OPERATIONAL MEMORY: PREDICTIVE SPECTRA AND REALIZATION

Anonymous authors
Paper under double-blind review

ABSTRACT

A state representation of a dynamical system is typically judged by how accurately it predicts future observations. In systems with memory, however, two histories with the same observed present can respond differently to the same intervention. A model may therefore predict trajectories accurately while discarding historical information that remains predictive of intervention responses. We call this information Operational Memory. We formalize this distinction and show that accurate trajectory prediction need not imply accurate prediction of response differences under intervention. To characterize Operational Memory, we introduce the Operational Predictive Memory Spectrum (OPMS), which decomposes the dependence of future intervention responses on history into predictive directions; a context-resolved form prevents cancellation when history–response relationships vary across matched contexts. We further develop Operation-Complete Predictive Realization (OCPR), which constructs a finite predictive state whose time and intervention updates can be identified and reused recursively. In two spin–boson environments, updates identified using responses with at most two interventions compose to predict responses after a third intervention, with Causal Skill 0.966/0.971. On superconducting qubit data, context resolution reveals dependence missed by the pooled spectrum; an independent quantum jump dataset supports the same principle.

1 INTRODUCTION

State representations of dynamical systems are commonly evaluated by their ability to predict future observations. For systems with memory, this criterion can be incomplete. Two histories can reach the same observed present and follow nearly identical unperturbed trajectories, yet respond differently to the same intervention. The difference reflects degrees of freedom that are unresolved by the observed present and remain coupled to future dynamics. Classical projection methods describe such dependence through memory kernels (Zwanzig, 1961; Mori, 1965), while process tensors describe future statistics as functions of intervention histories (Pollock et al., 2018a;b; Milz & Modi, 2021; Keeling et al., 2026). Modern recurrent and state-space models similarly compress history for future prediction (Gu et al., 2020; 2022; Gupta et al., 2022; Smith et al., 2023; Gu & Dao, 2024; Dao & Gu, 2024; Buitrago Ruiz et al., 2025). This raises a central question: does a learned state preserve the historical information that remains predictive under intervention?

We call this information Operational Memory. Consider two histories in the same matched context: they share the observed present and all other conditions fixed for the comparison, while differing in their preceding histories. Figure 1 depicts one future test, comprising interventions and readouts, applied to both histories. A difference in their responses then reflects information carried by history after the matched context has been fixed. We therefore define operational predictive equivalence with respect to a specified family of future tests: two matched histories are equivalent when every future test in the family produces the same response. We show that trajectory prediction and recovery of these response differences can diverge sharply. In particular, trajectory R^2 can approach one while recovery of response differences under intervention becomes arbitrarily poor. Operational Memory therefore captures predictive information that trajectory accuracy alone does not certify.

To characterize Operational Memory, we introduce the Operational Predictive Memory Spectrum (OPMS), which decomposes the dependence of future intervention responses on history into predic-

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.

tive directions. The spectrum describes the strength, accessible rank, and directional structure of this dependence. The history–response relationship can vary across matched contexts. In the pooled OPMS, these context-specific relationships are aggregated before the spectrum is formed, so differently oriented dependencies can cancel. This motivates a context-resolved OPMS that preserves the response structure within each context before aggregation. OPMS complements process-level operational notions of memory (Taranto et al., 2021) and instrument-specific notions of quantum memory (Taranto et al., 2019b; Guo et al., 2021) by resolving the predictive directions accessible under fixed observation and intervention conditions.

OPMS characterizes this predictive structure; recursive prediction asks whether it can be represented and updated through time and intervention. We therefore develop Operation-Complete Predictive Realization (OCPR), which constructs a finite predictive state with shared updates for free evolution and intervention. To make the realization operation-complete within a specified finite-dimensional intervention family, OCPR represents every intervention in exact coordinates over a fixed basis. We combine these coordinates with a generalized Hankel realization built from the centers and response differences of matched histories. The resulting center–gap Hankel preserves the common response structure needed for state prediction together with the directions that govern how matched histories separate under intervention. These shared updates allow the same predictive state to be propagated recursively through sequences of free evolution and intervention.

To test this recursive reuse, we separate the intervention depth used for identification from the depth used for evaluation. In the spin–boson experiments, intervention updates are identified from responses containing at most two interventions, while responses after a third intervention form the held-out evaluation. Predictions after the third intervention are produced by composing the same identified updates. We establish conditions for identifying these updates and derive a perturbation bound showing how restricted response geometry controls finite-horizon prediction error. We also evaluate state access at two levels: states obtained directly from response statistics isolate realization quality, while inference from observable history measures the additional state-identification problem.

OPMS is tested on controlled latent-memory systems, generalized Langevin dynamics, two spin–boson environments, and an additional PDE setting. OCPR is evaluated in the two spin–boson environments, with a non-quantum latent-memory system providing a separate realization test. In the spin–boson experiments, updates identified using responses with at most two interventions compose to predict responses after a third intervention with Causal Skill $0.966/0.971$, where a score of one corresponds to exact recovery of the response differences. The non-quantum system reproduces the same depth-two-to-three composition result to numerical precision. On released superconducting qubit data (Giarmatzi et al., 2025), the context-resolved OPMS reveals history–response dependence missed by the pooled OPMS; an independent quantum jump dataset (Gosling et al., 2026b) supports the same context-resolution principle under a different measurement setting. The hardware studies provide evidence for OPMS and context resolution, while recursive OCPR identification is evaluated in the controlled systems. Our contributions are as follows:

- We formalize a matched-context notion of Operational Memory as historical information that remains predictive of intervention responses after the present and other comparison conditions are matched. Through operational predictive equivalence, we show that arbitrarily accurate trajectory prediction can coexist with arbitrarily poor recovery of response differences under intervention.
- We develop the Operational Predictive Memory Spectrum (OPMS), which resolves this dependence into predictive directions. Its context-resolved form preserves context-specific history–response relationships and prevents cross-context cancellation.
- We develop Operation-Complete Predictive Realization (OCPR), which constructs a finite predictive state with updates identified from limited intervention depth. We characterize the restricted rank condition for identification, show that the full singular spectrum does not certify update stability, and give perturbation guarantees for recursive composition.
- We evaluate these ideas across controlled quantum and non-quantum systems and two independent superconducting datasets. Updates identified through depth two compose accurately to a third intervention in the controlled systems, while the hardware data independently support the context-resolution principle underlying OPMS.

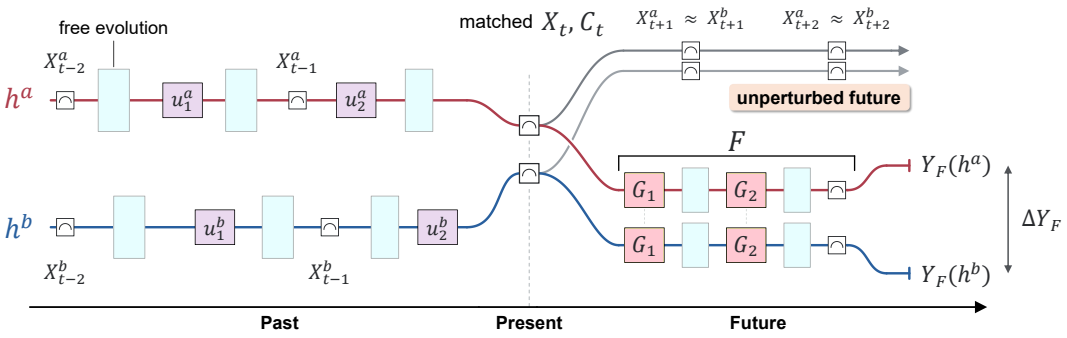

Figure 1: **Operational Memory under matched observation and intervention.** Two different histories can share the same observed present X_t and matched context C_t, and have nearly identical unperturbed futures, yet respond differently to the same future intervention test F. The resulting response difference ΔY_F exposes historical information that remains predictive under the fixed matched context.

2 Related Work

Operational notions of memory and predictive-state methods. Process-tensor theory represents a quantum process as a multilinear map from intervention sequences to outcome statistics (Pollock et al., 2018a;b; Milz & Modi, 2021), while quantum Markov order and instrument-specific memory make the role of the probing instrument explicit (Taranto et al., 2019a;b; Guo et al., 2021). Taranto et al. (2021) define process-level memory strength by optimizing future–history dependence over probing instruments. OPMS instead fixes the observation, matched context, intervention family, and future tests, and resolves the history–response dependence into predictive directions.

Predictive-state representations define state through future tests (Littman et al., 2001; Singh et al., 2004; Boots et al., 2011); observable operator models provide related finite-dimensional realizations (Jaeger, 2000), while ϵ-transducers group input–output histories by future output statistics conditioned on future inputs (Barnett & Crutchfield, 2015). These notions are closely related to operational predictive equivalence; our comparison is restricted to histories in the same matched context and to a specified family of future tests. Canonical variate analysis and spectral or supervised PSRs recover predictive coordinates from past–future relations (Larimore, 1990; Hefny et al., 2015; Sun et al., 2016), while recurrent predictive-state models reuse such states recursively (Downey et al., 2017; Hefny et al., 2018). Conditional CCA studies associations that vary with observed covariates (Cai & Sun, 2016). OPMS uses this canonical geometry to resolve matched-context history–response dependence, while its context-resolved form prevents cross-context cancellation. OCPR builds on classical Hankel and subspace realization (Ho & Kalman, 1966; Kung, 1978; Juang & Pappa, 1985; Van Overschee & De Moor, 1994), but identifies recursive updates from limited intervention depth. Core test and sparse sensing methods provide related ideas for reducing the response set (Littman et al., 2001; McCracken & Bowling, 2005; Hamilton et al., 2013; Huang et al., 2018; Manohar et al., 2018).

Learning non-Markovian dynamics. Tensor-network methods compress influence functionals and process tensors (Strathearn et al., 2018; Jørgensen & Pollock, 2019; Guo et al., 2020; Keeling et al., 2026), while process-tensor tomography can exploit low-memory structure (White et al., 2022); related work develops recursive auxiliary-memory models or compact temporal representations (Zhang et al., 2025; Luchnikov et al., 2026). Superconducting experiments have reconstructed multi-time processes or exposed long-lived environmental memory (Giarmatzi et al., 2025; Gosling et al., 2026b). In machine learning, structured state-space models and memory neural operators compress history for future prediction (Gu et al., 2022; Gu & Dao, 2024; Dao & Gu, 2024; Buitrago Ruiz et al., 2025). OCPR focuses on a finite predictive state whose intervention responses and updates are explicitly identified. We also compare with the open quantum evolution (OQE)/system–memory model of Zhang et al. (2025) as an external recursive baseline; Appendix G describes the adaptation and information access used for this comparison.

3 OPERATIONAL MEMORY AND ITS PREDICTIVE SPECTRUM

Let $H_t = (X_{\leq t}, U_{<t})$ denote an observable history, C_t the *matched context*, and F a *future test* specifying future interventions and readouts. Define

$$Y_F(H_t) = \mathbb{E}[f(X_{>t}) \mid H_t, F], \qquad \Delta Y_F = Y_F(H_t^a) - Y_F(H_t^b), \tag{1}$$

for histories H_t^a, H_t^b in the same matched context. We call historical information that remains predictive of these response differences *Operational Memory*. In the quantum experiments, a replacement intervention fixes the observed system state while preserving environmental history that can influence later dynamics; Section 4.1 gives its form, and Appendix C.1 gives the corresponding constructions for the other controlled processes.

Definition 1 (Operational predictive equivalence)**.** *Fix a matched-context map $m : \mathcal{H} \to \mathcal{C}$, reachable histories $\mathcal{H}$, and a family of future tests $\mathcal{T}$. Histories $H, H' \in \mathcal{H}$ are equivalent, written $H \equiv_{\mathcal{T}} H'$, when $m(H) = m(H')$ and $Y_\tau(H) = Y_\tau(H')$ for every $\tau \in \mathcal{T}$.*

Operational Memory is present when two reachable histories in the same matched context are not equivalent. The test family determines which historical distinctions must be preserved.

For matched-history pairs with nonzero total response-difference energy, we use

$$\text{Causal Skill} = 1 - \frac{\sum_n \|\Delta\widehat{Y}^n - \Delta Y^n\|_2^2}{\sum_n \|\Delta Y^n\|_2^2}, \tag{2}$$

where 1 is exact recovery and 0 is the zero-response-difference predictor.

Proposition 1 (Trajectory accuracy and response-difference recovery can diverge arbitrarily)**.** *For every $\eta > 0$ and $M \geq 0$, there exists a matched response family and a predictor with $R^2 > 1 - \eta$ but Causal Skill $= -M$.*

An $O(1)$ common component can dominate ordinary R^2 while an $O(\epsilon)$ response difference is lost; Appendix A gives the construction.

3.1 OPERATIONAL PREDICTIVE MEMORY SPECTRUM

Let $P_0 = \phi(H_t)$ be history features and $Q_0 = \psi(F, X_{>t})$ future-response features, and residualize both with respect to C_t: $P = P_0 - \mathbb{E}[P_0 \mid C_t]$ and $Q = Q_0 - \mathbb{E}[Q_0 \mid C_t]$. Following canonical-correlation geometry (Hotelling, 1936), for positive-definite residual covariances define

$$M_{\text{OP}} = \Sigma_{QQ}^{-1/2}\Sigma_{QP}\Sigma_{PP}^{-1/2} = U\,\text{diag}(\rho_1, \rho_2, \ldots)V^\top, \qquad S_{\text{OP}} = \sum_i \rho_i^2. \tag{3}$$

The singular values $\{\rho_i\}$ form the population *Operational Predictive Memory Spectrum (OPMS)*; S_{OP} summarizes total standardized linear dependence, while the spectrum retains accessible rank and directional structure in the chosen whitened feature coordinates. Finite-sample estimation is described in Appendix C.2.

