a_r$ for $r = y_0 y_1$. The instantiated zero-one population loss is

$$L_{\mathrm{TAB}} = \mathbb{P}[\phi(f_0(Y_0)) \neq f_1(Y_0, Y_1)],$$

and, since $\mathbb{P}(r = +1) = 5/8$ and $\mathbb{P}(r = -1) = 3/8$,

$$L_{\mathrm{TAB}} = \frac{5}{8}\mathbf{1}\{a_\phi a_0 \neq a_{+1}\} + \frac{3}{8}\mathbf{1}\{a_\phi a_0 \neq a_{-1}\}.$$

Exactly four models have zero loss; all four are persistent copies, and each has a within-fibre next-law total-variation collision gap $1/5$. This is an exact result for the declared one-transition, fixed-quotient-fibre class, not an exhaustive landscape result for the all-age theorem class.

P-NN1. The smooth model initializes $a_{-1}, a_{+1}, w \overset{\mathrm{iid}}{\sim} \mathcal{N}(0, 0.05^2)$ and uses Adam for exactly 2000 steps at learning rate 0.03. Its exact population terms are

$$L_{\mathrm{pred}} = \frac{3}{8}(w - a_{-1})^2 + \frac{5}{8}(w - a_{+1})^2,$$

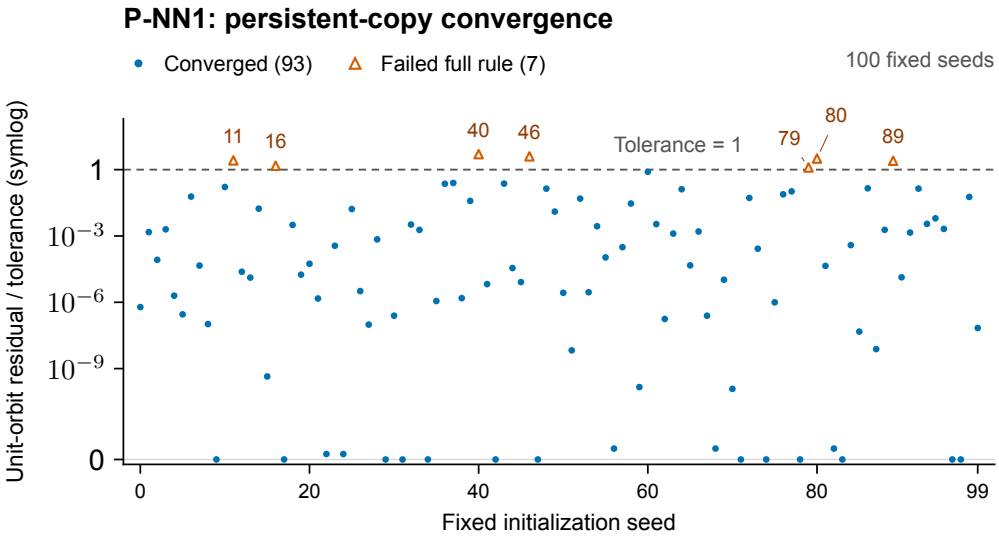


Figure 2: **Persistent-copy convergence in the fixed-seed P-NN1 control.** Each marker represents one of the 100 fixed initialization seeds after the previously specified 2000 Adam steps. The ordinate is $\max_{r\in\{-1,+1\}} |a_r^2 - 1|/10^{-3}$, with the fixed unit-orbit tolerance marked at 1. Filled circles are the 93 seeds satisfying the full convergence rule; open triangles are the seven failures, labeled by seed. All 100 seeds satisfy the sign and amplitude conditions. All seven failures exceed the unit-orbit tolerance; five also exceed the prediction-loss tolerance 10^{-6}. The base-10 symlog axis is linear for normalized residual magnitude at most 10^{-12} and preserves all 13 exact zeros without an added offset. These counts summarize the fixed initialization list in this finite control, not a population-frequency estimate or general optimizer-performance claim.

$$L_{\mathrm{norm}} = \frac{3}{8}(a_{-1}^2 - 1)^2 + \frac{5}{8}(a_{+1}^2 - 1)^2, \qquad L_{\mathrm{NN1}}(\lambda) = L_{\mathrm{pred}} + \lambda L_{\mathrm{norm}}.$$

The unit-orbit weight is zero through step 499, rises linearly to one over steps 500–999, and remains one. A registered convergence requires equal nonzero signs, $\min_r |a_r| \geq 0.99$, prediction loss at most 10^{-6}, and $\max_r |a_r^2 - 1| \leq 10^{-3}$. Confirmation requires at least 90 of the 100 fixed seeds; 93 pass. No seed is excluded, and the seven misses remain in the raw table. All 100 runs have equal nonzero signs, but only 14 have $a_{-1} = a_{+1}$ as exact binary64 values. The registered $1/5$ collision is therefore the collision of the pre-specified sign-quantized readout, not of every continuous latent at step 2000. This lane is an approach-to-manifold mechanism control.

P-NN2. For each seed, identical two-class gauge logits initialize the baseline and calibrated arms. The baseline objective is an explicit differentiable zero in the logits; the calibration adds cross-entropy toward the identity gauge. Both use Adam for exactly 200 steps at learning rate 0.05. The registered conjunction requires maximum baseline drift at most 10^{-12}, 30–70 baseline closure violations, at least 99 calibrated identity gauges, and zero exact physical-equivariance error. The observed values are, respectively, 0, 53, 100, and 0. The zero baseline gradient and physical equivariance are algebraic identities of the parameterization, while the calibrated arm directly supervises the identity gauge. This lane is a mechanism control, not evidence of a discovered optimization phenomenon.

Post-outcome claim audit. The three pre-specified numerical rulings pass their written decision rules, but those rules do not license a field-level claim about ordinary learned world models. P-NN1 uses the declared sign quantizer to create the exact collision, and P-NN2 compares an identically flat direction with direct gauge supervision. We therefore retain P-TAB as an exact one-transition, fixed-quotient-fibre landscape and report P-NN1 and P-NN2 only as scoped mechanism controls.

Implementation and replay. The standard-library implementation uses CPython 3.14.6 on a local arm64 CPU, exact `Fraction` arithmetic for P-TAB, and binary64 arithmetic for the optimization

lanes. The implementation SHA-256 is

`a29f1faf4669fb91b5c73fd9362570a995b6b60fd7cacfe080425f877c348bc6.`

Nine unit tests pass. A fresh optimized-interpreter replay produces all six data and report files byte-for-byte identically; their six-line result manifest has SHA-256

`4964679f6d0e9652ef46effcc41032444e384f9b645cda73cfc0ddb10cdb5834.`

F INITIAL LEARNED STUDY AND READOUT-ONLY CALIBRATION

We study how component calibration interacts with a learned autonomous predictor in a controlled $\mathrm{SO}(2)$ setting. Three successive studies distinguish held-out trajectories of training objects, unseen objects without adaptation, and unseen objects with labelled readout calibration. Only the first study contains a prospective population test; the two follow-ups describe the same 60 frozen model sets on fresh evaluation banks. They do not constitute new training replications. The negative persistence comparisons below are part of the result, not failed validity checks.

Generator, labels, and metric. Each procedural bank contains 16 objects with 24 intensity-marked planar points. Point radii, angles, and intensities are drawn independently from $[0.25, 1.05]$, $[0, 2\pi)$, and $[0.45, 1.55]$, respectively; an object is retained when its intensity-weighted mean position has norm at least 0.24. An observation rotates the object by a unit complex physical state q and adds independent coordinate noise uniform in a disk of radius 0.015; intensities are unchanged. Initial poses are uniform. Actions have uniform angle and magnitude uniform in $[0.2, 1]$. With $z = \overline{q}a$, the physical recurrence is

$$q^{+} = q \exp\big(i \left[0.38 \operatorname{Im} z + 0.16 \operatorname{Im}(z^2) + 0.11 \operatorname{Re} z \operatorname{Im} z\right]\big) . \tag{1}$$

The exact physical probe has four complex channels $y_k = q^k w_k$, $k = 1, \ldots, 4$, with amplitudes $(1, 0.72, 0.46, 0.31)$ and phases $(0.20, -0.73, 1.17, -1.41)$. At horizon h, the reported error is

$$\mathrm{NMSE}_h = \frac{\sum_{j=1}^{n} \sum_{k=1}^{4} |\widehat{y}_{j,h,k} - y_{j,h,k}|^2}{n \sum_{k=1}^{4} |w_k|^2}, \qquad \sum_{k=1}^{4} |w_k|^2 = 1.8261. \tag{2}$$

Lower is better. The constant zero-vector forecast has NMSE exactly one; this is an unconstrained squared-loss reference, not a physical emission or a claim of Bayes optimality. It is distinct from a rollout with zero actions.

