CAUSAL ECHOES: TASK PRECISION THROUGH PHYSICAL CANCELLATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Control actions can improve the measurements used to learn a system. For unknown linear systems with signal-dependent sensing, paired inputs nearly cancel in the plant and leave a small residual with small absolute readout noise. Repeating this physical cancellation sharpens the information needed for a future control task. Under radial relative noise and zero additive floors, K actionable packets yield minimax orders σ^K for the learning-terminal state and σ^{K+1} for a future command using a final upload; full physical calibration remains at order σ in dimension at least two. Matching lower bounds allow adaptive packet timing and silence. A calibrated-drift gain family gives a finite separation from every noninteractive acquisition, including randomized designs, under nonzero readout and deployment floors. Held-out matrix results are consistent with the predicted orders and show how conditioning and noise floors limit the useful refinement depth.

1 INTRODUCTION

A controller needs the right input, which may require much less knowledge than a calibrated physical model. Task-oriented exploration already uses this distinction to decide where to measure (Wagenmaker et al., 2021). We study a further possibility: an action can make the *next measurement itself more precise*.

The opportunity arises with signal-dependent sensing. A small state has small absolute measurement noise. Using a rough model, we apply two inputs whose effects nearly cancel. Their residual measures the error in the cancellation; because that residual is small, it can reveal the error more accurately than the measurements used to fit the rough model. We call this operation *physical cancellation*.

A sensor observes the unknown system privately and sends at most K packets to an actuator. We ask how accurately this interaction can control the learning trajectory, learn a command for a later run, and calibrate the physical model. The three objectives have different precision limits because information can arrive too late to change an action yet still improve a future plan.

We first present this distinction for fully actuated systems with relative sensing noise. We then extend the construction to controllable systems and prove matching lower bounds for adaptive packet times, including information conveyed by silence. A finite scalar guarantee isolates the value of adaptive acquisition; matrix studies examine conditioning and noise floors.

2 THE SIMPLEST CASE: THREE GOALS, ONE LEARNING RECORD

Consider a fully actuated system

$$x_{t+1} = Ax_t + Bu_t, \qquad Y_t = x_t + \sigma\|x_t\|Z_t, \qquad Z_t \overset{\text{iid}}{\sim} \mathcal{N}(0, I).$$

Here $x_0 \neq 0$ is known, $n \geq 2$, and the unknown square matrices A, B belong to a compact class with nonempty interior and uniform inverse bounds. There are no additive noise floors. Fix the class, a finite moment order p, a depth $K \geq 1$, and sufficient learning time and all-path input energy.

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.

The core result. The first objective is the terminal state of the learning run. The second is the error in a minimum-energy zeroing command for a new run from the same x_0. Both use K actionable packets; the command estimator also receives a final upload. The third objective estimates the physical matrices (A, B) and is allowed immediate feedback after every reading. As $\sigma \to 0$, the optimal worst-case L^p errors have orders

$$\underbrace{\sigma^K}_{\text{learning-terminal state}}, \qquad \underbrace{\sigma^{K+1}}_{\text{future command error}}, \qquad \underbrace{\sigma}_{\text{physical matrices}}. \tag{1}$$

Command errors use Euclidean norm; matrix errors use Frobenius norm. The command may target any fixed horizon $H \geq 1$ in this invertible setting. Theorem 1 and Corollary 2 give the general statements.

Table 1: **One timeline for learning and deployment.** Observe, send and act repeat along one learning trajectory. The final upload and a new preparation follow that trajectory.

Stage	Information available	Effect on actions
1. Observe	The state generates a private sensor reading Y_t before u_t.	Readings remain in its buffer.
2. Send	At most K packets, with unlimited payloads; timing and silence carry information.	The actuator can update its plan.
3. Act	The actuator uses messages and silences; the sensor keeps observing.	Inputs change the state and the precision of later readings.
4. Final upload	After the last learning action, the reporter receives the stored record.	No learning action remains; the record improves the future command.
5. Deploy	Prepare the same known x_0 anew and apply the learned H-step command.	Evaluate a separate future run.

The last private query batch improves the command after its information can no longer affect the learning inputs. Uploading that batch reveals existing readings; it adds no observation. Thus the command uses $K + 1$ disclosures in total. With a total budget of N, use $N - 1$ actionable packets and one upload to obtain order σ^N.

2.1 Paired inputs reveal the cancellation error

A scalar example. Take $x_{t+1} = .8x_t + u_t$ and a coarse drift estimate .79. Starting from $x = .01$, the pair $(1, -.79)$ ends at $.8^2(.01) + (.8 - .79) = .0164$. With $\sigma = .01$, its readout standard deviation is .000164. Ignoring readout noise in this calculation,

$$\widehat{a}_{\text{new}} = .79 + (.0164 - .79^2 \cdot .01) = .800159.$$

The drift error drops from .01 to .000159: subtracting the predicted starting-state contribution exposes the error in the coarse model. The matrix construction applies this update in each input direction.

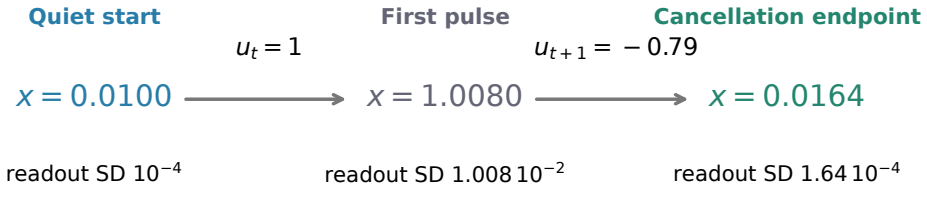


Figure 1: **Create a small residual before measuring it.** The first pulse moves the state away from zero; the second nearly cancels its effect. With relative readout noise, the final observation measures the cancellation error at high absolute precision.

In input coordinates $z_t = B^{-1}x_t$, the drift matrix becomes $M = B^{-1}AB$. Each pair tests one column of its current estimate. Suppose a prefix has supplied a coarse fit $(\bar{A}, \bar{B})$ whose errors are $O(\sigma)$, with bounded inverses. We first describe one refinement; the implementation safeguards follow below. The actuator first corrects using $-\bar{B}^{-1}\bar{A}Y$, with an all-path input cap. With current estimate $\widehat{M}$ of M and a chosen query amplitude $a > 0$, it executes the pair

$$u_t = ae_j, \qquad u_{t+1} = -a\widehat{M}e_j. \tag{2}$$

It freezes $\widehat{M}$ throughout the complete batch of n pairs. The exact endpoint is

$$x_{t+2} = A^2x_t + aB(M - \widehat{M})e_j. \tag{3}$$

When $\widehat{M}$ approximates M, the pair leaves the propagated starting state and a small cancellation error. After a correction, the start x_t and endpoint x_{t+2} are close to zero; we call these states *quiet*. Their small absolute readout noise makes the residual in Eq. (3) informative about $M - \widehat{M}$. The sensor combines the cached endpoint readings as

$$\widetilde{M}e_j = \widehat{M}e_j + a^{-1}\bar{B}^{-1}(Y_{t+2} - \bar{A}^2Y_t). \tag{4}$$

The subtraction has two errors: the coarse calibration error multiplies an already small state or cancellation error, and the endpoint reading has small absolute noise. If the state and coordinate errors have order σ^r, then

$$\text{next coordinate error } \lesssim \sigma\,\sigma^r + \sigma\,\sigma^r = O(\sigma^{r+1}). \tag{5}$$

Each nonfinal batch sends its update and a current reading in the next packet, enabling another coarse-model correction. Thresholded prefix regression, spectral clipping and a fixed fallback provide the bounded estimates used above (Appendix G.2). Appendix B proves the rate with dependent model and state errors, capped inputs and projection.

2.2 From the learning run to a future operation

After the last batch, the reporter turns the learned input-coordinate dynamics into the future command $\widehat{U}_H$ in Table 1.

The physical preparation x_0 is known, but the command calculation needs $z_0 = B^{-1}x_0$, which is unknown because B is unknown. Coarse decoding of x_0 has only $O(\sigma)$ accuracy. Instead, retain the quiet reading at $s = s_K$, just after the last correction. Since $x_s = O(\sigma^K)$, $\widehat{z}_s = \bar{B}^{-1}Y_s$ has error $O(\sigma^{K+1})$. The refined dynamics and actual input history recover z_0 by reversing $z_{t+1} = Mz_t + u_t$. For fixed s, uniform inverse bounds preserve this precision (Appendix B.8). Then

$$\widehat{U}_H = -\mathcal{C}_H(\widehat{M}_K, I)^\dagger \widehat{M}_K^H \widehat{z}_0, \qquad \mathcal{C}_H(M, I) = [M^{H-1}, \dots, I]. \tag{6}$$

The input-coordinate Gramian is at least I, so this final solve preserves the extra noise order. Spectral clipping and matrix solves cost $O((K + H)n^3 + Tn^2)$ after prefix regression (Proposition 16). The general construction in Section 5.2 uses forward powers to recover the needed task displacement even when A is singular.

3 General systems and observation models

An unknown pair $\theta = (A, B)$ lies in a known compact class $\Theta \subset \mathbb{R}^{n\times n} \times \mathbb{R}^{n\times m}$ of controllable plants. The initial state $x_0 \neq 0$ is known. For T transitions,

$$x_{t+1} = Ax_t + Bu_t + \tau\xi_t, \qquad Y_t = x_t + \sigma V_{\beta,d}(x_t)Z_t, \tag{7}$$

where ξ_t, Z_t are independent standard Gaussian vectors. We use either radial sensing, $V_{\beta,d}(x) = \sqrt{d^2 + \|x\|^{2\beta}}I$, or coordinatewise sensing, whose ith diagonal entry is $\sqrt{d^2 + |x_i|^{2\beta}}$. Here $\beta, d \in [0, 1]$; at $\beta = 0$ the state-power term is one, including at zero. The absolute additive readout standard deviation is σd; τ is the process-noise scale. The known observation mean is the physical state. At $\beta = 1$, independent Gaussian channel fluctuations in $Y = (I + \sigma G)x + \sigma dZ$ give the radial law; diagonal G gives its coordinatewise counterpart. Gain fluctuations and additive components coexist in electronic readers (Weinreb & Schleeh, 2014), and ADC jitter also motivates separating signal-dependent error from a floor (Brannon & MacLeod, 2009). We use an independent Gaussian full-state model for these two noise components. The signal-dependent component can be reduced by cancellation; the additive component sets a floor.

Information and control. At clocks $1, \dots, T-1$, the sensor receives Y_t before u_t is chosen. It may transmit or remain silent based on its private readings, past messages, clock and random seed. A packet has unlimited precision and may reveal every reading so far. The actuator selects commands from its seed, messages and all silences; between packets it can execute a previously specified word. Policies in $\Pi_K(E)$ satisfy

$$\#\{\text{packets}\} \le K, \qquad \sum_{t=0}^{T-1} \|u_t\|^2 \le E \quad \text{on every path.} \tag{8}$$

After the final action, the sensor uploads its record to an offline reporter for future command estimation. Corollary 2 counts this upload within the disclosure budget.

The task to be learned. For the general theorem, the future task has horizon $H = T \ge n$:

$$\mathcal{U}_T^*(\theta) = -\mathcal{C}_T(\theta)^\dagger A^T x_0, \qquad \mathcal{C}_T(\theta) = [A^{T-1}B, \dots, B]. \tag{9}$$

This minimum-energy word zeros a fresh noiseless preparation. The losses measure learning state, command, and physical-model error:

$$\begin{aligned} D_{K,p} &= \inf_{\pi \in \Pi_K(E)} \sup_{\Theta} \|x_T\|_{L^p}, \\ Q_{K,p} &= \inf_{\pi \in \Pi_K(E), \widehat{U}} \sup_{\Theta} \|\widehat{U} - \mathcal{U}_T^*\|_{L^p}, \\ F_p &= \inf_{\pi \in \Pi_{T-1}(E), \widehat{\theta}} \sup_{\Theta} \|\widehat{\theta} - \theta\|_{L^p}. \end{aligned} \tag{10}$$

Command errors use Euclidean norm and matrix errors use combined Frobenius norm. The physical-model estimator has immediate feedback after every reading. All moment orders $1 \le p < \infty$ are fixed.

Uniform geometry and resources. Choose $h \le n$ such that $\mathcal{C}_h(\theta)$ has a common positive least singular value on Θ. Compact controllability ensures this with $h = n$. A deterministic prefix of length L supplies a uniformly full-row-rank state–input regressor. Persistent excitation supplies such a prefix for every compact controllable class (van Waarde et al., 2020a); a shorter certified word may be used. We assume nonempty interior for the matching lower bounds. This supplies an interior gain segment (A_*, bB_*) with $A_*^T x_0 \ne 0$ and, for $n \ge 2$, a nontrivial similarity orbit fixing x_0. These conditions allow unstable systems and singular A or B.

4 The precision hierarchy

Put $e_1 = 1$, $e_{j+1} = 1 + \beta e_j$, and $R_j = \sigma^{e_j} + \sigma d + \tau$. Write $d_h = mh$ for the dimension of an h-step input word.

Theorem 1 (Learning a plan, executing control, and calibrating a model)**.** *Fix the class and preparation above, a common identifying prefix, and a supported depth*

$$T \ge L + Kh(d_h + n + 1) + 1. \tag{11}$$

There is a finite all-path allowance E_{echo}, independent of the supported depth. For every fixed finite $E \ge E_{\mathrm{echo}}$ and sufficiently small noise levels, both readers satisfy

$$D_{K,p} \asymp R_K, \qquad Q_{K,p} \asymp R_{K+1}. \tag{12}$$

For physical calibration, radial relative sensing ($\beta = 1$) and $n \ge 2$ give $F_p \asymp \sigma + \tau$. The same result holds for coordinatewise sensing if x_0 has a zero coordinate.

One K-packet policy attains all applicable upper bounds without knowing the noise parameters. Constants depend on the fixed class, resources and moment order, but not on σ, d, τ, β. All lower bounds allow adaptive packet times, private readings, silence and arbitrary payloads.

Reading Theorem 1. $D_{K,p}$ compares policies by the state they actually reach with K actionable packets. $Q_{K,p}$ estimates a minimum-energy zeroing command for a fresh run from the same x_0; it additionally uses the final stored record, disclosed after the learning actions end. F_p measures physical-matrix error even with immediate feedback; its stated law requires relative sensing and $n \geq 2$, plus a zero coordinate of x_0 for the coordinatewise reader. The task rates apply to both readers on the compact controllable class, with class geometry, time, energy and depth held fixed as noise decreases.

Under relative sensing, each actionable update adds one power of σ. For $0 < \beta < 1$, $e_K = (1 - \beta^K)/(1 - \beta)$; additive sensing gives exponent one. Refinement saturates at the common floor $\sigma d + \tau$.

The command statement uses $H = T$ and allows singular A. With uniformly invertible A, it also holds at any fixed $H \geq h$ (Appendix B.8). Proofs are in Appendices A, B and D.

Corollary 2 (A fixed total disclosure budget)**.** *With the corresponding sufficient fixed time and energy, restricting the command estimator to the complete K-packet public transcript gives $Q^{\mathrm{pub}}_{K,p} \asymp R_K$. The same law holds for $K - 1$ actionable packets followed by a charged final upload. Each estimator uses only the resulting public transcript, including packet times and silences.*

Use $K - 1$ corrections and transmit the last batch in packet K. The converse counts K stopped histories (Appendix G).

5 Which tasks and systems inherit the precision?

5.1 The task representation

For fully actuated systems with uniform inverse margins for A, B, the construction learns the dynamics and preparation in input coordinates:

$$\mathcal{Z} = (M, z_0) = (B^{-1}AB, B^{-1}x_0). \tag{13}$$

It is invariant under $(A, B) \mapsto (SAS^{-1}, SB)$, $Sx_0 = x_0$; equal quotients imply this relation with $S = B_2 B_1^{-1}$. A small family of recovery operations determines it completely.

Proposition 3 (A finite catalog determines the task representation)**.** *Prescribe a first input v from x_0 and let $c_\theta(v)$ be the unique second input that zeros the state. Then*

$$c_\theta(v) = -M^2 z_0 - Mv. \tag{14}$$

The $n + 1$ commands at $v = 0, e_1, \ldots, e_n$ are bi-Lipschitz equivalent to $\mathcal{Z}$ on the fixed class. Their minimax error is R_{K+1} with the final private record, or R_K with K public disclosures. Every representation determining this family determines $\mathcal{Z}$.

Subtracting $c_\theta(0)$ from the other commands recovers $-M$; then $z_0 = -M^{-2}c_\theta(0)$. The catalog consists of targets estimated from the original learning record (Appendix G.4). Every Lipschitz function $g = \psi(\mathcal{Z})$ inherits the upper rate, with a matching lower rate when its input-gain derivative is nonzero. This includes weighted input-energy zeroing and minimum-effort waypoints $x_{t_j} = c_j x_0$. Thus $\mathcal{Z}$ characterizes task responses even though equivalent systems may have different physical observation laws.

The representation also clarifies what a physical-model estimator can achieve. Given a coarse $\bar{B}$ and refined $(\widehat{M}, \widehat{z}_0)$, a rank-one adjustment enforcing $\widetilde{B}\widehat{z}_0 = x_0$, followed by $\widetilde{A} = \widetilde{B}\widehat{M}\widetilde{B}^{-1}$, produces exactly the canonical command whenever the adjusted $\widetilde{B}$ is invertible. Appendix G.5 gives the explicit adjustment and the additional margin for membership in the physical class. Thus accurate physical *calibration* and accurate physical-model *plug-in commands* have different requirements.

A task changing to first order along such a similarity has an $\Omega(\sigma)$ obstruction, even with immediate feedback. Physical-state quadratic penalties can change on this orbit; weighted input-energy zeroing does not. Appendix G.3 proves the criterion and gives an explicit quadratic-control example.

5.2 Underactuation and singular drift

An underactuated system can perform the same cancellation over several time steps. Set $F = A^h$ and $D = \mathcal{C}_h(A, B)$. An h-step input sequence $w \in \mathbb{R}^{d_h}$ produces $x^+ = Fx + Dw$, together with

accumulated process noise. Controllability supplies n columns of D that span the state space. Collect them with a selector J, set $C = DJ$, and work in the coordinates

$$\mathsf{M} = C^{-1}FC, \qquad \mathsf{N} = C^{-1}D, \qquad \mathsf{N}J = I. \tag{15}$$

A selected-column pair $aJe_j, -aJ\widehat{\mathsf{M}}e_j$ has endpoint $F^2x + aC(\mathsf{M} - \widehat{\mathsf{M}})e_j$. For each unselected column, the *single* word $a(e_\ell - J\widehat{\mathsf{N}}e_\ell)$ has endpoint $Fx + aC(\mathsf{N} - \widehat{\mathsf{N}})e_\ell$. There are n pairs and $d_h - n$ single words: a batch costs $h(d_h + n)$ transitions. The same endpoint refinement estimates both matrices to the next order using the retained coarse decoder.

Write $\mathsf{N} = [N_1, \dots, N_h]$ in blocks of m columns. The one-step dynamics in these coordinates are $A_c = C^{-1}AC$ and $B_c = C^{-1}B$. Shifting the reachability blocks gives

$$A_c\mathsf{N} = [\mathsf{M}N_h, N_1, \dots, N_{h-1}]. \tag{16}$$

Multiplying by J recovers A_c because $\mathsf{N}J = I$; the one-step input matrix is $B_c = N_h$. These identities use only nonnegative powers, so they also apply when A is singular. Joint projection of the two refined matrices onto their compact realizable image preserves the identities and controllability. Geometrically decreasing query amplitudes keep the total energy finite. This projection need not be convex: the general construction establishes statistical achievability, while the polynomial-time implementation applies to the fully actuated setting of Section 2.

The coarse fit selects and retains a maximum-volume minor. Higher prefix moments and bounded reports control poorly conditioned selections (Appendices B.2–B.5). Appendix C works through a singular three-state, two-input example with both types of query.

5.3 Recovering the command from a quiet reading

To compute the T-step command, we need the uncontrolled terminal displacement $A_c^T C^{-1}x_0$. A quiet reading stored after the last correction and the known input history recover it without recovering $C^{-1}x_0$ itself. Let s be that quiet clock. Using the coarse estimate $\bar{C}$, set $\widehat{z}_s = \bar{C}^{-1}Y_s$. The coarse decoder error multiplies a state of order R_K, so $z_s = C^{-1}x_s$ is estimated to order R_{K+1}. Subtracting the known input history gives, for $T \geq s$,

$$c_T = A_c^{T-s}z_s - \sum_{t<s} A_c^{T-1-t}B_c u_t = A_c^T C^{-1}x_0 + \text{process contribution}. \tag{17}$$

The refined dynamics give $\widehat{c}_T$ and the command $-\mathcal{C}_T(\widehat{A}_c, \widehat{B}_c)^\dagger \widehat{c}_T$ with error R_{K+1}. The process contribution is $O(\tau)$; all powers are nonnegative.

6 Lower bounds for adaptive interaction

The sensor can accumulate private readings and reveal information through both a transmission time and continued silence. Comparing only message contents would miss these channels. We instead compare the entire experiment up to the next packet, including its covariance and stopping information, and then retain the policy's true continuation.

A gain interval survives each disclosure. On an interior family (A_*, bB_*), a shared numerical history and subsequent silence produce states affine in b. The input budget bounds their slopes. A fixed positive portion of this experiment still cannot distinguish gains in a short interval after the next packet. An auxiliary comparison expresses this portion as a mixture of uniform intervals. Conditional on its center, each interval is independent of the disclosed history and has halfwidth

$$w^+ = \zeta \min\{w, \sigma(d + w^\beta)\}, \qquad \zeta > 0. \tag{18}$$

Here ζ depends on the resources. To repeat the argument, first append the actual controlled continuation, then reveal the auxiliary center. This order preserves the future information available to the policy while leaving the actual posterior unrestricted (Appendix A.4).

After K packets, the affine terminal response has mean norm proportional to the surviving width, giving R_K. The final private buffer adds one disclosure. Along the gain family, $\mathcal{U}_T^*(A_*, bB_*) = U_*/b$ with $U_* \neq 0$, so every estimate on an interval of halfwidth w has mean command error at least $\|U_*\|w/(2b_{\max}^2)$. After $K+1$ disclosures this gives R_{K+1}. Separate transition comparisons and the final independent innovation supply the command and terminal process-noise floors.

Physical calibration has a different obstruction. For $n \geq 2$, let $S = I + avv^\top$, where $\|v\| = 1$ and $v \perp x_0$. The plants (A_*, B_*) and (SA_*S^{-1}, SB_*) have equal zeroing commands but different physical matrices. Aligning their hidden states bounds transcript divergence by $Ca^2(\sigma^{-2} + 1)$ under every adaptive policy, independently of control amplitude. Testing gives physical-model error $\Omega(\sigma)$; perturbing the input matrix gives $\Omega(\tau)$ (Appendix D). The prefix fit matches their sum. Zeroing is invariant on this family, so its precision requires the separate task converse.

For a new preparation $y \notin \operatorname{span}\{x_0\}$, a similarity fixing x_0 can change the command for y. Even immediate feedback then leaves minimax command error $\sigma + \tau$ under radial relative sensing (Proposition 8). High precision is tied to the preparation used during learning.

7 Finite resources and numerical assessment

We examine noise orders, a finite acquisition guarantee, dimension and noise floors. Matrix studies summarize sampled systems using pooled RMS, with block standard errors conditional on each fixed panel; these are empirical summaries rather than continuous-class minimax guarantees. The scalar study gives uniform analytical bounds.

7.1 Noise orders on held-out two-state systems

We froze the policy, amplitudes and caps before drawing 64 two-state systems, independently of the 25-system design panel. All matrix entries vary in operator-norm balls of radius .05 (Appendix P). Each policy uses $T = 20$ and energy at most 7.168 on every path. Thirteen levels span $\sigma = .0002$ to .05, with $d = \tau = 0$ and four blocks of 128 antithetic paths per system. Slope fits use the first five levels, fixed before evaluation.

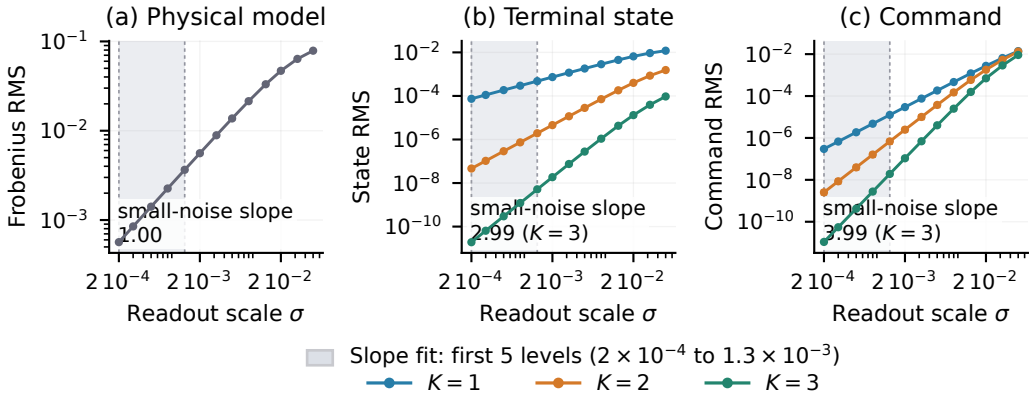


Figure 2: **Three losses from one learning record.** (a) Prefix model fit; (b) state at learning clock $T = 20$; (c) raw minimum-energy command error at $H = 3$. Axes show absolute errors under zero additive floors, pooled over the 64-system panel; error bars are block SEs. Shading marks the slope-fit interval.

At $K = 3$, the model, learning-terminal and command slopes are $1.00, 2.99, 3.99$ (Figure 2). In separate three-step deployments, the terminal-error slopes are $2.00, 2.99, 3.99$ for $K = 1, 2, 3$ at small noise, but $1.80, 2.34, 3.03$ on the last four levels. The integer orders emerge in the small-noise range; Appendix P gives the intermediate levels and guard events.

7.2 A finite scalar acquisition guarantee

Consider $x_{t+1} = x_t + bu_t$, $x_0 = 1$, $b \in [.5, 1.5]$, and $Y_t = x_t + \sigma\sqrt{d^2 + x_t^2}Z_t$. The drift is calibrated and learning has no process noise. Start with $\widehat{b}_0 = 1$, $u_0 = -1$. At packet clocks $j = 1, 2, 3$, use

$$\widehat{b}_j = \Pi_{[.5,1.5]}\{\widehat{b}_{j-1}(1 - cY_j)\}, \quad u_j = \widehat{b}_{j-1}^{-1} - \widehat{b}_j^{-1}, \qquad c = \frac{2}{1 + \sqrt{1 + 4\sigma^2}}. \tag{19}$$

The cumulative input is $-1/\widehat{b}_j$, so the next state is the relative gain error. At $d = 0$, cY_j is the unrestricted single-reading maximum-likelihood estimate of that state. At positive d the same shrinkage update admits a uniform mean-square error bound defined by

$$V_{j+1} = \{(1-c)^2 + c^2\sigma^2\}V_j + \tfrac{9}{4}c^2\sigma^2 d^2, \qquad V_0 = \tfrac{1}{4}, \qquad \mathbb{E}(\widehat{b}_j - b)^2 \le V_j. \tag{20}$$

The private reading at clock 4 updates only $\widehat{b}_4$; set $u_4 = 0$. Learning uses $T = 5$ transitions and energy at most $19/3$. Deploy three equal inputs $-1/(3\widehat{b}_4)$ from a new preparation, with energy at most $4/3$. At $\sigma = .05$, $d = .001$, preparation RMS 10^{-4} and independent deployment process scale 2×10^{-5}, the open-loop terminal RMS is at most 1.836×10^{-4}.

Every acquisition word chosen before sensing, including random designs independent of b, has worst-case open-loop deployment RMS exceeding 2.960×10^{-4}, with any decoder of its four readings and unlimited acquisition energy. The proof averages safe two-gain comparisons over a fixed prior independent of the design seed (Appendix J). Thus adaptive acquisition meets 2×10^{-4} in this calibrated scalar experiment, while every nonadaptive design misses it.

For fixed inputs, each state is $1 + bS_t$, where S_t is the cumulative input. The four observation means have at most four gain roots. For each design seed, a fixed-prior fraction of gain pairs stays away from these roots; adaptive inputs move the cancellation points toward the unknown gain.

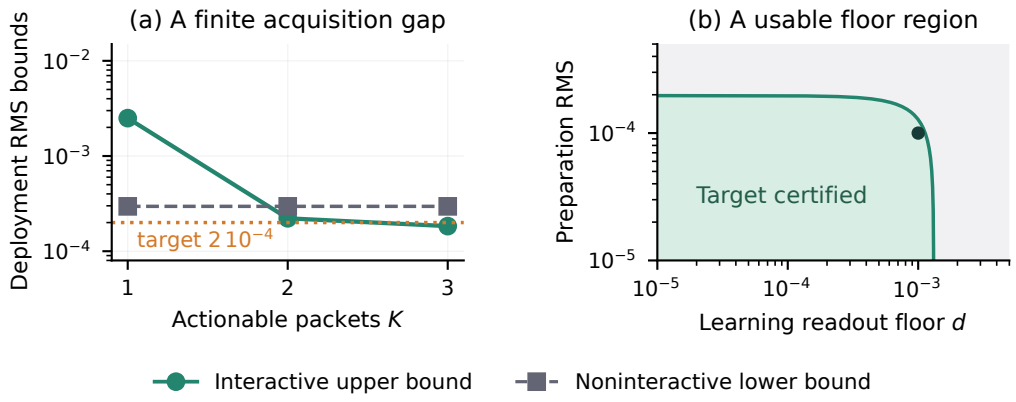


Figure 3: **A finite advantage over every noninteractive acquisition.** (a) Uniform bounds with four available readings, $T = 5$ and the common learning allowance $E = 19/3$. Interactive curves are sufficient upper bounds; the noninteractive line is a necessary lower bound, allowing randomization and unrestricted decoding. (b) Floor values for which the three-packet upper bound meets 2×10^{-4}; the dot is the example in the text. Here $\sigma = .05$ and deployment process scale is 2×10^{-5}. Both panels are analytical bounds.

7.3 Dimension and Reconstruction

We fixed the design using a development panel, then drew 16 test systems at each dimension $n = 4, 8, 16$. Both singular-vector frames vary, with $s(A) \subset [.995, 1.005]$, $s(B) \subset [.8, 1.2]$, and readout scale $\sigma = .01/\sqrt{n}$ (vector RMS scale .01).

Table 2: **Three-step deployment error on held-out dense systems.** Entries are pooled RMS (SE), in units of 10^{-3}. Each dimension uses 16 test systems and four independent blocks of 64 signed paths. SEs condition on the fixed panel. Learning uses $d = \tau = 0$; deployment starts exactly at x_0 without noise and adds energy at most two and one preparation to the learning budgets below (Appendix T).

Report	K	$n = 4$	$n = 8$	$n = 16$
Learning horizon T		44	84	164
Hard learning energy E		274	997	4032
Coarse-fit plan	any	20.95 (0.40)	24.94 (0.15)	28.97 (0.096)
Final LS ($K = 3$ record)	3	5.674 (0.032)	5.931 (0.0079)	6.505 (0.020)
Canonical refinement	1	28.18 (0.44)	94.89 (2.2)	216.2 (1.7)
Canonical refinement	2	2.342 (0.12)	10.77 (0.38)	27.25 (0.11)
Canonical refinement	3	**0.2602** (0.0076)	**1.587** (0.10)	**3.456** (0.023)

Three-packet refinement reduces deployment RMS relative to final LS on its $K = 3$ record by $21.8, 3.74, 1.88$ times at $n = 4, 8, 16$. The gain decreases with dimension; one-packet refinement is worse than the coarse-fit plan in all three cases.

7.4 Learning noise floors

Returning to the two-state policy in Figure 2, we use plants 0–15 from its 64-system test panel at $\sigma = .005$ and vary the readout and process floors separately; deployment starts exactly at x_0 without noise. At $d = \tau = 0$, three-packet deployment RMS is 5.26×10^{-6}. Setting $d = 10^{-5}$ raises it to 5.91×10^{-5}; at $d = .001$, the one-, two- and three-packet errors are $2.22 \times 10^{-4}, 5.02 \times 10^{-4}, 5.87 \times 10^{-3}$ (Appendix Q).

To isolate sensitivity to the stored reading, fix the true dynamics, exact decoder and input history. A perturbation ϵ in that reading changes the raw deployment endpoint by $-A^{H-s}\epsilon$. At the nominal system, the Frobenius amplification is $3.79, 48.26, 426.40$ at quiet clocks $s = 5, 10, 15$ (Appendix Q.1). Forward contraction in this system becomes error amplification when recovering the earlier preparation.

8 Related work and discussion

Task-oriented exploration and dual control. Task-optimal exploration designs experiments for downstream decisions (Wagenmaker et al., 2021), while data informativity can certify a task without identifying the system (van Waarde et al., 2020b; Baggio et al., 2019). These viewpoints motivate the task coordinates studied here. Dual control uses actions for both regulation and information (Klenske & Hennig, 2016; Ramesh et al., 2011). Iterative identification and numerical refinement correct residual errors (Campi et al., 2008; Carson & Higham, 2018); physical cancellation forms the residual in the plant before measurement, lowering its absolute readout noise.

Signal-dependent sensing and communication. State-dependent sensing couples control and estimation (Todorov, 2005; Ding et al., 2019). Multiplicative sensing can separate estimation from stabilization with known dynamics (Ranade & Sahai, 2013); here actions change readout precision while the dynamics are learned. Rate–cost theory prices information (Kostina & Hassibi, 2018), whereas event-triggered estimation uses silence as evidence (Han et al., 2015; Rabi et al., 2016), motivating our comparison of complete stopped histories. Batched learning also has round-dependent exponents (Gao et al., 2019); our rounds change the physical noise through cancellation.

Scope. The guarantees assume known full-state sensing. Unknown observation means, partial observation and finite-bit communication remain open.

Reproducibility statement

Section 3 defines the observation model, information interface and resource constraints. Appendices A–G give the theoretical arguments; Appendix G.4 details the operational task representation. Appendix H specifies the communication clocks, and Appendices I–I.1 give the finite certificates. Appendix J proves the scalar acquisition separation, and Appendix G.5 gives the task-preserving physical-model lift. The numerical appendices specify system classes, query inputs, guards, random streams, aggregation rules and uncertainty summaries. Appendices P and T describe the held-out calculations in the main figures and dimension table. Appendix Q gives the direct learning-floor study.

AI Use Statement

Generative AI tools were used solely to aid and polish the writing of this paper, including improving clarity, grammar, and readability. They were NOT used to originate research ideas, formulate mathematical claims, generate or analyze results, or draw scientific conclusions. All AI-assisted edits were reviewed by the authors. The authors take full responsibility for the final content of this work.

References

Giacomo Baggio, Vaibhav Katewa, and Fabio Pasqualetti. Data-driven minimum-energy controls for linear systems. *IEEE Control Systems Letters*, 3(3):589–594, 2019. doi: 10.1109/LCSYS.2019.2914090. URL https://arxiv.org/abs/1902.02228.

Xavier Bombois, Gérard Scorletti, Michel Gevers, Paul M. J. Van den Hof, and Roland Hildebrand. Least costly identification experiment for control. *Automatica*, 42(10):1651–1662, 2006. doi: 10.1016/j.automatica.2006.05.016. URL https://perso.uclouvain.be/michel.gevers/PublisMig/LCID_final.pdf.

Brad Brannon and Tom MacLeod. How ADIsimADC models an ADC. Technical Report AN-737, Rev. B, Analog Devices, 2009. URL https://www.analog.com/media/en/technical-documentation/application-notes/an_737.pdf.

Marco C. Campi, Toshiharu Sugie, and Fumitoshi Sakai. An iterative identification method for linear continuous-time systems. *IEEE Transactions on Automatic Control*, 53(7):1661–1669, 2008. doi: 10.1109/TAC.2008.929371.

Erin Carson and Nicholas J. Higham. Accelerating the solution of linear systems by iterative refinement in three precisions. *SIAM Journal on Scientific Computing*, 40(2):A817–A847, 2018. doi: 10.1137/17M1140819. URL https://eprints.maths.manchester.ac.uk/2629/1/cahi18.pdf.

Nicolas Chatzikiriakos, Kevin Jamieson, and Andrea Iannelli. High effort, low gain: Fundamental limits of active learning for linear dynamical systems. In Emtiyaz Khan, Yingzhen Li, Arno Solin, and Aaditya Ramdas (eds.), *Proceedings of the 29th International Conference on Artificial Intelligence and Statistics*, volume 300 of *Proceedings of Machine Learning Research*, pp. 3115–3123. PMLR, 2026. URL https://proceedings.mlr.press/v300/chatzikiriakos26a.html.

Milan S. Derpich and Serdar Yüksel. Dual effect, certainty equivalence, and separation revisited: A counterexample and a relaxed characterization for optimality. *IEEE Transactions on Automatic Control*, 68(2):1259–1266, 2023. doi: 10.1109/TAC.2022.3151189. URL https://doi.org/10.1109/TAC.2022.3151189.

Jian Ding, Yuval Peres, Gireeja Ranade, and Alex Zhai. When multiplicative noise stymies control. *The Annals of Applied Probability*, 29(4):1963–1992, 2019. doi: 10.1214/18-AAP1415. URL https://arxiv.org/abs/1612.03239.

Zijun Gao, Yanjun Han, Zhimei Ren, and Zhengqing Zhou. Batched multi-armed bandits problem. In *Advances in Neural Information Processing Systems*, volume 32, 2019. URL https://arxiv.org/abs/1904.01763.

Christopher Grimm, André Barreto, Satinder Singh, and David Silver. The value equivalence principle for model-based reinforcement learning. In *Advances in Neural Information Processing Systems*, volume 33, pp. 5541–5552. Curran Associates, Inc., 2020. URL https://papers.nips.cc/paper_files/paper/2020/hash/3bb585ea00014b0e3ebe4c6dd165a358-Abstract.html.

Duo Han, Yilin Mo, Junfeng Wu, Sean Weerakkody, Bruno Sinopoli, and Ling Shi. Stochastic event-triggered sensor schedule for remote state estimation. *IEEE Transactions on Automatic Control*, 60(10):2661–2675, 2015. doi: 10.1109/TAC.2015.2406975. URL https://yilinmo.github.io/public/papers/tac-13-event.pdf.

Håkan Hjalmarsson. From experiment design to closed-loop control. *Automatica*, 41(3):393–438, 2005. doi: 10.1016/j.automatica.2004.11.021. URL https://doi.org/10.1016/j.automatica.2004.11.021.

Edgar D. Klenske and Philipp Hennig. Dual control for approximate Bayesian reinforcement learning. *Journal of Machine Learning Research*, 17(127):1–30, 2016. URL https://jmlr.org/papers/v17/15-162.html.

Victoria Kostina and Babak Hassibi. Rate-cost tradeoffs in control. arXiv:1612.02126v3, 2018. URL https://arxiv.org/abs/1612.02126v3.

Béatrice Laurent and Pascal Massart. Adaptive estimation of a quadratic functional by model selection. *The Annals of Statistics*, 28(5):1302–1338, 2000. doi: 10.1214/aos/1015957395.

Horia Mania, Stephen Tu, and Benjamin Recht. Certainty equivalence is efficient for linear quadratic control. In *Advances in Neural Information Processing Systems*, volume 32, 2019. URL https://papers.neurips.cc/paper_files/paper/2019/hash/5dbc8390f17e019d300d5a162c3ce3bc-Abstract.html.

Ashutosh Nayyar, Aditya Mahajan, and Demosthenis Teneketzis. Decentralized stochastic control with partial history sharing: A common information approach. *IEEE Transactions on Automatic Control*, 58(7):1644–1658, 2013. doi: 10.1109/TAC.2013.2239000. URL https://arxiv.org/abs/1209.1695.