Proposition 2 (Linear interpretation of OPMS)**.** *The population OPMS is invariant under invertible linear reparameterizations of P and Q, and*

$$\text{rank}(M_{\text{OP}}) = \text{rank}(\Sigma_{QP}).$$

Let $\bar{P} = \Sigma_{PP}^{-1/2}P$, $\bar{Q} = \Sigma_{QQ}^{-1/2}Q$, and $q = \dim(Q)$. The minimum rank-k linear prediction error is

$$\min_{\text{rank}(L)\leq k} \mathbb{E}\|\bar{Q} - L\bar{P}\|_2^2 = q - \sum_{i=1}^{k} \rho_i^2.$$

Thus the cumulative spectrum gives the standardized future-response variance recoverable by a k-dimensional linear predictive state. The feature maps determine which linear dependence is accessible; proofs and the Gaussian mutual-information specialization are given in Appendix A.

Pooled and context-resolved OPMS. For repeated context strata β, let $\pi_\beta = \Pr(\beta)$ and let e_β denote the corresponding one-hot basis vector. Define $P^\circ = P_0 - \mathbb{E}[P_0 \mid \beta]$, $Q^\circ = Q_0 - \mathbb{E}[Q_0 \mid \beta]$, and $\Sigma_{QP|\beta} := \mathrm{Cov}(Q^\circ, P^\circ \mid \beta)$. The pooled cross-covariance is $\sum_\beta \pi_\beta \Sigma_{QP|\beta}$ and can cancel when context-specific history–response maps have different orientations. We therefore define

$$Q_{\mathrm{CR}} = e_\beta \otimes Q^\circ, \qquad \Sigma_{Q_{\mathrm{CR}}P^\circ} = \mathrm{stack}_\beta[\pi_\beta \Sigma_{QP|\beta}], \tag{4}$$

and obtain the context-resolved spectrum by whitening Q_{CR} and P° as in Equation (3); let S_{CR} denote its squared singular-value sum.

Proposition 3 (Context resolution prevents cross-context cancellation)**.** *If every retained stratum has $\pi_\beta > 0$ and the population covariance matrices are nonsingular, then*

$$S_{\mathrm{CR}} = 0 \quad \Longleftrightarrow \quad \Sigma_{QP|\beta} = 0 \textit{ for every } \beta,$$

whereas the pooled OPMS can vanish even when every conditional cross-covariance is nonzero.

Appendix A gives the proof and a controlled cancellation construction. If selected future-response features satisfy $Q = Cz + \epsilon$ for a d-dimensional linear predictive state with $\mathrm{Cov}(\epsilon, P) = 0$, then $\mathrm{rank}(M_{\mathrm{OP}}) \le d$. OPMS rank can therefore lower-bound linear realization order, while recursive closure may require additional state directions.

4 Operation-complete predictive realization

OCPR constructs a finite predictive state that can be updated under free evolution and intervention. It combines exact coordinates for a specified intervention family, a center–gap response realization, and shared updates identified from responses within a fixed intervention depth; the construction is detailed below.

4.1 Intervention coordinates and center–gap realization

Let $\mathcal{A}$ be the studied intervention family and $\mathcal{V} = \mathrm{span}(\mathcal{A})$ a $d_{\mathcal{A}}$-dimensional linear space with basis $\{B_i\}_{i=1}^{d_{\mathcal{A}}}$. For $G \in \mathcal{A}$, $G = \sum_i c_i(G) B_i$. By the standard multilinearity of multi-time process responses (Pollock et al., 2018a; Milz & Modi, 2021),

$$Y(G, G' \mid h) = \sum_{i,j} c_i(G) c_j(G') \Theta_{ij}(h), \qquad \Theta_{ij}(h) = Y(B_i, B_j \mid h). \tag{5}$$

Equation (5) therefore represents responses for all interventions in the specified family exactly. In the spin–boson experiments, replacement channels form a four-dimensional span:

$$\mathcal{R}_\varrho(X) = \varrho \, \mathrm{Tr}(X), \qquad \mathcal{R}_\varrho = \sum_{i=1}^{4} c_i(\varrho) B_i, \qquad \sum_i c_i(\varrho) = 1. \tag{6}$$

Further coordinate and error-propagation details are given in the appendix.

Let $\mathcal{F} = \{F_1, \ldots, F_D\}$ be a bank of future tests and $y(h) = [Y_{F_1}(h), \ldots, Y_{F_D}(h)]$. For a matched pair h^a, h^b, define

$$y_c = \tfrac{1}{2}[y(h^a) + y(h^b)], \qquad y_g = y(h^a) - y(h^b). \tag{7}$$

After separately normalizing the stacked center and response-difference blocks by their training RMS values, write

$$H = \begin{bmatrix} \bar{H}_c \\ \bar{H}_g \end{bmatrix} \approx LR, \qquad L = U_d \Sigma_d^{1/2}, \quad R = \Sigma_d^{1/2} V_d^\top, \qquad z(h) = y(h) V_d \Sigma_d^{-1/2}.$$

Before truncation, the normalized center–gap matrix is an invertible row transformation of the stacked individual-response matrix, so it preserves the exact response row space and any shared update closure; Appendix B gives the formal statement. In the main realization, the two normalized blocks are weighted equally; replacing $\bar{H}_g$ by $r\bar{H}_g$ defines the weighting sensitivity studied in Section 5.

4.2 DEPTH-RESTRICTED UPDATE IDENTIFICATION

Let L_{id} be the maximum intervention depth used to identify updates. For an update symbol σ, define

$$J_\sigma = \{s : \mathrm{depth}(\sigma s) \le L_{\mathrm{id}}\}, \tag{8}$$

$$A_\sigma = L^+ H^\sigma_{J_\sigma} (R_{J_\sigma})^+, \tag{9}$$

where $H^\sigma_{J_\sigma}[h, s] = H[h, \sigma s]$ uses the same row construction and normalization as H.

Proposition 4 (Exact depth-restricted update identification)**.** *If $H = LR$ is an exact rank-d realization, $H^\sigma_{J_\sigma} = LA_\sigma R_{J_\sigma}$, L has column rank d, and R_{J_σ} has row rank d, then A_σ is unique and Equation (9) recovers it.*

Under exact closure,

$$z(h\sigma_1 \cdots \sigma_k) = z(h) A_{\sigma_1} \cdots A_{\sigma_k},$$

so sequences deeper than L_{id} test composition of the same identified updates.

For the balanced factorization, define

$$\gamma_{\mathrm{shift}} = \min_\sigma \sigma_{\min}(R_{J_\sigma}), \qquad \delta_d(H) = \sigma_d(H) - \sigma_{d+1}(H). \tag{10}$$

A positive γ_{shift} is the restricted rank condition required above; it controls pseudoinverse amplification, which the full singular spectrum alone cannot, while δ_d controls subspace perturbation. Appendix B gives the perturbation and finite-horizon bounds.

Definition 2 (Depth-restricted response core)**.** *A coordinate subset S is a response core for a rank-d realization if $\mathrm{rank}(R_S) = d$ and $\mathrm{rank}(R_{S \cap J_\sigma}) = d$ for every identified update σ.*

In the exact rank-d setting, this condition preserves realization order and update identifiability; selection, conditioning, and measurement cost in the truncated setting are analyzed in Appendix F.2. Under exact update closure, multilinearity and uniqueness imply

$$A_{\mathcal{R}_\varrho} = \sum_i c_i(\varrho) A_{B_i}. \tag{11}$$

For the truncated empirical realization, we define $\widehat{A}_{\mathcal{R}_\varrho} := \sum_i c_i(\varrho) \widehat{A}_{B_i}$.

4.3 PREDICTIVE STATE INFERENCE FROM OBSERVABLE HISTORY

The realization above uses controlled future-response statistics. To access the same state from observable history, a temporal encoder maps the 64-step history to features $\phi(h) \in \mathbb{R}^{64}$, followed by an affine map $\widehat{z}(h) = [\phi(h); 1]W$. We fit W using individual-state, center, and response-difference terms over matched histories, with training-set normalization and ridge regularization. Full objectives and regularization details are given in the appendix. We report *response-derived predictive states*, which isolate realization quality, separately from *history-inferred predictive states*, which additionally test state inference.

5 EXPERIMENTS

5.1 EVALUATION SETTINGS

OPMS is evaluated on controlled latent-memory processes, a generalized Langevin process, two spin–boson environments, and an additional PDE setting in Appendix D.3. OCPR is studied primarily in the two spin–boson environments. The primary depth-three evaluation uses 30 held-out matched-history pairs per environment. The main realization comparisons use the same 504 scalar future-response coordinates, exact intervention coordinates, and a common state dimension $d = 16$, fixed independently of the held-out horizons and depth-three evaluation. Model selection uses only $\Delta \in \{2, 4, 8\}$; the held-out horizons and three-intervention endpoints are reserved for evaluation. Causal Skill is the primary response-difference metric, with trajectory R^2 reported separately. We also evaluate the OQE/system–memory model of Zhang et al. (2025) using response-derived identification. Its shared system–memory dynamics are selected over memory dimension $\chi \in \{1, \ldots, 6\}$ with five initializations per χ using development ordinary-response MSE only; $\Delta = 16$ and depth-three responses are excluded from fitting and state access. The absolute-response Hankel provides the realization comparison with matched response access, whereas OQE is treated as an external method comparison. Split sizes and further evaluation details are given in the appendix.

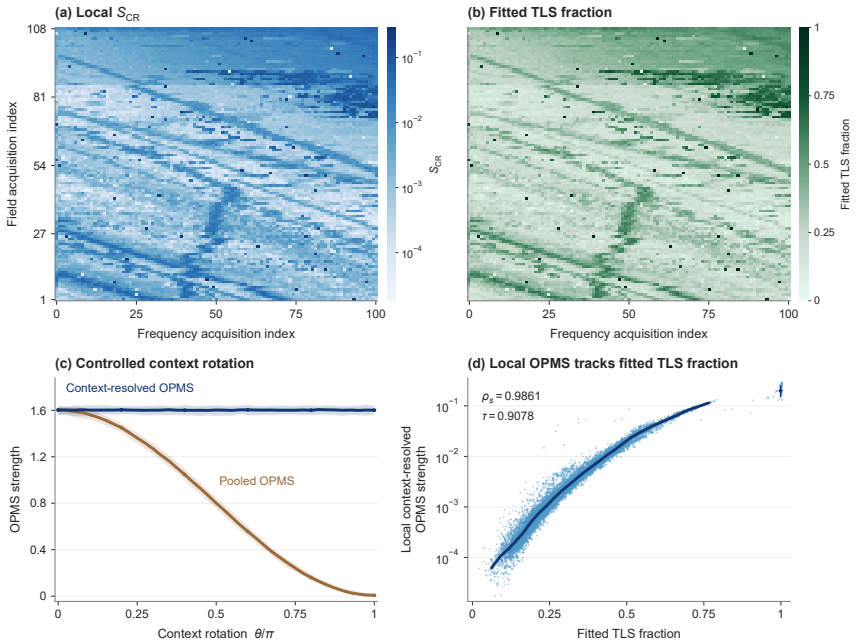

Figure 2: **Context resolution reveals structured Operational Memory.** **(a,b)** Local context-resolved OPMS strength and fitted two-level-system (TLS) fraction on the Gosling acquisition grid. **(c)** Context rotation suppresses pooled but not context-resolved OPMS. **(d)** Local OPMS strength tracks the fitted TLS fraction.

Table 1: **OPMS results across controlled and hardware settings.** Hardware rows report observed statistics relative to matched-null thresholds.

Evaluation	Result
Memoryless / latent 1-channel / latent 2-channel	$S_{\mathrm{OP}} \approx 0/0.994/1.970$
Oscillatory generalized Langevin	$S_{\mathrm{OP}} = 1.976$
Spin–boson, intermediate / structured	$S_{\mathrm{OP}} = 1.154/2.076$
Controlled conditions	Spearman $\rho = 0.695$, $p = 0.00370$
Giarmatzi, pooled / context-resolved	$0.0295 < 0.0888$ / $\mathbf{1.438 > 0.562}$
Gosling, context-resolved	$\mathbf{0.006256 > 0.000468}$; $\rho_s = 0.9861$, $\tau = 0.9078$

5.2 OPMS: MEMORY DIRECTIONS AND CONTEXT RESOLUTION

OPMS recovers the expected low-dimensional structure in the controlled systems: the memoryless control gives $S_{\mathrm{OP}} \approx 0$; the one- and two-channel latent-memory systems give $0.994/1.970$; the oscillatory generalized Langevin system gives 1.976; and the intermediate/structured spin–boson environments give $1.154/2.076$ with accessible ranks $2/3$. Across 20 controlled conditions, S_{OP} correlates with independently measured response-difference magnitude (domain-blocked Spearman $\rho = 0.695$, $p = 0.00370$). In the controlled context-rotation sweep, pooled OPMS decreases as the context-specific response maps become oppositely oriented, whereas context-resolved OPMS remains stable (Figure 2c).