Architecture and information budget. The encoder maps an observed cloud to a unit carrier U by multiplying its normalized weighted moment by a learned invariant phase. Its 24-dimensional descriptor d contains 16 radial features and four complex angular moments expressed relative to that observed moment. The encoder phase network is $24 \to 64 \to 64 \to 1$ with tanh activations (5,825 parameters). The transition predicts a phase increment from $(\operatorname{Re}(\overline{U}a), \operatorname{Im}(\overline{U}a), |a|, d)$ using a $27 \to 64 \to 64 \to 1$ tanh network (6,017 parameters). The decoder is $D_k(U, d) = U^k c_k(d)$, with a $24 \to 64 \to 8$ coefficient network (2,120 parameters). During autonomous prediction, only U_0, actions, and the initial observed descriptor d_0 are used; d_0 remains fixed at every step. No future cloud, label, canonical object, or object identifier enters these networks.

Stage A uses 1,600 updates of 160 examples, minimizing noisy-view consistency and known relative-rotation transport losses. REL uses these losses alone. ROOT adds correctly paired fiducial-root calibration; SHUFFLED uses the same root records with fixed object-dependent phase offsets. Both root arms add a loss weighted by 0.1, with one stored noisy pair per object (16 unique pairs, noise radius 0.004), replayed at every update: 25,600 pair exposures and 1,600 additional encoder-forward/backward contributions. All arms have 768,000 core Stage-A observation-forward uses. Thus ROOT and SHUFFLED have extra Stage-A information and computation relative to REL; incorrect correspondences can also actively harm SHUFFLED.

Stage B freezes the encoder and trains the transition and decoder for 1,000 updates of 256 examples. The transition receives a one-step latent cosine loss. The decoder receives exact current/next physical-probe labels on encoded observations, amounting to 512,000 vector-label uses per arm; this physical loss does not backpropagate through the transition or encoder. Both stages use Adam with learning rate 1.5×10^{-3}, default $(0.9, 0.999)$ moments, $\epsilon = 10^{-8}$, and zero weight decay. The comparison is therefore not equal-total-supervision or equal-compute learning.

Table 4: Mean H8 NMSE over 20 paired blocks. The upper panel gives separate evaluation studies. K0/A1/B1/B4 and controls use a third, fresh calibration-study bank set; K0 therefore need not equal the preceding unseen score. Persistence, permuted actions, and zero actions use the B1 readout.

Arm	Seen objects, fresh trajectories	Unseen objects, no adaptation
REL	.538654	3.195031
ROOT	.020244	1.866545
SHUFFLED	.602334	3.443697

Arm	K0	A1	B1	B4	Persist.	Perm.	Zero act.
REL	3.022531	2.168732	1.042801	1.042502	.653334	1.085315	1.110124
ROOT	1.832665	.895467	.889857	.873115	.673767	.990304	.696443
SHUFFLED	3.564461	2.365360	.954221	.936209	.674785	.956495	1.011519

Training-bank confirmation and frozen unseen-object evaluation. The training comparison uses 20 independently seeded procedural banks, each paired across the three arms in initialization and sampled data. A bank plus its initialization/training/evaluation streams is the replication unit. Evaluation uses 768 fresh trajectories per bank, at horizons 1, 2, 4, and 8, on the same 16 objects used in training. The internally frozen prospective protocol defines a ROOT win by $\mathrm{NMSE}_8(\mathrm{ROOT}) - \mathrm{NMSE}_8(\mathrm{REL}) < -10^{-12}$; ties are nonwins. Its sole inferential test is a one-sided exact binomial test of win probability at most $1/2$, with $N = 20$, $\alpha = 0.05$, and a threshold of 15 wins. ROOT wins 20/20, giving $p = 2^{-20} = 9.5367 \times 10^{-7}$. This tests win probability under the fixed procedural training distribution, not the sign or practical importance of a population mean difference. The mean paired difference, -0.518409, is descriptive. Development banks are not pooled, and no public preregistration is claimed.

The subsequent no-adaptation study reuses every saved encoder/transition/ decoder and evaluates a fresh bank of 16 unseen objects and 768 trajectories per paired block. All weights and initial-descriptor rules remain frozen. New identities are from the same generator, not a new domain. ROOT again has lower error than REL in all 20 blocks, but every one of the 60 H8 errors exceeds the post-hoc zero-vector reference of one (Table 4). ROOT's mean H1 error is 2.031863 and its static H8 readout error is 2.056662. This is consistent with a frame/readout transfer problem already at short horizons, but does not isolate that explanation from finite-capacity or extrapolation effects. Arbitrary object frames and action-dependent future uncertainty do not justify a universal H8 impossibility claim. No new significance test is attached to this follow-up.

Labelled readout-calibration protocol. The final diagnostic retains the same 60 frozen model sets and introduces another 20 fresh banks (320 objects). Each object supplies four independently posed calibration clouds with exact same-pose four-harmonic labels, plus 48 independent test trajectories of length eight. Calibration/test coordinate noise is 0.015 on both splits. The 1,280 unique calibration pairs and 15,360 unique test trajectories are shared across arms. One-record methods always use record zero (320 unique records); the four-record method uses all records. Known per-object association routes calibration data externally; association tokens and clean poses are not network inputs. These are strong supplied labels and known associations, not zero-shot or unsupervised adaptation.

K0 retains the original decoder. A1 estimates one output rotation from the first harmonic of the first calibration record: with $v = D_1(U_c, d_c)$ and $Z = y_{c,1}\overline{v}$, set $r = Z/|Z|$ and predict $r^k D_k(U, d_0)$. If $|y_{c,1}||v| = 0$ or $|Z| \le 10^{-12}|y_{c,1}||v|$, use the identity rotation and retain the degeneracy flag (zero such cases among 960 model/object fits). For B1 and B4, fit coefficients from $K = 1$ and $K = 4$ records, respectively:

$$\widehat{c}_k^{(K)} = \frac{1}{K}\sum_{c=1}^{K} \overline{U_c}^{\,k} y_{c,k}, \qquad \widehat{y}_k = U^k \widehat{c}_k^{(K)}. \tag{3}$$

This is unit-carrier least squares with no clipping, projection, or normalization. Under the clean harmonic-label contract, B1's eight stored real numbers obey one consistent phase relation and have the known probe amplitudes; B1 is not an unconstrained eight-parameter fit. Unlike A1, it can replace erroneous original coefficient amplitudes. B4 averages circular moments and can attenuate

Table 5: All primary calibration contrasts (REL, H8). Δ_0 =B1−K0, Δ_P =B1−persistence, and Δ_π =B1−permuted action. Negative gaps favor the B1 true-action forecast.

Block	B1 NMSE	Δ_0	Δ_P	Δ_π
100	1.011279	-2.212361	.317255	-.043581
101	1.314973	-1.509514	.631907	-.055722
102	.885547	-1.769152	.238440	-.053508
103	.978742	-1.700847	.338526	.014119
104	1.285289	-1.835019	.669271	-.073467
105	.722454	-2.420728	.102785	-.075953
106	.933758	-2.221454	.264906	-.036346
107	.925861	-3.498019	.287993	-.010350
108	.882115	-1.496717	.230982	-.122642
109	.905721	-2.663195	.236452	-.042308
110	1.173611	-1.759809	.516270	-.056590
111	1.200056	-2.275523	.543127	-.064343
112	1.152860	-2.423761	.466497	-.011283
113	1.046365	-1.926783	.375956	-.061068
114	.663027	-1.535206	-.001116	-.197337
115	1.165265	-.861153	.531411	.059582
116	1.158831	-1.645308	.499236	-.024048
117	1.574722	-.783502	.921493	-.027019
118	.906046	-3.820590	.293261	.087243
119	.969503	-1.235959	.324690	-.055655
Mean	1.042801	-1.979730	.389467	-.042514

Table 6: Mean B1 errors at each horizon and static H8 readout error. The final column is native relative-motion squared chord error at H8, not NMSE.

