FROM SMOOTHING TO NOISE CANCELLATION: MEMORY IN SEQUENTIAL ERROR DETECTION

Anonymous authors
Paper under double-blind review

ABSTRACT

Language models often solve mathematical problems through a sequence of calculations and deductions. An evaluator, called a verifier, assigns each step a score intended to indicate whether the reasoning is correct. These scores are imperfect, so a detector must decide how to use earlier scores alongside the current one. Earlier scores may provide evidence, but they may also repeat the same scoring error. We study when these scores should be added as evidence, subtracted to cancel persistent noise, or ignored. We develop a theory of this choice for causal linear detectors under correlated noise, optimizing a class-weighted signal-to-noise ratio (SNR). In a stationary binary Markov model, all historical weights are positive when the state persists longer than the noise, negative when the noise persists longer, and zero when the two persistences match. After normalizing the current weight to one, nonzero total absolute historical weight strictly decreases as observation quality improves. We characterize when this law extends beyond binary states. For finite traces, initialization can reverse the stationary rule or produce mixed signs. A study of 27 methods on 2,000 MATH and GSM8K problems finds that the model-derived kernel improves within-position SNR over current-only scores in all six settings, with simultaneous intervals excluding zero. Matched-SNR comparisons and conditional-moment diagnostics assess when the model-prescribed weights are accurate. The results explain both the role of history and the conditions under which these memory rules apply.

1 INTRODUCTION

Language models often solve mathematical problems by writing intermediate calculations and deductions before giving a final answer. Each calculation or deduction forms a reasoning step. To assess the solution beyond its final answer, an evaluator called a *verifier* checks these steps and assigns scores intended to indicate whether the reasoning is correct. These scores support the selection of candidate solutions and the training of language models (Lightman et al., 2024; Wang et al., 2024). Here, we study how to use them to detect errors while a solution is being generated (Lu et al., 2026). At each step, a detector receives another verifier score and must judge whether the reasoning is currently correct. Because an individual score can be unreliable, the detector may also use scores from earlier steps. Our question is how this history should influence the current judgment.

The difficulty is that persistence has two sources. The reasoning state can remain clean or erroneous across successive steps, making earlier scores evidence about the present. The verifier's scoring errors can also persist, for example because successive evaluations share context. Adding past scores may then reinforce the same noise. Subtracting them can instead remove a persistent nuisance component. Thus, deciding how to use history requires distinguishing persistence in the reasoning state from persistence in the scoring noise.

We develop a theory of this choice for linear detectors that combine current and past scores. Under a correlated-noise model, we ask *when should earlier scores be added, subtracted, or ignored, and how should their total weight change as individual scores become more informative?* The goal is to distinguish the current clean and erroneous states using only scores available so far. This differs from ranking completed solutions by a product or minimum of all step scores (Lightman et al., 2024; Wang et al., 2024). The memory belongs to the score aggregator; the verifier itself may already read the reasoning prefix.

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

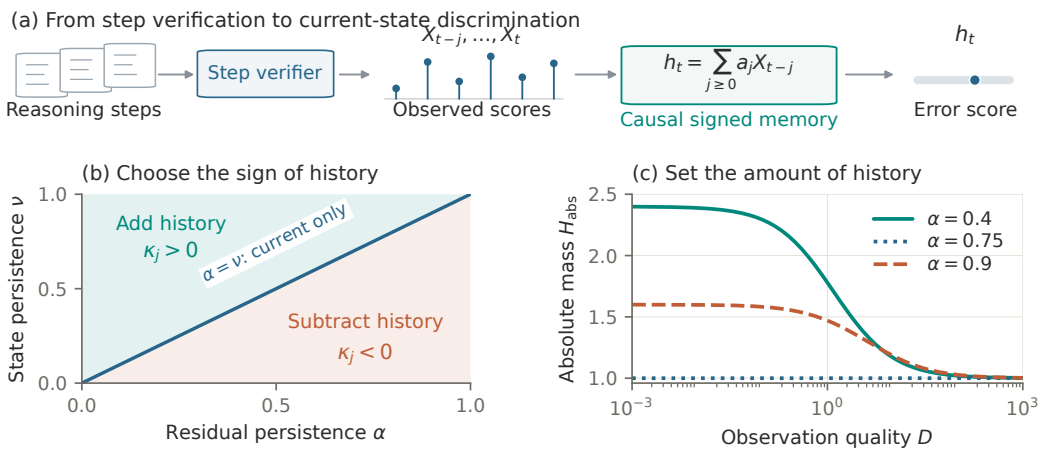


Figure 1: From verifier scores to signed memory. (a) A causal linear aggregator combines available scores to judge the current error state. (b) At stationarity, relative state and residual persistence determine whether history adds evidence or cancels noise. (c) Normalized absolute mass decreases with observation quality away from $\alpha = \nu$. Curves use $(\lambda, \varepsilon, w) = (0.2, 0.05, 0.5)$; $\kappa_j = a_j/a_0$.

Classical linear mean-square estimation and recursive filtering provide the estimation foundations; for a modern treatment, see Särkkä & Svensson (2023). The Fisher criterion gives an inverse-covariance direction (Murphy, 2022, Section 9.2.6), which we apply to a class-weighted signal-to-noise ratio (SNR). The inverse-covariance formula leaves a structural question open: how does each historical coefficient, and their total mass, depend on observation quality? In our binary Markov model, weighting the two current-state classes changes a nonstationary innovation profile over backward lags. Combining this profile with correlated residuals yields an explicit analysis beyond the constant-innovation Wiener specialization.

Our main result characterizes the entire signed kernel at stationarity. All historical coefficients share the sign of state persistence minus residual persistence. Nonzero total historical mass strictly decreases as observations improve, yet individual magnitudes need not decrease. We identify the exact innovation ordering responsible for this behavior, extending the result beyond the binary-chain profile and quantifying tolerable innovation perturbations. A complementary covariance criterion identifies when every nonnegative memory is inferior to the current observation alone.

Initialization changes the backward mean gap through the ratio of past and current state variances. We derive the resulting colored Fisher kernel and its time-dependent cancellation boundary. Identical chain and noise parameters can prescribe positive history at stationarity and negative history near initialization: truncating a stationary kernel does not solve the finite-start problem.

We test the theory under controlled moments and in six observation settings on 2,000 MATH and GSM8K problems. Local probes and a prefix-aware process reward model all benefit from history in within-position SNR. Matching the tuning objective and measuring conditional moments separate this benefit from the accuracy of the prescribed weights.

1. Signed memory and robustness. Theorems 1 and 2 characterize the optimal kernel and give a sharp innovation condition and quantitative perturbation guarantee.

2. Exact finite-start discrimination. Theorem 3 derives a time-dependent cancellation boundary and a linear-time kernel, including initialization-induced reversals and mixed signs.

3. Matched-target evidence. Section 5 tests model-prescribed and empirically selected memory on SNR and AUROC, with paired uncertainty and direct conditional-moment diagnostics.

2 Related Work

Process supervision and streaming detection. Human process labels (Lightman et al., 2024) and rollout-derived labels (Wang et al., 2024) support fine-grained reasoning feedback. Successful con-

tinuation from a prefix need not imply correctness of its last step (Zhang et al., 2025). ProcessBench evaluates first-error identification using expert annotations (Zheng et al., 2025); streaming detection studies evolving step and prefix judgments (Lu et al., 2026). Quickest-detection formulations instead optimize stopping delay and false alarms (Itkin, 2026). Our task is current-state discrimination. CUSUM serves as a score baseline, not a stopping rule.

Linear and nonlinear filtering. Our estimation interpretation builds on classical linear mean-square projection and recursive filtering; Särkkä & Svensson (2023) provide a modern filtering treatment. Under stationary class weighting, our SNR objective has the same optimizer as linear minimum-mean-square error; Gaussianity of the hidden state is unnecessary. In the stationary equal-correlation special case, cancellation follows directly from the classical linear-projection equations; Appendix A.2 gives the short calculation. Our analysis resolves coefficient signs and total-mass monotonicity for nonconstant, class-weighted innovations, and proves that individual magnitudes can grow as total history shrinks. The innovation criterion quantifies a broader validity range, while the finite-start result changes the cancellation boundary through conditional state variances. The covariance-inverse direction itself is classical Fisher discrimination (Murphy, 2022). Gaussian-process measurement models also extend filtering beyond autoregressive noise (Kurtz & Lin, 2019). Exact two-state HMM filtering is inexpensive; its stability concerns sensitivity to initial beliefs (McDonald & Yuksel, 2019). Linear discrimination isolates what conditional moments determine and makes each historical contribution interpretable; it offers no computational advantage over a two-state HMM.

Objectives and evaluation. Scores, probabilities, and likelihood ratios have different calibration requirements (Reid & Williamson, 2010; Guo et al., 2017). Optimizing a surrogate does not automatically optimize ranking (Agarwal, 2014). Balanced SNR gives a distribution-free AUROC bound, but the bound can be uninformative; our experiments measure AUROC directly. All selection occurs within outer training folds (Bischl et al., 2023). Paired bootstrap comparisons condition on fitted models and preserve problem clusters. Overlapping cross-validation fits are not independent replications (Bates et al., 2024); the conditional bootstrap intervals do not estimate variability across retrained models. Matching the tuning metric and controlling trajectory position are essential to interpreting empirical memory gains.

3 Model and Objective

State and observations. Let $z_t \in \{0, 1\}$ denote a clean/error state, with error coded as one, and set $S_t = 2z_t - 1$. The stationary binary Markov chain has

$$P = \begin{pmatrix} 1-\lambda & \lambda \\ \varepsilon & 1-\varepsilon \end{pmatrix}, \qquad \nu = 1 - \lambda - \varepsilon, \qquad \pi_1 = \frac{\lambda}{\lambda + \varepsilon}, \quad \pi_0 = 1 - \pi_1. \tag{1}$$

We assume $\lambda, \varepsilon > 0$ and $\lambda + \varepsilon < 1$. Then $0 < \nu < 1$, the chain is reversible, and its lag-j correlation is ν^j. Its geometric correlation horizon is $H_\nu = 1/(1-\nu)$.

The normalized observation stream obeys the path-conditional moment model

$$X_t = DS_t + \xi_t, \quad \mathbb{E}[\xi_t \mid z_{\mathbb{Z}}] = 0, \quad \mathbb{E}[\xi_t \xi_s \mid z_{\mathbb{Z}}] = 2D\alpha^{|t-s|}, \tag{2}$$

where $D > 0$ and $0 \le \alpha < 1$. State persistence ν carries historical evidence; residual persistence α carries shared observation noise. The white-noise specialization is

$$\mathbb{E}[\xi_t \xi_s \mid z_{\mathbb{Z}}] = 2D\mathbf{1}_{\{t=s\}} \qquad (\alpha = 0). \tag{3}$$

Independent Gaussian emissions suffice for the white model; state-independent AR(1) residuals suffice for the colored model. The theorems require conditional moments, not Gaussianity.

For raw scores Y_t with path-conditional means m_{z_t}, common residual variance σ^2, and residual correlation $\alpha^{|t-s|}$, the transformation

$$D = \frac{(m_1 - m_0)^2}{2\sigma^2}, \qquad X_t = \frac{m_1 - m_0}{\sigma^2}\left(Y_t - \frac{m_0 + m_1}{2}\right) \tag{4}$$

satisfies (2). Here D measures mean separation relative to residual variation. For equal-variance Gaussian emissions it equals the KL divergence between the two one-step emission distributions; it is not generally the divergence rate of a correlated stream. Marginal score moments alone cannot verify residual moments conditional on the complete state path.

Objective and memory. A stable causal linear score is $h_t(a) = \sum_{j\geq 0} a_j X_{t-j}$ with $\sum_j |a_j| < \infty$. The exponential family is the special case

$$h_t(\rho) = \sum_{j\geq 0} \rho^j X_{t-j} = \rho h_{t-1} + X_t, \qquad 0 \leq \rho < 1. \tag{5}$$

For fixed $w \in [0,1]$, define

$$\Delta_h = \mathbb{E}[h_t \mid z_t = 1] - \mathbb{E}[h_t \mid z_t = 0], \qquad \mathrm{SNR}_w(h) = \frac{\Delta_h}{\sqrt{w\mathrm{Var}(h_t \mid 1) + (1-w)\mathrm{Var}(h_t \mid 0)}}. \tag{6}$$

The class weight sets the relative contribution of each class variance; it does not specify a false-alarm cost. We optimize kernels with $\Delta_h > 0$, holding (P, α, w) fixed as D varies. For a kernel normalized by $a_0 = 1$, total absolute memory is $\sum_j |a_j|$. This measures historical coefficient mass, including subtraction, rather than an integer window length. The stationary-weight objective $w = \pi_1$ is equivalent to linear mean-square estimation; other weights change the discriminative objective. AUROC is measured separately in the experiments.

4 Optimal Discriminative Memory

4.1 The covariance seen by a discriminative filter

Scale observations by D and write $R = 2/D$ for the normalized residual variance. At lag j, the conditional mean difference is $2g_j$, where $g_j = \nu^j$. For a fixed class weight, let

$$\begin{aligned} U_j &= w\mathrm{Var}(S_{t-j} \mid z_t = 1) + (1-w)\mathrm{Var}(S_{t-j} \mid z_t = 0), \\ q_j &= U_j - \nu^2 U_{j-1} = c + \beta\nu^j, \qquad j \geq 1, \\ c &= 4\pi_0\pi_1(1-\nu^2), \qquad \beta = 4(\varepsilon - \lambda)(w - \pi_1). \end{aligned} \tag{7}$$

Here $U_0 = 0$, $q_j > 0$, and $q_{j+1} - \nu q_j = c(1-\nu) > 0$. Reversibility gives the weighted within-class covariance

$$N_{jk} = \nu^{|j-k|} U_{\min(j,k)} + R\alpha^{|j-k|}. \tag{8}$$

For a kernel a, its squared SNR is $4\langle a, g\rangle^2 / \langle a, Na\rangle$. The Fisher covariance-inverse rule therefore gives $a \propto N^{-1}g$ (Murphy, 2022). The explicit innovation profile in (7) lets us characterize every coefficient and its dependence on observation quality.

4.2 Smoothing, cancellation, and signed history

Theorem 1 (Optimal memory with correlated residuals)**.** *Assume (1) and (2), with $0 < \nu < 1$, $0 \leq \alpha < 1$, $D > 0$, and fixed $w \in [0,1]$. Among absolutely summable real causal kernels with positive conditional mean separation, the SNR optimum is unique up to positive scale. Its current coefficient is positive. Normalize it as $\kappa_0 = 1$ and define*

$$H_{\mathrm{abs}}(D) = \sum_{j\geq 0} |\kappa_j(D)|.$$

Every historical coefficient has the same sign:

$$\mathrm{sign}(\kappa_j) = \mathrm{sign}(\nu - \alpha), \qquad j \geq 1. \tag{9}$$

For $\alpha \neq \nu$, H_{abs} strictly decreases with D, and

$$H_{\mathrm{abs}}(0+) = 1 + \frac{|\nu - \alpha|}{1-\nu}, \qquad H_{\mathrm{abs}}(\infty) = 1. \tag{10}$$

At $\alpha = \nu$, the optimum uses only the current observation. The sign law and strict total-mass monotonicity also hold in every finite lag window containing at least one historical observation.

When the state persists longer than the residual noise, the optimal filter smooths by adding earlier scores. When the noise persists longer, it subtracts earlier scores to suppress predictable residual variation. Equal persistence cancels the value of linear history. The stationary cancellation is a direct consequence of the classical linear-projection equations (Appendix A.2); the theorem establishes the sign and absolute-memory structure for every fixed within-class variance weight, including weights that produce nonstationary reverse-time innovation profiles.

Proof mechanism. Set $s = R(1-\alpha^2)$. An auxiliary sequence $x_0 = 1$ minimizes

$$\sum_{j\geq 1} \frac{(x_j - \nu x_{j-1})^2}{q_j} + \sum_{j\geq 1} \frac{(x_j - \alpha x_{j-1})^2}{s}. \tag{11}$$

Its tridiagonal normal equations give an $O(L)$ backward recursion in a window of L lags. For $\alpha \neq \nu$, every ratio x_j/x_{j-1} lies strictly between α and ν and moves toward α as D increases. The Fisher kernel is recovered by

$$a_0 = \frac{1-\alpha x_1}{s}, \qquad a_j = \frac{(1+\alpha^2)x_j - \alpha x_{j-1} - \alpha x_{j+1}}{s}. \tag{12}$$

The innovation inequality $q_{j+1} > \alpha\nu q_j$ fixes the sign of every historical coefficient. If $X = \sum_{j\geq 1} x_j$, telescoping gives

$$\sum_{j\geq 0} \kappa_j = \frac{(1-\alpha)[1+(1-\alpha)X]}{1-\alpha x_1}.$$

For $\alpha < \nu$ this is H_{abs}; for $\alpha > \nu$ it is $2 - H_{\mathrm{abs}}$. Monotonicity of the auxiliary ratios then proves the total-memory law. Geometric bounds justify the infinite-history limit and summation. Appendix A supplies the complete construction, finite-window endpoint, and uniqueness proof.

The finite endpoint is $a_L = (x_L - \alpha x_{L-1})/s$. Appendix A gives the complete $O(L)$ recursion. The historical coefficient is a difference of auxiliary values, so individual magnitudes need not follow a single ratio; telescoping retains the strict aggregate order.

What does not survive from white noise. Total history can shrink while an individual historical coefficient grows. For $(\lambda, \varepsilon, w, \alpha) = (0.01, 0.89, 0.5, 0.9)$, increasing D from 0.1 to 1 changes $|\kappa_2|$ from 0.11194 to 0.26277, while H_{abs} decreases from 1.88402 to 1.85352. Appendix A gives an exact rational derivative certificate. Thus coefficientwise shortening is a property of the white-noise case, not a general interpretation of better observations.

The classical specialization remains inexpensive. When $w = \pi_1$ or $\lambda = \varepsilon$, $q_j = c$ is constant. Let $A = (s/c)(1+\nu^2) + 1 + \alpha^2$, $B = \nu s/c + \alpha$, and $\tau = 2B/(A + \sqrt{A^2 - 4B^2})$. Then $\kappa_j = (\tau - \alpha)\tau^{j-1}$ for $j \geq 1$, so the optimal stationary-weight score has the recursion

$$h_t = \tau h_{t-1} + X_t - \alpha X_{t-1}. \tag{13}$$

This classical smoothing-and-whitening filter subtracts the previous observation. Arbitrary class weights replace its constant auxiliary ratio with a lag-dependent recursion set by innovations.

4.3 The boundary for nonnegative memory

Proposition 1 (A residual-correlation criterion)**.** *Let the normalized residual covariance be RT, with $T_{00} = 1$, zero path-conditional residual means, and a bounded coercive total covariance N. Write $r_j = T_{0j}$. The current-only kernel uniquely maximizes SNR up to positive scale over nonnegative stable kernels if and only if $r_j \geq \nu^j$ at every historical lag. It maximizes SNR over signed stable kernels if and only if $r_j = \nu^j$ at every lag. The same statements hold on any fixed finite window.*

Normalize $\langle a, g\rangle = 1$ and expand the variance around e_0:

$$\langle a, Na\rangle - R = 2R\sum_{j\geq 1} a_j(r_j - g_j) + \langle a - e_0, N(a - e_0)\rangle. \tag{14}$$

Both terms are nonnegative under the stated inequalities. A perturbation at any violating lag proves necessity. For geometric residuals, the criterion reduces to $\alpha \geq \nu$. It therefore covers *all* nonnegative linear memories, not just exponential filters. A lag-one residual correlation alone cannot establish this boundary for an arbitrary covariance sequence.

The simplest signed witness uses one historical score. Writing $c_w = q_1$, its optimal coefficient and squared-SNR gain over current-only are

$$\kappa_1^{(1)} = \frac{2(\nu-\alpha)}{c_w D + 2(1-\alpha\nu)}, \qquad \Delta\mathrm{SNR}_w^2 = \frac{4(\nu-\alpha)^2}{c_w + R(1-\alpha^2)}. \tag{15}$$

Thus negative history can be useful precisely where positive smoothing fails.

4.4 WHICH INNOVATION DEPARTURES PRESERVE THE LAW?

The innovation ordering drives the proof. Characterizing that ordering yields an exact validity criterion beyond the affine profile in (7), together with a perturbation guarantee.

Theorem 2 (A sharp innovation condition)**.** *Keep $g_j = \nu^j$ and geometric residual covariance, but allow arbitrary positive innovations q_j in $U_j = \nu^2 U_{j-1} + q_j$. In infinite history assume q_j are bounded above and bounded away from zero. For $\alpha \neq \nu$, all historical coefficients have sign $\nu - \alpha$ for every $D > 0$ if and only if*

$$q_{j+1} \geq \alpha\nu q_j \quad \text{at every available adjacent lag.} \tag{16}$$

Under this condition, H_{abs} strictly decreases with D, with infinite-history limits given by (10).

Sufficiency follows from the auxiliary recursion, including equality cases. Necessity follows from its strong-signal expansion: an interior coefficient has leading term $R(\nu-\alpha)\alpha^{j-1}(q_j^{-1} - \alpha\nu/q_{j+1})$. A violated inequality therefore reverses that coefficient at sufficiently large D. Appendix B proves both directions and gives an exact mixed-sign example.

The binary-chain profile satisfies $q_{j+1} > \nu q_j$. Consequently, for $0 < \alpha < 1$, any fixed innovation perturbation obeying

$$|\widetilde{q}_j/q_j - 1| \leq \eta, \qquad \eta \leq \frac{1-\alpha}{1+\alpha} \tag{17}$$

preserves both laws beyond the binary-chain innovation family. The mean-gap decay and residual covariance remain fixed. In their Gaussian two-class realization, the same directions also maximize AUROC, since $\mathrm{AUROC} = \Phi(\mathrm{SNR}/\sqrt{2})$, with Φ the standard normal distribution function.

4.5 A FINITE START CHANGES THE CANCELLATION BOUNDARY

Now start the same chain at time zero with error probability $p_0 \in (0,1)$ and observe only times $0, \ldots, t$. Write $p_s = \pi_1 + (p_0 - \pi_1)\nu^s$ and $v_s = p_s(1-p_s)$. The backward mean-gap ratio is

$$g_{t,j} = \nu^j \frac{v_{t-j}}{v_t}, \qquad 0 \leq j \leq t. \tag{18}$$

It differs from ν^j because past uncertainty and current uncertainty need not be equal.

Theorem 3 (Finite-start discrimination)**.** *For $t \geq 1$, the exact colored, class-weighted covariance admits an $O(t)$ Fisher-kernel recursion under the initial law p_0. Among nonnegative kernels, the current score is optimal if and only if*

$$\alpha \geq \alpha_{\mathrm{crit}}(t) := \max_{1 \leq j \leq t} g_{t,j}^{1/j}. \tag{19}$$

Among signed kernels it is optimal if and only if $g_{t,j} = \alpha^j$ at every available historical lag. Neither boundary depends on observation quality D or the class weight w.

The reverse transition slope becomes $b_j = g_{t,j}/g_{t,j-1}$. Computing its class-weighted innovation variance gives the same covariance construction and auxiliary recursion with lag-dependent b_j; Appendix C derives them explicitly. The residual-column criterion then gives (19).

Initialization can reverse the stationary prescription. With $p_0 = \lambda = \varepsilon = 0.1$, $\nu = 0.8$, and $t = 1$, the boundary is $g_{1,1} = 20/41$. At $\alpha = 0.6$, the finite optimum subtracts history even though the stationary optimum smooths. For $w = 1/2, D = 1$, its normalized kernel is $(1, -1886/16295)$. At $t = 3$ the history has mixed signs. The finite-start optimum therefore changes the prescribed signs themselves, which truncating a stationary kernel cannot achieve.

At stationarity and $\alpha = 0$, the full kernel is geometric exactly when $w = \pi_1$ or $\lambda = \varepsilon$; the white-noise and scalar exponential results remain in Appendix D. The finite-start LMMSE and Gaussian HMM provide separate estimation references. Linearity specifies the discriminative class and moment requirements, rather than a computational advantage over a two-state HMM.

5 EXPERIMENTS

We evaluate the predictions under the stated moments and on reasoning traces. Model-predicted SNR, held-out SNR, and AUROC separate covariance-model transfer from objective mismatch.

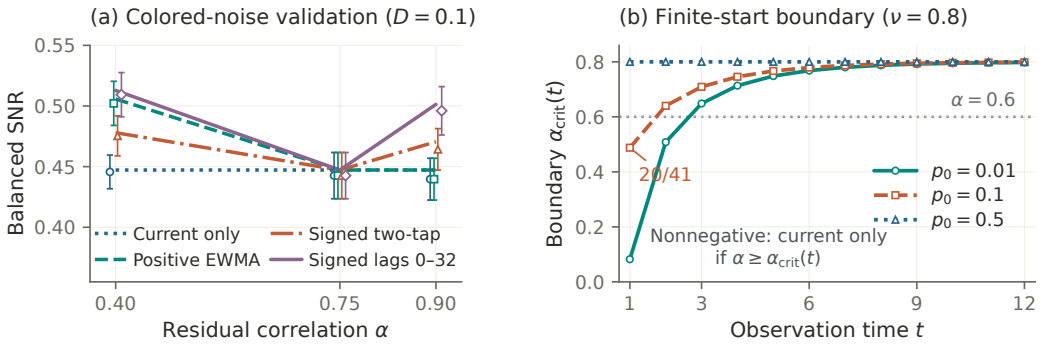


Figure 2: Two routes to cancellation. (a) Exact SNR and empirical mean $\pm$ one standard deviation over three seeds at $D = 0.1$: signed history remains useful when nonnegative smoothing fails. (b) The finite-start boundary $\alpha_{\text{crit}}(t)$ depends on p_0, with $\lambda = \varepsilon = 0.1$. At $\alpha = 0.6$, early nonnegative optima can be current-only despite a stationary boundary of $\nu = 0.8$.

5.1 Structural and finite-start checks

Independent dense covariance solves verify 2,000 colored-kernel configurations, with maximum coefficient discrepancy 2.96×10^{-14}. No sign or total-mass violation occurs; 116 comparisons exhibit growth of an individual historical magnitude. Another 108 geometric-noise and 48 general-covariance configurations test the nonnegative criterion. For the finite-start extension, enumeration of all binary paths in 720 configurations through $t = 5$ verifies the conditional moments and the $O(t)$ solver, with maximum error below 7.4×10^{-14}. Exhaustive nonnegative-QP face searches in 180 cases independently verify its boundary across the same model family.

A stochastic test fixes $(\lambda, \varepsilon, w) = (0.2, 0.05, 0.5)$ and crosses $D \in \{0.1, 1, 10\}$ with $\alpha \in \{0.4, 0.75, 0.9\}$. Each of three seeds retains 200,000 steps with state-independent Gaussian AR(1) residuals. We compare current-only, optimal nonnegative exponential, signed two-tap, and signed 32-lag scores. At $\alpha = \nu = 0.75$ all coincide. At $\alpha = 0.9, D = 0.1$, nonnegative memory gives no improvement, while the 32-lag signed kernel increases exact SNR from 0.4472 to 0.5013 (Figure 2a). The finite-start boundary in Figure 2b shows a different source of cancellation: initialization alone can change whether nonnegative history is useful. Appendices L and M report the complete checks and white-noise approximation scan, with all configurations and seeds.

5.2 Observation streams and evaluation targets

We use one Math-Shepherd trajectory for each of 1,000 MATH and 1,000 GSM8K problems, totaling 7,538 and 3,558 steps. Labels describe rollout-derived solvability, rather than expert judgments of logical correctness (Wang et al., 2024; Zhang et al., 2025). Two frozen encoders, Qwen2.5-Math-1.5B-Instruct and Qwen2.5-7B-Instruct, supply local problem-and-current-step embeddings to a linear probe. A frozen Qwen2.5-Math-PRM-7B supplies scores from each full causal reasoning prefix (Zhang et al., 2025). All observation systems share problems, labels and folds within each task; pretrained PRM overlap with these benchmarks is unknown.