The superconducting datasets further support context resolution. On the dataset of Giarmatzi et al. (2025), the pooled Bloch OPMS remains below its matched null threshold, whereas the context-resolved statistic is above its null. On the independent Gosling quantum jump dataset (Gosling et al., 2026b), context-resolved OPMS also exceeds its matched null. Across all $10{,}908$ retained Gosling contexts, local S_{CR} closely tracks the separately fitted TLS fraction (Spearman 0.9861, Kendall 0.9078; Figure 2d). This strong within-dataset correspondence supports the context-resolution interpretation, while the Giarmatzi and Gosling datasets provide independent tests. Table 1 summarizes the main quantitative OPMS results.

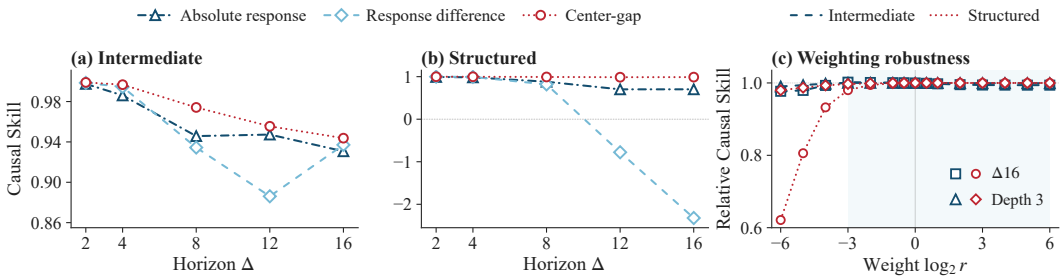

Figure 3: **Center–gap realization and weighting robustness.** **(a,b)** Causal Skill for absolute-response, response-difference, and center–gap realizations across horizons. **(c)** Center–gap sensitivity to r at $\Delta = 16$ and depth three, relative to $r = 1$.

5.3 OCPR COMPONENTS

Exact intervention coordinates. The replacement family and ordered two-intervention frame have full ranks 4 and 16, respectively, and direct multilinear contraction reconstructs held-out replacement responses to numerical precision (Appendix E.1). With basis responses predicted from observable history, exact coordinate contraction gives trajectory $R^2 = 0.999988$ and Causal Skill 0.959 on held-out replacements; generic continuous conditioning gives $R^2 = 0.953$ but Causal Skill $= -0.534$, and shuffling the physical coordinates reduces Causal Skill to -2.05.

Center–gap realization. At $\Delta = 16$, center–gap attains the highest Causal Skill among the three response-derived realizations in both environments, with the largest gain in the structured environment (Table 2; Figure 3). The result is stable over $r \in [2^{-3}, 2^6]$: all tested $\Delta = 16$ and depth-three endpoints retain at least 95% of their $r = 1$ performance. Center–gap also increases the restricted update margin γ_{shift} by $4.38\times$ / $4.84\times$, directly improving the conditioning relevant to update identification.

5.4 STATE INFERENCE AND RECURSIVE REUSE

The state-inference ablation shows that ordinary trajectory accuracy alone does not determine response-difference recovery. At $\Delta = 16$, all three variants that infer state from history have trajectory $R^2 > 0.9996$, yet their Causal Skill values differ substantially: direct state inference accounts for most of the gain over response projection, and the paired center–gap objective gives a further improvement in the intermediate environment (Table 2).

The same updates remain effective beyond the selection horizon and identification depth. On 30 held-out matched-history pairs per environment, the updates identified from responses with at most two interventions are composed through a third held-out intervention, giving history-inferred Causal Skill $0.966/0.971$ and response-derived Causal Skill $0.993/0.998$ in the intermediate/structured environments. Figure 4 shows coordinate-level agreement; the appendix reports additional compositional diagnostics.

External system–memory comparison. An adaptation of the OQE/system–memory model of Zhang et al. (2025) retains substantial ordinary-response fit but does not recover the smaller matched-history response differences. At $\Delta = 16$, OQE reaches trajectory $R^2 = 0.891/0.811$ but Causal Skill $= -74.30/-23.67$; after depth-two-to-three composition, $R^2 = 0.915/0.875$ while Causal Skill remains $-10.15/-1.06$ in the intermediate/structured environments. Across these four endpoints, response-difference RMSE is 0.0117–0.0195, correlation is 0.196–0.461, and multiplicative recovery is 1.39–8.90, so the negative Causal Skill is not explained by normalization alone. On the separate response-derived evaluation used for the OQE comparison, center–gap OCPR reaches Causal Skill $0.958/0.979$ at $\Delta = 16$ and $0.994/0.997$ at depth three. Appendix G gives the adaptation, information-access constraints, and full comparison.

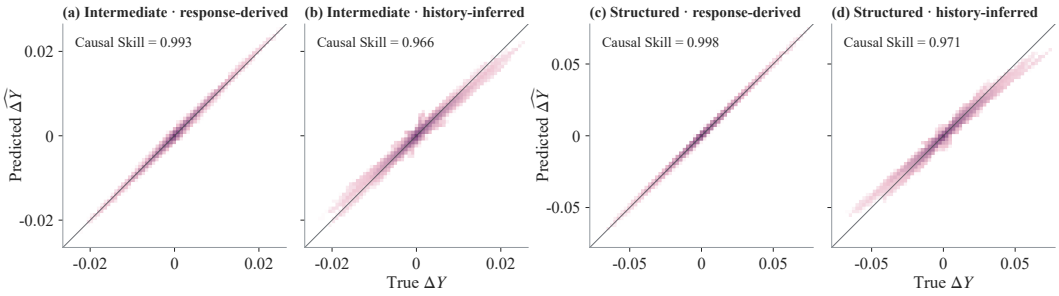

Figure 4: **Depth-three coordinate agreement.** Predicted versus true response differences on the 30-pair held-out evaluation for response-derived and history-inferred state access.

Table 2: **OCPR realization, state-access, and depth-three Causal Skill.** The first two blocks use a fixed 10-pair held-out subset per environment; the depth-three block uses the full 30-pair held-out evaluation set.

Evaluation	Method	Intermediate	Structured
$\Delta = 16$, realization	Absolute	0.931	0.703
	Difference	0.937	−2.322
	Center–gap	**0.944**	**0.989**
$\Delta = 16$, state access	Projection	0.501	0.700
	Direct state	0.892	0.959
	Paired center–gap	**0.912**	**0.961**
Depth three	Response-derived	0.993	0.998
	History-inferred	0.966	0.971

A non-quantum latent-memory control uses $d = 4$, matching its exact affine state $[1, x, z_1, z_2]$, and composes through all 6^3 held-out intervention triples with Causal Skill 1.000000, providing an exact non-quantum test of the same depth-two-to-three realization principle.

5.5 A COMPACT RESPONSE CORE REDUCES IDENTIFICATION COST

For the fixed $d = 16$ realization, a 16-coordinate response core retains rank 16 in the base realization and every update-identification block. Refitting the full and reduced response sets separately yields closely matched Causal Skill at $\Delta = 16$ and depth three. With the response core, the number of distinct scalar response expectations decreases from $391{,}800$ to $20{,}848$ per environment, a factor of $18.79\times$; comparisons at equal total budget also favor the compact response set (Appendix F.2).

6 DISCUSSION AND CONCLUSION

Operational Memory is historical information that remains predictive of intervention responses after the matched context is fixed. OPMS resolves this dependence into predictive directions, and context resolution prevents cross-context cancellation. OCPR constructs a finite predictive state with updates identified at bounded depth and reused recursively. Controlled quantum and non-quantum systems support depth-two-to-three composition, while superconducting datasets support the context-resolution principle. The OQE comparison shows that substantial ordinary-response fit need not preserve matched-history intervention responses under the tested objective and capacity range; it does not imply a generic limitation of learned auxiliary memory models. A compact response core reduces identification cost, while broader coherent interventions expose a remaining state-inference bottleneck.

AI USE STATEMENT

Generative AI tools were used for language polishing and assistance with writing portions of the code. The authors reviewed all AI-assisted material and take full responsibility for the manuscript, code, and scientific content.

REPRODUCIBILITY STATEMENT

Theoretical derivations and proofs are provided in Appendices A and B. Appendix C describes the main estimation, simulation, data-split, and state-inference procedures, while Appendices D–H provide additional experiments, baselines, robustness analyses, and computational characteristics.

REFERENCES

Shaojie Bai, J. Zico Kolter, and Vladlen Koltun. An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. *arXiv preprint arXiv:1803.01271*, 2018.

Nix Barnett and James P. Crutchfield. Computational mechanics of input–output processes: Structured transformations and the ϵ-transducer. *Journal of Statistical Physics*, 161(2):404–451, 2015. doi: 10.1007/s10955-015-1327-5.

Yoav Benjamini and Yosef Hochberg. Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: Series B*, 57(1):289–300, 1995.

Byron Boots, Sajid M. Siddiqi, and Geoffrey J. Gordon. Closing the learning-planning loop with predictive state representations. *The International Journal of Robotics Research*, 30(7):954–966, 2011. doi: 10.1177/0278364911404092.

Ricardo Buitrago Ruiz, Tanya Marwah, Albert Gu, and Andrej Risteski. On the benefits of memory for modeling time-dependent PDEs. In *International Conference on Learning Representations*, 2025.

Jia Cai and Hongwei Sun. Kernel-based conditional canonical correlation analysis via modified Tikhonov regularization. *Applied and Computational Harmonic Analysis*, 41(3):692–712, 2016. doi: 10.1016/j.acha.2015.04.006.

Tri Dao and Albert Gu. Transformers are SSMs: Generalized models and efficient algorithms through structured state space duality. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 10041–10071, 2024.

Carlton Downey, Ahmed Hefny, Boyue Li, Byron Boots, and Geoffrey J. Gordon. Predictive state recurrent neural networks. In *Advances in Neural Information Processing Systems 30*, 2017.

Gerald E. Fux, Piper Fowler-Wright, Joel Beckles, Eoin P. Butler, Paul R. Eastham, Dominic Gribben, Jonathan Keeling, Dainius Kilda, Peter Kirton, Ewen D. C. Lawrence, Brendon W. Lovett, Eoin O'Neill, Aidan Strathearn, and Roosmarijn de Wit. OQuPy: A Python package to efficiently simulate non-Markovian open quantum systems with process tensors. *The Journal of Chemical Physics*, 161(12):124108, 2024. doi: 10.1063/5.0225367.

Christina Giarmatzi, Tyler Jones, Alexei Gilchrist, Prasanna Pakkiam, Arkady Fedorov, and Fabio Costa. Multi-time quantum process tomography on a superconducting qubit. *Quantum*, 9:1952, 2025. doi: 10.22331/q-2025-12-18-1952.

Nicolas Gosling, Denis Bénâtre, Paul Kugler, Simon Geisert, Nicolas Zapata Gonzalez, Mitchell Field, Sumeru Hazra, Simon Günzler, Thomas Reisinger, Martin Spiecker, Mathieu Féchant, and Ioan Pop. Probing the memory of a superconducting qubit environment - main analysis code and data. Zenodo record 21908500, 2026a. Version v1, August 2026.

Nicolas Gosling, Denis Bénâtre, Nicolas Zapata, Paul Kugler, Mitchell Field, Sumeru Hazra, Simon Günzler, Thomas Reisinger, Martin Spiecker, Mathieu Féchant, and Ioan M. Pop. Probing the memory of a superconducting qubit environment. *arXiv preprint arXiv:2603.11889*, 2026b.

Albert Gu and Tri Dao. Mamba: Linear-time sequence modeling with selective state spaces. In *First Conference on Language Modeling*, 2024. URL https://openreview.net/forum?id=tEYskw1VY2.

Albert Gu, Tri Dao, Stefano Ermon, Atri Rudra, and Christopher Ré. HiPPO: Recurrent memory with optimal polynomial projections. In *Advances in Neural Information Processing Systems 33*, pp. 1474–1487, 2020.

Albert Gu, Karan Goel, and Christopher Ré. Efficiently modeling long sequences with structured state spaces. In *International Conference on Learning Representations*, 2022.

Chu Guo, Kavan Modi, and Dario Poletti. Tensor-network-based machine learning of non-Markovian quantum processes. *Physical Review A*, 102(6):062414, 2020. doi: 10.1103/PhysRevA.102.062414.

Yu Guo, Philip Taranto, Bi-Heng Liu, Xiao-Min Hu, Yun-Feng Huang, Chuan-Feng Li, and Guang-Can Guo. Experimental demonstration of instrument-specific quantum memory effects and non-Markovian process recovery for common-cause processes. *Physical Review Letters*, 126(23): 230401, 2021. doi: 10.1103/PhysRevLett.126.230401.

Ankit Gupta, Albert Gu, and Jonathan Berant. Diagonal state spaces are as effective as structured state spaces. In *Advances in Neural Information Processing Systems 35*, pp. 22982–22994, 2022.

William L. Hamilton, Mahdi Milani Fard, and Joelle Pineau. Modelling sparse dynamical systems with compressed predictive state representations. In *Proceedings of the 30th International Conference on Machine Learning*, volume 28 of *Proceedings of Machine Learning Research*, pp. 178–186. PMLR, 2013.