Arm	H1	H2	H4	H8	Static H8	Motion H8
REL	.172778	.341919	.629492	1.042801	.000367	.614638
ROOT	.251634	.406471	.622229	.889857	.044731	.466576
SHUFFLED	.196696	.335656	.583872	.954221	.045677	.490695

harmonics; a smaller error need not reflect better dynamics. Native carrier trajectories are identical across all four readouts.

Matched controls and the primary contrast. Persistence applies the same readout to U_0 at every horizon. Static scoring encodes future clouds solely for diagnosis, retaining d_0; it is unavailable to autonomous prediction. B1 additionally uses zero actions and a one-position cyclic shift of complete action sequences within each object's 48 trajectories, holding initial observations and true targets fixed. Each bank weights its 16 objects equally; reported means weight the 20 blocks equally. Comparisons use a numerical tolerance of 10^{-12}, not a practical-importance threshold. The primary descriptive contrast is REL B1 minus K0 at H8. Table 5 retains all 20 paired outcomes and both control gaps. The mean reduction is -1.979730 (median -1.802086), with 20 improvements and no ties. Nevertheless, REL B1's mean error is 1.042801 against persistence 0.653334; ROOT's is 0.889857 against 0.673767. Only REL block 114 beats matched persistence, by 0.001116; the other 59/60 arm/block cases lose.

All three arms improve over K0 in 20/20 blocks. Their mean B1-minus-persistence gaps are $+0.389467$, $+0.216090$, and $+0.279436$ for REL, ROOT, and SHUFFLED; mean B1-minus-permuted gaps are -0.042514, -0.100447, and -0.002274. True actions beat permutation in 17/20, 17/20, and 8/20 blocks, respectively. Action sensitivity does not establish adequate forecasting against persistence. After B1, ROOT beats REL in 13/20 blocks (mean difference -0.152944), and SHUFFLED in 10/20 (mean difference -0.064364). The original uncalibrated ranking is not uniform calibrated superiority.

What the diagnostic locates, and what remains unresolved. REL's B1 H0 and static H8 errors are 0.000366 and 0.000367, showing that supplied-label calibration repairs much of its readout mismatch on independent noisy observations. Its autonomous H8 error remains 1.042801. Table 6 also shows nonzero calibrated static errors for ROOT and SHUFFLED and changing arm order across horizons. The native motion metric is $\operatorname{mean}|\widehat{U}_h\overline{U}_0 - q_h\overline{q}_0|^2$; labels are evaluator-only, and readout calibration cannot change it. These errors retain noisy initial encoding and descriptor/transition interactions. Subtracting static from autonomous risk is not an additive causal decomposition and does not isolate a pure transition error. B4's mean coefficient-modulus/probe-amplitude ratios decrease from 0.998334 at harmonic one to 0.975392 at harmonic four (minimum 0.325281 at four), so its small H8 improvements can include attenuation of uncertain outputs.

All three studies used deterministic local CPU float64/complex128 execution, with immutable checkpoints, paired realized-data identities, and retained traces/results. The training batch took 204.264 seconds with three single-thread workers; the two frozen-model evaluations took about 4.14 and 10.77 seconds. Saved-data numerical replays reproduce the endpoints; these are verification of the same samples, not extra experimental evidence. Fresh seeds and object identity checks concern the accessible retained inventory only. The follow-ups use 20 paired fixed-model units, not 320 independent training replications, and add no p-values or confidence intervals. These three studies support calibration-sensitive representation/readout behavior in this synthetic setting. They do not establish useful unseen-object autonomous forecasting, equal-resource efficiency, necessity for arbitrary world models, or performance on real-world benchmarks.

G FUTURE-LAW SUFFICIENCY DOES NOT IMPLY FIXED-INNOVATION EMISSION

Relation to existing representation results. Representative-map factorization in Bloem-Reddy & Teh (2020, Lemma 8) supplies the algebraic mechanism in Theorem 4.2. On our product domain, equivariance gives $\mathcal{E}((u,v),j) = u \cdot B(v,j)$ with $B(v,j) = \mathcal{E}((e,v),j)$. Writing $b = \tau \circ B$, the actual emission identity implies $C_t = \Delta_{t+1} b(V_t, J_{t+1})^{-1}$, in that order. Pushing the common conditional kernel through this expression yields a Dirac law because C_t is already history-measurable. The Borel inverse of the Dirac embedding then gives $C_t = \chi(V_t)$. These are standard factorization and conditional-expectation mechanisms. The characterization concerns their application to a fixed actual innovation; it is not a new universal representation principle.

An explicit separation. Use additive $G = C_2$ on $O = G \times \{0,1\}$ with $g \cdot (q,w) = (g+q, w)$. Draw independent fair bits $Q_0, C, D_1, D_2, \ldots$ and set

$$W_0 = C, \quad W_t = 0 \ (t \geq 1), \quad Q_{t+1} = Q_t + D_{t+1}, \quad O_t = (Q_t, W_t).$$

The initial distribution is uniform on four states. The homogeneous transition sends every state to each of $(0,0),(1,0)$ with probability $1/2$. Both the initial law and the transition respect the free group action. The observation-aligned trivialization is $\mathrm{Rec}(w,q) = (q,w)$ with inverse $(\tau,\pi)(q,w) = (q,w)$. The total encoder on nonempty histories is

$$F((q_0,w_0),\ldots,(q_t,w_t)) = (q_t + w_0, \star).$$

It is Borel and equivariant even on impossible histories. On the process, $U_t = Q_t + C$, $V_t = \star$, $C_t = C$, and hard transport is exact. Furthermore

$$\mathcal{F}_t = \sigma(Q_0, C, D_1, \ldots, D_t), \qquad J_{t+1} = (0, D_{t+1}), \qquad K(\star) = \tfrac{1}{2}\delta_{(0,0)} + \tfrac{1}{2}\delta_{(0,1)}.$$

Thus all F+ class contracts hold with one common joint kernel.

Conditional on any history, the map from the next n increments to the next n poses is a bijection of G^n. Every pose word has probability 2^{-n} and every future quotient tag is zero. The finite cylinders determine the entire future law, so

$$\mathcal{L}(O_{t+1:\infty} \mid \mathcal{F}_t) = \mu := \big(\mathrm{Unif}(G) \otimes \delta_0\big)^{\otimes\mathbb{N}}.$$

The constant kernel $R_t(u,\star) = \mu$ proves full-future sufficiency. This law is equivariant and retains fresh randomness at every future time.

Nevertheless, $C_t = C$ is fair while V_t is constant, so point closure fails. Since C is also independent of (U_t, D_{t+1}) and $Q_{t+1} = U_t + C + D_{t+1}$,

$$\mathbb{P}(Q_{t+1} = a \mid U_t = u, J_{t+1} = (0,d)) = \tfrac{1}{2} \quad (a,u,d \in C_2).$$

No deterministic emitter, equivariant or otherwise, reproduces the next observation almost surely. Its success probability is at most $1/2$; $\mathcal{E}_0((u,\star),(w,d)) = (u+d,0)$ attains this bound. In particular, the atoms $(Q_0, C, D_1) = (0,0,0)$ and $(1,1,0)$ have equal allowed emitter inputs and different targets, each with probability $1/8$.

Almost-sure auxiliary-noise representation still exists. Let $B_{t+1} = C + D_{t+1}$. This bit is fair and independent of the entire $\mathcal{F}_t$: conditioning on a history fixes C but leaves D_{t+1} fair. Extend the probability space with independent $A_{t+1} \sim \mathrm{Unif}[0,1)$, and put $\eta_{t+1} = (B_{t+1} + A_{t+1})/2$. Then η_{t+1} is uniform and independent of $\mathcal{F}_t$, and the common pointwise-equivariant function

$$f(r,(u,\star)) = (u + \mathbf{1}_{[1/2,1]}(r), 0) \quad \text{satisfies} \quad O_{t+1} = f(\eta_{t+1}, (U_t, \star)) \quad \text{a.s.}$$

The input latent has an invariant law and the group is finite, matching the relevant compact-group hypotheses of Bloem-Reddy & Teh (2020, Theorem 9). The theorem permits an existentially chosen auxiliary-noise coupling; it does not promise emission from the world's already specified J. Indeed $B_{t+1} + D_{t+1} = C$: their pair reveals history information although each bit separately is independent of the history. Substituting B for the physical increment D changes the interface.