Five outer folds hold out complete problems. Three-fold inner cross-fitting supplies probe scores for calibration and model parameters; all memory selection and learned aggregation use only outer-training problems. Frozen PRM scores bypass probe fitting but retain downstream fitting restrictions. The 24-method AUROC study includes full white/colored kernels, AR(1)-corrected exponential memory, metric-tuned EWMA, learned lag/position aggregation, CUSUM, finite-start estimation and Gaussian HMM controls. The SNR study adds two SNR-tuned EWMAs and the exact finite-start colored Fisher score. For that score, we normalize the current coefficient to one and subtract the model's class midpoint at each lag. This fixes cross-position offsets without affecting within-position SNR or AUROC; it is a discriminative score, not a posterior probability.

	MATH			GSM8K		
Method	Gen. 7B	Math 1.5B	PRM 7B	Gen. 7B	Math 1.5B	PRM 7B
Per-step	0.737/0.720	0.725/0.708	0.787/0.762	0.740/0.701	0.735/0.693	0.874/0.889
Metric-tuned EWMA	0.758/0.752	0.741/0.736	0.808/0.781	0.750/0.720	0.744/0.716	0.894/0.906
White kernel	0.736/0.752	0.714/0.736	0.800/0.781	0.745/0.721	0.739/0.715	0.894/0.906
AR(1) exponential	0.736/0.752	0.716/0.736	0.799/0.780	0.750/0.719	0.744/0.712	0.891/0.903
Colored kernel	0.746/0.749	0.726/0.734	0.805/0.782	0.749/0.717	0.743/0.709	0.889/0.902
Lag logistic	0.756/0.746	0.739/0.730	0.795/0.777	0.751/0.717	0.744/0.709	0.886/0.905
Lag + position	0.773/0.750	0.756/0.734	0.822/0.782	0.773/0.719	0.763/0.712	0.924/0.907
Tuned fixed gain	0.773/0.753	0.749/0.736	0.822/0.781	0.774/0.720	0.766/0.715	0.923/0.906
Δ_{within}	-0.003	-0.003	$+0.001$	-0.003	-0.007	-0.004
95% lower bound	-0.007	-0.008	-0.002	-0.011	-0.018	-0.010
95% upper bound	$+0.002$	$+0.003$	$+0.004$	$+0.005$	$+0.004$	$+0.001$

Table 1: Held-out AUROC as pooled / within-position. Local probes use the problem and current step; PRM uses full prefixes. Kernel rows use balanced weighting. Metric-tuned EWMA selects memory separately for each reported metric. Final rows give colored minus within-position-tuned EWMA and its paired six-setting simultaneous interval. Complete results appear in Appendix M.

Ranking and moment separation. Pooled AUROC compares all steps within a fold. Within-position AUROC compares only opposite-label pairs at the same step index, weighting positions by eligible pair counts. With ω the fraction of pairs sharing a position, exactly

$$\mathrm{AUC}_{\text{pooled}} = \omega\,\mathrm{AUC}_{\text{within}} + (1-\omega)\,\mathrm{AUC}_{\text{cross}}. \tag{20}$$

Empirical balanced SNR uses the signed class-mean difference divided by the square root of the mean sample variance. Its within-position version averages position-specific SNRs using pair counts. Eligible positions must have at least 80 observations per class in the outer training set, retaining up to position 10 on MATH and position 4 on GSM8K; test counts remain at least 17 and 14 per class. The two extra EWMAs select memory by their corresponding training SNRs. All reported metrics are mean outer-fold values. Appendix N defines degenerate-score handling.

We use paired problem-bootstrap intervals conditional on fitted models. The AUROC study uses 2,000 draws and an exploratory family of 24 methods, two metrics, three references and six settings. The SNR study uses 20,000 attempted draws and five fixed contrasts across two SNR metrics and six settings. Neither family includes refitting or research-selection uncertainty. Learned lag/position methods and tuned fixed gain select by pooled AUROC; within-position results are secondary.

5.3 Separating target mismatch from failed moment transfer

History helps across all six streams; exact weights transfer unevenly. Relative to current-only, the colored kernel improves within-position SNR in all six settings, with simultaneous intervals excluding zero (Table 2). Its precise weight prescription transfers less consistently: on MATH it remains below SNR-tuned EWMA in pooled SNR. The differences for the two local probes are -0.0539 and -0.0574, with simultaneous intervals $[-0.0959, -0.0118]$ and $[-0.1040, -0.0109]$. Matching the selection objective therefore isolates a transfer failure beyond SNR–AUROC mismatch: the fitted covariance prescription does not reproduce the empirical optimum.

Finite-start correction has a measurable but nonuniform effect. On GSM8K PRM, finite-start correction increases within-position SNR by 0.0084, with interval $[0.0001, 0.0166]$. On MATH General-7B it decreases that metric by 0.0051, with interval $[-0.0099, -0.0003]$. The fitted finite moment model predicts the optimal finite direction to outperform a stationary prefix on its own SNR, but the corresponding empirical ordering need not hold. Exact initialization corrects the finite-trace covariance, while empirical ordering also depends on the observation model.

Observed dependence extends beyond lag one. For MATH General-7B, held-out class-centered residual correlations at lags two and five are 0.403 and 0.330, versus AR(1) predictions 0.220 and 0.023. Math-1.5B gives 0.464 and 0.366, versus 0.259 and 0.034. Residual means also depend

	Held-out within-position SNR				Finite minus stationary	
Setting	Current	Colored	Finite	SNR-tuned EWMA	Δ	Simultaneous 95% CI
MATH Gen. 7B	0.853	0.995	0.990	1.012	-0.0051	$[-0.0099, -0.0003]$
MATH Math 1.5B	0.801	0.933	0.930	0.956	-0.0037	$[-0.0088, +0.0014]$
MATH PRM 7B	1.024	1.154	1.153	1.146	-0.0007	$[-0.0070, +0.0057]$
GSM8K Gen. 7B	0.747	0.816	0.816	0.830	-0.0001	$[-0.0059, +0.0056]$
GSM8K Math 1.5B	0.708	0.765	0.762	0.783	-0.0031	$[-0.0084, +0.0022]$
GSM8K PRM 7B	1.903	2.042	2.050	2.086	$+0.0084$	$[+0.0001, +0.0166]$

Table 2: Held-out within-position balanced SNR on training-selected supported positions. SNR-tuned EWMA selects its memory by this same target. The final column compares the exact finite-start colored score with the stationary colored kernel, using the SNR study's paired simultaneous family. The full pooled-SNR and model-prediction comparison is in Appendix N.

on nearby transitions: General-7B's current residual averages -0.312 training standard deviations at clean-to-error transitions and $+0.127$ during continuing error. These deviations are visible in held-out conditional moments. Their association with filter performance does not identify which departure causes a particular ranking change. Appendix N reports mean gaps, variance diagnostics, and support counts for all six streams and all measured lags.

A direct application of the model-free first-column criterion to position-matched sample moments finds positive historical directions in 580 of 600 fold–position–lag cells. All 91 position–lag groups with consistent five-fold directions and sufficient class support are positive. This diagnostic supports adding history, without establishing a real-data cancellation effect.

AUROC and label dynamics provide complementary checks. The colored kernel has no within-position AUROC advantage over matched-metric EWMA with a simultaneous interval excluding zero. Its GSM8K PRM comparison against pooled-tuned EWMA is adverse; an independent 10,000-draw bootstrap check retains a negative interval, $[-0.00901, -0.00032]$. Position-aware controls improve pooled ranking without comparable within-position gains. Label dynamics also depart from stationarity: MATH has no recovery among 2,749 observed error-source transitions, and GSM8K's second-order next-label model improves held-out log loss by 0.0110 over first order. Matched targets and position control thus distinguish useful history from accurate weights.

6 Discussion and Conclusion

The innovation criterion gives a necessary and sufficient condition for common historical signs across observation qualities, with a quantitative robustness guarantee. The finite-start solution reveals a distinct cancellation boundary governed by backward variance ratios: initialization can reverse historical signs or produce mixed signs. Together, these results extend constant-innovation filtering to a broader innovation family and class-weighted finite traces. The coefficientwise counterexample further separates decreasing total historical mass from uniform shrinkage at every lag: an individual coefficient can grow as observations improve while total historical mass decreases.

The colored kernel improves within-position SNR over current-only in all six settings; direct covariance diagnostics also favor adding history. Matched-SNR comparisons distinguish this benefit from the accuracy of prescribed weights; conditional-moment measurements expose remaining model departures. The resulting filters and diagnostic conditions connect memory selection to mean gaps, covariance columns, and innovation profiles, with initialization treated explicitly.

The practical distinction is between selecting a historical direction and selecting all weights jointly. A covariance-column diagnostic addresses the former directly from score moments; the full Fisher kernel addresses the latter under a specified model. Keeping these decisions separate makes the theory useful even when a fitted stationary covariance fails to predict the best empirical weights.

Reproducibility statement

Code supplies dataset parsers, the nested 24-method AUROC evaluation and 27-method SNR study, simulations, covariance checks, and vector figures. Appendices L, M, and N specify preprocessing, seeds, estimation, and uncertainty. GPU inference uses a single RTX 5090 with separate project environments. Analytic proofs cover the colored-kernel results; the selected Lean certificates in Appendix K concern the earlier scalar exponential analysis. The complete reproducibility package, including code and Lean 4 certificates, is available in an anonymous Git repository at `https://anonymous.4open.science/r/signed-memory-8F83`.

AI use statement

We used generative AI tools for language polishing and supplementary checks of theoretical arguments. The theoretical results and proofs were developed and verified by the authors, with AI serving only as an additional checking aid. All AI-assisted edits and suggestions incorporated into the paper were reviewed by the authors. We take full responsibility for the final content of this work.

Ethics statement

To the best of our knowledge, this work does not raise specific ethical concerns requiring additional discussion under the ICLR Code of Ethics.

References

Shivani Agarwal. Surrogate regret bounds for bipartite ranking via strongly proper losses. *The Journal of Machine Learning Research*, 15(1):1653–1674, 2014.

Stephen Bates, Trevor Hastie, and Robert Tibshirani. Cross-validation: what does it estimate and how well does it do it? *Journal of the American Statistical Association*, 119(546):1434–1445, 2024.

Bernd Bischl, Martin Binder, Michel Lang, et al. Hyperparameter optimization: Foundations, algorithms, best practices, and open challenges. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 13(2):e1484, 2023.

Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q. Weinberger. On calibration of modern neural networks. In Doina Precup and Yee Whye Teh (eds.), *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, volume 70 of *Proceedings of Machine Learning Research*, pp. 1321–1330. PMLR, 2017.

Igor Itkin. Quickest detection of hallucination onset: Delay bounds and learned cusum statistics, 2026. arXiv preprint arXiv:2606.12476.

Vince Kurtz and Hai Lin. Kalman filtering with gaussian processes measurement noise, 2019. arXiv preprint arXiv:1909.10582.

Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let's verify step by step. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net, 2024.

Haolang Lu, Minghui Pan, Ripeng LI, Guoshun Nan, Jialin Zhuang, Zijie Zhao, Zhongxiang Sun, Kun Wang, and Yang Liu. Streaming hallucination detection in long chain-of-thought reasoning. In Maria Liakata, Viviane P. Moreira, Jiajun Zhang, and David Jurgens (eds.), *Findings of the Association for Computational Linguistics: ACL 2026*, pp. 21157–21183, San Diego, California, United States, 2026. Association for Computational Linguistics. ISBN 979-8-89176-395-1. doi: 10.18653/v1/2026.findings-acl.1064.

Curtis McDonald and Serdar Yuksel. Exponential filter stability via dobrushin's coefficient, 2019. arXiv preprint arXiv:1910.08463.

Kevin P Murphy. *Probabilistic machine learning: an introduction*. MIT press, 2022.

Mark D Reid and Robert C Williamson. Composite binary losses. *The Journal of Machine Learning Research*, 11:2387–2422, 2010.

Simo Särkkä and Lennart Svensson. *Bayesian filtering and smoothing*, volume 17. Cambridge university press, 2023.

Peiyi Wang, Lei Li, Zhihong Shao, Runxin Xu, Damai Dai, Yifei Li, Deli Chen, Yu Wu, and Zhifang Sui. Math-shepherd: Verify and reinforce LLMs step-by-step without human annotations. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 9426–9439, Bangkok, Thailand, 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.510.

Zhenru Zhang, Chujie Zheng, Yangzhen Wu, Beichen Zhang, Runji Lin, Bowen Yu, Dayiheng Liu, Jingren Zhou, and Junyang Lin. The lessons of developing process reward models in mathematical reasoning. In Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and Mohammad Taher Pilehvar (eds.), *Findings of the Association for Computational Linguistics: ACL 2025*, pp. 10495–10516, Vienna, Austria, 2025. Association for Computational Linguistics. ISBN 979-8-89176-256-5. doi: 10.18653/v1/2025.findings-acl.547.

Chujie Zheng, Zhenru Zhang, Beichen Zhang, Runji Lin, Keming Lu, Bowen Yu, Dayiheng Liu, Jingren Zhou, and Junyang Lin. ProcessBench: Identifying process errors in mathematical reasoning. In Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and Mohammad Taher Pilehvar (eds.), *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 1009–1024, Vienna, Austria, 2025. Association for Computational Linguistics. ISBN 979-8-89176-251-0. doi: 10.18653/v1/2025.acl-long.50.

Appendix Contents

A Optimal causal kernels with correlated residuals **14**

A.1 Covariance structure and statement 14

A.2 The classical stationary cancellation special case 15

A.3 An auxiliary quadratic problem 15

A.4 Bounds and monotonicity of the auxiliary ratios 16

A.5 Recovering the discriminative kernel 16

A.6 Signs of the historical coefficients 17

A.7 Infinite-history construction and strict monotonicity 17

A.8 Total absolute history and its limits 18

A.9 Finite-window algorithm 18

A.10 Individual coefficients need not shorten 19

A.11 A criterion for general residual covariance 19

B Innovation conditions and robustness **20**

B.1 Relative perturbations of the innovations 22

B.2 A sign reversal outside the innovation condition 23

C Finite-start kernels with correlated residuals **23**

C.1 Finite-start moments and optimality 23

C.2 Score normalization and class midpoints 26
C.3 Initialization effects and stationary limits 26

D Exponential and white-noise results, and finite-start recursions 27
D.1 Mean separation and path uncertainty 27
D.2 A memory law for every fixed class weight 28
D.3 Beyond exponential memory: the optimal causal linear kernel 28
D.4 The stationary weighting recovers causal linear estimation 29
D.5 When correlated residuals remove the benefit of memory 29
D.6 Finite starts: exact moments and time-varying memory 30

E Proofs of the stationary memory law 30
E.1 Existence of the stationary sum 30
E.2 Proof of Lemma 1 30
E.3 Proof of Theorem 8 31

F Optimal causal kernels for arbitrary class weights 31
F.1 Weighted covariance and the Fisher optimum 32
F.2 Finite information recursion and its infinite limit 32
F.3 Dependence on observation quality 33
F.4 Geometric tails and exact geometric kernels 33

G The stationary-weight linear estimator 34

H Correlated observation residuals 35
H.1 Exact objective and the zero-memory threshold 35
H.2 Positive memory when residuals mix faster than the state 36

I Finite-history calculations 37
I.1 Proof of Proposition 5 37
I.2 Proof of the finite linear estimator 37
I.3 Score origin and finite accumulation 38

J Finite-history and perturbation bounds 38

K Scope of formal verification 39

L Baseline measurements and supporting experiments 40
L.1 Trajectory measurements 40
L.2 Problem sampling and representations 41
L.3 Nested fitting and working-model calibration 41
L.4 Score definitions 42

L.5 Common metrics and training diagnostics 42
L.6 Synthetic model validation 43
L.7 Resources and reproduction 44
L.8 Class weights and correlated residuals 44
L.9 Sensitivity to transition smoothing 46
L.10 Verification of the unrestricted optimal kernel 46

M Evaluation on local and prefix-aware observation streams **46**
M.1 Controlled covariance and emission experiments 46
M.2 Samples, observations, and problem splits 48
M.3 Working-model estimates and score conventions 48
M.4 Aggregation methods and selection 49
M.5 Metrics and a unified uncertainty family 49
M.6 Transfer results and observed limits 50
M.7 Diagnostics for the label process 51

N From working-model SNR to held-out discrimination **51**
N.1 Signed SNR, position support, and training selection 51
N.2 Finite-start colored discrimination 54
N.3 Predicted and observed objectives 54
N.4 Prespecified comparisons and uncertainty 55
N.5 Residual covariance and conditional mean diagnostics 55
N.6 Direct first-column directions without an AR(1) fit 56

A Optimal causal kernels with correlated residuals

We prove Theorem 1 by representing the optimal kernel through a tridiagonal quadratic problem. The representation determines the sign of every historical coefficient and the change in total absolute history with observation quality. It also yields a finite-window algorithm.

A.1 Covariance structure and statement

Fix $\lambda, \varepsilon > 0$ with $\lambda + \varepsilon < 1$, and write

$$\nu = 1 - \lambda - \varepsilon, \qquad \pi_1 = \frac{\lambda}{\lambda + \varepsilon}, \qquad \pi_0 = 1 - \pi_1, \qquad R = \frac{2}{D}, \qquad s = R(1 - \alpha^2), \tag{21}$$

where $D > 0$ and $0 \le \alpha < 1$. Thus $0 < \nu < 1$ and $s > 0$. Fix a class weight $w \in [0, 1]$ and define

$$\begin{aligned} c &= 4\pi_0\pi_1(1 - \nu^2), \qquad \beta = 4(\varepsilon - \lambda)(w - \pi_1), \\ q_j &= c + \beta\nu^j \quad (j \ge 1), \qquad U_0 = 0, \qquad U_j = \nu^2 U_{j-1} + q_j. \end{aligned} \tag{22}$$

Here U_j is the weighted conditional variance of the binary state $S_{t-j} = 2z_{t-j} - 1$ given z_t, and q_j is its reverse-time innovation variance. For $r_{j-1} = \pi_1 + (w - \pi_1)\nu^{j-1} \in [0, 1]$, the equivalent expression

$$q_j = 4\big[(1 - r_{j-1})\lambda(1 - \lambda) + r_{j-1}\varepsilon(1 - \varepsilon)\big] \tag{23}$$

gives $0 < q_{\min} \le q_j \le q_{\max} < \infty$, where $q_{\min} = 4\min\{\lambda(1 - \lambda), \varepsilon(1 - \varepsilon)\}$ and $q_{\max} = 4\max\{\lambda(1 - \lambda), \varepsilon(1 - \varepsilon)\}$. A second identity will control the kernel signs:

$$q_{j+1} - \nu q_j = c(1 - \nu) > 0, \qquad q_{j+1} > \alpha\nu q_j. \tag{24}$$

For normalized observations indexed by lag, suppose the conditional mean difference is $2g_j = 2\nu^j$. The residuals have zero mean conditional on the entire state path and conditional covariance $R\alpha^{|j-k|}$. The weighted within-class observation covariance is

$$N = K + RT_\alpha, \qquad K_{jk} = \nu^{|j-k|}U_{\min(j,k)}, \qquad (T_\alpha)_{jk} = \alpha^{|j-k|}, \qquad g_j = \nu^j. \tag{25}$$

In particular, $Ke_0 = 0$ and $N_{00} = R$. For a nonzero real causal kernel a, the squared signal-to-noise ratio is

$$\mathrm{SNR}_w(a)^2 = \frac{4\langle a, g\rangle^2}{\langle a, Na\rangle}. \tag{26}$$

We choose the positive orientation $\langle a, g\rangle > 0$. Finite kernels use lags $0, \dots, L$, with $L \ge 1$; infinite stable kernels belong to $\ell^1(\mathbb{N}_0)$.

Theorem 4 (Signs and total history for finite and infinite kernels)**.** *Under (21)–(25), the positive-oriented optimal stable causal kernel is unique up to positive scale and has $a_0 > 0$. Normalize it by $\kappa_j = a_j/a_0$ and set*

$$H_{\mathrm{abs}}(D) = \sum_{j \ge 0} |\kappa_j(D)|.$$

If $\alpha \ne \nu$, every historical coefficient satisfies

$$\mathrm{sgn}(\kappa_j) = \mathrm{sgn}(\nu - \alpha), \qquad j \ge 1, \tag{27}$$

and H_{abs} strictly decreases with D, with limits

$$\lim_{D \downarrow 0} H_{\mathrm{abs}}(D) = 1 + \frac{|\nu - \alpha|}{1 - \nu}, \qquad \lim_{D \to \infty} H_{\mathrm{abs}}(D) = 1. \tag{28}$$

If $\alpha = \nu$, the optimal kernel is e_0 at every D.

The same sign law and strict total-history monotonicity hold for every finite window $0, \dots, L$. Its weak-signal limit is

$$\lim_{D \downarrow 0} H_{\mathrm{abs},L}(D) = 1 + |\nu - \alpha| \left[\frac{1 - \nu^{L-1}}{1 - \nu} + \frac{\nu^{L-1}}{1 - \alpha\nu}\right], \tag{29}$$

and its strong-signal limit is one.

The covariance-inverse direction in (26) is the Fisher optimum (Murphy, 2022). The proof below constructs that direction and establishes its coefficient structure. Positivity below the threshold means that the unrestricted optimum also solves the nonnegative-kernel problem there. Above the threshold, every historical coefficient of the unrestricted optimum is negative.

A.2 THE CLASSICAL STATIONARY CANCELLATION SPECIAL CASE

For completeness, the stationary equal-correlation case has a direct linear-projection interpretation. This is a specialization of classical mean-square estimation, not a new cancellation principle; see Särkkä & Svensson (2023) for modern filtering background. Set $w = \pi_1$, $v = 4\pi_0\pi_1$, and write the centered, normalized observation as $Y_t = T_t + \eta_t$, where $T_t = S_t - \mathbb{E}S_t$ and $R = 2/D$. At $\alpha = \nu$, the model assumptions give, for every $j \geq 0$,

$$\mathrm{Cov}(T_t, Y_{t-j}) = v\nu^j, \qquad \mathrm{Cov}(Y_t, Y_{t-j}) = (v+R)\nu^j.$$

Consequently $\widehat{T}_t = \frac{v}{v+R} Y_t$ satisfies

$$\mathrm{Cov}(T_t - \widehat{T}_t, Y_{t-j}) = v\nu^j - \frac{v}{v+R}(v+R)\nu^j = 0 \quad \text{for every } j \geq 0.$$

The estimation residual is therefore orthogonal to the closed linear span of all current and past observations, so no linear history term improves the mean-square estimate. This argument needs only the stated second moments, not a Gaussian hidden state. The stationary-weight SNR equivalence in Appendix G gives the same current-only discriminative direction. Theorem 4 additionally covers arbitrary fixed class weights and determines the coefficient signs and total-memory behavior away from equality.

A.3 AN AUXILIARY QUADRATIC PROBLEM

Set $x_0 = 1$ and consider

$$\mathcal{E}_s(x) = \sum_{j\geq 1} \frac{(x_j - \nu x_{j-1})^2}{q_j} + \frac{1}{s}\sum_{j\geq 1}(x_j - \alpha x_{j-1})^2. \tag{30}$$

For a finite window, both sums stop at L. For infinite history, minimize over $(x_j)_{j\geq 1} \in \ell^2$. The quadratic part is bounded and coercive. Indeed, if $v_0 = 0$, the unilateral shift bound gives

$$\sum_{j\geq 1} \frac{(v_j - \nu v_{j-1})^2}{q_j} \geq \frac{(1-\nu)^2}{q_{\max}} \sum_{j\geq 1} v_j^2. \tag{31}$$

The boundary $x_0 = 1$ contributes a bounded linear functional. Thus (30) has a unique minimizer.

Define

$$B_j = \frac{\nu}{q_j} + \frac{\alpha}{s}, \qquad A_j = \frac{1}{q_j} + \frac{\nu^2}{q_{j+1}} + \frac{1+\alpha^2}{s}. \tag{32}$$

At an interior index the normal equation is

$$B_j x_{j-1} - A_j x_j + B_{j+1} x_{j+1} = 0. \tag{33}$$

In a finite window the terminal diagonal is instead $A_L = 1/q_L + 1/s$.

The finite minimizer can be constructed without solving a dense system. Starting from

$$\tau_{L-1} = \frac{\nu s/q_L + \alpha}{s/q_L + 1}, \tag{34}$$

compute, backward for $j = L-1, \ldots, 1$,

$$\tau_{j-1} = F_j(s, \tau_j), \qquad F_j(s,t) = \frac{\nu s/q_j + \alpha}{s/q_j + \nu^2 s/q_{j+1} + 1 + \alpha^2 - (\nu s/q_{j+1} + \alpha)t}. \tag{35}$$

Then $x_j = \prod_{k=0}^{j-1} \tau_k$ satisfies the terminal and interior normal equations. The following bounds make the construction positive and identify how it changes with D.

A.4 Bounds and monotonicity of the auxiliary ratios

Write $\Delta = \nu - \alpha$, $m = \min\{\nu, \alpha\}$, and $M = \max\{\nu, \alpha\} < 1$. For $\Delta \neq 0$, the terminal ratio (34) is a strict weighted average of ν and α. We show that every backward update maps $[m, M]$ strictly into (m, M).

Let $d_j(s,t)$ denote the denominator of $F_j(s,t)$. Its noise part, $1 + \alpha^2 - \alpha t$, is positive on $[m, M]$. If $\nu \geq \alpha$, its state part is at least s/q_j. If $\alpha > \nu$, (24) gives

$$\frac{1}{q_j} + \frac{\nu(\nu - \alpha)}{q_{j+1}} = \frac{q_{j+1} - \alpha\nu q_j + \nu^2 q_j}{q_j q_{j+1}} > 0,$$

which bounds that state part from below. Hence $d_j(s,t) > 0$ on the interval, and F_j is strictly increasing in t. Endpoint substitution yields

$$\begin{aligned} F_j(s,\alpha) - \alpha &= \frac{s\Delta(1/q_j - \alpha\nu/q_{j+1})}{d_j(s,\alpha)}, \\ F_j(s,\nu) - \nu &= -\frac{\Delta(1 - \alpha\nu)}{d_j(s,\nu)}. \end{aligned} \tag{36}$$

Their signs and monotonicity in t establish the interval claim. Backward induction therefore proves

$$m < \tau_j < M \quad (\alpha \neq \nu), \qquad \tau_j = \nu \quad (\alpha = \nu). \tag{37}$$

For fixed t, differentiation with respect to s gives

$$\frac{\partial F_j}{\partial s}(s,t) = \frac{E_j(t)}{d_j(s,t)^2}, \qquad E_j(t) = \frac{\nu(1 + \alpha^2 - \alpha t) - \alpha}{q_j} + \frac{\alpha\nu(t - \nu)}{q_{j+1}}. \tag{38}$$

The affine function E_j has endpoint values

$$E_j(\alpha) = \Delta\left(\frac{1}{q_j} - \frac{\alpha\nu}{q_{j+1}}\right), \qquad E_j(\nu) = \frac{\Delta(1 - \alpha\nu)}{q_j}. \tag{39}$$

Both have the sign of Δ by (24). The terminal ratio also has derivative with respect to s of that sign. Since each update increases in its successor, backward induction proves that every finite τ_j strictly increases with s when $\Delta > 0$ and strictly decreases with s when $\Delta < 0$. As $s = 2(1 - \alpha^2)/D$, the directions reverse with respect to D.