Samet Oymak and Necmiye Ozay. Non-asymptotic identification of LTI systems from a single trajectory. *arXiv preprint arXiv:1806.05722*, 2019. URL https://arxiv.org/abs/1806.05722.

Maben Rabi, Chithrupa Ramesh, and Karl Henrik Johansson. Separated design of encoder and controller for networked linear quadratic optimal control. arXiv:1405.0135v4, 2016. URL https://arxiv.org/abs/1405.0135v4.

Chithrupa Ramesh, Henrik Sandberg, Lei Bao, and Karl Henrik Johansson. On the dual effect in state-based scheduling of networked control systems. In *Proceedings of the American Control Conference*, pp. 2216–2221, 2011. doi: 10.1109/ACC.2011.5990925. URL https://people.kth.se/~kallej/papers/networked_control_ramesh_acc11.pdf.

Gireeja Ranade and Anant Sahai. Non-coherence in estimation and control. In *51st Annual Allerton Conference on Communication, Control, and Computing*, pp. 189–196, 2013. URL https://people.eecs.berkeley.edu/~gireeja/Papers/allerton13camerareadу.pdf.

Tuhin Sarkar and Alexander Rakhlin. Near optimal finite time identification of arbitrary linear dynamical systems. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 5610–5618, 2019. URL https://proceedings.mlr.press/v97/sarkar19a.html.

Max Simchowitz, Horia Mania, Stephen Tu, Michael I. Jordan, and Benjamin Recht. Learning without mixing: Towards a sharp analysis of linear system identification. In *Proceedings of the 31st Conference On Learning Theory*, volume 75 of *Proceedings of Machine Learning Research*, pp. 439–473, 2018. URL https://proceedings.mlr.press/v75/simchowitz18a.html.

Emanuel Todorov. Stochastic optimal control and estimation methods adapted to the noise characteristics of the sensorimotor system. *Neural Computation*, 17(5):1084–1108, 2005. doi: 10.1162/0899766053491887.

Henk J. van Waarde, Claudio De Persis, M. Kanat Camlibel, and Pietro Tesi. Willems' fundamental lemma for state-space systems and its extension to multiple datasets. arXiv:2002.01023, 2020a. URL https://arxiv.org/abs/2002.01023.

Henk J. van Waarde, Jaap Eising, Harry L. Trentelman, and M. Kanat Camlibel. Data informativity: A new perspective on data-driven analysis and control. *IEEE Transactions on Automatic Control*, 65(11):4753–4768, 2020b. doi: 10.1109/TAC.2020.2966717. URL https://arxiv.org/abs/1908.00468.

Andrew Wagenmaker and Kevin Jamieson. Active learning for identification of linear dynamical systems. In *Proceedings of Thirty Third Conference on Learning Theory*, volume 125 of *Proceedings of Machine Learning Research*, pp. 3487–3582. PMLR, 2020. URL https://proceedings.mlr.press/v125/wagenmaker20a.html.

Andrew J. Wagenmaker, Max Simchowitz, and Kevin Jamieson. Task-optimal exploration in linear dynamical systems. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 10641–10652. PMLR, 2021.

Sander Weinreb and Joel Schleeh. Multiplicative and additive low-frequency noise in microwave transistors. *IEEE Transactions on Microwave Theory and Techniques*, 62(1):83–91, 2014. doi: 10.1109/TMTT.2013.2293123.

Hongyi Zhou, Jingwei Li, and Jingzhao Zhang. Finite sample analyses for continuous-time linear systems: System identification and online control. In *Advances in Neural Information Processing Systems*, volume 38, 2025. doi: 10.52202/085713-1072. URL https://papers.neurips.cc/paper_files/paper/2025/hash/2df20ebc2bdaca5b8b48966cbe0ab3cf-Abstract-Conference.html.

Guide to the supplement. For a first technical reading, follow the construction in Appendix B, the stopped-history converse in Appendix A.4, and the finite scalar guarantee in Appendix J.

Main-text evidence. Figure 2: two-state noise orders in Appendix P; Figure 3: analytical bounds in Appendix J; Table 2: dense systems in Appendix T; Section 7: learning floors and reconstruction sensitivity in Appendix Q, including Q.1.

Claim to check	Proof route and essential condition
Learning-terminal upper bound	Appendices A.2–A.3: unconditional prefix moments, clipped block corrections and the energy bound.
Command upper bound	Appendices B.2–B.5: fitted chart, frozen queries, refinement moments and the bad-chart complement; the general construction requires $H \geq \max\{s, h\}$.
Adaptive disclosure lower bounds	Appendices A.4–A.5 and B.6: full stopped buffers, silence and actual controlled continuations on the gain segment.
Physical-model accuracy barrier	Appendices D.1–D.4: preparation-preserving similarity and input-gain witnesses, with the sensing and geometry conditions stated separately.
Final upload counted in the budget	Appendix G.1: public transcripts and charged uploads with the same total disclosure count.
Spectral computation and task scope	Appendix G.2 gives arithmetic cost under full actuation; Appendices G.3–G.4 give the task representation and operation catalog. Short horizons use Appendix B.8; another preparation uses Appendix B.9.
Task-preserving physical lift	Appendix G.5: exact command equality and physical-class feasibility margins.
Finite acquisition separation	Appendix J: clipped recurrence and a fixed-prior lower bound covering randomized noninteractive designs.
Finite sufficient certificate	Appendices I–I.1: radius recurrences, clipping and energy checks, and contraction of the comparison matrix.

Symbol crosswalk. The command risk carries a horizon subscript only when H differs from the learning horizon T. Block coordinates describe the same construction over h transitions.

Notation	Meaning and location
$D_{K,p}$	Learning-terminal minimax norm; main text and Appendix B.
$Q_{K,p} = Q_{K,p,T}$	Offline-command minimax error at $H = T$; Appendix B writes $Q_{K,p,H}$ for a general target horizon.
M, z_0	Fully actuated task coordinates: $M = B^{-1}AB$, $z_0 = B^{-1}x_0$. Superscripts identify particular reconstructions.
$\mathsf{M}, \mathsf{N}, A_c, B_c$	General controllability-chart coordinates in Appendix B.2; full actuation uses $C = B$, $\mathsf{M} = M$, $\mathsf{N} = I$.
T, H, h, K	Learning transitions, future task transitions, controllability block length, and actionable packet budget.
$d, \sigma d, \tau$	Relative floor parameter, absolute additive readout standard deviation, and process-noise scale.
q; $d_h = mh$	Prefix regressor margin in Appendix A; block input dimension in Appendix B.
Q, X, U, D in the held-out slope table	Respectively joint task-coordinate error, learning-terminal norm, command error and separately prepared deployment error; these local labels differ from the minimax risks above.

Additional calculations. Design-panel and posthoc deployment studies are in Appendices K–L; the prefix certificate, late-packet controls and one-input example follow in Appendices M–O. Appendix S examines likelihood fitting on fixed records. Appendix H collects the clocks and resource counts.

Statistical conventions. Evaluations retain all specified plants and records, including failures, fallbacks, projections and clips; diagnostic subsets are explicit. Signed paths are antithetic. Independent blocks give Monte Carlo SEs conditional on the fixed panel and policy, without adjustment for selecting maxima. Panel summaries are not continuous-class guarantees. Design, deployment and underactuated tables maximize plantwise RMS; held-out and dense tables pool squared losses before square roots; the likelihood diagnostic pools sixteen records. Deterministic bounds certify energy feasibility; sampled maxima summarize usage.

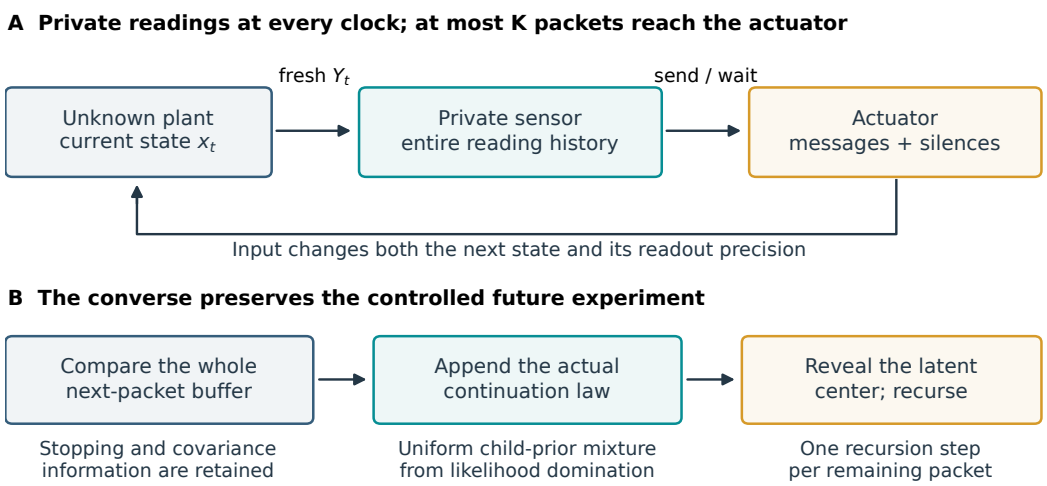


Figure 4: **Feedback depth counts packets, while retaining silence.** A: the sensor chooses whether to transmit after each private reading. Every command changes the state that produces later observations; the constructive policy uses only a subset of these opportunities. B: the converse compares the complete next-packet experiment and appends the policy's actual future kernel before granting an auxiliary center. This ordering permits a recursion per packet without replacing the future data by a fixed reference experiment.

A Depth on a compact controllable class

A.1 Experiment, resources, and a uniform statement

Finite-sample identification from trajectories is studied by Simchowitz et al. (2018); Sarkar & Rakhlin (2019); Oymak & Ozay (2019), and for sampled continuous-time dynamics by Zhou et al. (2025).

Let Θ be a known compact class of pairs $\theta = [A\ B] \in \mathbb{R}^{n\times(n+m)}$, every one of which is controllable. Starting from a known $x_0 \neq 0$, consider

$$x_{t+1} = Ax_t + Bu_t + \tau\xi_t, \qquad Y_t = x_t + \sigma V_{\beta,d}(x_t)Z_t, \qquad V_{\beta,d}(x) = \operatorname{diag}_{i=1}^{n}\sqrt{d^2 + r_i(x)^{2\beta}}. \tag{21}$$

The ξ_t, Z_t are independent standard n-Gaussians, independent of policy seeds. Readings arrive privately at clocks $t = 1, \dots, T-1$, before u_t. We consider either $r_i(x) = \|x\|$ for every coordinate (*radial sensing*) or $r_i(x) = |x_i|$ (*coordinatewise sensing*). At $\beta = 0$ define $V_{0,d}(x) = \sqrt{1+d^2}I_n$, including at zero. Throughout, $\beta, d \in [0,1]$, $0 < \sigma \leq 1/10$, and $0 \leq \tau \leq 1/10$. The sensor can send at most K packets, with arbitrary private-history dependent times and payloads. The actuator observes all packets and silences, and no other measurement. Every execution has $\sum_{t<T}\|u_t\|^2 \leq E$. When prescribed model accuracy is required, the final offline report must satisfy $\sup_{\theta\in\Theta}\mathbb{P}_\theta\{\|\widehat{\theta} - \theta\|_F > \epsilon\} \leq \delta$, where $0 < \epsilon, \delta \leq 1$. This is an overall randomized-policy condition, not a condition on each seed or transmission branch.

For this class let $D_{K,p}$ be the infimum, over energy-feasible policies, of $\sup_\theta(\mathbb{E}_\theta\|x_T\|^p)^{1/p}$, $1 \leq p < \infty$. The theorem below additionally certifies the prescribed report whenever its noise condition holds. The converse uses a segment

$$\{(A_*, bB_*) : b \in [1-w_0, 1+w_0]\} \subset \Theta, \qquad 0 < w_0 \leq 1/4, \qquad A_*^T x_0 \neq 0. \tag{22}$$

Nonempty interior of Θ suffices for such a segment; an explicit construction appears in Appendix E. Full affine hull alone is not used as a substitute for Eq. (22).

Fix a deterministic prefix of L commands, and denote its noiseless states at θ by $\bar{x}_t$. Assume the certified margins

$$W_\theta = \begin{bmatrix} \bar{x}_0 & \cdots & \bar{x}_{L-1} \\ u_0 & \cdots & u_{L-1} \end{bmatrix}, \quad q \le \min_\Theta s_{\min}(W_\theta),$$
$$\mathcal{C}_h(\theta) = [A^{h-1}B, \dots, B], \qquad s_h \le \min_\Theta s_{\min}(\mathcal{C}_h(\theta)), \qquad q, s_h > 0. \tag{23}$$

All stacks retain chronological column order. Taking $h = n$ always gives a positive margin by compactness. No stability assumption is needed for the fixed finite horizon below. Fix also a known nonempty bounded closed convex set $\mathcal{X}$ containing every noiseless terminal prefix state $\bar{x}_L$, and a number $r_0 > 0$. These are design choices, independent of the noise. An origin-centered ball containing the prefix states and $r_0 = 1$ will recover the default construction.

Theorem 4 (Compact-class depth law with explicit resource bounds)**.** *For either sensing model, fix the class, preparation, prefix, margins, $\mathcal{X}$ and r_0 above, and $1 \le K \le \lfloor (T-L)/h \rfloor$. The formulas below give finite E_0, C_{tail} and positive finite constants such that, for each* fixed finite *budget $E \ge E_0$, if $C_{\mathrm{tail}}(\sigma+\tau)\sqrt{\log(2L/\delta)} \le \epsilon$, then*

$$c_{\Theta,T,K,E} R_K \le D_{K,p} \le C_{\Theta,T,K,p} R_K, \qquad R_k = \sigma^{e_k} + \sigma d + \tau, \quad e_0 = 0, \quad e_{k+1} = 1 + \beta e_k. \tag{24}$$

The constants may depend on these fixed design choices and are uniform in the four noise parameters. One policy attains the upper bounds without knowing them, and uses at most E_0 energy on every record, independently of depth. The lower bound drops the report restriction and allows arbitrary adaptive schedules on all $T-1$ clocks. Its constant depends on the declared budget E.

A common prefix exists. Put $k = n+1$, $L = (m+1)k - 1$, and set $u_{jk-1} = e_j$ for $j = 1, \dots, m$, with every other prefix input zero. The order-k input Hankel matrix has mk rows and $L - k + 1 = mk$ columns. A window beginning at $s \in [(j-1)k, jk-1]$ contains exactly the impulse e_j; its block-row position runs through all k possibilities. The Hankel matrix is therefore a permutation of I_{mk}. The state-space fundamental lemma, with trajectory length one, gives full row rank of W_θ for every controllable plant (van Waarde et al., 2020a, Theorem 1(i)). Continuity and compactness give a positive common q. The prefix uses exactly m energy. Its regressor stops at $\bar{x}_{L-1}$; its response stack includes the additional reading Y_L, available at the first packet. A shorter prefix with certified Eq. (23) is equally valid.

A.2 Finite geometry, an all-record fit, and a report

Choose finite certified power and parameter bounds

$$\bar{A} \ge \sup_\Theta \|A\|, \quad \bar{B} \ge \sup_\Theta \|B\|, \quad P_j \ge \sup_\Theta \|A^j\| \ (0 \le j \le T), \qquad P_0 = 1.$$

For example, $P_j = \max\{1, \bar{A}\}^j$ is sufficient. Let $E_{\mathrm{pref}} = \sum_{t<L} \|u_t\|^2$ and define

$$H_s = \left(\sum_{j=0}^{s-1} P_j^2 \right)^{1/2}, \quad H_0 = 0, \qquad X = \max\left\{ 1, \max_{j \le L} P_j \|x_0\| + \bar{B} H_L \sqrt{E_{\mathrm{pref}}} \right\},$$
$$M = \sqrt{\bar{A}^2 + \bar{B}^2}, \qquad \Delta = \mathrm{diam}_F(\Theta), \qquad F_0 = \max\{4(1+M)/q,\ 2\Delta/q\}. \tag{25}$$

Thus $\|\bar{x}_t\| \le X$ on the prefix, and its process perturbation has covariance at most $\tau^2 H_L^2 I$. One may use $\Delta \le 2\sqrt{n}\, M$.

Set $Y_0 = x_0$ and form the observed stacks $W_Y = [Y_0, \dots, Y_{L-1}; u_0, \dots, u_{L-1}]$ and $X_Y^+ = [Y_1, \dots, Y_L]$. Fix $\theta_{\mathrm{ref}} \in \Theta$. If $s_{\min}(W_Y) \ge q/2$, take a measurable nearest minimizer in Θ of $X_Y^+ W_Y^\dagger$; otherwise return θ_{ref}. Compactness ensures that a minimizer exists. This rule may require nonconvex optimization, but is defined on every record and uses no noise parameter.

Write $\mathcal{H}^2 = \sum_{t=1}^L \|Y_t - \bar{x}_t\|^2$. On the first branch, the exact noiseless regression identity and $W_Y W_Y^\dagger = I_{n+m}$ imply $\|X_Y^+ W_Y^\dagger - \theta\|_F \le 2(1+M)\mathcal{H}/q$. A nearest projection onto a nonconvex set costs a factor of two: the true θ is an available comparison point. On the fallback branch, Weyl's

inequality gives $\mathcal{H} \geq \|W_Y - W_\theta\| > q/2$, while its parameter error is at most Δ. Consequently every record satisfies

$$\|\widehat{\theta} - \theta\|_F \leq F_0 \mathcal{H}. \tag{26}$$

Equation (26) holds on both the regression and fallback branches, with dependence between regressor and response errors retained.

For either sensing model the pointwise envelope

$$\|V_{\beta,d}(x)z\| \leq v_{\beta,d}(x)\|z\| \leq (d + \|x\|^\beta)\|z\|, \qquad v_{\beta,d}(x) = \sqrt{d^2 + \|x\|^{2\beta}}, \tag{27}$$

holds, with the declared constant convention at $\beta = 0$. The readout error is conditionally centered, so its cross term with the prefix process perturbation vanishes. Concavity of the radial power in this envelope gives

$$\mathbb{E}\|\widehat{\theta} - \theta\|_F^2 \leq F_0^2 L n\{H_L^2(1 + n\sigma^2)\tau^2 + (1 + X^2)\sigma^2\}. \tag{28}$$

Indeed $\mathbb{E}\|x_t - \bar{x}_t\|^2 \leq nH_L^2\tau^2$ and $\mathbb{E}v_{\beta,d}(x_t)^2 \leq 1 + X^2 + nH_L^2\tau^2$. An explicit sufficient reporting constant is

$$C_{\mathrm{rep}} = \max\{n,\ F_0^2 L n \max(2H_L^2, 1 + X^2)\}. \tag{29}$$

The alternative condition $C_{\mathrm{rep}}(\sigma^2 + \tau^2) \leq \delta\epsilon^2$ implies $n\sigma^2 \leq 1$, so Eq. (28) is at most $C_{\mathrm{rep}}(\sigma^2 + \tau^2)$; Markov's inequality proves the report.

A confidence bound for the same fit. The Gaussian prefix structure gives a sharper confidence dependence. Set $v_\delta = \log(2L/\delta)$ and $a_\delta = \sqrt{n} + \sqrt{2v_\delta}$. Then

$$\mathbb{P}_\theta\Big\{\|\widehat{\theta} - \theta\|_F > F_0\sqrt{L}\left[(\sigma(d + X) + \tau H_L)a_\delta + \sigma\tau H_L a_\delta^2\right]\Big\} \leq \delta. \tag{30}$$

To prove this, the prefix process perturbation $g_t = \tau\sum_{s<t} A^{t-1-s}\xi_s$ has covariance bounded by $\tau^2 H_L^2 I$. The Gaussian norm inequality (Laurent & Massart, 2000, Lemma 1) $\mathbb{P}\{\|Z\| > \sqrt{n} + \sqrt{2v}\} \leq e^{-v}$ and a union bound give $\max_{t\leq L}\|g_t\| \leq \tau H_L a_\delta$ with failure at most $\delta/2$. This step does not require independence across prefix times. Condition on the full prefix process vector. Because the prefix commands are deterministic, all states are now fixed. On the event just obtained, the operator norm of either reader scale is at most $d + X + \tau H_L a_\delta$, uniformly in $\beta \in [0, 1]$. A conditional Gaussian norm bound and a second union bound therefore bound all reader errors by $\sigma(d + X + \tau H_L a_\delta)a_\delta$, with conditional failure at most $\delta/2$. Integrating over the process event proves a joint probability of at least $1 - \delta$. Adding the process and reader errors, stacking the L columns, and using Eq. (26) proves Eq. (30). Singular reader covariances, including point masses at zero, satisfy the same norm bound. The mixed term is retained because unconditional signal-dependent reader errors need not be Gaussian.

For an explicit simpler constant, define

$$\begin{gathered} c_n = \sqrt{n/\log 2} + \sqrt{2}, \qquad C_g = F_0\sqrt{L}\max\{c_n \max(X, H_L), c_n^2 H_L\}, \\ C_{\mathrm{tail}} = \max\{1, 9C_g/4\}. \end{gathered} \tag{31}$$

The error threshold in Eq. (30) is at most $C_g[\{\sigma(1 + d) + \tau\}\sqrt{v_\delta} + \sigma\tau v_\delta]$. Write $r = (\sigma + \tau)\sqrt{v_\delta}$. For $d \leq 1$, $\epsilon \leq 1$ and $C_{\mathrm{tail}} r \leq \epsilon$, this threshold is at most $C_g(2r + r^2/4) \leq 9C_g r/4 \leq \epsilon$. This proves the reporting condition in Theorem 4. Equivalently, for $0 < \epsilon \leq 1$,

$$\sup_\theta \mathbb{P}_\theta\{\|\widehat{\theta} - \theta\|_F > \epsilon\} \leq \min\{1, 2L\exp[-\epsilon^2/(C_{\mathrm{tail}}^2(\sigma + \tau)^2)]\}.$$

The same fit and controller therefore satisfy an unconditional risk bound, with reporting failures included in the terminal-error calculation.

The same estimator satisfies the following moment bounds. For $r \geq 1$ put $r_* = \max\{r, 2\}$ and

$$\begin{gathered} z_r = \left\{2^{r_*/2}\frac{\Gamma((n + r_*)/2)}{\Gamma(n/2)}\right\}^{1/r_*}, \qquad J_r = X + \tfrac{1}{10}H_L z_r, \qquad S_r = z_r(1 + J_r), \\ F_r = F_0\sqrt{L}(H_L z_r + S_r), \qquad \|\widehat{\theta} - \theta\|_{L^r(F)} \leq F_r(\sigma + \tau). \end{gathered} \tag{32}$$

Fresh Gaussian independence bounds the readout moment by σS_r, and the prefix process moment by $\tau H_L z_r$. These bounds retain their displayed dimension factors for coordinatewise sensing. The

pointwise envelope transfers the upper proof; no ordering of the two controlled statistical experiments is asserted. First lift to r_*, apply Minkowski to the sum of squared stack norms, then lower the moment order. Lyapunov's inequality also covers $\beta r < 1$. Thus the square-root factor in L remains valid for $r < 2$.

A.3 Block corrections, energy, and the upper rate

At time L retain this fit and define

$$G = \widehat{\mathcal{C}}_h^{\dagger} \widehat{A}^h, \qquad \Gamma_h = P_h / s_h, \qquad \|G\| \le \Gamma_h. \tag{33}$$

Send at $L, L+h, \dots, L+(K-1)h$. The first packet carries the prefix fit and the current reading; subsequent packets carry the current reading. Commit the first h-input word $-G\Pi_{\mathcal{X}}(Y_L)$, where $\Pi_{\mathcal{X}}$ is Euclidean projection. Later words are $-G\Pi_{c_j}(Y)$ with radial projection and $c_j = r_0\, 2^{-j}$, $j = 1, 2, \dots$. Intervening readings do not modify a committed word. All inputs after the last word are zero, and the report is the retained prefix fit. Choose any finite certified bound

$$U_{\mathrm{pref}} \ge \sup_{\widehat{\theta} \in \Theta,\ z \in \mathcal{X}} \|\mathcal{C}_h(\widehat{\theta})^{\dagger} \widehat{A}^h z\|.$$

Boundedness of $\mathcal{X}$ and the controllability margin make this supremum finite. Every observation record obeys

$$\sum_{t<T} \|u_t\|^2 \le E_0 := E_{\mathrm{pref}} + U_{\mathrm{pref}}^2 + \Gamma_h^2 r_0^2 / 3. \tag{34}$$

Put $\ell = \max\{1, \bar{A}\}$ and use the computable bounds

$$L_C = \left\{ \sum_{j=0}^{h-1} (\ell^j + j\bar{B}\ell^{j-1})^2 \right\}^{1/2}, \qquad L_G = h\ell^{h-1} + L_C\Gamma_h, \qquad \kappa = \bar{B} H_h \Gamma_h. \tag{35}$$

The $j = 0$ second summand is zero. Telescoping powers preserves the order of matrix factors and gives

$$A^h - \mathcal{C}_h G = (A^h - \widehat{A}^h) + (\widehat{\mathcal{C}}_h - \mathcal{C}_h)\widehat{\mathcal{C}}_h^{\dagger}\widehat{A}^h, \qquad \|A^h - \mathcal{C}_h G\| \le L_G \|\widehat{\theta} - \theta\|_F. \tag{36}$$

This uses $\widehat{\mathcal{C}}_h \widehat{\mathcal{C}}_h^{\dagger} = I_n$, not instantaneous invertibility of B. Also $\|\mathcal{C}_h G\| \le \kappa$, and the block process innovation has covariance at most $H_h^2 I$.

Taking $\mathcal{X} = \{z : \|z\| \le X\}$, $r_0 = 1$ and $U_{\mathrm{pref}} = \Gamma_h X$ recovers the origin-ball energy formula. For the first word, $\bar{x}_L \in \mathcal{X}$. Projection nonexpansivity gives

$$\|x_L - \Pi_{\mathcal{X}} Y_L\| \le \mathrm{dist}(x_L, \mathcal{X}) + \|Y_L - x_L\| \le \|x_L - \bar{x}_L\| + \|Y_L - x_L\|.$$

Hölder, Eq. (32), and fresh block innovations bound its residual by a constant times $\sigma + \tau$. For later words write $M_{j,p} = \sup_{\theta} \|x_{L+jh}\|_{L^p}$. The exact block response and the global inequality $\|y - \Pi_c y\| = (\|y\| - c)_+ \le \|y\|^2/(4c)$ give

$$\begin{aligned} M_{j+1,p} \le{}& L_G F_{2p} (\sigma + \tau) M_{j,2p} + \kappa\sigma z_p (d + M_{j,p}^{\beta}) + \tau H_h z_p \\ &+ \frac{\kappa}{4c_j} [M_{j,2p} + \sigma z_{2p}(d + M_{j,2p}^{\beta})]^2. \end{aligned} \tag{37}$$

Only the new readout and physical innovations require independence; the fit and current state may be correlated. All moments exist because the finite-horizon bounds on matrix powers and hard input budget bound the controlled state by a constant plus a finite Gaussian convolution.

For R_j in Eq. (24), $j \ge 1$, uniformly over the stated noise ranges,

$$(\sigma + \tau) R_j \le R_{j+1}, \qquad R_j^2 \le 3R_{j+1}, \qquad \sigma R_j^{\beta} \le 3R_{j+1}, \qquad R_{j+1} \le R_j. \tag{38}$$

For the third inequality split $x = \sigma d$ or τ at $x = \sigma^{e_j}$ to obtain $\sigma x^{\beta} \le \sigma^{e_{j+1}} + x$. The endpoints follow directly. The square bound uses $2e_j \ge e_{j+1}$. An explicit induction is

$$\begin{aligned} B_{1,p} &= \max\{1, L_G F_{2p} J_{2p} + \kappa(H_L z_p + S_p) + H_h z_p\}, \\ B_{j+1,p} &= L_G F_{2p} B_{j,2p} + \kappa z_p (1 + 3B_{j,p}) + H_h z_p \\ &\quad + \frac{3\kappa}{4c_j} [B_{j,2p} + z_{2p}(1 + 3B_{j,2p})]^2. \end{aligned} \tag{39}$$

Equations (37)–(38) give $M_{j,p} \le B_{j,p} R_j$. For $s = T - L - Kh$, coasting adds every remaining physical innovation, so one may take

$$C_{\Theta,T,K,p} = P_s B_{K,p} + H_s z_p. \tag{40}$$

If a uniform exponential bound on matrix powers is available, both coasting factors can instead be bounded independently of its length.

A.4 A NEXT-PACKET COMPARISON FOR THE COMPLETE PRIVATE RECORD

Common histories and synthetic continuations. Fix the finite declared budget E, including energy already spent. Project proposed inputs onto the remaining-energy ball at every numerical history. This preserves a feasible policy's actual law and ensures feasibility also on every counterfactual history used below. Restrict to Eq. (22) and use the same initial uniform gain prior for every original policy seed. For a lower bound, grant the complete primitive process vector $\xi = (\xi_0, \dots, \xi_{T-1})$ and all original random seeds to both nodes before action. These variables are gain-independent. Future readout innovations are never granted.

We also grant the complete accumulated private buffer at every packet. At a recursive epoch, a *complete common history* H specifies the clock, all earlier commands and their spent energy, the common messages and silences, the private buffer through the preceding packet, the fixed seeds and process vector, and any previously granted auxiliary centers. The initial empty epoch is included. All new private readings occur after this common buffer. The actual actuator and sensor continuation maps remain arbitrary measurable maps of their allowed histories.

For a fixed numerical H, consider a *synthetic* experiment with gain uniform on any interval $I = [c_0 - w, c_0 + w] \subset [1 - w_0, 1 + w_0]$. This initializes the auxiliary continuation; the posterior of the original experiment may have a different form. The physical state corresponding to a comparison gain is determined by the recorded commands and that gain. Successive silence fixes the entire no-next-packet command word. Its reference states have the form

$$\begin{aligned} P_t(b) &= v_t + bQ_t, \qquad v_t = A_*^t x_0 + \tau \sum_{s<t} A_*^{t-1-s} \xi_s, \\ \|Q_t\| &\le M_E := \sqrt{E\, \lambda_{\max}\left(\sum_{j=0}^{T-1} A_*^j B_* B_*^{\mathsf{T}} (A_*^j)^{\mathsf{T}} \right)}. \end{aligned} \tag{41}$$

The Gramian bound uses the whole recorded and future command word and the actual declared budget. It is uniform in H and ξ. No restriction is placed on any intercept v_t.

Let $N = T - 1$, $D = nN$, $\overline{M} = \max\{1, M_E\}$, and define

$$\eta = \frac{1}{4D}, \qquad \zeta = \frac{1}{2048 D^2 \overline{M}}, \qquad \ell(w) = \zeta \min\{w, \sigma(d + w^\beta)\}. \tag{42}$$

All constants here are independent of the fixed process vector.

Lemma 5 (Uniform next-packet comparison)**.** *In either sensing model, fix a complete common history, remaining budget, seeds, primitive process vector and uniform interval as above. Let $\mathcal{K}_b^j(dH)$ be the actual numerical private-prefix kernel up to its next packet at j, including its stopping branch and payload. Use $j = \infty$ for no further packet, in which case H contains the remaining no-packet record. There exist a measurable center set $G_I \subset [c_0 - w/2, c_0 + w/2]$, with $|G_I| \ge w/2$, and subprobability kernels $\Lambda_c^j(dH)$ such that*

$$\frac{\mathbf{1}_I(b)}{2w}\, db\, \mathcal{K}_b^j(dH) \;\ge\; \frac{2/3}{2w} \int_{G_I} \frac{\mathbf{1}_{\{|b-c| \le \ell(w)\}}}{2\ell(w)}\, db\, \Lambda_c^j(dH)\, dc. \tag{43}$$

The measure on the right has total mass at least $1/8$ after summing all next-packet branches. In each component, conditional on its center, the gain is uniform on that child interval and independent of the numerical prefix. The inequality can first be extended by the policy's actual controlled continuation kernel; afterward the center may be revealed to enlarge the continuation class.

Proof. We give a simultaneous comparison for the entire remaining reference buffer, including every coordinate's covariance information.

Regular centers. For radial sensing and $Q_t \neq 0$, exclude the closest point of the affine line to the origin,

$$r_t = -\langle v_t, Q_t\rangle/\|Q_t\|^2, \qquad \|P_t(b)\|^2 = \|Q_t\|^2(b-r_t)^2 + \|v_t + r_t Q_t\|^2.$$

For coordinatewise sensing and $Q_{t,i} \neq 0$, exclude its scalar root $r_{t,i} = -v_{t,i}/Q_{t,i}$. There are at most D excluded locations in either case. Remove a radius ηw around each from the middle half of I; the remaining set has length at least $w - 2D\eta w = w/2$. At every retained center,

$$\begin{cases} \|P_t(c)\| \geq \|Q_t\|\eta w, & \text{radial,} \\ |P_{t,i}(c)| \geq |Q_{t,i}|\eta w, & \text{coordinatewise.} \end{cases} \tag{44}$$

For radial sensing a zero vector Q_t makes the entire reading law gain-independent. A single zero coordinate of Q_t does *not* make its radial variance gain-independent; that coordinate is retained in the comparison. For coordinatewise sensing, a zero $Q_{t,i}$ does make that coordinate's reading kernel common to all gains.

Common kernels remain in the experiment and may affect its trigger. They include a point mass at zero when $\beta > 0, d = 0$ and the corresponding constant state is zero. Their likelihood factors are one on their common support. Gaussian densities are used only on the positive-variance coordinates. One can use a product of Lebesgue measure plus δ_0 as a common dominating measure for these mixed coordinates. Nonconstant coordinates in the comparison have positive variance throughout every child, because $\ell(w) < \eta w$. At $\beta = 0$ all scales are constant.

Coordinate likelihood bounds. Write $q_t = \|Q_t\|$ for a nonconstant radial block, or $q_{t,i} = |Q_{t,i}|$ for a nonconstant scalar block. In either case, Eq. (44) lower-bounds its state magnitude by $q\eta w$. For $|b - c| \leq \ell = \ell(w)$, the logarithmic derivative of its scale is bounded by q divided by that magnitude. Thus each log scale ratio $a_i = \log(s_i(b)/s_i(c))$ satisfies

$$|a_i| \leq \frac{\ell}{\eta w - \ell} \leq \frac{1}{512D\overline{M} - 1} \leq \frac{1}{256D} =: \varepsilon_D. \tag{45}$$

For $\beta = 0$ it is zero. The standardized mean shift in each coordinate also obeys the uniform bound

$$\begin{aligned} \frac{q(d + w^\beta)}{\sqrt{d^2 + (q\eta w)^{2\beta}}} &\leq q + q^{1-\beta}\eta^{-\beta} \leq \frac{2\overline{M}}{\eta}, \\ |\Delta_i| := \left|\frac{P_i(b) - P_i(c)}{\sigma s_i(c)}\right| &\leq \frac{2\overline{M}\zeta}{\eta} = \frac{1}{256D}. \end{aligned} \tag{46}$$

For a radial block, a coordinate mean shift is at most its vector mean shift, so the same display applies. The bound remains valid as a nonzero slope tends to zero. Zero-slope coordinates use the common-kernel branch, which avoids evaluating 0^0. If $d = 0$, the first summand in the scale bound is simply zero.

One event for every prefix. Under the center law, standardize all nondegenerate coordinates of the remaining no-packet buffer. There are at most D independent standard normal variables. The event

$$\mathcal{A}_c = \left\{\sum_i z_i^2 \leq 4D\right\}$$

has probability at least $3/4$ by Markov's inequality. Common point-mass coordinates need no innovation in this event. For any temporal prefix, on its common support the exact density ratio is

$$L_{b,c} = \prod_i e^{-a_i} \exp\left\{\frac{1}{2}\sum_i z_i^2 - \frac{1}{2}\sum_i e^{-2a_i}(z_i - \Delta_i)^2\right\}.$$

The determinant factors are retained even for coordinates whose mean does not change. On $\mathcal{A}_c$, every prefix satisfies $\|z\| \leq 2\sqrt{D}$ and $\|\Delta\| \leq 1/(256\sqrt{D})$. Expansion and Cauchy–Schwarz give

$$\begin{aligned} \log L_{b,c} &\geq -D\varepsilon_D - \tfrac{1}{2}(e^{2\varepsilon_D} - 1)\|z\|^2 - e^{2\varepsilon_D}\|z\|\|\Delta\| - \tfrac{1}{2}e^{2\varepsilon_D}\|\Delta\|^2 \\ &\geq -\left\{\frac{1}{256} + \frac{3}{127} + \frac{1}{130048D}\right\} > -.028 > \log(2/3). \end{aligned} \tag{47}$$

For the rational bound use $e^{1/(128D)} \le 128/127$. It holds simultaneously for every prefix and every point of the child, not merely for a pair of expected divergences.

Stopping and marginalization. Apply the actual stopping and payload rule to the complete reference stream, with the fixed common history. On a branch $J = j < \infty$, this stream is an actual execution only through that packet. Define Λ_c^j by restricting the center reference law to $\mathcal{A}_c \cap \{J = j\}$ and marginalizing its unused future coordinates. The latter marginalization multiplies a prefix measure by a conditional probability at most one. Whenever that probability is positive, the observed prefix already satisfies the squared-norm bound used in Eq. (47). Hence the *prefix* likelihood ratio gives $\mathcal{K}_b^j \ge (2/3)\Lambda_c^j$ for every child gain. At a common numerical stream, the same trigger selects the same branch under both comparison gains. This retains all information in silence.

Integrate over centers. At fixed b the weights $\mathbf{1}_{\{|b-c|\le \ell\}}/(2\ell)$ integrate to at most one; this proves Eq. (43). Summing its branches gives mass at least

$$\frac{2/3}{2w}\int_{G_I} \mathbb{P}_c(\mathcal{A}_c)\, dc \ge \frac{2}{3}\frac{3}{4}\frac{1}{4} = \frac{1}{8} =: m_*. \tag{48}$$

Its displayed product form gives the conditional independence asserted in the lemma.

The physical continuation. First append the original policy's actual future kernel given b, H, the fixed process vector, seeds, and recorded common history. This kernel retains all remaining energy constraints and can depend on b through the future states and sensing variances. It does not depend on the new latent center, so it preserves measure domination. Only afterward reveal that center to both nodes and apply a lower bound uniform over the resulting enlarged continuation class. Thus the comparison continues with the actual post-correction measurements; the unused reference future has already been marginalized. This proves the final assertion. □

A.5 Terminal contrast and the packet recursion

Fix the granted process vector and put $a_\xi = \|v_T(\xi)\|$. At a common history with no packets left, the remaining command word is fixed. If $a_\xi = 0$ the next bound is zero. Otherwise projection onto v_T/a_ξ, for gain uniform on $[c-w, c+w] \subset I_0$, gives

$$\mathbb{E}_b\|v_T + bQ_T\| \ge \inf_{q\in\mathbb{R}} \mathbb{E}_b|a_\xi + bq| = \frac{a_\xi w}{\sqrt{c^2+w^2}+c} \ge \frac{a_\xi w}{2b_{\max}}, \qquad b_{\max} = 1 + w_0. \tag{49}$$

To check the equality, nonnegative q cannot improve on zero. For $q < 0$ put $r = -a_\xi/q$. Within $[c-w, c+w]$ the normalized loss is $\{(r-c)^2 + w^2\}/(2wr)$, minimized at $r = \sqrt{c^2+w^2}$; outside the interval it improves toward an endpoint. The positive gain interval gives $\sqrt{c^2+w^2} + c \le 2b_{\max}$.

Iterate Lemma 5 with $w_{j+1} = \ell(w_j)$. Each finite next-packet branch reduces the remaining packet budget by one, but may advance by any number of physical clocks. A never-packet branch, or a branch with no actionable clocks remaining, uses Eq. (49) immediately. Its width is larger and it needs no additional mixture factor. Backward induction therefore gives, for every fixed process vector and original seed, under the *same* initial uniform gain prior,

$$\mathbb{E}_{b,\text{readouts}}\|x_T\| \ge \frac{\|v_T(\xi)\|}{2b_{\max}}\, 8^{-K} w_K. \tag{50}$$

This includes policies using fewer than K packets.