The construction reuses the disappearing-tag mechanism of Theorem 3.2; unlike the sharpness example in Appendix B.2, it has an invariant input law and an equivariant output law for the matched comparison. It establishes a contract separation, not a learned forecasting advantage or a universal necessity of point closure. Bare F+ also contains no recursive-update assumption, so emission alone does not imply the full-future conclusion of Theorem 5.2.

H Calibration routing, training interventions, and information limits

The studies below extend the learned diagnostic in Appendix F. They distinguish changing an inference interface, retraining a transition, changing encoder supervision, and changing the available observations. Positive and adverse comparisons use their own declared populations; results from different studies are not pooled. The harmonic target, generator, encoder, descriptor, and NMSE normalization are as defined in Appendix F, unless specified otherwise. The statistical unit is a paired model/data block, not an individual object, trajectory, horizon, or repeated method evaluation.

H.1 Routing a supplied calibration into a frozen transition

Information and intervention. Each object supplies one noisy calibration cloud and its four exact same-pose physical harmonic labels, with known object association. The selected record is always the first of four generated records; the other three are not used for fitting or selection. This is privileged orientation-bearing information. Let U_c be its encoded unit carrier, $\beta_{i,k} = \overline{U}_c^{\,k} y_{c,k}$, b the fixed training-fiducial direction, and $w_1 = \exp(0.20i)$. The routing factor is $r_i = b\beta_{i,1}/w_1$, without projection or normalization. For a frozen transition T and initial descriptor d_0,

$$\widetilde{U}_0 = r_i U_0, \qquad \widetilde{U}_{h+1} = T(\widetilde{U}_h, a_h, d_0), \qquad \widehat{U}_h = \overline{r}_i \widetilde{U}_h, \qquad \widehat{y}_{h,k} = \beta_{i,k} \widehat{U}_h^k. \tag{4}$$

Here $\widetilde{U}$ is the routed carrier, not the theory's invariant code V. Actions retain their global coordinates. Native prediction uses $r = 1$; association placebo uses the next object's factor cyclically while retaining the correct β_i. Action permutation shifts complete eight-step action sequences by one trajectory within each object's 48 trajectories, keeping initial observations and scoring targets fixed. Persistence applies the same readout to U_0. Future observations and physical states enter scoring only. Routing changes the inputs to the frozen network; it is not a conjugate coordinate re-expression that would leave predictions unchanged.

Exploratory screen and prospective fresh-data check. The initial screen reuses the previously inspected calibration-study data and all 20 ROOT models. A subsequent internally frozen protocol evaluates the same models on 320 fresh objects and 15,360 fresh H8 trajectories, with 320 supplied calibration records. Both studies average objects within a block and weight the 20 blocks equally. Their descriptive criterion requires the mean and median of routed-minus-{native, persistence, placebo, routed-permuted} H8 gaps all to be below -10^{-12}. The fresh-data check meets that

Table 7: Frozen-model routing, mean H8 NMSE over 20 blocks. Every path uses the correct per-object readout. Fresh-data paired gaps for routed true minus native, persistence, placebo true, and routed permuted are respectively $-.821550$, $-.600782$, $-.836649$, and $-.920535$, each favorable in 20/20 blocks. The two columns use distinct evaluation data and are not pooled.

Path	Inspected-data screen	Fresh-data check
Native, true actions	.889857	.897903
Native, permuted actions	.990304	.999728
Routed, true actions	.076403	.076353
Routed, permuted actions	.993594	.996888
Association placebo, true actions	.890219	.913002
Association placebo, permuted actions	.986199	.997838
Persistence	.673767	.677135

criterion; the observed 20/20 favorable blocks for each contrast are reported outcomes, not an additional test. It is conditional on these models and the same generator, not new training replication or public preregistration. No p-value or confidence interval is attached to either routing follow-up.

These results locate a useful inference-interface intervention for the retained models. They do not establish its necessity, a unique explanation of prior error, equal historical ROOT/REL supervision, or unlabelled frame recovery. The original native-interface failures remain valid.

H.2 One-step failure and the matched four-arm intervention

Matched one-step experiment. Each of the same 20 frozen ROOT encoders received 16 fresh training and 16 disjoint evaluation objects. Two fresh 6,017-parameter transition heads shared initialization, 256-example batches, exact physical next-step targets, and 1,000 Adam updates. One used native coordinates and the other Equation 4; neither used a learned decoder. Both had first-record K1 calibration on training and evaluation objects. Only the transition received gradients from energy-normalized physical squared error. Training used the current observed carrier and current descriptor; autonomous evaluation fixed d_0. Adam used learning rate $.0015$, moments $(.9, .999)$, $\epsilon = 10^{-8}$, and no weight decay. The final checkpoint was evaluated without selection on 16 objects $\times$ 48 H8 trajectories per block.

The prescribed useful-prediction conjunction failed. Native, routed, and persistence H8 means were 1.431634, 1.156480, and $.683852$. Routed-minus-native was $-.275154$ (18/20 favorable), but routed-minus-persistence was $+.472628$ (only 2/20 favorable). Routed-minus-placebo and routed-minus-permuted were $-.279185$ and $-.294755$ (20/20 each). Thus action/association sensitivity did not imply a useful H8 forecast. H1/H2/H4 routed means were $.107399/.242633/.591771$, versus persistence $.166617/.267627/.434088$; the shorter horizon cannot replace H8.

Four matched training arms. The subsequent experiment retains every one of the 20 encoders and trains four fresh, identically initialized heads per block: native/routed crossed with teacher forcing (TF)/autoregressive training (AR). All four heads use the same 32 H8 sequences per update, 1,000 updates, the optimizer above, fixed d_0 at all eight steps, and loss summed over all eight physical targets and four harmonics, divided by target energy. TF consumes observed carriers during training; AR consumes its own predictions with full H8 backpropagation through time. Both are evaluated autonomously. Equal raw sequences and transition applications do not imply equal active observation access, FLOPs, or optimization difficulty. Within each training regime, native and routed heads have the same supplied calibration information.

There are 320 new training objects and 320 disjoint evaluation objects, with 15,360 held-out trajectories. One K1 fit per object is shared across all four heads. All 80 final heads are retained; there is no seed or checkpoint selection. The primary descriptive conjunction requires mean and median of routed AR minus routed TF and routed AR minus persistence to be negative at the fixed 10^{-12} tolerance. Both are negative in every block.

Native AR remains worse than persistence in mean and median; the interaction does not favor AR in two blocks. The result supports this joint training and routing recipe, conditional on ROOT encoders

Table 8: Four-arm experiment: all H8 paths, mean NMSE over 20 paired blocks. TF/AR describe training only. Both persistence entries are the same forecasts.

Evaluation path	TF	AR
Native, true actions	1.221177	.707873
Native, permuted actions	1.222761	.727712
Routed, true actions	.964155	.127500
Routed, permuted actions	1.358783	1.005090
Association placebo, true actions	1.320647	.901203
Association placebo, permuted actions	1.359146	1.003065
Persistence	.675110	.675110

Table 9: Every predeclared four-arm H8 contrast. Negative favors the first term. Counts are lower/tied/higher across paired blocks; the interaction is $(\mathrm{routed}_{AR} - \mathrm{native}_{AR}) - (\mathrm{routed}_{TF} - \mathrm{native}_{TF})$. No inferential test is added.

Contrast	Mean gap	Median gap	Counts
Routed AR − routed TF	-.836656	-.847244	20/0/0
Native AR − native TF	-.513304	-.484129	20/0/0
Routed AR − persistence	-.547611	-.565279	20/0/0
Routed TF − native TF	-.257022	-.200527	16/0/4
Routed AR − native AR	-.580373	-.591933	20/0/0
Routed AR − placebo AR	-.773704	-.771609	20/0/0
Routed AR true − permuted	-.877590	-.912378	20/0/0
Interaction	-.323351	-.379417	18/0/2

and supplied calibration. The earlier one-step recipe used a different descriptor schedule, so the cross-study improvement does not uniquely explain its failure. Multi-step training itself is not claimed as a new method.