A.5 Recovering the discriminative kernel

A Gaussian auxiliary model evaluates the quadratic form associated with N. Let

$$G_0 = \theta, \qquad G_j = \nu G_{j-1} + u_j, \qquad Y_j = G_j + e_j, \qquad e_j = \alpha e_{j-1} + v_j, \tag{40}$$

where θ is deterministic, the innovations are mutually independent Gaussian variables, and $\mathrm{Var}(u_j) = q_j$, $\mathrm{Var}(e_0) = R$, and $\mathrm{Var}(v_j) = s$. Conditional on θ, the vector Y has mean θg and covariance N. Gaussianity is used to calculate these finite quadratic forms; the original observations require only the stated moments.

For a finite window, whitening the residuals gives the joint negative log density, up to a constant,

$$\frac{(Y_0 - \theta)^2}{2R} + \frac{1}{2}\sum_{j=1}^{L} \frac{(G_j - \nu G_{j-1})^2}{q_j} + \frac{1}{2s}\sum_{j=1}^{L} \big(Y_j - \alpha Y_{j-1} - G_j + \alpha G_{j-1}\big)^2. \tag{41}$$

At $Y = 0$, the minimizing latent path is $G_j = \theta x_j$, where x minimizes (30). Gaussian integration over $G_1, \ldots, G_L$ differs from minimization by a log determinant independent of (θ, Y). Differentiating the minimized quadratic with respect to Y at $Y = 0$ therefore identifies its mixed θ, Y term. The log-likelihood coefficient linear in θ is

$$\frac{Y_0}{R} + \sum_{j=1}^{L} d_j(Y_j - \alpha Y_{j-1}), \qquad d_j = \frac{x_j - \alpha x_{j-1}}{s}. \tag{42}$$

The direct Gaussian likelihood has linear coefficient $g^T N^{-1} Y$. Thus (42) gives exactly $a = N^{-1} g$, with

$$\begin{aligned} a_0 &= \frac{1 - \alpha x_1}{s}, \\ a_j &= \frac{(1+\alpha^2) x_j - \alpha x_{j-1} - \alpha x_{j+1}}{s}, \quad 1 \le j < L, \\ a_L &= \frac{x_L - \alpha x_{L-1}}{s}. \end{aligned} \tag{43}$$

The first equality uses $s = R(1 - \alpha^2)$. In infinite history the middle formula applies at every $j \ge 1$.

A.6 Signs of the historical coefficients

Since $x_1 = \tau_0 < M < 1$, (43) gives $a_0 > 0$. At an interior historical index, substitute (33) into the numerator of the coefficient formula to obtain

$$1 + \alpha^2 - \frac{\alpha}{\tau_{j-1}} - \alpha \tau_j = \frac{E_j(\tau_j)}{B_j}. \tag{44}$$

Indeed, multiplying the left-hand side by B_j and using $B_j / \tau_{j-1} = A_j - B_{j+1} \tau_j$ cancels all terms involving $1/s$ and leaves (38)'s expression for E_j. By (39), $E_j(\tau_j)$ has the sign of Δ. The factors B_j and x_j / s are positive. Consequently every interior a_j has sign Δ. At the finite terminal index, $a_L = x_{L-1}(\tau_{L-1} - \alpha)/s$ has the same sign by (34).

If $\alpha = \nu$, all ratios equal ν, so $x_j = \nu^j$. Equation (43) then gives $a_0 = 1/R$ and $a_j = 0$ for $j \ge 1$.

A.7 Infinite-history construction and strict monotonicity

Let $x^{(L)}$ be the finite minimizer, extended by zero beyond L. The ratio bounds imply

$$0 < x_j^{(L)} \le M^j \quad (0 \le j \le L). \tag{45}$$

The geometric tail gives compactness in both ℓ^1 and ℓ^2: any sequence of window lengths has a coordinatewise convergent subsequence, and dominated tails upgrade this convergence to either norm. At every fixed positive index, the finite normal equation eventually has the interior form (33). Each subsequential limit therefore solves the infinite normal equations. Coercivity in (31) makes that solution unique, so the full sequence $x^{(L)}$ converges in ℓ^1 and ℓ^2 to the minimizer of (30).

The limiting coordinates are positive. When $\alpha > 0$, the finite ratio lower bound $\tau_j \ge m > 0$ gives $x_j \ge m^j$. When $\alpha = 0$, the nonnegative limit cannot have a zero coordinate: a first zero in (33) would force its positive predecessor to vanish because $B_j > 0$ and its successor is nonnegative. Thus the limiting ratios are well defined and lie in $[m, M]$. Their recursion (35) maps this interval strictly into its interior whenever $\Delta \neq 0$, proving the strict bounds in infinite history as well.

Finite monotonicity passes to weak monotonicity of each limiting ratio. The dependence of the local update on s makes it strict. For example, if $\Delta > 0$ and $s_2 > s_1$, then

$$\begin{aligned} \tau_{j-1}(s_2) &= F_j(s_2, \tau_j(s_2)) \\ &\ge F_j(s_2, \tau_j(s_1)) > F_j(s_1, \tau_j(s_1)) = \tau_{j-1}(s_1). \end{aligned} \tag{46}$$

For $\Delta < 0$, both inequalities reverse. Hence every infinite ratio has the same strict monotonicity as its finite counterparts.

The coefficient formulas and (45) imply that the finite kernels, extended by zero, converge in ℓ^1 to the infinite coefficients. The terminal formula changes at a moving index, whose contribution is bounded by a constant times M^{L-1}. The limit is therefore absolutely summable.

To identify its optimality, first note that K and T_α are bounded self-adjoint operators on ℓ^2. Their row sums are bounded by $(1+\nu)/(1-\nu)$ and $(1+\alpha)/(1-\alpha)$, respectively, because $0 \le U_j \le 1$. The Toeplitz spectral density

$$f_\alpha(\theta) = \frac{1 - \alpha^2}{1 - 2\alpha \cos\theta + \alpha^2} \ge \frac{1 - \alpha}{1 + \alpha} \tag{47}$$

gives $N \succeq R(1-\alpha)/(1+\alpha)I$. This identity follows from summing the geometric Fourier series for T_α; the quadratic-form bound extends from finite sequences by density. For any fixed row, the finite normal equations $N^{(L)}a^{(L)} = g^{(L)}$ pass to the limit by coefficient convergence and boundedness. Hence $Na = g$. Cauchy–Schwarz in the N inner product gives

$$\frac{4\langle b, g\rangle^2}{\langle b, Nb\rangle} \le 4\langle g, N^{-1}g\rangle, \tag{48}$$

with equality exactly when b is a scalar multiple of a. Positive orientation restricts the scalar to be positive. The constructed $a \in \ell^1$ thus attains the optimum over all stable causal kernels, and the sign argument applies to every historical index.

A.8 Total absolute history and its limits

Let $X = \sum_{j\ge1} x_j < \infty$. Summing the infinite coefficient formulas gives

$$H_{\mathrm{sgn}} := \sum_{j\ge0} \kappa_j = \frac{(1-\alpha)[1+(1-\alpha)X]}{1-\alpha x_1}. \tag{49}$$

One way to obtain this identity is to sum (42): $\sum_j a_j = 1/R + (1-\alpha)\sum_{j\ge1} d_j$ and $\sum_{j\ge1} d_j = [(1-\alpha)X - \alpha]/s$. For a finite window the corresponding identity is

$$H_{\mathrm{sgn},L} = \frac{(1-\alpha)[1+(1-\alpha)\sum_{j=1}^{L-1} x_j + x_L]}{1-\alpha x_1}. \tag{50}$$

When $\nu > \alpha$, all kernel coefficients are positive, so $H_{\mathrm{abs}} = H_{\mathrm{sgn}}$. Increasing D strictly decreases every auxiliary ratio and hence every x_j. The numerator of (49) strictly decreases, while its positive denominator increases. Therefore H_{abs} strictly decreases. If $\alpha = 0$, the denominator is constant and the numerator still strictly decreases.

When $\nu < \alpha$, all historical coefficients are negative and $\kappa_0 = 1$, so $H_{\mathrm{abs}} = 2 - H_{\mathrm{sgn}}$. Increasing D strictly increases every x_j. The positive numerator in (49) increases and its denominator decreases, making H_{sgn} strictly increase. Thus H_{abs} again strictly decreases. The same reasoning with (50) proves finite-window monotonicity.

For the weak-signal limit, $D \downarrow 0$ is equivalent to $s \to \infty$. Compare the minimizer of (30) with the path $x_j = \nu^j$, whose first energy term is zero. Its second unscaled sum is finite, so

$$\sum_{j\ge1} \frac{(x_j(s) - \nu x_{j-1}(s))^2}{q_j} \longrightarrow 0.$$

Since $x_0(s) = 1$, induction over fixed indices yields $x_j(s) \to \nu^j$. For the strong-signal limit, multiply (30) by s and compare with $x_j = \alpha^j$. Its second energy term is zero and its first is finite by $q_j \ge q_{\min}$; hence

$$\sum_{j\ge1} (x_j(s) - \alpha x_{j-1}(s))^2 \longrightarrow 0 \quad (s \downarrow 0).$$

This gives $x_j(s) \to \alpha^j$ at every fixed index. The bound (45), uniform in s, permits dominated summation in both limits. Substitution into (49) yields

$$\lim_{D\downarrow0} H_{\mathrm{sgn}} = \frac{1-\alpha}{1-\nu}, \qquad \lim_{D\to\infty} H_{\mathrm{sgn}} = 1. \tag{51}$$

The sign law converts these into (28). The same finite energy comparisons and substitution into (50) give (29). This completes the proof of Theorem 4 and the infinite-history statement in Theorem 1.

A.9 Finite-window algorithm

The backward ratios and forward coefficient pass give the following $O(L)$ construction. The common factor $1/s$ can be omitted because the output is normalized by the current coefficient.

1. Given $(\lambda, \varepsilon, w, \alpha, D, L)$, compute $\nu, q_1, \dots, q_L$, and s from (21)–(22).
2. Set τ_{L-1} by (34). For $j = L-1, L-2, \dots, 1$, set $\tau_{j-1} = F_j(s, \tau_j)$ using (35).
3. Set $x_0 = 1$ and, for $j = 1, \dots, L$, compute $x_j = \tau_{j-1} x_{j-1}$.
4. Form

$$\begin{aligned} \widetilde{a}_0 &= 1 - \alpha x_1, \\ \widetilde{a}_j &= (1+\alpha^2) x_j - \alpha x_{j-1} - \alpha x_{j+1}, \quad 1 \le j < L, \\ \widetilde{a}_L &= x_L - \alpha x_{L-1}, \qquad \kappa_j = \frac{\widetilde{a}_j}{\widetilde{a}_0}, \qquad 0 \le j \le L. \end{aligned} \tag{52}$$

The denominator is positive by (37). Each pass is linear in L, and storing the ratios, response, and output requires $O(L)$ memory. At exact cancellation $\alpha = \nu$, the output is directly $\kappa = e_0$.

A.10 Individual coefficients need not shorten

Total absolute history can decrease while the magnitude of a particular coefficient increases. Consider

$$\lambda = \frac{1}{100}, \qquad \varepsilon = \frac{89}{100}, \qquad \nu = \frac{1}{10}, \qquad w = \frac{1}{2}, \qquad \alpha = \frac{9}{10}. \tag{53}$$

An exact finite example already occurs at $L = 2$. Solving the three-dimensional covariance equations gives

$$\kappa_2(1/10) = -\frac{287020000}{2308036101}, \qquad \kappa_2(1) = -\frac{11602000}{36659301}. \tag{54}$$

The second coefficient has larger magnitude at the larger D.

For an infinite-history calculation, put $\zeta = 1/s = D/[2(1-\alpha^2)]$ and differentiate the auxiliary minimizer at $\zeta = 0$, where $x_j = \nu^j$. The prior precision on historical coordinates is the bounded, coercive operator $Q_\nu = K_{\mathrm{h}}^{-1}$, where K_{h} is K restricted to indices $j \ge 1$. The noise precision is a bounded tridiagonal operator Q_α. The normal equations have the form $(Q_\nu + \zeta Q_\alpha)x = b_\nu + \zeta b_\alpha$, with the boundary contributions b_ν, b_α determined by $x_0 = 1$. A bounded-inverse expansion at $\zeta = 0$ justifies differentiation. The noise-gradient forcing at $x_j = \nu^j$ is $\Delta(1-\alpha\nu)\nu^{j-1}$. Therefore, with $\dot{x}_0 = 0$,

$$\dot{x}_j := \left.\frac{dx_j}{d\zeta}\right|_{\zeta=0} = -\Delta(1-\alpha\nu)\nu^{j-1}\left[\sum_{k=1}^{j-1} U_k + \frac{U_j}{1-\nu^2}\right]. \tag{55}$$

To evaluate this expression, apply K_{h} to $(\nu^{k-1})_{k\ge1}$ and split the sum at $k = j$:

$$\sum_{k\ge1} K_{jk}\nu^{k-1} = \nu^{j-1}\sum_{k=1}^{j} U_k + \frac{U_j\nu^{j+1}}{1-\nu^2} = \nu^{j-1}\left[\sum_{k=1}^{j-1} U_k + \frac{U_j}{1-\nu^2}\right].$$

Differentiating the normalized coefficient formula gives

$$\left.\frac{d\kappa_2}{dD}\right|_{D=0+} = \frac{(1+\alpha^2)\dot{x}_2 - \alpha\dot{x}_1 - \alpha\dot{x}_3 + \Delta\nu\alpha\dot{x}_1}{2(1-\alpha^2)(1-\alpha\nu)}. \tag{56}$$

For (53), exact substitution yields

$$\kappa_2(0+) = -\frac{2}{25}, \qquad \left.\frac{d\kappa_2}{dD}\right|_{D=0+} = -\frac{922973}{2671875} < 0. \tag{57}$$

Thus $|\kappa_2(D)|$ strictly increases on a sufficiently small positive interval. The auxiliary-ratio monotonicity in (38) does not imply coefficientwise shortening of the discriminative kernel.

A.11 A criterion for general residual covariance

The decision to use nonnegative history can be characterized without geometric residual correlations. The relevant quantities are the residual covariances with the current observation.

Theorem 5 (Residual-column criterion)**.** *Let $g_0 = 1$ and $g_j > 0$ at each historical index. Suppose $K \succeq 0$, $Ke_0 = 0$, T is a residual covariance with $T_{00} = 1$, and $N = K + RT$ with $R > 0$. In finite dimension assume $N \succ 0$. In infinite dimension assume N is bounded and coercive on ℓ^2, $g \in \ell^2$, and optimize over nonzero ℓ^1 kernels. Write $r_j = T_{0j}$. Then e_0 is the unique optimum up to positive scale over nonnegative kernels if and only if*

$$r_j \geq g_j \qquad \text{for every } j \geq 1. \tag{58}$$

It is the unique positive-oriented optimum up to positive scale over signed kernels if and only if $r_j = g_j$ at every historical index. If $r_j < g_j$, a sufficiently small positive coefficient at lag j improves on e_0; if $r_j > g_j$, a sufficiently small negative coefficient improves on e_0.

Proof. For a nonnegative kernel, normalize $\langle a, g\rangle = 1$ and put $d = a - e_0$. Then $\langle d, g\rangle = 0$ and $Ne_0 = Rr$, with $r_0 = g_0 = 1$. Expanding the quadratic form gives

$$\langle a, Na\rangle - R = 2R \sum_{j\geq 1} a_j(r_j - g_j) + \langle d, Nd\rangle. \tag{59}$$

Under (58), the right-hand side is nonnegative and is strictly positive unless $a = e_0$. This proves sufficiency and uniqueness. For necessity, direct differentiation gives

$$\frac{d}{db}\mathrm{SNR}_w(e_0 + be_j)^2\Big|_{b=0} = \frac{8}{R}(g_j - r_j). \tag{60}$$

A violated inequality therefore supplies a strictly improving positive perturbation. If all entries satisfy $r_j = g_j$, then $Ne_0 = Rg$, and Cauchy–Schwarz in the N inner product gives

$$\langle a, g\rangle^2 = \frac{\langle a, Ne_0\rangle^2}{R^2} \leq \frac{\langle a, Na\rangle}{R}.$$

Equality holds exactly at scalar multiples of e_0; positive orientation selects positive multiples. If any entry differs, (60) gives an improving signed perturbation with sufficiently small magnitude to preserve positive orientation. In infinite dimension, boundedness of N and membership in $\ell^1 \subset \ell^2$ justify these pairings and the quadratic expansion. The row $r = Ne_0/R$ belongs to ℓ^2, so the displayed scalar sums also converge. □

For stationary residuals, r_j is their lag-j correlation. The criterion compares this correlation with the current-state mean separation at that lag. Substituting $r_j = \alpha^j$ and $g_j = \nu^j$ gives the threshold $\alpha \geq \nu$ for nonnegative kernels and the equality $\alpha = \nu$ for signed cancellation. Correlations between two historical residuals determine the size of an improvement, but they do not change this criterion for whether an improvement exists.

B Innovation conditions and robustness

The sign law depends on how quickly successive state innovation variances can decrease. This section characterizes that dependence beyond the binary-chain profile in (22). It also gives a relative perturbation bound under which the sign law and total-history monotonicity persist.

Keep $0 < \nu < 1$, $0 \leq \alpha < 1$, $R = 2/D$, and the geometric mean difference $2g_j = 2\nu^j$. Replace the binary-chain innovations by positive numbers q_j and define

$$U_0 = 0, \qquad U_j = \nu^2 U_{j-1} + q_j, \qquad K_{jk} = \nu^{|j-k|}U_{\min(j,k)}, \qquad N = K + RT_\alpha. \tag{61}$$

The residual correlation remains $(T_\alpha)_{jk} = \alpha^{|j-k|}$. These moments are realized by the Gaussian state and residual recursions in (40), with class label $\theta \in \{-1, 1\}$.

In that equal-covariance Gaussian realization, $\mathrm{AUROC}(a) = \Phi(\mathrm{SNR}(a)/\sqrt{2})$ for positive-oriented scores, where Φ is the standard normal distribution function. The same covariance-optimal direction therefore maximizes AUROC in this realization. This implication does not extend to the path mixtures of the original binary-state model without further distributional assumptions.

Theorem 6 (Innovation criterion)**.** *Consider either a finite window $0, \dots, L$, with $L \geq 1$ and $q_1, \dots, q_L > 0$, or infinite history with $0 < q_{\min} \leq q_j \leq q_{\max} < \infty$. Fix the innovation profile independently of D. For the covariance in (61), the direction $a = N^{-1}g$ has $a_0 > 0$; in infinite history it belongs to ℓ^1 and attains the optimum over stable causal kernels. Normalize $\kappa_j = a_j/a_0$.*

If $\alpha \neq \nu$, every historical coefficient has sign $\nu - \alpha$ for every $D > 0$ if and only if

$$q_{j+1} \geq \alpha\nu q_j \quad \text{at every adjacent pair of historical indices.} \tag{62}$$

The index range is $1 \leq j < L$ in a finite window and $j \geq 1$ in infinite history. Under (62), $H_{\mathrm{abs}}(D) = \sum_{j \geq 0} |\kappa_j(D)|$ strictly decreases with D. Its finite and infinite limits are respectively (29) and (28), with strong-signal limit one in both cases. If $\alpha = \nu$, the normalized optimum is e_0 for every positive innovation profile and every $D > 0$.

Proof. Normalization and existence. The auxiliary energy (30), its normal equations, and the coefficient representation (43) hold for arbitrary positive innovations. Positivity of a_0 does not require (62). To see this, minimize the finite energy terms beyond index j with x_j fixed. Their minimum is $c_j x_j^2$ for some $c_j \geq 0$. Minimizing the remaining local term with x_{j-1} fixed gives

$$\frac{x_j}{x_{j-1}} = \frac{\nu/q_j + \alpha/s}{1/q_j + 1/s + c_j}, \qquad s = R(1 - \alpha^2). \tag{63}$$

Starting from $x_0 = 1$, this proves $0 < x_j/x_{j-1} \leq M$, where $M = \max\{\nu, \alpha\} < 1$. In particular, $a_0 = (1 - \alpha x_1)/s > 0$.

In infinite history, coercivity follows from (31). The bound $x_j^{(L)} \leq M^j$ permits the same finite-window convergence argument used in Appendix A. Every limiting coordinate is positive: a first zero in (33) would contradict its positive predecessor and nonnegative successor. The coefficient representation therefore yields an ℓ^1 direction with $a_0 > 0$. To identify it with $N^{-1}g$, the covariance row bound in that argument becomes

$$\sup_j \sum_k |K_{jk}| \leq \frac{q_{\max}}{1 - \nu^2} \frac{1 + \nu}{1 - \nu}.$$

Boundedness of N and the coercive residual bound (47) then give $Na = g$ by passing to the limit in each finite covariance equation. Optimality follows from (48).

Sufficiency in finite windows. Let $\Delta = \nu - \alpha \neq 0$ and $m = \min\{\nu, \alpha\}$. Under (62), the denominator $d_j(s, t)$ of the ratio update (35) is positive for $t \in [m, M]$. When $\alpha > \nu$, the required lower bound for its state contribution is

$$\frac{1}{q_j} + \frac{\nu(\nu - \alpha)}{q_{j+1}} = \frac{q_{j+1} - \alpha\nu q_j + \nu^2 q_j}{q_j q_{j+1}} > 0;$$

when $\nu > \alpha$, that contribution is at least $1/q_j$. The noise contribution is positive in both cases.

The endpoint identities (36) continue to hold. The endpoint α can map to itself when $q_{j+1} = \alpha\nu q_j$, whereas the endpoint ν maps strictly toward the interval interior. Since F_j strictly increases in its successor argument, it maps (m, M) into (m, M). The terminal ratio (34) is strictly interior, so backward induction gives $m < \tau_j < M$ at every finite index.

The affine derivative numerator E_j in (38) satisfies

$$E_j(\alpha) = \Delta \left(\frac{1}{q_j} - \frac{\alpha\nu}{q_{j+1}} \right), \qquad E_j(\nu) = \frac{\Delta(1 - \alpha\nu)}{q_j}.$$

The first endpoint has weak sign Δ and the second has strict sign Δ. Thus $E_j(t)$ has strict sign Δ for every $t \in (m, M)$, even at equality in (62). The terminal derivative and backward recursion prove the same strict monotonicity in s as in Appendix A. The coefficient identity (44) gives the strict sign of every interior historical coefficient; the terminal coefficient follows from its ratio.

Every x_j therefore decreases with D when $\Delta > 0$ and increases when $\Delta < 0$. The finite signed-sum formula (50) is strictly increasing in these coordinates. Since $H_{\mathrm{abs}} = H_{\mathrm{sgn}}$ for $\Delta > 0$ and $H_{\mathrm{abs}} = 2 - H_{\mathrm{sgn}}$ for $\Delta < 0$, total absolute history strictly decreases in both cases.

Sufficiency in infinite history. The finite ratio bounds pass to $\tau_j \in [m, M]$. The endpoint ν is excluded by (36). If a limiting ratio equals α, strict increase of F_j in its successor forces that successor to equal α and the local innovation inequality to be an equality. Repeating this implication would give

$$q_{j+n} = (\alpha\nu)^n q_j \longrightarrow 0,$$

contrary to $q_{\min} > 0$. For $\alpha = 0$, the endpoint-map inequality at zero is already strict. Thus all infinite ratios lie in (m, M).

Finite monotonicity passes to weak monotonicity of the limiting ratios. The two-point comparison (46) makes it strict, because E_j has strict sign at each interior successor. The infinite coefficient identity and signed-sum formula (49) now prove the sign law and strict total-history monotonicity. In both finite and infinite history, the energy comparisons with $x_j = \nu^j$ and $x_j = \alpha^j$ give the weak- and strong-signal limits as in Appendix A. The bound $x_j \le M^j$ justifies summation in the infinite case.

Necessity. Partition the current and historical coordinates. Write $r_j = \alpha^j$, $v_j = \nu^j$ for $j \ge 1$, let $K_{\mathrm{h}}, T_{\mathrm{h}}$ be the historical covariance blocks, and put $C_R = K_{\mathrm{h}} + R(T_{\mathrm{h}} - rr^T)$. The two block equations for $Na = g$ give

$$\begin{aligned} a_{\mathrm{h}} &= C_R^{-1}(v - r), & a_0 &= \frac{1}{R} - r^T C_R^{-1}(v - r), \\ \kappa_{\mathrm{h}} &= \frac{RC_R^{-1}(v - r)}{1 - Rr^T C_R^{-1}(v - r)} = RK_{\mathrm{h}}^{-1}(v - r) + O(R^2). \end{aligned} \tag{64}$$

The denominator retains the normalization by the current coefficient. In finite dimension, the expansion follows from $K_{\mathrm{h}} \succ 0$. In infinite dimension, let S be the unilateral shift on historical coordinates. Then

$$K_{\mathrm{h}}^{-1} = (I - \nu S)^* \operatorname{diag}(q_j^{-1})(I - \nu S)$$

is bounded on ℓ^2. A bounded-inverse expansion gives (64) in ℓ^2, hence at every fixed index.

For $w_j = \nu^j - \alpha^j$ and $w_0 = 0$, the innovation difference is $w_j - \nu w_{j-1} = \Delta\alpha^{j-1}$. Applying the tridiagonal state precision yields

$$\begin{aligned} [K_{\mathrm{h}}^{-1}(v - r)]_j &= \Delta\alpha^{j-1}\left(\frac{1}{q_j} - \frac{\alpha\nu}{q_{j+1}}\right), \quad j < L \text{ or infinite history}, \\ [K_{\mathrm{h}}^{-1}(v - r)]_L &= \frac{\Delta\alpha^{L-1}}{q_L}, \qquad\qquad \text{finite history}. \end{aligned} \tag{65}$$

A violation of (62) requires $\alpha > 0$ and reverses the leading sign at the violated index. Equation (64) then gives a coefficient of the wrong sign for all sufficiently large D. This proves necessity. When $\alpha = 0$, the condition holds for every positive profile. Finally, if $\alpha = \nu$, the identity $Ne_0 = Rg$ gives $a = e_0/R$ for every profile. □

B.1 Relative perturbations of the innovations

The criterion provides a quantitative margin for changes in the innovation variances. For a profile satisfying (62), define

$$\eta_*(q) = \inf_j \frac{q_{j+1} - \alpha\nu q_j}{q_{j+1} + \alpha\nu q_j}, \tag{66}$$

where the infimum ranges over adjacent historical indices; for $L = 1$, set $\eta_*(q) = 1$.

Corollary 1 (Relative innovation perturbations)**.** *Let q satisfy the hypotheses and innovation condition of Theorem 6. Suppose a fixed positive profile $\widetilde{q}$ satisfies*

$$\left|\frac{\widetilde{q}_j}{q_j} - 1\right| \le \eta, \qquad 0 \le \eta < 1, \qquad \eta \le \eta_*(q). \tag{67}$$

With the mean difference and AR residual correlation unchanged, the perturbed profile satisfies (62). For $\alpha \ne \nu$, its optimal kernel has the same historical signs for every $D > 0$, and its total absolute history strictly decreases as D increases. For the binary-chain innovations in (22), a sufficient bound uniform over the chain parameters and class weight is

$$\eta \le \frac{1 - \alpha}{1 + \alpha}. \tag{68}$$

Proof. At each adjacent pair,

$$\begin{aligned}\widetilde{q}_{j+1} - \alpha\nu\widetilde{q}_j &\geq (1-\eta)q_{j+1} - \alpha\nu(1+\eta)q_j \\ &= q_{j+1} - \alpha\nu q_j - \eta(q_{j+1} + \alpha\nu q_j) \geq 0.\end{aligned}$$

In infinite history, $\eta < 1$ also preserves the bounds $(1-\eta)q_{\min} \leq \widetilde{q}_j \leq (1+\eta)q_{\max}$. Theorem 6 therefore applies to the fixed perturbed profile. For the binary chain, (24) gives $q_{j+1} > \nu q_j$, and hence

$$\frac{q_{j+1} - \alpha\nu q_j}{q_{j+1} + \alpha\nu q_j} \geq \frac{1-\alpha}{1+\alpha}.$$

Taking the infimum proves (68). □

For $\alpha = 0$, every positive innovation profile satisfies the criterion, including perturbations beyond (67) whenever the theorem's finite or infinite positivity assumptions remain satisfied. The perturbation result concerns innovation variances with the geometric mean difference and AR residual covariance held fixed. It does not assert robustness to arbitrary covariance perturbations or monotonicity along profiles that change with D.