Ahmed Hefny, Carlton Downey, and Geoffrey J. Gordon. Supervised learning for dynamical system learning. In *Advances in Neural Information Processing Systems 28*, 2015.

Ahmed Hefny, Zita Marinho, Wen Sun, Siddhartha Srinivasa, and Geoffrey Gordon. Recurrent predictive state policy networks. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 1949–1958. PMLR, 2018.

B. L. Ho and R. E. Kalman. Effective construction of linear state-variable models from input/output functions. *Regelungstechnik*, 14(12):545–548, 1966. doi: 10.1524/auto.1966.14.112.545.

Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963. doi: 10.1080/01621459.1963.10500830.

Harold Hotelling. Relations between two sets of variates. *Biometrika*, 28(3/4):321–377, 1936. doi: 10.1093/biomet/28.3-4.321.

Chunqing Huang, Yisheng An, Sun Zhou, Zhezheng Hong, and Yunlong Liu. Basis selection in spectral learning of predictive state representations. *Neurocomputing*, 310:183–189, 2018. doi: 10.1016/j.neucom.2018.04.079.

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

Mathias R. Jørgensen and Felix A. Pollock. Exploiting the causal tensor network structure of quantum processes to efficiently simulate non-Markovian path integrals. *Physical Review Letters*, 123 (24):240602, 2019. doi: 10.1103/PhysRevLett.123.240602.

Jer-Nan Juang and Richard S. Pappa. An eigensystem realization algorithm for modal parameter identification and model reduction. *Journal of Guidance, Control, and Dynamics*, 8(5):620–627, 1985. doi: 10.2514/3.20031.

Jonathan Keeling, E. Miles Stoudenmire, Mari-Carmen Bañuls, and David R. Reichman. Process tensor approaches to non-Markovian quantum dynamics. *Physical Review X*, 16(2):020502, 2026. doi: 10.1103/1ncg-11hz.

S. Y. Kung. A new identification and model reduction algorithm via singular value decomposition. In *Proceedings of the 12th Asilomar Conference on Circuits, Systems and Computers*, pp. 705–714, 1978.

Wallace E. Larimore. Canonical variate analysis in identification, filtering, and adaptive control. In *Proceedings of the 29th IEEE Conference on Decision and Control*, volume 2, pp. 596–604, 1990. doi: 10.1109/CDC.1990.203665.

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

Ilia A. Luchnikov, Michael Sonner, and Dmitry A. Abanin. Scalable tomography of many-body quantum environments with low temporal entanglement. *Physical Review A*, 113(5):052611, 2026. doi: 10.1103/78jc-wcqb.

Krithika Manohar, Bingni W. Brunton, J. Nathan Kutz, and Steven L. Brunton. Data-driven sparse sensor placement for reconstruction: Demonstrating the benefits of exploiting known patterns. *IEEE Control Systems Magazine*, 38(3):63–86, 2018. doi: 10.1109/MCS.2018.2810460.

Peter McCracken and Michael Bowling. Online discovery and learning of predictive state representations. In *Advances in Neural Information Processing Systems 18*, 2005.

Simon Milz and Kavan Modi. Quantum stochastic processes and quantum non-Markovian phenomena. *PRX Quantum*, 2(3):030201, 2021. doi: 10.1103/PRXQuantum.2.030201.

Hazime Mori. Transport, collective motion, and Brownian motion. *Progress of Theoretical Physics*, 33(3):423–455, 1965. doi: 10.1143/PTP.33.423.

Felix A. Pollock, César Rodríguez-Rosario, Thomas Frauenheim, Mauro Paternostro, and Kavan Modi. Non-Markovian quantum processes: Complete framework and efficient characterization. *Physical Review A*, 97(1):012127, 2018a. doi: 10.1103/PhysRevA.97.012127.

Felix A. Pollock, César Rodríguez-Rosario, Thomas Frauenheim, Mauro Paternostro, and Kavan Modi. Operational Markov condition for quantum processes. *Physical Review Letters*, 120(4): 040405, 2018b. doi: 10.1103/PhysRevLett.120.040405.

Satinder Singh, Michael R. James, and Matthew R. Rudary. Predictive state representations: A new theory for modeling dynamical systems. In *Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence*, pp. 512–519, 2004.

Jimmy T. H. Smith, Andrew Warrington, and Scott W. Linderman. Simplified state space layers for sequence modeling. In *International Conference on Learning Representations*, 2023.

Alex Strathearn, Peter Kirton, Dainius Kilda, Jonathan Keeling, and Brendon W. Lovett. Efficient non-Markovian quantum dynamics using time-evolving matrix product operators. *Nature Communications*, 9:3322, 2018. doi: 10.1038/s41467-018-05617-3.

Wen Sun, Arun Venkatraman, Byron Boots, and J. Andrew Bagnell. Learning to filter with predictive state inference machines. In *Proceedings of the 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pp. 1197–1205, 2016.

Philip Taranto, Simon Milz, Felix A. Pollock, and Kavan Modi. Structure of quantum stochastic processes with finite Markov order. *Physical Review A*, 99(4):042108, 2019a. doi: 10.1103/PhysRevA.99.042108.

Philip Taranto, Felix A. Pollock, Simon Milz, Marco Tomamichel, and Kavan Modi. Quantum Markov order. *Physical Review Letters*, 122(14):140401, 2019b. doi: 10.1103/PhysRevLett.122.140401.

Philip Taranto, Felix A. Pollock, and Kavan Modi. Non-Markovian memory strength bounds quantum process recoverability. *npj Quantum Information*, 7:149, 2021. doi: 10.1038/s41534-021-00481-4.

Peter Van Overschee and Bart De Moor. N4SID: Subspace algorithms for the identification of combined deterministic-stochastic systems. *Automatica*, 30(1):75–93, 1994. doi: 10.1016/0005-1098(94)90230-5.

Per-Åke Wedin. Perturbation bounds in connection with singular value decomposition. *BIT Numerical Mathematics*, 12(1):99–111, 1972. doi: 10.1007/BF01932678.

Gregory A. L. White, Felix A. Pollock, Lloyd C. L. Hollenberg, Kavan Modi, and Charles D. Hill. Non-Markovian quantum process tomography. *PRX Quantum*, 3(2):020344, 2022. doi: 10.1103/PRXQuantum.3.020344.

Xinfang Zhang, Zhihao Wu, Gregory A. L. White, Zhongcheng Xiang, Shun Hu, Zhihui Peng, Yong Liu, Dongning Zheng, Xiang Fu, Anqi Huang, Dario Poletti, Kavan Modi, Junjie Wu, Mingtang Deng, and Chu Guo. Learning and forecasting open quantum dynamics with correlated noise. *Communications Physics*, 8:29, 2025. doi: 10.1038/s42005-025-01944-2.

Robert Zwanzig. Memory effects in irreversible thermodynamics. *Physical Review*, 124(4):983–992, 1961. doi: 10.1103/PhysRev.124.983.

A THEORY OF OPERATIONAL MEMORY AND OPMS

Let $P \in \mathbb{R}^p$ and $Q \in \mathbb{R}^q$ denote the population residualized history and future-response features from Section 3.1, with positive-definite covariance matrices. The population canonical matrix is $M_{\mathrm{OP}} = \Sigma_{QQ}^{-1/2}\Sigma_{QP}\Sigma_{PP}^{-1/2}$.

A.1 POPULATION OPMS GEOMETRY

Proof of Proposition 2. For invertible transforms $P' = AP$ and $Q' = BQ$, define

$$O_P = \Sigma_{P'P'}^{-1/2} A \Sigma_{PP}^{1/2}, \qquad O_Q = \Sigma_{Q'Q'}^{-1/2} B \Sigma_{QQ}^{1/2}.$$

Because $O_P O_P^\top = I$ and $O_Q O_Q^\top = I$, both are orthogonal and

$$M'_{\mathrm{OP}} = O_Q M_{\mathrm{OP}} O_P^\top.$$

Thus the population canonical correlations are invariant under nonsingular linear reparameterizations (Hotelling, 1936). Invertible whitening also preserves matrix rank, giving

$$\mathrm{rank}(M_{\mathrm{OP}}) = \mathrm{rank}(\Sigma_{QP}).$$

Let $\bar{P} = \Sigma_{PP}^{-1/2} P$ and $\bar{Q} = \Sigma_{QQ}^{-1/2} Q$. Since $\mathrm{Cov}(\bar{P}) = I_p$, $\mathrm{Cov}(\bar{Q}) = I_q$, and $\mathrm{Cov}(\bar{Q}, \bar{P}) = M_{\mathrm{OP}}$,

$$\mathbb{E}\|\bar{Q} - L\bar{P}\|_2^2 = q - 2\,\mathrm{tr}(LM_{\mathrm{OP}}^\top) + \mathrm{tr}(LL^\top) \tag{12}$$

$$= q - \|M_{\mathrm{OP}}\|_F^2 + \|L - M_{\mathrm{OP}}\|_F^2. \tag{13}$$

Thus the population least-squares map is M_{OP}, and the rank-k problem reduces to the best rank-k Frobenius approximation of this matrix. By the Eckart–Young theorem,

$$\min_{\mathrm{rank}(L)\le k} \mathbb{E}\|\bar{Q} - L\bar{P}\|_2^2 = q - \sum_{i=1}^{k} \rho_i^2. \tag{14}$$

The canonical history coordinates $V_{1:k}^\top \bar{P}$ therefore form the minimum-error k-dimensional linear predictive state for the selected future-response features. These statements concern linear dependence in the chosen feature spaces; nonlinear dependence can require nonlinear feature maps.

A.2 CONTEXT RESOLUTION AND GAUSSIAN INFORMATION

Proof of Proposition 3. Using $\pi_\beta = \Pr(\beta)$, $P^\circ = P_0 - \mathbb{E}[P_0 \mid \beta]$, $Q^\circ = Q_0 - \mathbb{E}[Q_0 \mid \beta]$, and $\Sigma_{QP|\beta} = \mathrm{Cov}(Q^\circ, P^\circ \mid \beta)$,

$$\Sigma_{Q_{\mathrm{CR}}P^\circ} = \mathbb{E}[(e_\beta \otimes Q^\circ)(P^\circ)^\top] = \mathrm{stack}_\beta[\pi_\beta \Sigma_{QP|\beta}]. \tag{15}$$

The row blocks occupy disjoint coordinates and every π_β is positive, so the stacked cross-covariance is zero exactly when every $\Sigma_{QP|\beta}$ is zero. Moreover, since $\mathbb{E}[Q^\circ \mid \beta] = 0$, with $\Sigma_{QQ|\beta} := \mathrm{Cov}(Q^\circ, Q^\circ \mid \beta)$,

$$\Sigma_{Q_{\mathrm{CR}}Q_{\mathrm{CR}}} = \mathrm{blockdiag}_\beta\big(\pi_\beta \Sigma_{QQ|\beta}\big)\,.$$

Hence positive-definite retained conditional covariances make the lifted covariance nonsingular, and whitening preserves the zero/nonzero equivalence. By contrast, the pooled cross-covariance $\sum_\beta \pi_\beta \Sigma_{QP|\beta}$ can vanish through sign or orientation cancellation; the two-context choice $\Sigma_{QP|1} = M$ and $\Sigma_{QP|2} = -M$ is the simplest example. The lift therefore preserves nonzero conditional residual cross-covariance without adding a context main effect.

Corollary 1 (Gaussian information decomposition)**.** *If (P, Q) are nondegenerate jointly Gaussian and $\rho_i < 1$ for every canonical pair, then*

$$I(P; Q) = -\frac{1}{2}\sum_i \log(1 - \rho_i^2). \tag{16}$$

Proof. An invertible canonical-coordinate transform decomposes jointly Gaussian (P, Q) into independent correlated scalar pairs with correlations ρ_i plus uncorrelated residual coordinates. Mutual information is invariant to the transform and additive across canonical pairs, which gives Equation (16). For example, $(0.8, 0)$ and $(0.8/\sqrt{2}, 0.8/\sqrt{2})$ have the same $S_{\mathrm{OP}} = 0.64$ but different accessible ranks and different Gaussian mutual information.

A.3 Trajectory separation and predictive realization order

Relation to predictive realization order. If $Q = Cz + \epsilon$ and $\mathrm{Cov}(\epsilon, P) = 0$ for a d-dimensional predictive state z, then $\Sigma_{QP} = C\Sigma_{zP}$ and

$$\mathrm{rank}(M_{\mathrm{OP}}) = \mathrm{rank}(\Sigma_{QP}) \leq \mathrm{rank}(\Sigma_{zP}) \leq d. \tag{17}$$

This is a one-way relation: recursive closure can require state directions that are weak or absent in a single conditioned history–response cross-covariance.