The width recursion satisfies

$$w_j \ge w_0 \zeta^j \sigma^{e_j}, \qquad w_j \ge w_0 \zeta^j \sigma d. \tag{51}$$

For the power term, the cap branch uses $e_{j+1} \ge e_j$; the other branch uses $w_0^\beta \ge w_0$, $\zeta^{j\beta} \ge \zeta^j$ and $e_{j+1} = 1 + \beta e_j$. The noncap floor branch is at least $\zeta\sigma d$, and the cap branch propagates its previous bound. Initially $\sigma d \le 1$. Thus both endpoints of β are included without a β-dependent small-noise threshold.

The coefficients in the comparison and its width recursion are independent of ξ. Average the *whole original* process law and then the original seeds. Unconditional Jensen gives

$$\mathbb{E}_\xi\|v_T(\xi)\| \ge \|\mathbb{E}_\xi v_T(\xi)\| = \|A_*^T x_0\| =: a_T > 0. \tag{52}$$

This averages the full process law and data history. Worst-plant mean norm is at least the common-prior mean. Consequently the readout lower coefficient is

$$A_K = \frac{a_T w_0}{2b_{\max}}(\zeta/8)^K, \qquad D_{K,p} \geq A_K \max\{\sigma^{e_K}, \sigma d\}. \tag{53}$$

Here we use $L^p \geq L^1$, for $p \geq 1$.

Separately, in the *original causal experiment*, the final physical innovation is independent of the last chosen input. Conditional absolute Gaussian loss in any fixed coordinate is minimized at zero translation, so $\mathbb{E}\|x_T\| \geq \sqrt{2/\pi}\,\tau$. This independence statement is not made after granting future process innovations. Combining the two lower arguments proves Eq. (24) with

$$c_{\Theta,T,K,E} = \tfrac{1}{3}\min\{A_K, \sqrt{2/\pi}\}. \tag{54}$$

The budget enters M_E, hence ζ and this constant. The comparison mass $1/8$ is independent of n, T, but the child widths and the deterministic terminal contrast need not be. Neither a dimension-optimal constant nor a useful finite separation for every geometry is asserted. The same converse works for $K = T - 1$; only the construction is restricted by L and h. If an available constructive depth satisfies $\sigma^{e_K} \leq \sigma d + \tau$, its value and full feedback both have order $\sigma d + \tau$ at the fixed geometry, horizon and budget.

Reading saturation as a packet budget. Suppose the floor scales as $f(\sigma) \asymp \sigma^a$ with $a > 1$. The first depth whose exponent reaches that floor is

$$K_{\text{floor}} = \begin{cases} \lceil a \rceil, & \beta = 1, \\ \left\lceil \dfrac{\log(1 - a(1-\beta))}{\log \beta} \right\rceil, & 0 < \beta < 1, \quad a < (1-\beta)^{-1}. \end{cases}$$

This conclusion requires $L + hK_{\text{floor}} \leq T$. For $\beta = .5$ and $a = 1.8$, four rounds suffice in order; for $\beta = 1$, two suffice. If $0 < \beta < 1$ and $a \geq (1-\beta)^{-1}$, no fixed finite depth reaches that exponent. When $d > 0$ is fixed and $\tau = 0$, the floor instead has order σ, so additional packets can improve constants but not the noise exponent. These are consequences of the rate law; the controller itself uses no floor estimate.

B Learning a minimum-energy command beyond full calibration

B.1 Target, protocol, and statement

Let $\Theta \subset \mathbb{R}^{n\times n} \times \mathbb{R}^{n\times m}$ be a known compact class of controllable pairs, with preparation $x_0 \neq 0$. We use the physical experiment in Eq. (21): independent Gaussian process innovations, private readings at every clock $1, \ldots, T-1$, either radial or coordinatewise sensing, at most K causal packets, and a hard input-energy cap E on every record. The ranges are $\beta, d \in [0, 1]$, $0 < \sigma \leq .1$ and $0 \leq \tau \leq .1$. All packet times, payloads and silence may depend on the private history. The final estimator has the complete private buffer, the command history and the policy seeds. Its report is offline; it is not an additional command transmitted to the actuator.

For a fixed planning horizon H define the chronological reachability matrix and the target command

$$\mathcal{C}_H(A, B) = [A^{H-1}B, \ldots, B], \qquad U_H^*(A, B) = -\mathcal{C}_H(A, B)^\dagger A^H x_0. \tag{55}$$

When $\mathcal{C}_H$ has full row rank, this is the unique minimum-energy input word taking the original preparation x_0 exactly to zero in the noiseless plant. It is a deterministic function of the plant, estimated from the noisy experiment; it need not be the word executed in that experiment. No reset is part of the observation protocol.

For $1 \leq p < \infty$, let $Q_{K,p,H}$ be the infimum over all legal policies and offline estimators of $\sup_{\theta\in\Theta}(\mathbb{E}_\theta\|\widehat{U} - U_H^*(\theta)\|^p)^{1/p}$. At $H = T$, write $Q_{K,p} = Q_{K,p,T}$. Define $D_{K,p}$ analogously using $\|x_T\|$, and $\mathcal{R}_{K,p}^\theta$ using the Frobenius error of a full-model report with K packets; the main-text F_p

instead allows immediate feedback. These separate infima impose the same physical information and energy constraints; the construction below attains all stated upper bounds simultaneously.

Choose a deterministic prefix of length L with a positive uniform noiseless regressor margin, as in Eq. (23). Write E_{pref} for its energy and choose R containing all noiseless prefix endpoints. Choose h such that

$$F = A^h, \quad D = \mathcal{C}_h(A, B) \in \mathbb{R}^{n \times d_h}, \quad d_h = mh, \qquad \min_{\Theta} s_{\min}(D) \geq \kappa > 0. \tag{56}$$

One can always take $h = n$ by compactness and controllability. A common prefix exists by the construction preceding Eq. (26). No stability, square-input, or invertible-A assumption is imposed.

Theorem 6 (Optimal command precision on a compact controllable class)**.** *Fix the class, preparation, prefix and margins above. Put $b_h = h(mh + n + 1)$ and choose an integer $K \geq 1$ with*

$$T \geq L + Kb_h + 1, \qquad s = L + (K-1)b_h + h, \qquad H \geq \max\{s, h\}. \tag{57}$$

Any spare clocks are zero inputs after the stored quiet reading at s and before the final private batch. The resource bound below gives a finite E_0 such that, for each fixed finite $E \geq E_0$, one policy independent of the noise parameters satisfies

$$\begin{aligned} &\sup_{\Theta} \|x_T\|_{L^p} \leq C_p R_K, \qquad \sup_{\Theta} \|\widehat{U} - U_H^*\|_{L^p} \leq C_p R_{K+1}, \\ &\sup_{\Theta} \|\widehat{\theta} - \theta\|_{L^p} \leq C_p(\sigma + \tau), \qquad R_j = \sigma^{e_j} + \sigma d + \tau, \quad e_0 = 0,\ e_{j+1} = 1 + \beta e_j. \end{aligned} \tag{58}$$

For the command converse, suppose the class contains a segment

$$\{(A_*, bB_*) : |b - 1| \leq w_0\} \subset \Theta, \qquad 0 < w_0 \leq 1/4, \qquad A_*^H x_0 \neq 0. \tag{59}$$

Then $Q_{K,p,H} \geq c_E R_{K+1}$ for arbitrary adaptive packet policies. If also $A_*^T x_0 \neq 0$, then $D_{K,p} \geq c_E R_K$. Nonempty Euclidean interior of Θ suffices for both contrasts. For radial sensing, $\beta = 1$, $n \geq 2$, and nonempty interior, one additionally has $\mathcal{R}^{\theta}_{K,p} \geq c_E(\sigma + \tau)$, even if every reading is actionable. This full-model statement also holds under coordinatewise relative sensing when x_0 has a zero coordinate. All constants concern fixed finite geometry, horizons, depth, energy and p, and are uniform in σ, d, τ, β in the stated ranges; dimension, computational cost and sharp constants are outside this fixed-resource asymptotic statement.

The command target remains meaningful when some plants have $A^H x_0 = 0$ and hence $U_H^* = 0$. A positive minimax lower bound uses the nondegenerate family in Eq. (59), not merely controllability. In particular, the upper does not require every plant's A to be invertible. The choice $H = T$ meets the planning-horizon condition, but it is a long-horizon task when the experiment is long.

B.2 A FINITE FITTED ATLAS AND BOUNDED CANONICAL COORDINATES

The projected prefix fit $(\bar{A}, \bar{B}) \in \Theta$ in Eq. (26) has all-record moments

$$\|\|\bar{A} - A\| + \|\bar{B} - B\|\|_{L^r} \leq C_r(\sigma + \tau), \qquad 1 \leq r < \infty. \tag{60}$$

This follows from Eq. (32), including its fallback branch. The bounded fixed prefix has all finite moments, uniformly in the stated noise ranges. We retain this physical fit throughout the experiment.

Take $U \geq \sup_{\Theta} \|D\|$ and define

$$\eta = \frac{\kappa^n}{\sqrt{\binom{q}{n}}\, U^{n-1}} > 0. \tag{61}$$

Cauchy–Binet gives a column minor of D with determinant magnitude at least $\kappa^n / \sqrt{\binom{q}{n}}$. Its other singular values are at most U, so its smallest is at least η. Choose a maximum-volume n-column minor of $\bar{D} = \mathcal{C}_h(\bar{A}, \bar{B})$, with fixed tie breaking. Its selector $J \in \mathbb{R}^{d_h \times n}$ consists of columns of I_q; write $\bar{C} = \bar{D}J$. Thus $J^{\mathsf{T}} J = I$ and $s_{\min}(\bar{C}) \geq \eta$ on every record. This choice is included in the first packet and is then frozen.

For each selector with nonempty set define

$$\begin{aligned} &\Theta_J = \{\theta \in \Theta : s_{\min}(D_\theta J) \geq \eta/2\}, \qquad C = DJ, \\ &\mathsf{M} = C^{-1} F C, \qquad \mathsf{N} = C^{-1} D, \qquad \Gamma_J = \{(\mathsf{M}, \mathsf{N}) : \theta \in \Theta_J\}. \end{aligned} \tag{62}$$

These are compact sets and $\mathsf{N}J = I$. Whenever J is selected, $\bar{\theta} \in \Theta_J$, so its projection set is nonempty. All refined pairs are projected jointly onto Γ_J, using a measurable nearest-Frobenius projection. The projection can be nonconvex; only existence and its factor-two error inequality are used.

Let N_B denote the last m columns of N. The physical dynamics in these coordinates are recovered without an inverse of A or M:

$$\begin{aligned} A_c &= C^{-1}AC = [\mathsf{M}\mathsf{N}_B,\ \mathsf{N}_{\text{first (h}-1\text{) blocks}}]J, \\ B_c &= C^{-1}B = \mathsf{N}_B, \qquad \mathcal{C}_h(A_c, B_c) = \mathsf{N}. \end{aligned} \tag{63}$$

Indeed $A_c\mathsf{N}$ is the displayed shifted block matrix, and $\mathsf{N}J = I$. Consequently $\mathsf{N}\mathsf{N}^{\mathsf{T}} \succeq I$. Every projected canonical pair therefore has a reachability Gramian at least I at every planning horizon $H \geq h$.

Only the analysis uses the event

$$\mathcal{G} = \{\|\bar{D} - D\| \leq \eta/2\}. \tag{64}$$

On this event the true plant belongs to the selected Θ_J. The event is an analysis device; the algorithm selects its chart from the fitted matrix and has no access to the unknown true D. For $a = \max\{1, \sup_\Theta \|A\|\}$ and $b_* = \sup_\Theta \|B\|$, a finite explicit Lipschitz constant is

$$L_D = \left[\sum_{j=0}^{h-1} (a^j + j b_* a^{j-1})^2\right]^{1/2}, \quad \mathbb{P}(\mathcal{G}^c) \leq C_r (2L_D/\eta)^r (\sigma + \tau)^r. \tag{65}$$

The second summand at $j = 0$ is zero. The bound follows by telescoping the powers in D and applying Eq. (60). It holds for every finite r under the full law of the prefix fit.

B.3 Physical cancellation queries and the energy bound

An h-step word $w = (u_t, \ldots, u_{t+h-1})$ has endpoint $Fx_t + Dw + \tau g_t$, where $g_t = \sum_{i=0}^{h-1} A^{h-1-i}\xi_{t+i}$. Every intermediate reading still occurs. The construction commits the whole word and uses its endpoint reading.

At the first packet send the physical fit, J, and the current reading. Put $\bar{F} = \bar{A}^h$. A correction of radius c is the word

$$w = -J\bar{C}^{-1}\bar{F}\Pi_c(Y), \tag{66}$$

with first radius R and later radii $2^{-(r-1)}$ at packet r. Initialize the actuator's pair (M_o, N_o) from the coarse physical fit. In batch r freeze this pair and take $a_r = 2^{-(r-1)}$. For each selected coordinate $j = 1, \ldots, n$, execute two successive words

$$w_1 = a_r J e_j, \qquad w_2 = -a_r J M_o e_j, \qquad x^+ = F^2 x + a_r C(\mathsf{M} - M_o)e_j + \tau g^{(2)}. \tag{67}$$

For each nonselected column ℓ of I_q, execute the single word

$$w = a_r(e_\ell - JN_o e_\ell), \qquad x^+ = Fx + a_r C(\mathsf{N} - N_o)e_\ell + \tau g^{(1)}. \tag{68}$$

The two summands in this last command are physically added within the same word. They are not two successive experiments. The selected columns of N_o are exactly I and need no query.

For their respective endpoint readings, the sensor forms

$$\begin{aligned} \widetilde{M}e_j &= M_o e_j + a_r^{-1}\bar{C}^{-1}(Y^+ - \bar{F}^2 Y), \\ \widetilde{N}e_\ell &= N_o e_\ell + a_r^{-1}\bar{C}^{-1}(Y^+ - \bar{F}Y). \end{aligned} \tag{69}$$

After all queries, jointly project the pair onto Γ_J. For $r < K$, the next packet sends it and the final endpoint reading; apply Eq. (66) and start the next batch. After the Kth correction, store Y_s and insert $T - L - Kb_h - 1$ zero inputs before starting the final batch. Its last endpoint is then $T - 1$. Retain the refined pair privately and set $u_{T-1} = 0$. No query uses a within-batch refined column.

Each batch uses $2hn + h(d_h - n) = h(d_h + n)$ physical clocks, explaining Eq. (57). Its exact energy is $a_r^2(d_h + \|M_o\|_F^2 + \|N_{o,\mathrm{ns}}\|_F^2)$, where "ns" denotes the nonselected columns. Orthogonality of e_ℓ to the range of J gives the last term. Define the finite chart bounds

$$\begin{aligned} G_* &= \max_J \sup_{\Theta_J} \|JC^{-1}F\|, \\ Q_* &= \max_J \sup_{\Gamma_J} \{d_h + \|\mathsf{M}\|_F^2 + \|\mathsf{N}_{\mathrm{ns}}\|_F^2\}, \\ E_0 &= E_{\mathrm{pref}} + G_*^2(R^2 + 1/3) + (4/3)Q_*. \end{aligned} \tag{70}$$

Every record satisfies this budget, since the squared later correction radii sum to at most $1/3$ and $\sum_r a_r^2 \le 4/3$. For explicit coarse bounds one may use $G_* \le 2a^h/\eta$ and $Q_* \le d_h + 4nU^2(a^{2h} + 1)/\eta^2$. The bound is uniform in the number of completed batches for the fixed prefix and geometry. These constants may be large.

B.4 Moment refinement with chart and clipping failures

All following canonical quantities under an indicator $\mathbf{1}_{\mathcal{G}}$ are defined only on $\mathcal{G}$; no inverse of the true selected chart is taken on its complement. Constants are uniform over the finitely many charts. Let $f = \sigma d + \tau$ and $R_r = \sigma^{e_r} + f$. The first correction has endpoint moments $O(R_1)$ by Eq. (60) and nonexpansivity of Π_R about the noiseless prefix endpoint. The coarse canonical pair has the same error order on $\mathcal{G}$.

Suppose every batch-r endpoint and the old canonical-pair error have indicator-weighted L^v norm at most $C_{r,v}R_r$ for every finite v. The exact residual identities in Eqs. (67)–(68) propagate this bound through the finite query list, even though intermediate physical states may be of order a_r. Writing $\epsilon_t = Y_t - x_t$, a representative updated-column error is

$$\begin{aligned} &(\bar{C}^{-1}C - I)(\text{true column} - \text{old column}) \\ &\quad + a_r^{-1}\bar{C}^{-1}(F^j - \bar{F}^j)x + a_r^{-1}\bar{C}^{-1}(\epsilon^+ - \bar{F}^j\epsilon + \tau g^{(j)}), \qquad j \in \{1, 2\}. \end{aligned} \tag{71}$$

Thus coarse decoder error multiplies a small endpoint or parameter error. The event $\mathcal{G}$ and the selector are determined by the prefix. At later clocks the current readout innovation remains fresh conditional on the past. For either sensing model its moment is bounded by $C_v\sigma(d + \|x\|_{L^{\beta v}}^{\beta})$, with indicator versions on $\mathcal{G}$; at $\beta = 0$ the state term is one. If $\beta v < 1$, this is simply a moment bounded by a larger available moment, not an application of a quasi-norm triangle inequality. Hölder controls the correlated product terms in Eq. (71). Joint projection changes the error by at most a factor two on $\mathcal{G}$. The new pair therefore has moments

$$C_{r,v}\{(\sigma + \tau)R_r + \sigma R_r^{\beta} + f\} \le C'_{r,v}R_{r+1}. \tag{72}$$

For the nonlinear term, split $f \le \sigma^{e_r}$ and $f > \sigma^{e_r}$; in the latter case $e_r(1 - \beta) \le 1$ implies $\sigma f^{\beta} \le f$. Also $R_r^2 \le CR_{r+1}$ because $2e_r \ge 1 + \beta e_r$. These bounds include both endpoints of β.

The next correction has exact endpoint

$$(F - C\bar{C}^{-1}\bar{F})x + C\bar{C}^{-1}\bar{F}(x - \Pi_c Y) + \tau g.$$

Its first term is controlled by the same Hölder argument. For the clipping tail use the pointwise inequality $\|Y - \Pi_c Y\| \le c^{-1}\|Y\|^2$. Its indicator moments are $O(R_r^2)$, hence $O(R_{r+1})$. This closes the induction using finitely many higher moments for each fixed depth and target moment. The optional zero-input interval before the final batch preserves $O(R_K)$ state moments: it has at most T steps and adds a bounded process convolution of order τ. The induction includes clipping and fitting tails. These moment bounds are uniform in $\beta \in [0, 1]$: Gaussian moments use the common state-norm envelope, $e_j \le j$, and raising any fixed constant to β remains bounded. After the last batch and the final zero command,

$$\|\mathbf{1}_{\mathcal{G}}x_s\|_{L^v} + \|\mathbf{1}_{\mathcal{G}}x_T\|_{L^v} \le C_v R_K, \quad \|\mathbf{1}_{\mathcal{G}}[(\widehat{\mathsf{M}}, \widehat{\mathsf{N}}) - (\mathsf{M}, \mathsf{N})]\|_{L^v} \le C_v R_{K+1}. \tag{73}$$

B.5 Estimating the command with nonnegative powers

Store the reading at the quiet clock s in Eq. (57). Set $\widehat{z}_s = \bar{C}^{-1}Y_s$. On $\mathcal{G}$ its error from $z_s = C^{-1}x_s$ is $(\bar{C}^{-1} - C^{-1})x_s + \bar{C}^{-1}\epsilon_s$, whose indicator moments are $O(R_{K+1})$ by the preceding induction and

Hölder. Use the final refined pair in Eq. (63) to obtain $(\widehat{A}_c, \widehat{B}_c)$. This polynomial map is Lipschitz on the compact chart images, so its errors have the same order.

For $H \geq s$ the exact preparation contrast is

$$\begin{aligned} c_H := C^{-1}A^H x_0 = A_c^{H-s} z_s - \sum_{t<s} A_c^{H-1-t} B_c u_t \\ - \tau \sum_{t<s} A_c^{H-1-t} C^{-1} \xi_t. \end{aligned} \tag{74}$$

Define $\widehat{c}_H$ by plugging estimates into the first two terms. Every exponent is nonnegative. Telescoping powers on bounded matrix sets, $\sum_{t<s} \|u_t\| \leq \sqrt{sE}$, and Hölder give $\|\mathbf{1}_{\mathcal{G}}(\widehat{c}_H - c_H)\|_{L^v} \leq C_v R_{K+1}$. The omitted process convolution contributes $O(\tau)$ to this bound; it is not assumed to vanish or to be independent of the fitted chart.

Form the raw command and its final bounded report as

$$\begin{aligned} &\widetilde{U} = -\mathcal{C}_H(\widehat{A}_c, \widehat{B}_c)^{\dagger} \widehat{c}_H, \\ &\widehat{U} = \Pi_{\mathcal{U}_H(\Theta)}(\widetilde{U}), \qquad \mathcal{U}_H(\Theta) = \{U_H^*(\theta) : \theta \in \Theta\}. \end{aligned} \tag{75}$$

The raw expression exists on every record: its canonical Gramian is at least I. On $\mathcal{G}$, matrix inversion and multiplication are Lipschitz on these compact sets. Multiplying a terminal constraint by C^{-1} preserves its solution set and its minimum-norm solution, so the target is the physical command in Eq. (55). Thus the report has indicator error $O(R_{K+1})$ on $\mathcal{G}$. The final projection costs at most a factor two there.

The set $\mathcal{U}_H(\Theta)$ is compact and bounded because $H \geq h$ and the physical reachability margin is uniform. Consequently the report error on $\mathcal{G}^c$ is bounded by a fixed constant. Its L^p contribution is at most $C\mathbb{P}(\mathcal{G}^c)^{1/p}$. For the terminal state, finite-horizon bounded-input dynamics give the pathwise bound $\|x_T\| \leq C(1 + \tau\|\xi\|)$, regardless of the selected true chart. Its bad-event contribution is at most $C_p\mathbb{P}(\mathcal{G}^c)^{1/(2p)}$ by Hölder. Use Eq. (65) with any $r \geq 2p(K+1)$ and note

$$(\sigma + \tau)^{K+1} \leq C_K(\sigma^{e_{K+1}} + \tau), \qquad e_{K+1} \leq K + 1.$$

Both complements are absorbed by Eq. (58). This proves the unconditional upper bounds, with the bad-chart contribution controlled by the complement estimates above. The full-model upper is the retained prefix fit. None of the estimator or actuator formulas uses the moment order or a noise parameter.

B.6 A FINAL PRIVATE DISCLOSURE AND THE COMMAND CONVERSE

On the gain segment in Eq. (59),

$$U_H^*(b) = U_{H,*}^*/b, \qquad U_{H,*}^* = -\mathcal{C}_H(A_*, B_*)^{\dagger} A_*^H x_0 \neq 0. \tag{76}$$

Write $I_0 = [1 - w_0, 1 + w_0]$. Grant the process innovations and policy seeds as in Lemma 5, and grant the complete private prefix at every actual packet. Hard-energy projection at every numerical history makes its counterfactual no-send comparisons feasible. Let $N = T - 1$ and use the same explicit constants

$$\begin{aligned} M_E &= \sqrt{E\lambda_{\max}\left(\sum_{j=0}^{T-1} A_*^j B_* B_*^{\mathsf{T}} (A_*^j)^{\mathsf{T}}\right)}, \\ \zeta &= [2048(nN)^2 \max\{1, M_E\}]^{-1}, \qquad \ell(w) = \zeta \min\{w, \sigma(d + w^{\beta})\}. \end{aligned} \tag{77}$$

The next-packet lemma gives a shared submeasure of mass at least $1/8$ with a uniform child gain interval of halfwidth $\ell(w)$, independent of its numerical prefix conditional on the auxiliary center. It retains every readout coordinate, silence, and the actual trigger. Append the policy's actual continuation kernel before revealing that center, exactly as in that lemma.

For estimation, the next disclosure is either the next actual packet or, when no such packet occurs, the final complete private record. The latter is precisely the lemma's $j = \infty$ branch, including all

remaining no-send readings. It receives one buffer comparison before the estimation base is applied. There is no extra physical observation. When no new reading remains, applying the base directly is stronger.

After a final disclosure, conditional on its component center and record, the comparison gain is uniform on $[c-w, c+w] \subset I_0$. For any report u, projecting onto the direction of $U^*_{H,*}$ and using the median of this monotone scalar transform gives

$$\begin{aligned}\mathbb{E}_b\|u - U^*_{H,*}/b\| &\geq \inf_{a\in\mathbb{R}} \mathbb{E}_b|a - \|U^*_{H,*}\|/b| \\ &= \mathbb{E}_b|\|U^*_{H,*}\|/c - \|U^*_{H,*}\|/b| \geq \frac{\|U^*_{H,*}\|w}{2b^2_{\max}}, \quad b_{\max} = 1 + w_0.\end{aligned} \tag{78}$$

This bound is uniform over the numerical history, all granted process innovations, and all revealed centers.

Backward induction counts at most K actual disclosures plus this one offline disclosure. At an actual branch the genuine continuation is retained and its remaining packet budget decreases by one. At a never-send branch only the final comparison and the displayed base are needed. Since $\ell(w) \leq w$ and the mixture mass is at most one, branches using fewer packets give a stronger bound. Hence, with $w_{j+1} = \ell(w_j)$ under the same initial uniform prior for every seed,

$$\mathbb{E}\|\widehat{U} - U^*_H(b)\| \geq \frac{\|U^*_{H,*}\|}{2b^2_{\max}} 8^{-(K+1)} w_{K+1}. \tag{79}$$

The bound can be averaged over each full process law. The width induction in Eq. (51) then gives

$$Q_{K,p,H} \geq \frac{\|U^*_{H,*}\|w_0}{2b^2_{\max}} (\zeta/8)^{K+1} \max\{\sigma^{e_{K+1}}, \sigma d\}. \tag{80}$$

Worst-plant loss is at least the common-prior loss, and $L^p \geq L^1$. This lower makes no chart or constructive-schedule restriction.

For the process term, separately give the original causal experiment its physical states causally. Compare gains 1 and $1+v$, with $v = \min\{w_0, \tau/(2\|B_*\|\sqrt{E})\}$ when $\tau > 0$. Conditional on common states, readout kernels agree; transition means differ by vB_*u_t. The chain rule and hard budget give $\mathrm{KL}(P_1\|P_{1+v}) \leq v^2\|B_*\|^2E/(2\tau^2) \leq 1/8$. Pinsker and two-point estimation therefore yield

$$Q_{K,p,H} \geq \frac{3\|U^*_{H,*}\|}{16 b_{\max}} \min\{w_0, \tau/(2\|B_*\|\sqrt{E})\}. \tag{81}$$

At $\tau = 0$ the term is zero. At fixed E and the declared range of τ, it is a positive constant times τ. Combining Eqs. (80)–(81) proves the command converse. The terminal-state converse is the unrestricted lower in Eq. (53) together with its final-innovation bound. An interior point with invertible A_* supplies both nonzero contrasts and a gain segment; such points exist in any nonempty Euclidean interior.

For radial relative sensing in dimension at least two, the similarity-orbit comparison of Theorem 9, at fixed confidence, gives full-model L^p error $\Omega(\sigma)$ even under full feedback. The process comparison in Proposition 12 gives $\Omega(\tau)$ on an interior input segment. Together with the prefix fit these prove the last assertion of Theorem 6. If x_0 has a zero coordinate, the diagonal similarity comparison in the same reporting theorem gives the coordinatewise relative-sensing conclusion.

B.7 Fixed charts, full actuation, and the meaning of the task

If a single known selector has a uniform true margin, it can be fixed in advance; the atlas selection and bad-chart split are unnecessary. For $m = n$, invertible B and $h = 1$, take $J = I$, $C = B$, $\mathsf{M} = B^{-1}AB$ and $\mathsf{N} = I$. Only the paired queries remain:

$$u_t = ae_j, \quad u_{t+1} = -aM_oe_j, \qquad x_{t+2} = A^2x_t + aB(\mathsf{M} - M_o)e_j + \tau(A\xi_t + \xi_{t+1}).$$

The quiet endpoints provide the improved private estimate while the large intermediate state is not used in its innovation.

If A is also uniformly invertible in this special case, the same stored quiet state allows the additional report

$$\widehat{z}_0 = \widehat{\mathsf{M}}^{-s}\widehat{z}_s - \sum_{t<s} \widehat{\mathsf{M}}^{-(t+1)} u_t, \qquad z_0 = B^{-1}x_0.$$

The process convolution contributes $O(\tau)$, and inverse-power Lipschitz bounds give joint (M, z_0) error $O(R_{K+1})$. The gain-ray base with $z_{0,*}/b$ proves its matching joint lower. It does not prove a lower for M alone, which is constant on that ray. Invertibility is needed for this optional quotient report, not for the general command theorem.

The invariant command in Eq. (55) is a control-relevant statistical target: similarities fixing x_0 preserve its terminal constraint and unique minimum-energy solution. It does not encode every physical-coordinate quadratic loss or every tracking task. Moreover the finite lower constant contains $\|U^*_{H,*}\|$, which can be small for strongly stable plants at large H. The theorem describes fixed-resource precision orders and retains these geometry costs; it does not claim a useful finite separation for every class.

B.8 Fixed short planning horizons on nonsingular classes

Proposition 7 (An offline short-horizon command)**.** *Under Theorem 6, additionally suppose* $\min_{\theta\in\Theta} s_{\min}(A) > 0$. *Then its command upper bound holds for every fixed* $H \geq h$*, without requiring* $H \geq s$*, using the same learning trajectory, packets and energy. If Eq.* (59) *holds, the matching command lower bound also holds. Constants may depend on the inverse margin and both horizons.*

Proof. Every canonical $A_c = C^{-1}AC$, including the refined estimate obtained from the compact chart image, has a uniformly bounded inverse. Use the cached quiet reading and the final refined matrices to define

$$\widehat{z}_{c,0} = \widehat{A}_c^{-s}\widehat{z}_s - \sum_{t<s} \widehat{A}_c^{-(t+1)}\widehat{B}_c u_t, \qquad z_{c,0} = C^{-1}x_0. \tag{82}$$

The exact identity uses the true matrices and subtracts the process term

$$\tau \sum_{t<s} A_c^{-(t+1)} C^{-1}\xi_t.$$

Inverse powers are Lipschitz on the finitely many compact chart images. Thus the quiet-state decoding and refined-pair bounds in Eq. (73), Hölder's inequality, and $\sum_{t<s}\|u_t\| \leq \sqrt{sE}$ give $\|\mathbf{1}_{\mathcal{G}}(\widehat{z}_{c,0} - z_{c,0})\|_{L^p} \leq C_p R_{K+1}$, including the process convolution of order τ. Set $\widehat{c}_H = \widehat{A}_c^H \widehat{z}_{c,0}$ and use the raw command and final compact-set projection in Eq. (75). Its Gramian is uniformly positive for $H \geq h$; the same good-event Lipschitz argument and bad-chart moment bound prove the unconditional upper. Nonsingularity and $x_0 \neq 0$ ensure $A_*^H x_0 \neq 0$, so the unchanged gain-segment proof gives the lower for this H. □

B.9 Precision for a different preparation

Write $U_H^*(\theta; y) = -\mathcal{C}_H(\theta)^\dagger A^H y$. The learning trajectory still starts from x_0; y is the known preparation of a subsequent task and may be supplied to the learning policy in advance.

Proposition 8 (A preparation-specific precision advantage)**.** *Let* Θ *be a compact controllable class with nonempty interior,* $n \geq 2$*, and a uniformly identifying prefix supported by* T, E*. Under radial relative sensing, for every fixed* $H \geq h$ *and* $y \notin \operatorname{span}\{x_0\}$*,*

$$\inf_{\pi\in\Pi_{T-1}(E),\widehat{U}} \sup_{\theta\in\Theta} \|\widehat{U} - U_H^*(\theta; y)\|_{L^p} \asymp \sigma + \tau. \tag{83}$$

Here $1 \leq p < \infty$*, resources and* y *are fixed,* $0 \leq d \leq 1$*, and noise levels are sufficiently small. The constants can depend on the angle between* y *and* x_0*. Immediate access to every learning reading is allowed. The upper bound also holds with no actionable packets, since the reporter receives the fixed prefix record.*

Proof. The prefix fit has all finite moments of order $\sigma + \tau$. Uniform reachability and compactness make $\theta \mapsto U_H^*(\theta; y)$ Lipschitz, yielding the upper bound.

For the reader lower bound, choose an interior pair (A_*, B_*) with A_* invertible, which exists because singular matrices have empty interior. Let q be the normalized projection of y onto $x_0^\perp$, put $P = qq^\top$ and $S_a = I + aP$. For sufficiently small $a \geq 0$, the transformed pair $\theta_a = (S_a A_* S_a^{-1}, S_a B_*)$ stays in the class and fixes the learning preparation. Multiplying the terminal constraint by S_a^{-1} preserves its solution set and minimum-norm solution; therefore, with $C_* = \mathcal{C}_H(A_*, B_*)$,

$$U_H^*(\theta_a; y) = -C_*^\dagger A_*^H S_a^{-1} y, \tag{84}$$

$$\|U_H^*(\theta_a; y) - U_H^*(\theta_0; y)\| = \frac{a}{1+a} \|C_*^\dagger A_*^H P y\|. \tag{85}$$

The norm on the right is positive: $Py \neq 0$, A_* is invertible, and $C_*^\dagger$ is injective on $\mathbb{R}^n$. Lemma 10 applies without changing the physical reader and bounds the full adaptive transcript KL by $a^2\{(T-1)/(2\sigma^2) + n(T-1) + T\mathbf{1}_{\{\tau>0\}}\}$. Taking a a sufficiently small constant times σ and applying binary testing gives an $\Omega(\sigma)$ mean command error, uniformly over all policies. This also bounds every L^p risk.

For the process lower bound, use the interior input-gain segment (A_*, bB_*). Even after granting the true physical states causally, the transcript KL between $b = 1$ and $b = 1 + a$ is at most $a^2 \|B_*\|_{\mathrm{op}}^2 E/(2\tau^2)$. Conditional readout laws are common given the states. The target is $U_H^*(\theta_0; y)/b$ and is nonzero. Taking a proportional to τ proves an $\Omega(\tau)$ error. The two lower bounds imply $\Omega(\sigma+\tau)$ at fixed resources. □

For $y = cx_0$, in contrast, $U_H^*(\theta; y) = cU_H^*(\theta; x_0)$, so any command precision bound for the original preparation transfers by scaling. The distinction is the direction of the preparation, not whether its numerical value was known before learning.

C A SINGULAR UNDERACTUATED EXAMPLE WITH A NONSELECTED QUERY

The general construction includes both paired queries of M and single-word queries of nonselected columns of N. The following exact example makes both operations explicit. Take $n = 3$, $m = 2$, $h = 2$, $x_0 = e_1$, and

$$A = \begin{pmatrix} 7/10 & 1/5 & 0 \\ 0 & 3/5 & 1 \\ 0 & 0 & 0 \end{pmatrix}, \qquad B = \begin{pmatrix} 0 & 0 \\ 1 & 0 \\ 0 & 1 \end{pmatrix}.$$

Here $\operatorname{rank} A = 2$, $\operatorname{rank} B = 2 < 3$, and $A^H x_0 = (7/10)^H x_0 \neq 0$ for every finite integer $H \geq 0$. Thus the drift is singular, the system is underactuated, and its uncontrolled task endpoint remains nonzero.

The chart and recovered dynamics. With the chronological word ordering (u_t, u_{t+1}), set $D = [AB, B]$ and choose the selector $J = [e_1^{(4)}, e_3^{(4)}, e_4^{(4)}]$. Direct multiplication gives

$$D = \begin{pmatrix} 1/5 & 0 & 0 & 0 \\ 3/5 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \quad C = DJ = \begin{pmatrix} 1/5 & 0 & 0 \\ 3/5 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad \mathsf{N} = C^{-1}D = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}.$$

The determinant $\det C = 1/5$ proves $\operatorname{rank} D = 3$. This selector is a maximum-volume minor, tied with columns $(1, 2, 4)$. Moreover, $\mathsf{N}J = I_3$ and $\mathsf{N}\mathsf{N}^\mathsf{T} = \operatorname{diag}(1, 2, 1)$. The exact canonical matrices are

$$\mathsf{M} = C^{-1}A^2C = \begin{pmatrix} 127/100 & 13/10 & 1 \\ -273/500 & -21/50 & 0 \\ 0 & 0 & 0 \end{pmatrix}.$$

$$A_c = C^{-1}AC = \begin{pmatrix} 13/10 & 1 & 0 \\ -21/50 & 0 & 1 \\ 0 & 0 & 0 \end{pmatrix}, \qquad B_c = C^{-1}B = \begin{pmatrix} 0 & 0 \\ 1 & 0 \\ 0 & 1 \end{pmatrix}.$$

Write $\mathsf{N} = [N_1, N_2]$ in two-column blocks. The last block is $N_2 = \mathsf{N}_B = B_c$, so the ordering in Eq. (63) yields

$$A_c = [\mathsf{M}N_2, N_1]J, \qquad [A_cB_c, B_c] = \mathsf{N}, \qquad A_c^2 = \mathsf{M}.$$

These identities recover the one-step dynamics using matrix products, including when A and M are singular.

A nonzero word whose input effect cancels. Column 2 is the single nonselected column. For a frozen estimate $\widehat{N}$, its query is $w = a(e_2^{(4)} - J\widehat{N}e_2^{(4)})$. At the exact representation, $\mathsf{N}e_2^{(4)} = e_2^{(3)}$ and hence

$$w = a(e_2^{(4)} - e_3^{(4)}), \qquad (u_t, u_{t+1}) = ((0, a), (-a, 0)), \qquad Dw = 0.$$

Starting at zero, its intermediate state is ae_3 and its endpoint is zero. From a general state, the endpoint is A^2x_t; with an inexact frozen representation it becomes $A^2x_t + aC(\mathsf{N} - \widehat{N})e_2^{(4)}$ in the noiseless system. Thus the nonselected query is an actual nonzero input word and reveals the corresponding representation error.

Each of the three selected-coordinate pairs uses two words of length two. The one nonselected query uses one word of length two. A batch therefore takes $3 \cdot 2 \cdot 2 + 1 \cdot 2 = 14$ transitions. An additional correction word takes two transitions, giving $16 = h(mh + n + 1)$ per complete round, before any prefix or padding. All displayed identities were checked with exact rational arithmetic; this example verifies the construction's algebra and resource count.

D A SEPARATE ACCURACY FRONTIER FOR THE FINAL REPORT

The recovery objective and the final parameter report impose different statistical requirements. This appendix proves necessary report conditions that hold at every feedback depth. The reader converse uses a similarity orbit that fixes the preparation; the recovery converse in Appendix A uses an input-gain segment. These are different subfamilies, and the similarity orbit need not be difficult for recovery.

D.1 REPORTING EXPERIMENT AND CLASS GEOMETRY

Write $\theta = [A\ B]$ and use Frobenius distance on this concatenated matrix. Consider

$$x_{t+1} = Ax_t + Bu_t + \tau\xi_t, \qquad Y_t = x_t + \sigma\sqrt{d^2 + \|x_t\|^2}\, Z_t, \tag{86}$$

with fixed known $x_0 \neq 0$, $0 < \sigma$, $d, \tau \geq 0$, and independent standard n-Gaussian process and reader innovations, independent also of policy seeds. There are T transitions and at most $N \geq 1$ vector readings. In the paper's interface $N = T - 1$: the private readings are at $1, \ldots, T-1$. Any additional terminal or repeated reading available to the reporter must be counted in N.

The final offline reporter may use the entire private buffer and the command history. The bounds below even permit every reading to reach the actuator immediately. They therefore cover every packet budget, all private-history-dependent transmissions, and the information conveyed by silence. Feedback is based on the specified noisy readings; the reader converse does not grant noiseless state feedback or additional observations between clocks.