B.2 A SIGN REVERSAL OUTSIDE THE INNOVATION CONDITION

Positive innovations alone do not imply the sign law. Take $L = 2$, $\nu = 3/4$, $\alpha = 2/5$, $q_1 = 1$, $q_2 = 1/20$, and $D = 10$. Here $q_2 < \alpha\nu q_1 = 3/10$. Direct solution of the covariance equations gives

$$\kappa_1 = -\frac{2317}{68255}, \qquad \kappa_2 = \frac{1841}{13651}. \tag{69}$$

The historical coefficients have opposite signs even though $\nu > \alpha$. This example belongs to the Gaussian conditional-moment construction following (61); the stationary binary chain satisfies the stronger inequality (24).

C FINITE-START KERNELS WITH CORRELATED RESIDUALS

A sequence observed from its initialization has a time-dependent reverse transition law. This changes the historical class mean differences, even when its forward transition matrix is homogeneous. We derive the optimal kernel for the available observations $X_0, \ldots, X_t$ and identify its current-only boundary. The residual moment model and class-weighted objective remain those of (2) and (6).

C.1 FINITE-START MOMENTS AND OPTIMALITY

Let $\Pr(z_0 = 1) = p_0 \in (0,1)$, retain $\lambda, \varepsilon > 0$ and $\nu = 1 - \lambda - \varepsilon \in (0,1)$, and write $\pi_1 = \lambda/(\lambda+\varepsilon)$. The marginal error probability and variance at time s are

$$p_s = \pi_1 + (p_0 - \pi_1)\nu^s, \qquad v_s = p_s(1-p_s). \tag{70}$$

Fix $t \geq 1$. Below, lag-indexed quantities depend on this observation time, which we suppress in their notation. Define the mean-gap ratio and the slope of the one-step reverse conditional mean by

$$g_j = \nu^j \frac{v_{t-j}}{v_t} \quad (0 \leq j \leq t), \qquad b_j = \nu \frac{v_{t-j}}{v_{t-j+1}} = \frac{g_j}{g_{j-1}} \quad (1 \leq j \leq t). \tag{71}$$

Thus $g_0 = 1$. Bayes' rule gives the two reverse transition probabilities

$$\begin{aligned}a_{0j} &= \Pr(z_{t-j} = 1 \mid z_{t-j+1} = 0) = \frac{p_{t-j}\varepsilon}{1 - p_{t-j+1}}, \\ a_{1j} &= \Pr(z_{t-j} = 1 \mid z_{t-j+1} = 1) = \frac{p_{t-j}(1-\varepsilon)}{p_{t-j+1}}.\end{aligned} \tag{72}$$

They satisfy $a_{1j} - a_{0j} = b_j \in (0,1)$.

For a class weight $w \in [0,1]$, the mixture of the two classes at time t has reverse-time marginal r_{j-1}. Its expected reverse innovation variance for $S_s = 2z_s - 1$ is q_j, where

$$\begin{aligned} r_{j-1} &= p_{t-j+1} + (w - p_t) g_{j-1}, \\ q_j &= 4\big[(1 - r_{j-1}) a_{0j}(1 - a_{0j}) + r_{j-1} a_{1j}(1 - a_{1j})\big], \\ U_0 &= 0, \qquad U_j = b_j^2 U_{j-1} + q_j. \end{aligned} \tag{73}$$

Here $r_{j-1} \in [0,1]$ and $q_j > 0$, including the endpoint weights $w = 0, 1$. Define matrices on indices $0, \dots, t$ by

$$K_{jk} = \frac{g_{\max(j,k)}}{g_{\min(j,k)}} U_{\min(j,k)}, \qquad (T_\alpha)_{jk} = \alpha^{|j-k|}, \qquad N = K + RT_\alpha, \quad R = \frac{2}{D}. \tag{74}$$

Theorem 7 (Finite-start colored Fisher kernel)**.** *Under (70)–(74), $D > 0$, $0 \le \alpha < 1$, and the path-conditional residual moments in (2), the observations X_{t-j}/D have class mean difference $2g_j$ and weighted within-class covariance $N \succ 0$. The positive-oriented Fisher optimum over signed kernels on $0, \dots, t$ is unique up to positive scale and equals $a = N^{-1}g$, with $a_0 > 0$. It can be computed in $O(t)$ time and normalized as $\kappa_j = a_j/a_0$.*

The optimum over nonnegative kernels is current-only if and only if

$$\alpha \ge \alpha_{\mathrm{crit}}(t), \qquad \alpha_{\mathrm{crit}}(t) = \max_{1 \le j \le t} g_j^{1/j} = \nu \max_{1 \le j \le t} \left(\frac{v_{t-j}}{v_t} \right)^{1/j}. \tag{75}$$

The signed optimum is current-only if and only if $g_j = \alpha^j$ for every $1 \le j \le t$. Both current-only criteria are independent of D and w.

Proof. Conditional moments. For $s \le t$, the forward conditional mean is $\mathbb{E}[z_t \mid z_s] = \pi_1 + \nu^{t-s}(z_s - \pi_1)$, so $\mathrm{Cov}(z_s, z_t) = \nu^{t-s} v_s$. For a binary conditioning variable, covariance divided by its variance is the difference between the two conditional means. Consequently,

$$\mathbb{E}[z_{t-j} \mid z_t = c] = p_{t-j} + (c - p_t) g_j, \qquad c \in \{0, 1\}. \tag{76}$$

This proves the mean-gap formula and, after mixing with weights $1 - w, w$, the expression for r_{j-1}.

The reverse process is Markov. At step j, write

$$S_{t-j} = 2a_{0j} + b_j - 1 + b_j S_{t-j+1} + \eta_j,$$

where the conditional mean of η_j given $z_{t-j+1}, \dots, z_t$ is zero. Its conditional variance is $4a_{cj}(1 - a_{cj})$ when $z_{t-j+1} = c$. The innovations are therefore orthogonal within either class $z_t = c$, and their class-weighted variances are q_j. Centering separately within each current class gives the variance recursion for U_j. Iterating the reverse conditional mean from lag j to lag $k > j$ gives

$$\mathrm{Cov}(S_{t-j}, S_{t-k} \mid z_t = c) = \left(\prod_{\ell=j+1}^{k} b_\ell \right) \mathrm{Var}(S_{t-j} \mid z_t = c).$$

Since the product is g_k/g_j, the class-weighted covariance is K. In particular, $Ke_0 = 0$. On historical coordinates, the innovation representation is an invertible triangular transformation of positive innovation variances, so its principal block K_{h} is positive definite.

The residuals divided by D have covariance RT_α within either class and zero covariance with the states. This follows by conditioning on the entire state path in (2). Hence their addition gives N. The matrix T_α is positive definite: it is the covariance, after division by R, of a stationary residual recursion with initial variance R and innovation variance $R(1 - \alpha^2) > 0$. Thus $N \succ 0$.

Fisher direction. For any nonzero kernel u, Cauchy–Schwarz in the N inner product yields

$$\mathrm{SNR}_w(u)^2 = \frac{4(u^T g)^2}{u^T N u} \le 4 g^T N^{-1} g. \tag{77}$$

Equality holds exactly when u is a scalar multiple of $N^{-1}g$. Since $g^T N^{-1} g > 0$, positive orientation selects positive multiples. The construction below identifies this direction and proves that its current coefficient is positive.

Linear-time construction. Set $s = R(1-\alpha^2)$ and fix $x_0 = 1$. The quadratic problem

$$\min_{x_1,\dots,x_t} \sum_{j=1}^{t} \frac{(x_j - b_j x_{j-1})^2}{q_j} + \frac{1}{s}\sum_{j=1}^{t}(x_j - \alpha x_{j-1})^2 \tag{78}$$

has a unique minimizer because its historical quadratic part is positive definite. Its tridiagonal coefficients are

$$\begin{aligned} B_j &= \frac{b_j}{q_j} + \frac{\alpha}{s}, \\ A_j &= \frac{1}{q_j} + \frac{b_{j+1}^2}{q_{j+1}} + \frac{1+\alpha^2}{s} \quad (1 \le j < t), \\ A_t &= \frac{1}{q_t} + \frac{1}{s}. \end{aligned} \tag{79}$$

Backward elimination followed by a forward pass gives

$$\begin{aligned} \tau_{t-1} = \frac{B_t}{A_t}, \quad \tau_{j-1} = \frac{B_j}{A_j - B_{j+1}\tau_j} \quad (j = t-1,\dots,1), \\ x_0 = 1, \qquad x_j = \tau_{j-1} x_{j-1} \qquad (j = 1,\dots,t). \end{aligned} \tag{80}$$

To see that every denominator is positive and the ratios are bounded, minimize the terms beyond step j with x_j fixed. Their minimum is $P_j x_j^2$ for some $P_j \ge 0$. The minimizing local ratio is therefore

$$\frac{x_j}{x_{j-1}} = \frac{b_j/q_j + \alpha/s}{1/q_j + 1/s + P_j}, \qquad 0 < \frac{x_j}{x_{j-1}} \le \frac{b_j s + \alpha q_j}{s + q_j} < 1. \tag{81}$$

This also proves $0 < x_1 < 1$.

To recover the Fisher coefficients, introduce independent Gaussian innovations in the auxiliary model

$$G_0 = \theta, \qquad G_j = b_j G_{j-1} + u_j, \qquad Y_j = G_j + e_j, \qquad e_j = \alpha e_{j-1} + v_j,$$

with $\mathrm{Var}(u_j) = q_j$, $\mathrm{Var}(e_0) = R$, and $\mathrm{Var}(v_j) = s$. Conditional on deterministic θ, its mean is θg and its covariance is N. Its joint negative log density, up to a constant, is

$$\frac{(Y_0 - \theta)^2}{2R} + \frac{1}{2}\sum_{j=1}^{t}\frac{(G_j - b_j G_{j-1})^2}{q_j} + \frac{1}{2s}\sum_{j=1}^{t}(Y_j - \alpha Y_{j-1} - G_j + \alpha G_{j-1})^2. \tag{82}$$

At $Y = 0$, its minimizing latent path is $G_j = \theta x_j$. Gaussian integration over historical latent coordinates differs from minimization by a log determinant independent of (θ, Y). Differentiating the minimized quadratic in Y therefore identifies the likelihood term linear in θ as

$$\frac{Y_0}{R} + \sum_{j=1}^{t}\frac{x_j - \alpha x_{j-1}}{s}(Y_j - \alpha Y_{j-1}).$$

The direct Gaussian likelihood has linear term $g^T N^{-1} Y$. Equating coefficients gives

$$\begin{aligned} a_0 &= \frac{1 - \alpha x_1}{s}, \\ a_j &= \frac{(1+\alpha^2)x_j - \alpha x_{j-1} - \alpha x_{j+1}}{s} \quad (1 \le j < t), \\ a_t &= \frac{x_t - \alpha x_{t-1}}{s}, \qquad \kappa_j = \frac{a_j}{a_0}. \end{aligned} \tag{83}$$

Thus these coefficients equal $N^{-1}g$, and $a_0 > 0$ follows from (81). Gaussianity evaluates these quadratic forms. The observation model itself requires only the stated moments. Computing the marginal, innovation, backward, and forward passes takes $O(t)$ operations and $O(t)$ storage. At $t = 0$, the normalized kernel is directly $\kappa_0 = 1$.

Current-only boundary. All g_j are positive. Normalize a nonnegative kernel u by $u^T g = 1$, and set $d = u - e_0$. Since $Ne_0 = R(1, \alpha, \ldots, \alpha^t)^T$,

$$u^T N u - R = 2R \sum_{j=1}^{t} u_j(\alpha^j - g_j) + d^T N d. \tag{84}$$

If $\alpha^j \geq g_j$ for every lag, the right-hand side is positive unless $u = e_0$. Conversely,

$$\frac{d}{dc}\mathrm{SNR}_w(e_0 + ce_j)^2\Big|_{c=0} = \frac{8}{R}(g_j - \alpha^j). \tag{85}$$

A violated inequality therefore gives an improving positive perturbation. This proves (75). Over signed kernels, any unequal pair g_j, α^j gives an improving perturbation of one sign. If they agree at every lag, $Ne_0 = Rg$, and the equality condition in (77) proves exact cancellation. This is the finite-start specialization of Theorem 5. □

C.2 Score normalization and class midpoints

The normalization $\kappa_0 = 1$ fixes the arbitrary positive scale of the Fisher direction. To align its origin across observation times, let ℓ_s denote the working log-likelihood-ratio score on the normalized observation scale, so that $\ell_s = X_s$ under the moment model. Equation (76) gives the midpoint of its two current-class conditional means at lag j:

$$m_{t,j} = \frac{1}{2}\big(\mathbb{E}[X_{t-j} \mid z_t = 0] + \mathbb{E}[X_{t-j} \mid z_t = 1]\big) = D\big[2p_{t-j} - 1 + g_j(1 - 2p_t)\big]. \tag{86}$$

The resulting score is

$$h_t = \sum_{j=0}^{t} \kappa_j\big(\ell_{t-j} - m_{t,j}\big). \tag{87}$$

Under the model its two class means are $\pm D\sum_j \kappa_j g_j$. Midpoint subtraction is deterministic at a fixed t and leaves both the class mean difference and the within-class variances unchanged. It therefore preserves the same-time Fisher objective. It supplies an origin convention when comparing scores across positions. The theorem establishes optimality separately at each observation time.

C.3 Initialization effects and stationary limits

The variance ratio in (75) determines how initialization shifts the current-only boundary. If $v_0 < \cdots < v_t$, then $\alpha_{\mathrm{crit}}(t) < \nu$; if $v_0 > \cdots > v_t$, then $\alpha_{\mathrm{crit}}(t) > \nu$. For example, variance increases when p_0 and π_1 lie on the same side of $1/2$ and π_1 is strictly closer to $1/2$.

An exact early-time example separates the finite-start and stationary predictions. Set

$$p_0 = \lambda = \varepsilon = \frac{1}{10}, \qquad \nu = \frac{4}{5}, \qquad \alpha = \frac{3}{5}, \qquad w = \frac{1}{2}, \qquad D = 1. \tag{88}$$

At $t = 1$, $p_1 = 9/50$, $g_1 = 20/41$, and the reverse probabilities are $a_{01} = 1/82$ and $a_{11} = 1/2$. Thus $q_1 = 881/1681$, and the two-coordinate covariance equations give

$$\kappa_1 = \frac{R(g_1 - \alpha)}{q_1 + R(1 - \alpha g_1)} = -\frac{1886}{16295}. \tag{89}$$

The nonnegative optimum is current-only, while the signed optimum subtracts the previous observation. The stationary theorem predicts positive history because $\alpha < \nu$. Changing the residual persistence to $\alpha = 20/41$ gives exact signed cancellation at $t = 1$ for every D, w, still with $\alpha < \nu$.

At $t = 3$, keeping all parameters in (88), the kernel instead has mixed historical signs:

$$\kappa \approx (1,\ 0.11522975,\ 0.07247532,\ -0.06533947). \tag{90}$$

The constant reverse slope in Appendix A is essential to its sign argument. Appendix B also retains a constant slope and geometric mean gap while varying the innovation profile. Neither its common-sign conclusion nor its total-absolute-history monotonicity follows by replacing ν with the varying b_j.

Stationarity is recovered exactly at $p_0 = \pi_1$: then $g_j = \nu^j$, $b_j = \nu$, and (73) reduces to (22). More generally, as $t \to \infty$, the moments on any fixed lag window converge to their stationary counterparts. For fixed D, α, w, positive definiteness then gives convergence of the optimal kernel computed using only that fixed window.

The threshold for the full available history also converges:

$$\alpha_{\mathrm{crit}}(t) \longrightarrow \nu. \tag{91}$$

To prove this, v_s is bounded above and away from zero and converges to $\pi_1(1-\pi_1)$. There is a finite constant C such that $|\log(v_s/v_u)| \le C$ for all $s, u \ge 0$. For lags $j \ge J$, the logarithm of the variance-ratio factor in (75) has absolute value at most C/J. For $1 \le j < J$, it tends uniformly to zero as $t \to \infty$, since this is a finite set of lags. Taking $t \to \infty$ and then $J \to \infty$ proves (91). Hence $\alpha < \nu$ eventually permits a strict nonnegative-history improvement, whereas $\alpha > \nu$ eventually gives a current-only nonnegative optimum. This threshold convergence does not assert convergence of every coefficient when the window itself grows.

For completeness, finite-horizon signal-strength limits remain available without a common-sign law. At fixed t,

$$\lim_{D\downarrow 0} \kappa = \frac{T_\alpha^{-1} g}{(T_\alpha^{-1} g)_0}, \qquad \lim_{D\to\infty} \sum_{j=0}^{t} |\kappa_j| = 1. \tag{92}$$

The first follows by dividing N by R and taking $R \to \infty$; its normalizing entry is $(1-\alpha g_1)/(1-\alpha^2) > 0$. For the second, write $r = (\alpha, \dots, \alpha^t)^T$ and let T_{h} be the historical residual block. The block equations for $Na = g$ give

$$a_{\mathrm{h}} = [K_{\mathrm{h}} + R(T_{\mathrm{h}} - rr^T)]^{-1}(g_{\mathrm{h}} - r), \qquad Ra_0 = 1 - Rr^T a_{\mathrm{h}}.$$

Since $K_{\mathrm{h}} \succ 0$, a_{h} has a finite limit and $Ra_0 \to 1$ as $R \downarrow 0$. Thus every normalized historical coefficient tends to zero. These endpoint limits do not impose monotonicity between them.

Finite-path verification. We checked (71)–(74) by enumerating all binary paths for 720 parameter cases with horizons $t = 0, \dots, 5$. Direct conditioning on the terminal state gave a maximum absolute discrepancy of 7.364×10^{-14} in the mean-gap and covariance entries. Comparing (80)–(83) with dense solves using these enumerated moments gave a maximum coefficient error, divided by one plus the largest absolute dense coefficient, of 7.154×10^{-14}. For 180 cases, enumeration of every nonempty coordinate face of the nonnegative Fisher problem also verified (75). The checks include $w = 0, 1, 1/2$, interior weights, white and colored residuals, and signal strengths sampled logarithmically between 0.01 and 100. The reference implementation is `scripts/finite_colored.py`.

D Exponential and white-noise results, and finite-start recursions

D.1 Mean separation and path uncertainty

Additional history carries a weaker version of the current-state signal because the chain can switch between the past observation and the present. Its contribution can be calculated exactly.

Lemma 1 (Stationary conditional moments)**.** *Under* (1)–(3)*,*

$$\Delta_h(\rho) = \frac{2D}{1-\rho\nu}, \qquad \mathrm{Var}(h_t \mid z_t = i) = \frac{2D}{1-\rho^2} + D^2 V_i(\rho), \tag{93}$$

where $V_i = \mathrm{Var}(\sum_{k\ge 0} \rho^k S_{t-k} \mid z_t = i)$*. These path variances uniquely solve*

$$\begin{aligned} V_1 &= \rho^2 \left[(1-\varepsilon) V_1 + \varepsilon V_0 + \frac{4\varepsilon(1-\varepsilon)}{(1-\rho\nu)^2} \right], \\ V_0 &= \rho^2 \left[\lambda V_1 + (1-\lambda) V_0 + \frac{4\lambda(1-\lambda)}{(1-\rho\nu)^2} \right]. \end{aligned} \tag{94}$$

Reversibility gives a conditional mean difference $2D\nu^k$ for the observation at lag k; summing yields (93). The variance equations follow by conditioning on the first backward state. The observation-noise contribution scales as D, whereas uncertainty about the latent path scales as D^2. Thus increasing observation quality changes the relative cost of including obsolete evidence.

Mean separation increases with ρ but approaches the finite ceiling $2D/(1-\nu)$. This is a statement about accumulation of a fixed observation stream. It does not describe how changing a language model's input context changes the observation distribution.

D.2 A MEMORY LAW FOR EVERY FIXED CLASS WEIGHT

Let $V_w = wV_1 + (1-w)V_0$. Maximizing (6) is equivalent to minimizing

$$Q_{D,w}(\rho) = A(\rho) + DB_w(\rho), \qquad A(\rho) = \frac{2(1-\rho\nu)^2}{1-\rho^2},$$
$$B_w(\rho) = (1-\rho\nu)^2 V_w(\rho) = \frac{4\rho^2(a_w - b_w\rho^2)}{(1-\rho^2)(1-\nu\rho^2)}, \tag{95}$$

where

$$\begin{aligned} a_w &= w\varepsilon(1-\varepsilon) + (1-w)\lambda(1-\lambda), \\ b_w &= \nu\{w\varepsilon(1-\lambda) + (1-w)\lambda(1-\varepsilon)\}. \end{aligned} \tag{96}$$

The optimum balances decreasing observation-noise cost up to $\rho = \nu$ against increasing path uncertainty.

Theorem 8 (Class-conditional memory law)**.** *Fix $\lambda, \varepsilon > 0$ with $\lambda + \varepsilon < 1$, and fix $w \in [0,1]$. For every $D > 0$, SNR_w has a unique maximizer $\rho_w^\star(D) \in (0,\nu)$. It is smooth and strictly decreasing in D, with*

$$\lim_{D\downarrow 0} \rho_w^\star(D) = \nu, \qquad \lim_{D\to\infty} \rho_w^\star(D) = 0, \qquad \rho_w^\star(D) \sim \frac{2\nu}{c_w D}, \quad c_w = 4a_w. \tag{97}$$

These conclusions hold for every fixed weight within the stated family.

Proof outline. The key identity is $a_w - b_w = \lambda\varepsilon(1+\nu) > 0$. The power-series expansion of B_w has strictly positive coefficients, giving $B_w'(\rho) > 0$ for $\rho > 0$ and $B_w''(\rho) > 0$. Direct differentiation also gives $A'' > 0$, $A'(0) = -4\nu$, and $A'(\nu) = 0$. Therefore $Q_{D,w}$ is strictly convex with opposite derivative signs at 0 and ν, giving a unique interior minimizer. Its derivative with respect to D is

$$\frac{d\rho_w^\star}{dD} = -\frac{B_w'(\rho_w^\star)}{A''(\rho_w^\star) + DB_w''(\rho_w^\star)} < 0.$$

Comparison with $Q(\nu)$ and $Q(0)$ yields the two limits. Finally, $A'(0) = -4\nu$ and $B_w''(0) = 8a_w$ give the strong-signal rate from the first-order condition. Appendix E gives the complete argument. □

The effective geometric history $1/(1-\rho_w^\star)$ decreases from H_ν to one as the observations improve. A memoryless detector is never exactly SNR-optimal at finite D in this model, even though its optimality gap vanishes in the strong-signal limit. The exact optimizer and rate constant depend on w; the monotonic direction and endpoints do not. Dropping path uncertainty would remove the DB_w term and leave $\rho = \nu$ optimal at every signal level, which identifies uncertainty about past states as the mechanism behind decreasing optimal memory.

D.3 BEYOND EXPONENTIAL MEMORY: THE OPTIMAL CAUSAL LINEAR KERNEL

The same shortening principle holds without constraining the filter to be geometric. Divide observations by D and center them: $Y_j = X_{t-j}/D - (2\pi_1 - 1)$, with residual variance $R = 2/D$. For any stable (absolutely summable) causal score $H = \sum_{j\geq 0} u_j Y_j$, extend SNR_w by using its conditional mean gap and weighted conditional variance. Write $g_j = \nu^j$ and let N_w be the weighted within-class covariance matrix of these observations. Its entries have the form

$$(N_w)_{jk} = \nu^{|j-k|} U_{w,\min(j,k)} + R\mathbf{1}_{\{j=k\}},$$

where $U_{w,j}$ is the weighted conditional variance of S_{t-j} given z_t. Its innovation profile is

$$q_{w,j} = U_{w,j} - \nu^2 U_{w,j-1} = 4\pi_0\pi_1(1-\nu^2) + 4(\varepsilon-\lambda)(w-\pi_1)\nu^j > 0. \tag{98}$$

Theorem 9 (All-causal linear memory law)**.** *Under (3), every fixed $w \in [0,1]$ has a unique SNR-optimal stable causal linear kernel up to positive scale. Normalize its current coefficient to $\kappa_{w,0} = 1$. Its coefficients are strictly positive and satisfy*

$$J_{w,j} = \frac{D}{2} + \frac{\nu^2 J_{w,j+1}}{1 + q_{w,j+1} J_{w,j+1}}, \qquad \frac{\kappa_{w,j}}{\kappa_{w,j-1}} = \frac{\nu}{1 + q_{w,j} J_{w,j}}, \quad j \geq 1, \tag{99}$$

The information recurrence also defines $J_{w,0}$ and has a unique bounded nonnegative solution. Every adjacent coefficient ratio strictly decreases with D, from ν to zero. The normalized total history weight $H_w(D) = \sum_{j\geq 0} \kappa_{w,j}(D)$ strictly decreases, with

$$H_w(D) \longrightarrow \frac{1}{1-\nu} \; (D \downarrow 0), \qquad H_w(D) - 1 \sim \frac{2\nu}{c_w D} \; (D \to \infty). \tag{100}$$

The entire kernel is geometric if and only if $w = \pi_1$ or $\lambda = \varepsilon$. All weights share the stationary Wiener decay (101) asymptotically in the remote tail.

The covariance-inverse Fisher rule gives the optimal direction $N_w^{-1} g$ (Murphy, 2022). A moment-matched Gaussian state model computes that direction by the backward information recursion in (99), without requiring Gaussian observations. Increasing D strictly increases every $J_{w,j}$, which shortens every adjacent coefficient ratio. Appendix F proves the infinite-history construction, uniqueness, and limits, and gives an $O(L)$ algorithm for a finite lag window. Class weighting changes the recent part of the kernel, while its remote tail has a common decay. The optimized exponential approximation has the same leading strong-signal history weight; its numerical SNR loss is small on the fixed grid used for covariance verification.

D.4 The stationary weighting recovers causal linear estimation

One weight has a canonical estimation interpretation. Set $w = \pi_1$, $v = 4\pi_0\pi_1$, $c = v(1-\nu^2)$, and $C = 1 + \nu^2 + Dc/2$.

Proposition 2 (Wiener specialization)**.** *For stationary class weighting,*

$$\rho^\star_{\pi_1} = \frac{2\nu}{C + \sqrt{C^2 - 4\nu^2}}. \tag{101}$$

Up to positive scaling, this exponential kernel also maximizes stationary-weight SNR among stable causal linear filters with positive conditional mean separation. After optimal affine calibration, it is the unique linear minimum-mean-square estimate of the current state.

At this weight, $b_w = \nu a_w$ and $4a_w = c$. Maximizing SNR is equivalent to minimizing affine estimation error by total variance. Appendix G verifies the Wiener–Hopf orthogonality equations directly. This specialization and the estimation interpretation are classical (Särkkä & Svensson, 2023); Theorem 9 characterizes how other weights change the kernel. Balanced SNR maximizes a distribution-free AUROC lower bound, although it need not maximize AUROC itself (Appendix G).