Proof of Proposition 1. For this proposition, ordinary multivariate R^2 is the pooled coefficient of determination

$$R^2 = 1 - \frac{\sum_{n,s\in\{+,-\}} \|\widehat{y}_n^s - y_n^s\|_2^2}{\sum_{n,s\in\{+,-\}} \|y_n^s - \bar{y}\|_2^2},$$

where $\bar{y}$ is the mean over all paired responses. Choose paired vector responses $y_n^+ = c_n + \frac{\epsilon}{2} g_n$ and $y_n^- = c_n - \frac{\epsilon}{2} g_n$, with $C = \sum_n \|c_n - \bar{c}\|_2^2 > 0$ and $G = \sum_n \|g_n\|_2^2 > 0$. For a fixed $M \geq 0$, let $\kappa = \sqrt{M+1} - 1$ and predict $\hat{y}_n^+ = c_n - \frac{\kappa\epsilon}{2} g_n$, $\hat{y}_n^- = c_n + \frac{\kappa\epsilon}{2} g_n$. The true and predicted response differences are ϵg_n and $-\kappa\epsilon g_n$, hence Equation (2) gives

$$\text{Causal Skill} = 1 - (1+\kappa)^2 = -M. \tag{18}$$

The ordinary sum of squared errors is $\frac{1}{2}(M+1)\epsilon^2 G$, whereas the total sum of squares over both members of each pair is $2C + \frac{1}{2}\epsilon^2 G$. Therefore

$$R^2 = 1 - \frac{(M+1)\epsilon^2 G/2}{2C + \epsilon^2 G/2} \longrightarrow 1 \qquad (\epsilon \to 0). \tag{19}$$

For any $\eta > 0$, sufficiently small ϵ gives $R^2 > 1 - \eta$ while Causal Skill remains $-M$. The proposition separates the two evaluation objectives; it does not assert failure of a particular learning algorithm.

Corollary 2 (Finite-rank predictive realization under update closure)**.** *Let $\mathfrak{H}_0$ be a matched-context fiber of reachable histories and let $\overline{\mathfrak{H}}$ be its closure under the history extensions considered for realization. Let $\mathcal{T}$ be closed under the corresponding future-test prefixes, define $\mathcal{H} : \overline{\mathfrak{H}} \times \mathcal{T} \to \mathbb{R}$ by $\mathcal{H}(h, \tau) = Y_\tau(h)$, and assume $\mathcal{H}(h\sigma, \tau) = \mathcal{H}(h, \sigma\tau)$ whenever the indexed extension and future test are defined. If $\mathrm{rank}(\mathcal{H}) = d < \infty$, then $\mathcal{H}$ admits a d-dimensional exact linear realization with shared state updates; no exact linear realization closed under these updates has smaller dimension, and any two minimal realizations are related by an invertible similarity transform.*

Proof. Take a rank factorization $\mathcal{H} = LR$ on $\overline{\mathfrak{H}} \times \mathcal{T}$, with row state $\ell(h)$ and future-test vector $v(\tau)$. If $\sum_i \alpha_i \ell(h_i) = 0$, then

$$\sum_i \alpha_i \mathcal{H}(h_i, \tau) = 0 \qquad \text{for every } \tau \in \mathcal{T}.$$

For any allowed update σ, prefix closure gives $\sigma\tau \in \mathcal{T}$ whenever the shifted entry is defined, and therefore

$$\sum_i \alpha_i \mathcal{H}(h_i\sigma, \tau) = \sum_i \alpha_i \mathcal{H}(h_i, \sigma\tau) = 0.$$

Hence every linear relation among row states is preserved by the shift, so $\ell(h) \mapsto \ell(h\sigma)$ is a well-defined linear map on the d-dimensional row-state span. Any exact d'-dimensional realization factorizes $\mathcal{H}$, hence $d \leq d'$. If two minimal factorizations satisfy $L' = LT$ and $R' = T^{-1}R$ for an invertible T, their update operators satisfy $A'_\sigma = T^{-1} A_\sigma T$. This is the standard finite-rank Hankel realization consequence (Ho & Kalman, 1966; Kung, 1978; Jaeger, 2000; Singh et al., 2004; Boots et al., 2011) applied to the intervention-response object. OCPR is its finite-data, depth-restricted, and generally truncated counterpart.

A.4 Relation to process-level memory strength

Process-level comparison. The memory strength of Taranto et al. (2021) is an operational quantity defined at the process level by optimizing measured conditional mutual information over future and history instruments for a specified intermediate instrument, and it controls process recoverability. OPMS instead conditions on a chosen observation and intervention setting and resolves the resulting second-order history–response dependence into canonical feature directions. Its population object is indexed by (ϕ, ψ, C_t, F) rather than by an optimization over all probing instruments. Proposition 2 gives the accessible linear rank and optimal rank-constrained predictive gain, while Corollary 1 gives the Gaussian mutual-information specialization. OCPR separately imposes recursive closure under free evolution and intervention.

B OCPR realization theory

B.1 Exact intervention coordinates

Vectorize the real Pauli-transfer representation of a replacement channel as $b(\varrho) \in \mathbb{R}^{16}$. Its nonzero degrees of freedom are the trace coordinate and three Bloch coordinates. If $B = [b(\varrho_1), \dots, b(\varrho_4)]$ has rank four, then

$$c(\varrho) = B^{+} b(\varrho), \qquad b(\varrho) = Bc(\varrho) \tag{20}$$

for every physical qubit state ϱ. The selected basis consists of $|0\rangle$, $|1\rangle$, $|+\rangle$, and a fourth non-coplanar pure state; its condition number is 3.010278 and maximum coordinate residual is 1.37×10^{-15}. The ordered two-intervention coordinate frame used in the response construction has rank 16. Substitution into a bilinear two-slot process response yields Equation (5); the maximum direct response reconstruction error is 7.051×10^{-15}.

Negative basis coefficients are allowed because the linear combination is taken in the linear span of the intervention maps; the final summed channel is the physical replacement. For m replacement slots the explicit response table contains 4^m basis entries, while the realized representation composes four basis-intervention updates instead of storing the full table at inference time.

Correspondence with classical predictive-state realization. OCPR retains low-rank Hankel factorization and shared state updates from predictive-state and subspace realization, while changing the response object to matched-history center and response-difference statistics, representing interventions in exact coordinates, and restricting update identification to the future-test coordinates available at the chosen intervention depth. Response-core selection is the corresponding reduced-response analogue.

B.2 Center–gap equivalence

Lemma 1 (Center–gap equivalence under exact realization)**.** *Let H_{indiv} stack the two individual response rows of every matched-history pair. For positive fixed normalization scales s_c, s_g, let H_{cg} stack the corresponding normalized center and response-difference rows. Then*

$$H_{\mathrm{cg}} = \mathcal{T} H_{\mathrm{indiv}},$$

where $\mathcal{T}$ is block diagonal with pairwise blocks

$$\begin{bmatrix} s_c^{-1} & 0 \\ 0 & s_g^{-1} \end{bmatrix} \begin{bmatrix} 1/2 & 1/2 \\ 1 & -1 \end{bmatrix},$$

and is invertible. Hence $\mathrm{rowspan}(H_{\mathrm{cg}}) = \mathrm{rowspan}(H_{\mathrm{indiv}})$ *and* $\mathrm{rank}(H_{\mathrm{cg}}) = \mathrm{rank}(H_{\mathrm{indiv}})$*. If an exact individual-response realization satisfies $H_{\mathrm{indiv}} = L_{\mathrm{indiv}} R$ and $H^{\sigma}_{\mathrm{indiv},J} = L_{\mathrm{indiv}} A_\sigma R_J$, then the center–gap construction satisfies*

$$H_{\mathrm{cg}} = (\mathcal{T} L_{\mathrm{indiv}}) R, \qquad H^{\sigma}_{\mathrm{cg},J} = (\mathcal{T} L_{\mathrm{indiv}}) A_\sigma R_J,$$

with the same right factor and update matrices.

Proof. For one pair,

$$\begin{bmatrix}\bar{y}_c\\ \bar{y}_g\end{bmatrix} = \begin{bmatrix}s_c^{-1} & 0\\ 0 & s_g^{-1}\end{bmatrix}\begin{bmatrix}1/2 & 1/2\\ 1 & -1\end{bmatrix}\begin{bmatrix}y(h^a)\\ y(h^b)\end{bmatrix}.$$

The 2×2 center–difference matrix has determinant -1, and the normalization matrix is invertible because $s_c, s_g > 0$. Stacking the same transformation over pairs gives an invertible $\mathcal{T}$, which preserves row space and rank. Applying that fixed row transformation to both the unshifted and shifted response matrices yields the two factorization identities above. Equivalently, $y(h^a) = s_c\bar{y}_c + \frac{s_g}{2}\bar{y}_g$ and $y(h^b) = s_c\bar{y}_c - \frac{s_g}{2}\bar{y}_g$. Thus center–gap changes the geometry used by truncation and weighting without adding information or changing exact realization order; its empirical advantage can arise through the retained subspace and the conditioning of the restricted update blocks.

B.3 Depth-restricted update identification

The response row contains 504 scalar coordinates: 24 free-evolution readouts, 48 readouts after one basis replacement, and 432 readouts after an ordered pair of basis replacements. For a future-test suffix s and update symbol σ, define the depth-restricted set J_σ by Equation (8). The corresponding response block satisfies

$$H^\sigma_{J_\sigma}[h, s] = H[h, \sigma s], \qquad s \in J_\sigma. \tag{21}$$

Using the rank-d truncated SVD $H \approx U_d\Sigma_d V_d^\top$, define $L = U_d\Sigma_d^{1/2}$ and $R = \Sigma_d^{1/2}V_d^\top$. Then

$$A_\sigma = L^+ H^\sigma_{J_\sigma}(R_{J_\sigma})^+. \tag{22}$$

If L has column rank d and R_{J_σ} has row rank d, this is the unique least-squares update on the realized predictive state. In the exact update-closed setting of Proposition 4,

$$L^+ H^\sigma_{J_\sigma}(R_{J_\sigma})^+ = L^+ L A_\sigma R_{J_\sigma}(R_{J_\sigma})^+ = A_\sigma, \tag{23}$$

because $L^+L = I_d$ and $R_{J_\sigma}(R_{J_\sigma})^+ = I_d$. The same identity proves uniqueness: any second update producing the same depth-restricted block has zero difference after left and right inversion. Repeated application of $z(h\sigma) = z(h)A_\sigma$ then gives $z(h\sigma_1\cdots\sigma_k) = z(h)A_{\sigma_1}\cdots A_{\sigma_k}$ without using responses beyond the identification depth to fit the updates.

B.4 Perturbation and restricted-update stability

Proposition 5 (Full singular spectrum does not certify restricted update stability)**.** *For every $K > 0$ and $\eta > 0$, there exist rank-two right factors $R, \widetilde{R} \in \mathbb{R}^{2\times 3}$ with identical full singular values, $RR^\top = \widetilde{R}\widetilde{R}^\top = I_2$, and row-space projectors satisfying $\|\Pi_R - \Pi_{\widetilde{R}}\|_2 < \eta$, while for the same two-column restriction J,*

$$\frac{\sigma_{\min}(R_J)}{\sigma_{\min}(\widetilde{R}_J)} > K. \tag{24}$$

Hence full-spectrum energy and global subspace proximity do not lower-bound restricted update stability.

Proof. For $0 < t < 1$, let

$$R_t = \begin{bmatrix}1 & 0 & 0\\ 0 & t & \sqrt{1-t^2}\end{bmatrix}. \tag{25}$$

Its rows are orthonormal for every t, so all R_t have the same two nonzero singular values. With $J = \{1, 2\}$, $\sigma_{\min}((R_t)_J) = t$. Comparing R_t with R_{t^2} gives a restricted-margin ratio $1/t$, while their second row vectors—and therefore their row-space projectors—converge to one another as $t \to 0$. Choosing t sufficiently small makes $1/t > K$ and the projector distance below η.

Proposition 6 (Perturbed depth-restricted update recovery)**.** *Fix an update σ and an exact realized factorization $H^\sigma_J = LA_\sigma R_J$. After aligning an estimated realization by an invertible gauge T, write $\bar{L} = LT$, $\bar{R}_J = T^{-1}R_J$, and $\bar{A}_\sigma = T^{-1}A_\sigma T$. Let $\epsilon_L = \|\widehat{L} - \bar{L}\|_2$, $\epsilon_R = \|\widehat{R}_J - \bar{R}_J\|_2$, and $\epsilon_\sigma = \|\widehat{H}^\sigma_J - H^\sigma_J\|_2$. If $\ell = \sigma_{\min}(\bar{L}) > \epsilon_L$ and $\gamma_\sigma = \sigma_{\min}(\bar{R}_J) > \epsilon_R$, then*

$$\|\widehat{A}_\sigma - \bar{A}_\sigma\|_2 \le \frac{\epsilon_\sigma}{(\ell - \epsilon_L)(\gamma_\sigma - \epsilon_R)} \tag{26}$$

$$+ \|\bar{A}_\sigma\|_2\left[\frac{\epsilon_L}{\ell - \epsilon_L}\left(1 + \frac{\epsilon_R}{\gamma_\sigma - \epsilon_R}\right) + \frac{\epsilon_R}{\gamma_\sigma - \epsilon_R}\right]. \tag{27}$$

For balanced SVD factors compared in an orthogonally aligned gauge with $\delta_d(H) > 0$, Wedin's theorem together with Weyl's singular-value bound controls ϵ_L, ϵ_R through the perturbation magnitude, the retained spectral separation $\delta_d(H)$, and the retained singular scale $\sigma_d(H)$. Thus δ_d, σ_d, and γ_{shift} govern complementary aspects of factor and restricted-update stability.