For a class Θ and fixed experiment and budget, define the reporting quantile

$$\mathcal{R}_{\mathrm{id}}(\delta) = \inf_{\pi, \widehat{\theta}} \inf \left\{ \epsilon \geq 0 : \sup_{\theta \in \Theta} \mathbb{P}_\theta^\pi \{ \|\widehat{\theta} - \theta\|_F > \epsilon \} \leq \delta \right\}, \qquad 0 < \delta < \tfrac{1}{2}. \tag{87}$$

The infimum is over the declared causal policies, with $\sum_{t<T} \|u_t\|^2 \leq E$ on every execution, and measurable reports. Our reader lower bound holds for the larger policy class with immediate access to all N noisy readings, so it also holds when this infimum is restricted by any packet cap.

Assume $n \geq 2$ and choose a unit vector $v \perp x_0$. Set

$$P = vv^\top, \quad S_s = I + sP, \quad \theta_s = [S_sA_*S_s^{-1}\ \ S_sB_*] \in \Theta \quad (0 \leq s \leq \bar{s}), \qquad \bar{s} > 0. \tag{88}$$

This is an explicit class assumption. Nonempty interior around a controllable nominal pair supplies a sufficiently short segment; controllability or compactness alone does not. Similarity preserves the spectrum and controllability. No stability assumption is required for this finite-horizon comparison. The reader in Eq. (86) remains fixed across the family.

Since $Px_0 = 0$, both preparations are the same. Moreover,

$$A_s - A_* = s[P, A_*]S_s^{-1}, \qquad B_s - B_* = sPB_*,$$
$$\gamma_*^2 = \|[P, A_*]\|_F^2 + \|PB_*\|_F^2, \qquad c_{\mathrm{orb}} = \frac{\gamma_*}{1+\bar{s}}, \qquad \|\theta_s - \theta_0\|_F \geq \frac{\gamma_* s}{1+s} \geq c_{\mathrm{orb}} s. \tag{89}$$

If (A_*, B_*) is controllable, then $\gamma_* > 0$: otherwise $[P, A_*] = 0$ and $PB_* = 0$ would imply $PA_*^j B_* = 0$ for every j, contradicting the full row rank of the controllability matrix.

Theorem 9 (Necessary accuracy for the report)**.** *Suppose Eq.* (88) *holds. Define*

$$L_\delta = (1-2\delta)\log\frac{1-\delta}{\delta}, \qquad H_{\mathrm{rad}} = \frac{N}{2\sigma^2} + nN + T\mathbf{1}_{\{\tau>0\}}. \tag{90}$$

For radial sensing with $\beta = 1$, every report-feasible policy satisfies

$$\epsilon \;\geq\; \mathcal{R}_{\mathrm{id}}(\delta) \;\geq \frac{c_{\mathrm{orb}}}{2}\min\left\{\bar{s}, \sqrt{\frac{L_\delta}{H_{\mathrm{rad}}}}\right\}. \tag{91}$$

The bound is independent of the packet budget and the input-energy cap. For coordinatewise $\beta = 1$ sensing, it remains valid when $P = e_j e_j^\top$ and $x_{0,j} = 0$, with

$$H_{\mathrm{coord}} = \frac{N}{2\sigma^2} + N + T\mathbf{1}_{\{\tau>0\}} \tag{92}$$

in place of H_{rad}.

For $0 < \sigma \leq 1$, the radial consequence

$$\mathcal{R}_{\mathrm{id}}(\delta) \geq \frac{c_{\mathrm{orb}}}{2}\min\left\{\bar{s},\ \sigma\sqrt{\frac{L_\delta}{N(n+\frac{1}{2}) + T\mathbf{1}_{\{\tau>0\}}}}\right\} \tag{93}$$

makes the resource dependence explicit. When T/N is bounded and the segment is unsaturated, this is an $\sigma\sqrt{\log(1/\delta)/N}$ obstruction with fixed geometry and dimension constants for $\delta \leq 1/4$. Indeed, $L_\delta \geq \frac{1}{4}\log(1/\delta)$ on that interval, because $1-2\delta \geq 1/2$ and $1-\delta \geq \sqrt{\delta}$. The interval minimum is necessary on a bounded parameter class.

D.2 A COMPARISON ON THE AUGMENTED HIDDEN-STATE LAW

Lemma 10 (Uniform transcript comparison)**.** *For any common causal policy π, let P_s^π be its full visible transcript law at θ_s, including private readings, commands, packets, and silences. Under radial sensing,*

$$\mathrm{KL}(P_0^\pi \| P_s^\pi) \leq s^2 H_{\mathrm{rad}}. \tag{94}$$

Under the diagonal coordinatewise assumptions of Theorem 9, replace H_{rad} by H_{coord}.

Proof. Include the common independent policy seeds in both experiments. Under θ_s, introduce the auxiliary hidden coordinate $z_t = S_s^{-1}x_t$; under θ_0, put $z_t = x_t$. This is a change of variables in the augmented law, not an observation granted to the policy. In both experiments $z_0 = x_0$.

Conditional on the same numerical hidden-state history and the same numerical reading history, the policy issues the same commands. Packet times and contents, including every silence, are generated by the same functions of these readings and seeds. Thus both hidden transition kernels have conditional mean $A_* z_t + B_* u_t$. For $\tau > 0$ their respective covariances are $\tau^2 I_n$ and $\tau^2 S_s^{-1} S_s^{-\top}$. The conditional KL is

$$\frac{1}{2}\{\mathrm{tr}(S_s^\top S_s) - n - \log\det(S_s^\top S_s)\} = s + \frac{s^2}{2} - \log(1+s) \leq s^2, \tag{95}$$

using $\log(1+s) \geq s - s^2/2$ for $s \geq 0$. Only one eigenvalue is stretched, so no dimension factor is needed. At $\tau = 0$ the two pulled-back transitions are identical Dirac kernels and their KL contribution is zero.

At a radial reading, write $a^2 = d^2 + \|z\|^2$ and $b_s^2 = d^2 + \|S_s z\|^2$. If $a = 0$, then $d = 0, z = 0$, and both readings are the same point mass. Otherwise both covariances are positive and

$$1 \leq b_s^2/a^2 \leq (1+s)^2, \qquad \lambda = \log(b_s^2/a^2) \in [0, 2s].$$

The base-to-alternative conditional KL is

$$\frac{n}{2}(e^{-\lambda} - 1 + \lambda) + \frac{s^2\|Pz\|^2}{2\sigma^2 b_s^2} \leq ns^2 + \frac{s^2}{2\sigma^2}. \tag{96}$$

Here $e^{-\lambda} - 1 + \lambda \leq \lambda^2/2$ for $\lambda \geq 0$, and $b_s^2 \geq \|z\|^2 \geq \|Pz\|^2$. The inequality is pointwise and uses no state moment or excitation condition.

For coordinatewise sensing and $P = e_j e_j^\top$, only the jth reader coordinate changes at the same hidden state. Its scalar Gaussian calculation is Eq. (96) with $n = 1$. Every other coordinate is a common conditional kernel, including any zero-variance point mass. A non-diagonal similarity would not preserve this coordinatewise argument.

Apply the chain rule on the common augmented space, interleaving the T transitions and at most N readings in their actual order. Commands and communications add no KL because their conditional kernels are common. Adaptive acquisition and stopping can be padded with common dummy outcomes after stopping. The bounds above yield $s^2 H_{\mathrm{rad}}$ or $s^2 H_{\mathrm{coord}}$ for the augmented law. Its reading marginal is the actual experiment at each plant. Data processing onto the visible transcript proves the lemma. In particular, we have not assumed that the visible history itself has Gaussian conditional kernels. □

Proof of Theorem 9. If $\|\theta_s - \theta_0\|_F > 2\epsilon$, selecting whichever point is closer to the report gives a test whose error probability is at most δ under either plant. Binary data processing implies

$$\mathrm{KL}(P_0^\pi \| P_s^\pi) \geq (1-\delta)\log\frac{1-\delta}{\delta} + \delta\log\frac{\delta}{1-\delta} = L_\delta. \tag{97}$$

Consequently $\|\theta_s - \theta_0\|_F \leq 2\epsilon$ whenever $0 < s < \min\{\bar{s}, \sqrt{L_\delta/H_{\mathrm{rad}}}\}$. Use Eq. (89) and take s to that endpoint. The argument applies to every admissible policy and report, proving Eq. (91); the coordinatewise case is identical. □

D.3 An exact scale experiment inside coordinatewise sensing

For the coordinatewise subfamily there is a sharper interpretation when $d = \tau = 0$. It explains why a small state need not provide an accurate full-parameter report.

Lemma 11 (Reduction to independent scale observations)**.** *Assume Eq.* (88) *with $P = e_j e_j^\top$, $x_{0,j} = 0$, coordinatewise $\beta = 1$ sensing, and $d = \tau = 0$. Grant the nominal pair (A_*, B_*), common policy seeds, and numerical command history to the reporter. For every causal policy, its complete transcript is a parameter-independent garbling of N independent observations*

$$V_\ell = c(1 + \sigma Z_\ell), \qquad c = 1 + s, \qquad Z_\ell \overset{\mathrm{iid}}{\sim} N(0,1). \tag{98}$$

If exactly N readings occur and each corresponding base state has nonzero jth coordinate almost surely, the two experiments are Blackwell equivalent, up to parameter-independent randomness.

Proof. From the command history compute $z_{t+1} = A_* z_t + B_* u_t$, $z_0 = x_0$. The physical state is $x_t = S_s z_t$. Whenever $z_{t,j} \neq 0$, the normalized reading satisfies

$$Y_{t,j}/z_{t,j} = c(1 + \sigma Z_t'). \tag{99}$$

The state and its sign are predictable before this reading; freshness and symmetry of its Gaussian innovation imply that Z_t' is standard normal and independent of the preceding numerical history.

Conversely, a simulator supplied with Eq. (98) and independent noises for all other coordinates runs the original policy. At each reading it sets $Y_{t,j} = z_{t,j} V_\ell$ and generates the other coordinates with

their common reader kernel. It consumes a V_ℓ even when $z_{t,j} = 0$, in which case this reading is zero. It then continues the policy, including its packets, silences, and subsequent commands. All operations of this simulator are independent of c. This proves domination by the independent scale experiment. Under the additional nonzero-state condition, every V_ℓ is recovered from the complete transcript and the granted controls by Eq. (99), proving the reverse garbling. No equivalence is claimed for a compressed final report alone. □

The scale parameter is identifiable: a single observation in Eq. (98) has Fisher information $(\sigma^{-2} + 2)/c^2 > 0$, and more readings permit consistent estimation. Making a nonzero state smaller does not improve this relative calibration precision; a zero state carries no scale information. At the same time, if a nominal open-loop word fits the energy cap and sends z_T to zero, it sends $x_T = S_s z_T$ to zero for every s. Thus this reporting obstruction does not imply a recovery obstruction.

D.4 A SEPARATE CONVERSE FROM PROCESS NOISE AND INPUT ENERGY

Proposition 12 (Energy-limited report accuracy)**.** *Suppose, independently of Eq.* (88)*, that the class contains*

$$\widetilde{\theta}_h = [A_* \;\; B_* + hF], \quad 0 \le h \le \bar{h}, \qquad \bar{h} > 0, \quad \|F\|_{\mathrm{op}} = \|F\|_F = 1. \tag{100}$$

For $E > 0, \tau > 0$, every report-feasible policy satisfies

$$\mathcal{R}_{\mathrm{id}}(\delta) \ge \frac{1}{2} \min\left\{ \bar{h}, \tau\sqrt{\frac{2L_\delta}{E}} \right\}. \tag{101}$$

This holds for either reader and any packet budget. It even holds when each actual state is supplied causally to the controller and the complete state trajectory is supplied to the final reporter.

Proof. In this more informative experiment, commands are functions of the same numerical causal state/reading history and common seeds. Their values therefore agree in the two conditional transition kernels. The means differ by hFu_t and the covariances are the same $\tau^2 I_n$. Reader kernels conditional on the actual state are common. The chain rule gives

$$\mathrm{KL}(\widetilde{P}_0^\pi \| \widetilde{P}_h^\pi) = \frac{h^2}{2\tau^2} \mathbb{E}_0^\pi \sum_{t<T} \|Fu_t\|^2 \le \frac{h^2 E}{2\tau^2}. \tag{102}$$

The hard cap here is the actual energy cap of the compared class. The parameter distance is h, so apply the testing argument in Eq. (97) and take a limit from below. The genie is causal for control and does not reveal future process innovations before an input is chosen. □

If $E = 0$, the trajectory never depends on B, giving $\mathcal{R}_{\mathrm{id}}(\delta) \ge \bar{h}/2$ directly for $\delta < 1/2$. At $\tau = 0$, Eq. (101) supplies only the trivial zero bound; a Gaussian transition formula is not applied to unequal Dirac means. When both geometric assumptions hold, take the maximum of Eqs. (91) and (101). The two necessary conditions need not use the same hard plant.

D.5 SIMULTANEOUS REPORTING AND RECOVERY AT FIXED RESOURCES

Corollary 13 (Two precision scales on one class)**.** *Assume the class, prefix, and resource conditions of Theorem 4, with $\beta = 1$, $n \ge 2$, nonempty interior of Θ, radial sensing, and a fixed finite $E \ge E_0$. There are constants $c, C, r_* > 0$, independent of the noise levels, δ, and ϵ, with the following property. For $0 < \delta \le 1/4$, $d \in [0, 1]$, and $\rho = (\sigma + \tau)\sqrt{\log(1/\delta)} \le r_*$:*

1. *If $\epsilon < c\rho$, the report requirement is infeasible even with immediate feedback after every private reading.*
2. *If $C\rho \le \epsilon \le 1$, the retained prefix fit meets the report requirement, and for each $1 \le K \le \lfloor (T - L)/h \rfloor$ the corresponding depth-K controller attains $D_{K,p} \asymp \sigma^K + \sigma d + \tau$ for every fixed $1 \le p < \infty$.*

For coordinatewise sensing the same conclusion holds if x_0 has at least one zero coordinate. The constants depend on the fixed geometry, horizon, energy cap, and construction choices.

Proof. Choose a controllable interior point of Θ. A sufficiently short similarity segment (88) and a sufficiently short rank-one input segment (100) lie in its interior ball. For radial sensing take any unit $v \perp x_0$; for coordinatewise sensing take the coordinate projection at a zero coordinate of x_0. Controllability gives $c_{\mathrm{orb}} > 0$ in both cases. Put $\ell_\delta = \log(1/\delta)$ and define

$$Q = N(n + \tfrac{1}{2}) + T, \qquad a = \frac{c_{\mathrm{orb}}}{4\sqrt{Q}}, \qquad b = \frac{1}{2\sqrt{2E}}, \qquad c = \tfrac{1}{2}\min\{a, b\},$$
$$B_L = \sqrt{1 + \frac{\log(2L)}{\log 4}}, \qquad C = C_{\mathrm{tail}} B_L, \tag{103}$$
$$r_* = \min\left\{ \tfrac{1}{10}\sqrt{\log 4},\ 2\bar{s}\sqrt{Q},\ \bar{h}\sqrt{2E},\ C^{-1} \right\}.$$

Here C_{tail} is the explicit constant in Eq. (31). The construction has $E_0 > 0$; in particular $E > 0$. The first term in r_* ensures $\sigma, \tau \le 1/10$. Since $L_\delta \ge \ell_\delta/4$, Eqs. (93) and (101) imply

$$\mathcal{R}_{\mathrm{id}}(\delta) \ge \frac{c_{\mathrm{orb}}}{2}\min\{\bar{s}, \sigma\sqrt{\ell_\delta}/(2\sqrt{Q})\}, \qquad \mathcal{R}_{\mathrm{id}}(\delta) \ge \frac{1}{2}\min\{\bar{h}, \tau\sqrt{\ell_\delta/(2E)}\}.$$

The second and third terms in r_* make the displayed linear branches no larger than their interval radii. Thus $\mathcal{R}_{\mathrm{id}}(\delta) \ge \max\{a\sigma, b\tau\}\sqrt{\ell_\delta} \ge c\rho$. When $\tau = 0$, the process bound is read as zero and the same conclusion follows from the reader bound alone. Using the larger radial value of Q is valid for coordinatewise sensing as well.

Conversely, $\log(2L/\delta) \le B_L^2 \ell_\delta$ for $\delta \le 1/4$. Consequently $C\rho \le \epsilon \le 1$ implies $C_{\mathrm{tail}}(\sigma + \tau)\sqrt{\log(2L/\delta)} \le \epsilon$. Equations (30)–(31) give the report guarantee for the unchanged prefix fit. Theorem 4 then supplies the risk bounds for every supported K, with $e_K = K$. The report and recovery guarantees hold on the same trajectory; no conditioning on report success is used. □

This corollary fixes the resource counts. It does not assert optimal joint dependence when dimension, measurement count, horizon, or energy are allowed to grow. The explicit interval minima in the preceding converses remain available outside its small-noise regime.

E Certificates and explicit instances of plant geometry

E.1 Certifying the class margins and its lower witness

The margins in Appendix A are minima over the declared class, not sampled minima. Here is one route to finite certificates. Cover Θ by cells with centers $\theta_c \in \Theta$ and radii $\|A - A_c\|, \|B - B_c\| \le \eta_c$. Let $S_L = \sum_{j=0}^{L-1} P_j$, using its finite bounds on matrix powers. For the fixed prefix, the exact difference convolution and a zero-filling right-column shift give

$$\|W_\theta - W_{\theta_c}\| \le \eta_c S_L(\|X_{c,-}\| + \|U\|). \tag{104}$$

Thus a positive lower bound on every cell is $s_{\min}(W_{\theta_c}) - \eta_c S_L(\|X_{c,-}\| + \|U\|)$. For the universal impulse prefix $\|U\| = 1$. Likewise $s_{\min}(\mathcal{C}_h(\theta_c)) - \eta_c L_C$, with L_C from Eq. (35), bounds the controllability singular value on that cell. Outward interval enclosures or other rigorous norm bounds certify these expressions. Their minimum gives valid q, s_h when positive. The required covering computation can be expensive; positivity by compactness alone is not a numerical certificate.

For the lower witness, suppose a Frobenius ball of radius $\rho > 0$ around (A_c, B_c) lies inside Θ. Among the $n + 1$ matrices

$$A_c + \delta_j I, \qquad \delta_j = \frac{j\rho}{4(n+1)\sqrt{n}}, \qquad j = 1, \dots, n+1, \tag{105}$$

at least one is invertible: $\det(A_c + zI)$ is a monic polynomial of degree n. Choose it as A_* and set $B_* = B_c$ and $w_0 = \min\{1/4, \rho/(4\max\{1, \|B_c\|_F\})\}$. The entire segment (A_*, bB_*), $|b - 1| \le w_0$, remains inside the ball, since both component displacements are at most $\rho/4$. Invertibility and $x_0 \ne 0$ ensure $A_*^t x_0 \ne 0$ at every finite clock. This establishes Eq. (22) for classes with nonempty interior, including disconnected classes. Full affine hull alone is insufficient for this construction. For a class with fixed nilpotent drift and $\tau = 0$, sufficiently long zero-input coasting can erase every earlier probe. The positive sensor-noise lower term established here uses the prepared gain contrast.

F ROBUSTNESS TO THE SENSOR SCALE PROFILE

Replace the reader in Eq. (7) by $Y = x+\sigma s(x)Z$, where the profile is fixed throughout an experiment. Assume uniform constants $0 < c_- \le c_+ < \infty$ and $L_s < \infty$ with

$$c_-\sqrt{d^2 + \|x\|^{2\beta}} \le s(x) \le c_+\sqrt{d^2 + \|x\|^{2\beta}}, \tag{106}$$

$$|\log s(x') - \log s(x)| \le L_s \frac{\|x' - x\|}{\|x\|} \quad \text{if } x \ne 0, \quad \|x' - x\| \le \|x\|/2. \tag{107}$$

The endpoint convention at $\beta = 0$ is the same as in the main model. When $\beta > 0, d = 0$, set $s(0) = 0$ and interpret that conditional reader as a point mass. With a positive floor, Eq. (106) requires $s(0) \ge c_- d$.

Corollary 14 (Uniform profile robustness)**.** *Under Eqs.* (106)–(107)*, Theorem 4 holds uniformly over the declared family of profiles, with constants also depending on c_-, c_+, L_s. The attaining controller needs neither the profile nor these three constants in its fit and action rule; its hard energy and packet count are unchanged. The sufficient reporting constant is adjusted below.*

For the upper bound, fresh conditional centering and $s(x) \le c_+(d + \|x\|^\beta)$ replace every readout moment coefficient by c_+ times its original value. The pathwise fit inequality, parameter projection and control words are unchanged. In the notation of Appendix A, the squared fit bound becomes

$$F_0^2 Ln\{H_L^2(1 + nc_+^2\sigma^2)\tau^2 + c_+^2(1 + X^2)\sigma^2\}.$$

Consequently the explicit sufficient condition uses

$$C'_{\mathrm{rep}} = \max\{nc_+^2, F_0^2 Ln \max(2H_L^2, c_+^2(1 + X^2))\}.$$

It implies $nc_+^2\sigma^2 \le 1$, controlling the mixed variance term. The same moment induction gives the identical exponent and both floors.

For the lower bound use the radial affine reference and closest-point exclusion in Lemma 5. Put $q_t = \|Q_t\|$, $D = n(T-1)$ and $\eta = 1/(4D)$. Each retained center has $\|P_t(c)\| \ge q_t\eta w$. Replace ζ in Eq. (42) by

$$\zeta' = \zeta \min\{1, c_-, 1/\max(1, L_s)\}.$$

The child displacement is below half the center's state norm, so Eq. (107) gives a log-scale ratio at most $L_s\zeta'/\eta \le 1/(512D)$. Its standardized mean shift is at most

$$\frac{\zeta'}{c_-}\frac{q_t(d + w^\beta)}{\sqrt{d^2 + (q_t\eta w)^{2\beta}}} \le \frac{2\max(1, M_E)\zeta'}{c_-\eta} \le \frac{1}{256D}.$$

The same whole-buffer event and prefix likelihood inequality therefore give comparison mass $1/8$. Zero-slope vector kernels are common and retained, including point masses. Stopping marginalization, actual continuation and center revelation proceed in the order of that lemma. The width recursion and unconditional process averaging give the same lower rates with ζ', without a favorable process event.

For a concrete family, let the singular values of an unknown invertible Q lie in $[\ell, u]$, where $0 < \ell \le u < \infty$, and set $s(x) = \sqrt{d^2 + \|Qx\|^{2\beta}}$. One can take $c_- = \min(1, \ell)$, $c_+ = \max(1, u)$ and $L_s = 2u/\ell$. The norm comparisons prove the first condition. On the segment from x to x' in the second condition, every point y has norm at least $\|x\|/2$, and $\|\nabla \log s(y)\| \le \beta\|Q\|/\|Qy\| \le 2u/(\ell\|x\|)$. Integrating gives the stated finite-difference bound. At $\beta = 0$ the profile is constant. The observation mean is still x and its conditional covariance still isotropic; this example does not require estimating Q or imply a result for an unknown observation map.

G DISCLOSURE ACCOUNTING, A COMPUTABLE SUBCLASS, AND THE TASK QUOTIENT

G.1 WHEN THE FINAL UPLOAD IS CHARGED

The resource in Theorem 1 counts packets that can change subsequent inputs. Its offline report uses an additional transfer of the actual private record. This transfer creates no observation, but it is an information disclosure. Here we count it explicitly.

Let $Q^{\mathrm{pub}}_{K,p,H}$ restrict the estimator to the complete public transcript: the at most K transmitted packets, their times, every silence, and the command history. Independent seeds may also be granted; private readings not contained in the transcript are unavailable. Let $Q^{\mathrm{up}}_{K,p,H}$ instead allow $K-1$ actionable packets and one mandatory final upload of the complete private buffer, charged as the Kth packet. The upload is after the learning inputs and permits no further correction. There are still only $T-1$ physical readings.

Corollary 15 (Public reports and charged final uploads)**.** *Under the compact-class, nonempty-interior, fixed-resource assumptions of Theorem 6, for $K \geq 1$ and enough time and energy for the construction below,*

$$Q^{\mathrm{pub}}_{K,p,H} \asymp R_K, \qquad Q^{\mathrm{up}}_{K,p,H} \asymp R_K. \tag{108}$$

One may take $H = T$. More generally, for $K \geq 2$ use the time and planning-horizon conditions of that theorem with $K-1$ in place of K; for $K = 1$, a common prefix and $H \geq h$ suffice. If A is uniformly invertible, any fixed $H \geq h$ is allowed, as in Appendix B.8. The constants concern fixed resources.

Proof. For the charged-upload upper bound with $K \geq 2$, run the $(K-1)$-actionable-packet construction of Theorem 6. Its final physical query batch is observed privately. The charged upload then delivers that buffer and hence the R_K command estimate. For the public-transcript upper, place this last batch so that it ends at clock $T-1$ and transmit the completed estimate or its sufficient record at that clock. The remaining input is zero. This uses K actual packets and the same physical queries and energy; the last packet need not cause another correction. When $K = 1$, transmit the common prefix fit; smoothness of the command functional on the uniformly controllable class gives $O(\sigma + \tau) = O(R_1)$.

For the public converse, use the gain segment (A_*, bB_*) of Eq. (59), with target U_*/b. At each next actual packet, grant its complete stopped private buffer rather than its possibly compressed payload. If fewer than K packets occur, complete the next buffer at the final reading clock; this additional genie reveals all residual silences and can only help the estimator. There are at most K such buffer disclosures: after K actual transmissions the cap forces all subsequent silences, and no private observation reaches the public transcript. Remaining positions can be padded by empty buffers. Apply the same stopped-buffer comparison of Appendix A at these at most K disclosures, appending the actual continuation before revealing each auxiliary center. Every actual public packet and silence is measurable from this enlarged transcript. The final gain interval has width of order $\sigma^{e_K} + \sigma d$, and its conditional median loss for U_*/b has the positive constant in the command converse. The separate causal transition comparison supplies $\Omega(\tau)$. This proves the lower bound against all public-transcript policies, not just fixed times. For the charged-upload interface, the same induction has $K-1$ actionable buffers and one final complete buffer, giving the identical bound. □

Thus K actionable packets plus an uncharged offline upload yield R_{K+1}, whereas K total report-bearing disclosures yield R_K. The final upload supplies the last physical query batch to the reporter; counting that upload among the disclosures gives the second rate.

G.2 POLYNOMIAL-TIME REFINEMENT UNDER FULL ACTUATION

For the fully actuated subclass below, spectral clipping replaces the realizable-set projection of the general compact-class result. Suppose $m = n$, $s_{\min}(B) \geq b_- > 0$, $\|B\| \leq b_+$ and $\|A\| \leq a_+$ on the class. A common identifying prefix has a specified positive regressor margin. Form the thresholded least-squares fit and project $\bar{A}$ onto the operator-norm ball of radius a_+. Clip the singular values of $\bar{B}$ to $[b_-, b_+]$, using a fixed admissible fallback on a low-margin record.

Set $M = B^{-1}AB$ and $m_+ = a_+b_+/b_-$. In each complete batch use Eq. (4), then clip the singular values of its update to $[0, m_+]$. The control-coordinate input matrix is exactly I. No independently estimated pair has to be made jointly realizable. If short-horizon reconstruction is desired and also $s_{\min}(A) \geq a_- > 0$, use the interval

$$[m_-, m_+], \qquad m_- = a_-b_-/b_+, \tag{109}$$

and the same spectral clipping for $\bar{A}$ to $[a_-, a_+]$. The final report is projected onto the Euclidean ball of radius $a_+^H\|x_0\|/b_-$. The true command lies in this ball, since using only the last input gives a feasible zeroing word with at most that norm. This replaces the task-image projection in the general proof by an explicit radial clipping operation.

Proposition 16 (Computable refinement)**.** *These spectral operations, clipped corrections, cached quiet reading and command reconstruction attain the state and command upper bounds of Theorem 6 on this subclass, with the same fixed-resource qualifications and the explicit sufficient energy allowance below. The singular-A case uses nonnegative powers; the short-H case uses the stated inverse margin. After the prefix fit, the arithmetic cost is $O((K + H)n^3 + Tn^2)$, in addition to storing or streaming the actual observations. No nonconvex search oracle or controllability-minor enumeration is used.*

For query amplitudes $a_j = 2^{-j}$, $j = 0, \ldots, K-1$, correction caps $(a_+/b_-)R$ at the first correction and $(a_+/b_-)2^{-j}$ for $j = 1, \ldots, K-1$ thereafter, a conservative all-path allowance is

$$E_{\mathrm{comp}} = E_{\mathrm{pref}} + (a_+/b_-)^2(R^2 + 1/3) + (4/3)n(1 + m_+^2).$$

Here R exceeds the noiseless prefix-state bound. As in the original clipping proof, fixed positive caps preserve the moment rates at each fixed depth; they need not equal the minimal feasible energy.

Proof. Singular-value clipping is a Frobenius-nearest matrix in the specified spectral interval: unitary alignment follows from the trace inequality, then each singular value solves its one-dimensional projection. Because the true matrix lies in this set, its distance from the projected estimate is at most twice its distance from the raw estimate. This argument does not require a convex lower-singular-value constraint. The operator-norm-ball projection is likewise computable by one SVD. Hence the coarse L^q bounds and every refinement bound survive these operations. The projected $\bar{B}$ has inverse norm at most b_-^{-1}; all projected M are bounded, and, when needed, their inverse norms are at most m_-^{-1}.

In control coordinates the dynamics are $z^+ = Mz + u$. Their H-step reachability matrix is $G_H = [M^{H-1}, \ldots, I]$, so $G_H G_H^{\mathsf{T}} \succeq I$ for every projected M, independently of realizability as a member of the original class. The physical-pair identity, moment recursion, quiet-state reconstruction and clipping-tail argument therefore apply without a joint projection. The final Euclidean ball projection cannot increase error to the true command and removes the task-image projection as well. These are the only places for which the original proof needed boundedness, invertibility of the decoder and a reachability margin. One SVD per coarse matrix and per batch costs $O(Kn^3)$; the reachability solve costs $O(Hn^3)$ and the query/history products cost $O(Tn^2)$. Prefix regression has its usual polynomial matrix-solve cost. The statement is an arithmetic-cost bound; it does not assert a bit-complexity bound uniform in noise or conditioning. □

G.3 A COMPLETE TASK QUOTIENT AND SMOOTH TASK FUNCTIONALS

A related principle is value equivalence, which groups models by selected Bellman updates (Grimm et al., 2020). Here the operations of interest are prepared control tasks.

Continue with full actuation and the uniform inverse margins for A, B. Define

$$\mathcal{Z}(\theta) = (M, z_0), \qquad M = B^{-1}AB, \quad z_0 = B^{-1}x_0. \tag{110}$$

This is a complete invariant for the action $(A, B) \mapsto (SAS^{-1}, SB)$ of invertible matrices satisfying $Sx_0 = x_0$. Indeed, it is unchanged by this action. Conversely, if two pairs have the same (M, z_0), put $S = B_2B_1^{-1}$. Equality of M gives $A_2 = SA_1S^{-1}$, and equality of z_0 gives $Sx_0 = x_0$. Thus no physical calibration direction along this similarity orbit is retained. This algebraic task quotient is not a sufficient statistic for the full physical observation likelihood.

Proposition 17 (Precision of the task quotient)**.** *On a compact fully actuated, uniformly nonsingular class with nonempty interior and the preceding resources, the minimax L^p error for (M, z_0) is $\asymp R_{K+1}$ with K actionable packets and the final private record. The spectral construction is an attaining estimator with polynomial arithmetic cost. With only the K-packet public transcript, the quotient rate is $\asymp R_K$ under the corresponding time conditions.*

Proof. The final batch estimates M with error $O(R_{K+1})$. At the stored quiet clock s, let $\widehat{z}_s = \bar{B}^{-1}Y_s$. The exact history identity is

$$z_0 = M^{-s}z_s - \sum_{t<s} M^{-1-t}u_t - \tau \sum_{t<s} M^{-1-t}B^{-1}\xi_t. \tag{111}$$

The quiet-coordinate error is $O(R_{K+1})$; the inverse maps and their derivatives are bounded at the fixed inverse margin and horizon. The input sum is bounded by the hard budget, and the process sum has L^p norm $O(\tau)$. Substituting the final refined M and the stored quiet record into Eq. (111) therefore gives the same order for z_0. Project its estimate onto the ball $\|z\| \le \|x_0\|/b_-$ if needed.

Along an interior input-gain segment (A_*, bB_*), M is constant and $z_0 = z_*/b$, with $z_* \neq 0$. The stopped-buffer estimation base used for U_*/b applies verbatim to z_*/b. It gives the matching reading/depth terms, and the causal transition comparison gives the process term. The public version uses one fewer refinement cycle for its upper and exactly K disclosures for its lower, as above. □

Any task $g(\theta) = \psi(\mathcal{Z}(\theta))$ that is Lipschitz on the bounded estimated-coordinate domain consequently has an $O(R_{K+1})$ estimator. A smooth map defined on a neighborhood of that compact domain is sufficient. This is an upper-bound criterion, not a universal matching rate for every functional: constant or locally flat tasks can be easier. If a scalar contrast of $g(A_*, bB_*)$ has derivative bounded away from zero on an interior gain interval, the same interval argument supplies the matching depth and floor terms.

For example, with any known positive-definite input metric W, the weighted minimum-energy zeroing word is

$$U^*_{H,W} = -W^{-1}G_H^{\mathsf{T}}(G_H W^{-1} G_H^{\mathsf{T}})^{-1} M^H z_0, \tag{112}$$

and inherits the rate. Minimum-input-energy commands satisfying a fixed set of feasible waypoints $x_{t_j} = c_j x_0$ do so as well: in control coordinates the constraints are $z_{t_j} = c_j z_0$. The stacked input constraint matrix has full row rank, since successive waypoint blocks have an identity input block at their latest input. Its weighted least-norm solution is smooth on a fixed compact domain. Thus the quotient supports a family of prepared operations, not only one zero terminal word.

A task can require the discarded physical calibration. With radial relative sensing, if a smooth task changes to first order along the rank-one similarity used in Eq. (94), that bound gives an $\Omega(\sigma)$ error even with immediate feedback. Ordinary physical quadratic state costs need not factor through $\mathcal{Z}$. For example, take $x_0 = e_1$, $B = I$ and

$$A = \begin{pmatrix} 1/2 & 0 \\ 1 & 1/2 \end{pmatrix}, \qquad S = \operatorname{diag}(1, 1+a).$$

The one-step problem $\min_u\{\|Ax_0 + Bu\|^2 + \|u\|^2\}$ has, on the transformed plant, $u_2^* = -(1+a)^2/[1+(1+a)^2]$. Its derivative at zero is $-1/2$, although (M, z_0) is unchanged. This task has the calibration obstruction. The quotient criterion therefore distinguishes which control objectives inherit physical cancellation precision; it does not assert the higher order for all LQR or tracking problems.

A first-order obstruction along any preparation-preserving orbit. The following statement concerns radial sensing with $\beta = 1$. It requires a local varying task contrast, not merely global noninvariance on a possibly disconnected plant class.

Proposition 18 (Local orbit obstruction for a task)**.** *Let $g : \Theta \to \mathbb{R}^r$ be a task, and suppose Θ contains*

$$\theta_s = (S_s A_* S_s^{-1}, S_s B_*), \qquad S_s = I + s\mathsf{H}, \quad \mathsf{H}x_0 = 0, \qquad 0 \le s \le \bar{s},$$

where $\bar{s}\|\mathsf{H}\|_{\mathrm{op}} \le 1/2$. Assume $s \mapsto g(\theta_s)$ is differentiable at zero with $a_g = \|\frac{d}{ds} g(\theta_s)|_{s=0}\| > 0$. There is $s_g \in (0, \bar{s}]$ such that, with N readings and T transitions, every fixed $1 \le p < \infty$ satisfies

$$\inf_{\pi, \widehat{g}} \sup_{\theta \in \Theta} \|\widehat{g} - g(\theta)\|_{L^p} \ge \frac{3a_g}{16} \min\left\{ s_g, \frac{\sigma}{\sqrt{8C_{\mathsf{H}}}} \right\}, \qquad 0 < \sigma \le 1, \tag{113}$$

where

$$C_{\mathsf{H}} = N(2 + 4n)\|\mathsf{H}\|_{\mathrm{op}}^2 + \frac{3}{2} T \|\mathsf{H}\|_F^2.$$

The infimum may allow immediate feedback after every reading, arbitrary measurable policies and unrestricted input energy. The result holds for every $d, \tau \ge 0$ in the radial model. In this paper $N = T - 1$. In particular it gives an $\Omega(\sigma)$ task obstruction at fixed geometry and resources. An interior plant and a nonzero directional derivative along such an orbit supply the stated local assumptions.

Proof. Let P_s^π denote the complete visible transcript law, including the private readings. Include common policy seeds. Augment the base law by $z = x$ and the alternative law by $z = S_s^{-1} x$. Because $S_s x_0 = x_0$, their preparations coincide. At the same numerical hidden and reading histories, the policy takes the same action, even if its maps are discontinuous. Both hidden transition means are $A_* z + B_* u$. For $\tau > 0$ the covariance matrices are respectively $\tau^2 I$ and $\tau^2 S_s^{-1} S_s^{-\top}$. Writing q_j for the singular values of S_s, the base-to-alternative transition KL is

$$\frac{1}{2} \sum_j (q_j^2 - 1 - 2 \log q_j) \le \frac{3}{2} \sum_j (q_j - 1)^2 \le \frac{3}{2} s^2 \|\mathsf{H}\|_F^2.$$

The scalar inequality holds on $[1/2, 3/2]$, and the second inequality is the singular-value perturbation bound. At $\tau = 0$ both transitions are the same Dirac kernel, so this contribution is zero.

At a reading put $v^2 = d^2 + \|z\|^2$ and $v_s^2 = d^2 + \|S_s z\|^2$. If $v = 0$, both readings are the same point mass. Otherwise, with $h = s\|\mathsf{H}\|_{\mathrm{op}} \le 1/2$, $r = v_s / v \in [1 - h, 1 + h]$, and the conditional reading KL is

$$\frac{n}{2}(r^{-2} - 1 + 2 \log r) + \frac{s^2 \|\mathsf{H} z\|^2}{2\sigma^2 v_s^2} \le 4nh^2 + \frac{2h^2}{\sigma^2}.$$

Here $\frac{1}{2}(r^{-2} - 1 + 2 \log r) \le 4(r - 1)^2$ on $[1/2, 3/2]$ and $v_s^2 \ge (1 - h)^2 \|z\|^2$. Thus all covariance information is retained, including for a nonsymmetric generator H. The chain rule on the augmented law and data processing yield

$$\mathrm{KL}(P_0^\pi \| P_s^\pi) \le s^2 \left\{ N(2\sigma^{-2} + 4n)\|\mathsf{H}\|_{\mathrm{op}}^2 + \frac{3}{2} T \mathbf{1}_{\{\tau > 0\}} \|\mathsf{H}\|_F^2 \right\} \le \frac{s^2 C_{\mathsf{H}}}{\sigma^2}. \tag{114}$$

Messages and silence are functions of the same readings and seeds and add no KL. The common-history argument applies to arbitrary measurable policies because their actions agree at each common numerical history.