D.5 When correlated residuals remove the benefit of memory

The white-residual assumption has a sharp consequence. Replace it by a stationary, zero-mean residual process independent of the state path, with $\mathrm{Cov}(\xi_t, \xi_{t-k}) = 2D\alpha^{|k|}$ for $0 \leq \alpha < 1$. The mean separation and path variance are unchanged, but the objective becomes

$$Q_{D,w,\alpha}(\rho) = \frac{2(1-\rho\nu)^2(1+\alpha\rho)}{(1-\rho^2)(1-\alpha\rho)} + DB_w(\rho). \tag{102}$$

Proposition 3 (A boundary for useful exponential memory)**.** *Under this geometric residual covariance model, if $\alpha \geq \nu$, then $\rho = 0$ uniquely maximizes SNR_w over $[0,1)$ for every $D > 0$ and every $w \in [0,1]$.*

Thus persistent residual noise can make positive exponential memory unhelpful even when individual observations are weak. If $\alpha < \nu$, every choice of global optima strictly decreases across signal levels, and its strong-signal rate is $2(\nu - \alpha)/(c_w D)$. Appendix H proves these claims and gives the weak-signal limit. Fixed-signal uniqueness in this extension is established at stationary weighting. This extension concerns a specified covariance model, rather than identifying it from a single empirical residual correlation. This boundary does not preclude improvements from signed filters that suppress temporally persistent observation noise.

D.6 Finite starts: exact moments and time-varying memory

For $h_0 = 0$ and an arbitrary initial state law, the stationary backward transition no longer gives finite-start moments. Proposition 5 in Appendix I propagates state probabilities and joint first and second moments forward, giving exact conditional means and variances in $O(t)$ time per memory value. At the first step, all memory values produce the same score. The appendix also gives finite-history approximation and objective-perturbation bounds.

An estimator can also adapt its gain over time. For $Y_t = m_0 + \delta z_t + \eta_t$, with common residual variance R, let $p_t = \Pr(z_t = 1)$ and $q_t = p_t(1 - p_t) - \nu^2 p_{t-1}(1 - p_{t-1})$. The best affine linear estimate follows

$$\begin{aligned} \widehat{z}_t^- &= \lambda + \nu \widehat{z}_{t-1}, & C_t^- &= \nu^2 C_{t-1} + q_t, \\ K_t &= \frac{\delta C_t^-}{\delta^2 C_t^- + R}, & \widehat{z}_t &= \widehat{z}_t^- + K_t(Y_t - m_0 - \delta \widehat{z}_t^-), \\ C_t &= (1 - K_t\delta) C_t^-. \end{aligned} \tag{103}$$

Initialize $\widehat{z}_0 = p_0$ and $C_0 = p_0(1 - p_0)$. The effective coefficient on the previous estimate is $\nu(1 - K_t\delta)$, so finite-start filtering uses a time-varying memory. This classical linear-projection recursion is valid for the binary chain because $z_t - \lambda - \nu z_{t-1}$ is orthogonal to past observations; Gaussian states are unnecessary. Here C_t is unconditional linear estimation error, not a random posterior variance. Appendix I proves the recursion and its connection to (101).

E Proofs of the stationary memory law

E.1 Existence of the stationary sum

Under (3), $\mathbb{E}[X_t^2] = D^2 + 2D$. In stationarity, $\mathrm{Var}(X_t) = 2D + 4D^2\pi_0\pi_1$; the term $2D$ alone is the conditional observation-noise variance. Minkowski's inequality gives

$$\left\| \sum_{k=m}^{n} \rho^k X_{t-k} \right\|_2 \le \sqrt{D^2 + 2D} \sum_{k=m}^{n} \rho^k \longrightarrow 0$$

as $m, n \to \infty$. Thus the series defines an L^2 random variable. It also converges absolutely almost surely because $\mathbb{E}\sum_{k\ge 0} \rho^k |X_{t-k}| = \mathbb{E}|X_0|/(1 - \rho) < \infty$. These arguments do not require independent observations.

E.2 Proof of Lemma 1

Detailed balance holds because $\pi_0\lambda = \pi_1\varepsilon$. Therefore $\Pr(z_{t-k} = j \mid z_t = i) = P_{ij}^k$. Writing Π for the matrix with both rows equal to π, $P = \Pi + \nu(I - \Pi)$ and $P^k = \Pi + \nu^k(I - \Pi)$. In particular,

$$P_{11}^k = \pi_1 + \pi_0\nu^k, \qquad P_{01}^k = \pi_1(1 - \nu^k).$$

The two conditional means of X_{t-k} differ by $2D\nu^k$. The L^2 convergence above licenses conditional expectation and geometric summation, proving $\Delta_h = \sum_{k\ge 0} 2D(\rho\nu)^k = 2D/(1 - \rho\nu)$.

Let $H_t = \sum_{k\ge 0} \rho^k S_{t-k}$. Conditional on the whole state path, the residual sum has mean zero and variance $2D/(1-\rho^2)$, so total variance gives $\mathrm{Var}(h_t \mid z_t = i) = 2D/(1-\rho^2) + D^2\mathrm{Var}(H_t \mid z_t = i)$. Given z_t, the current S_t is fixed. In the recursion $H_t = S_t + \rho H_{t-1}$, condition on z_{t-1}. By the

Markov property, the conditional distribution of the earlier past given (z_{t-1}, z_t) depends only on z_{t-1}. The difference between the two conditional means of H_{t-1} is $2/(1-\rho\nu)$. Applying total variance yields (94).

Solving that system gives

$$V_1(\rho) = \frac{4\varepsilon\rho^2[\lambda(1+\nu\rho^2)+\nu(1-\rho^2)]}{(1-\rho^2)(1-\nu\rho)^2(1-\nu\rho^2)},$$
$$V_0(\rho) = \frac{-4\lambda\rho^2[\lambda\nu\rho^2+\lambda+\nu^2\rho^2-1]}{(1-\rho^2)(1-\nu\rho)^2(1-\nu\rho^2)}.$$

The system is invertible because its coefficient matrix is $I-\rho^2 P$ up to state ordering, whose eigenvalues are $1-\rho^2$ and $1-\nu\rho^2$. Substitution and collection of terms give (95)–(96).

E.3 Proof of Theorem 8

Since the conditional mean gap is positive, $\mathrm{SNR}_w^2 = 4D/Q_{D,w}$. Hence minimizing Q is equivalent to maximizing SNR. Direct substitution of $\nu = 1-\lambda-\varepsilon$ gives

$$a_w - b_w = \lambda\varepsilon(1+\nu) > 0, \qquad a_w > 0.$$

For $|\rho| < 1$, write $y = \rho^2$ and $c_k = \sum_{j=0}^{k}\nu^j$. Expanding the denominator in (95),

$$B_w/4 = a_w y + \sum_{m\geq 2}[(a_w - b_w)c_{m-2} + a_w\nu^{m-1}]y^m.$$

Every coefficient is positive. Termwise differentiation is valid on compact subsets of $(-1,1)$, giving $B_w'(\rho) > 0$ on $(0,1)$ and $B_w''(\rho) > 0$ on $[0,1)$. In particular $B_w(0) = B_w'(0) = 0$ and $B_w''(0) = 8a_w$.

Direct differentiation gives

$$A'(\rho) = \frac{4(\rho-\nu)(1-\rho\nu)}{(1-\rho^2)^2}, \qquad A''(\rho) = \frac{4[(1-\nu)^2(1+3\rho^2)+2\nu(1-\rho)^3]}{(1-\rho^2)^3} > 0.$$

Thus $Q'' = A'' + DB_w'' > 0$. Since $Q'(0) = -4\nu < 0$ and $Q'(\nu) = DB_w'(\nu) > 0$, Q has a unique stationary point in $(0,\nu)$. Strict convexity makes it the unique global minimizer over $[0,1)$. Smoothness follows from the implicit function theorem because Q'' is strictly positive. Differentiating $Q'(\rho_w^\star) = 0$ gives the formula and sign in the main proof outline.

For the weak-signal limit, all optimizers lie in the compact interval $[0,\nu]$, and

$$A(\rho_w^\star) \leq Q_{D,w}(\rho_w^\star) \leq A(\nu) + DB_w(\nu) \longrightarrow A(\nu).$$

A has its unique minimum at ν, so every convergent subsequence of optimizers tends to ν. For the strong-signal limit,

$$DB_w(\rho_w^\star) \leq Q_{D,w}(\rho_w^\star) \leq Q_{D,w}(0) = 2.$$

On $[0,\nu]$, B_w vanishes only at zero. Therefore $\rho_w^\star \to 0$. Using the first-order condition once more,

$$D\rho_w^\star = -\frac{A'(\rho_w^\star)}{B_w'(\rho_w^\star)/\rho_w^\star} \longrightarrow \frac{4\nu}{8a_w} = \frac{2\nu}{c_w}.$$

These limits keep the chain and variance weight fixed. They make no claim about joint limits in which the chain approaches a boundary at the same time as D.

F Optimal causal kernels for arbitrary class weights

We prove Theorem 9 and derive the optimal SNR and tail behavior. Fix w and suppress its subscript on q_j, J_j, and κ_j in this appendix. Set $m = 2\pi_1 - 1$, $v = 4\pi_0\pi_1$, and $R = 2/D$. For lag $j \geq 0$, normalize and center the observations as

$$Y_j = X_{t-j}/D - m, \qquad g_j = \nu^j.$$

Their class-conditional mean difference is $2g_j$ and their residual variance is R. The unconditional covariance is $v\nu^{|j-k|} + R\mathbf{1}_{\{j=k\}}$, a bounded Toeplitz operator on ℓ^2. Thus every square-summable coefficient sequence defines an L^2 limit of centered causal scores.

F.1 Weighted covariance and the Fisher optimum

Define the weighted conditional state variance at lag j by

$$U_j = w\mathrm{Var}(S_{t-j} \mid z_t = 1) + (1-w)\mathrm{Var}(S_{t-j} \mid z_t = 0), \qquad U_0 = 0.$$

For $j \le k$, reversibility and the Markov property give

$$\mathbb{E}[S_{t-k} \mid S_{t-j}, z_t] = m + \nu^{k-j}(S_{t-j} - m).$$

Consequently the weighted within-class covariance of the observations is

$$N_{jk} = \nu^{|j-k|}U_{\min(j,k)} + R\mathbf{1}_{\{j=k\}}. \tag{104}$$

The transition probabilities $p_{11}(j) = \pi_1 + \pi_0\nu^j$ and $p_{01}(j) = \pi_1 - \pi_1\nu^j$ yield

$$\begin{aligned} U_j &= 4\{wp_{11}(j)[1-p_{11}(j)] + (1-w)p_{01}(j)[1-p_{01}(j)]\} \\ &= v(1-\nu^{2j}) - 4m(w-\pi_1)(\nu^j - \nu^{2j}). \end{aligned} \tag{105}$$

Taking $q_j = U_j - \nu^2 U_{j-1}$ proves (98). Positivity also follows directly by initializing the reverse chain with class probability w:

$$q_j = 4\mathbb{E}_w[(1-z_{j-1})\lambda(1-\lambda) + z_{j-1}\varepsilon(1-\varepsilon)] \ge q_{\min} > 0,$$

where $q_{\min} = 4\min\{\lambda(1-\lambda), \varepsilon(1-\varepsilon)\}$.

Because $0 \le U_j \le 1$, the state-covariance part of N has absolute row sums at most $(1+\nu)/(1-\nu)$. Hence N is bounded and self-adjoint on ℓ^2, and $N \succeq RI$. For any nonzero coefficient sequence a, its squared SNR is

$$\mathrm{SNR}_w(a)^2 = \frac{4\langle a, g\rangle^2}{\langle a, Na\rangle}.$$

Cauchy–Schwarz in the N inner product gives

$$\mathrm{SNR}_w(a)^2 \le 4\langle g, N^{-1}g\rangle, \qquad a \propto N^{-1}g. \tag{106}$$

Equality holds exactly in this direction. For positive mean separation, the scaling constant is positive. It remains to compute this direction and show that it is itself a stable causal kernel.

F.2 Finite information recursion and its infinite limit

Consider first lags $0, \dots, L$. An auxiliary Gaussian model represents exactly the moments in (104):

$$x_0 = \theta, \qquad x_j = \nu x_{j-1} + u_j, \qquad y_j = x_j + e_j,$$

where θ is deterministic, the u_j are independent zero-mean Gaussians with variance q_j, and the e_j are independent zero-mean Gaussians with variance R and independent of the u_j. The covariance recursion gives $\mathrm{Var}(x_j \mid \theta) = U_j$, so the observation mean and covariance are θg and N restricted to these lags. Gaussianity is used here to evaluate this finite quadratic form, not as an assumption on the original observations.

Write the likelihood of $y_j, \dots, y_L$ as a function of x_j as $\exp(-J_j^{(L)}x_j^2/2 + \ell_j^{(L)}x_j + \mathrm{const})$. Completing the square when integrating x_{j+1} gives

$$\begin{aligned} J_j^{(L)} &= \frac{1}{R} + \frac{\nu^2 J_{j+1}^{(L)}}{1 + q_{j+1}J_{j+1}^{(L)}}, \\ \ell_j^{(L)} &= \frac{y_j}{R} + \frac{\nu\ell_{j+1}^{(L)}}{1 + q_{j+1}J_{j+1}^{(L)}}, \end{aligned} \tag{107}$$

with $J_L^{(L)} = 1/R$ and $\ell_L^{(L)} = y_L/R$. Expanding the linear coefficient at lag zero yields

$$a_0^{(L)} = 1/R, \qquad \frac{a_j^{(L)}}{a_{j-1}^{(L)}} = \frac{\nu}{1 + q_j J_j^{(L)}} \quad (1 \le j \le L). \tag{108}$$

The same likelihood computed directly from its covariance has quadratic coefficient $g^\top N^{-1} g$ and linear coefficient $g^\top N^{-1} y$. Thus $a^{(L)} = N^{-1} g$ on the finite index set and $J_0^{(L)} = g^\top N^{-1} g$. The backward information pass followed by the coefficient pass costs $O(L)$ operations.

For fixed j, adding older observations increases $J_j^{(L)}$. The recursion also gives the uniform bounds

$$\frac{1}{R} \le J_j^{(L)} \le \frac{1}{R(1-\nu^2)}. \tag{109}$$

Therefore $J_j = \lim_{L\to\infty} J_j^{(L)}$ exists and satisfies (99). The update map has derivative $\nu^2/(1 + q_{j+1}J)^2 \le \nu^2 < 1$ in its nonnegative argument. Taking sup norms shows that this is the unique bounded nonnegative solution.

Let $a_0 = 1/R$ and take the limit of (108). Every adjacent ratio belongs to $(0, \nu)$, and $0 < a_j \le \nu^j/R$. After extending each finite coefficient sequence by zero, these bounds give convergence to a in ℓ^1 and ℓ^2. The finite normal equations pass to the limit because N is bounded, giving $Na = g$. Thus the infinite kernel attains (106), and normalizing by a_0 gives the coefficients κ_j in (99). In particular, the optimizer is strictly positive and absolutely summable. The information identity also passes to the limit:

$$\max_a \mathrm{SNR}_w(a)^2 = 4J_0. \tag{110}$$

F.3 Dependence on observation quality

Fix the chain and class weight, so the innovation profile does not depend on D. The update in (107) increases in its successor and has additive term $D/2$. For $D_2 > D_1$, backward induction and then the infinite-history limit give

$$J_j(D_2) - J_j(D_1) \ge \frac{D_2 - D_1}{2} > 0.$$

Since $q_j > 0$, every adjacent ratio $\kappa_j/\kappa_{j-1} = \nu/(1 + q_j J_j)$ strictly decreases with D. The bound (109) implies $J_j(D) \to 0$ as $D \downarrow 0$. For strong signal, the recursion gives

$$\frac{D}{2} \le J_j(D) \le \frac{D}{2} + \frac{\nu^2}{q_{j+1}},$$

and hence

$$\lim_{D\downarrow 0} \frac{\kappa_j}{\kappa_{j-1}} = \nu, \qquad \frac{\kappa_j}{\kappa_{j-1}} \sim \frac{2\nu}{q_j D} \quad (D \to \infty). \tag{111}$$

These statements hold at every fixed positive lag.

Each κ_j for $j \ge 1$ is a product of strictly decreasing positive ratios, while $\kappa_0 = 1$ and $\kappa_j \le \nu^j$. Consequently $H_w(D) = \sum_{j\ge 0} \kappa_j(D)$ strictly decreases, and dominated convergence gives its limits $1/(1-\nu)$ and 1. Moreover, uniformly over lags, $\kappa_j/\kappa_{j-1} \le \nu/(1 + q_{\min} D/2)$. The mass beyond lag one is therefore $O(D^{-2})$. Since $q_1 = 4a_w = c_w$, the leading excess history is

$$H_w(D) - 1 \sim \kappa_1(D) \sim \frac{2\nu}{c_w D}. \tag{112}$$

This also identifies the strong-signal constant in the exponential-family theorem as the first reverse-time innovation variance.

F.4 Geometric tails and exact geometric kernels

Write $q_j = c + \beta\nu^j$, where $c = v(1-\nu^2)$ and $\beta = 4(\varepsilon - \lambda)(w - \pi_1)$. Let J_∞ be the unique positive solution of

$$J_\infty = 1/R + \frac{\nu^2 J_\infty}{1 + cJ_\infty}.$$

Comparing this equation with the recursion for J_j gives

$$|J_j - J_\infty| \le \nu^2 |J_{j+1} - J_\infty| + \nu^2 J_\infty^2 |\beta| \nu^{j+1}.$$

Iteration removes the bounded remainder and yields

$$|J_j - J_\infty| \le \frac{J_\infty^2 |\beta| \nu^{j+3}}{1-\nu^3}. \tag{113}$$

Thus every weight has the same limiting adjacent ratio

$$r_\infty = \frac{\nu}{1+cJ_\infty} = \frac{2\nu}{C+\sqrt{C^2-4\nu^2}}, \qquad C = 1+\nu^2+Dc/2. \tag{114}$$

The fixed-point equation gives the quadratic $\nu r_\infty^2 - C r_\infty + \nu = 0$; its root in $(0,\nu)$ is (114). The bound (113) also implies $\kappa_j/\kappa_{j-1} = r_\infty + O(\nu^j)$. The resulting relative errors are summable, so their positive infinite product converges. Hence $\kappa_j \sim A_w r_\infty^j$ for some $A_w > 0$ at each fixed D.

If $(\varepsilon - \lambda)(w - \pi_1) = 0$, the innovation profile is constant and the unique bounded information solution is constant. The full kernel is then geometric. Conversely, suppose all adjacent ratios equal r. Then $q_j J_j = \nu/r - 1$ for every $j \ge 1$, and the information recursion becomes $J_j = 1/R + \nu r J_{j+1}$. Its only bounded solution is constant, as follows by repeated substitution and $\nu r < 1$. Therefore every q_j is constant, forcing $(\varepsilon - \lambda)(w - \pi_1) = 0$. This proves the geometricity criterion and completes Theorem 9.

G THE STATIONARY-WEIGHT LINEAR ESTIMATOR

At $w = \pi_1$, direct algebra yields $4a_w = c = 4\pi_0\pi_1(1-\nu^2)$ and $b_w = \nu a_w$. The objective becomes $Q = [2(1-\rho\nu)^2 + Dc\rho^2]/(1-\rho^2)$. Its first-order equation is

$$\nu\rho^2 - C\rho + \nu = 0, \qquad C = 1+\nu^2+Dc/2.$$

The polynomial is positive at zero and negative at ν, where its value is $-Dc\nu/2$. The two roots multiply to one, so the smaller root lies in $(0,\nu)$ and the larger root exceeds one. Rationalizing the smaller root gives (101).

We use "stable causal linear filter" for an absolutely summable coefficient sequence. The projection argument below also covers the closed linear span of the observations. For a direct all-causal proof, center and rescale the observations: $Y_t = (X_t - \mathbb{E}X_t)/D = T_t + \eta_t$, where $T_t = S_t - \mathbb{E}S_t$, $\mathrm{Cov}(T_t, T_{t-j}) = v\nu^{|j|}$, $v = 4\pi_0\pi_1$, and the residual variance is $R = 2/D$. Let r denote (101) and $K = 1 - r/\nu$. We show that

$$\widehat{T}_t = K \sum_{k\ge 0} r^k Y_{t-k}$$

is the orthogonal projection of T_t onto the closed linear span of past and present centered observations. For every $j \ge 0$,

$$\sum_{k\ge 0} r^k \nu^{|j-k|} = \frac{\nu^{j+1} - r^{j+1}}{\nu - r} + \frac{r^{j+1}\nu}{1-r\nu}.$$

Thus the coefficient of ν^j in $\mathrm{Cov}(\widehat{T}_t, Y_{t-j})$ is $Kv\nu/(\nu - r) = v$. Its coefficient of r^j is

$$K\left[R + vr\left(-\frac{1}{\nu - r} + \frac{\nu}{1-r\nu}\right)\right] = 0,$$

because the quadratic equation is equivalent to $R(\nu - r)(1 - r\nu) = vr(1-\nu^2)$. Therefore $\mathrm{Cov}(T_t - \widehat{T}_t, Y_{t-j}) = 0$ for every j. Continuity extends this orthogonality to the closed span, and Pythagoras proves optimality. The positive-variance residual component, uncorrelated across time and orthogonal to the state process, makes distinct square-summable coefficient sequences distinguishable in L^2, so the calibrated linear estimator is unique.

For any scalar score H with conditional mean gap Δ_H, let $p = \pi_0\pi_1$ and $N = \pi_1 \mathrm{Var}(H \mid 1) + \pi_0 \mathrm{Var}(H \mid 0)$. Total variance and binary-state covariance give

$$\mathrm{Var}(H) = N + p\Delta_H^2, \qquad \mathrm{Cov}(H, z_t) = p\Delta_H.$$

$$\mathrm{Corr}(H, z_t)^2 = \frac{\pi_0\pi_1 \mathrm{SNR}_{\pi_1}(H)^2}{1 + \pi_0\pi_1 \mathrm{SNR}_{\pi_1}(H)^2}. \tag{115}$$

The identity follows from these two relations. The minimum MSE after affine calibration of H is $p(1 - \mathrm{Corr}(H, z_t)^2)$. Since the projection above is itself a stable geometric filter, it maximizes squared SNR over the causal linear class. For positive mean separation this also maximizes SNR. Uncalibrated scores are unique only up to scale; negative scalings reverse the class orientation.

Balanced SNR and ranking. Balanced SNR also has a ranking interpretation. For positive mean separation and independent draws from the two conditional score distributions, Cantelli's inequality gives

$$\mathrm{AUROC}(H) \geq \frac{\mathrm{SNR}_{1/2}(H)^2}{\mathrm{SNR}_{1/2}(H)^2 + 2}. \tag{116}$$

When $\mathrm{SNR}_{1/2}^2 \leq 2$, this bound is at most $1/2$ and does not certify better-than-chance ranking. In particular, the fitted local-probe current-score models have $\mathrm{SNR}_{1/2}^2 = 2D < 2$. We do not use this bound to explain their observed AUROC. Maximizing balanced SNR maximizes this lower bound, not necessarily AUROC itself. Even with Gaussian emissions, the conditional accumulator distributions mix over latent paths and need not be Gaussian.

Let H_1, H_0 be independent draws from the respective conditional score laws and assume $\Delta_H = \mathbb{E}H_1 - \mathbb{E}H_0 > 0$. The difference has variance $V_1 + V_0$. Cantelli's inequality gives

$$\Pr(H_1 - H_0 \leq 0) \leq \frac{V_1 + V_0}{V_1 + V_0 + \Delta_H^2}.$$

Hence (116) holds for AUROC, which assigns half credit to ties and is at least $\Pr(H_1 > H_0)$. The positive-orientation assumption is necessary; a negative score orientation cannot be repaired by squaring its SNR in this bound.

H Correlated observation residuals

Assume the hidden chain is stationary and $X_t = DS_t + \xi_t$, where $S_t = 2z_t - 1$ and $D > 0$. The zero-mean stationary residual process is independent of the entire state path and satisfies $\mathrm{Cov}(\xi_t, \xi_{t-k}) = 2D\alpha^{|k|}$ for $0 \leq \alpha < 1$. These assumptions preserve the conditional signal and path variance while changing the noise accumulated from history.

H.1 Exact objective and the zero-memory threshold

Splitting the noise covariance sum into its diagonal and two off-diagonal parts gives

$$\begin{aligned}\mathrm{Var}\left(\sum_{k\geq 0} \rho^k \xi_{t-k}\right) &= 2D \sum_{j,k\geq 0} \rho^{j+k}\alpha^{|j-k|} \\ &= \frac{2D}{1-\rho^2}\left(1 + 2\sum_{\ell \geq 1}(\alpha\rho)^\ell\right) = \frac{2D(1+\alpha\rho)}{(1-\rho^2)(1-\alpha\rho)}.\end{aligned} \tag{117}$$

The geometric weights make these sums absolutely convergent. Independence from the state path makes (117) the noise variance within either class, with zero covariance between noise and the state contribution. The conditional mean separation remains $2D/(1-\rho\nu)$. Combining these quantities yields (102), with $\mathrm{SNR}_w^2 = 4D/Q_{D,w,\alpha}$ and the same B_w as in the uncorrelated-residual model.

Proof of Proposition 3. Write $f_\alpha(\rho) = (1+\alpha\rho)/(1-\alpha\rho)$. At $\alpha = \nu$ the noise term simplifies to

$$A_\nu(\rho) = \frac{2(1-\nu^2\rho^2)}{1-\rho^2}, \qquad A_\nu'(\rho) = \frac{4\rho(1-\nu^2)}{(1-\rho^2)^2} > 0 \quad (\rho > 0).$$

For $\alpha \geq \nu$, the factorization $A_\alpha = A_\nu f_\alpha / f_\nu$ has a positive nondecreasing ratio because

$$\frac{d}{d\rho}\log\frac{f_\alpha}{f_\nu} = \frac{2\alpha}{1-\alpha^2\rho^2} - \frac{2\nu}{1-\nu^2\rho^2} \geq 0.$$

The inequality follows from monotonicity of $x/(1-x^2\rho^2)$ in $x \in [0,1)$. Hence A_α is strictly increasing for positive memory. Since $B_w'(\rho) > 0$ for every $w \in [0,1]$ and $\rho > 0$, the entire objective increases strictly away from zero. Thus $\rho = 0$ is its unique minimizer for every $D > 0$, including the boundary $\alpha = \nu$. □

H.2 Positive memory when residuals mix faster than the state

When $\alpha < \nu$, the weak-signal optimum is determined by both correlation rates. Define

$$G_\alpha(\rho) = \frac{\rho(1-\alpha\rho)}{1-\alpha\rho^3}, \qquad \mathcal{R}_w(D) = \operatorname*{arg\,min}_{0\le\rho<1} Q_{D,w,\alpha}(\rho).$$

The set notation allows the comparison across signal levels to be stated for arbitrary class weights without assuming uniqueness at a fixed level.

Proposition 4 (Memory selection below the correlation threshold)**.** *Suppose $0 \le \alpha < \nu$ and fix $w \in [0,1]$. There is a unique $r_0 \in (0,1)$ satisfying*

$$G_\alpha(r_0) = \frac{\nu-\alpha}{1-\alpha\nu}. \tag{118}$$

It satisfies $r_0 = \nu$ when $\alpha = 0$ and $r_0 < \nu$ when $\alpha > 0$. For every $D > 0$, the set $\mathcal{R}_w(D)$ is nonempty and contained in $(0, r_0)$. If $D_2 > D_1$, every $\rho_2 \in \mathcal{R}_w(D_2)$ and $\rho_1 \in \mathcal{R}_w(D_1)$ satisfy $\rho_2 < \rho_1$. Every selection $\rho^\star_w(D) \in \mathcal{R}_w(D)$ obeys

$$\lim_{D\downarrow 0} \rho^\star_w(D) = r_0, \qquad \lim_{D\to\infty} \rho^\star_w(D) = 0, \qquad \rho^\star_w(D) \sim \frac{2(\nu-\alpha)}{c_w D}, \tag{119}$$

where $c_w = 4a_w$. For stationary weighting $w = \pi_1$, the minimizer is unique, differentiable in D, and has strictly negative derivative.