Fixed-model shape and calibration-error sensitivity. A separate study uses the original frozen models, without the newly trained heads, on the original object family and one ellipse family. The latter has 24 points with $\theta_j = 2\pi(j+\phi)/24$, $x_j = R_\gamma(a\cos\theta_j, b_e\sin\theta_j)$ and $I_j = 1 + \kappa\cos\theta_j$, where independent $a \sim U[.96, 1.02]$, $b_e \sim U[.38, .62]$, $\kappa \sim U[.52, .54]$, $\phi \sim U[0,1]$ and $\gamma \sim U[0, 2\pi]$ require no rejection. The dynamics, actions, probe and .015 coordinate noise are unchanged. For each object, one draw $z \sim U[-1, 1]$ is reused across angular-error bounds $\delta \in \{0, 5, 15\}$ degrees; supplied labels are $\widetilde{y}_k = \exp(ik\delta\pi z/180)y_k$. All paths share this noisy readout. These are coherent per-object errors, not independent harmonic noise. Two families contain 640 unique objects and 30,720 unique trajectories in total; reuse across error levels does not triple the sample count.

H.3 Encoder supervision under two downstream regimes

This experiment tests whether the earlier ROOT advantage persists after accounting for extra relative-supervision exposure and changing downstream training. It trains 20 new paired blocks with three Stage-A parents each: REL, EXTRA_REL, and ROOT. Each block has three disjoint 16-object banks for Stage A/original B, calibrated AR training, and common held-out evaluation (960 objects total). Stage A retains the architecture and 1,600-update recipe in Appendix F. EXTRA_REL adds 16 stored relative pairs, weighted .1, replayed at every update; ROOT instead adds 16 correctly paired fiducial targets with the same weight. REL and EXTRA_REL never receive root targets. Per-parent encoder uses are 768,000/819,200/793,600 for REL/EXTRA_REL/ROOT. Thus equal auxiliary pair counts do not equalize information or computation.

Each frozen parent has exactly two children. Original B retains its 6,017-parameter transition, 2,120-parameter decoder, 1,000 updates of batch 256, latent cosine objective, and physical static loss reaching the decoder but not the predicted transition. It receives no test calibration. Calibrated AR uses one 6,017-parameter head, 1,000 updates of 32 H8 sequences, full BPTT, fixed d_0, the physical sequence objective, and K1 routing recomputed for that encoder and split. Its training API receives only initial carriers, actions, descriptors, calibration, and physical targets; future observations are

Table 10: All six sensitivity conditions, mean NMSE over the fixed 20 models. The four routed H8 paired gaps have negative mean and median in every condition (20/20 favorable blocks per contrast). These are dependent descriptive conditions, not six independent replications.

Family	δ	H0	Native	Routed	Persist.	Placebo	Routed perm.
Original	0°	.038998	.893425	.065567	.667392	.896870	.984685
Original	5°	.047135	.895877	.077122	.670348	.899801	.986071
Original	15°	.110001	.916230	.155977	.696539	.921094	1.000635
Ellipse	0°	.026105	.885403	.064746	.666821	.897999	.983474
Ellipse	5°	.033304	.887948	.074831	.669102	.900741	.984377
Ellipse	15°	.097011	.910257	.153902	.694634	.922421	.999566

Table 11: All six frozen encoder-supervision contrasts. Intervals are marginal approximate 95% Student-t intervals on 20 paired block differences, $\bar{d} \pm t_{.975,19} s_d/\sqrt{20}$; they are neither simultaneous nor distribution-free. Counts give negative/positive differences (no ties).

Regime/panel	Contrast	Mean	95% interval	Counts
AR/held-out (primary)	ROOT − EXTRA_REL	.140583	[.094076,.187091]	0/20
AR/held-out	EXTRA_REL − REL	.000006060	[-.000033772,.000045892]	8/12
B/same-object	ROOT − EXTRA_REL	-.523911	[-.562030,-.485791]	20/0
B/same-object	EXTRA_REL − REL	.000107546	[-.000408347,.000623439]	8/12
B/held-out	ROOT − EXTRA_REL	-1.351227	[-1.668041,-1.034414]	20/0
B/held-out	EXTRA_REL − REL	.003693206	[-.002848330,.010234742]	9/11

excluded. Initialization and raw schedules are paired across Stage-A arms within each regime. All final heads are retained.

Original B's same-object panel uses 768 empirically sampled trajectories per block (realized counts 29–66 per object). The common held-out panel uses 16 objects × 48 trajectories and is shared by all six children. The 20 paired blocks, not these repeated 15,360 held-out trajectories, determine uncertainty. All 60 parents and 120 children completed their budgets without replacement or excluded pairs.

The primary result is adverse: ROOT loses to EXTRA_REL in all 20 AR blocks. ROOT also loses to REL in every AR block, while outperforming both in each B panel. The EXTRA_REL–REL intervals cross zero; this is not equivalence. These regimes differ in objective, training observations and calibration, so their reversal does not isolate a single causal mechanism. In particular, ROOT's low same-object B error does not solve held-out prediction, whose mean remains 1.888265.

System identification comparator. Each block additionally fits one six-feature quadratic transition by unregularized SVD (cutoff 10^{-12}), shared across Stage-A arms. The features are $(1, x, y, x^2, xy, y^2)$ for $x + iy = \overline{q}a$. Only allowed physical labels $q_1, \ldots, q_8$ and actions $a_1, \ldots, a_7$ supply adjacent pairs: 224,000 rows per block, excluding the hidden initial q_0. All 20 fits have rank six and near-machine-precision training residuals under this source law. The calibrated observed-moment SID comparator outperforms all three learned AR heads. This simple strong baseline and its privileged physical-label, K1, and structural information are part of the result.

H.4 History and calibration in an unknown-mounting control task

Original task and observation interface. We evaluate the registered `Pendulum-v1` reset, dynamics, reward and 200-step TimeLimit. Reset draws $\theta_0 \sim U[-\pi, \pi]$ and $v_0 \sim U[-1, 1]$. Every episode receives the fixed float32 torque prefix $(0, .5, -.5, 1, -1, 0, .25, -.25)$, then selects one open-loop H16 sequence at time eight, with no later observation or reward access. The ordered 27-arm library contains all triples in $\{-2, 0, 2\}^3$, applying the three torques for 5, 5, and 6 steps; zero torque is first, followed by the other triples in lexicographic order. Ties use the lowest index. Cost is the sum

Table 12: Mean H8 physical NMSE for the encoder-supervision experiment. SID uses privileged physical labels and a fixed structural feature family. The observed-moment row is one shared baseline, not three independent fits.

Method/panel	REL	EXTRA_REL	ROOT
Original B, same-object	.544368	.544475	.020564
Original B, held-out	3.235799	3.239492	1.888265
Calibrated AR, held-out	.012055	.012061	.152644
Calibrated carrier persistence	.654148	.654150	.671643
SID with calibrated encoder	.000476051	.000476647	.053792222
SID with calibrated observed moment	.000377631	.000377631	.000377631

of the original pre-transition costs,

$$J(u; s_8) = \sum_{t=8}^{23} \left[\text{angle_normalize}(\theta_t)^2 + .1v_t^2 + .001\,\text{clip}(u_t, -2, 2)^2\right]. \tag{5}$$

Regret subtracts the minimum of the 27 evaluator-only true-state costs; it is not regret against an unrestricted optimal controller.

An episode-fixed mounting angle is independently uniform on $[-\pi, \pi)$. The sensor first rounds $(\cos\theta, \sin\theta, v)$ to float32, rotates the first two coordinates by mounting plus independent uniform angular noise, then rounds them again. Velocity has rounding but no added noise. Noise radii $0°$ and $5°$ share coupled base draws. The eight cells cross noise with current/history and uncalibrated/calibrated input. Current is one three-coordinate record; history is nine records and the eight issued torques. Calibration is an independently corrupted upright record from $(1, 0, 0)$ and its known-pose declaration. The physical state after the prefix has the reset law's pushforward distribution; it is not assumed Haar or independent of velocity.

Training, decisions, and units. Twelve independent blocks each contain 4,096 training, 512 validation, and 128 test episodes: 49,152/6,144/1,536 physical episodes, 56,832 total. Neither those 56,832 episodes across splits nor their repeated views are 56,832 independent training replications. The inferential unit is the block ($n = 12$). Both fitted families share training physical-state labels at all 25 times and actions, and the same frozen supervised angle-density estimator in each cell. These are privileged training labels; no physical labels, seeds, rewards or mounting values enter test packets.