Choose s_g small enough that $\|g(\theta_s) - g(\theta_0)\| \ge a_g s/2$ for $0 < s \le s_g$. At $s = \min\{s_g, \sigma/\sqrt{8C_{\mathsf{H}}}\}$, the KL is at most $1/8$, hence the total variation distance is at most $1/4$. Integrating the triangle inequality against the overlap measure of the two transcript laws gives, for any estimator,

$$\max_{i \in \{0, s\}} \mathbb{E}_i \|\widehat{g} - g(\theta_i)\| \ge \frac{1}{2} \|g(\theta_s) - g(\theta_0)\| (1 - \mathrm{TV}(P_0^\pi, P_s^\pi)) \ge \frac{3a_g s}{16}.$$

Taking the infimum and using $L^p \ge L^1$ proves the claim. □

G.4 AN OPERATIONALLY COMPLETE CATALOG OF TASKS

Continue with the compact fully actuated class of Proposition 17, including the uniform inverse margins for A and B. Consider a two-step noiseless operation from x_0: the first command is prescribed as $v \in \mathbb{R}^n$, and the second command $c_\theta(v)$ must make the final state zero. Since B is invertible, this last command exists and is unique. Direct substitution gives

$$c_\theta(v) = -B^{-1}A^2x_0 - B^{-1}ABv = -M^2z_0 - Mv. \tag{115}$$

This is a family of command targets, evaluated as functionals of the unknown plant. Defining or reporting these targets supplies no additional observation or preparation to the learning experiment.

Proposition 19 (Operational completeness)**.** *Let $\mathcal{V}(\theta) = [c_\theta(0), c_\theta(e_1), \dots, c_\theta(e_n)]$. The finite catalog $\mathcal{V}$ and the task quotient $\mathcal{Z} = (M, z_0)$ determine each other through Lipschitz maps on the stated compact class. Their minimax L^p errors therefore have the same order: R_{K+1} with the final private record and R_K with K public disclosures, under the corresponding resources. Any representation that determines all the response commands $c_\theta(v)$ determines $\mathcal{Z}$.*

Proof. Equation (115) maps $\mathcal{Z}$ to the catalog by a polynomial map, hence is Lipschitz on its bounded domain. Conversely,

$$M = -[c_\theta(e_1) - c_\theta(0), \dots, c_\theta(e_n) - c_\theta(0)], \qquad z_0 = -M^{-2}c_\theta(0). \tag{116}$$

The first map is linear; the second is Lipschitz at the uniform inverse margin. For example, if $\|M_i^{-1}\| \le \mu$ then $\|M_1^{-2} - M_2^{-2}\| \le 2\mu^3\|M_1 - M_2\|$ by the inverse identity and telescoping. This proves bi-Lipschitz equivalence using Frobenius norm for the catalog and the Euclidean product norm for $\mathcal{Z}$.

For risk transfer from an arbitrary catalog estimate, form its matrix of negative differences and clip its singular values to the known interval $[m_-, m_+]$. Its matrix error is at most twice the unprojected error. Apply Eq. (116) and project z_0 onto its known Euclidean ball. The resulting quotient error is bounded by a fixed constant times the catalog error. The reverse transfer simply evaluates Eq. (115) on the bounded quotient estimate. These two constructions, combined with Proposition 17, prove the stated minimax orders without assuming an estimator is realizable. Finally, any representation sufficient to determine all responses determines these $n + 1$ responses and hence the quotient by Eq. (116). □

Thus a learned quotient supports more than repetition of one optimal word: it gives the correction after any prescribed first input. On $\|v\| \le V$, a catalog estimate also yields a uniform response estimate with error at most a constant depending on V times the catalog error. The qualification on A^{-1} matters: when M is singular, the catalog need not reveal components of z_0 killed by M^2, so this operational minimality statement does not extend unchanged to singular drift.

The representation is minimal for this family of input commands. It is not a minimal representation of the full observation likelihood, nor does it encode arbitrary costs in physical state coordinates. The quadratic-cost example above supplies an explicit task outside the family.

G.5 A PHYSICAL-MODEL PLUG-IN CAN PRESERVE THE TASK EXACTLY

The physical-model lower bound concerns error in the complete model. Its unresolved directions can coexist with an accurate command. The following construction makes this distinction explicit on an unchanged acquisition record. Related distinctions arise between model error and LQ control loss (Mania et al., 2019), and between certainty equivalence and absence of a dual effect (Derpich & Yüksel, 2023).

Proposition 20 (Task-preserving physical lift)**.** *Let $x_0 \neq 0$, let $\bar{B} \in \mathbb{R}^{n\times n}$ be invertible, and let $(\widehat{M}, \widehat{z})$ be a reported task representation with $\widehat{z} \neq 0$. Define*

$$\widetilde{B} = \bar{B} + \frac{(x_0 - \bar{B}\widehat{z})\widehat{z}^{\mathsf{T}}}{\|\widehat{z}\|^2}. \tag{117}$$

The matrix $\widetilde{B}$ is invertible exactly when $\widehat{z}^{\mathsf{T}}\bar{B}^{-1}x_0 \neq 0$. In particular,

$$\|x_0 - \bar{B}\widehat{z}\| < s_{\min}(\bar{B})\|\widehat{z}\| \tag{118}$$

is sufficient. Whenever the lift is invertible, put $\widetilde{A} = \widetilde{B}\widehat{M}\widetilde{B}^{-1}$. *For every integer* $H \geq 1$, *its minimum-energy physical-model plug-in command equals the canonical report exactly:*

$$-\mathcal{C}_H(\widetilde{A}, \widetilde{B})^\dagger \widetilde{A}^H x_0 = -C_H(\widehat{M})^\dagger \widehat{M}^H \widehat{z}, \qquad C_H(M) = [M^{H-1}, \ldots, I]. \tag{119}$$

This equality also holds after applying the same command guard to both reports.

Proof. Equation (117) gives $\widetilde{B}\widehat{z} = x_0$. The matrix determinant lemma yields

$$\det \widetilde{B} = \det \bar{B}\, \frac{\widehat{z}^\mathsf{T} \bar{B}^{-1} x_0}{\|\widehat{z}\|^2},$$

proving the exact criterion. Moreover, $\|\widetilde{B} - \bar{B}\| = \|x_0 - \bar{B}\widehat{z}\|/\|\widehat{z}\|$, so Eq. (118) and the singular-value perturbation bound give invertibility. For an invertible lift, $\mathcal{C}_H(\widetilde{A}, \widetilde{B}) = \widetilde{B} C_H(\widehat{M})$ and $\widetilde{A}^H x_0 = \widetilde{B}\widehat{M}^H \widehat{z}$. Multiplying the physical zeroing constraint by $\widetilde{B}^{-1}$ therefore gives exactly $C_H(\widehat{M})U = -\widehat{M}^H \widehat{z}$. The matrix C_H has full row rank because its last block is I; both representations have the same unique minimum-norm solution. A shared deterministic guard preserves equality. □

Class membership and profiling. For a physical report constrained to the original class Θ, the lift must also pass that class's membership conditions. A useful sufficient check follows from its interior margins. Let $\bar{M} = \bar{B}^{-1}\bar{A}\bar{B}$, $\Delta B = \widetilde{B} - \bar{B}$, and $\varepsilon = \|\Delta B\| < s_{\min}(\bar{B})$. Then

$$\|\widetilde{A} - \bar{A}\| \leq \frac{\varepsilon(\|\widehat{M}\| + \|\bar{A}\|) + \|\bar{B}\|\, \|\widehat{M} - \bar{M}\|}{s_{\min}(\bar{B}) - \varepsilon}. \tag{120}$$

Indeed, $(\widetilde{A} - \bar{A})\widetilde{B} = \Delta B \widehat{M} - \bar{A}\Delta B + \bar{B}(\widehat{M} - \bar{M})$. Together with the distance of $(\bar{A}, \bar{B})$ from the physical class boundary, Eq. (120) certifies membership whenever those margins exceed the two perturbation bounds.

A likelihood fit can also profile the nuisance matrix B while holding $(\widehat{M}, \widehat{z})$ fixed, imposing $B\widehat{z} = x_0$ and setting $A = B\widehat{M}B^{-1}$. Every invertible feasible candidate gives the same command in Eq. (119), regardless of the profile's objective value. This is a task-preserving physical-model baseline; global full-model likelihood optimization is a separate estimation choice. At a class boundary, independently projecting A and B can move the task representation, so the resulting command change belongs in the projection diagnostics.

The practical comparison is therefore between acquisition policies under a specified final reporter. A lower bound for physical-model Frobenius error alone supplies no lower bound for all physical-model plug-in commands. The finite gain example in Appendix J makes acquisition decisive even when a physical-model plug-in exactly matches the direct report.

H Information timing for the constructive policies

Private and shared histories also underlie the common-information approach to decentralized control (Nayyar et al., 2013).

A clock t labels the state before input u_t: the sensor obtains Y_t, may send a packet, and the actuator then chooses u_t. The transition produces x_{t+1}. The initial preparation is known; $Y_0 = x_0$ introduces no noisy observation. Each Y_t, $1 \leq t < T$, is generated once and cached. A shared endpoint uses this same cached reading, including when it is a query endpoint, the next query's start, or the stored quiet reading. There is no Y_T. An input interval $[a, b)$ means commands $u_a, \ldots, u_{b-1}$ and endpoint state x_b.

General construction. Put $q = mh$, $b_h = h(q + n + 1)$ and $B_h = h(q + n)$. For the constructive schedule in Appendix B,

$$t_r = L + (r-1)b_h, \qquad s = t_K + h, \qquad f = T - 1 - B_h, \qquad f - s = T - L - Kb_h - 1 \geq 0. \tag{121}$$

Table 3: Information flow for the general construction. The physical fit and chart are retained throughout; only the canonical estimate is refined. Computation of a private update does not make it actionable.

Phase	Clock/input interval	Available information and operation
Acquisition	$[0, L)$; reading Y_L	The sensor fits $(\bar{A}, \bar{B})$ from $Y_{0:L}$ and the executed prefix, selects J, and initializes (M_o, N_o).
First packet	$t_1 = L$	Send the fit, chart and Y_L. The actuator can now publish the correction and the ensuing query word.
Correction	$[t_r, t_r + h)$	Use the retained fit and the transmitted Y_{t_r} in Eq. (66); commit all h inputs. At $s = t_K + h$, cache Y_s for the later command report.
Frozen batch	$[t_r + h, t_{r+1})$, $r < K$; or $[f, T - 1)$	Freeze (M_o, N_o) for every column. Each selected column uses two h-step words; each nonselected column uses one. Intermediate readings remain private. Form the raw endpoint updates, then jointly project only after the entire batch.
Next packet	t_{r+1}, $r < K$	Send the completed canonical update and the endpoint reading. Only then can the new estimate determine the next query word. No partially updated column enters a running batch.
Final private stage	$[s, f)$, then endpoint $T - 1$	Use zero inputs on $[s, f)$. Complete the last frozen batch privately, retain its update, and apply the predetermined $u_{T-1} = 0$. There is no $(K + 1)$st actionable packet.
Offline report	After the learning trajectory	Combine the final refined coordinates, cached Y_s and actual inputs $u_0, \ldots, u_{s-1}$ to estimate the command. Clipped inputs, not proposed inputs, enter reconstruction. This changes no past action.

The rth packet is at t_r. Its correction occupies $[t_r, t_r + h)$. Batches $r < K$ occupy $[t_r + h, t_{r+1})$; the final batch occupies $[f, T - 1)$. Table 3 states which computations can affect inputs. This is an attaining schedule, not a restriction on the adaptive policy class in the lower bounds.

For a batch starting at a, selected coordinate j uses the endpoint pair $a + 2h(j - 1)$ and $a + 2hj$. After those $2hn$ inputs, nonselected columns use consecutive h-step words. Every start or endpoint is an existing physical reading. The last stored quiet reading is obtained *before* any optional padding; it is not replaced by an independently generated reading at the final batch. When $f = s$, it is also that batch's first cached endpoint. The general amplitude is $2^{-(r-1)}$. The two-state precision and underactuated protocols use 2^{-r}; the dense refinement protocol uses 2^{2-r} after its fixed development comparison.

Explicit finite schedules. Table 4 collects the fixed canonical policies in Appendices K, O and T. A listed batch start a stands for $[a, a + B_h)$. In every row, the *last* listed batch is private, its endpoint is $T - 1$, and the final input is zero. Each preceding batch ends exactly at the next packet. The Kth correction ends at s; all spare inputs between s and the final batch are zero.

Table 4: All fixed-policy clock schedules. Precision uses $(L, h, B_h, T) = (4, 1, 4, 20)$; underactuated uses $(5, 2, 8, 36)$; dense dimension n uses $(4n, 1, 2n, 10n + 4)$. The first packet includes the acquisition fit. s is a stored reading, not an additional message.

Protocol	K	Packet clocks	Batch starts	s
Precision	1	4	15	5
	2	4, 9	5, 15	10
	3	4, 9, 14	5, 10, 15	15
Underactuated	1	5	27	7
	2	5, 15	7, 27	17
	3	5, 15, 25	7, 17, 27	27
Dense $n = 4$	1	16	35	17
	2	16, 25	17, 35	26
	3	16, 25, 34	17, 26, 35	35
Dense $n = 8$	1	32	67	33
	2	32, 49	33, 67	50
	3	32, 49, 66	33, 50, 67	67
Dense $n = 16$	1	64	131	65
	2	64, 97	65, 131	98
	3	64, 97, 130	65, 98, 131	131

For the dense construction the clock convention is explicitly $t_r = 4n + (r - 1)(2n + 1)$, $f = 8n + 3$ and $u_{10n+3} = 0$. It uses $4n + K(2n + 1) + 1$ nonpadding clock slots and $(3 - K)(2n + 1)$ zero-padding slots, totaling $10n + 4$. Thus all nine declared dimension/depth combinations fit the same per-dimension horizon. Padding and the final zero input do not add input energy. These fixed schedules do not describe the adaptive comparison policies.

Public-only and charged-upload reports. In Corollary 2, use $K - 1$ actionable packets and the final physical batch shown above. A charged upload after learning is the Kth disclosure. For a public-only report, instead send its completed estimate at $T - 1$ and keep $u_{T-1} = 0$. The same physical record then yields an R_K command estimate within K public packets. With $K = 1$, the prefix fit suffices. Both interfaces use the same physical readings and allow the same total number of public disclosures.

Reporting versus a separately prepared deployment. The general theorem estimates U_H^* after learning; its resource budget covers learning and reporting; execution is a separate operation. General reconstruction requires $H \geq \max\{s, h\}$, while Appendix B.8 permits a fixed shorter $H \geq h$ on a

uniformly invertible-A class. The precision study reports both $H = 20$ and $H = 3$. Its guarded three-step deployment in Appendix L starts a *new* preparation: learning $(T, E) = (20, 10)$ plus deployment $(3, 1)$ gives 23 transitions and total input energy at most 11. The underactuated study reports $H = 36$ without executing that word; its budget remains 36 transitions and energy at most 50. Dense refinement adds three deployment transitions with energy at most two to learning $(10n+4, E_n)$, where $E_4 = 274$, $E_8 = 997$ and $E_{16} = 4032$ are the explicit bounds for every record in Appendix T. The combined accounting is $(10n + 7, E_n + 2)$. Both executed deployment protocols use two preparations and no feedback during deployment. Preparation cost is separate from input energy. Each learning experiment is a single uninterrupted trajectory.

I A FINITE SUFFICIENT CERTIFICATE FOR ECHO REFINEMENT

We derive computable confidence radii for a specified fully actuated policy on a well-conditioned continuous class. These analytic bounds certify a tolerance when the corresponding radius lies below it; a larger radius leaves that tolerance unresolved. Comparing sufficient radii alone also leaves the ordering of actual risks unresolved. The example concerns the class specified here, separately from the wider numerical class in Appendix K. Noise-level bounds are inputs to the certificate; the policy can operate without them.

One event and one correction. Use radial relative sensing, $n = m$, and $N = T - 1$ private readings. Allocate $\delta_Z + \delta_\xi = \delta$. On an event of probability at least $1 - \delta$, bound all primitive innovations by $\|Z_t\| \le g_Z$, $\|\xi_t\| \le g_\xi$, and put $\nu = \sigma g_Z$, $w = \tau g_\xi$. In dimension two the exact Gaussian norm tail gives

$$g_Z = \sqrt{2\log(N/\delta_Z)}, \qquad g_\xi = \sqrt{2\log(T/\delta_\xi)}.$$

At $\tau = 0$ set $w = 0$, $\delta_Z = \delta$; otherwise use equal allocations. Thus $\|Y_t - x_t\| \le \nu(d + \|x_t\|)$ simultaneously, including reused readings and adaptive states. No independence between fitting and refinement errors is needed.

Suppose the retained fit obeys $\|A\|, \|\bar{A}\| \le A_+$, $\|B\| \le B_+$, $s_{\min}(B), s_{\min}(\bar{B}) \ge b > 0$ and $\|A - \bar{A}\| \le \epsilon_A$, $\|B - \bar{B}\| \le \epsilon_B$. All unmarked matrix norms are operator norms. Set

$$G_+ = A_+/b, \quad J_B = B_+ G_+, \quad \gamma = \epsilon_A + \epsilon_B A_+/b + J_B\nu.$$

For an incoming state bound S, the proposed correction $-\bar{B}^{-1}\bar{A}Y$ has norm at most $G_+[(1+\nu)S + \nu d]$. If this is below its command cap, its endpoint satisfies

$$S^+ \le \gamma S + J_B\nu d + w. \tag{122}$$

The implemented command is clipped on every record; the displayed test only certifies absence of clipping on the noise event.

A frozen batch and an explicit contraction test. Let $M = B^{-1}AB$ and $e \ge \|M - M_o\|_F$. Freeze M_o throughout the n pairs $ae_j, -aM_oe_j$. Define $g_0 = h_0 = 0$ and

$$\begin{aligned} g_j &= \left(\sum_{k<j} A_+^{4k}\right)^{1/2}, \qquad & h_j &= \sum_{k<j} A_+^{2k}, \\ S_j &= A_+^{2j} S_0 + aB_+ e g_j + w(1 + A_+)h_j. \end{aligned} \tag{123}$$

These bound the successive pair endpoints: expand $x_j = A^2 x_{j-1} + aB(M - M_o)e_j + \tau(A\xi + \xi')$ and apply Cauchy–Schwarz to the columns of the same frozen error matrix. For the update $\widetilde{M}e_j = M_oe_j + a^{-1}\bar{B}^{-1}(Y_j^+ - \bar{A}^2 Y_j^-)$, let $D_2 = 2A_+\epsilon_A$. Then

$$\begin{aligned} e^+ \le \frac{\epsilon_B e}{b} + \frac{1}{ba}\Bigg[& D_2\left(\sum_{j<n} S_j^2\right)^{1/2} + w(1 + A_+)\sqrt{n} \\ & + \nu\left(\sum_{j<n}[d + S_{j+1} + A_+^2(d + S_j)]^2\right)^{1/2}\Bigg]. \end{aligned} \tag{124}$$

Indeed the error terms are $\bar{B}^{-1}(B-\bar{B})(M-M_o)$, the drift difference times the starting states, and the reader/process errors. Taking the Frobenius norm after stacking columns gives the displayed sums. Nearest-Frobenius projection onto a closed convex set containing the truth preserves this bound; a nonconvex projection would cost a factor two. Zero-input padding of length ℓ replaces S by $A_+^\ell S + w\sum_{j<\ell} A_+^j$.

For $d=\tau=0$, set $q=S_0/a$ and $a^+=\rho a$. With $P=g_n$, $J=(\sum_{j<n} g_j^2)^{1/2}$ and $Q=(\sum_{j<n}(g_{j+1}+A_+^2 g_j)^2)^{1/2}$, consecutive batches separated by a correction satisfy

$$\begin{pmatrix} q^+ \\ e^+ \end{pmatrix} \le \mathcal{L}\begin{pmatrix} q \\ e \end{pmatrix}, \qquad \mathcal{L}=\begin{pmatrix} \gamma A_+^{2n}/\rho & \gamma B_+ g_n/\rho \\ (D_2+2\nu A_+^2)P/b & [\epsilon_B + B_+(D_2 J+\nu Q)]/b \end{pmatrix}. \tag{125}$$

Thus $\|\mathcal{L}\|_\infty < 1$ is a sufficient contraction condition. The floors in Eq. (124) instead introduce terms of size $(\nu d+w)/a$; ever smaller queries cannot remove them.

Quiet history and the command radius. Let e now be the final radius from Eq. (124). Suppose the truth and refined matrix have norm at most m_+ and, when inverse powers occur, smallest singular value at least $m_- > 0$. For integer k put

$$P_k=\begin{cases} m_+^k & k\ge 0, \\ m_-^k & k<0, \end{cases} \qquad L_k=\begin{cases} k m_+^{k-1} & k>0, \\ 0 & k=0, \\ |k| m_-^{k-1} & k<0. \end{cases}$$

Then $\|\widehat{M}^k - M^k\| \le L_k e$ by telescoping, including the inverse identity. At the stored quiet clock s, let

$$Z_s = S_s/b, \qquad D_z = \epsilon_B S_s/b^2 + \nu(d+S_s)/b.$$

These bound $\|B^{-1}x_s\|$ and the error of $\widehat{z}_s = \bar{B}^{-1}Y_s$. Substituting estimates in the exact identity

$$M^H B^{-1}x_0 = M^{H-s}z_s - \sum_{t<s} M^{H-1-t}u_t - \tau\sum_{t<s} M^{H-1-t}B^{-1}\xi_t$$

gives the contrast radius

$$D_c = P_{H-s}D_z + e\left[L_{H-s}Z_s + \sum_{t<s} L_{H-1-t}b_t\right] + \frac{w}{b}\sum_{t<s} P_{H-1-t}, \tag{126}$$

where b_t bounds the norm of the actual executed input. It can be the observed history or a valid preceding envelope. For $H \ge s$ all exponents are nonnegative. For shorter tasks the inverse margin and the length of the quiet history are explicit.

For $C_H(M)=[M^{H-1},\dots,I]$, write $G_H=\sum_{j<H} m_-^{2j}$, $D_H=(\sum_{j<H} L_j^2)^{1/2}$ and $C_0 = m_+^H\|x_0\|/b$. The canonical reachability Gramian is at least $G_H I$. The full-row pseudoinverse difference identity gives $\|\widehat{C}_H^\dagger - C_H^\dagger\| \le 2D_H e/G_H$: its two terms are bounded respectively by $\|\widehat{C}_H - C_H\|/G_H$. Consequently $\widehat{U} = -\widehat{C}_H^\dagger \widehat{c}_H$ has error at most

$$D_U = G_H^{-1/2}D_c + 2D_H e C_0/G_H. \tag{127}$$

This is a confidence radius on the original joint noise event; no expected-risk conclusion is inferred from it.

A fully specified continuous-class evaluation. Take $n=2$, $x_0=e_1$, $\|A-I\|,\|B-I\|\le r$, and prefix $(-e_1,.5e_2,0,.5e_1)$ of energy 1.5. With nominal states $x_{t,0}$, let $d_0=0$ and $d_{t+1}=(1+r)d_t + r(\|x_{t,0}\|+\|u_t\|)$. The nominal regressor has $s_{\min} > .2312$, verified by positive rational LDL pivots of $W_0W_0^\top - .2312^2 I$; hence $q=.2312-(d_1^2+d_2^2+d_3^2)^{1/2}$ is a uniform noiseless prefix margin. Initialize $S_4=\|x_{4,0}\|+d_4+w\sum_{j<4}A_+^j$. Fit on $s_{\min}(W_Y)\ge q/2$, use (I,I) otherwise, and project both blocks onto their original convex operator balls. Every record then has $\epsilon_A,\epsilon_B\le 2r$. Use $A_+=B_+=1+r$, $b=1-r$ and the convex refinement set $\{M' : \|M'-I\|\le r_M\}$, where r_M is $r(1+r)/(1-r)$ rounded upward to six decimals. It contains the truth and $M_o=\bar{B}^{-1}\bar{A}\bar{B}$; take $m_\pm = 1\pm r_M$, $e_0 = 2\sqrt{2}r_M$.

Fix $T = 20$, $H = 3$ and $K = 1, 2, 3$. Packets are at 4, $(4, 9)$, or $(4, 9, 14)$, and quiet clocks are $5, 10, 15$. Each batch freezes its old matrix, uses $a_j = \alpha 2^{-j}$, and has two pairs. Put all spare zero commands before the last batch, which starts at 15 and ends at 19; set $u_{19} = 0$ and use no Y_{20}. Correction-command caps are $2, .5, .25$. The last batch is private. Packets convey the current reading and completed refinement; the first also conveys the retained fit. For b_t use the exact prefix norms, query bounds $a_j, a_j m_+$ and the verified correction bounds above. For every record, including all Gaussian tails,

$$E_{\mathrm{learn}} \le 1.5 + 4 + \tfrac{1}{3} + \tfrac{2\alpha^2}{3}(1 + m_+^2). \tag{128}$$

All displayed tuples pass the separate no-clipping tests on the noise event. The principal case is $r = .005$, $\sigma = .001$, $d = \tau = 0$, $\alpha = 1$, $\delta = .05$: $q > .2091$, $b = .995$, $m_- = .994949$, $g_Z < 3.447$, and $\|\mathcal{L}\|_\infty < .1156$. At $K = 3$, $S_s < 3.107\,10^{-5}$, $P_{3-15} < 1.063$ and $\sum_{t<s} L_{2-t} b_t < 24.08$. This explicitly favorable inverse geometry is part of the certificate.

Prepared execution and interpretation. Clip the reported word to the unit ball: the true command norms are less than $.608$ in every displayed case ($.593$ in the principal case), so its error cannot increase. An additional known preparation at x_0 and three noiseless transitions give the certified deployment radius

$$D_{\mathrm{dep}} = B_+ \sqrt{1 + A_+^2 + A_+^4}\, D_U.$$

The totals are 23 transitions, learning energy at most 8 plus deployment energy at most 1, and one additional known preparation. Preparation and delivery of the learned word are additional resources; deployment time and input energy are counted separately from learning. The uncontrolled three-step state norm is at least $(1 - r)^3 > .985$ in the principal case. With independent deployment process noise τ_{dep}, add $\tau_{\mathrm{dep}} \sqrt{1 + A_+^2 + A_+^4}\, \sqrt{2 \log(1/\delta_{\mathrm{dep}})}$ and budget total failure $\delta + \delta_{\mathrm{dep}}$; the tables use noiseless deployment.

Table 5: Analytic-bound evaluation; all radii are rounded upward. The principal case uses the same $T = 20, E = 8$ for every depth. The lower panel keeps $K = 3$, $\delta = .05$ and the same schedule. Unspecified values are $r = .005, \sigma = .001, \alpha = 1, d = \tau = 0$.

Principal depth	D_U	D_{dep}	energy bound
$K = 1$	$8.017\,10^{-3}$	$1.403\,10^{-2}$	7.174
$K = 2$	$1.624\,10^{-3}$	$2.841\,10^{-3}$	7.174
$K = 3$	$2.387\,10^{-4}$	$4.175\,10^{-4}$	7.174
Changed geometry/noise/amplitude		D_{dep}	energy bound
$r = .0025,\ \sigma = .001$		$4.765\,10^{-5}$	7.171
$r = .0025,\ \sigma = .002$		$1.386\,10^{-4}$	7.171
$r = .005,\ \sigma = .002$		$8.213\,10^{-4}$	7.174
$r = .01,\ \sigma = .001$		$5.215\,10^{-3}$	7.181
$r = .01,\ \sigma = .002$		$7.738\,10^{-3}$	7.181
$\alpha = .5$		$4.451\,10^{-4}$	6.169
$\alpha = .75$		$4.184\,10^{-4}$	6.588
$\alpha = 1.25$		$4.271\,10^{-4}$	7.928
$d = .001,\ \tau = 10^{-7}$		$3.018\,10^{-3}$	7.174
$d = .01,\ \tau = 10^{-6}$		$2.628\,10^{-2}$	7.174

The third principal design therefore certifies a 10^{-3} deployment tolerance at 95% confidence. The other displayed upper bounds alone cannot establish failure, actual-risk dominance, or a minimum required packet count. The floor rows retain the same total 5% failure budget, split equally between process and reader envelopes. Their larger radii and the geometry rows delimit what this sufficient rule certifies.

I.1 A FINITE CRITERION FOR SELECTING THE QUERY RATIO

In the zero-floor comparison of Eq. (125), write

$$\mathcal{L}_\rho = \begin{pmatrix} f_{11}/\rho & f_{12}/\rho \\ f_{21} & f_{22} \end{pmatrix}, \qquad a_{j+1} = \rho a_j, \qquad 0 < \rho < 1,$$

where all four f_{ij} are the nonnegative coefficients specified by that equation and are independent of ρ. Ordinary row-sum contraction can be unnecessarily restrictive because the normalized state radius and matrix-error radius have different natural weights.

Proposition 21 (Exact contraction window for the comparison)**.** *The following statements are equivalent: (i) the spectral radius of $\mathcal{L}_\rho$ is less than one; (ii) it contracts some weighted supremum norm with strictly positive weights; and (iii)*

$$f_{22} < 1, \qquad \rho > \rho_{\min} := f_{11} + \frac{f_{12} f_{21}}{1 - f_{22}}. \tag{129}$$

Consequently a geometrically decaying query sequence with this property exists exactly when $f_{22} < 1$ and $\rho_{\min} < 1$; any $\rho \in (\rho_{\min}, 1)$ is admissible for the comparison.

Proof. A nonnegative 2×2 matrix L has spectral radius less than one exactly when both diagonal entries are less than one and $\det(I - L) > 0$. One direction follows from the nonnegative inverse $(I - L)^{-1} = \sum_{k \geq 0} L^k$; the converse follows from the quadratic formula for its largest eigenvalue. For $L = \mathcal{L}_\rho$ these conditions are $f_{22} < 1$, $\rho > f_{11}$ and $(\rho - f_{11})(1 - f_{22}) > f_{12} f_{21}$, which reduce to Eq. (129), including zero off-diagonal entries.

For the norm statement, when the spectral radius is below one choose $w = (I - L)^{-1}\mathbf{1} > 0$ and $\|v\|_w = \max_i |v_i|/w_i$. Then $Lw = w - \mathbf{1}$, so $\|L\|_w = \max_i (Lw)_i / w_i < 1$. Conversely any induced contraction norm bounds the spectral radius by its norm. The last assertion follows by intersecting the strict interval in Eq. (129) with $(0, 1)$. □

Equation (129) exactly characterizes contraction of the two-radius comparison matrix. A complete experiment certificate additionally requires the prefix-conditioning, correction-cap, time and energy checks. With amplitudes $a_j = a_0 \rho^j$ and spectral bound $\|M_j\| \leq m_+$, the query energy over J batches obeys

$$E_{\text{query}} \leq n(1 + m_+^2) a_0^2 \frac{1 - \rho^{2J}}{1 - \rho^2}. \tag{130}$$

Thus a small ratio saves query energy but can destroy normalized-state contraction; a ratio near one spends more energy while relaxing that constraint. This tradeoff can be evaluated before any record is observed.

For a finite candidate set of depths, ratios and amplitudes, propagate Eqs. (122)–(124) on the same simultaneous Gaussian event, verify all correction caps and the all-path energy bound, and compute the resulting command/deployment radius. Selecting the smallest certified radius among these feasible candidates retains the stated $1 - \delta$ guarantee. Selection is deterministic from the class and noise bounds, before a record is observed. This optimizes the displayed certificate over that candidate set, not the actual risk over all admissible policies. Positive readout or process floors require the affine recurrences, because the normalized floor terms grow as amplitudes shrink.

The displayed comparison window. Figure 5(a) uses $n = 2$, $A_+ = B_+ = 1 + r$, $b = 1 - r$, $\epsilon_A = \epsilon_B = 2r$, $N = 19$, $\delta = .05$ and $\nu = \sigma\sqrt{2\log(N/\delta)}$ in Eq. (125). Each of the 101 radii in $[.001, .1]$ and 101 noise scales in $[.0001, .03]$ is logarithmically spaced. The coarse bounds are uniform class-diameter bounds, not a fitted small-noise rate. Of 10,201 grid points, 8,555 contract at $\rho = 1/2$, 738 admit only a larger ratio, and 908 admit no ratio below one. The ordinary row-sum test certifies 8,085 points at $\rho = 1/2$. At the marked point $(r, \sigma) = (.005, .001)$, $\rho_{\min} = .02543$ and the spectral radius at $\rho = 1/2$ is $.09506$. The plot classifies this homogeneous envelope alone. For example, 54 half-ratio contraction points have a nonpositive prefix-margin lower bound under the separate sufficient prefix test; that test cannot certify them. A fixed-N Gaussian event supports the stated finite experiment, not an infinite-time probability guarantee.

J A BROAD GAIN INTERVAL WITH AN EXPLICIT FINITE BUDGET

This example isolates the operational effect of changing acquisition. It assumes calibrated drift $A = 1$, no learning process noise, and one unknown scalar actuator gain. The multidimensional physical-model obstruction is a separate result. All displayed constants evaluate analytic bounds.

Plant, reader, and resource budget. Let $x_{t+1} = x_t + bu_t$, $x_0 = 1$, and $b \in [\underline{b}, \overline{b}]$ with $0 < \underline{b} < 1 < \overline{b}$. The learning run has no process noise. Its reader is

$$Y_t = x_t + \sigma\sqrt{d^2 + x_t^2}\, Z_t, \qquad Z_t \overset{\text{iid}}{\sim} N(0,1).$$

Set $\widehat{b}_0 = 1$ and execute $u_0 = -1$. For $j = 1, \dots, K$, the packet at clock j permits

$$\begin{aligned} c_\sigma &= \frac{2}{1 + \sqrt{1 + 4\sigma^2}}, \\ \widehat{b}_j &= \mathrm{clip}_{[\underline{b},\overline{b}]}\{\widehat{b}_{j-1}(1 - c_\sigma Y_j)\}, \\ u_j &= \widehat{b}_{j-1}^{-1} - \widehat{b}_j^{-1}. \end{aligned} \tag{131}$$

All subsequent inputs are zero. At a fixed final private clock $s \geq K + 1$, apply the same estimator once more to Y_s, with $\widehat{b}_K$ replacing $\widehat{b}_{j-1}$, to obtain $\widehat{b}_{K+1}$. This update does not affect any learning action. Take $T = s + 1$ and $u_s = 0$, so the protocol uses no reading at clock T. Deployment is open loop and takes no new observations. Its three-step word has three equal entries $-1/(3\widehat{b}_{K+1})$. For general H, use H equal entries $-1/(H\widehat{b}_{K+1})$.

Proposition 22 (Finite gain refinement with a reader floor)**.** *Write*

$$r_0 = \max\{1 - \underline{b}, \overline{b} - 1\}, \quad \lambda = (1 - c_\sigma)^2 + c_\sigma^2\sigma^2, \quad r_d = c_\sigma^2\sigma^2\overline{b}^2 d^2,$$

and define $V_0 = r_0^2$, $V_{j+1} = \lambda V_j + r_d$. If $\lambda < 1$, equivalently use $V_j = r_0^2\lambda^j + r_d(1 - \lambda^j)/(1 - \lambda)$. Uniformly over the entire gain interval,

$$\begin{aligned} \mathbb{E}|\widehat{b}_j - b|^2 &\leq V_j, \\ \|x_T\|_{L^2} &\leq \sqrt{V_K}/\underline{b}, \\ \|\widehat{U}_H - U_H^*\|_{L^2} &\leq \sqrt{V_{K+1}}/(\sqrt{H}\,\underline{b}^2), \\ \|x_H^{\text{dep}}\|_{L^2} &\leq \sqrt{V_{K+1}/\underline{b}^2 + \delta_{\text{prep}}^2 + H\tau_{\text{dep}}^2}. \end{aligned} \tag{132}$$

The last line assumes an independent zero-mean preparation error of RMS δ_{prep} and independent zero-mean deployment process increments of variance τ_{dep}^2. On every observation history, the learning and deployment energies obey

$$E_{\text{learn}} \leq 1 + K(\underline{b}^{-1} - \overline{b}^{-1})^2, \qquad E_{\text{dep}} \leq (H\underline{b}^2)^{-1}. \tag{133}$$

Proof. The input increments telescope. Just before packet j the cumulative input is $-1/\widehat{b}_{j-1}$, so $x_j = 1 - b/\widehat{b}_{j-1}$. Write $e = \widehat{b}_{j-1} - b$. Before interval projection the new error is

$$(1 - c_\sigma)e - c_\sigma\sigma\widehat{b}_{j-1}\sqrt{d^2 + (e/\widehat{b}_{j-1})^2}\, Z_j.$$

Conditionally on the preceding history its squared expectation equals $\lambda e^2 + c_\sigma^2\sigma^2\widehat{b}_{j-1}^2 d^2$. Projection onto an interval containing b cannot increase the squared error. This gives the recurrence, including the independent final private reading. The terminal state is $(\widehat{b}_K - b)/\widehat{b}_K$. The deployment command has noiseless terminal state $(\widehat{b}_{K+1} - b)/\widehat{b}_{K+1}$, while its Euclidean command error is $|\widehat{b}_{K+1} - b|/(\sqrt{H}\,b\widehat{b}_{K+1})$. The preparation and process contributions have zero cross terms by independence and zero means. Finally, each reciprocal increment has magnitude at most $\underline{b}^{-1} - \overline{b}^{-1}$ and the uniform deployment word has squared norm $1/(H\widehat{b}_{K+1}^2)$. These bounds include every Gaussian tail. □

Table 6: Uniform RMS upper bounds on $[.5, 1.5]$, rounded upward. All methods share five learning and three deployment transitions. The final column includes the stated preparation and deployment-process floors. There is no learning process noise.

K	Learning-energy bound	Command-induced deployment RMS	Total deployment RMS
1	$25/9$	2.499×10^{-3}	2.501×10^{-3}
2	$41/9$	1.949×10^{-4}	2.218×10^{-4}
3	$19/3$	1.500×10^{-4}	1.836×10^{-4}

Likelihood interpretation of the update. At $d = 0$, for one unrestricted scalar mean x and a nonzero observation $Y \sim N(x, \sigma^2 x^2)$, the global maximum-likelihood estimate is $\widehat{x} = c_\sigma Y$. Indeed an optimum has the sign of Y; differentiating its negative log likelihood gives $\sigma^2 x^2 + Yx - Y^2 = 0$, whose root of that sign is $c_\sigma Y$. At $Y = 0$, define the estimate as zero. Equation (131) therefore uses the exactly solved unrestricted one-reading likelihood followed by a safety projection when $d = 0$. This one-reading interpretation differs from a constrained full-history likelihood fit. At $d > 0$, the update is justified by the moment bound above. The physical-model plug-in $(\widehat{A}, \widehat{B}) = (1, \widehat{b}_{K+1})$ gives exactly the same word as the direct report. The comparisons below therefore vary acquisition, allowing each noninteractive design an unrestricted final decoder.

A strong fixed-acquisition comparator first applies $u_0 = -1$, then uses only zero inputs and receives all $N = T - 1$ readings. They are iid with mean $x = 1 - b$ and variance $\sigma^2(d^2 + x^2)$. For $d > 0$, let $S = \sum_i Y_i$ and $Q = \sum_i Y_i^2$. Its full-record constrained global maximum-likelihood fit is obtained by evaluating the negative log likelihood at the interval endpoints and at every real root, inside $[1 - \overline{b}, 1 - \underline{b}]$, of

$$N\sigma^2 x^3 + Sx^2 + [N(1 + \sigma^2)d^2 - Q]x - Sd^2 = 0. \tag{134}$$

This finite candidate rule has no initialization or convergence choice. The comparator has the same time and energy allowances, although its actual learning energy is only one. We refer to this specified word as the fixed-null acquisition.

A finite-accuracy example with noise floors. Take $[\underline{b}, \overline{b}] = [.5, 1.5]$, $\sigma = .05$, $d = .001$, $T = 5$, $s = 4$, $H = 3$, $\delta_{\mathrm{prep}} = 10^{-4}$, and $\tau_{\mathrm{dep}} = 2 \times 10^{-5}$. The absolute reader-noise floor is $\sigma d = 5 \times 10^{-5}$. A common allowance $E_{\mathrm{learn}} = 19/3$ supports $K = 1, 2, 3$; deployment adds $E_{\mathrm{dep}} = 4/3$, three transitions, and one independent preparation having the stated RMS accuracy. Thus the complete budget is eight transitions, energy at most $23/3$, and that preparation. The preparation tolerance is an explicit external resource assumption. Its creation requires a preparation facility with separately specified time and energy costs. The displayed quantities are rigorous uniform RMS upper bounds, rounded upward from the closed-form recurrence.