Proof. Let $\delta = \nu - \alpha > 0$ and $b = 1 - \alpha\nu$. Differentiating the noise term gives

$$A'_\alpha(\rho) = -\frac{4(1-\nu\rho)F(\rho)}{(1-\rho^2)^2(1-\alpha\rho)^2}, \qquad F(\rho) = \delta(1-\alpha\rho^3) - b\rho(1-\alpha\rho).$$

The sign change of A'_α is therefore governed by G_α. Its derivative is

$$G'_\alpha(\rho) = \frac{1 - 2\alpha\rho + 2\alpha\rho^3 - \alpha^2\rho^4}{(1-\alpha\rho^3)^2} > 0.$$

To see positivity, its numerator decreases in $\alpha \in [0,1]$ and equals $(1-\rho)^3(1+\rho) > 0$ at $\alpha = 1$. Since $G_\alpha(0) = 0$, $G_\alpha(1) = 1$, and $\delta/b \in (0,1)$, (118) has exactly one solution. Thus A_α decreases up to r_0 and increases afterwards. For $\alpha = 0$, its minimizer is ν. For $\alpha > 0$, write $A_\alpha = A_0 f_\alpha$ and use $A'_0(\nu) = 0$ and $f'_\alpha > 0$ to obtain $A'_\alpha(\nu) > 0$, hence $r_0 < \nu$.

Because $B'_w > 0$, the full objective increases on $[r_0, 1)$. It attains a minimum on $[0, r_0]$, while $Q'_{D,w,\alpha}(0) = -4\delta < 0$ and $Q'_{D,w,\alpha}(r_0) = DB'_w(r_0) > 0$ exclude both endpoints. For minimizers at $D_1 < D_2$, adding their two minimizing inequalities yields

$$(D_2 - D_1)\big[B_w(\rho_2) - B_w(\rho_1)\big] \le 0.$$

Strict increase of B_w implies $\rho_2 \le \rho_1$. Equality would contradict the two interior first-order conditions because $(D_2 - D_1)B'_w(\rho_1) > 0$.

As $D \downarrow 0$, comparison with r_0 gives

$$A_\alpha(\rho^\star_w) \le Q_{D,w,\alpha}(\rho^\star_w) \le A_\alpha(r_0) + DB_w(r_0) \longrightarrow A_\alpha(r_0).$$

Continuity on $[0, r_0]$ and uniqueness of the minimum of A_α imply convergence to r_0. As $D \to \infty$, comparison with zero gives $DB_w(\rho^\star_w) \le Q_{D,w,\alpha}(0) = 2$. Since B_w vanishes only at zero, the minimizers converge to zero. Their first-order condition then yields

$$D\rho^\star_w = -\frac{A'_\alpha(\rho^\star_w)}{B'_w(\rho^\star_w)/\rho^\star_w} \longrightarrow \frac{4(\nu-\alpha)}{8a_w},$$

which proves the rate in (119).

For $w = \pi_1$, put $c = 4\pi_0\pi_1(1-\nu^2)$, so $B_w = c\rho^2/(1-\rho^2)$. On $(0, r_0)$ the first-order condition is $D = T(\rho)$, where

$$T(\rho) = -\frac{A'_\alpha(\rho)}{B'_w(\rho)} = \frac{2}{c}\frac{1-\nu\rho}{1-\alpha\rho}\left[\frac{\delta}{G_\alpha(\rho)} - b\right].$$

Both factors are positive and strictly decreasing on this interval. Moreover, $T(\rho) \to \infty$ at zero and $T(\rho) \to 0$ at r_0. Thus there is exactly one stationary point, and the derivative of the objective changes from negative to positive there. Since $T' < 0$, implicit differentiation gives $d\rho^\star_{\pi_1}/dD = 1/T' < 0$. □

The uncorrelated-residual convexity argument does not extend to this model: at $(\nu, \alpha, \rho) = (0.999, 0.5, 0.9)$, direct differentiation gives $A''_\alpha(\rho) = -4104324183/1141166125 < 0$. The stationary-weight proof instead uses the single crossing of T.

I FINITE-HISTORY CALCULATIONS

For this section the score residual $\eta_t = Y_t - m_{z_t}$ satisfies

$$\mathbb{E}[\eta_t \mid z_{0:T}] = 0, \quad \mathbb{E}[\eta_t^2 \mid z_{0:T}] = s_{z_t}^2, \quad \mathbb{E}[\eta_t \eta_u \mid z_{0:T}] = 0 \quad (t \neq u).$$

The moments are conditional on the entire path, not merely marginal moments given the current state.

Now let $h_0 = 0$, specify any initial state law, and observe a finite trajectory. For a score Y_t with path-conditional mean m_{z_t}, residual variance $s_{z_t}^2$, and zero cross-time residual covariances conditional on the full path, define row vectors

$$p_t(j) = \Pr(z_t = j), \quad a_t(j) = \mathbb{E}[h_t \mathbf{1}_{\{z_t = j\}}], \quad b_t(j) = \mathbb{E}[h_t^2 \mathbf{1}_{\{z_t = j\}}].$$

These joint moments avoid using the stationary backward transition at a nonstationary start.

Proposition 5 (Finite-start moment recursion)**.** *Initialize $a_0 = b_0 = 0$ and $p_0 = (1-p, p)$. For $h_t = \rho h_{t-1} + Y_t$, let $u = p_{t-1}P$, $f = a_{t-1}P$, $g = b_{t-1}P$. Then*

$$\begin{aligned} p_t(j) &= u_j, \qquad a_t(j) = \rho f_j + m_j u_j, \\ b_t(j) &= \rho^2 g_j + 2\rho m_j f_j + (m_j^2 + s_j^2) u_j. \end{aligned} \tag{120}$$

Consequently $M_{t,j} = a_t(j)/p_t(j)$ and $V_{t,j} = b_t(j)/p_t(j) - M_{t,j}^2$ give exact conditional means and variances in $O(t)$ time for each memory parameter.

Under a stationary state start but finite accumulator history, the normalized model gives

$$\Delta_{h,t} = 2D \frac{1 - (\rho\nu)^t}{1 - \rho\nu}, \qquad V_{\mathrm{noise},t} = 2D \frac{1 - \rho^{2t}}{1 - \rho^2}. \tag{121}$$

At $t = 1$, every ρ gives the same score. At later times, (120) computes the finite objective without assuming that a fixed number of discarded steps establishes stationarity. Appendix J gives an explicit uniform finite-history error bound on compact memory intervals and a deterministic bound for perturbed objective selection.

I.1 PROOF OF PROPOSITION 5

Multiply $h_t = \rho h_{t-1} + Y_t$ by the indicator of $z_t = j$ and take expectations. The state transition gives $\mathbb{E}[h_{t-1}\mathbf{1}_{\{z_t=j\}}] = (a_{t-1}P)_j$. The residual has zero expectation conditional on all states, so the observation contribution is $m_j p_t(j)$. Squaring the recursion and applying the same conditioning gives

$$\mathbb{E}[h_t^2 \mathbf{1}_{\{z_t=j\}}] = \rho^2 (b_{t-1}P)_j + 2\rho m_j (a_{t-1}P)_j + (m_j^2 + s_j^2) p_t(j).$$

Each cross term with a previous residual vanishes by the stated conditional cross-covariance assumption. This proves (120). Positive transition entries ensure $p_t(j) > 0$ for $t \geq 1$, allowing division to obtain conditional moments.

If the state starts in stationarity, reversibility gives the lag-k conditional signal $2D\nu^k$ for every $k < t$. Summing the finite geometric series gives (121); the residual variance follows from summing $2D\rho^{2k}$.

I.2 PROOF OF THE FINITE LINEAR ESTIMATOR

For common variance R and mean gap δ, write $u_t = z_t - \lambda - \nu z_{t-1}$. Its conditional expectation given all earlier states is zero. Consequently it is orthogonal to earlier states and, under the path-conditional observation assumptions, to earlier observations. Its variance is

$$\begin{aligned} q_t &= p_t(1 - p_t) - \nu^2 p_{t-1}(1 - p_{t-1}) \\ &= (1 - p_{t-1})\lambda(1 - \lambda) + p_{t-1}\varepsilon(1 - \varepsilon) > 0. \end{aligned}$$

The current observation residual is orthogonal to all states and to other residuals.

Suppose inductively that $\widehat{z}_{t-1}$ is the affine projection onto previous observations, with error variance C_{t-1}. The prediction error is $\nu(z_{t-1} - \widehat{z}_{t-1}) + u_t$, orthogonal to the previous observation span, and has variance $C_t^- = \nu^2 C_{t-1} + q_t$. The innovation $e_t = Y_t - m_0 - \delta\widehat{z}_t^-$ is orthogonal to that span, has variance $\delta^2 C_t^- + R$, and covariance δC_t^- with the prediction error. Projecting onto this new orthogonal direction yields K_t and (103). The resulting residual remains orthogonal to all observations, closing the induction.

When the first available state probability is specified as p_1 directly, the first update starts with mean p_1 and variance $p_1(1-p_1)$; no extra transition is applied. The linear estimate need not remain in $[0,1]$, and we do not clip it when evaluating rankings. At stationary steady state, the coefficient $\nu(1 - K_t\delta)$ is constant and reduces to (101).

I.3 Score origin and finite accumulation

For any constant c, a zero-initialized accumulator obeys

$$h_t(Y + c) = h_t(Y) + c\frac{1-\rho^t}{1-\rho}.$$

At a fixed position t this is a common offset and leaves conditional SNR and AUROC unchanged. Across positions it can change rankings. In the infinite stationary sum it is the same constant $c/(1-\rho)$ at every position. This identity explains why the input-score origin and initialization must be specified when transferring a stationary objective to finite trajectories.

J Finite-history and perturbation bounds

Proposition 6 (Uniform finite-history approximation)**.** *Assume a stationary state start and the normalized moment model, and let $t \geq 1$. Fix a common memory interval $I = [0, \bar{\rho}]$ with $\bar{\rho} < 1$, and let*

$$Q_t(\rho) = \frac{4D[wV_{t,1} + (1-w)V_{t,0}]}{\Delta_{h,t}(\rho)^2}, \qquad M = \frac{\sqrt{D^2 + 2D}}{(1-\bar{\rho})\sqrt{\min(\pi_0, \pi_1)}}.$$

Here $V_{t,i}$ includes observation noise. Uniformly on I,

$$|Q_t - Q_\infty| \leq \left(\frac{2M^2}{D} + \frac{2M^3}{D^2}\right)\bar{\rho}^t =: L_t.$$

If r_t, r_∞ minimize the two objectives over I, then $0 \leq Q_t(r_\infty) - Q_t(r_t) \leq 2L_t$.

Proof. Couple finite and infinite accumulators on the same stationary path. Conditional L^2 norms of either sum are at most M, and that of the omitted tail is at most $M\bar{\rho}^t$. Conditional standard deviations differ by at most the centered-tail norm, itself no larger than the tail norm. Their sum is at most 2M, so the conditional variances differ by at most $2M^2\bar{\rho}^t$. The mean gaps differ by at most $2M\bar{\rho}^t$, both gaps are at least 2D by (121), and all variances are at most M^2. On $[2D, \infty)$, the derivative magnitude of x^{-2} is at most $1/(4D^3)$. Applying these bounds to the variance-to-squared-gap ratio and multiplying by 4D gives L_t. The regret bound follows by adding the two uniform errors around the minimizing inequality for Q_∞. □

The cap $\bar{\rho} = \nu$ contains the stationary optimum. The comparison remains against the finite optimum on that same capped interval. The bound is conservative for rare classes or long memory; the exact finite moments are preferable for numerical selection.

Proposition 7 (Stability of memory selection)**.** *On a common compact memory interval, let r_* minimize $Q = Q_{D,w}$ and let $\widehat{r}$ globally minimize a differentiable $\widehat{Q}$. Put $m = 4(1-\nu)^2$. If $\sup|\widehat{Q}' - Q'| \leq \eta$, then*

$$|\widehat{r} - r_*| \leq \eta/m, \qquad Q(\widehat{r}) - Q(r_*) \leq \eta^2/(2m).$$

If instead only $\sup|\widehat{Q} - Q| \leq \delta$ is known, the regret is at most 2δ.

Proof. The explicit A" formula implies $A'' \geq 4(1-\nu)^2$, and $B''_w \geq 0$, so Q is m-strongly convex. Add the first-order variational inequalities at the two constrained minimizers and use strong monotonicity of Q' to obtain the distance bound. Strong convexity at $\widehat{r}$ also gives

$$Q(\widehat{r}) - Q(r_*) \leq Q'(\widehat{r})(\widehat{r} - r_*) - \frac{m}{2}|\widehat{r} - r_*|^2.$$

The variational inequality for $\widehat{Q}$ bounds the first term by $\eta|\widehat{r} - r_*|$. Maximizing the resulting scalar quadratic yields the regret bound. The uniform-value-error case follows from two approximation errors and optimality of $\widehat{r}$ for $\widehat{Q}$. □

For errors only in D and w, a computable derivative bound is

$$\eta \leq |\widehat{D} - D| \sup_I |B'_{\widehat{w}}| + D|\widehat{w} - w| \sup_I |B'_1 - B'_0|.$$

This is a deterministic sensitivity statement. It does not assert that a real score distribution is close to the assumed moment model.

K SCOPE OF FORMAL VERIFICATION

The Lean development checks seven scalar declarations supporting the earlier exponential-memory analysis. The colored-kernel, innovation-robustness, and finite-start discrimination theorems are proved analytically; they are not formalized in Lean. The certificate source is `formal/MemoryMatching.lean`.

Checked declaration	Exact role
`a_sub_b`	Difference of the two class-weighted path coefficients.
`stationary_b_eq_nu_a`	Stationary-weight coefficient cancellation.
`stableRoot_quadratic`	The rationalized root satisfies its quadratic equation under the discriminant hypotheses.
`stableRoot_mem_Ioo`	The selected scalar root belongs to $(0, 1)$.
`hasDerivAt_A`	A Mathlib `HasDerivAt` proof of the observation-noise term's first derivative.
`hasDerivAt_A_first`	A separate checked derivative of that first-derivative expression.
`A_second_expression_pos`	Positivity of the resulting second-derivative formula on the admissible interval.

The definitions of `nu`, `a`, and `b` translate (1) and (96). The generic Lean variable `v` denotes the manuscript's ν, and `r` denotes ρ. The calculus certificates concern

$$A(\rho) = \frac{2(1-\rho\nu)^2}{1-\rho^2}, \qquad A'(\rho) = \frac{4(\rho-\nu)(1-\rho\nu)}{(1-\rho^2)^2},$$

$$A''(\rho) = \frac{4[(1-\nu)^2(1+3\rho^2) + 2\nu(1-\rho)^3]}{(1-\rho^2)^3}.$$

The derivative proofs apply product, power and quotient rules with explicit nonzero-denominator hypotheses. The positivity certificate uses $0 < \nu < 1$ and $0 \leq \rho < 1$. These declarations verify calculus statements, not just algebraic equality of externally differentiated expressions. The positive-series argument for the path term, global convexity and optimizer uniqueness remain in the analytic proof.

The certificates were checked with Lean 4.34.0-rc2 and Mathlib revision `85e3a25e006c35636f0e53b0e9296caca2685bc0` (tag `v4.34.0-rc2`). The supplied runner uses an existing Mathlib project:

```
python formal/check.py --mathlib-project /path/to/mathlib-project
```

It invokes `lake env lean`, records the compiler and exit status, and prints each declaration's dependencies. The saved check completed with exit code zero. The dependencies are the standard

foundations `propext`, `Classical.choice`, and `Quot.sound`; there is no `sorryAx` dependency or added axiom. The source, environment record and compiler log specify the exact checked scope. Dense covariance solves and finite-path enumeration separately check the implementations of the main analytic results.

L BASELINE MEASUREMENTS AND SUPPORTING EXPERIMENTS

This appendix specifies dataset reconstruction, the common local-probe protocol, and the original restricted baseline and simulation suites. The expanded six-setting, 24-method evaluation is in Appendix M. The per-comparison percentile intervals below describe the original baseline contrasts; the expanded comparisons use the simultaneous family specified in that appendix.

L.1 TRAJECTORY MEASUREMENTS

The descriptive measurements use the PRM800K phase-2 training file and the local first shard of Math-Shepherd. We encode a negative step label as $z_t = 1$. For a collection of observed trajectories, let N_{ij} count adjacent labels (i, j) and let I_i count trajectories whose first observed label is i. The descriptive estimates are

$$\widehat{\lambda} = \frac{N_{01}}{N_{00} + N_{01}}, \qquad \widehat{\varepsilon} = \frac{N_{10}}{N_{10} + N_{11}}, \qquad \widehat{p}_{\text{first}} = \frac{I_1}{I_0 + I_1}.$$

These ratios weight observed transitions or trajectories. For uncertainty, we resample exact problem-text clusters, retaining all trajectories and annotations of each sampled problem. We recompute the ratios in 2,000 bootstrap replicates with seed 20260912 and report the 2.5th and 97.5th percentiles. The intervals describe sampling variation within these sources; source selection and annotation mechanisms determine the population to which the ratios refer.

PRM800K reconstruction. We exclude screening and quality-control records, removing 20,751 of 97,782 rows. At each annotated step, we match completion text exactly to the corresponding original generated step (`pre_generated_steps`). Multiple exact matches must have the same rating. Reconstruction stops at a missing match, an invalid or absent rating, conflicting duplicate ratings, or the end of the original generated sequence. We retain the current rating before stopping when a human continuation or a different chosen completion would change the prefix for subsequent ratings. A null `chosen_completion` does not remove an original negative step: 64,110 such negative ratings are retained.

Of the 77,031 remaining records, 75,922 contain a nonempty observed prefix, spanning 10,806 problem clusters. The primary encoding treats rating zero as non-error. It gives $N_{00} = 374{,}175$, $N_{01} = 58{,}838$, and $\widehat{\lambda} = 0.13588$ with interval $[0.13446, 0.13735]$. Treating neutral ratings as errors gives 0.16854 with interval $[0.16687, 0.17030]$; censoring at the first neutral rating gives 0.13173 with interval $[0.13022, 0.13327]$. These are adjacent-label hazards in selected annotated prefixes. Under the primary encoding, $N_{10} = N_{11} = 0$ because annotation stops at an error, so a recovery probability is not identifiable. Transitions created by recoding neutral ratings do not resolve that censoring.

Math-Shepherd parsing. The measurement parser requires that replacing terminal positive and negative label tags with the input's step marker reconstructs the input text exactly. It also checks consecutive step numbering, agreement between marker and label counts, and a nonempty problem prefix. Of 329,763 rows, it retains 329,536: 157,351 MATH and 172,185 GSM8K trajectories. Exclusions comprise 193 numbering or count mismatches, 33 text mismatches, and one empty problem. Clustering uses task and exact problem text, giving 37,291 MATH and 7,481 GSM8K clusters.

The retained MATH stratum contains 16 observed recoveries among 477,853 transitions out of the error class, whereas GSM8K contains 6,788 among 203,578. Their recovery ratios are 3.3483×10^{-5} and 0.03334, with cluster intervals $[1.8758, 5.0569] \times 10^{-5}$ and $[0.03214, 0.03455]$, respectively. Math-Shepherd labels measure rollout-derived solvability. A negative-to-positive label transition therefore measures recovery of this label, rather than establishing a logically valid self-correction. The two data sources are measured separately, and their rates are not combined into one transition matrix.

L.2 Problem sampling and representations

For encoder evaluation, we first group candidate Math-Shepherd trajectories by exact problem text within each task, requiring a nonempty problem and at least two extracted, labeled steps. With NumPy's default random generator and seed 20260912, we sample 1,000 distinct problems without replacement from the sorted problem list, then select one trajectory uniformly within each chosen problem. The candidate pools contain 37,203 MATH problems with 154,270 trajectories and 7,481 GSM8K problems with 171,670 trajectories. The selected manifests record row identifiers, problem text, step text, labels, and positions. All selected trajectories pass the stricter measurement parser. The sampling candidate pool is formed by marker extraction before that validation; it is therefore distinct from the full-shard measurement population above. Validation failure stops extraction without substituting another trajectory.

The MATH sample has 7,538 steps and the GSM8K sample has 3,558. Each task uses the identical problem sample for Qwen2.5-Math-1.5B-Instruct and Qwen2.5-7B-Instruct. The initial protocol specified MATH. After the MATH evaluation, the difference in full-shard label dynamics motivated a separate GSM8K validation. Its configuration was fixed before extracting GSM8K features or evaluating its scores, and it uses the same hyperparameters and all eight baselines.

Both checkpoints remain frozen and run in evaluation mode with bfloat16 weights. Each input is the literal concatenation of `Problem:`, the problem text, a newline, `Current reasoning step:`, and the current step. We extract the hidden vector at index -6 of the model's returned hidden-state tuple, at the final nonpadding token. The input contains no other solution step. We use right padding and attention masks, sort batches by token length, and restore vectors to their original trajectory positions. Inputs exceeding the checkpoint's context limit raise an error rather than being truncated. The maximum observed input length is 1,953 tokens for MATH and 1,348 for GSM8K. Feature caches contain vectors, labels, problem identifiers, and step indices.

L.3 Nested fitting and working-model calibration

We partition a seed-20260912 permutation of problem identifiers into five outer folds. Each fold holds out 200 problems and trains on 800. Within outer fold $k \in \{0, \dots, 4\}$, a separate permutation with seed $20260912+k+1$ defines three inner folds by problem. Every probe consists of featurewise standardization followed by logistic regression on the full vector, with inverse regularization strength $C = 0.01$, the L-BFGS solver, tolerance 10^{-4}, and at most 1,000 iterations. Each inner validation score is produced by a scaler and probe fitted on the other inner folds.

Let Y_t denote the probe's decision score. Pooling the inner held-out scores within an outer training set gives class means $\widehat{m}_0, \widehat{m}_1$ and class variances $\widehat{v}_0, \widehat{v}_1$, computed with the population-variance denominator. We set

$$\widehat{\delta} = \widehat{m}_1 - \widehat{m}_0, \qquad \widehat{R} = \frac{\widehat{v}_0 + \widehat{v}_1}{2}, \qquad \widehat{D} = \frac{\widehat{\delta}^2}{2\widehat{R}}, \qquad X_t = \frac{\widehat{\delta}}{\widehat{R}} \left(Y_t - \frac{\widehat{m}_0 + \widehat{m}_1}{2} \right).$$

We then refit the scaler and probe on the complete outer training set and apply this fixed inner-derived transformation to its outer-test scores. Thus moment calibration is transferred from cross-fitted scores to a refitted probe. No held-out labels or pooled test-score centering enter that transfer. The equal-variance Gaussian model makes X_t a log-likelihood ratio; on the learned score stream it is a working-model calibration.

Transition and initial probabilities use the outer training trajectories:

$$\widehat{\lambda} = \frac{N_{01} + 1/2}{N_{00} + N_{01} + 1}, \qquad \widehat{\varepsilon} = \frac{N_{10} + 1/2}{N_{10} + N_{11} + 1}, \qquad \widehat{p}_{\mathrm{first}} = \frac{I_1 + 1/2}{I_0 + I_1 + 1}.$$

The half-counts are Jeffreys pseudocounts. The 1,000-problem MATH sample has no observed $1 \to 0$ transition, so all fitted MATH recovery probabilities arise from these pseudocounts. Across outer folds, $\widehat{\varepsilon}$ ranges from 2.18×10^{-4} to 2.32×10^{-4}, while $\widehat{\lambda}$ ranges from 0.1046 to 0.1164 and $\widehat{p}_{\mathrm{first}}$ from 0.2016 to 0.2129. The GSM8K sample contains 43 observed recoveries. These task-specific fits make the MATH experiment a particularly strong finite-start departure from the stationary class distribution.

L.4 SCORE DEFINITIONS

Every score increases toward the error class. The per-step baseline uses Y_t, which has the same ranking as X_t because $\widehat{\delta} > 0$ in the fitted folds. Four accumulator baselines use

$$h_0 = 0, \qquad h_t = \rho h_{t-1} + X_t.$$

Their memories are $\rho = 0.5$, the stationary optimum for $w = 1/2$, the stationary optimum for $w = \widehat{\pi}_1$, and an inner-selected value, where $\widehat{\pi}_1 = \widehat{\lambda}/(\widehat{\lambda} + \widehat{\varepsilon})$. The theory optima minimize $Q_{\widehat{D},w}$ over $[0, \widehat{\nu}]$, with $\widehat{\nu} = 1 - \widehat{\lambda} - \widehat{\varepsilon}$, using bounded scalar minimization with tolerance 10^{-10}. For stationary weighting this is numerically equivalent to the closed-form Wiener memory. The inner-selected value maximizes all-position AUROC of the inner held-out accumulator scores over $\{0, 0.05, \ldots, 0.95, 0.98\}$; ties select the first grid value. Accumulators reset at each problem and retain every observed position. In particular, the stationary-theory memories are deployed with zero initialization, rather than with their infinite stationary histories.

Finite linear MMSE. For compactness, suppress hats on train-fitted parameters in the following recurrences. Let $p_1 = p_{\text{first}}$ and $p_t = \lambda + \nu p_{t-1}$. Initialize the predicted mean and error variance as $a_1 = p_1$ and $V_1 = p_1(1 - p_1)$. At step t, update

$$K_t = \frac{\delta V_t}{\delta^2 V_t + R}, \qquad \widehat{z}_t = a_t + K_t(Y_t - m_0 - \delta a_t), \qquad C_t = (1 - \delta K_t)V_t.$$

For the next step, use

$$a_{t+1} = \lambda + \nu \widehat{z}_t, \qquad V_{t+1} = \nu^2 C_t + p_{t+1}(1 - p_{t+1}) - \nu^2 p_t(1 - p_t).$$

The score is $\widehat{z}_t$, without clipping. The process variance depends on the unconditional finite-time marginal p_t, as required by the affine projection.

Gaussian HMM and position prior. The HMM starts at $b_1 = p_{\text{first}}$, updates $u_t = \text{sigmoid}(\text{logit}(b_t) + X_t)$, and predicts $b_{t+1} = \lambda + \nu u_t$. Its score is u_t. The implementation clips the probability passed to the logit to $[10^{-12}, 1 - 10^{-12}]$. The position-only baseline scores each step with p_t and uses no encoder observation. Both filters reset at each problem. Together with the five per-step or accumulator scores, these give the eight initially specified methods.

Attribution controls. After the initial four task–encoder evaluations, independent review motivated two additional controls for real-data attribution. They reuse every original split, probe, calibration moment, transition estimate, and bootstrap draw, with no test-dependent parameter selection. The first combines the current score with the position prior but uses no earlier observations:

$$s_t^{\text{point}} = p_t + \frac{\delta p_t(1 - p_t)}{\delta^2 p_t(1 - p_t) + R}(Y_t - m_0 - \delta p_t).$$

The second replaces the finite linear estimator's gain schedule with the stationary gain $K_\infty = (1 - \rho_{\text{stat}}/\nu)/\delta$, where ρ_{stat} is the fitted stationary-weight memory. It initializes $a_1 = p_{\text{first}}$ and updates

$$s_t^{\text{fixed}} = a_t + K_\infty(Y_t - m_0 - \delta a_t), \qquad a_{t+1} = \lambda + \nu s_t^{\text{fixed}}.$$

Neither control is clipped. The fixed-gain control matches the finite estimator's initial mean and Markov prediction, so their difference is the observation-gain schedule. These two methods already appear in the finite-start synthetic experiment below; their addition here extends the comparison to learned scores. All original and additional results are retained. In addition to comparisons with per-step and tuned EWMA scores, we report paired bootstrap differences against both attribution controls.