Each density estimator has two width-128 tanh layers and an eight-component von Mises mixture, fitted for 2,000 Adam updates at batch 128. The learned transition has two width-128 tanh layers, input $(\cos\theta, \sin\theta, v/8, u/2)$ and normalized direction/clipped velocity outputs, fitted for 4,000 updates at batch 256. Both use float64 Adam with learning rate .001, moments $(.9, .999)$, $\epsilon = 10^{-8}$, no weight decay, gradient-norm clipping at 5, and final checkpoints without selection. System-ID instead fits A, B by rank-two least squares on unclipped training transitions in $v' = \text{clip}(v + .05(A\sin\theta + Bu), -8, 8)$, $\theta' = \theta + .05v'$. Its known integration structure and clipping bounds are additional priors. All 96 density fits, 12 learned transitions and 12 SID fits completed.

Fitted consumers average costs over 256 fixed angle midpoints weighted by the estimated density, using the rounded observed velocity. A 512-grid diagnostic on eight fixed test indices per block/cell does not alter choices. It records one argmin disagreement across the 1,536 method/cell diagnostics, and maximum arm-cost discrepancy .018874. Neither grid is an exact posterior or a certified Bayes computation.

For block differences d_b, the adjusted critical level is $1 - .05/(2 \cdot 8)$ with standard error $s_d/\sqrt{12}$. All eight intervals exclude zero positively, but none is wholly within the predeclared practical-matching interval $[-.16, .16]$. History is useful in these fitted consumers and does not fully substitute for calibration. The additional bounded-block Hoeffding intervals use $M = 16(\pi^2 + 6.4 + .004) = 260.377670\ldots$, $d_b \in [-M, M]$, and half-width $M\sqrt{2\log(2/(.05/8))/12} \simeq 255.30118$. All eight include zero. Thus the reported directional evidence is not an assumption-free conclusion. The 104 frozen secondary contrasts have unadjusted descriptive intervals, not simultaneous coverage.

Table 13: All eight fitted-method cells: mean H16 cumulative cost over 12 blocks. C/H means current/history; U/K means uncalibrated/calibrated. Learned minus SID is positive in all eight cells. Their shared library-oracle cost is 89.826775; subtracting it gives mean realized library regret.

Noise	Information	SID	Learned	Learned − SID
0°	C/U	96.381584	96.610201	.228617
0°	H/U	92.891888	93.128358	.236470
0°	C/K	89.856474	90.025285	.168811
0°	H/K	89.894207	90.047591	.153384
5°	C/U	96.722696	96.916446	.193750
5°	H/U	92.719744	92.970815	.251071
5°	C/K	90.002402	90.152493	.150091
5°	H/K	89.899285	90.064973	.165688

Table 14: Every primary control-task contrast, in cumulative cost units. History value is C/U minus H/U; residual calibration value is H/U minus H/K. Intervals are Bonferroni-adjusted approximate 95% Student-t intervals for the eight-contrast family, using 12 paired block means and 11 degrees of freedom. They rely on the sample-mean t approximation.

Family	Noise	Contrast	Estimate	Adjusted interval
SID	0°	History	3.489696	[1.387324,5.592067]
SID	0°	Calibration	2.997681	[1.367501,4.627861]
SID	5°	History	4.002952	[1.690012,6.315892]
SID	5°	Calibration	2.820459	[.917700,4.723218]
Learned	0°	History	3.481844	[1.840173,5.123514]
Learned	0°	Calibration	3.080767	[1.530952,4.630583]
Learned	5°	History	3.945631	[1.443809,6.447453]
Learned	5°	Calibration	2.905842	[.707154,5.104531]

Strong controls and unresolved certification. SID has lower mean cost than the learned transition in every cell; the four calibrated marginal descriptive intervals favor SID, whereas the four uncalibrated ones cross zero. Zero torque costs 107.915097 (regret 18.088322); assuming zero mounting costs $108.393345/108.406507$ at $0^\circ/5^\circ$. The known-model calibration plug-in has regret zero at 0° and $.136713$ at 5°, and intentionally ignores history. The common prefix costs 36.923641; it is additional to every H16 cost.

The original-runtime outer-set reference retains unresolved contractor and numerical status in all 12,288 method/cell rows and selects zero torque. A separate conditional-envelope reference assumes exact-real dynamics with per-step state errors at most 2^{-40}, cost error at most 2^{-28}, and sensor angular enclosure error at most 2^{-20}. Its source-to-envelope assumptions, including a host trigonometric error bound, remain unverified. It optimizes a relaxed outer support, not the exact posterior or full feasible set, at cumulative budgets 4,096/16,384/65,536 boxes. Its current-uncalibrated mean cost 111.967340 is worse than zero torque. All four history cells and zero-noise current-calibrated cells have zero retained realized library regret; the five repeated cells do not constitute 7,680 independent physical episodes. At 5°, current-calibrated regret $.340952$ exceeds the plug-in's $.136713$. Conditional statuses are 9,180 outer-set optimization certificates, 23 selected-arm-cost certificates, and 3,085 numerical-unresolved rows. Observed coverage of all retained truths does not validate universal source-to-envelope inclusion. Rounded, otherwise unnoised velocity histories provide strong information in this setting; the result establishes no general control optimum or real sensor acquisition-cost tradeoff.

H.5 Continuous-prior information and finite-horizon Bayes risk

This subsection states the existing continuous-law result with its full assumptions and proof. It concerns a fresh random test object before conditioning on its realized canonical coordinates or test bank. Old training evidence and learned parameters may be held fixed provided that the fresh law is independent of them. They are not an additional test reference.

Let λ be normalized Haar measure on S^1. A canonical object $C = ((z_j, w_j))_{j=1}^{24}$ is sampled from independent radii uniform on $[1/4, 21/20)$, angles with law λ, and intensities uniform on $[9/20, 31/20)$, accepting only when

$$\left| \frac{\sum_j w_j z_j}{\sum_j w_j} \right| \geq \frac{6}{25}. \tag{6}$$

Let $Q_0 \sim \lambda$, initial additive coordinate noise ε, and future global-frame actions $A = (a_0, \ldots, a_{H-1})$ be mutually independent and independent of C. Each noise point has magnitude at most $3/200$; the source uses independent uniform-disk draws. Define

$$X_0 = ((Q_0 z_j + \varepsilon_j, w_j))_{j=1}^{24}. \tag{7}$$

The horizon $H \geq 0$ is finite. Each action is $a_t = r_t e^{i\phi_t}$ with $1/5 \leq r_t \leq 1$; the endpoint one closes the source support. The exact-real transition is

$$Q_{t+1} = Q_t \exp\left(i \left[\frac{19}{50} \operatorname{Im}(\overline{Q_t} a_t) + \frac{43}{100} \operatorname{Re}(\overline{Q_t} a_t) \operatorname{Im}(\overline{Q_t} a_t) \right] \right). \tag{8}$$

For deterministic a, write its finite composition as $\Phi_a : S^1 \to S^1$, with $\Phi_\varnothing$ the identity. The target is $Y_H = (c_k Q_H^k)_{k=1}^4$, with constant squared norm $S = \sum_{k=1}^4 |c_k|^2 > 0$. The nonzero c_k have magnitudes $(1, .72, .46, .31)$ and phases $(.20, -.73, 1.17, -1.41)$. A predictor is any measurable $f(X_0, A) \in \mathbb{C}^4$, with loss $\|f - Y_H\|_2^2/S$; no unit-phase constraint is imposed. All spaces are standard Borel. A deterministic PRNG sequence, its finite floating-point image, realized test banks, model-fitting output, and a predictor allowed cross-scenario adaptation are not assigned this law.

Theorem H.1 (Continuous fresh-object risk boundary)**.** *Under the preceding law:*

1. *Q_0 is independent of (X_0, A). Conditional on the allowed information, Q_H has kernel $(\Phi_A)_\# \lambda$.*

2. *With $\mu_k(a) = \int_{S^1} \Phi_a(q)^k \, d\lambda(q)$, a Bayes-optimal predictor and its conditional normalized risk are*

$$f_k^*(X_0, A) = c_k \mu_k(A), \qquad R_H^*(a) = 1 - \frac{\sum_{k=1}^4 |c_k|^2 |\mu_k(a)|^2}{S}. \tag{9}$$

   The unconditional optimal risk is $\mathbb{E}_A R_H^*(A)$. The observation X_0 cannot improve this optimum.