The third design certifies the target 2×10^{-4} throughout a three-to-one gain interval. The first two upper bounds alone do not show that those depths fail the target.

Proposition 23 (Every deterministic nonadaptive word misses the target)**.** *For the same gain interval, reader, and four private readings, every fixed deterministic acquisition word and every possibly randomized report estimator satisfy*

$$\sup_{b \in [.5, 1.5]} \|x_H^{\mathrm{dep}}\|_{L^2} > 6.685 \times 10^{-4}.$$

This holds regardless of the acquisition energy and already holds without preparation or deployment process errors. It therefore applies under the common energy allowance and with those independent errors. Proposition 24 extends the conclusion to randomized nonadaptive designs with a smaller lower bound.

Proof. For a fixed input word let $S_t = \sum_{k<t} u_k$, so each observed mean is $1 + bS_t$. The nonzero S_t, $t = 1, \ldots, 4$, supply at most four roots $r_t = -1/S_t$. Partition $[.5, 1.5]$ at the roots inside it. One gap has length at least $1/5$. Its midpoint b_0 lies in $[.6, 1.4]$ and is at distance at least $R = .1$ from every root, including roots outside the gain interval. Choose $h = .0015$ and $b_\pm = b_0 \pm h$.

For a reading with $S_t \neq 0$, the two means differ by $2hS_t$. Both mean magnitudes are at least $|S_t|(R - h)$. Put $q = (R + h)/(R - h) > 1$. The variance ratio lies between q^{-2} and q^2; adding

the common nonnegative $\sigma^2 d^2$ moves that ratio toward one. Moreover, the squared mean difference divided by either variance is at most $4h^2/[\sigma^2(R-h)^2]$. Since $z-1-\log z$ is largest at $z=q^2$ on $[q^{-2}, q^2]$, the KL divergence of the four independent readings in either direction is at most

$$\overline{\mathcal{K}} = \frac{4}{2}\left[q^2 - 1 - 2\log q + \frac{4h^2}{\sigma^2(R-h)^2}\right] < .74577.$$

Readings with $S_t = 0$ contribute zero divergence. Thus the total variation distance between the two full-record laws obeys $\mathrm{TV} \le \sqrt{\overline{\mathcal{K}}/2}$.

Let a denote the sum of a proposed deployment word. At every record and every independent estimator seed,

$$(1+b_-a)^2 + (1+b_+a)^2 \ge \frac{2h^2}{b_0^2+h^2},$$

by minimizing the left side over $a \in \mathbb{R}$. Integrating against the minimum of the two record densities and using that its mass is $1-\mathrm{TV}$ gives

$$\max_{b\in\{b_-,b_+\}} \mathbb{E}|1+ba|^2 \ge \frac{h^2}{b_0^2+h^2}(1-\mathrm{TV}).$$

Use $b_0 \le 1.4$ and $\overline{\mathcal{K}} < .74577$ to obtain

$$\sup_b \|x_H^{\mathrm{dep}}\|_{L^2} \ge \frac{.0015}{\sqrt{1.4^2+.0015^2}}\sqrt{1-\sqrt{.74577/2}} > 6.685\times 10^{-4}.$$

Independent zero-mean preparation and deployment errors add nonnegative squared risk. The construction of $b_\pm$ depends on the fixed acquisition word. Proposition 24 instead averages over a fixed gain prior to handle randomized acquisition. □

The interactive uniform upper bound is therefore more than 3.6 times smaller than a lower bound for every deterministic nonadaptive acquisition, even with unrestricted final decoding.

For the particular fixed-null acquisition, an even larger lower bound is available, including for its globally optimized likelihood report. To see this, compare $b_0 = .50$ and $b_1 = .51$, and put $\mu_i = 1-b_i$, $v_i = \sigma^2(d^2+\mu_i^2)$. The KL divergence between its four-reading laws is

$$\mathcal{K} = \frac{4}{2}\left[\log(v_1/v_0) + \frac{v_0+(\mu_0-\mu_1)^2}{v_1} - 1\right] < .33485.$$

For any estimated deployment word, let $\widehat{a}$ be the sum of its entries. The noiseless terminal loss at gain b_i is $b_i|\widehat{a}+1/b_i|$. Testing the two target sums at their midpoint and using Pinsker's inequality yields

$$\sup_{b\in[.5,1.5]} \|x_H^{\mathrm{dep}}\|_{L^2} \;\ge\; .5\left|\frac{1}{.50}-\frac{1}{.51}\right|\sqrt{\frac{1-\sqrt{\mathcal{K}/2}}{8}} \;>\; 5.328\times 10^{-3}. \tag{135}$$

The same lower bound holds after adding the independent zero-mean preparation and process errors, which add nonnegative squared risk. For this specified fixed-null word with unrestricted decoding, the lower bound exceeds the interactive upper bound by a factor greater than 29.

Proposition 24 (Randomized nonadaptive acquisition)**.** *With the same gain interval, four readings, $\sigma = .05$, and any $d \ge 0$, every nonadaptive acquisition design whose random seed is independent of the gain and observation noise satisfies*

$$\sup_{b\in[.5,1.5]} \|x_H^{\mathrm{dep}}\|_{L^2} > 2.960\times 10^{-4}.$$

The estimator may observe the design seed and use additional independent randomness. Acquisition energy is unrestricted in this lower bound.

Proof. Condition on any design seed. Independently draw a base gain v uniformly on $[.51, 1.49]$, and consider the two true gains $b_\pm = v \pm h$ with $h = .0012$. Let $R = .08$. The four roots exclude at most $8R = .64$ total length from the base-gain interval when neighborhoods of radius R are removed.

Consequently a fraction at least $1 - .64/.98 = 17/49$ of its base gains has distance at least R from every root, for every design seed. On these good pairs of base gain and seed, the preceding Gaussian KL calculation applies. Because $h/R = .015$ as before, it gives the same bound $\overline{\mathcal{K}} < .74577$. For each such pair, the density-overlap calculation gives

$$\mathbb{E}_{v-h}[|1 + (v-h)a|^2 \mid \text{seed}] + \mathbb{E}_{v+h}[|1 + (v+h)a|^2 \mid \text{seed}] \geq \frac{2h^2}{v^2 + h^2}(1 - \sqrt{.74577/2}).$$

For the remaining pairs the sum is nonnegative. Average over the base gain and the design seed, and use $v \leq 1.49$. The left side is a sum of two Bayes risks for fixed shifted uniform priors independent of the seed; each is at most the worst-case risk. Therefore

$$\sup_b \mathbb{E}|1 + ba|^2 \geq \frac{h^2}{1.49^2 + h^2} \frac{17}{49}(1 - \sqrt{.74577/2}).$$

Its square root exceeds 2.960×10^{-4}. This rounding can be checked with rational arithmetic: $q = 203/197$, the sixth-order alternating lower bound for $\log(1 + 6/197)$ gives $\overline{\mathcal{K}} < .74577$, hence $\mathrm{TV} < .610644$; substitution gives a squared lower bound greater than $(.000296)^2$. The good subset may depend on the seed, but the prior averaged on the left does not. Thus this argument does not select a true gain after observing the design seed. Independent zero-mean deployment errors again add nonnegative squared risk. □

This lower bound also excludes the 2×10^{-4} target for every randomized nonadaptive acquisition design.

K Design-panel precision simulation and implementation checks

We estimate expectations by antithetic Monte Carlo simulation to compare the retained physical fit, executed terminal state and offline command on the same trajectory. The calculation evaluates the selected policy on a finite plant panel. Policy choices were fixed after the exploratory design calculations described below and before this evaluation; the resulting panel summaries do not estimate a continuous-class supremum or an optimum over policies.

K.1 Class, panel and common acquisition

Use $x_0 = (1, 0)^\top$ and the full eight-parameter class

$$A_0 = \begin{pmatrix} .65 & .22 \\ -.08 & .58 \end{pmatrix}, \qquad B_0 = \begin{pmatrix} 1 & .15 \\ .08 & .9 \end{pmatrix}, \qquad \|A - A_0\|_{\mathrm{op}},\ \|B - B_0\|_{\mathrm{op}} \leq .05.$$

All matrix entries are unknown. The panel contains the center (index zero) and 24 additional plants. For $j = 1, \ldots, 8$, form the matrices

$$\widetilde{D}_{A,j} = \begin{pmatrix} \sin(j + .37) & \sin(2j + .37) \\ \sin(3j + .37) & \sin(4j + .37) \end{pmatrix}, \qquad \widetilde{D}_{B,j} = \begin{pmatrix} \cos(j + .71) & \cos(2j + .71) \\ \cos(3j + .71) & \cos(4j + .71) \end{pmatrix}.$$

Normalize each separately in operator norm to obtain $D_{A,j}, D_{B,j}$. For $f = -1, -1/2, 1$, in that order, include $(A_0 + .05 f D_{A,j}, B_0 + .05 f D_{B,j})$; its index is $3(j-1) + 1, 2, 3$ respectively. The same 25 plants are used in Appendix N.

The common first four inputs, stated numerically to make the fixed policy reproducible, are the rows of

$$P = \begin{pmatrix} -.6722972972972974 & .14864864864864866 \\ -.1855529694408922 & .7266843162829836 \\ 0 & 0 \\ .7266843162829836 & .1855529694408922 \end{pmatrix}. \tag{136}$$

Their energy is 1.59908007669832. This prefix was selected in an exploratory screen of 36 distinct candidates. Starting with $c = -B_0^{-1} A_0 x_0$, that screen included the original $(0, e_2, 0, e_1)$ and words $(\gamma c, a R_\phi e_2, 0, a R_\phi e_1)$ for $\gamma \in \{0, .5, 1, 1.25\}$, $a \in \{.5, .75, 1\}$, $\phi \in \{-.25, 0, .25\}$ with energy below 2.71, excluding the original duplicate, together with the three words $(c, a e_1, 0, a e_2)$. The

criterion was the worst of the two readers' 25-plant, first-order terminal-sensitivity ratios to the original six-input completion. Candidates also had to pass a positive uniform regressor-margin bound, a sufficient completion-energy bound, and nominal completion residual below 10^{-9}. The selected parameters are $(\gamma, a, \phi) = (1, .75, .25)$. We selected this prefix on the design panel and then fixed it for evaluation.

For the precision simulation, $T = 20$, $E = 10$, $\beta = 1$, $d = \tau = 0$, and sensing is radial: $Y_t = x_t + \sigma\|x_t\|Z_t$. The initial reading $Y_0 = x_0$ is exact. Every physical Y_t, $1 \le t \le 19$, is generated once and cached; reusing an endpoint reuses the same reading. No Y_{20} is available. Let $Y_- = [Y_0, \dots, Y_3]$, $Y_+ = [Y_1, \dots, Y_4]$ and $W = [Y_-^\top, P]^\top$, where equivalently $W = [Y_-; P^\top]$. On $s_{\min}(W) \ge .174$ use $Y_+W^\dagger$; otherwise use $[A_0, B_0]$. Project each resulting block onto its original operator ball. Explicitly, if $D = U\,\mathrm{diag}(s_i)V^\top$, replace it by $U\,\mathrm{diag}(\min\{s_i, .05\})V^\top$. Denote the result by $(\widehat{A}, \widehat{B})$ and retain it throughout the trajectory. The offline full-model report is this same pair, with error $F^2 = \|\widehat{A} - A\|_F^2 + \|\widehat{B} - B\|_F^2$. All four evaluated noise levels are within the common report certificate for $(\epsilon, \delta) = (.025, .1)$; this certificate follows from the analytical prefix bound in Appendix M.

K.2 Queries, packets and reports

Write $M = B^{-1}AB$, $z_0 = B^{-1}x_0$, and $M_0 = B_0^{-1}A_0B_0$. Initialize $M^{(0)} = \widehat{B}^{-1}\widehat{A}\widehat{B}$. Every refinement is projected onto $\{M' : \|M' - M_0\|_{\mathrm{op}} \le .2\}$ by the same singular-value clipping rule. The true class lies inside this ball: a uniform radius bound is .165133. Throughout the ball, $\|M'\|_{\mathrm{op}} \le .923029$ and $\|(M')^{-1}\|_{\mathrm{op}} \le 2.892181$. The three schedules are

K	packet clocks	batch starts	stored quiet clock s
1	4	15	5
2	4, 9	5, 15	10
3	4, 9, 14	5, 10, 15	15

At packet number p, execute $-\widehat{B}^{-1}\widehat{A}Y_t$, clipped to a Euclidean ball of radius two for $p = 1$ and 2^{1-p} thereafter. The packet includes the retained fit and current reading; later packets also convey the completed refinement. At batch $r = 1, \dots, K$, set $a_r = 2^{-r}$ and freeze M_{old}. For $j = 1, 2$, execute the two inputs $a_re_j, -a_rM_{\mathrm{old}}e_j$. If their endpoint clocks are $t, t+2$, set

$$M_{\mathrm{new}}e_j = M_{\mathrm{old}}e_j + a_r^{-1}\widehat{B}^{-1}(Y_{t+2} - \widehat{A}^2Y_t).$$

Only after both columns are obtained is the whole matrix projected. Intermediate noisy readings do not enter this update. Insert zero commands before the last batch until it starts at 15. This final batch is private and uses no packet; $u_{19} = 0$. Every query and correction is charged to energy. In particular, with $\overline{M} = \|M_0\|_{\mathrm{op}} + .2$,

$$\sum_{t=0}^{19}\|u_t\|^2 \le \|P\|_F^2 + 4 + \frac{1}{3} + \frac{2(1+\overline{M}^2)}{3} = 7.167068015 < 10$$

on every record, including all Gaussian tails. This deterministic bound establishes the hard-budget guarantee.

Let $M^\sharp$ be the final private refinement and store the reading immediately after the last correction. With $\widehat{z}_s = \widehat{B}^{-1}Y_s$, reconstruct

$$z_0^\sharp = (M^\sharp)^{-s}\left[\widehat{z}_s - \sum_{t<s}(M^\sharp)^{s-1-t}u_t\right]. \tag{137}$$

The inputs here are the actual executed inputs, including clipping. The auxiliary final-state reconstruction instead uses $\widehat{z}_{20} = \widehat{B}^{-1}\widehat{A}Y_{19}$ and all twenty inputs. Setting that terminal estimate to zero defines a separate, explicitly labeled ablation; it is not the estimator used for the precision claim. The quotient error is $Q^2 = \|M^\sharp - M\|_F^2 + \|z_0^\sharp - z_0\|^2$.

For horizon H, let $C_H(M) = [M^{H-1}, \dots, M, I]$ and $U_H^* = -C_H(M)^\dagger M^H z_0$. This is the minimum-energy zeroing command for the original preparation, reported offline rather than ex-

ecuted on this trajectory. The primary $H = 20$ report avoids inverse powers:

$$\widehat{c}_{20} = (M^\sharp)^{20-s}\widehat{z}_s - \sum_{t<s}(M^\sharp)^{19-t}u_t,$$
$$\widehat{U}_{20} = -C_{20}(M^\sharp)^\dagger \widehat{c}_{20}.$$

Also compute the algebraically equivalent plug-in using $z_0^\sharp$, a three-step plug-in U_3, and their counterparts based only on $(M^{(0)}, \widehat{B}^{-1}x_0)$. The latter is a specified coarse estimator, not the best full-fit estimator. Relative command error divides each plant's squared command error by its own $\|U_{20}^*\|^2$ before aggregation.

K.3 Monte Carlo aggregation and diagnostics

The four fixed levels are $\sigma = .0004, .0007, .001, .0013$. Eight independent blocks use PCG64 seeds $47111000 + b$, $b = 0, \ldots, 7$. Each draws 256 independent arrays of shape 20×2 of standard Gaussian coordinates, then appends their negatives, giving 512 signed paths. Clock-zero primitives are unused. The same primitives are paired across plants, levels and depths; each controlled trajectory generates its own readings. All records remain in the expectations. There are $25 \cdot 4 \cdot 3 \cdot 8 = 2400$ block rows, plus 75 zero-noise arithmetic checks.

For each metric and plant p, let m_{pb} be its block MSE. First compute $r_p = (8^{-1}\sum_b m_{pb})^{1/2}$, then display $\max_p r_p$. If p_* maximizes it, the displayed standard error is

$$\mathrm{SE}(r_{p_*}) = \frac{\mathrm{sd}(m_{p_*1}, \ldots, m_{p_*8})}{\sqrt{8}\, 2r_{p_*}}.$$

The block SE measures Monte Carlo variability at the selected plant, conditional on the fixed panel; it is not adjusted for selecting the maximum. The supplement-wide statistical conventions apply. Slopes are ordinary least-squares slopes of log panel RMS against log σ, using all four levels and an intercept.

The observed maximum energy is 3.078546; prefix fallbacks, full-report failures, correction clips and refined-matrix projections all number zero. The original A, B projections at boundary plants remain included.

Implementation and arithmetic checks. The checks vary the implementation or arithmetic on specified records:

Target	Variation	Coverage and reported result
Clock and algebra	A second clock-by-clock implementation using a direct inverse fit, spectral-calculus projection and a Gram-matrix command solve	Reproduced 18,432 paths in 36 complete blocks.
Fallback and clipping branches	Synthetic extreme readings	Eighteen records exercised nine fallbacks, eighteen correction clips and eighteen refinement projections.
Cancellation arithmetic	70-digit decimal arithmetic	Seven complete replays, including smallest-noise $K = 3$ and zero-noise checks.

These checks concern the specified implementation branches and records; the simulation summaries continue to use the original sample.

The exact twenty-step command norm ranges from $2.175792\,10^{-5}$ to $2.412145\,10^{-4}$ across the panel. The largest sampled coordinate difference between direct and plug-in command reports is $5.693\,10^{-18}$. At $K = 3$, the panel maxima at zero noise are $5.019\,10^{-14}$ for the quiet quotient, $1.169\,10^{-17}$ for terminal error and $4.832\,10^{-18}$ for direct command error. The two smallest-noise high-precision command spots differ from double precision by about .041% and .072%; thus arbitrarily small errors should not be interpreted as exact arithmetic.

Primary precision trends. The three main objectives can be read directly from the following rows of Table 7. The retained model remains first order, the learning terminal gains one order per actionable packet, and the offline command gains the final private-batch order. Tables 8–10 report both reconstructions, the zero-terminal ablation, and the direct–plug-in arithmetic comparison.

Finite-range slope	$K = 1$	$K = 2$	$K = 3$
Physical-model error F	1.00000	1.00000	1.00000
Learning-terminal norm x_{20}	0.99790	1.99723	2.99593
Offline command $U_{20,\mathrm{direct}}$	1.99644	2.99658	3.99386

Table conventions. Each cell gives the panel RMS above $(100\,\mathrm{SE}/\mathrm{RMS}\%); p_*$: parentheses contain the relative block SE, followed by the maximizing plant index. Powers of ten use standard scientific notation. The tables separate model errors, task-coordinate and decoding errors, state norms, command errors, relative command errors, and arithmetic agreement. Here $z_0 = B^{-1}x_0$ is the task's initial coordinate; $z^{\sharp}_{0|s}$ and $z^{\sharp}_{0|20}$ reconstruct it using the quiet and final state, respectively. The $\widehat{z}_s, \widehat{z}_{20}$ rows measure decoding error against the actual $B^{-1}x$ at those clocks. The "zero" row reconstructs z_0 using the zero-terminal ablation; "coarse" uses the retained full fit. State rows report $\|x_{20}\|$ and $\|x_s\|$, rather than estimation error. The final, separate block reports direct–plug-in arithmetic disagreement; its slope describes roundoff behavior rather than a precision law.

Table 7: Descriptive log-log slopes of precision diagnostics over all four declared levels.

Metric	$K = 1$	$K = 2$	$K = 3$
Model errors			
F	1.00000	1.00000	1.00000
$M^{(0)}$	1.00001	1.00001	1.00001
$M^{\sharp}$	2.00004	2.99456	3.99910
Task coordinates and decoding			
$z^{\sharp}_{0\|s}$	1.99690	2.99547	3.98400
$z^{\sharp}_{0\|20}$	1.99438	3.00236	3.99311
Q_s	1.99689	2.99547	3.98400
Q_{20}	1.99438	3.00236	3.99311
$\widehat{z}_s$	1.99839	2.98851	3.98027
$\widehat{z}_{20}$	1.99185	2.99262	4.00440
$z_{0,\mathrm{zero}}$	0.99527	1.99995	2.99346
State norms			
x_{20}	0.99790	1.99723	2.99593
x_s	1.00205	1.99466	2.99072
Command errors			
$U_{3,s}$	1.99782	2.99855	3.98877
$U_{3,20}$	1.99489	3.00596	3.99167
$U_{3,\mathrm{coarse}}$	1.00020	1.00020	1.00020
$U_{20,\mathrm{direct}}$	1.99644	2.99658	3.99386
$U_{20,\mathrm{plug}}$	1.99644	2.99658	3.99386
$U_{20,\mathrm{coarse}}$	0.99603	0.99603	0.99603
Relative command errors			
$U^{\mathrm{rel}}_{20,\mathrm{direct}}$	1.99484	2.99123	3.99262
$U^{\mathrm{rel}}_{20,\mathrm{plug}}$	1.99484	2.99123	3.99262
$U^{\mathrm{rel}}_{20,\mathrm{coarse}}$	0.99951	0.99951	0.99951
Arithmetic agreement (not risk)			
$U_{20,\mathrm{gap}}$	0.01048	0.01189	0.00938

Table 8: Precision diagnostics by metric type, $K = 1$. Entries follow the conventions above; eight blocks, 25 plants, $T = 20, E = 10$.

	Readout scale σ			
Metric	0.0004	0.0007	0.001	0.0013
Model errors				
F	1.166×10^{-3} (0.44%); 11	2.041×10^{-3} (0.44%); 11	2.915×10^{-3} (0.44%); 11	3.790×10^{-3} (0.44%); 11
$M^{(0)}$	6.899×10^{-4} (0.68%); 17	1.207×10^{-3} (0.68%); 17	1.725×10^{-3} (0.68%); 17	2.242×10^{-3} (0.68%); 17
$M^{\sharp}$	7.375×10^{-7} (1.00%); 17	2.259×10^{-6} (1.00%); 17	4.609×10^{-6} (1.00%); 17	7.790×10^{-6} (1.00%); 17
Task coordinates and decoding				
$z^{\sharp}_{0\|s}$	1.009×10^{-5} (1.66%); 22	3.085×10^{-5} (1.65%); 22	6.288×10^{-5} (1.65%); 22	1.061×10^{-4} (1.64%); 22
$z^{\sharp}_{0\|20}$	9.926×10^{-3} (2.08%); 9	3.033×10^{-2} (2.10%); 9	6.176×10^{-2} (2.10%); 9	1.041×10^{-1} (2.11%); 9
Q_s	1.011×10^{-5} (1.65%); 22	3.092×10^{-5} (1.65%); 22	6.302×10^{-5} (1.64%); 22	1.064×10^{-4} (1.64%); 22
Q_{20}	9.926×10^{-3} (2.08%); 9	3.033×10^{-2} (2.10%); 9	6.176×10^{-2} (2.10%); 9	1.041×10^{-1} (2.11%); 9
$\widehat{z}_s$	5.601×10^{-7} (2.31%); 10	1.714×10^{-6} (2.31%); 10	3.496×10^{-6} (2.30%); 10	5.904×10^{-6} (2.30%); 10
$\widehat{z}_{20}$	2.246×10^{-7} (3.03%); 10	6.855×10^{-7} (3.03%); 10	1.394×10^{-6} (3.02%); 10	2.349×10^{-6} (3.02%); 10
$z_{0,\mathrm{zero}}$	7.076×10^{0} (1.46%); 9	1.236×10^{1} (1.46%); 9	1.762×10^{1} (1.46%); 9	2.287×10^{1} (1.46%); 9
State norms				
x_{20}	1.649×10^{-4} (0.72%); 7	2.882×10^{-4} (0.72%); 7	4.114×10^{-4} (0.72%); 7	5.344×10^{-4} (0.71%); 7
x_s	5.714×10^{-4} (1.04%); 10	1.001×10^{-3} (1.04%); 10	1.431×10^{-3} (1.04%); 10	1.862×10^{-3} (1.04%); 10
Command errors				
$U_{3,s}$	1.423×10^{-6} (1.54%); 22	4.353×10^{-6} (1.53%); 22	8.877×10^{-6} (1.53%); 22	1.499×10^{-5} (1.52%); 22
$U_{3,20}$	1.240×10^{-3} (2.32%); 9	3.789×10^{-3} (2.33%); 9	7.717×10^{-3} (2.34%); 9	1.301×10^{-2} (2.34%); 9
$U_{3,\mathrm{coarse}}$	3.509×10^{-4} (0.65%); 13	6.142×10^{-4} (0.65%); 13	8.775×10^{-4} (0.64%); 13	1.141×10^{-3} (0.64%); 13
$U_{20,\mathrm{direct}}$	2.354×10^{-9} (1.13%); 7	7.199×10^{-9} (1.13%); 7	1.467×10^{-8} (1.13%); 7	2.476×10^{-8} (1.13%); 7
$U_{20,\mathrm{plug}}$	2.354×10^{-9} (1.13%); 7	7.199×10^{-9} (1.13%); 7	1.467×10^{-8} (1.13%); 7	2.476×10^{-8} (1.13%); 7
$U_{20,\mathrm{coarse}}$	1.888×10^{-6} (0.45%); 7	3.299×10^{-6} (0.45%); 7	4.705×10^{-6} (0.45%); 7	6.107×10^{-6} (0.45%); 7
Relative command errors				
$U^{\mathrm{rel}}_{20,\mathrm{direct}}$	1.657×10^{-5} (1.88%); 12	5.064×10^{-5} (1.88%); 12	1.031×10^{-4} (1.87%); 12	1.739×10^{-4} (1.87%); 12
$U^{\mathrm{rel}}_{20,\mathrm{plug}}$	1.657×10^{-5} (1.88%); 12	5.064×10^{-5} (1.88%); 12	1.031×10^{-4} (1.87%); 12	1.739×10^{-4} (1.87%); 12
$U^{\mathrm{rel}}_{20,\mathrm{coarse}}$	1.391×10^{-2} (0.82%); 12	2.433×10^{-2} (0.82%); 12	3.475×10^{-2} (0.82%); 12	4.517×10^{-2} (0.82%); 12

Continued on next page

Table 8 continued: precision diagnostics, $K = 1$.
Each cell: RMS above (relative SE%); maximizing plant index.

	Readout scale σ			
Metric	0.0004	0.0007	0.001	0.0013
Arithmetic agreement (not risk)				
$U_{20,\mathrm{gap}}$	1.749×10^{-19} (0.88%); 7	1.757×10^{-19} (0.59%); 7	1.759×10^{-19} (0.63%); 7	1.773×10^{-19} (0.96%); 7

Table 9: Precision diagnostics by metric type, $K = 2$. Entries follow the conventions above; eight blocks, 25 plants, $T = 20, E = 10$.

	Readout scale σ			
Metric	0.0004	0.0007	0.001	0.0013
Model errors				
F	1.166×10^{-3} (0.44%); 11	2.041×10^{-3} (0.44%); 11	2.915×10^{-3} (0.44%); 11	3.790×10^{-3} (0.44%); 11
$M^{(0)}$	6.899×10^{-4} (0.68%); 17	1.207×10^{-3} (0.68%); 17	1.725×10^{-3} (0.68%); 17	2.242×10^{-3} (0.68%); 17
$M^{\sharp}$	2.106×10^{-9} (3.13%); 7	1.126×10^{-8} (3.12%); 7	3.276×10^{-8} (3.11%); 7	7.182×10^{-8} (3.10%); 7
Task coordinates and decoding				
$z^{\sharp}_{0\|s}$	1.433×10^{-7} (4.07%); 22	7.665×10^{-7} (4.08%); 22	2.230×10^{-6} (4.08%); 22	4.891×10^{-6} (4.09%); 22
$z^{\sharp}_{0\|20}$	1.026×10^{-5} (4.30%); 9	5.503×10^{-5} (4.31%); 9	1.606×10^{-4} (4.33%); 9	3.532×10^{-4} (4.35%); 9
Q_s	1.433×10^{-7} (4.07%); 22	7.665×10^{-7} (4.08%); 22	2.231×10^{-6} (4.08%); 22	4.892×10^{-6} (4.09%); 22
Q_{20}	1.026×10^{-5} (4.30%); 9	5.503×10^{-5} (4.31%); 9	1.606×10^{-4} (4.33%); 9	3.532×10^{-4} (4.35%); 9
$\widehat{z}_s$	3.118×10^{-10} (3.95%); 10	1.664×10^{-9} (3.93%); 10	4.828×10^{-9} (3.92%); 10	1.056×10^{-8} (3.90%); 10
$\widehat{z}_{20}$	2.377×10^{-10} (3.91%); 10	1.270×10^{-9} (3.91%); 10	3.693×10^{-9} (3.92%); 10	8.089×10^{-9} (3.92%); 10
$z_{0,\mathrm{zero}}$	5.905×10^{-3} (1.59%); 9	1.808×10^{-2} (1.60%); 9	3.690×10^{-2} (1.62%); 9	6.236×10^{-2} (1.63%); 9
State norms				
x_{20}	2.068×10^{-7} (1.68%); 7	6.326×10^{-7} (1.69%); 7	1.290×10^{-6} (1.70%); 7	2.177×10^{-6} (1.71%); 7
x_s	2.757×10^{-7} (2.81%); 10	8.425×10^{-7} (2.80%); 10	1.716×10^{-6} (2.80%); 10	2.893×10^{-6} (2.79%); 10

Continued on next page

Table 9 continued: precision diagnostics, $K = 2$.
Each cell: RMS above (relative SE%); maximizing plant index.

	Readout scale σ			
Metric	0.0004	0.0007	0.001	0.0013
Command errors				
$U_{3,s}$	2.187×10^{-8} (3.43%); 9	1.171×10^{-7} (3.46%); 9	3.413×10^{-7} (3.48%); 9	7.494×10^{-7} (3.50%); 9
$U_{3,20}$	1.461×10^{-6} (5.24%); 12	7.851×10^{-6} (5.27%); 12	2.294×10^{-5} (5.31%); 12	5.052×10^{-5} (5.34%); 12
$U_{3,\mathrm{coarse}}$	3.509×10^{-4} (0.65%); 13	6.142×10^{-4} (0.65%); 13	8.775×10^{-4} (0.64%); 13	1.141×10^{-3} (0.64%); 13
$U_{20,\mathrm{direct}}$	1.571×10^{-11} (4.13%); 7	8.409×10^{-11} (4.13%); 7	2.448×10^{-10} (4.13%); 7	5.371×10^{-10} (4.14%); 7
$U_{20,\mathrm{plug}}$	1.571×10^{-11} (4.13%); 7	8.409×10^{-11} (4.13%); 7	2.448×10^{-10} (4.13%); 7	5.371×10^{-10} (4.14%); 7
$U_{20,\mathrm{coarse}}$	1.888×10^{-6} (0.45%); 7	3.299×10^{-6} (0.45%); 7	4.705×10^{-6} (0.45%); 7	6.107×10^{-6} (0.45%); 7
Relative command errors				
$U^{\mathrm{rel}}_{20,\mathrm{direct}}$	1.417×10^{-7} (3.79%); 12	7.570×10^{-7} (3.78%); 12	2.199×10^{-6} (3.77%); 12	4.814×10^{-6} (3.75%); 12
$U^{\mathrm{rel}}_{20,\mathrm{plug}}$	1.417×10^{-7} (3.79%); 12	7.570×10^{-7} (3.78%); 12	2.199×10^{-6} (3.77%); 12	4.814×10^{-6} (3.75%); 12
$U^{\mathrm{rel}}_{20,\mathrm{coarse}}$	1.391×10^{-2} (0.82%); 12	2.433×10^{-2} (0.82%); 12	3.475×10^{-2} (0.82%); 12	4.517×10^{-2} (0.82%); 12
Arithmetic agreement (not risk)				
$U_{20,\mathrm{gap}}$	6.626×10^{-19} (0.68%); 7	6.772×10^{-19} (0.57%); 7	6.650×10^{-19} (0.73%); 7	6.765×10^{-19} (0.69%); 7

Table 10: Precision diagnostics by metric type, $K = 3$. Entries follow the conventions above; eight blocks, 25 plants, $T = 20, E = 10$.

	Readout scale σ			
Metric	0.0004	0.0007	0.001	0.0013
Model errors				
F	1.166×10^{-3} (0.44%); 11	2.041×10^{-3} (0.44%); 11	2.915×10^{-3} (0.44%); 11	3.790×10^{-3} (0.44%); 11
$M^{(0)}$	6.899×10^{-4} (0.68%); 17	1.207×10^{-3} (0.68%); 17	1.725×10^{-3} (0.68%); 17	2.242×10^{-3} (0.68%); 17
$M^{\sharp}$	4.204×10^{-12} (6.79%); 11	3.942×10^{-11} (6.78%); 11	1.641×10^{-10} (6.78%); 11	4.686×10^{-10} (6.77%); 11

Continued on next page

Table 10 continued: precision diagnostics, $K = 3$.
Each cell: RMS above (relative SE%); maximizing plant index.

	Readout scale σ			
Metric	0.0004	0.0007	0.001	0.0013
Task coordinates and decoding				
$z^{\sharp}_{0\|s}$	1.142×10^{-9} (4.71%); 22	1.064×10^{-8} (4.68%); 22	4.404×10^{-8} (4.66%); 22	1.250×10^{-7} (4.65%); 22
$z^{\sharp}_{0\|20}$	7.365×10^{-9} (4.29%); 9	6.889×10^{-8} (4.30%); 9	2.861×10^{-7} (4.31%); 9	8.149×10^{-7} (4.32%); 9
Q_s	1.142×10^{-9} (4.71%); 22	1.064×10^{-8} (4.68%); 22	4.404×10^{-8} (4.66%); 22	1.250×10^{-7} (4.65%); 22
Q_{20}	7.365×10^{-9} (4.29%); 9	6.889×10^{-8} (4.30%); 9	2.861×10^{-7} (4.31%); 9	8.149×10^{-7} (4.32%); 9
$\widehat{z}_s$	2.613×10^{-13} (8.17%); 10	2.431×10^{-12} (8.14%); 10	1.005×10^{-11} (8.09%); 10	2.846×10^{-11} (8.03%); 10
$\widehat{z}_{20}$	1.943×10^{-13} (6.53%); 24	1.827×10^{-12} (6.70%); 24	7.622×10^{-12} (6.86%); 24	2.179×10^{-11} (7.01%); 24
$z_{0,\mathrm{zero}}$	3.986×10^{-6} (3.27%); 9	2.131×10^{-5} (3.28%); 9	6.196×10^{-5} (3.30%); 9	1.358×10^{-4} (3.32%); 9
State norms				
x_{20}	1.553×10^{-10} (4.67%); 24	8.314×10^{-10} (4.71%); 24	2.420×10^{-9} (4.74%); 24	5.306×10^{-9} (4.77%); 24
x_s	2.231×10^{-10} (3.59%); 10	1.191×10^{-9} (3.56%); 10	3.460×10^{-9} (3.52%); 10	7.574×10^{-9} (3.49%); 10
Command errors				
$U_{3,s}$	2.264×10^{-10} (4.48%); 22	2.114×10^{-9} (4.47%); 22	8.766×10^{-9} (4.47%); 22	2.492×10^{-8} (4.48%); 22
$U_{3,20}$	1.131×10^{-9} (6.71%); 9	1.057×10^{-8} (6.70%); 9	4.388×10^{-8} (6.68%); 9	1.249×10^{-7} (6.67%); 9
$U_{3,\mathrm{coarse}}$	3.509×10^{-4} (0.65%); 13	6.142×10^{-4} (0.65%); 13	8.775×10^{-4} (0.64%); 13	1.141×10^{-3} (0.64%); 13
$U_{20,\mathrm{direct}}$	8.437×10^{-14} (3.91%); 7	7.894×10^{-13} (3.92%); 7	3.280×10^{-12} (3.91%); 7	9.343×10^{-12} (3.91%); 7
$U_{20,\mathrm{plug}}$	8.437×10^{-14} (3.91%); 7	7.894×10^{-13} (3.92%); 7	3.280×10^{-12} (3.91%); 7	9.343×10^{-12} (3.91%); 7
$U_{20,\mathrm{coarse}}$	1.888×10^{-6} (0.45%); 7	3.299×10^{-6} (0.45%); 7	4.705×10^{-6} (0.45%); 7	6.107×10^{-6} (0.45%); 7
Relative command errors				
$U^{\mathrm{rel}}_{20,\mathrm{direct}}$	1.281×10^{-9} (4.16%); 4	1.198×10^{-8} (4.18%); 4	4.976×10^{-8} (4.19%); 4	1.417×10^{-7} (4.20%); 4
$U^{\mathrm{rel}}_{20,\mathrm{plug}}$	1.281×10^{-9} (4.16%); 4	1.198×10^{-8} (4.18%); 4	4.976×10^{-8} (4.19%); 4	1.417×10^{-7} (4.20%); 4
$U^{\mathrm{rel}}_{20,\mathrm{coarse}}$	1.391×10^{-2} (0.82%); 12	2.433×10^{-2} (0.82%); 12	3.475×10^{-2} (0.82%); 12	4.517×10^{-2} (0.82%); 12
Arithmetic agreement (not risk)				
$U_{20,\mathrm{gap}}$	2.231×10^{-18} (1.44%); 21	2.237×10^{-18} (1.05%); 7	2.229×10^{-18} (0.89%); 10	2.266×10^{-18} (0.57%); 10

L A SEPARATELY PREPARED THREE-STEP DEPLOYMENT DIAGNOSTIC

This posthoc calculation evaluates the three-step command reports already computed in Appendix K, using its unchanged learning trajectories and reported outputs. For deployment, both the refined and coarse reports are projected onto the Euclidean unit ball in $\mathbb{R}^6$. This guard was added after the learning calculation and forms part of the deployment rule assessed here.

Task and resources. After the twenty-transition learning experiment, prepare a separate trajectory at the same $x_0 = (1, 0)^\top$ under the same unknown plant. Execute the three-command report open loop, with no new observations or feedback. This requires two preparations and 23 transitions in total, with hard input budgets 10 for learning and 1 for deployment (total 11). Preparation cost is separate from input energy, and learning uses one uninterrupted trajectory. These resources differ from the single-trajectory $T = 10, E = 2.76$ matched-controller study.

For the full matrix class, $\|A\|_{\mathrm{op}} \le .7613$ and $s_{\min}(B) \ge .77$. The last-input-only word $(0, 0, -B^{-1}A^3x_0)$ is feasible for exact zeroing and has norm at most $.7613^3/.77 < .574$. Consequently the true minimum-energy three-step word U_3^* lies inside the unit ball. If $\widehat{U}_3$ is either reported word, the deployed word is

$$\widetilde{U}_3 = \Pi_{\{U:\|U\|\le 1\}}\widehat{U}_3.$$

It satisfies the deployment energy constraint on every record. Convex projection also gives $\|\widetilde{U}_3 - U_3^*\| \le \|\widehat{U}_3 - U_3^*\|$. This does not assert that projection reduces the physical terminal residual on each individual record. Both reported words use only the learning record: the refined word uses the final $M^\sharp$ and quiet-state $z_0^\sharp$; the comparator uses the retained coarse fit.

Physical target and normalization. In physical coordinates define

$$G_3 = [A^2B, AB, B], \qquad c_3 = A^3x_0, \qquad U_3^* = -G_3^\top(G_3G_3^\top)^{-1}c_3.$$

Because $G_3 = B[M^2, M, I]$ and $c_3 = BM^3z_0$, this oracle equals the canonical-coordinate minimum-norm word used in the learning report. The deployed terminal state is $c_3 + G_3\widetilde{U}_3 = G_3(\widetilde{U}_3 - U_3^*)$. The calculation evaluates this identity and the three chronological physical state transitions. The oracle enters evaluation only; no true parameter is supplied to the deployed policy.