We prove the bound directly. Write $\Delta_L = \widehat{L} - \bar{L}$, $\Delta_R = \widehat{R}_J - \bar{R}_J$, and $\Delta_H = \widehat{H}_J^{\sigma} - H_J^{\sigma}$. Since $\widehat{L}$ has full column rank and $\widehat{R}_J$ has full row rank,

$$\widehat{A}_\sigma - \bar{A}_\sigma = \widehat{L}^{+}\Delta_H(\widehat{R}_J)^{+} + \widehat{L}^{+}\bar{L}\bar{A}_\sigma\bar{R}_J(\widehat{R}_J)^{+} - \bar{A}_\sigma. \tag{28}$$

Using $\widehat{L}^{+}\bar{L} - I = -\widehat{L}^{+}\Delta_L$ and $\bar{R}_J(\widehat{R}_J)^{+} - I = -\Delta_R(\widehat{R}_J)^{+}$, the second term is bounded by

$$\|\bar{A}_\sigma\|_2 \left[\|\widehat{L}^{+}\|_2\epsilon_L(1 + \epsilon_R\|(\widehat{R}_J)^{+}\|_2) + \epsilon_R\|(\widehat{R}_J)^{+}\|_2\right]. \tag{29}$$

Weyl's inequality gives $\sigma_{\min}(\widehat{L}) \geq \ell - \epsilon_L$ and $\sigma_{\min}(\widehat{R}_J) \geq \gamma_\sigma - \epsilon_R$, yielding Equation (27). When the exact and estimated realizations are both taken in the balanced SVD gauge, aligned by an orthogonal transformation, and $\delta_d(H) > 0$, $\ell = \sqrt{\sigma_d(H)}$. Wedin's theorem then controls movement of the retained left/right singular subspaces through the spectral separation $\delta_d(H)$, while Weyl's inequality controls the singular-value perturbation; the square-root scaling of balanced factors also depends on the retained singular scale $\sigma_d(H)$ (Wedin, 1972). Thus total response-energy retention alone is not a stability certificate: a small subspace perturbation can substantially change γ_σ and hence the restricted pseudoinverse appearing in the depth-restricted update estimator.

At the fixed state dimension $d = 16$, free evolution uses all 504 suffix columns. Each replacement update uses 72 columns (free-evolution and one-replacement future-test suffixes), so prefixing a replacement never exceeds intervention depth two. The restricted designs retain rank 16 with condition numbers 992.264 and 617.950 in the intermediate and structured environments, compared with 95.158 and 71.720 for all 504 columns.

A known finite-dimensional linear system recovers its state update with relative error 3.91×10^{-15} and seven-step composed-response error 2.70×10^{-14} under Equation (22). Perturbing response entries excluded from replacement-update identification by independent noise six orders of magnitude larger than the response scale leaves the identified updates unchanged to machine precision. Thus the depth-three evaluation uses updates identified without depth-three response statistics. If v_F is the realized column associated with future test F, prediction takes the form

$$\widehat{Y}_F = z(h)A_{\sigma_1}\cdots A_{\sigma_k}v_F. \tag{30}$$

The resulting free-evolution spectral radii are 1.042844 and 1.000002 in the intermediate and structured environments. Recursive validity is evaluated as a finite-horizon property at $\Delta = 12, 16$ and at intervention depth three.

B.5 Finite-horizon composition

All vector and operator norms in this subsection are 2-norms. For row-state dynamics define $e_t = \widehat{z}_t - z_t$, $\widehat{z}_{t+1} = \widehat{z}_t\widehat{A}_t$, and $z_{t+1} = z_tA_t$. Then

$$\|e_{t+1}\| \leq \|\widehat{A}_t\|\|e_t\| + \|\widehat{A}_t - A_t\|\|z_t\|, \tag{31}$$

$$\|e_k\| \leq \Big(\prod_{t=0}^{k-1}\|\widehat{A}_t\|\Big)\|e_0\| \tag{32}$$

$$+ \sum_{j=0}^{k-1}\Big(\prod_{t=j+1}^{k-1}\|\widehat{A}_t\|\Big)\|\widehat{A}_j - A_j\|\,\|z_j\|. \tag{33}$$

Under exact closure, multilinearity and Proposition 4 give $A_G = \sum_i c_i(G)A_{B_i}$. For the truncated empirical realization we use $\widehat{A}_G := \sum_i c_i(G)\widehat{A}_{B_i}$, so

$$\|\widehat{A}_G - A_G\| \leq \sum_i |c_i(G)|\,\|\widehat{A}_{B_i} - A_{B_i}\| \leq \|c(G)\|_2\Big(\sum_i\|\widehat{A}_{B_i} - A_{B_i}\|^2\Big)^{1/2}. \tag{34}$$

For a linear readout v_F, the additional output error is at most $\|e_k\|\|v_F\|$ plus readout approximation error.

Corollary 3 (Finite-horizon OCPR error)**.** *Suppose over a k-step composition that $\|\widehat{A}_t\|_2 \le \alpha$, $\|\widehat{A}_t - A_t\|_2 \le \epsilon_A$, $\|z_t\|_2 \le Z$, and $\|e_0\|_2 \le \epsilon_0$. Then*

$$\|e_k\|_2 \le \alpha^k \epsilon_0 + Z\epsilon_A \sum_{j=0}^{k-1} \alpha^j.$$

Equivalently, the second term is $Z\epsilon_A(\alpha^k - 1)/(\alpha - 1)$ for $\alpha \ne 1$ and $kZ\epsilon_A$ for $\alpha = 1$. If the final readout has approximation error at most ϵ_{read}, then

$$|\widehat{Y}_F - Y_F| \le \|v_F\|_2 \left(\alpha^k \epsilon_0 + Z\epsilon_A \sum_{j=0}^{k-1} \alpha^j \right) + \epsilon_{\mathrm{read}}.$$

The bound follows directly from Equation (33). Proposition 6 controls the basis-update errors entering ϵ_A, while Equation (34) transfers those errors to held-out interventions through their exact coordinates. Together, these bounds link response and factor perturbations, restricted conditioning, and finite-horizon output error. The depth-three experiment evaluates the same composition rule one intervention beyond the maximum physical intervention depth used to identify replacement updates.

B.6 Response-core identifiability

The response-core condition in Definition 2 has a direct exact-realization interpretation.

Corollary 4 (Exact response-core sufficiency)**.** *Let $H = LR$ be an exact rank-d realization and let S be a coordinate subset such that $\mathrm{rank}(R_S) = d$ and $\mathrm{rank}(R_{S\cap J_\sigma}) = d$ for every identified update σ. Then the restricted response bank preserves realization order d, and every identified update is uniquely recovered from*

$$H^\sigma_{S\cap J_\sigma} = LA_\sigma R_{S\cap J_\sigma}.$$

Proof. Because R_S has row rank d, $H_S = LR_S$ has rank d and therefore preserves the exact realization dimension. For every identified σ, the matrix $R_{S\cap J_\sigma}$ also has row rank d; Proposition 4 therefore gives uniqueness and exact recovery of A_σ from the restricted shifted block. The result concerns exact rank preservation; finite-sample conditioning of a selected response core is evaluated separately in Appendix F.2.

C Estimation and evaluation setup

C.1 Matched-context constructions

Matched histories share the observed present and all subsequent interventions while retaining distinct historical information that may affect future responses. In the latent-memory process, x_t is matched while z_t is preserved. In the generalized Langevin process (Zwanzig, 1961; Mori, 1965), visible coordinate and velocity are matched while the convolutional memory term is preserved. In the quantum process, the same replacement fixes the system marginal while preserving environmental history.

C.2 Estimation and statistics

Finite-sample OPMS uses regularized CCA. Held-out analyses fit all preprocessing and canonical directions on the training fold and apply them unchanged to held-out data; the controlled spectra in Fig. 5 use Eq. (3) directly. We report $\widehat{S}_{\mathrm{OP}} = \sum_i \widehat{\rho}_i^2$. Domain-blocked permutations repeat the full estimation; PDE multiplicity is controlled by Benjamini–Hochberg (Benjamini & Hochberg, 1995) and a max-statistic null.

For context-resolved OPMS, both history and future-response features are centered within each context stratum. The future-response residuals are then placed in disjoint blocks as in Equation (4). Blockwise whitening is algebraically equivalent to constructing the full lifted feature; the underlying samples, response coordinates, and null construction are unchanged.

Table 3: **Spin–boson and evaluation settings.**

Item	Setting
Intermediate environment	$\alpha = 0.10$, cutoff 4, temperature 0.5
Structured environment	$\alpha = 0.20$, cutoff 2, temperature 0.2
Observed history	64 state–control steps
Replacement split	$20/5/30$ train/validation/evaluation pairs
Main selection horizons	$\Delta \in \{2, 4, 8\}$
Main held-out horizons	$\Delta = 12, 16$ and intervention depth three
State dimension	$d = 16$
Coherent-intervention split	$20/5/10$ train/validation/test pairs

For the hardware analyses, the matched-null threshold is the empirical 95th percentile of the corresponding statistic under its null distribution, computed separately for each dataset, condition, and statistic. Giarmatzi nulls use 10,000 conditional binomial resamples: the future-outcome probability is pooled over the six histories within each fixed context and readout axis, while each observed selected-branch count is preserved. Gosling nulls use 5,000 replicates with independent circular shifts of the history feature relative to future responses within each context. For a context with n_β eligible anchors, the integer shift is sampled uniformly from $\{m_\beta, \ldots, n_\beta - m_\beta\}$, where

$$m_\beta = \min\{100, \max\{1, \lfloor n_\beta/4 \rfloor\}\}.$$

One-sided Monte Carlo p-values are

$$p = \frac{1 + \sum_{b=1}^{B} \mathbf{1}\{S_b^{\mathrm{null}} \geq S_{\mathrm{obs}}\}}{B + 1}.$$

For the primary hardware analysis, a positive result requires $p \leq 0.01$ and a positive 2.5th percentile of the replicatewise $S_b^{\mathrm{boot}} - S_b^{\mathrm{null}}$ differences. For the split-shot robustness analysis, the corresponding criterion uses $p \leq 0.05$ and a positive 2.5th percentile of the split-shot-minus-null differences. Bootstrap uncertainty uses 10,000 multinomial resamples of the four-outcome counts from the Giarmatzi protocol and 5,000 circular moving-block resamples with 100 eligible anchors per block within each context for Gosling.

All OCPR realization comparisons use the same 504-coordinate future-response bank and the same depth-restricted update access. Causal Skill is the primary response-difference metric; trajectory R^2, response-difference RMSE/Pearson correlation, and multiplicative recovery are reported only when they clarify failure mode or scale. The matched-history pair is the statistical unit.

For endpoints with nonzero true response-difference energy, multiplicative recovery measures the aggregate amplitude of the predicted response differences relative to the true response differences:

$$R_{\mathrm{mult}} = \frac{\mathrm{RMS}(\widehat{\Delta Y})}{\mathrm{RMS}(\Delta Y)} = \left(\frac{\sum_{p,j} (\widehat{\Delta Y}_{p,j})^2}{\sum_{p,j} (\Delta Y_{p,j})^2} \right)^{1/2}.$$

Here p indexes matched-history pairs and j runs over all evaluated intervention targets, temporal configurations, future times, and readout components within the reported endpoint. The squared responses are pooled before taking the ratio. Values above or below one indicate aggregate amplitude inflation or attenuation, respectively; a value of one does not by itself imply correct response directions.

C.3 Simulation and evaluation settings

The spin–boson process uses

$$H(t) = \tfrac{1}{2}\epsilon\sigma_z + \tfrac{1}{2}(\delta + u_x(t))\sigma_x + \tfrac{1}{2}u_y(t)\sigma_y, \tag{35}$$

with $\epsilon = \delta = 1$ and initial state $|0\rangle$. Propagation uses PT–TEMPO (Strathearn et al., 2018; Jørgensen & Pollock, 2019) through OQuPy (Fux et al., 2024). The evaluated environments and data splits are summarized in Table 3.

The 30-pair replacement evaluation set comprises a fixed 10-pair subset used for the realization and state-access ablations in Table 2, together with 20 further held-out matched-history pairs per environment.

The history-to-state model and generic conditional-regression control use the same causal temporal encoder family (Bai et al., 2018); preprocessing is fitted on training data only.

C.4 Predictive state inference from observable history

Let $X \in \mathbb{R}^{n\times 65}$ stack normalized history features and an intercept, let $Z \in \mathbb{R}^{n\times 16}$ contain realized predictive states, and let n_p denote the number of matched-history pairs. Pair-incidence matrices $C, G \in \mathbb{R}^{n_p\times n}$ compute pair centers and state differences: $C_{p,a_p} = C_{p,b_p} = 1/2$ and $G_{p,a_p} = 1$, $G_{p,b_p} = -1$. We fit

$$\frac{\|XW - Z\|_F^2}{nd\, s_a^2} + \frac{\|CXW - CZ\|_F^2}{n_p d\, s_c^2} + \frac{\|GXW - GZ\|_F^2}{n_p d\, s_g^2} + \lambda \|W_{1:64}\|_F^2, \tag{36}$$

with training RMS scales s_a, s_c, s_g and an unregularized intercept. Hyperparameter selection uses validation data only.

D Supplementary OPMS results

D.1 Controlled systems

The controlled spectra recover the expected low-dimensional memory structure (Figure 5).

D.2 Superconducting hardware

On the Giarmatzi data (Giarmatzi et al., 2025), pooled OPMS remains below its matched-null threshold whereas context-resolved OPMS exceeds its threshold. Across nine IBM timing settings, the context-resolved statistic exceeds its matched-null threshold in all nine settings, with the same conclusion under the split-shot analysis. On the Gosling quantum-jump data (Gosling et al., 2026a;b), context-resolved OPMS exceeds its matched-null threshold and closely follows the separately fitted two-level-system (TLS) structure.