3. *Put*

$$M(r) = \begin{cases} 1 + \dfrac{19}{50} r - \dfrac{43}{100} r^2, & 1/5 \leq r \leq 19/86, \\ 1 + \dfrac{43}{100} r^2 + \dfrac{361}{8600}, & 19/86 \leq r \leq 1, \end{cases} \qquad m(a) = \prod_{t=0}^{H-1} M(r_t)^{-1}. \tag{10}$$

   The empty product is one. For every allowed finite action sequence,

$$R_H^*(a) \geq 2m(a) - m(a)^2 \geq 2 \left(\frac{8600}{12659} \right)^H - \left(\frac{8600}{12659} \right)^{2H} > 0. \tag{11}$$

   At $H = 0$, $R_0^* = 1$ exactly.

The formula defines a conditional-law version for all deterministic a; its interpretation for random A is almost sure. The bound is not claimed sharp as a function of H or the whole action sequence. A smaller predictor class has optimal risk at least this unrestricted optimum.

Proof. Accepted object law and observation independence. The acceptance event has positive probability: an open neighborhood of objects with all coordinates near the positive real number $1/2$ and weights near one satisfies the strict moment inequality. The pre-acceptance law is invariant under $C \mapsto gC$, which multiplies every coordinate by $g \in S^1$ and leaves intensities unchanged. Each angle is uniform, so a common shift preserves their joint distribution; radii and intensities are unchanged.

The moment rotates by g, leaving its norm and the acceptance event unchanged. Conditioning on this positive-probability event preserves rotation invariance.

Let ν denote the accepted object law. For bounded measurable b on the observation space and fixed $q \in S^1$,

$$\mathbb{E}_{C,\varepsilon}[b(qC+\varepsilon)] = \mathbb{E}_{C,\varepsilon}[b(C+\varepsilon)]. \tag{12}$$

The notation includes unaltered intensities. The equality follows from $qC \overset{d}{=} C$ and independence of ε; noise isotropy is not required. For bounded measurable h and k,

$$\mathbb{E}[h(Q_0)b(X_0)k(A)] = \left(\int h\, d\lambda\right) \mathbb{E}[b(C+\varepsilon)]\, \mathbb{E}[k(A)]. \tag{13}$$

Indicator rectangles and a monotone-class argument give joint independence. The conditional law of Q_0 given (X_0, A) is λ, not a point estimate of canonical object orientation. Applying the deterministic composition yields the stated conditional law of Q_H.

The argument holds for any rotation-invariant object law independent of the other variables. It does not hold by this argument after supplying the realized canonical object or conditioning on an arbitrary fixed bank. A shape fingerprint in X_0 does not reveal an independent unobserved global orientation of a fresh canonical object under this law. The bounded-noise margin adds no orientation-dependent selection: the weighted perturbation mean has magnitude at most $3/200$, so the observed moment norm is at least $6/25 - 3/200 = 9/40$ for every allowed noise realization.

Exact Bayes-risk identity. The map $(a, q) \mapsto \Phi_a(q)$ is continuous, so each $\mu_k(a)$ is measurable and bounded by one. The conditional law gives $\mathbb{E}[Y_{H,k} \mid X_0, A] = c_k \mu_k(A)$. For any square-integrable predictor f, conditional-expectation orthogonality in complex Euclidean space gives

$$\mathbb{E}[\|Y_H - f\|_2^2 \mid X_0, A] = \mathbb{E}[\|Y_H - f^*\|_2^2 \mid X_0, A] + \|f - f^*\|_2^2. \tag{14}$$

Expanding the square makes each conditional cross term vanish. For finite-risk predictors, square integrability follows from bounded Y_H and $\|f\|^2 \le 2\|f - Y_H\|^2 + 2\|Y_H\|^2$; infinite-risk predictors cannot lower the infimum. Because $\|Y_H\|^2 = S$ in every realization, the first term divided by S is $1 - \|f^*(a)\|^2/S = 1 - \sum_k |c_k|^2 |\mu_k(a)|^2/S$. The bounded measurable f^* attains it; integration over A proves the unconditional statement. At H0 every nonzero Haar harmonic integrates to zero: rotating its integral by an element with $g^k \neq 1$ forces it to vanish. Hence $f^* = 0$ and $R_0^* = 1$. The conditional mean need not belong to the physical four-harmonic phase manifold; it is a relaxed Bayes comparator, not necessarily an output of a trained decoder.

Circle dynamics and the derivative maximum. For $a = re^{i\phi}$, a real lift is

$$F_a(\theta) = \theta + \frac{19}{50} r \sin(\phi - \theta) + \frac{43}{200} r^2 \sin(2(\phi - \theta)). \tag{15}$$

It satisfies $F_a(\theta + 2\pi) = F_a(\theta) + 2\pi$ and

$$F_a'(\theta) = 1 - \frac{19}{50} r \cos(\phi - \theta) - \frac{43}{100} r^2 \cos(2(\phi - \theta)) \ge 1 - \frac{19}{50} r - \frac{43}{100} r^2 \ge \frac{19}{100}. \tag{16}$$

It is strictly increasing, onto the real line, and has a continuously differentiable inverse. It induces an orientation-preserving circle diffeomorphism, as does every finite composition. With $x = \cos(\phi - \theta) \in [-1, 1]$, the derivative equals

$$D_r(x) = 1 + \frac{43}{100} r^2 - \frac{19}{50} rx - \frac{86}{100} r^2 x^2. \tag{17}$$

This strictly concave quadratic has vertex $x_* = -19/(86r)$. For $r \le 19/86$, the vertex is at or left of -1, so the interval maximum is $D_r(-1) = 1 + 19r/50 - 43r^2/100$. Otherwise the vertex is in the interval; completing the square gives $1 + 43r^2/100 + 361/8600$. The formulas agree at the boundary. Every $x \in [-1, 1]$ is attained by the cosine, so these are exact one-step maxima. The first branch is increasing on its interval since its derivative is $19/50 - 43r/50 > 0$; the second is increasing for $r > 0$. Thus

$$M(r) \le M(1) = 1 + \frac{43}{100} + \frac{361}{8600} = \frac{12659}{8600}. \tag{18}$$

Also $M(r) \ge 1$, because the mean of F_a' over a period is one. For the composition lift Ψ_a, the chain rule gives $0 < \Psi_a'(\theta) \le \prod_t M(r_t) = m(a)^{-1}$. Only the product upper bound is needed; separate maxima need not occur along a common trajectory.

Density minorization and normalized risk. Under uniform initial angle, the terminal density relative to normalized Haar measure is

$$p_a(e^{i\psi}) = \frac{1}{\Psi_a'(\Psi_a^{-1}(\psi))}. \tag{19}$$

Periodicity makes the formula independent of lifts; one-to-one change of variables over one period proves it. Hence $p_a \geq m(a)$ and $0 < m(a) \leq 1$. If $m(a) < 1$, decompose the terminal probability law as $m(a)\lambda + (1 - m(a))\nu_a$, where ν_a has density $(p_a - m(a))/(1 - m(a))$. Vanishing Haar harmonics imply, for $k = 1, \ldots, 4$,

$$|\mu_k(a)| = (1 - m(a)) \left| \int q^k \, d\nu_a(q) \right| \leq 1 - m(a). \tag{20}$$

If $m(a) = 1$, then $p_a = 1$ almost everywhere and the same inequality holds with both sides zero. Substitution into the exact identity yields $R_H^*(a) \geq 1 - (1 - m(a))^2 = 2m(a) - m(a)^2$. Since $m(a) \geq (8600/12659)^H$ and $2m - m^2$ is increasing on $[0, 1]$, the uniform bound follows. It is strictly positive at every finite horizon and tends to zero as the horizon grows; no horizon-independent positive floor is asserted. □

The conditional Bayes rule uses the action sequence alone but can exploit the nonuniform terminal law created by the dynamics; its risk need not be one for positive horizons. Object-specific calibration, a supplied canonical object, informative observed history, or a prior with externally fixed canonical angle changes the experiment. Additional observations of the same test object or parameters correlated with it require a different conditional analysis. No numerical floor for a realized finite C06 sample mean, attribution of all learned error, or general world-model impossibility follows.

H.6 Fixed-action numerical enclosures for the continuous law

The following computation evaluates the unrestricted risk in Theorem H.1 for 16 already retained H8 action sequences and their H0/H1/H2/H4 prefixes. It does not assign that theorem's prior to a realized floating-point test bank. Stored float64 actions are interpreted as exact binary rationals and checked in the full annulus. All 16 sequences and all five prefixes are retained; they are fixed conditional queries, not 16 iid scientific replications.