Use all 25 plants and the same four noise levels, three packet counts, eight blocks and 512 antithetic paths per block from Appendix K, with its unchanged seeds. Deployment here is noiseless; the learning model has radial $\beta = 1, d = \tau = 0$. For each plant p, normalize its deployment MSE by $\|A_p^3x_0\|^2$, average its eight blocks, take the square root, and then maximize over the panel. Thus the displayed statistic is the fractional residual relative to that plant's passive three-step terminal norm. Selected-plant block SEs use the same formula as the learning tables, under the supplement-wide statistical conventions.

For each declared σ, the complete tolerance grid is $\eta_j = 10^{-2-6j/240}$, $j = 0, \dots, 240$. Define the minimum among the three evaluated depths as the smallest $K \in \{1, 2, 3\}$ whose tabulated refined panel RMS is at most η_j, and report "none" if that set is empty. The table supplies all thresholds needed to reproduce this step function and the coarse-comparator pass/fail result at every grid point. Its scope is this finite set of policies on this panel.

The passive three-step terminal norm ranges from .209688 to .312821 across the panel, and true minimum zeroing energy from .0351296 to .0720766. Maximum sampled deployment energies are .0721101 for the refined command and .0726631 for the coarse command; neither has a sampled guard activation. The pointwise guard establishes feasibility, while these maxima summarize sampled usage. An estimated command can spend less than the oracle energy by leaving a nonzero residual; the oracle minimizes energy subject to exact zeroing.

Table 11: Guarded three-step deployment: normalized terminal RMS. Below each value are relative block SE in percent and maximizing plant index. The coarse column is identical across learning packet counts.

σ	refined $K=1$	refined $K=2$	refined $K=3$	coarse fit
0.0004	6.958×10^{-6} $(1.50\%; 22)$	1.129×10^{-7} $(3.07\%; 9)$	1.253×10^{-9} $(4.39\%; 22)$	1.603×10^{-3} $(0.56\%; 15)$
0.0007	2.129×10^{-5} $(1.49\%; 22)$	6.046×10^{-7} $(3.09\%; 9)$	1.171×10^{-8} $(4.39\%; 22)$	2.807×10^{-3} $(0.56\%; 15)$
0.001	4.341×10^{-5} $(1.48\%; 22)$	1.762×10^{-6} $(3.11\%; 9)$	4.859×10^{-8} $(4.40\%; 22)$	4.014×10^{-3} $(0.56\%; 15)$
0.0013	7.331×10^{-5} $(1.48\%; 22)$	3.867×10^{-6} $(3.13\%; 9)$	1.383×10^{-7} $(4.41\%; 22)$	5.222×10^{-3} $(0.56\%; 15)$

Deployment process floor and the longer horizon. If a newly prepared deployment instead has independent isotropic process innovations of standard deviation τ_{dep}, its exact conditional expected terminal squared norm is

$$\|c_3 + G_3\widetilde{U}_3\|^2 + \tau_{\mathrm{dep}}^2 \sum_{j=0}^{2} \|A^j\|_F^2.$$

Average the first term over learning records, then divide both terms by $\|A^3x_0\|^2$ for the normalized risk. The denominator and finite sum are specified for every plant by the panel definition; no deployment noise level was selected as a numerical result. The independent deployment noise is additional to the zero-floor learning model and does not change its data or report.

A twenty-step prepared-deployment diagnostic is also available for the original direct command report. Its passive terminal norm is only $2.301182\,10^{-5}$ to $2.949894\,10^{-4}$ on this stable panel, which motivates the shorter task above. At $\sigma = .001$, its normalized residuals are $1.323134\,10^{-4}$, $2.505096\,10^{-6}$ and $5.883497\,10^{-8}$ for $K = 1, 2, 3$, versus $.04539771$ for the coarse comparator. These longer-horizon numbers are posthoc diagnostics with separately measured deployment energy; the unit-energy deployment guarantee in this appendix concerns the guarded three-step task.

M A CONTINUOUS-CLASS CERTIFICATE FOR THE NUMERICAL ACQUISITION

This certificate concerns the full matrix class and selected four-input word in Appendix K. It applies to both numerical studies despite their different later controls. It certifies the retained model report and hard resource rules, not the continuous-class terminal risk of a numerical controller. Control-oriented and active input design address related choices of excitation (Hjalmarsson, 2005; Bombois et al., 2006; Wagenmaker & Jamieson, 2020; Chatzikiriakos et al., 2026).

Throughout the two operator balls, outward bounds are

$$\|A\|_{\mathrm{op}} \le a = .7613, \qquad \|A\|_F \le .973, \qquad s_{\min}(B) \ge .77.$$

Let $x_{t,0}$ denote the nominal noiseless prefix and $\bar{x}_t$ the true noiseless prefix. Nominal state norms are bounded by $(1, 10^{-12}, .6438, .3877, .9971)$, and command norms by $(.6886, .7501, 0, .7501)$. With these respective bounds b_t, c_t, define

$$d_0 = 0, \qquad d_{t+1} = ad_t + .05(b_t + c_t).$$

The dynamics give $\|\bar{x}_t - x_{t,0}\| \le d_t$, with $d_1 = .08443$, $d_2 = .10178155900005$ and $d_3 = .109676300866738065$. The nominal regressor has $s_{\min}(W_0) > .5201$; hence

$$s_{\min}(W_\theta) \ge .5201 - \sqrt{d_1^2 + d_2^2 + d_3^2} \ge .348 =: q. \tag{138}$$

The noiseless state envelopes at clocks one through four are $v = (.0845, .7456, .4974, 1.1375)$. These bounds hold for every dense perturbation in the class, rather than only the displayed plant panel. The four-input energy is less than 1.5991.

Keep the shared regression errors. Set $e_t = Y_t - x_t$, $E_- = [0, e_1, e_2, e_3]$, $E_+ = [e_1, e_2, e_3, e_4]$ and $\Xi = [\xi_0, \xi_1, \xi_2, \xi_3]$. On the full-rank branch,

$$\widehat{\theta}_{\mathrm{raw}} - \theta = (E_+ - AE_- + \tau\Xi)W_Y^{\dagger}. \tag{139}$$

Neighboring residual columns share the same reading error. Conditional on the prefix process vector, the reader innovations remain independent. If $\mathrm{Cov}(e_t)/\sigma^2 \preceq w_t^2 I_2$, the vectorized reader residual $R = (E_+ - AE_-)/\sigma$ has covariance trace at most

$$2\sum_{t=1}^{4} w_t^2 + .973^2 \sum_{t=1}^{3} w_t^2,$$

and covariance norm at most $\|H(w)\|^2$, where

$$H(w) = \begin{pmatrix} w_1 & 0 & 0 & 0 \\ aw_1 & w_2 & 0 & 0 \\ 0 & aw_2 & w_3 & 0 \\ 0 & 0 & aw_3 & w_4 \end{pmatrix}.$$

For E_-/σ, the corresponding bounds are $2\sum_{t=1}^{3} w_t^2$ and $\max_{t\le 3} w_t^2$. The block operator is norm-dominated by $H(w)$, so these estimates retain the residual correlations. They also cover singular coordinatewise reader covariance.

A centered Gaussian vector with covariance trace s and covariance norm λ satisfies

$$\mathbb{P}\{\|g\|^2 > s + 2\sqrt{s\lambda t} + 2\lambda t\} \le e^{-t}.$$

The following outward constants give the needed events.

Domain	Vector	s	λ	t	Norm bound
$d = \tau = 0$	R	5	1.49	2.313	4.51
	E_-/σ	1.625	.557	6.908	3.80
Both floors	$R \mid \Xi$	5.13	1.523	2.324	4.56
	$E_-/\sigma \mid \Xi$	1.68	.572	6.908	3.85
Process	Ξ	8	1	6.908	6.06

For the zero-floor rows use $w = v$. For the floor rows, take $d \le .01$, $\tau \le 10^{-5}$ and first restrict to $\|\Xi\|_F \le J_* = 6.06$. Since

$$\sqrt{\sum_{j=0}^{t-1} a^{2j}} \le 1.453 \quad (t \le 4), \qquad S = 1 + a + a^2 = 2.34087769,$$

the process perturbations satisfy $\|x_t - \bar{x}_t\| \le \tau(1.453)J_*$ and $\|X_- - \bar{X}_-\|_F \le \tau S J_*$. Thus $w_t = v_t + .0100880518$ suffices uniformly on the process event. The displayed spectral bounds can be checked by positive LDL pivots of $\lambda I - H^\top H$; the scalar threshold checks reduce to $(Q^2 - s - 2\lambda t)^2 \ge 4s\lambda t$ and $Q^2 - s - 2\lambda t \ge 0$. All constants were rounded outwards.

With the applicable residual and regressor thresholds R_*, E_*, the joint event gives

$$\|\widehat{\theta}_{\mathrm{rep}} - \theta\|_F \le \frac{\sigma R_* + \tau J_*}{q - \sigma E_* - \tau S J_*}. \tag{140}$$

For zero floors, the union failure is less than $.099 + .001 = .1$; with both floors it is less than $.098 + .001 + .001 = .1$. No independence between these events is needed. The denominator is also a lower bound on $s_{\min}(W_Y)$ and exceeds .3428 on the whole stated rectangle. Thus the actual fit threshold $q_{\mathrm{design}}/2 < .1742$ cannot trigger its fallback on this event. Frobenius projection onto the two convex operator balls cannot increase the error.

At $\sigma = .001, d = \tau = 0$, Eq. (140) is at most $4.51\sigma/(.348 - 3.80\sigma) < .013103$. Uniformly for $\sigma \le .0013, d \le .01, \tau \le 10^{-5}$, it is less than .017467. Both are below the numerical tolerance .025, so the retained report has failure probability less than .1 on the entire continuous class. Later adaptive inputs cannot change the law of this already completed acquisition.

The late-update study enforces its hard budget by projecting every feedback command onto the remaining energy ball. Before its first packet, the nominal nine-input word has energy $2.75 < 2.76$; after that packet, every remaining clock is a feedback clock. Therefore no future fixed command can violate a previously allocated cap. The precision study instead uses its separate geometric resource bound in Appendix K. Neither hard-budget argument conditions on report success or a bounded-noise event.

N Learning-time noise floors for matched late-packet controllers

This separate calculation holds $T = 10$, $E = 2.76$ and the report accuracy requirement $(\epsilon, \delta) = (.025, .1)$ fixed. It tests six explicit late-packet algorithms sharing the acquisition design. These controllers differ from the command-refinement policy in Appendix P; their floor experiments assess learning-terminal control. Use the full matrix class, all 25 panel plants, and the first four commands in Appendix K.1. The report remains the projected first-four fit on every trajectory. Its fallback threshold here is $.34838610645518703/2$. The analytical common-prefix certificate applies to the entire class for both readers, $\sigma \le .0013$, $d \le .01$ and $\tau \le 10^{-5}$; sampled report success is only a diagnostic.

N.1 Fixed plan, policy families and causal fits

The common precomputed plan consists of the four rows of (136), followed by

$$\begin{pmatrix} u_4^\top \\ u_5^\top \\ u_6^\top \\ u_7^\top \\ u_8^\top \end{pmatrix} = \begin{pmatrix} -.6769727153232308 & -.4969764169238549 \\ -.3342940006883367 & -.21115139561696636 \\ .12581399794622194 & .08323492545327249 \\ .41934450407988855 & .21406274943532166 \\ -.19645614861723287 & -.0792465093372276 \end{pmatrix}. \tag{141}$$

The full nine-input plan has energy 2.75. Its last five commands were chosen before noisy validation by minimizing $\sum_{p=0}^{24} \|A_p^5 \bar{x}_{4,p} + [A_p^4 B_p, \dots, B_p](u_4, \dots, u_8)\|^2$ subject to total plan energy at most 2.75. This noiseless optimization and the 36-candidate acquisition screen used the known design panel. The resulting coefficients were fixed before noisy evaluation; the literal values above fully specify the evaluated policy.

For each of two fit families, evaluate $K = 1, 2, 3$ with packets at $10 - K, \dots, 9$. Before the first packet execute the fixed plan. At every packet clock t, replace the planned command by

$$u_t = \Pi_{\{u:\|u\|^2 \le R_t\}}(-\widehat{B}_t^{-1}\widehat{A}_t Y_t), \qquad R_t = \max\left\{0, 2.76 - \sum_{s<t} \|u_s\|^2\right\}.$$

The retained family, denoted R_K, always uses the first-four fit. The causal-refit family, denoted C_K, recomputes at each packet from

$$W_t = \begin{bmatrix} Y_0 & \cdots & Y_{t-1} \\ u_0 & \cdots & u_{t-1} \end{bmatrix}, \qquad [Y_1, \dots, Y_t] W_t^\dagger.$$

Here the u_s are the actual executed inputs of that policy, including previous corrections and clipping. Use the same threshold as above, nominal fallback, and separate operator-ball projections for A, B. The projected B is invertible on every record. No current command, future reading or true parameter enters this fit. All six policies carry the same retained first-four report in their first packet; control refitting does not alter that report. The hard-energy projection includes all tails. The underlying at-most-K policy classes are nested, but these are forced-update algorithms, whose individual risks need not be monotone in K.

N.2 Readers, floors and paired Monte Carlo simulation

The state recursion is $x_{t+1} = Ax_t + Bu_t + \tau\xi_t$. For radial sensing use $Y_t = x_t + \sigma\sqrt{d^2 + \|x_t\|^2} Z_t$; for coordinatewise sensing its ith coordinate is $x_{t,i} + \sigma\sqrt{d^2 + x_{t,i}^2} Z_{t,i}$. All primitive Gaussian

coordinates are independent before pairing; $Y_0 = x_0$ is exact. Configuration IDs used in the tables are

ID	σ	d	τ	ID	σ	d	τ
0	.0001	0	0	7	.001	.002	0
1	.0002	0	0	8	.001	.01	0
2	.0004	0	0	9	.001	0	10^{-8}
3	.0008	0	0	10	.001	0	10^{-6}
4	.001	0	0	11	.001	0	10^{-5}
5	.0013	0	0	12	.001	.002	10^{-6}
6	.001	.0002	0	13	.001	.01	10^{-5}

Four independent blocks use PCG64 seeds $47102000 + b$, $b = 0, 1, 2, 3$. Each draws 256 arrays of shape 10×2 for ξ, then 256 such arrays for Z, and appends the simultaneous negatives. There are 512 signed paths per block and 1024 independent antithetic pairs per cell. Primitives are paired across all plants, readers, configurations and policies. Subsequent observations belong to each policy's own controlled trajectory. Only the common prefix observations are shared as observations. Z_0 is unused and no Y_{10} is supplied.

Physical innovations $\xi_0, \ldots, \xi_8$ are generated chronologically. All policies last act at 9, so only ξ_9 is integrated analytically: its conditional terminal MSE is $\|Ax_9 + Bu_9\|^2 + 2\tau^2$. Earlier innovations remain in the controlled histories; they are not treated as independent of later commands. All report-failure, fallback, projection and clipping records remain in the expectation. The complete output contains $2 \cdot 14 \cdot 25 = 700$ plant/configuration rows with four blocks and all six policies each.

For each plant and policy, take the square root of the mean of its four block MSEs, then maximize over the 25 plants. Attach the maximizing plant's own block SE, calculated as $\mathrm{sd}(m_{p1}, \ldots, m_{p4})/(\sqrt{4}\, 2r_p)$. Fit each log-log slope over precisely IDs 0–5. The SEs and slopes describe Monte Carlo variability and finite-range behavior under the supplement-wide statistical conventions.

N.3 Results and implementation checks

The maximum sampled energy over all policies is 2.751866703. Input clips, fit fallbacks and report failures number zero in the simulation; operator-ball projections at boundary plants are included. The largest sampled report error is $.008136001$, the smallest common-prefix regressor singular value is $.478050211$, and all fitted B singular values are at least $.779466847$. Source and chronological checks covered 552 scalar histories and 4416 comparisons, including 24 forced clips and six forced fallbacks. A separate aggregation recomputed every cell, selected plant, SE and slope across all 700 rows and checked their correspondence to the specified source and protocol.

Tables below retain every declared configuration and policy. In this particular panel no successive forced-depth maximum increases. At positive floors the differences can become very small; this is not evidence that an extra forced correction is always beneficial. For example, the shared-prefix design and late timing already make the one-packet curve far smaller than the older fixed-time-four baseline. Thus a comparison with that older acquisition would mix acquisition design and feedback opportunity. The present six curves share both acquisition and the precomputed plan, while still representing specific policies rather than the optimal risks of the nested classes.

Table conventions. As above, $a(b) = a \times 10^b$; beneath each RMS is $(100\,\mathrm{SE}/\mathrm{RMS}\%; p_*)$. Columns R_K use the retained fit and C_K use the causal refit. Configuration IDs are defined above. All six values, their selected-plant SEs and their maximizing plants are reported, including every floor configuration.

Table 12: Matched late-packet terminal RMS: radial sensing. Four independent blocks, 25 plants, $T = 10, E = 2.76$.

ID	R_1	R_2	R_3	C_1	C_2	C_3
	$3.989(-7)$	$1.048(-8)$	$2.100(-12)$	$4.035(-7)$	$4.729(-9)$	$9.534(-13)$
0	$(1.26\%; 10)$	$(1.90\%; 10)$	$(9.02\%; 10)$	$(1.21\%; 10)$	$(1.82\%; 13)$	$(2.32\%; 13)$
	$7.983(-7)$	$4.191(-8)$	$1.677(-11)$	$8.072(-7)$	$1.892(-8)$	$7.629(-12)$
1	$(1.26\%; 10)$	$(1.90\%; 10)$	$(9.03\%; 10)$	$(1.21\%; 10)$	$(1.82\%; 13)$	$(2.33\%; 13)$
	$1.599(-6)$	$1.674(-7)$	$1.337(-10)$	$1.615(-6)$	$7.570(-8)$	$6.107(-11)$
2	$(1.27\%; 10)$	$(1.88\%; 10)$	$(9.04\%; 10)$	$(1.21\%; 10)$	$(1.82\%; 13)$	$(2.36\%; 13)$
	$3.208(-6)$	$6.675(-7)$	$1.063(-9)$	$3.233(-6)$	$3.029(-7)$	$4.890(-10)$
3	$(1.28\%; 10)$	$(1.85\%; 10)$	$(9.05\%; 10)$	$(1.21\%; 10)$	$(1.84\%; 13)$	$(2.41\%; 13)$
	$4.016(-6)$	$1.041(-6)$	$2.068(-9)$	$4.043(-6)$	$4.734(-7)$	$9.556(-10)$
4	$(1.29\%; 10)$	$(1.83\%; 10)$	$(9.03\%; 10)$	$(1.21\%; 10)$	$(1.84\%; 13)$	$(2.43\%; 13)$
	$5.232(-6)$	$1.756(-6)$	$4.519(-9)$	$5.258(-6)$	$8.004(-7)$	$2.101(-9)$
5	$(1.30\%; 10)$	$(1.80\%; 10)$	$(8.99\%; 10)$	$(1.20\%; 10)$	$(1.85\%; 13)$	$(2.47\%; 13)$
	$4.020(-6)$	$1.058(-6)$	$1.880(-7)$	$4.047(-6)$	$5.063(-7)$	$1.880(-7)$
6	$(1.29\%; 10)$	$(1.82\%; 10)$	$(1.89\%; 7)$	$(1.21\%; 10)$	$(1.79\%; 13)$	$(1.90\%; 7)$
	$4.424(-6)$	$2.125(-6)$	$1.879(-6)$	$4.449(-6)$	$1.930(-6)$	$1.880(-6)$
7	$(1.36\%; 10)$	$(1.88\%; 10)$	$(1.90\%; 7)$	$(1.28\%; 10)$	$(1.77\%; 7)$	$(1.90\%; 7)$
	$1.010(-5)$	$9.448(-6)$	$9.397(-6)$	$1.011(-5)$	$9.408(-6)$	$9.398(-6)$
8	$(1.75\%; 10)$	$(1.85\%; 7)$	$(1.90\%; 7)$	$(1.67\%; 10)$	$(1.87\%; 7)$	$(1.90\%; 7)$
	$4.016(-6)$	$1.042(-6)$	$1.429(-8)$	$4.043(-6)$	$4.736(-7)$	$1.417(-8)$
9	$(1.29\%; 10)$	$(1.83\%; 10)$	$(0.19\%; 10)$	$(1.21\%; 10)$	$(1.84\%; 13)$	$(0.01\%; 13)$
	$4.257(-6)$	$1.756(-6)$	$1.414(-6)$	$4.283(-6)$	$1.491(-6)$	$1.414(-6)$
10	$(1.15\%; 10)$	$(0.64\%; 10)$	$(0.00\%; 10)$	$(1.07\%; 10)$	$(0.19\%; 13)$	$(0.00\%; 13)$
	$1.470(-5)$	$1.418(-5)$	$1.414(-5)$	$1.471(-5)$	$1.415(-5)$	$1.414(-5)$
11	$(0.10\%; 10)$	$(0.01\%; 10)$	$(0.00\%; 15)$	$(0.09\%; 10)$	$(0.00\%; 13)$	$(0.00\%; 21)$
	$4.645(-6)$	$2.553(-6)$	$2.352(-6)$	$4.668(-6)$	$2.393(-6)$	$2.352(-6)$
12	$(1.23\%; 10)$	$(1.30\%; 10)$	$(1.21\%; 7)$	$(1.16\%; 10)$	$(1.15\%; 7)$	$(1.21\%; 7)$
	$1.738(-5)$	$1.701(-5)$	$1.698(-5)$	$1.738(-5)$	$1.699(-5)$	$1.698(-5)$
13	$(0.59\%; 10)$	$(0.57\%; 7)$	$(0.58\%; 7)$	$(0.56\%; 10)$	$(0.57\%; 7)$	$(0.58\%; 7)$
Slope	1.00339	1.99659	2.99229	1.00090	2.00052	3.00111

Table 13: Matched late-packet terminal RMS: coordinatewise sensing. Four independent blocks, 25 plants, $T = 10$, $E = 2.76$.

ID	R_1	R_2	R_3	C_1	C_2	C_3
	$2.855(-7)$	$7.078(-9)$	$1.436(-12)$	$2.989(-7)$	$2.328(-9)$	$2.916(-13)$
0	$(2.13\%; 10)$	$(4.04\%; 10)$	$(9.78\%; 10)$	$(1.90\%; 10)$	$(0.76\%; 21)$	$(4.04\%; 10)$
	$5.713(-7)$	$2.829(-8)$	$1.147(-11)$	$5.980(-7)$	$9.312(-9)$	$2.334(-12)$
1	$(2.13\%; 10)$	$(4.03\%; 10)$	$(9.77\%; 10)$	$(1.91\%; 10)$	$(0.76\%; 21)$	$(4.03\%; 10)$
	$1.144(-6)$	$1.130(-7)$	$9.157(-11)$	$1.196(-6)$	$3.727(-8)$	$1.868(-11)$
2	$(2.14\%; 10)$	$(4.03\%; 10)$	$(9.77\%; 10)$	$(1.91\%; 10)$	$(0.76\%; 21)$	$(4.03\%; 10)$
	$2.293(-6)$	$4.510(-7)$	$7.291(-10)$	$2.393(-6)$	$1.492(-7)$	$1.497(-10)$
3	$(2.14\%; 10)$	$(4.03\%; 10)$	$(9.75\%; 10)$	$(1.91\%; 10)$	$(0.76\%; 21)$	$(4.03\%; 10)$
	$2.869(-6)$	$7.037(-7)$	$1.421(-9)$	$2.992(-6)$	$2.332(-7)$	$2.925(-10)$
4	$(2.14\%; 10)$	$(4.03\%; 10)$	$(9.75\%; 10)$	$(1.91\%; 10)$	$(0.77\%; 21)$	$(4.03\%; 10)$
	$3.735(-6)$	$1.187(-6)$	$3.109(-9)$	$3.891(-6)$	$3.945(-7)$	$6.433(-10)$
5	$(2.14\%; 10)$	$(4.04\%; 10)$	$(9.74\%; 10)$	$(1.91\%; 10)$	$(0.77\%; 21)$	$(4.03\%; 10)$
	$2.875(-6)$	$7.281(-7)$	$1.880(-7)$	$2.998(-6)$	$2.978(-7)$	$1.880(-7)$
6	$(2.14\%; 10)$	$(3.87\%; 10)$	$(1.90\%; 7)$	$(1.91\%; 10)$	$(0.73\%; 21)$	$(1.90\%; 7)$
	$3.416(-6)$	$1.990(-6)$	$1.880(-6)$	$3.521(-6)$	$1.894(-6)$	$1.880(-6)$
7	$(1.98\%; 10)$	$(2.01\%; 7)$	$(1.90\%; 7)$	$(1.88\%; 10)$	$(1.85\%; 7)$	$(1.90\%; 7)$
	$9.694(-6)$	$9.423(-6)$	$9.398(-6)$	$9.732(-6)$	$9.402(-6)$	$9.398(-6)$
8	$(1.85\%; 10)$	$(1.89\%; 7)$	$(1.90\%; 7)$	$(1.82\%; 10)$	$(1.89\%; 7)$	$(1.90\%; 7)$
	$2.869(-6)$	$7.039(-7)$	$1.421(-8)$	$2.993(-6)$	$2.337(-7)$	$1.415(-8)$
9	$(2.14\%; 10)$	$(4.03\%; 10)$	$(0.10\%; 10)$	$(1.91\%; 10)$	$(0.76\%; 21)$	$(0.00\%; 10)$
	$3.198(-6)$	$1.580(-6)$	$1.414(-6)$	$3.310(-6)$	$1.433(-6)$	$1.414(-6)$
10	$(1.72\%; 10)$	$(0.80\%; 10)$	$(0.00\%; 10)$	$(1.56\%; 10)$	$(0.02\%; 21)$	$(0.00\%; 21)$
	$1.443(-5)$	$1.416(-5)$	$1.414(-5)$	$1.446(-5)$	$1.414(-5)$	$1.414(-5)$
11	$(0.08\%; 10)$	$(0.01\%; 10)$	$(0.00\%; 15)$	$(0.08\%; 10)$	$(0.00\%; 21)$	$(0.00\%; 21)$
	$3.697(-6)$	$2.442(-6)$	$2.352(-6)$	$3.794(-6)$	$2.364(-6)$	$2.352(-6)$
12	$(1.69\%; 10)$	$(1.33\%; 7)$	$(1.21\%; 7)$	$(1.62\%; 10)$	$(1.19\%; 7)$	$(1.21\%; 7)$
	$1.715(-5)$	$1.699(-5)$	$1.698(-5)$	$1.717(-5)$	$1.698(-5)$	$1.698(-5)$
13	$(0.59\%; 10)$	$(0.58\%; 7)$	$(0.58\%; 7)$	$(0.59\%; 10)$	$(0.58\%; 7)$	$(0.58\%; 7)$
Slope	1.00238	1.99711	2.99455	1.00051	2.00107	3.00152

O A BOUNDED UNDERACTUATED PRECISION CALCULATION

This calculation tests the command-refinement mechanism with one input and two states. Antithetic Monte Carlo simulation evaluates the stated Gaussian model on a fixed finite panel. The class, schedule and amplitudes were fixed before evaluation; the supplement-wide statistical conventions apply.

Class and geometric bounds. Set $x_0 = (1, 0)^\top$ and

$$A_0 = \begin{pmatrix} .94 & .4 \\ -.15 & .90 \end{pmatrix}, \qquad B_0 = \begin{pmatrix} .1 \\ .9 \end{pmatrix}, \qquad \|A - A_0\|_{\mathrm{op}},\ \|B - B_0\|_2 \le .015.$$

All six entries are unknown. Use radial readouts $Y_t = x_t + \sigma\|x_t\|Z_t$, without process noise or a reader floor. The scalar prefix $(0, 0, 1, 0, 0)$ has length five and energy one. Let x_t^0 be its nominal states, $a = \|A_0\|_{\mathrm{op}} + .015$, and $e_0 = 0$, $e_{t+1} = ae_t + .015\|x_t^0\| + .015|u_t|$. Weyl's inequality gives a uniform prefix margin $q_{\mathrm{PE}} \ge .4980971 - (\sum_{t=0}^{4} e_t^2)^{1/2} > .3895$; the prefix endpoint norm is below 1.139956. For $C = [AB, B]$, the perturbation bound

$$\|C - C_0\|_{\mathrm{op}} \le \sqrt{(.015\|B_0\| + .015\|A_0\|_{\mathrm{op}} + .015^2)^2 + .015^2} < .033640$$

gives $s_{\min}(C) > .2273$. The inverse perturbation identity around C_0 gives $\|C^{-1}A^2\| < 4.5$ throughout the class.

Here $h = 2$, $N = I$, $B_c = e_2$, and $A_c = \left(\begin{smallmatrix} p & 1 \\ q & 0 \end{smallmatrix}\right)$, where $p = \operatorname{tr} A$, $q = -\det A$, and $M = A_c^2$. The bounds $p \in [1.81, 1.87]$ and $q \in [-.935031640, -.876968360]$ follow from $|\det A - \det A_0| \le .015\|A_0\|_* + .015^2$. Project a raw M by clipping its second column to these intervals (enlarged by 10^{-12} in computation), then reconstructing A_c, M. This Lipschitz map fixes every true pair and always satisfies $[A_c e_2, e_2] = I$, even when its rectangle contains points outside the physical class. Separate convexity gives $\|M\|_F^2 < 14$ on the rectangle.

Fixed causal policy. At clock five, fit $[A\ B] = Y_{1:5}W^\dagger$, with $W = [Y_{0:4}; u_{0:4}]$, using the nominal pair if $s_{\min}(W) < .3895/2$. Project the two components to the declared matrix and vector balls. Retain this fit for the full-model report and for $\bar{C} = [\bar{A}\bar{B}, \bar{B}]$, $\bar{F} = \bar{A}^2$. Initialize M from this coarse pair. Each packet is followed by the two-input correction $-\bar{C}^{-1}\bar{F}\Pi_{R_r}Y$, where $R_1 = 1.25$, $R_r = 2^{-(r-1)}$ subsequently. At batch r, freeze M_{r-1} and, for each $j = 1, 2$, execute $a_r e_j, -a_r M_{r-1} e_j$, with $a_r = 2^{-r}$. Each word contains two scalar inputs. The raw column update is

$$M_{r-1}e_j + a_r^{-1}\bar{C}^{-1}(Y_{\text{end}} - \bar{F}^2 Y_{\text{start}}).$$

Apply the rectangle map only after both columns. No controller input uses the true plant. Every numerical record obeys the analytic energy bound

$$E_{\text{learn}} \le 1 + 4.5^2(1.25^2 + 1/3) + (2 + 14)/3 < 44.724 < 50. \tag{142}$$

For $K = 1, 2, 3$, packets occur at the first K clocks in $(5, 15, 25)$. Cache the last quiet reading at $s = 7, 17, 27$, respectively. The first $K - 1$ batches run immediately; coast with zero inputs after the last correction until clock 27, then run the final private batch to 35 and set $u_{35} = 0$. All policies have $T = H = 36$, energy cap 50, and the same 35 available noisy readings; Y_{36} is never used. Using the final canonical estimate, form

$$\widehat{c}_{36} = \widehat{A}_c^{36-s}\bar{C}^{-1}Y_s - \sum_{t<s}\widehat{A}_c^{35-t}e_2 u_t, \qquad \widehat{U} = -\mathcal{C}_{36}(\widehat{A}_c, e_2)^\dagger \widehat{c}_{36}.$$

The sum uses the actual past inputs. The retained physical fit supplies a separate plug-in command comparator. Both are offline reports; their execution would require additional time and energy.

Panel, Monte Carlo simulation, and outcomes. Use the center and the four predeclared plants $(A_0 + .015 s_A I, B_0 + .015 s_B e_1)$ for $s_A, s_B \in \{-1, 1\}$. Four independent PCG64 seeds $47124000 + b$, $b = 0, \ldots, 3$, each generate 128 arrays of 35 two-dimensional standard normals and their negatives. Within a block the primitives are shared across all plants, noise levels and depths. All $5 \times 4 \times 3 \times 4 = 240$ cells of 256 signed paths are retained. For each plant, average block MSEs, take the square root, then maximize over the panel. Table 14 reports this order of aggregation.

At $\sigma = .001$, the four block relative-command RMS values at the maximizing plant range from .0022341 to .0032834 for $K = 1$, from $7.2775\,10^{-5}$ to $9.9228\,10^{-5}$ for $K = 2$, and from $4.4957\,10^{-6}$ to $9.0991\,10^{-6}$ for $K = 3$. These ranges are descriptive, not simultaneous confidence intervals. The four-level command-error slopes for the three depths are 1.9967, 2.8429 and 3.7150; terminal-error slopes are .9747, 1.9112 and 2.8774. The full-fit slope is 1.0000. Finite-range departures from integer exponents are retained. There were 45,995 rectangle-clip events across all batches, with no prefix fallback or correction clip. Maximum observed learning energy was 22.9642. The target norms $\|U_{36}^*\| \in [.03466, .07369]$ and $\|A^{36}x_0\| \in [.09775, .28143]$ state the absolute task scale.

All 15 noiseless plant/depth checks had relative command error below $5.3\,10^{-14}$. Extreme records exercised fallback and all clipping branches while respecting Eq. (142). A separate scalar least-squares replay checked 30 signed records, and direct regrouping checked all 240 cells. This example demonstrates underactuation with $mh = n$ and $N = I$, so the nonselected-column queries required when $mh > n$ are absent. Singular drift matrices, chart switching and noise floors remain untested here; the fixed policy is not an optimized finite-sample benchmark.

Table 14: Five-plant maxima. Command RMS is normalized by the true $\|U_{36}^*\|$ separately for each plant. The fit comparator is one retained-fit plug-in policy, not an optimal command estimator.

σ	K	Terminal RMS	Full-fit RMS	Relative command RMS	Fit comparator
.0004	1	$1.1750\,10^{-3}$	.0013190	$4.5358\,10^{-4}$	.013835
	2	$1.5603\,10^{-5}$	.0013190	$5.9765\,10^{-6}$	.013835
	3	$1.0300\,10^{-7}$	.0013190	$2.2572\,10^{-7}$	.013835
.0007	1	$2.0356\,10^{-3}$	.0023082	$1.3897\,10^{-3}$	.023912
	2	$4.6165\,10^{-5}$	.0023082	$3.0139\,10^{-5}$	.023912
	3	$5.2646\,10^{-7}$	.0023082	$1.8937\,10^{-6}$	.023912
.0010	1	$2.8787\,10^{-3}$	.0032974	$2.8321\,10^{-3}$	.033732
	2	$9.0911\,10^{-5}$	.0032974	$8.2462\,10^{-5}$	.033732
	3	$1.4609\,10^{-6}$	.0032974	$7.0264\,10^{-6}$	.033732
.0013	1	$3.7043\,10^{-3}$	.0042867	$4.7696\,10^{-3}$	.043300
	2	$1.4807\,10^{-4}$	.0042867	$1.6978\,10^{-4}$	.043300
	3	$3.0497\,10^{-6}$	.0042867	$1.7872\,10^{-5}$	.043300

P HELD-OUT NOISE-ORDER EXPERIMENT: POLICY, DEPLOYMENT AND UNCERTAINTY

This experiment compares the retained physical fit, learning-terminal state, offline command and separately prepared three-step deployment across noise levels. The policy and reporting rules are specified below. The 25-plant design panel was used to select the acquisition prefix; Appendix K.1 records that exploratory screen. The held-out panel distribution, seeds, noise levels, four-block budget and aggregation rules were fixed before drawing this panel, and its outcomes did not enter policy selection. The deployment guard was first introduced in the posthoc design-panel study of Appendix L and then fixed for this evaluation.

Fixed policy and resources. Learning uses twenty transitions, at most ten units of input energy, and the indicated number of actionable packets. The eight-parameter class has center A_0, B_0, radius .05, and exact preparation $x_0 = e_1$:

$$A_0 = \begin{pmatrix} .65 & .22 \\ -.08 & .58 \end{pmatrix}, \qquad B_0 = \begin{pmatrix} 1 & .15 \\ .08 & .9 \end{pmatrix}, \qquad \|A - A_0\|_{\mathrm{op}},\ \|B - B_0\|_{\mathrm{op}} \le .05.$$

All entries are unknown. The four inputs, selected from the prior 36-candidate design-panel screen, are the rows of the fixed matrix

$$P = \begin{pmatrix} -.6722972972972974 & .14864864864864866 \\ -.1855529694408922 & .7266843162829836 \\ 0 & 0 \\ .7266843162829836 & .1855529694408922 \end{pmatrix}.$$

At clock four, form $Y_- = [Y_0, \dots, Y_3]$, $Y_+ = [Y_1, \dots, Y_4]$ and $W = [Y_-; P^\top]$. Use $Y_+ W^\dagger$ when its regressor singular value is at least .174, and $[A_0, B_0]$ otherwise. Project each fitted block's offset from its class center by clipping the offset singular values at .05; retain this pair $(\widehat{A}, \widehat{B})$ throughout learning. The full-model metric is $F^2 = \|\widehat{A} - A\|_F^2 + \|\widehat{B} - B\|_F^2$. Write $M = B^{-1}AB$, $z_0 = B^{-1}x_0$, $M_0 = B_0^{-1}A_0B_0$ and initialize $M^{(0)} = \widehat{B}^{-1}\widehat{A}\widehat{B}$. Project each completed refinement onto the .2 operator ball for M centered at M_0, again by clipping the offset singular values. The schedules are

K	packet clocks	batch starts	quiet clock s
1	4	15	5
2	4, 9	5, 15	10
3	4, 9, 14	5, 10, 15	15

At packet number p, send the retained fit, current reading and any completed matrix update. Execute $-\widehat{B}^{-1}\widehat{A}Y_t$, clipped to radius two for $p = 1$ and 2^{1-p} thereafter. Each batch r uses amplitude 2^{-r}

and a frozen incoming matrix; for each $j = 1, 2$, execute $a_r e_j, -a_r M_{\mathrm{old}} e_j$ with $a_r = 2^{-r}$. Only the two endpoint innovations update its two columns:

$$M_{\mathrm{new}} e_j = M_{\mathrm{old}} e_j + a_r^{-1} \widehat{B}^{-1} (Y_{t+2} - \widehat{A}^2 Y_t).$$

Project after both columns are available. Unscheduled inputs are zero. The last batch is private, $u_{19} = 0$, and each $Y_1, \ldots, Y_{19}$ is generated once and cached; reuse means the same reading. Sensing is radial with $\beta = 1, d = \tau = 0$: $x_{t+1} = Ax_t + Bu_t$, $Y_t = x_t + \sigma \|x_t\| Z_t$ for independent standard two-dimensional Gaussian Z_t, with exact $Y_0 = x_0$ and no terminal reading Y_{20}.

At the stored quiet clock s, set $\widehat{z}_s = \widehat{B}^{-1} Y_s$. The final matrix $M^\sharp$ gives the quiet-state reconstruction from Eq. (137), explicitly

$$z_0^\sharp = (M^\sharp)^{-s} \left[\widehat{z}_s - \sum_{t<s} (M^\sharp)^{s-1-t} u_t \right], \qquad \widehat{U}_3 = -[(M^\sharp)^2, M^\sharp, I]^\dagger (M^\sharp)^3 z_0^\sharp.$$

The sum uses actual executed inputs, including clipped corrections. The comparison command uses the same retained coarse fit, replacing $(M^\sharp, z_0^\sharp)$ by $(M^{(0)}, \widehat{B}^{-1} x_0)$. Both three-step reports receive the same Euclidean unit-ball guard. A final upload supplies the complete private record to the reporter after learning, in addition to the actionable packets. These are offline reports; the learning terminal metric is $\|x_{20}\|$. The matrix ball has inverse bound 2.892181, supporting this fixed short-horizon reconstruction. Every record obeys the learning bound

$$\|P\|_F^2 + 4 + \tfrac{1}{3} + \tfrac{2}{3}\{1 + (\|M_0\|_{\mathrm{op}} + .2)^2\} = 7.167068015 < 10.$$

The added deployment has energy at most one and starts from a separate exact preparation. Thus the accounting is 23 transitions, input-energy cap 11 and two preparations; preparation cost and offline word delivery are additional resources. Learning uses one uninterrupted trajectory.