L.5 COMMON METRICS AND TRAINING DIAGNOSTICS

The AUROC definitions, paired problem-bootstrap protocol, and simultaneous comparison family are specified in Appendix M.5. Appendix N defines the matched empirical SNR study and its distinct family. Both studies average metrics over outer folds, preserving problem clusters and keeping

scores from separately fitted models out of a common ranking. Tables 5 and 6 consolidate the baseline and extended AUROC results.

The following training diagnostics use inner held-out scores. They report $\widehat{D}$, the class variance ratio $\widehat{v}_1/\widehat{v}_0$, adjacent-residual Pearson correlation, and selected memories. These quantities assess marginal calibration and one-lag dependence. The held-out higher-lag and conditional-mean diagnostics in Appendix N address additional implications of the path-conditional model.

Outer-training quantity	MATH 1.5B	MATH 7B	GSM8K 1.5B	GSM8K 7B
$\widehat{\lambda}$	0.105–0.116	0.105–0.116	0.302–0.318	0.302–0.318
$\widehat{\varepsilon}$	0.00022–0.00023	0.00022–0.00023	0.03285–0.03850	0.03285–0.03850
$\widehat{p}_1$	0.202–0.213	0.202–0.213	0.295–0.315	0.295–0.315
$\widehat{D}$	0.301–0.383	0.310–0.395	0.347–0.406	0.370–0.469
$\widehat{v}_1/\widehat{v}_0$	1.104–1.242	1.005–1.294	0.921–1.131	1.097–1.210
Residual lag-one correlation	0.516–0.550	0.479–0.518	0.347–0.384	0.317–0.351
ρ, balanced stationary	0.820–0.828	0.817–0.828	0.575–0.589	0.559–0.585
ρ, frequency stationary	0.882–0.894	0.882–0.894	0.604–0.619	0.597–0.619
ρ, validation-tuned	0.400–0.450	0.400–0.500	0.250–0.400	0.250–0.350

Table 3: Ranges across the five outer training folds. Score moments and residual correlations use inner out-of-fold predictions.

L.6 Synthetic model validation

Stationary moments and memory. We use transition pairs $(0.2, 0.05)$ and $(0.1, 0.1)$. Each pair is evaluated at $D \in \{0.02, 0.1, 0.5, 2, 10, 50\}$. Observations satisfy $X_t = D(2z_t - 1) + \sqrt{2D}\eta_t$, with independent η_t drawn either from $\mathcal{N}(0,1)$ or uniformly from $[-\sqrt{3}, \sqrt{3}]$. Each condition uses three independent chain/noise replicates. A replicate starts the chain at its stationary law, discards 3,000 accumulator steps, and measures the next 200,000 steps, giving 14.4 million retained steps in total. We evaluate $\rho \in \{0, 0.01, \ldots, 0.95\}$, compute both conditional means and variances, and compare class-balanced SNR with its exact curve. The exact optimizer uses bounded minimization over $[0, \nu]$ with tolerance 10^{-13}. Error bars represent standard deviations across the three replicates. The seeds are $202609120 + 1000c + 100d + 10e + r$, where c, d, e index the listed chain, strength, and emission settings from zero and $r \in \{0, 1, 2\}$.

Finite-start estimation. We generate 10,000 independent trajectories of length 40 using seed 2026091201, $(\lambda, \varepsilon, p_{\text{first}}, D) = (0.12, 0.002, 0.1, 0.2)$, and observations $Y_t = z_t + \eta_t$ with $\eta_t \sim \mathcal{N}(0, 1/(2D))$. Parameters are known to all estimators. We compare finite linear MMSE, exact Gaussian HMM, stationary Wiener filtering initialized at π_1, the optimal per-position affine estimator from the current observation, and a fixed-gain filter initialized at the true first-step prior. The last method uses the stationary Wiener gain with the same Markov prediction and initial mean as finite linear MMSE, separating the gain schedule from the initial mean.

For these observations, let $R = 1/(2D)$, $q = \pi_1(1-\pi_1)(1-\nu^2)$, $c = 1 + \nu^2 + q/R$, and $\rho = 2\nu/(c + \sqrt{c^2 - 4\nu^2})$. The stationary Wiener gain is $K = 1 - \rho/\nu$ and its update is $s_t = \pi_1 + \rho(s_{t-1} - \pi_1) + K(Y_t - \pi_1)$, with $s_0 = \pi_1$. The memoryless affine estimate is $p_t + \frac{p_t(1-p_t)}{p_t(1-p_t)+R}(Y_t - p_t)$. We evaluate per-position mean squared error and paired excess error relative to finite linear MMSE. Standard errors use the 10,000 independent trajectories. For a mean across positions, we first average each trajectory's losses and then compute its standard error across trajectories.

Exact numerical checks. The finite joint-moment recursions are compared with enumeration over 144 combinations of three transition pairs, four initial laws, four horizons, and three memories, including unequal conditional emission variances. The largest conditional-moment discrepancy is 1.14×10^{-13}. The finite linear estimator is compared with direct dense covariance projection, with maximum coefficient discrepancy 3.06×10^{-16}. Across 300 additional parameter sets, the closed-form stationary memory agrees with numerical minimization within 1.11×10^{-8}, and the largest

checked Wiener–Hopf residual is 3.42×10^{-12}. Eighty parameter sets additionally check finite-horizon and perturbation bounds. These executable checks complement the analytic proofs.

L.7 Resources and reproduction

Feature extraction uses one RTX 5090, processing checkpoints sequentially. Batch sizes are 12 for the 1.5B model and four for the 7B model. For MATH, extraction takes approximately 60 and 107 seconds, with peak allocated CUDA memory of 15.0 and 19.8 GiB, respectively. For GSM8K, extraction takes approximately 16 and 45 seconds, with peaks of 6.9 and 16.0 GiB. These extraction timings begin after loading and tokenization and exclude checkpoint loading. CPU probe fitting uses four numerical threads; simulation uses one. The feature and evaluation JSON files record per-run timing and configuration, and the prediction archives retain labels, problem identifiers, positions, folds, and all ten scores.

From the repository root, the following commands regenerate measurements and synthetic evidence:

```
python scripts/measurement_utils.py --n-boot 2000
python -m unittest discover -s tests -v
python scripts/verify_strengthened_theory.py
python scripts/simulate_memory_law.py
```

For each task and checkpoint, feature extraction and evaluation use

```
python scripts/extract_corrected_features.py \
  --model CHECKPOINT --name RUN --task TASK \
  --n 1000 --batch-size BATCH
python scripts/evaluate_corrected.py --name RUN --bootstrap 2000
```

Here `CHECKPOINT` is `Qwen/Qwen2.5-Math-1.5B-Instruct` or `Qwen/Qwen2.5-7B-Instruct`; `TASK` is `MATH` or `GSM8K`. The corresponding run names are `math1p5b`, `general7b`, `gsm_math1p5b`, and `gsm_general7b`. Checkpoints must already be available locally because extraction uses local-files-only loading. Existing sample manifests fix the sampled problems across checkpoint runs. The protocol is recorded in `EXPERIMENT_PROTOCOL.md`; results and caches are written under `results/corrected/`.

L.8 Class weights and correlated residuals

A second stationary experiment directly tests the class-weight extension and correlated-residual boundary. We fix $(\lambda, \varepsilon) = (0.2, 0.05)$, so $\nu = 0.75$ and $\pi_1 = 0.8$, and cross $w \in \{0, 0.5, 0.8, 1\}$ with $\alpha \in \{0, 0.4, 0.75, 0.9\}$ and $D \in \{0.02, 0.1, 0.5, 2, 10, 50\}$. Independent Gaussian AR(1) residuals have unit stationary variance before multiplication by $\sqrt{2D}$. The first residual is drawn from its stationary law. Each chain/noise configuration has three independent replicates, 200,000 retained observations, and a 3,000-step burn-in. All weights use the same simulated conditional moments within a replicate, for 14.4 million retained observations in total. These settings were fixed as a diagnostic after manuscript review and do not fit any real-data outcome.

We evaluate 96 memories from zero to 0.95. Exact conditional means and variances combine Lemma 1 with (117); the continuous optimum is computed from the exact objective using a dense global grid followed by local refinement and explicit endpoint comparison. This numerical localization is not used as a proof of uniqueness for general weights. Empirical optima maximize the sampled SNR curve.

Figure 3 compares analytic and empirical optima. Across all weights and signal levels, the maximum absolute optimum deviation is 0.0224 for $\alpha = 0$ and 0.0386 for $\alpha = 0.4$. At the boundary $\alpha = \nu$, finite-sample optima can be as large as 0.05, consistent with a shallow objective at zero; for $\alpha = 0.9$, every empirical optimum is zero. The largest root-mean-square relative conditional-variance discrepancy is 2.14% across the simulated configurations and replicates. The record includes every weight, noise correlation, and signal level, including the near-boundary sampling deviations.

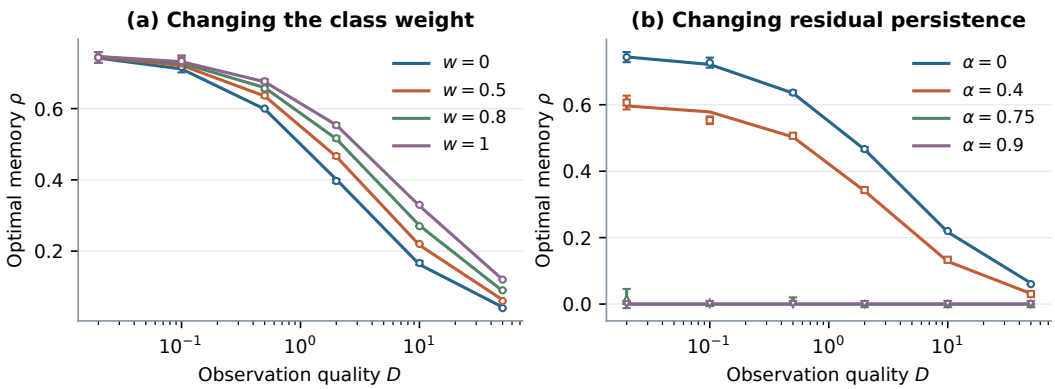


Figure 3: Direct validation of the theory extensions at $\nu = 0.75$. Lines show numerical optima of the exact objective; markers and bars show the mean and one standard deviation of three empirical grid optima. (a) Four class weights with uncorrelated residuals. (b) Four residual correlations at balanced weight. The two curves with $\alpha \geq \nu$ coincide at zero in the population.

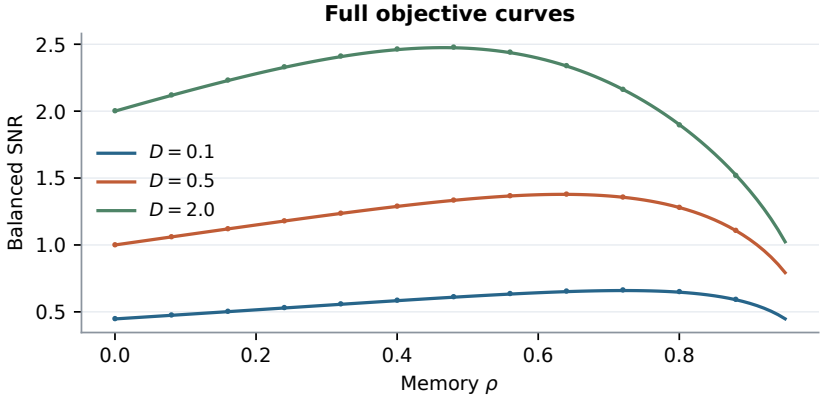


Figure 4: Full balanced-SNR objective curves for the original Gaussian experiment, with $(\lambda, \varepsilon) = (0.2, 0.05)$. Lines are exact values; dots are empirical averages across three independent replicates.

L.9 SENSITIVITY TO TRANSITION SMOOTHING

The MATH sample's fitted recovery rate is entirely determined by smoothing. We therefore repeat the train-label transition and initial-prior estimates with symmetric pseudocounts $\beta \in \{0.1, 0.5, 1\}$, replacing $1/2$ in the numerator and 1 in the denominator by β and 2β. We keep the outer folds, probe predictions, score moments, and calibration fixed. Each setting recomputes the stationary optima, finite LMMSE, Gaussian HMM, and position prior; no value of β is selected using these results. At $\beta = 0.5$, the scores reproduce the primary AUROCs exactly. Table 4 reports all finite-filter results. The largest finite-LMMSE AUROC change relative to the primary setting is below 0.00003. These small changes show stability over this smoothing range, while the lack of observed MATH recoveries remains a limitation on interpreting the fitted recovery probability.

Task and encoder	Filter	$\beta = 0.1$	$\beta = 0.5$	$\beta = 1$
MATH 1.5B	Finite LMMSE	0.75336	0.75337	0.75337
MATH 1.5B	Gaussian HMM	0.75487	0.75487	0.75487
MATH 7B	Finite LMMSE	0.77074	0.77077	0.77079
MATH 7B	Gaussian HMM	0.77218	0.77222	0.77225
GSM8K 1.5B	Finite LMMSE	0.76708	0.76708	0.76707
GSM8K 1.5B	Gaussian HMM	0.76664	0.76664	0.76663
GSM8K 7B	Finite LMMSE	0.77511	0.77511	0.77510
GSM8K 7B	Gaussian HMM	0.77393	0.77392	0.77392

Table 4: Mean outer-fold AUROC under transition and initial-prior smoothing. Probe predictions and moment calibration remain fixed.

L.10 VERIFICATION OF THE UNRESTRICTED OPTIMAL KERNEL

We independently compare the backward-information kernel with a dense covariance solve. The fixed grid uses $(\lambda, \varepsilon) \in \{(0.2, 0.05), (0.1, 0.1)\}$, $D \in \{0.2, 2, 20\}$, and $w \in \{0, 0.5, \pi_1, 1\}$. For lags zero through 80, we construct each covariance matrix directly from the conditional transition probabilities, rather than from the innovation factorization. Across 24 configurations, the maximum absolute coefficient discrepancy is below 1.34×10^{-15}, and the largest normal-equation residual is below 4.45×10^{-16}. The identity $q_{w,j} = U_{w,j} - \nu^2 U_{w,j-1}$ and the equality between $J_{w,0}$ and $g^\top N_w^{-1} g$ are checked separately.

A second fixed grid uses the asymmetric chain, all four weights, 300 adjacent ratios, and 15 logarithmically spaced signal levels from 10^{-7} to 10^{7}. Every adjacent ratio decreases strictly between consecutive signal levels. Figure 5 shows the shared tail and the shortening of normalized history weight. The plotted fixed-signal kernels use a 512-lag truncation. At $(\lambda, \varepsilon, D) = (0.2, 0.05, 2)$, the full-kernel SNR improvement over the best exponential approximation ranges from zero to 0.017% across the four weights. This comparison illustrates the structural distinction without implying a material predictive gain from replacing the optimized exponential kernel on this grid.

M EVALUATION ON LOCAL AND PREFIX-AWARE OBSERVATION STREAMS

The empirical evaluation asks how the population SNR prescriptions transfer to finite reasoning traces. We compare 24 aggregation rules on six observation settings, retain both pooled and within-position discrimination, and estimate all aggregation parameters from training problems. The stronger comparison includes EWMA tuned for the metric being evaluated and trained aggregators with the same position information.

M.1 CONTROLLED COVARIANCE AND EMISSION EXPERIMENTS

An independent covariance construction checks the colored-kernel ratio recursion in 2,000 finite configurations. We draw 200 combinations of persistence $\nu \in [0.03, 0.985]$, stationary prevalence $\pi_1 \in [0.001, 0.999]$, class weights including 0, $1/2$, 1, and random interior values, and residual correlations including zero and $\alpha = \nu$. Each combination uses ten logarithmically spaced signal

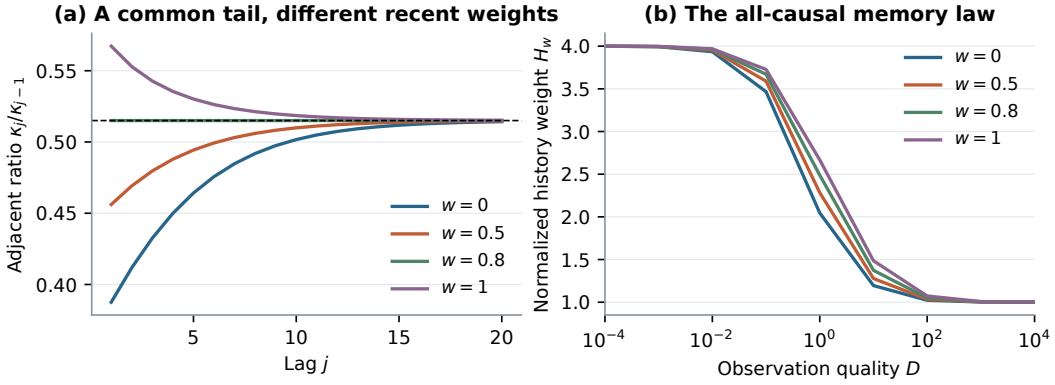


Figure 5: Structure of the optimal causal linear kernel at $(\lambda, \varepsilon) = (0.2, 0.05)$. (a) At $D = 2$, class weights other than π_1 change the recent coefficient ratios, which converge to a common Wiener decay (dashed). Stationary weighting $w = 0.8$ is exactly geometric. (b) The total coefficient weight after normalizing the current coefficient to one decreases from four to one for every weight. Curves evaluate the information recursion, not a fitted real-data model.

levels from 10^{-3} to 10^{4} and a last lag chosen from $\{1, 2, 8, 50\}$. Direct inversion of the conditional covariance and the independent ratio recursion agree to maximum normalized coefficient error 2.96×10^{-14}. The maximum normal-equation residual is 1.84×10^{-11}. No tested case violates the historical-coefficient sign law, total absolute-memory monotonicity, or auxiliary-ratio monotonicity. Individual coefficient magnitudes nevertheless increase in 116 adjacent signal comparisons, consistent with the counterexample in Appendix A. These finite checks assess the formulas; the infinite-history claims follow from the proofs.

The nonnegative-kernel boundary is also checked by constrained quadratic optimization in 108 AR(1) configurations. The chains are $(\lambda, \varepsilon) \in \{(0.2, 0.05), (0.1, 0.1), (0.015, 0.003)\}$, with $w \in \{0, 1/2, 1\}$, $D \in \{0.1, 1, 10\}$, and $\alpha \in \{0, \nu/2, \nu, (1+\nu)/2\}$, using lags 0 through 32. The two-tap coefficient and squared-SNR identities agree with covariance solutions to 5.11×10^{-15} and 7.11×10^{-15}, respectively. At exact cancellation, the largest historical coefficient is 1.17×10^{-14}. A further 48 positive-definite residual-covariance cases use 13 coefficients and normalized first columns that dominate, equal, fall below, or cross the conditional mean-gap vector. Their normalized quadratic identity has maximum error 1.43×10^{-14}.

A separate emission experiment uses $(\lambda, \varepsilon, w) = (0.2, 0.05, 1/2)$, $\alpha \in \{0.4, 0.75, 0.9\}$, and $D \in \{0.1, 1, 10\}$. Each of the nine configurations has three independent seeds, 200,000 retained steps and a 3,000-step burn-in, giving 27 runs. State-independent Gaussian AR(1) residuals have stationary variance $2/D$. Common random numbers pair configurations within a seed. We compare exact and empirical balanced SNR for the current score, optimal nonnegative exponential memory, the signed two-tap rule, and the unrestricted kernel on lags 0 through 32. The maximum absolute relative deviation of an individual empirical SNR from its exact value is 5.86%. At $\alpha = \nu = 0.75$, all four exact scores have the same SNR. At $\alpha = 0.9$ and $D = 1$, the current score and nonnegative exponential score have exact SNR 1.41421, while the two-tap and unrestricted scores attain 1.45369 and 1.53663. The corresponding unrestricted empirical mean is 1.53201, with across-seed standard deviation 0.02407.

Finally, a deterministic 5,000-point scan measures the white-kernel advantage over the best exponential approximation. It crosses $\nu \in \{0.1, 0.5, 0.8, 0.95, 0.99\}$, $\pi_1 \in \{0.001, 0.01, 0.1, 0.3, 0.5, 0.9, 0.99, 0.999\}$, $w \in \{0, 0.1, 0.5, 0.9, 1\}$, and 25 logarithmically spaced values of D from 10^{-3} to 10^{3}. The maximum SNR gain is 0.11753%, attained at $(\nu, \pi_1, w, D) = (0.95, 0.001, 0.9, 0.17783)$. The maximum is 0.05472% over the 816 grid points satisfying $\nu \leq 0.95$, $\pi_1 \in [0.1, 0.9]$, $w \in [0.1, 0.9]$, and $D \in [0.01, 100]$. The scan demonstrates that nongeometricity can have a small numerical cost for this white model. These are attained grid maxima, not uniform approximation guarantees or empirical AUROC gains.

M.2 SAMPLES, OBSERVATIONS, AND PROBLEM SPLITS

We use the same 1,000 distinct Math-Shepherd problems and one trajectory per problem in each task throughout the evaluation. MATH contains 7,538 labeled steps and GSM8K contains 3,558. A label indicates the dataset's rollout-derived solvability assessment; an error-to-clean label transition therefore need not represent a logically repaired proof. The three observation systems within a task share every problem, step, label, and outer-fold assignment.

Two local observation systems use frozen Qwen2.5-7B-Instruct and Qwen2.5-Math-1.5B-Instruct encoders. Each input contains the problem and current step. The probe reads the last nonpadding token at hidden-state index -6. Standardization and an ℓ_2-regularized logistic probe with inverse regularization $C = 0.01$ are fitted on training features. The optimizer is L-BFGS with at most 1,000 iterations and tolerance 10^{-4}. The cached embeddings are unchanged across aggregation methods.

The third observation system is the frozen Qwen2.5-Math-PRM-7B process reward model. We use its official chat template and `<extra_0>` step separators, taking the error logit minus the correct logit at each separator. Each score can use the full causal reasoning prefix. The complete trajectories fit within the 4,096-token context budget: maximum lengths are 2,082 tokens for MATH and 1,394 for GSM8K. All 2,000 problems and their labeled steps are retained without context truncation. Inference uses bfloat16 and batch size one. Comparing selected separator logits against separately evaluated truncated prefixes gives a maximum absolute difference of 0.0078125, consistent with finite-precision execution. The PRM weights are never fitted in this evaluation. Its pretraining overlap with the evaluation data is unknown, so this setting measures aggregation of a fixed prefix-aware score stream rather than generalization to a verified unseen benchmark. It also changes the observation system: overlapping prefixes can induce residual dependence beyond the fitted AR(1) approximation.

Five outer folds hold out complete problems. The split seed is 20260912. Within outer fold $f \in \{0, \dots, 4\}$, three problem folds generated with seed $20260912 + f + 1$ provide cross-fitted local probe scores. No test problem enters feature standardization, probe fitting, calibration, transition estimation, or hyperparameter selection. For the frozen PRM, the raw scores are fixed; downstream parameter estimation uses the same outer-training partitions.

M.3 WORKING-MODEL ESTIMATES AND SCORE CONVENTIONS

Let r_t denote a raw observation score, and let m_i and v_i be its training class mean and variance. For local probes these moments use the three-fold cross-fitted scores. We set

$$R_{\mathrm{raw}} = \frac{v_0 + v_1}{2}, \qquad \delta = m_1 - m_0, \qquad \widehat{D} = \frac{\delta^2}{2R_{\mathrm{raw}}}, \qquad \ell_t = \frac{\delta}{R_{\mathrm{raw}}}\left(r_t - \frac{m_0 + m_1}{2}\right).$$

Thus ℓ_t is the shared Gaussian working log-likelihood ratio used by the linear memory rules. Each transition count receives Jeffreys smoothing of $1/2$. The initial error probability is $(n_{\mathrm{initial},1} + 1/2)/(n_{\mathrm{problems}} + 1)$. Both transition rates and this initial probability use outer-training labels.

For the colored model, form residuals $e_t = r_t - m_{z_t}$ on training scores and fit the no-intercept slope

$$\widehat{\alpha}_{\mathrm{raw}} = \frac{\sum e_{t-1} e_t}{\sum e_{t-1}^2}, \qquad \widehat{\alpha} = \min\{1 - 10^{-6}, \max\{0, \widehat{\alpha}_{\mathrm{raw}}\}\},$$

where sums include adjacent steps from the same problem. We retain the unclipped slope and Pearson correlation as diagnostics. Class centering and one fitted slope specify a working covariance model; they do not establish state-independent AR(1) residuals or equal class variances.

All stationary kernels use current coefficient one and score $h_t = \sum_{j=0}^{t-1} \kappa_j \ell_{t-j}$, with zero inputs before the observed trace. The coefficients are prefixes of the infinite stationary kernel. We do not recenter each lag by a position-dependent mean. This convention makes the stationary white kernel agree with its exponential specialization on every finite trace, while leaving finite-start effects in pooled rankings. The finite-start estimators instead use their fitted initial state distribution.

The colored unrestricted kernels use the auxiliary ratio recursion in Appendix A. The remote ratio is initialized at the constant-innovation fixed point. We extend the horizon until $\nu^{\mathrm{padding}} \le 10^{-13}$ and compare the retained prefix against a doubled horizon, requiring maximum coefficient discrepancy

at most 10^{-12}. The retained coefficients agree at floating-point precision in all six settings. For stationary weighting, the independent identity

$$\kappa_0 = 1, \qquad \kappa_j = (\tau - \alpha)\tau^{j-1} \quad (j \geq 1)$$

gives the equivalent recurrence $h_t = \tau h_{t-1} + \ell_t - \alpha \ell_{t-1}$, with zero starts. Dense finite covariance solves independently verify both white and colored kernel implementations.

M.4 Aggregation methods and selection

Tables 5 and 6 report all 24 methods. The per-step score is the raw observation. Fixed EWMA uses $\rho = 0.5$. White exponential rules minimize their balanced or stationary-weight SNR objectives. White unrestricted rules use the corresponding optimal causal kernels. AR(1) exponential and unrestricted rules use the same two weights and the training residual slope. For an AR(1) exponential rule, we include $\rho = 0$ in the optimization, bracket local minima on a 1,001-point grid in $[0, 1)$, and refine them numerically. If $\widehat{\alpha} \geq \widehat{\nu}$, its memory is zero.

Pooled-tuned EWMA selects ρ from $\{0, 0.05, \ldots, 0.95, 0.98\}$ by pooled AUROC on training cross-fitted scores. Position-tuned EWMA uses the identical grid and selects by within-position AUROC. These are distinct fitted rules. The signed two-tap rule scores $\ell_t + b\ell_{t-1}$, where

$$b = \frac{2(\widehat{\nu} - \widehat{\alpha})}{c_{1/2}\widehat{D} + 2(1 - \widehat{\alpha}\widehat{\nu})}.$$

It uses the same fitted parameters as the colored kernels and zero missing history.

Finite LMMSE and the Gaussian-emission HMM use the same fitted moments, transitions, and initial prior. The position prior uses the Markov marginal alone. The current-score-plus-prior rule applies a memoryless linear update using that marginal. Fixed Wiener gain uses the stationary gain with the fitted initial prior. Tuned fixed gain retains this finite-start affine recursion but selects its memory coefficient from 21 equally spaced values in $[0, \widehat{\nu}]$ by training pooled AUROC. It is a stronger control for initialization and gain choice than a fixed stationary gain.