For a step with $a = A + iB$, write $x = A\cos\alpha + B\sin\alpha$, $y = B\cos\alpha - A\sin\alpha$, and $F(\alpha) = \alpha + \lambda_0 y + \kappa_0 xy$, with $\lambda_0 = 19/50$, $\kappa_0 = 43/100$. The previously established bounds are

$$|F'| \leq M = 1 + \kappa_0 r^2 + \lambda_0^2/(8\kappa_0), \qquad |F''| \leq B_2 = \lambda_0 r + 2\kappa_0 r^2. \tag{21}$$

Indeed $F' = 1 - \lambda_0 x + \kappa_0(y^2 - x^2) \geq 19/100$; completing the square gives the upper bound, and $F'' = -\lambda_0 y - 4\kappa_0 xy$ with $2|xy| \leq r^2$ gives the second. For the composition, set $D_0 = 1, E_0 = 0$ and update $D_{h+1} = M_h D_h$, $E_{h+1} = B_{2,h} D_h^2 + M_h E_h$. The chain rule bounds its first two derivatives by D_h, E_h. Each real component of $\exp(ik\Psi_h)$ has second derivative bounded by $C_{kh} = kE_h + k^2 D_h^2$. Midpoint Taylor expansion cancels its linear term and bounds the integral remainder on a cell of length Δ by $C_{kh}\Delta^3/24$. Consequently the normalized composite-midpoint error is

$$e_{kh} = \frac{(2\pi)^2}{24N^2}(kE_h + k^2 D_h^2). \tag{22}$$

The computation fixes $N = 8192$ and 96-bit directed arithmetic, encloses midpoints and every flow/trigonometric operation, then adds $[-e_{kh}, e_{kh}]$ to each component. Squared-moment intervals propagate through the exact risk identity; rational endpoints remain authoritative. Numerical intervals are deterministic enclosures, not statistical confidence intervals.

The analytic H8 lower bound $\simeq .088686$ is neither the exact optimum nor the reported numerical mean. Individual risks need not decrease with horizon: sequence 01 rises from H4 $[.94205799, .94205976]$ to H8 $[.97651739, .97652399]$. These values do not measure the excess error of a trained model and do not apply to a calibrated or history-augmented task.

H.7 Remaining numerical and reproduction limits

Two bounded implementations of the quantized ideal ellipse/K1 consumer are reported separately: the information-preserving predecessor Q and the targeted successor S. They use the same law,

Table 15: Fixed-library risk summaries. Displayed endpoints are rounded outwards. The range encloses all 16 risks; the mean is the equal-weight mean of these fixed queries, not an action-population estimate. All 64 positive-horizon widths meet the predeclared .002 target.

H	Enclosing range	Mean enclosure	Maximum raw width
0	[1,1]	[1,1]	0
1	[.96476801,.99890692]	[.98484780,.98484790]	.000000223361
2	[.90558982,.99866902]	[.97049820,.97049846]	.000000823405
4	[.89354267,.98382293]	[.94735200,.94735285]	.000002194257
8	[.73704901,.98708800]	[.90285044,.90285799]	.000033883106

finite-library task and 59 previously seen engineering declarations, comprising 55 positive-horizon declarations and four H0 controls. Resources were fixed separately before each run. Here excess is the conditional expected $[0, 4]$-valued loss of a fixed comparator minus the minimum conditional expected loss over its declared finite library. The numerical target is an excess-enclosure width of at most $1/500$; a NONTRIVIAL decision instead certifies excess greater than $1/500$ under positive event evidence. Neither criterion establishes an operational floating-observation posterior.

Quantity	Q predecessor	S successor
Declared $H > 0$ / H0	55 / 4	55 / 4
Completed / hard timeout	59 / 0	56 / 3
Positive evidence $H > 0$ / H0	50 / 4	48 / 4
Unresolved / encoding-unresolved / zero	3 / 1 / 1	2 / 1 / 1
No accepted task interval	5	7
Positive $H > 0$ width $\leq 1/500$	0/50	0/48
ADEQUATE / NONTRIVIAL / INCONCLUSIVE	4 / 6 / 44	4 / 0 / 48
Minimum positive-horizon width (approx.)	0.12897325	2.48340986

All ADEQUATE decisions and width-target attainments were H0 controls. Q's six NONTRIVIAL certificates concern empty-label-mask or X0-only declarations (017, 024, 025, 032, 033 and 040), not successful labelled-calibration inference or independent confirmations. S returned INCONCLUSIVE for all six; this does not refute the earlier Q certificates.

S's unavailable task intervals comprise evidence-unresolved declarations 011/013, encoding-unresolved declaration 015, zero-evidence declaration 016, and hard timeouts 012/030/038. No internal result is accepted for the three timed-out cases. Q had positive evidence for 030/038; 012 was already unresolved in Q after an earlier implementation had certified positive evidence.

Among the 52 declarations with intervals in both versions, S returned 32 larger widths, 19 equal widths, and one smaller width (006); seven comparisons remain unavailable. The 19 equal widths include the four H0 controls. Both versions therefore missed the positive-horizon width target under their fixed budgets. These descriptive comparisons neither identify a cause of the adverse changes nor establish impossibility at other budgets. The retained enclosures remain conditional on the stated law and numerical-enclosure premises; no best-of-Q/S result is substituted.

All studies used local CPU computation. Completed training runs retained final checkpoints. The A1 records retain every declared terminal outcome, including unresolved declarations and hard timeouts; S's three timed-out cases have no accepted internal result. The matched one-step/sensitivity run took 28.19 seconds; four-arm training/evaluation took 64.19 seconds with three single-thread workers and 20,480,000 transition applications. The encoder-supervision run took 255.29 seconds with three workers and 216,000 Adam updates. These times exclude preparation and review and do not establish equal computation. The fixed continuous-law integration took 280.09 seconds with four single-thread workers. Separate end-to-end and per-component control task timings were not all recorded, so no complete efficiency claim is made.

Retained-array replay and independent reductions checked reported aggregates; they did not repeat optimizer training or generate independent scientific replications. Seed/object exclusion applies only to the accessible historical inventory. The anonymous supporting material provides the retained result tables, exact numerical intervals and aggregate-replay scope; it does not claim to contain every

checkpoint, multi-gigabyte trace or a full empirical rerun package. None of these results establishes uncalibrated operational identifiability, real-world transfer, general noisy-calibration sufficiency, or learned-model benchmark superiority.

I Reproducibility statement

All theorem-side counterexamples are finite and use rational probabilities. The proofs and finite conditional-law constructions are contained in this appendix. Outcome-blind specifications distinguish independent blocks from repeated method views. The anonymous supplement supplies source/protocol material, retained numerical summaries including adverse results, a relative-path manifest, and checks of the reported aggregate statistics. Its README distinguishes those checks from full experiment reproduction. Large learned-model checkpoints, raw tensor corpora and the full A1 certificate corpus are not included in the compact package; its replay does not retrain models, reproduce every forecast or establish an operational posterior. Numerical interval claims retain the law and implementation dependencies stated in Appendix H. For S, the anonymous materials contain all 59 declarations and a selected retained-summary dataset. The supplied reducer checks exact interval arithmetic, outcome counts and Q/S comparisons; it does not regenerate the enclosures, re-evaluate numerical primitives, reconstruct the full certificate trace, or establish an operational posterior.

J Statement on the use of AI tools

AI tools, including large language models, were used substantially as research assistants: to refine hypotheses and experimental designs, develop theoretical formulations and proof/counterexample arguments, implement synthetic-data generators and methods, interpret results, draft and refactor LaTeX and verification code, and perform adversarial consistency reviews. They were not treated as sources of mathematical or bibliographic authority. The manuscript supplies explicit proofs, exact finite constructions, primary-source citations and scoped empirical evidence rather than relying on model agreement. Responsibility for all claims, citations, code and reported results remains with the author.

K Ethics statement

This work develops mathematical results and controlled synthetic studies that use no human participants, personal data, external dataset, or deployed decision system. We identified no domain-specific ethical issue that requires an additional empirical analysis. The results may nevertheless inform representation-learning systems used in downstream applications; claims about safety, fairness, or reliability in such applications require separate evidence and are not made here.