The reference values reported with the hardware statistics in Table 4 are the matched-null 95th percentiles defined in Appendix C.2. The hardware results and robustness checks are summarized in the same table. Together, the two datasets support the context-resolution principle. Neither hardware analysis identifies OCPR updates.

Table 4: **Superconducting hardware results.**

Dataset	Analysis	Result	Reference / robustness
Giarmatzi	Pooled	0.0295	0.0888
	Context-resolved	1.4378	0.5621
	Timing robustness	9/9 primary criterion	9/9 split-shot criterion
Gosling	Full grid	0.006256	0.000468
	First / second half	0.005279 / 0.005620	0.000251 / 0.000345

D.3 PDE dynamics

Cross-fitted OPMS detects delayed-memory structure in the medium and strong Burgers settings; weak Burgers and Kuramoto–Sivashinsky are not resolved by the selected feature and future-test family.

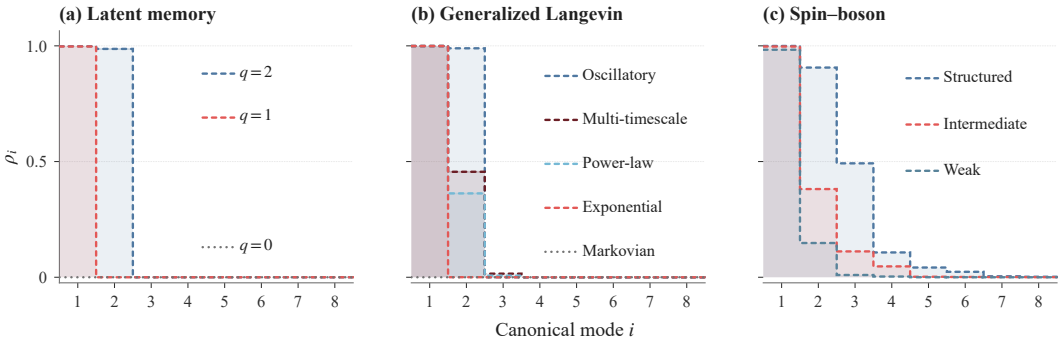

Figure 5: **Controlled OPMS spectra.** Regularized empirical canonical correlations ρ_i for **(a)** latent-memory controls, **(b)** generalized Langevin controls, and **(c)** spin–boson environments. In panel (a), q denotes the number of latent-memory channels.

E Supplementary OCPR results

E.1 Intervention representation and realization geometry

Exact coordinates reconstruct held-out replacement responses to numerical precision; mixed and maximally mixed replacement targets give the same Causal Skill to three decimals.

In Table 5, $\kappa_{\text{shift}} := \max_\sigma \kappa_2(R_{J_\sigma})$ is the worst restricted-update condition number and $\delta_{16} := \sigma_{16}(H) - \sigma_{17}(H)$ is the retained spectral separation at $d = 16$.

Table 5: **Restricted-update stability at rank 16.**

Environment	γ_{shift} Abs.→C–G	κ_{shift} Abs.→C–G	δ_{16} Abs.→C–G
Intermediate	$0.00420 \to 0.01843$ (4.38×)	$2264 \to 992$ (2.28×)	$0.00408 \to 0.02239$
Structured	$0.00589 \to 0.02848$ (4.84×)	$1616 \to 618$ (2.61×)	$0.01023 \to 0.02521$

The singular spectra are shown in Figure 6; realization geometry and weighting robustness are reported in Figure 3, with restricted-update conditioning summarized in Table 5. Center–gap improves restricted-update conditioning under the same response access, and the gain persists across the tested weighting range. At the fixed $d = 16$, the realization retains more directions than are needed to capture most response energy. This illustrates why response energy alone does not determine recursive stability: $\delta_d(H)$ controls the retained subspace, while γ_{shift} controls the restricted update geometry.

E.2 History-to-state inference

The same ordering holds at $\Delta = 8$ and $\Delta = 16$: direct state inference provides most of the gain, while paired center–gap supervision adds a smaller improvement.

E.3 Temporal and compositional generalization

Temporal generalization is shown in Figure 7. On the 30-pair depth-three evaluation, history-inferred response-difference correlation is $0.984/0.987$ and multiplicative recovery is $0.939/0.931$ in the intermediate/structured environments.

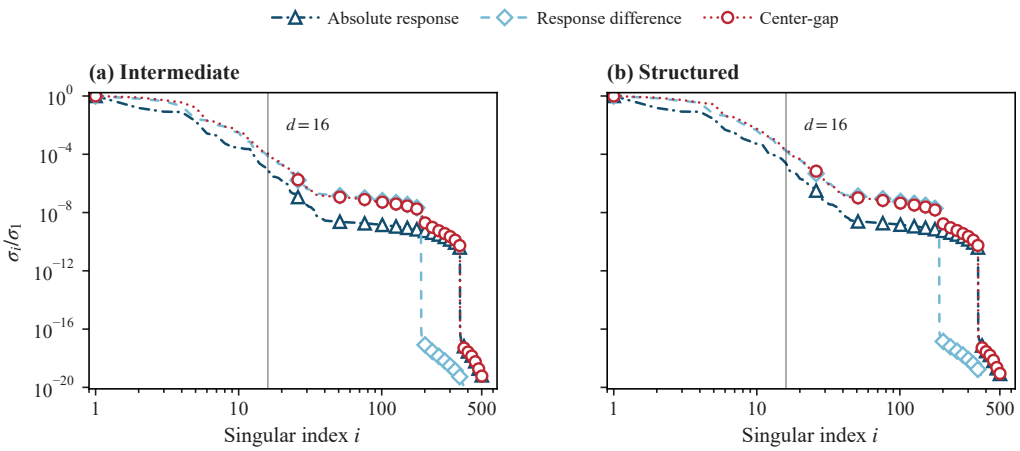

Figure 6: **Hankel singular spectra.** Absolute-response, response-difference, and center–gap realizations; the vertical line marks $d = 16$.

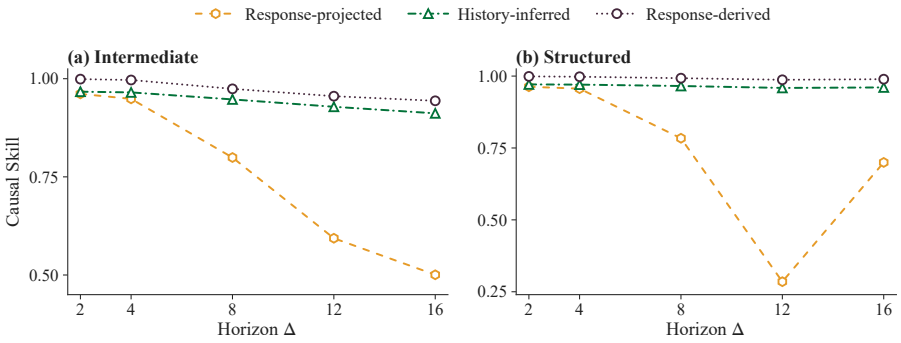

Figure 7: **Temporal generalization.** Causal Skill across tested prediction horizons.

E.4 NON-QUANTUM REALIZATION TEST

The two-channel non-quantum latent-memory control has exact affine state $[1, x, z_1, z_2]$, uses three basis replacements, and identifies updates from tests of depth at most two. We evaluate three instances of this control; each uses $d = 4$, matching the exact affine state, and composes through all 6^3 held-out intervention triples with Causal Skill 1.000000.

E.5 COHERENT-INTERVENTION BREADTH

The continuous y-axis rotation family has channel-span rank three and reconstructs held-out rotations to numerical precision. Response-derived center–gap Causal Skill remains positive across the tested environments and horizons (Figure 8). The coherent-intervention setting also exposes a remaining limitation of state inference from observable history: history-inferred states do not consistently preserve both trajectory accuracy and response-difference recovery across the two environments.

E.6 HELD-OUT REPLACEMENT EVALUATION

On the 30-pair held-out replacement evaluation, history-inferred Causal Skill is $0.930/0.958$ at $\Delta = 16$ and $0.966/0.971$ at depth three, while response-derived Causal Skill is $0.949/0.988$ and $0.993/0.998$, respectively, in the intermediate/structured environments.

Coordinate-level agreement is shown in Figure 4; the history-inferred response atlas is shown in Figure 9.

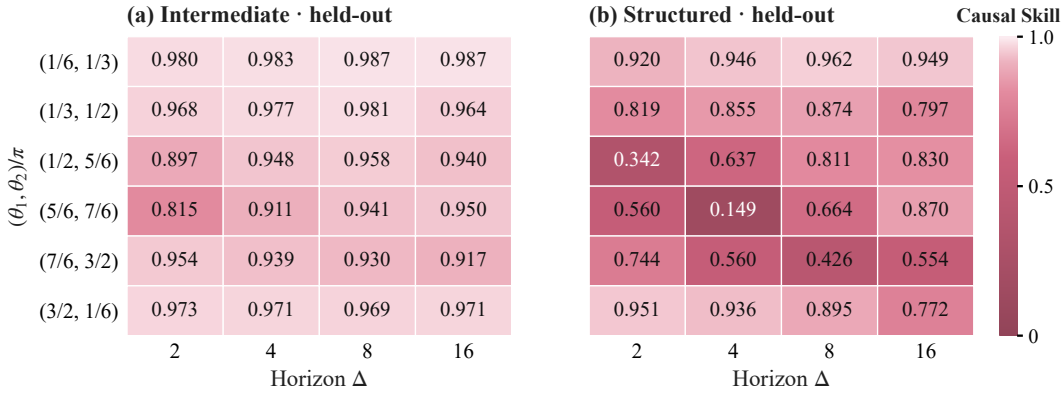

Figure 8: **Coherent-intervention breadth.** Response-derived Causal Skill across held-out coherent rotations and prediction horizons in both environments.

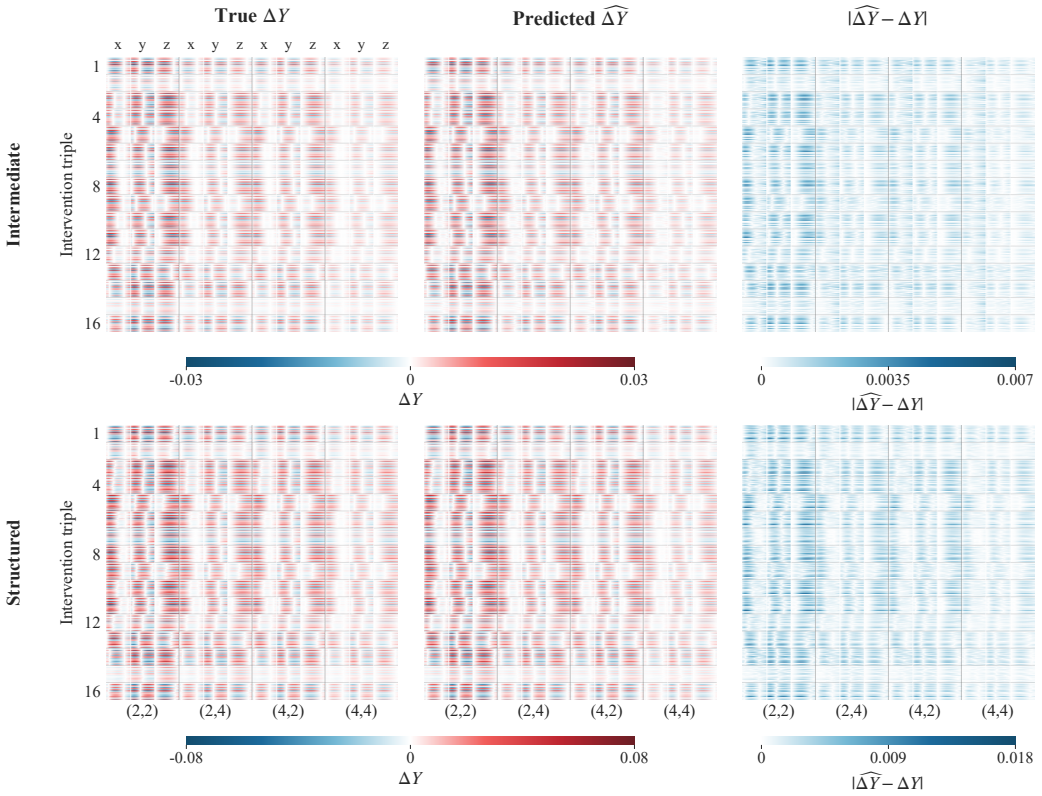

Figure 9: **Depth-three history-inferred response atlas.** True response differences, OCPR predictions, and absolute errors on the 30-pair held-out evaluation in the intermediate and structured environments.

F MEASUREMENT EFFICIENCY

F.1 FINITE-MEASUREMENT ROBUSTNESS

For a D-coordinate response row y estimated from m independent binary measurements per coordinate, Hoeffding's inequality (Hoeffding, 1963) gives

$$\Pr(\|\widehat{y} - y\|_\infty > \epsilon) \le 2D \exp(-m\epsilon^2/2), \tag{37}$$

and therefore

$$m \ge 2\epsilon^{-2} \log \frac{2D}{\delta} \tag{38}$$

suffices for simultaneous coordinate error at most ϵ with probability $1 - \delta$. For a fixed spectral basis,

$$\|\widehat{z} - z\|_2 \le \frac{\|\widehat{y} - y\|_2}{\sqrt{\sigma_d}}. \tag{39}$$

When the basis is estimated, subspace perturbation and restricted-update conditioning additionally depend on δ_d, σ_d, and γ_{shift} as analyzed in Proposition 6.