Held-out plants and antithetic Monte Carlo simulation. Use NumPy PCG64 with plant seed 48100101. For each plant $p = 0, \ldots, 63$, draw its A offset and then its B offset. For each, first draw $U \sim \mathrm{Unif}[0, 1)$, then four independent standard normal coordinates in a row-major 2×2 matrix G. Set

$$D = .05\, U^{1/4} G / \|G\|_{\mathrm{op}}, \qquad A = A_0 + D_A, \quad B = B_0 + D_B.$$

These offsets are independent. This rule defines the test distribution: the radius scales as $U^{1/4}$ and the direction is a Gaussian matrix normalized in operator norm. All 64 plants enter the evaluation. The noise grid is

$$\sigma \in \{.0002, .0003, .0005, .0008, .0013, .002, .0032, .005, .008, .013, .02, .032, .05\}.$$

Four independent blocks $b = 0, 1, 2, 3$ use PCG64 seeds $48110000 + b$. Each draws a standard normal array of shape $64 \times 20 \times 2$, then appends its negative, giving 128 signed paths (64 antithetic pairs). Clock-zero primitives are unused. The same block primitives pair all plants, levels and depths, while every trajectory uses its own physical states. All $64 \cdot 13 \cdot 3 \cdot 4 = 9984$ block cells are retained. The independent numerical-uncertainty units are the four blocks, not the signed paths or repeated coarse fits across K.

Risks, normalization and uncertainty. The joint task-coordinate error has $Q^2 = \|M^\sharp - M\|_F^2 + \|z_0^\sharp - z_0\|^2$, and the raw command error is $\|\widehat{U}_3 - U_3^*\|$. For a metric with block MSE m_{pb}, define

$$r_p = \sqrt{\tfrac{1}{4} \sum_b m_{pb}}, \qquad R_{\mathrm{pool}} = \sqrt{\tfrac{1}{64} \sum_p r_p^2}, \qquad R_{\max} = \max_p r_p.$$

The primary table uses pooled RMS. The pooled delta-method block SE is $\mathrm{sd}_b(64^{-1} \sum_p m_{pb}) / (2\sqrt{4} R_{\mathrm{pool}})$. For $R_{\max}$ use the same formula at the selected plant. The supplement-wide statistical conventions apply: four blocks give limited Monte Carlo precision conditional on this panel, and selected maxima have no selection-adjusted coverage. Absolute deployment risk is the RMS of $A^3 x_0 + [A^2 B, AB, B] \widetilde{U}_3$. The normalized variant divides each plant's MSE by its own $\|A^3 x_0\|^2$ before pooling or maximizing. For evaluation, the physical oracle is $U_3^* = -[A^2 B, AB, B]^\dagger A^3 x_0$; the residual is also computed by three chronological state transitions

Table 15: Held-out panel: absolute pooled RMS. The model column is the common coarse fit; other columns are noiseless three-step deployment residuals from a separate preparation. All noise levels are retained.

σ	full model	coarse deployment	$K=1$	$K=2$	$K=3$
0.0002	$5.655\,10^{-4}$	$1.595\,10^{-4}$	$3.395\,10^{-7}$	$2.640\,10^{-9}$	$1.287\,10^{-11}$
0.0003	$8.481\,10^{-4}$	$2.392\,10^{-4}$	$7.635\,10^{-7}$	$8.904\,10^{-9}$	$6.473\,10^{-11}$
0.0005	$1.412\,10^{-3}$	$3.979\,10^{-4}$	$2.118\,10^{-6}$	$4.110\,10^{-8}$	$4.971\,10^{-10}$
0.0008	$2.255\,10^{-3}$	$6.358\,10^{-4}$	$5.413\,10^{-6}$	$1.676\,10^{-7}$	$3.252\,10^{-9}$
0.0013	$3.654\,10^{-3}$	$1.032\,10^{-3}$	$1.425\,10^{-5}$	$7.151\,10^{-7}$	$2.258\,10^{-8}$
0.002	$5.599\,10^{-3}$	$1.587\,10^{-3}$	$3.360\,10^{-5}$	$2.587\,10^{-6}$	$1.259\,10^{-7}$
0.0032	$8.896\,10^{-3}$	$2.537\,10^{-3}$	$8.562\,10^{-5}$	$1.050\,10^{-5}$	$8.152\,10^{-7}$
0.005	$1.373\,10^{-2}$	$3.962\,10^{-3}$	$2.078\,10^{-4}$	$3.939\,10^{-5}$	$4.696\,10^{-6}$
0.008	$2.143\,10^{-2}$	$6.325\,10^{-3}$	$5.255\,10^{-4}$	$1.559\,10^{-4}$	$2.909\,10^{-5}$
0.013	$3.315\,10^{-2}$	$1.023\,10^{-2}$	$1.356\,10^{-3}$	$6.175\,10^{-4}$	$1.826\,10^{-4}$
0.02	$4.689\,10^{-2}$	$1.551\,10^{-2}$	$3.081\,10^{-3}$	$1.909\,10^{-3}$	$8.130\,10^{-4}$
0.032	$6.374\,10^{-2}$	$2.317\,10^{-2}$	$7.216\,10^{-3}$	$5.747\,10^{-3}$	$3.384\,10^{-3}$
0.05	$7.874\,10^{-2}$	$3.077\,10^{-2}$	$1.533\,10^{-2}$	$1.438\,10^{-2}$	$1.084\,10^{-2}$

from the separate preparation. True parameters enter these evaluations only. The last-input-only zeroing word has norm at most $.7613^3/.77 < .574$, so the unit-ball guard contains the oracle. Its Euclidean command-error guarantee does not imply a smaller physical residual on each record.

The full-model report certificate at $(\epsilon, \delta) = (.025, .1)$ applies to the first five levels, $\sigma \le .0013$. Larger levels measure behavior beyond its certified regime. At .05, the refined deployment advantage over the coarse report has narrowed and full-report exceedances are nearly universal. The calculations at that level therefore describe performance outside the joint report-feasibility guarantee.

Finite-range slopes and events. Least-squares slopes of $\log R_{\mathrm{pool}}$ against $\log \sigma$, with an intercept, use the predeclared first five levels and last four levels, respectively. In the following table each entry is "first five / last four." The full-model slopes are $.997/.643$ and coarse-deployment slopes $.997/.821$. Here Q is the joint (M, z_0) error, X the executed learning terminal state, U the raw three-command error, and D guarded deployment.

Table 16: Prespecified small- and larger-noise slopes of pooled RMS (first-five / last-four levels).

K	Q	X	U	D
1	$1.996/1.806$	$0.997/0.726$	$1.997/1.818$	$1.996/1.801$
2	$2.991/2.261$	$1.994/1.593$	$2.992/2.314$	$2.992/2.337$
3	$4.001/2.991$	$2.989/2.302$	$3.993/3.013$	$3.992/3.031$

The event table counts prefix-fit projections and full-report errors once per level, since the fit is identical across K. Each denominator is $64 \cdot 4 \cdot 128 = 32768$ signed paths; the entries are not independent Bernoulli trials. Across all levels and depths, prefix fallbacks, correction clips, refined-M projections and deployment-guard clips are all zero. The nonzero A, B projection counts remain included. Maximum sampled learning energy is 3.257800; maximum deployment energy is $.123326$ for refinement and $.075639$ for the coarse report. The deterministic caps establish feasibility; these maxima summarize the energy used by the sampled trajectories.

Table 17: Retained prefix events per 32,768 signed paths, identical across K. An exceedance means $F > .025$, not a rejected trajectory.

σ	A projections	B projections	report exceedances	fraction
0.0002	392	257	0	0.00%
0.0003	481	274	0	0.00%
0.0005	685	391	0	0.00%
0.0008	1,032	612	0	0.00%
0.0013	1,630	1,064	0	0.00%
0.002	2,320	1,732	0	0.00%
0.0032	3,345	2,915	0	0.00%
0.005	4,814	4,761	547	1.67%
0.008	7,182	7,921	8,071	24.63%
0.013	10,860	12,917	21,912	66.87%
0.02	15,599	18,934	29,971	91.46%
0.032	22,245	25,817	32,209	98.29%
0.05	28,152	30,384	32,706	99.81%

Arithmetic and preparation sensitivity. A second chronological implementation reproduced selected complete blocks spanning the smallest, certified-upper and largest noise levels. At the smallest level, $K = 3$, an 80-digit computation of all 128 paths at plant 0/block 0 gives deployment RMS $1.320325388\,10^{-11}$, versus $1.320356903\,10^{-11}$ in binary64, a .00239% difference. Individual nearly cancelled residuals can have much larger relative error: among this block and two additional paths the maximum was 25.6%, with absolute difference at most $1.89\,10^{-14}$. The reported simulation values use the original arithmetic throughout. This replay measures rounding on the selected records, with no uniform arithmetic-error bound.

As a separately labeled posthoc analytic sensitivity, suppose deployment starts at $x_0 + \eta$, with η independent of learning, mean zero and covariance $\delta^2_{\mathrm{prep}} I/2$. The vector RMS preparation error is δ_{prep}. Independent deployment process noise has coordinate standard deviation τ_{dep}. Conditional on a plant, the exact extra terminal MSE is

$$\delta^2_{\mathrm{prep}} \frac{\|A^3\|_F^2}{2} + \tau^2_{\mathrm{dep}} \sum_{j=0}^{2} \|A^j\|_F^2.$$

The preparation coefficient has panel mean .07711943 and maximum .10533754; the process coefficient has mean 3.15980592 and maximum 3.31944768. Add the mean coefficients to the squared pooled base risk before taking its square root. For a panel maximum, add terms separately for each plant before maximizing; separate maxima only give an upper bound. Divide by each plant's passive terminal squared norm first for normalized risks. A biased or learning-dependent preparation can add a cross term and is outside this formula. This sensitivity is an analytic calculation without additional simulated trajectories; preparation accuracy remains a separate resource.

Q Learning-time floors for the command-refinement policy

This calculation applies learning-time noise floors directly to the command-refinement policy. It uses a fixed subset of an existing test panel, with the protocol fixed before the calculation.

Fixed policy and evaluation. Take plants 0–15 from the stored 64-plant panel of Appendix P, set $\sigma = .005$, and evaluate $K = 1, 2, 3$. Replace only the simulated physical model by

$$x_{t+1} = Ax_t + Bu_t + \tau\xi_t, \qquad Y_t = x_t + \sigma\sqrt{\|x_t\|^2 + d^2}\, Z_t, \qquad \xi_t, Z_t \sim N(0, I_2),$$

with independent process and reader primitives. Here $Y_0 = x_0$ is exact; the noisy readout formula applies at $t = 1, \ldots, 19$. The table varies the two floors separately. The twenty-transition policy, exact initial preparation, prefix, .174 fallback threshold, correction caps, spectral projections, frozen query batches, cached readings and quiet-state reconstruction are unchanged. In particular, reconstruction uses the actual input history and receives no latent process innovations. The final private batch

reaches clock 19, the last input is zero, and there is no terminal readout. A final upload supplies the offline report.

Four independent blocks use reader PCG64 seeds $53091000 + b$ and process seeds $53092000 + b$, $b = 0, 1, 2, 3$. Each stream draws an array of shape $16 \times 20 \times 2$ and appends its negative; the joint antithetic partner negates both arrays. Reader clock zero is unused and all twenty process innovations are applied. Each block bank is shared across plants, floor settings and depths. Every setting/depth includes 2,048 signed trajectories; the full calculation contains 43,008. All paths, including report exceedances and guard activations, enter the summaries.

The three metrics are the learning-terminal norm $\|x_{20}\|$, raw three-command error $\|\widehat{U}_3 - U_3^*\|$, and the terminal norm from a fresh exact preparation executing the unit-ball-guarded report in a noiseless three-step deployment. Within each block, squared losses are averaged equally over plants and signed paths. The displayed RMS is $R = \sqrt{m}$ for the four block MSEs, with Monte Carlo standard error $s_m/(2\sqrt{4}R)$. These errors condition on the selected panel and policy; the four blocks do not supply a population confidence claim.

Table 18: Direct learning-floor simulation at $\sigma = .005$. Entries are pooled RMS, followed by relative block SE in percent. Deployment is separately prepared and noiseless.

d	τ	K	Learning terminal	Raw command	Deployment
0	0	1	$1.775\,10^{-3}$ (2.0)	$2.006\,10^{-4}$ (6.4)	$2.156\,10^{-4}$ (5.9)
		2	$2.720\,10^{-5}$ (8.5)	$3.906\,10^{-5}$ (12.4)	$4.259\,10^{-5}$ (11.7)
		3	$2.723\,10^{-7}$ (19.4)	$4.498\,10^{-6}$ (15.9)	$5.260\,10^{-6}$ (16.4)
10^{-5}	0	1	$1.775\,10^{-3}$ (2.0)	$2.006\,10^{-4}$ (6.4)	$2.156\,10^{-4}$ (5.9)
		2	$2.720\,10^{-5}$ (8.5)	$3.943\,10^{-5}$ (12.2)	$4.300\,10^{-5}$ (11.5)
		3	$2.732\,10^{-7}$ (19.2)	$4.880\,10^{-5}$ (8.1)	$5.912\,10^{-5}$ (8.2)
10^{-3}	0	1	$1.775\,10^{-3}$ (2.0)	$2.067\,10^{-4}$ (5.9)	$2.224\,10^{-4}$ (5.4)
		2	$2.728\,10^{-5}$ (8.5)	$4.555\,10^{-4}$ (5.7)	$5.024\,10^{-4}$ (4.0)
		3	$2.571\,10^{-6}$ (5.1)	$4.844\,10^{-3}$ (8.0)	$5.869\,10^{-3}$ (8.1)
10^{-2}	0	1	$1.775\,10^{-3}$ (2.0)	$5.132\,10^{-4}$ (5.0)	$5.549\,10^{-4}$ (3.6)
		2	$3.412\,10^{-5}$ (5.8)	$4.511\,10^{-3}$ (5.9)	$4.982\,10^{-3}$ (4.2)
		3	$2.570\,10^{-5}$ (5.2)	$4.845\,10^{-2}$ (8.0)	$5.869\,10^{-2}$ (8.1)
0	10^{-7}	1	$1.775\,10^{-3}$ (2.0)	$2.006\,10^{-4}$ (6.4)	$2.156\,10^{-4}$ (5.9)
		2	$2.720\,10^{-5}$ (8.5)	$4.025\,10^{-5}$ (10.9)	$4.398\,10^{-5}$ (10.6)
		3	$3.233\,10^{-7}$ (11.3)	$1.054\,10^{-4}$ (4.8)	$1.237\,10^{-4}$ (5.9)
0	10^{-5}	1	$1.774\,10^{-3}$ (2.1)	$2.248\,10^{-4}$ (6.8)	$2.463\,10^{-4}$ (5.5)
		2	$3.272\,10^{-5}$ (7.2)	$1.062\,10^{-3}$ (4.1)	$1.158\,10^{-3}$ (4.7)
		3	$1.833\,10^{-5}$ (2.8)	$1.056\,10^{-2}$ (4.7)	$1.239\,10^{-2}$ (5.7)
0	10^{-3}	1	$2.455\,10^{-3}$ (5.2)	$1.073\,10^{-2}$ (3.7)	$1.202\,10^{-2}$ (4.2)
		2	$1.858\,10^{-3}$ (3.9)	$1.065\,10^{-1}$ (4.1)	$1.161\,10^{-1}$ (4.7)
		3	$1.834\,10^{-3}$ (2.8)	$1.112\,10^{0}$ (4.7)	$8.911\,10^{-1}$ (3.4)

Finite-depth sensitivity. The zero-floor deployment RMS decreases across the three depths, while even a small learning floor can reverse that ordering for the command report. At $d = 10^{-5}$, the three-packet deployment RMS is $5.912\,10^{-5}$, versus $5.260\,10^{-6}$ at zero floor; at $d = 10^{-3}$ it is $5.869\,10^{-3}$ despite a learning-terminal RMS of $2.571\,10^{-6}$. This policy reconstructs the preparation through inverse powers at progressively later quiet clocks. A small learning terminal therefore need not imply an accurate short-horizon command when a floor perturbs that reconstruction. These are finite-policy outcomes; the fixed-resource rate bounds retain their geometry and noise-floor constants.

At $\tau = 10^{-3}$, the three-packet raw-command RMS is 1.112 and 763 of 2,048 report words activate the deployment guard. Every prefix fallback, correction clip and refined-matrix projection count is zero. For each floor setting the prefix fit is common across depths: A, B projections number 381 and 280 per 2,048 records, and report errors above .025 number 51, except at $\tau = 10^{-3}$ where the counts are 394, 293 and 73. These exceedances remain in every risk. The chosen $\sigma = .005$ is outside the earlier sufficient joint-report certificate; the calculation evaluates the fixed policy at this noise level.

Resources and execution record. Process innovations change the states but not the bounded input rules. The fixed prefix, capped corrections and projected frozen query matrices therefore retain the pathwise learning-energy bound $7.167068015 < 10$. The guard gives deployment energy at most one: total accounting is 23 transitions, input-energy cap 11 and two preparations, with preparation cost and word delivery separate. The largest sampled learning energy is 3.099844; deployment reaches its unit cap (up to floating-point rounding).

All 192 zero-floor plant/block/depth cells match the unchanged policy's complete returned outputs on identical reader primitives. Reversing the two observation and transition changes recovers that policy exactly. This verifies implementation preservation in the perturbation comparison.

Q.1 Sensitivity to the stored quiet reading

The reconstruction in Eq. (6) uses a late reading to recover an earlier preparation. Its sensitivity to that reading helps explain the depth reversal in Table 18. For this calculation, use the true invertible A, B, set $M = B^{-1}AB$, and hold the executed input history and deployment preparation fixed. Thus both the canonical dynamics and the physical-to-input-coordinate decoder are exact. Write $C_H(M) = [M^{H-1}, \dots, I]$. Perturbing only the stored physical reading by ϵ changes the raw command by the exact finite difference

$$\delta U_{\mathrm{raw}} = -C_H(M)^{\dagger} M^{H-s} B^{-1} \epsilon. \tag{143}$$

Since $C_H(M)C_H(M)^{\dagger} = I$ and $\mathcal{C}_H(A, B) = BC_H(M)$, its physical deployment endpoint changes by

$$\delta x_H = \mathcal{C}_H(A, B)\delta U_{\mathrm{raw}} = -BM^{H-s}B^{-1}\epsilon = -A^{H-s}\epsilon. \tag{144}$$

These identities apply to the raw report and locally to the guarded report while the guard remains inactive. They are partial sensitivities with the other arguments frozen. In particular, the cancellation of B in Eq. (144) uses the exact decoder: retaining $\bar{B}^{-1}$ instead gives $-A^{H-s}B\bar{B}^{-1}\epsilon$ even when M is exact.

At the nominal matrix $A_0 = \left(\begin{smallmatrix} .65 & .22 \\ -.08 & .58 \end{smallmatrix}\right)$, the policy uses $H = 3$ and quiet clocks $s = 5, 10, 15$ for $K = 1, 2, 3$. Table 19 evaluates the resulting inverse powers. For an isolated isotropic perturbation with covariance $(\sigma d)^2 I_2$, the raw endpoint RMS is exactly $\sigma d \|A_0^{H-s}\|_F$. This Frobenius expression already includes both state coordinates.

Table 19: Frozen-report sensitivity at the nominal matrix, with $H = 3$. The last column is the isolated quiet-reading calculation at $\sigma = .005$, $d = .001$. Entries evaluate matrix norms, rather than new simulated risks.

K	s	$\|A_0^{3-s}\|_{\mathrm{op}}$	$\|A_0^{3-s}\|_F$	$\sigma d\|A_0^{3-s}\|_F$
1	5	3.22675	3.79121	1.89560×10^{-5}
2	10	45.9977	48.2573	2.41286×10^{-4}
3	15	385.636	426.402	2.13201×10^{-3}

A smaller state at a later clock can therefore coexist with greater sensitivity of the short-horizon report to an additive reader floor. The reader law admits independent relative and additive Gaussian components with the same conditional covariance, motivating the isolated perturbation above. In the actual policy, however, readings also affect estimated matrices and subsequent inputs; those errors can be dependent. The table consequently diagnoses an amplification mechanism at one nominal plant, rather than decomposing the pooled risks in Table 18 into independent variance contributions. The matrix identities were checked in exact rational arithmetic, and the norms were evaluated deterministically.

R Execution floors for fixed learned commands

Learning saturates at its readout and process floors (Appendix N). Execution additionally requires accurate preparation. Let deployment start at $x_0 + \eta$, where η is zero mean with covariance Σ_{prep},

independent of the learning record and deployment noise. Independent process increments during deployment have scale τ_{dep}. Then

$$\mathbb{E}\|x_H^{\mathrm{dep}}\|^2 = \underbrace{\mathbb{E}\|\mathcal{C}_H(\widehat{U}_H - \mathcal{U}_H^*)\|^2}_{\text{learned-command error}} + \underbrace{\mathrm{tr}(A^H \Sigma_{\mathrm{prep}} (A^H)^\top)}_{\text{preparation floor}} + \underbrace{\tau_{\mathrm{dep}}^2 \,\mathrm{tr} \sum_{j=0}^{H-1} A^j (A^j)^\top}_{\text{process floor}}. \tag{145}$$

Independence and zero means eliminate the cross terms even for a biased command estimate. Figure 5(b,c) applies this identity to the unchanged learned commands. Once preparation or process noise dominates, improving the command has little effect on terminal accuracy.

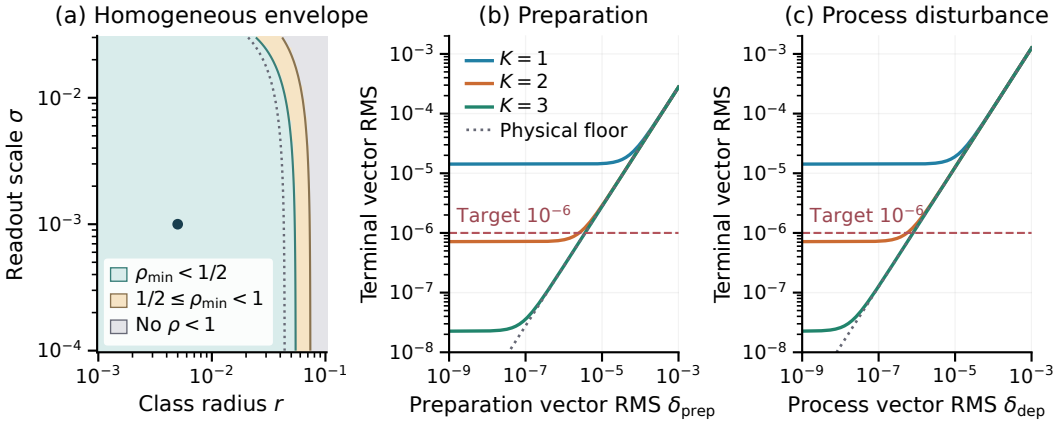


Figure 5: **Finite contraction and execution floors.** (a) Homogeneous comparison on $\|A - I\|, \|B - I\| \le r$, with coarse-error bounds $2r$. Dotted curve: the unweighted row-sum test at $\rho = 1/2$; dot: the certified example $(r, \sigma) = (.005, .001)$. Complete certification also uses the checks in the text. (b,c) Preparation and process floors for fixed commands on 64 systems at $\sigma = .0013$, varying one source at a time. Horizontal axes use vector RMS, per step for process noise ($\tau_{\mathrm{dep}} = \delta_{\mathrm{dep}}/\sqrt{2}$). Floors follow Eq. (145); dashed lines mark 10^{-6}.

S LIKELIHOOD OPTIMIZATION AND PROJECTION ON FIXED RECORDS

We process the same sixteen echo-generated records with the refined reconstruction and a physical-model likelihood fit, examining how initialization and physical-class projection affect task error. All estimators use the unchanged learning commands. These finite runs have neither a constrained-stationarity nor a global-optimality certificate; the results describe the implemented optimization procedure.

Records, target and accounting. Use plants 0 and 24 of Appendix K, $\sigma = .005$, $d = \tau = 0$, radial sensing and $K = 3$. The twenty-transition policy, four-input prefix, correction caps and cached readings are exactly those of that appendix. Generate the PCG64 stream with seed 47111000 in an array of shape $(256, 20, 2)$. Take its first four paths and their negatives, in that order, giving eight records per plant. Every estimator receives the same actual $Y_1, \ldots, Y_{19}$ and $u_0, \ldots, u_{19}$, including readings unused by the refinement rule. Adding record capture preserves every original output exactly. No true parameter or task value enters initialization, optimization or selection of a fit.

The target is the three-step zeroing word. Every reported word is projected onto the same unit ball before a separately prepared noiseless deployment, as in Appendix L. Learning uses twenty transitions and input energy at most ten; deployment adds three transitions, a second preparation and energy at most one. The pooled relative deployment RMS in Table 20 is

$$\left[\frac{1}{16}\sum_{i=1}^{16} \frac{\|A_i^3 x_0 + \mathcal{C}_3(A_i, B_i)\widehat{U}_i\|^2}{\|A_i^3 x_0\|^2}\right]^{1/2}. \tag{146}$$

Table 20: **Likelihood optimization diagnostic; convergence not certified.** RMS is Eq. (146). NLL and gradient columns concern the minimum-NLL selected physical fit; "stopped" counts selected runs that reached the ten-unsuccessful-iteration rule by that checkpoint. The remaining selected runs reached the allowance.

Report	Steps/start	Relative RMS	Mean NLL	Median $\|D^{-1}g\|_\infty$	Stopped
Refined reconstruction	—	3.053×10^{-6}	—	—	—
Coarse plug-in	0	1.442×10^{-2}	—	—	—
Selected likelihood fit	100	1.541×10^{-3}	23105	.0233	6/16
Selected likelihood fit	500	1.319×10^{-3}	22710	.0218	7/16
Selected likelihood fit	2000	1.324×10^{-3}	22642	.0177	8/16

This is a pooled relative error over the sixteen records. The physical rollout and the displayed algebraic residual agree to numerical precision. The pilot provides descriptive diagnostics without a confidence interval.

Full-record likelihood and coordinates. For a candidate plant, unroll its states from x_0 using the actual recorded commands. The causal reading likelihood, up to a constant, is

$$\ell(A,B)=\sum_{t=1}^{19}\left[\log v_t+\frac{\|Y_t-x_t(A,B)\|^2}{2v_t}\right],\qquad v_t=\sigma^2\|x_t(A,B)\|^2+10^{-30}. \tag{147}$$

The likelihood estimator knows σ; the learning policy and direct reconstruction do not use it. The numerical commands are held fixed when differentiating this causal likelihood. This does not condition on the event that an adaptive policy produced its entire input sequence. The added variance is numerical regularization, not a physical floor.

Use eight coordinates $q=(\operatorname{vec}M,z_0,b_2)$, where $z_0=B^{-1}x_0$ and b_2 is the second column of B. Since $x_0=e_1$,

$$B=\left[(x_0-b_2z_{0,2})/z_{0,1},\ b_2\right],\qquad A=BMB^{-1}.$$

After each trial, project $A-A_0$ and $B-B_0$ onto their operator-norm balls of radius .05 and re-express the candidate in q. Three fixed starts are used: the prefix fit; that fit perturbed by $\Delta A=.002\left(\begin{smallmatrix}1&-1\\1&1\end{smallmatrix}\right)$, $\Delta B=.002\left(\begin{smallmatrix}-1&1\\1&1\end{smallmatrix}\right)$ and projected; and a refined-coordinate warm start. For the latter, first set

$$B_w=\bar B+\frac{(x_0-\bar B\widehat z_0)\widehat z_0^{\mathsf T}}{\|\widehat z_0\|^2},\qquad A_w=B_w\widehat M B_w^{-1},$$

then apply the physical class projection. Both its pre-projection and post-projection diagnostics are retained. The known nominal center was removed after a preliminary two-record implementation check because it coincides with plant 0; it is not an optimization start in this pilot.

Finite optimization budget. We use projected damped Fisher descent with an analytic gradient and state Jacobian. If g and F are the gradient and Gaussian Fisher matrix, put $D_{jj}=\sqrt{\max\{F_{jj},10^{-30}\}}$ and take

$$\Delta q=-D^{-1}(D^{-1}FD^{-1}+\lambda I)^{-1}D^{-1}g.$$

Cap $\|\Delta q\|_\infty$ at .05. Try at most sixteen successive halvings, accepting a projected trial only if its NLL decreases by more than $10^{-13}\max\{1,|\ell|\}$. Initially $\lambda=10^{-6}$; an accepted step halves it down to 10^{-10}, and an unsuccessful iteration multiplies it by ten up to 10^{12}. Ten consecutive unsuccessful iterations terminate that run. Allowances are 100, 500 and 2000 iterations *per start*, with the same runs continued between checkpoints. At each allowance choose the smallest NLL among the three starts, never the smallest true loss. The average total objective calls per record across all three starts are 757, 1915 and 4986, respectively.

Initialization, projection, and task error. The refined warm start supplies the selected fit on all sixteen records. At 2000 steps, twelve individual deployment residuals improve on the direct report, while the pooled RMS remains larger. Thus decreasing NLL need not improve the pooled task

loss. On three boundary-plant records, the initial physical projection alone changes the estimated quotient by 4.39×10^{-4} to 1.75×10^{-3} and produces relative deployment residuals 4.57×10^{-4} to 1.50×10^{-3}. The other thirteen warm starts preserve the refined coordinates to roundoff. The largest selected unconstrained normalized gradient $\|D^{-1}g\|_\infty$ is 72.4.

There is also a singular-start mechanism: a square-prefix fit can interpolate its readings and predict zero at its own cancellation, although the observed quiet residual is nonzero. Its pure-relative likelihood is then singular. The numerical variance floor makes this start finite without resolving its conditioning. Analytic gradients agree with an independent complex-step likelihood calculation within 1.77×10^{-14} relative error. At the final reported candidates, the floor accounts for at most 4.36×10^{-11} of a reading variance; reducing it to 10^{-36} changes NLL by at most 1.40×10^{-7}.

On these fixed records, physical-class projection changes three refined warm starts, while exact interpolation makes the prefix starts poorly conditioned for pure-relative likelihood.

T Dense refinement with an explicit control budget

Class, target, and development choice. For $n = m \in \{4, 8, 16\}$, consider

$$.995 \le s_{\min}(A) \le \|A\| \le 1.005, \qquad .8 \le s_{\min}(B) \le \|B\| \le 1.2, \qquad x_0 = e_1.$$

The target is the minimum-energy three-step zeroing word for a separately prepared x_0, with no deployment feedback. Its passive terminal norm is in $[.995^3, 1.005^3]$; a last-input-only zeroing word has norm at most $1.005^3/.8 < \sqrt{2}$. Every reported word is therefore projected onto the radius-$\sqrt{2}$ ball, which contains the target. Sensing is $Y_t = x_t + \sigma\|x_t\|Z_t$ with independent standard Gaussian Z_t, $\sigma = .01/\sqrt{n}$, and $d = \tau = 0$.

A development diagnostic illustrates the importance of correction authority. On plant 3 of the development panel below, the length-16 prefix with time-row commands $2Q$ and correction caps $(\sqrt{32}, 1, .5)$ requires a noiseless first correction of norm $7.350297 > \sqrt{32}$. With query amplitudes $(.5, .25, .125)$ and the clock schedule below, its three corrected-state norms are $2.03213, .92543, .36791$. The endpoint identities still hold at zero noise, but these states do not become quiet as σ decreases. This concerns finite policy resources, not the uniform rate theorem.

The 40-plant development panel uses PCG64 seed 47064002. For plant indices $i = 0, \dots, 39$, generate four positive-diagonal QR frames as specified below. For $i < 8$, bit 0 chooses the constant lower or upper A spectrum and bit 1 does the same for B; if bit 2 is set, replace both by alternating lower/upper endpoints, beginning with the lower endpoint. These eight plants consume no random singular-value draws. For $i \ge 8$, draw independent uniform singular values for A and then B after that plant's four Gaussian arrays. In particular, plant 3 has $A = 1.005Q_1Q_2^\mathsf{T}$ and $B = 1.2Q_3Q_4^\mathsf{T}$.

Development keeps all six combinations of prefix commands $2sQ$, $s \in \{.5, 1\}$, and query amplitudes $(a, a/2, a/4)$, $a \in \{.5, 1, 2\}$. It uses the fitting and reporting rules below, correction caps $(26, 2, 1)$ and common $(T, E) = (44, 767)$. Prefix seed is 47080004. A single PCG64 stream with seed 47070000 draws a C-order Gaussian array with axes $(40 \text{ plants}, 43 \text{ clocks}, 128 \text{ base paths}, 4)$; appending negatives gives 256 signed paths at $\sigma = .005$. The selected combination $(s, a) = (.5, 2)$ has the smallest pooled $K = 3$ deployment MSE in this grid. We fixed the fresh-panel seeds, caps, horizons and sample counts below before generating the evaluation outcomes, then evaluated the selected policy unchanged.

Exact matrices and random streams. Let $L = 4n$. Generate a $4n \times n$ standard Gaussian matrix with NumPy PCG64 seed $47080000 + n$ and reduced QR factor Q. Multiply each column of Q by the sign of the corresponding diagonal of R (zero has positive sign). Prefix command u_t is row t of Q, hence its total energy is n. For each dimension, a PCG64 stream with seed $48100900 + 100n$ generates 16 fresh plants. At each plant draw four $n \times n$ Gaussian arrays successively and apply the same positive-diagonal QR convention to obtain $Q_1, \dots, Q_4$. Then draw n independent uniform singular values in $[.995, 1.005]$ and n in $[.8, 1.2]$; set $A = Q_1 \operatorname{diag}(a_i) Q_2^\mathsf{T}$ and $B = Q_3 \operatorname{diag}(b_i) Q_4^\mathsf{T}$. These panels contain no prescribed corner plants.

For block $b = 0, 1, 2, 3$, use seed $48110900 + 1000n + b$ to draw a C-order array with axes $(16 \text{ plants}, T - 1 \text{ reading clocks}, 32 \text{ base paths}, n)$. Append its negative on the path axis, giving

64 signed paths. No unused process-noise draws are consumed. Banks are paired across K and all reports; commands generate each policy's own physical trajectory. Readings $Y_1, \dots, Y_{T-1}$ are generated once and cached; $Y_0 = x_0$ is exact and no Y_T is used.

A common resource bound for every record. Write $m_+ = 1.005(1.2/.8) = 1.5075$ and

$$R_L = 1.005^L + 1.2 \sum_{t=0}^{L-1} 1.005^{L-1-t} \|u_t\|,$$
$$c_1 = \lceil 1.05(1.005/.8) R_L \rceil, \qquad c_2 = \sqrt{n}, \qquad c_3 = \sqrt{n}/2.$$

The prefix-state bound R_L holds uniformly over the stated class. Every projected coarse fit satisfies $\|\bar{B}^{-1}\bar{A}\| \le 1.005/.8$, so c_1 exceeds every noiseless first proposal even when the fit is inaccurate. All frozen query matrices have norm at most m_+. Consequently the common learning allowance for all three depths is

$$T = 10n + 4, \qquad E = \left\lceil n + c_1^2 + \frac{5n}{4} + \frac{21n}{4}(1 + m_+^2) \right\rceil. \tag{148}$$

The unrounded energy bounds are 273.72368125, 996.4473625, and 4031.894725; (c_1, E) equals $(14, 274)$, $(29, 997)$, and $(61, 4032)$. Input projection against the remaining budget is retained on every history. These are conservative sufficient energy budgets derived from the uniform prefix-state and query bounds. Deployment adds three transitions, at most two input-energy units, and one separately charged preparation after the learning trajectory.

Causal refinement and the three reports. At $t = L$, fit $[\bar{A}, \bar{B}] = X_+ W^\dagger$ from $Y_{0:L}$ and the actual prefix. Use (I, I) if $s_{\min}(W) < 10^{-3}$; otherwise clip the two matrices' singular values to their class intervals. The pseudoinverse uses relative singular-value cutoff 10^{-15}. Retain this fit and initialize $M_o = \bar{B}^{-1}\bar{A}\bar{B}$. Packet r occurs at $t_r = L + (r-1)(2n+1)$, sends the available update and Y_{t_r}, and applies $-\bar{B}^{-1}\bar{A}Y_{t_r}$ capped at c_r. For each coordinate, its frozen batch uses

$$u_1 = a_r e_j, \quad u_2 = -a_r M_o e_j, \qquad \widetilde{M} e_j = M_o e_j + a_r^{-1}\bar{B}^{-1}(Y^+ - \bar{A}^2 Y), \quad a_r = 2^{2-r}. \tag{149}$$

Only after the whole batch is the updated matrix clipped by SVD to $[.995/1.5, 1.5075]$. Its next actionable use follows another packet. Cache the last corrected reading at $s = t_K + 1$. Earlier batches start immediately after correction; the last starts at $8n+3$ and ends at $T-1$, with zero inputs between s and that start. The final input is zero. The last matrix update is private and improves only the offline report. These clocks also follow Appendix H.

For $G_3(M) = [M^2, M, I]$, the refined report uses the actual, possibly clipped input history:

$$\widehat{z}_0 = \widehat{M}^{-s}\bar{B}^{-1}Y_s - \sum_{t<s} \widehat{M}^{-1-t} u_t, \qquad \widehat{U} = -G_3(\widehat{M})^\dagger \widehat{M}^3 \widehat{z}_0. \tag{150}$$

The *prefix plug-in* instead computes the physical command from $(\bar{A}, \bar{B})$ and the known x_0. The *same-record final LS* refits a physical model using every pair (Y_t, u_t, Y_{t+1}), $0 \le t < T-1$, with the same threshold and spectral projection, then computes its plug-in command. All three reports share the acquired record and final word guard. Final LS is ordinary projected least squares, not a likelihood optimum.

Aggregation and finite outcomes. Every dimension/depth has $16 \cdot 4 \cdot 64 = 4096$ paths. Table 2 pools squared deployment norms with equal plant and path weights before taking a square root. For the four block MSEs m_b, its descriptive delta-method Monte Carlo standard error is $s_m/(2\sqrt{4}\sqrt{\bar{m}})$, where s_m is their sample standard deviation. This conditions on the fixed panel and selected policy, with four independent blocks under the supplement-wide statistical conventions. Panel maxima in Table 21 first average all blocks within each plant.

One-packet refinement is worse than both plug-ins; two packets beat final LS only at $n = 4$. Three packets improve on this LS estimator, but the $n = 8, 16$ pooled losses exceed 10^{-3}. The largest observed $\|\widehat{M}^{3-s}\|$ is 1.628: matrix error can matter through the long input history without catastrophic inverse conditioning.

Table 21: Fresh-panel diagnostics. Quiet and matrix-error columns are RMS; matrix error uses Frobenius norm. LS is the same-record final physical fit. The final column is the maximum of the 16 plant-level refined deployment RMS values, not a continuous-class bound.

n	K	Quiet state	Final M error	Final LS deployment	Refined panel max
4	1	.06968	.003442	.006666	.05208
4	2	.006477	.0002758	.005851	.004812
4	3	.0003305	.00002804	.005674	.0006805
8	1	.1344	.009463	.006782	.2670
8	2	.01550	.001002	.006073	.03433
8	3	.0009435	.0001277	.005931	.005719
16	1	.2017	.02082	.007349	.3561
16	2	.02954	.002600	.006653	.04497
16	3	.001809	.0003218	.006505	.006369

Observed energy maxima are $57.117, 112.406, 214.687$, distinct from the hard limits in Eq. (148). Correction, budget, final-word, refined-matrix clipping and both fallback counts are zero. Prefix physical projections number $4080, 4096, 4096$ per 4096 paths. Minimum prefix singular values are $.3520, .3649, .4447$. The singular values and energy maxima describe this finite panel; the class-wide feasibility guarantee comes from Eq. (148).