Finite measurements are simulated by drawing binary outcomes for each response expectation and refitting the realization and state map. We use 20 realizations for each tested measurement count.

Finite-measurement results are shown in Figure 10. Performance improves with measurement count; depth-three composition is more sensitive to response noise.

F.2 RESPONSE-CORE SELECTION AND EQUAL-BUDGET IDENTIFICATION

A common 16-coordinate response core is selected from the training responses of both environments by column-pivoted QR (Manohar et al., 2018), subject to rank preservation in the base realization and every update-identification block. The same selected core is then used for all reported equal-budget evaluations.

Equal-budget response-core results are shown in Figure 11. The core preserves the main $\Delta = 16$ and depth-three results while reducing subsequent response expectations from $391{,}800$ to $20{,}848$ per environment ($18.79\times$), and gives higher mean history-inferred Causal Skill at every tested equal budget.

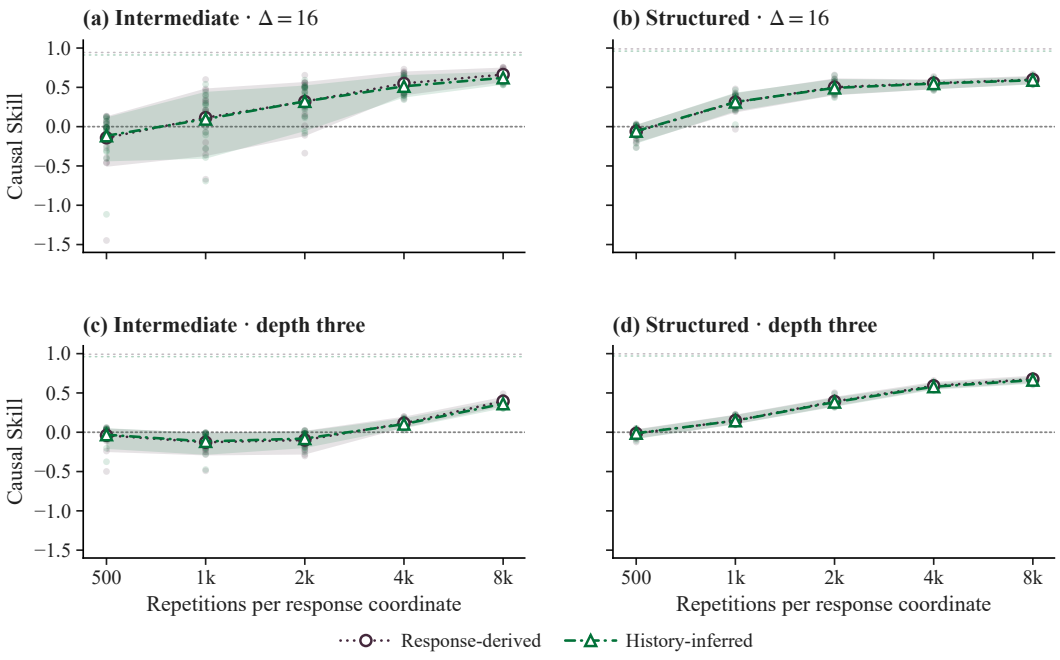

Figure 10: **Finite-measurement robustness.** Causal Skill versus repetitions per response coordinate. Curves show means over 20 realizations; shading summarizes variability, faint points show individual realizations, and dotted lines mark noise-free references.

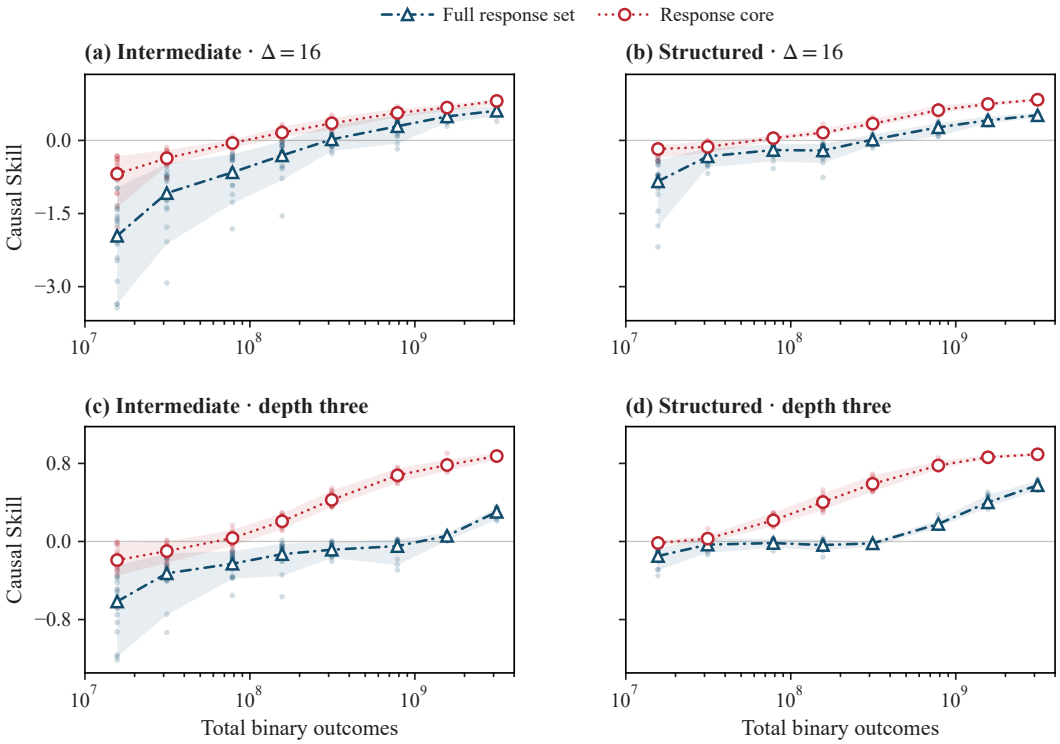

Figure 11: **Response-core efficiency.** History-inferred Causal Skill under equal total measurement budgets. Curves show means over 20 realizations; shading summarizes variability and faint points show individual realizations.

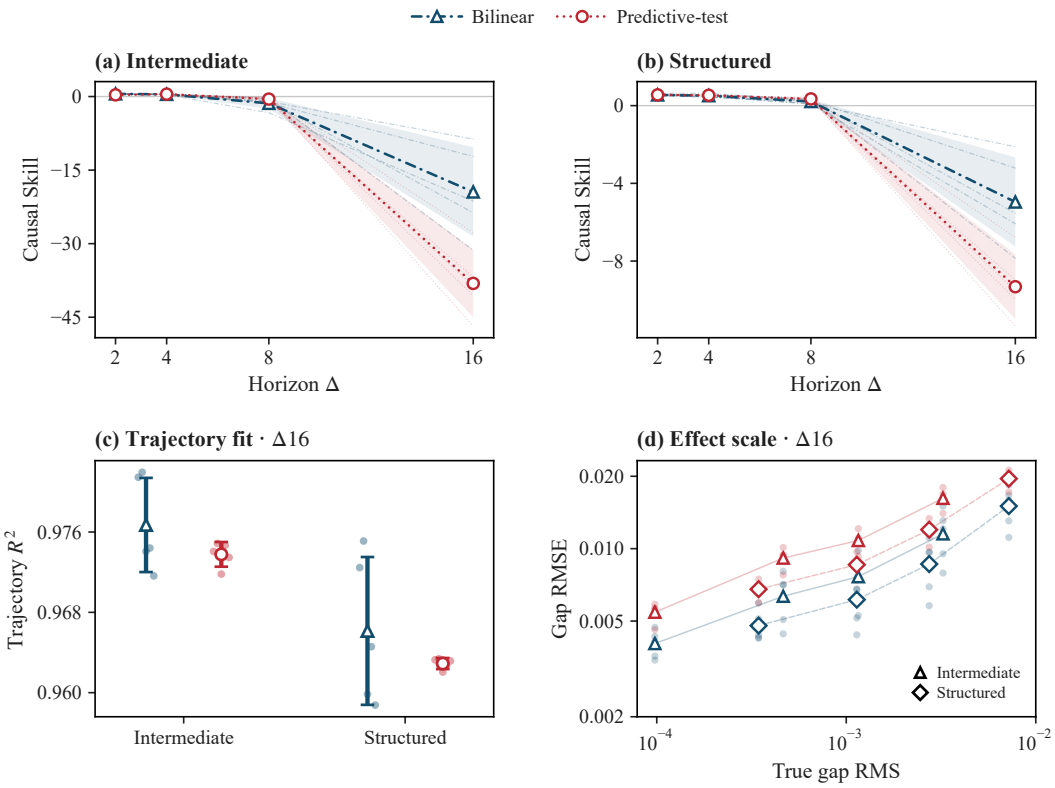

Figure 12: **Learned recurrent controls. (a,b)** Causal Skill, **(c)** trajectory R^2 at $\Delta = 16$, and **(d)** response-difference error versus effect scale. Means are over five initializations; shading/error bars show one sample standard deviation, and faint traces show individual initializations. Panel (d) uses grouped effect-size summaries.

G Recursive-state baselines

G.1 Learned recurrent controls

Both learned controls use 16-dimensional affine–bilinear recurrent states, continuous intervention descriptors, and the same validation horizons as OCPR. The predictive-test variant additionally predicts future-test coordinates.

Held-out results are shown in Figure 12. Both controls retain high trajectory R^2 while losing matched-history response-difference geometry at held-out horizons.

G.2 External OQE/system–memory baseline

The open quantum evolution (OQE)/system–memory model of Zhang et al. (2025) uses shared stationary system–memory evolution. We adapt its replacement intervention as $\rho_{SM} \mapsto \varrho_S \otimes \mathrm{Tr}_S(\rho_{SM})$, preserving the inferred memory marginal while fixing the system marginal. We test memory dimension $\chi \in \{1, \ldots, 6\}$ with five initializations and select by development ordinary-response MSE only; $\Delta = 16$ and depth-three targets are excluded from fitting and state access. Because OQE uses a different response representation and information interface, it is treated as an external method comparison rather than a realization ablation with matched response access.

In Table 6, the "Hardest $(4, 4)$" endpoint is the held-out depth-three timing pattern with four time steps between each pair of successive interventions.

For compactness, the "OQE Skill" and "C–G Skill" columns in Table 6 report Causal Skill as defined in Equation (2).

Table 6: **External OQE/system–memory comparison on the 30-pair response-derived evaluation.** Center–gap OCPR is evaluated on the same comparison set.

Environment	Endpoint	OQE R^2	OQE Skill	Gap RMSE	Corr.	Recovery	C–G Skill
Intermediate	$\Delta 16$	0.891	−74.30	0.01949	0.273	8.90	0.958
	Depth 3	0.915	−10.15	0.01257	0.461	3.67	0.994
	Hardest $(4, 4)$	0.911	−48.69	0.01322	0.428	7.41	0.975
Structured	$\Delta 16$	0.811	−23.67	0.01717	0.196	5.07	0.979
	Depth 3	0.875	−1.06	0.01171	0.317	1.39	0.997
	Hardest $(4, 4)$	0.867	−5.62	0.01179	0.087	2.45	0.995

The complete OQE comparison is reported in Table 6. OQE retains substantial ordinary-response fit while Causal Skill remains negative at the reported endpoints; both environments select the largest tested $\chi = 6$.

H Additional analyses

H.1 Downstream decision relevance

Decision relevance is evaluated on an independent held-out set of 40 matched-history pairs, comprising 20 pairs from each environment. For six intervention targets, we report decision-value recovery

$$\mathrm{DVR} = 1 - \frac{\sum(U^\star - U_{\mathrm{model}})}{\sum(U^\star - U_{\mathrm{blind}})}, \tag{40}$$

where $U^\star$, U_{model}, and U_{blind} are the oracle, model-selected, and blind-reference decision values. Mean regret averages $U^\star - U_{\mathrm{model}}$ over all evaluated history–target decisions, and action accuracy is the fraction of those decisions for which the model selects the oracle action. Reversal recovery is computed only on matched-history/target units whose oracle-optimal actions differ between the two histories, and reports the fraction for which both oracle actions are recovered. Decision results are summarized in Table 7. At $\Delta = 16$, response-derived and history-inferred OCPR recover $12/12$ and $11/12$ such units, respectively.

Table 7: **Intervention selection at $\Delta = 16$ over 40 held-out matched-history pairs (20 per environment) and six targets per pair.**

Method	DVR	Mean regret	Action acc.	Reversal recovery
Generic conditioning	−49.740	2.29×10^{-3}	0.738	0.000
Absolute-response realization	0.826	7.89×10^{-6}	0.985	0.833
History-inferred OCPR	0.829	7.73×10^{-6}	0.988	0.917
Response-derived OCPR	0.998	1.08×10^{-7}	0.998	1.000

H.2 Computational characteristics

The history encoder has 88,392 parameters and the affine state map has 1,040 coefficients; each recursive update acts on a 16-dimensional predictive state.