The trained lag logistic aggregator takes the current raw score and up to $L - 1$ previous scores, resetting at each problem and zero-padding missing history. We select $L \in \{1, 2, 4, 8\}$ and $C \in \{0.01, 0.1, 1, 10\}$. Three grouped meta-validation folds evaluate each candidate by mean pooled AUROC. Each meta-validation probe excludes its validation problems. Its meta-training raw scores are themselves cross-fitted over two further problem folds, with seed $20260912 + 1000 + 10f + j$ for inner fold j. Thus meta-validation labels cannot enter meta-training features through probe fitting. The selected aggregator is fitted on all outer-training cross-fitted scores using standardization and L-BFGS logistic regression, at most 1,000 iterations and tolerance 10^{-5}. The frozen PRM bypasses probe fitting at every level.

Lag logistic with position adds 18 features: indicators for zero-based positions $0, \ldots, 14$ and 15 or greater, numeric position, and $\log(1 + \text{position})$. Per-step logistic with position fixes $L = 1$ and tunes C. EWMA logistic with position jointly selects the EWMA grid and C under the same meta-validation procedure. Its memory input is the raw score centered by the meta-training mean, and it uses the same position features. These controls isolate the contribution of temporal aggregation from information available through trace position.

The final rule is a continuous CUSUM score, $h_t = \max\{0, h_{t-1} + \ell_t - k\}$ with $h_0 = 0$. Its drift is selected from $\{0, 0.1, 0.25, 0.5, 1, 2, 4\}\sqrt{2\widehat{D}}$ using training pooled AUROC. Its output is evaluated at every labeled step. This experiment measures current-state discrimination, not stopping time or detection delay.

M.5 Metrics and a unified uncertainty family

Pooled AUROC compares all positive–negative pairs within an outer fold. Within-position AUROC compares steps only at the same observed position and averages position-specific AUROCs with weights $n_{1,t}n_{0,t}$. Both reported metrics are unweighted means over the five outer folds; scores from separately fitted folds are not ranked against each other. These metrics answer different questions because error prevalence varies with position.

For each fold, let $\omega = \sum_t n_{1,t} n_{0,t} / (n_1 n_0)$. The exact pair decomposition is

$$\mathrm{AUROC}_{\mathrm{pooled}} = \omega\, \mathrm{AUROC}_{\mathrm{within}} + (1-\omega)\, \mathrm{AUROC}_{\mathrm{cross}},$$

where the last term uses pairs at different positions. Mean ω is 0.08323 on MATH and 0.20044 on GSM8K. Most eligible pairs therefore compare different positions, which explains why position calibration can change pooled AUROC substantially without a similar within-position change.

We use 2,000 paired bootstrap draws of whole problems within the fixed outer folds. A draw is shared by all methods and both metrics; models evaluated on the same task also share the problem multiplicities. Marginal and paired percentile intervals are retained. The final simultaneous family contains all 24 methods, both metrics, and differences from three references: per-step scoring, pooled-tuned EWMA, and position-tuned EWMA, across all six settings. Within each setting, we standardize each bootstrap-centered difference by its bootstrap standard deviation and take the maximum absolute value over the method–metric–reference family. Zero-variance identity contrasts are omitted from this maximum. The $1 - 0.05/6$ quantile gives the critical value for that setting; Bonferroni combines the six settings without assuming independence. Every simultaneous interval reported here uses this unified family.

These are exploratory bootstrap intervals conditional on the observed problems, fitted models, selected hyperparameters, and fixed folds. Model fitting and hyperparameter selection are not repeated in the bootstrap. The intervals therefore do not measure refitting variability or remove the effect of adaptive research choices. Comparisons of SNR-prescribed filters with AUROC-selected filters test transfer between objectives, rather than optimality under a common empirical objective.

M.6 Transfer results and observed limits

Prefix-aware scoring increases the fitted signal separation. Mean $\widehat{D}$ is 0.62603 for MATH PRM and 1.50123 for GSM8K PRM, compared with 0.33355–0.34991 and 0.36735–0.41998 for the local probes. The PRM per-step AUROCs are 0.78744 and 0.87401; its within-position AUROCs are 0.76195 and 0.88854. This stronger observation stream still benefits from aggregation, but does not establish empirical superiority of the SNR-optimal colored kernel.

Full colored balanced kernels have within-position AUROCs 0.74946, 0.73365, 0.71717, and 0.70900 in the four local settings. Position-tuned EWMA attains 0.75230, 0.73620, 0.72035, and 0.71617. The colored-minus-position-tuned differences are negative in all four settings, and all four simultaneous intervals contain zero. White balanced kernels and tuned fixed gain are also competitive under this metric. The full colored model therefore does not improve the matched-metric comparison consistently.

On MATH PRM, full colored balanced scoring attains pooled/within-position AUROC $0.80529/0.78163$, compared with $0.80836/0.78110$ for pooled-tuned EWMA. On GSM8K PRM the corresponding values are $0.88937/0.90159$ and $0.89422/0.90626$. Its GSM8K within-position difference from pooled-tuned EWMA is -0.00466, with simultaneous interval $[-0.00887, -0.00046]$. Against position-tuned EWMA, the difference is -0.00440 with interval $[-0.00996, 0.00115]$. Thus the adverse point estimate persists, while its interval depends on which fitted comparator is evaluated. Against the metric-matched comparator, none of the six colored within-position differences excludes zero.

Position-aware learned controls and tuned fixed gain achieve higher pooled discrimination than the stationary prescriptions. For example, tuned fixed gain attains pooled AUROC 0.77265 and 0.77427 for the General-7B local streams, and 0.82236 and 0.92327 for the two PRM streams. Its within-position scores are much closer to the kernel methods. The stationary-gain ablation alone consequently understates what a fitted finite-start control can achieve. Trained lag aggregation also does not dominate tuned EWMA: on MATH PRM their pooled AUROCs are 0.79544 and 0.80836. Continuous CUSUM is not a uniformly competitive current-state score; for the GSM8K Math-1.5B local stream its within-position AUROC is 0.67155, below the per-step value 0.69320.

The fitted residual slopes do not cross the cancellation threshold. Mean $\widehat{\nu} - \widehat{\alpha}$ is 0.42224, 0.38259, 0.35879, and 0.32924 for the four local settings, and 0.43122 and 0.38361 for the two PRM settings. Every outer fold has $\widehat{\alpha} < \widehat{\nu}$. On MATH, AR(1) exponential memory remains 0.80693 and 0.80208 for the local probes, whereas pooled-tuned memory is 0.42 and 0.43. Hence residual AR(1) correction does not explain their much shorter empirical pooled-AUROC optima. The stronger PRM

reduces prescribed AR(1) memory to 0.77120 on MATH and 0.33735 on GSM8K, yet pooled-tuned memories increase to 0.59 and 0.50. Changing observation systems is not a controlled variation of D with all other distributional properties fixed. Direct evidence for the theoretical memory law therefore comes from the controlled moment-model experiments, while these six settings assess its empirical transfer.

M.7 Diagnostics for the label process

We separately evaluate the homogeneous first-order assumption by predicting the next observed label from preceding gold labels. This diagnostic uses the same problem folds and does not modify the score-aggregation experiments. Jeffreys-smoothed categorical models condition on the current label; the current and preceding label with a start sentinel; the current label and source-position bins 1–3, 4–6, 7–10, 11+; or the current label, observed run-age bins with the same boundaries, and an initial-run indicator. Unseen refined contexts back off to first order. No hyperparameters are tuned.

MATH contains 6,538 observed transitions, including 411 clean-to-error transitions and no error-to-clean transition. GSM8K contains 2,558 transitions, including 421 clean-to-error and 43 error-to-clean transitions across 42 recovering problems. MATH's positive fitted recovery probability is therefore supplied entirely by smoothing. Final runs are right censored; only observed next-step transitions enter the exposure counts. Initial-run age measures time since the trajectory was observed, not its unobserved pre-trajectory dwell age.

On GSM8K, second-order conditioning reduces held-out next-label log loss from 0.40203 to 0.39105, a difference of -0.01098 with paired percentile interval $[-0.01663, -0.00556]$. Run-age conditioning gives 0.39281, a difference of -0.00922 with interval $[-0.01410, -0.00420]$. On MATH, first-order log loss is 0.19950; the second-order, position-dependent, and run-age models attain 0.19888, 0.19823, and 0.19862, with all three log-loss difference intervals including zero. Small MATH Brier improvements for the position and run-age models have unadjusted intervals below zero. These diagnostic intervals are paired problem-bootstrap percentile intervals, separate from the simultaneous AUROC family.

The GSM8K results show that first-order homogeneous transitions leave predictive structure in the observed labels. The MATH results do not establish geometric dwell times: recoveries are absent, refined contexts are sparse, and trajectory termination may depend on the labels. These losses describe next-label prediction conditional on a next label being observed. They cannot identify the latent transition process beyond the dataset's labeling and termination mechanism.

N From working-model SNR to held-out discrimination

A matched-SNR evaluation separates failure of the fitted moment prescription from a difference between SNR and AUROC objectives. We reuse every raw observation, outer-training cross-fitted score, and held-out prediction from the six-setting evaluation in Appendix M. The original 24 methods are unchanged. Two additional EWMA rules select memory by training SNR, and a finite-start colored Fisher rule uses the fitted initial distribution. No encoder, probe, PRM, or original aggregator is retrained.

N.1 Signed SNR, position support, and training selection

For a held-out score s, we estimate signed balanced SNR by

$$\widehat{\mathrm{SNR}}(s) = \frac{\overline{s}_1 - \overline{s}_0}{\sqrt{(\widehat{v}_1 + \widehat{v}_0)/2}},$$

where class variances use denominator $n_i - 1$. The direction remains error-minus-clean; test labels never select its sign. A constant score with zero mean gap has SNR zero. A nonzero gap with zero class variance is undefined, as is a cell with fewer than two observations in a class. We do not cap such ratios. Undefined draws are recorded from their class support before evaluating any contrast.

Pooled SNR uses all steps in an outer fold. Within-position SNR is the weighted mean of the SNR estimated separately at each eligible position, with weight $n_{1,t}n_{0,t}$. It is not the SNR of a

	MATH			GSM8K		
Method	Gen. 7B	Math 1.5B	PRM 7B	Gen. 7B	Math 1.5B	PRM 7B
Per-step	0.7371	0.7250	0.7874	0.7398	0.7350	0.8740
EWMA, fixed 0.5	0.7583	0.7403	0.8081	0.7481	0.7424	0.8942
White exponential, balanced	0.7340	0.7122	0.7981	0.7453	0.7393	0.8936
White exponential, stationary	0.7201	0.6979	0.7737	0.7433	0.7375	0.8941
EWMA, pooled-tuned	0.7579	0.7407	0.8084	0.7503	0.7438	0.8942
Finite LMMSE	0.7708	0.7534	0.8212	0.7751	0.7671	0.9201
Gaussian HMM	0.7722	0.7549	0.8197	0.7739	0.7666	0.9203
Position prior	0.6578	0.6578	0.6578	0.6897	0.6897	0.6897
Current score + prior	0.7410	0.7318	0.7882	0.7613	0.7540	0.9023
Fixed Wiener gain	0.6789	0.6775	0.6810	0.7634	0.7537	0.9202
AR(1) exponential, balanced	0.7365	0.7159	0.7989	0.7497	0.7440	0.8912
AR(1) exponential, stationary	0.7278	0.7073	0.7828	0.7491	0.7435	0.8933
White kernel, balanced	0.7360	0.7141	0.7997	0.7454	0.7395	0.8935
White kernel, stationary	0.7201	0.6979	0.7737	0.7433	0.7375	0.8941
Lag logistic	0.7559	0.7391	0.7954	0.7513	0.7442	0.8859
Lag logistic + position	0.7725	0.7556	0.8215	0.7731	0.7633	0.9238
Per-step logistic + position	0.7496	0.7395	0.7953	0.7606	0.7550	0.9061
Continuous CUSUM	0.7472	0.7347	0.7961	0.7481	0.7305	0.9016
Signed two-tap	0.7551	0.7385	0.8000	0.7497	0.7449	0.8884
EWMA, position-tuned	0.7443	0.7175	0.8070	0.7422	0.7264	0.8924
Tuned fixed gain	0.7727	0.7493	0.8224	0.7743	0.7658	0.9233
EWMA logistic + position	0.7716	0.7556	0.8210	0.7743	0.7649	0.9237
Colored kernel, balanced	0.7458	0.7259	0.8053	0.7490	0.7427	0.8894
Colored kernel, stationary	0.7403	0.7207	0.7977	0.7488	0.7425	0.8916

Table 5: Pooled AUROC for all 24 aggregation methods; higher is better. Each value is the mean over five outer problem folds. Gen. 7B and Math 1.5B denote local frozen-feature probes; PRM 7B denotes the frozen prefix-aware process reward model. The two tuned EWMA rows select memory by their stated training metric.

	MATH			GSM8K		
Method	Gen. 7B	Math 1.5B	PRM 7B	Gen. 7B	Math 1.5B	PRM 7B
Per-step	0.7200	0.7076	0.7619	0.7007	0.6932	0.8885
EWMA, fixed 0.5	0.7475	0.7321	0.7804	0.7201	0.7136	0.9063
White exponential, balanced	0.7523	0.7365	0.7802	0.7206	0.7153	0.9057
White exponential, stationary	0.7521	0.7360	0.7769	0.7207	0.7161	0.9063
EWMA, pooled-tuned	0.7448	0.7301	0.7811	0.7159	0.7078	0.9063
Finite LMMSE	0.7522	0.7362	0.7808	0.7202	0.7141	0.9058
Gaussian HMM	0.7513	0.7365	0.7808	0.7197	0.7129	0.9060
Position prior	0.5000	0.5000	0.5000	0.5000	0.5000	0.5000
Current score + prior	0.7200	0.7076	0.7619	0.7007	0.6932	0.8885
Fixed Wiener gain	0.7521	0.7360	0.7769	0.7207	0.7161	0.9063
AR(1) exponential, balanced	0.7523	0.7364	0.7802	0.7191	0.7117	0.9034
AR(1) exponential, stationary	0.7520	0.7362	0.7781	0.7197	0.7126	0.9053
White kernel, balanced	0.7524	0.7363	0.7805	0.7207	0.7153	0.9056
White kernel, stationary	0.7521	0.7360	0.7769	0.7207	0.7161	0.9063
Lag logistic	0.7456	0.7304	0.7771	0.7172	0.7094	0.9053
Lag logistic + position	0.7499	0.7341	0.7816	0.7192	0.7121	0.9066
Per-step logistic + position	0.7200	0.7076	0.7619	0.7007	0.6932	0.8885
Continuous CUSUM	0.7165	0.7133	0.7527	0.6920	0.6715	0.8895
Signed two-tap	0.7411	0.7261	0.7744	0.7155	0.7091	0.9013
EWMA, position-tuned	0.7523	0.7362	0.7809	0.7204	0.7162	0.9060
Tuned fixed gain	0.7526	0.7364	0.7814	0.7204	0.7150	0.9062
EWMA logistic + position	0.7496	0.7335	0.7803	0.7197	0.7137	0.9063
Colored kernel, balanced	0.7495	0.7337	0.7816	0.7172	0.7090	0.9016
Colored kernel, stationary	0.7506	0.7344	0.7812	0.7180	0.7096	0.9037

Table 6: Within-position AUROC for all 24 aggregation methods; higher is better. Each value is the mean over five outer problem folds. Gen. 7B and Math 1.5B denote local frozen-feature probes; PRM 7B denotes the frozen prefix-aware process reward model. The two tuned EWMA rows select memory by their stated training metric.

	Pooled: colored minus pooled-tuned			Within: colored minus position-tuned		
Setting	Difference	Lower	Upper	Difference	Lower	Upper
MATH Gen. 7B	−0.0121	−0.0262	+0.0020	−0.0028	−0.0073	+0.0016
MATH Math 1.5B	−0.0148	−0.0285	−0.0010	−0.0025	−0.0081	+0.0030
MATH PRM 7B	−0.0031	−0.0094	+0.0033	+0.0007	−0.0024	+0.0038
GSM8K Gen. 7B	−0.0013	−0.0050	+0.0024	−0.0032	−0.0110	+0.0047
GSM8K Math 1.5B	−0.0011	−0.0051	+0.0028	−0.0072	−0.0181	+0.0038
GSM8K PRM 7B	−0.0049	−0.0095	−0.0002	−0.0044	−0.0100	+0.0012

Table 7: Full colored balanced kernel relative to EWMA tuned for the evaluated metric. Bounds are simultaneous 95% intervals from the final family of 24 methods, two metrics, three reference rules, and six settings, as defined in Appendix M.5. All quantities are AUROC differences, not percentage changes.

pooled position-residualized score. Position eligibility requires at least 80 examples of each class in the outer-training set, before evaluating test scores. This retains MATH positions 1–10 in outer folds 0, 1, 3, and 4, and positions 1–9 in fold 2, where position 10 has 77 training clean examples. GSM8K positions 1–4 are eligible in every fold; the smallest held-out class cells contain 17 and 14 examples, respectively. Eligibility is fixed for every method and bootstrap draw. Both SNR metrics are averaged equally over the five outer folds. The reported within-position AUROC uses these same eligible positions, so it can differ slightly from the earlier metric that included every position with an eligible pair.

The new EWMA rules use the original grid $\rho \in \{0, 0.05, \ldots, 0.95, 0.98\}$ and the original working LLR stream. One maximizes pooled signed SNR on outer-training cross-fitted scores; the other maximizes within-position signed SNR on those scores. The frozen PRM scores bypass probe fitting as before. Test SNR and AUROC are then evaluated without further selection. The learned and prior-aware methods in the original suite retain their earlier AUROC-based training objectives; their SNR results describe transfer to a new metric.

N.2 Finite-start colored discrimination

For the fitted initial error probability p_0, transition rates, residual slope, and signal level, the finite-start moment calculation gives the conditional mean-gap vector $2g^{(t)}$ and balanced covariance $N^{(t)}$ at position t. We compute $a^{(t)} = (N^{(t)})^{-1} g^{(t)}$ by the linear-time recursion and normalize $\kappa_j^{(t)} = a_j^{(t)} / a_0^{(t)}$. With lag zero denoting the current observation, the score is

$$h_t^{\text{finite}} = \sum_{j=0}^{t} \kappa_j^{(t)} \left[\ell_{t-j} - \widehat{D} \left\{ 2p_{t-j} - 1 + g_j^{(t)} (1 - 2p_t) \right\} \right].$$

The subtracted vector is the model midpoint of the two conditional class means. All quantities in it come from training estimates and the observed position, not future labels. The current coefficient is one. Coefficients and midpoints are cached by position and reused across problems. Independent path enumeration verifies the adapter against directly calculated conditional means and dense Fisher solutions.

The midpoint convention affects comparisons across positions but leaves within-position SNR and AUROC unchanged. Thus the pooled finite-minus-stationary comparison includes both the changed kernel and the specified position offset. It is not a comparison of posterior probabilities or an isolated estimate of the effect of initialization.

N.3 Predicted and observed objectives

For each principal linear score, we evaluate its actual finite coefficient vector under two training-fitted moment models. The stationary AR(1) model uses $g_j = \widehat{\nu}^j$ and its stationary class-conditional covariance; the finite-start AR(1) model uses $g^{(t)}$ and $N^{(t)}$. The predicted SNR is $2g^\top \kappa / \sqrt{\kappa^\top N \kappa}$. Predictions are averaged over the same supported positions and held-out pair weights as empirical within-position SNR. This comparison holds the available history window fixed; it does not replace missing pre-trace observations by an infinite observed history.

Table 8 follows the chain from stationary working-model SNR to empirical SNR and then AUROC. Their differences reveal distinct aspects of transfer. For example, the MATH General-7B stationary model predicts within-position SNR 0.983 for the colored kernel and 0.973 for position-SNR-tuned EWMA, whereas their observed values are 0.995 and 1.012. On GSM8K PRM, the corresponding predictions are 1.801 and 1.773, while observed SNR is 2.042 and 2.086. The model's ranking can therefore disagree with the empirical ranking even before changing the evaluation objective to AUROC.

Using the fitted finite initial law does not fully close this gap. On MATH General-7B, the finite moment model predicts SNR 0.943 for the stationary colored score and 0.944 for the finite Fisher score. Their held-out values are 0.995 and 0.990. On GSM8K PRM, the finite-model predictions are 1.793 and 1.795, while the held-out values are 2.042 and 2.050. Correcting initialization changes a specific component of the working model, but its predictions still depend on the fitted emission and transition assumptions.

N.4 PRESPECIFIED COMPARISONS AND UNCERTAINTY

The bridge uses 20,000 paired whole-problem bootstrap resamples within the fixed outer folds. The five primary contrasts are stationary colored minus current-only, stationary colored minus white kernel, stationary colored minus metric-matched SNR-tuned EWMA, finite colored minus stationary colored, and finite colored minus metric-matched SNR-tuned EWMA. Each is evaluated for pooled and within-position SNR in all six settings. A bootstrap max-t statistic covers the ten contrasts within a setting; its $1 - 0.05/6$ quantile and Bonferroni correction cover the six settings without assuming independence. Position eligibility is held fixed and problem draws are shared across methods. AUROC is reported descriptively for the added methods and is not inserted into the earlier AUROC family.

These intervals condition on the fitted models, selected hyperparameters, fixed folds, and sampled problems. They do not include model refitting or turn an adaptive follow-up study into a confirmatory experiment. The bootstrap is conditional on jointly defined statistics. In each GSM8K setting, draw 13744 (zero-based) gives outer fold 4, position 4 only one resampled clean example and 75 error examples, so its unbiased class variance is undefined. We remove this same draw jointly across the entire contrast family in all three GSM8K settings, leaving 19,999 valid draws. All 20,000 MATH draws remain valid. Eligibility and the estimator are unchanged; no variance floor or contrast-dependent deletion is used. The draw index, fold, position, and counts are retained in the metadata. Marginal intervals for all 27 methods and per-position SNR values accompany the predictions in the evaluation artifacts.

The pooled SNR comparison rules out a purely AUROC-objective explanation for two local failures. On MATH General-7B, stationary colored SNR is 0.9559 versus 1.0098 for SNR-tuned EWMA, a difference of -0.0539 with simultaneous interval $[-0.0959, -0.0118]$. On MATH Math-1.5B the values are 0.8701 and 0.9276, giving -0.0574 with interval $[-0.1040, -0.0109]$. Here the fitted prescription falls short on the matched empirical SNR objective itself. This establishes a transfer failure of the working model and finite-score convention, without identifying a single cause.

Finite-start correction has mixed effects. It increases GSM8K PRM within-position SNR by 0.0084 relative to stationary colored filtering, with simultaneous interval $[0.0001, 0.0166]$. Its difference from position-SNR-tuned EWMA remains -0.0363 with interval $[-0.0826, 0.0101]$. On MATH General-7B it decreases within-position SNR by 0.0051, with interval $[-0.0099, -0.0003]$. The other four finite-minus-stationary within-position intervals include zero. These results support a limited improvement on the GSM8K PRM stream, not uniform superiority of the finite-start prescription.

An independent bootstrap-resolution check repeats the earlier GSM8K PRM AUROC comparison with 10,000 draws and a different seed. The stationary-colored-minus-pooled-tuned within-position AUROC interval is $[-0.009012, -0.000317]$, retaining the direction of the original near-zero adverse interval. This stability check uses the earlier AUROC family and does not replace it or alter the separate five-contrast SNR family.

N.5 RESIDUAL COVARIANCE AND CONDITIONAL MEAN DIAGNOSTICS

We examine the observation moments directly rather than infer their correctness from AUROC. For every fold, class means, residual variance, α, ν, and the current raw-score mean gap δ are fitted on training scores. Training and held-out diagnostics then compare residual covariance and correlation at lags 1–5 with $R\alpha^j$ and α^j. Historical raw-score mean gaps condition on the current label and are compared with $\delta\nu^j$. We also report the gap relative to the current-score gap on the same eligible anchor set. Pair counts and class support accompany every estimate.

Table 9 shows that fitting lag one does not recover the higher-lag shape. The MATH General-7B held-out residual correlations at lags two and five are 0.403 and 0.330, compared with AR(1) predictions 0.220 and 0.023. For Math-1.5B they are 0.464 and 0.366, versus 0.259 and 0.034. MATH PRM also has a heavier empirical residual tail: its lag-five correlation is 0.158 versus 0.021. These are pair-count-weighted averages of fold estimates. Covariances normalized by training residual variance, including lag zero, are reported alongside correlations in the artifacts.

Conditional mean-gap behavior also differs across streams. On MATH General-7B the lag-five gap, divided by the fitted current gap, is 0.800, compared with $\nu^5 = 0.562$ averaged over folds.

On MATH PRM it is 0.490 against the same prediction. Thus different streams exhibit different departures from the geometric mean-gap prescription. GSM8K lag-four and lag-five conditional-gap cells contain as few as eight and two examples in one anchor class within a fold. We retain their descriptive values and support flags, but do not use these sparse cells to draw a general conclusion about mean-gap decay.

After class centering, residual means still vary with the observed transition. In training-standard-deviation units, the MATH General-7B current residual averages -0.312 on clean-to-error transitions and 0.127 on error-to-error transitions. The MATH PRM pattern reverses, with means 0.534 and -0.030. These cells contain 411 and 2,749 transitions, respectively. On GSM8K PRM, error-to-clean transitions have mean residual 1.105 training standard deviations, but only 43 such transitions occur across 42 problems. Position-by-class means and excess means relative to the held-out current-class mean are retained as separate diagnostics, so calibration offsets can be distinguished from variation across those subsets.

The observed higher-lag covariance and conditional-mean patterns depart from the fitted path-conditional moment model. Together with the matched-SNR shortfalls, they show that objective mismatch alone is insufficient to explain the empirical transfer results. They do not causally attribute a performance difference to a particular covariance lag, transition effect, or initialization choice. No intervention isolates those departures here, and the finite-start correction changes the model prescription while leaving the observed emission process unchanged.

N.6 Direct first-column directions without an AR(1) fit

We also apply the first-column Fisher criterion directly to held-out moments. At a training-eligible current position and lag $j \le 5$, let Δ_0 and Δ_j be the current and historical raw-score gaps between current-label classes. Let N_{00} and N_{0j} be entries of the balanced within-current-class covariance matrix on the same paired observations, with sample variances and covariances using denominator $n_i - 1$. For positive Δ_0 and N_{00}, define

$$\widehat{\gamma}_j = \frac{\Delta_j}{\Delta_0} - \frac{N_{0j}}{N_{00}}.$$

The derivative of the sample signed SNR of $r_t + b r_{t-j}$ at $b = 0$ is

$$\frac{\Delta_j N_{00} - \Delta_0 N_{0j}}{N_{00}^{3/2}}.$$

Thus positive $\widehat{\gamma}_j$ indicates an estimated local SNR improvement from adding that lag, and negative $\widehat{\gamma}_j$ indicates improvement from subtracting it. This calculation does not fit α, assume geometric moments, or infer the coefficients of the full multivariate optimum.

The analysis includes every lag $1 \le j \le \min(5, t-1)$ at a training-eligible one-based position t, pairing exact positions within a problem. All 600 fold–position–lag cells have a positive current gap and valid covariance estimates. Of their directional estimates, 580 are positive, 20 negative, and none zero. We report every cell. As a descriptive support flag, a position–lag group must have valid, unanimous signs in all five folds and at least 20 paired examples per class in each fold. There are 91 such groups, all positive: 30, 25, 3, 2, 28, and 3 in the order MATH General-7B, MATH Math-1.5B, GSM8K General-7B, GSM8K Math-1.5B, MATH PRM, and GSM8K PRM. The twenty negative cell estimates do not form a group satisfying this support rule.

This is a sample-moment application of the covariance criterion, not a significance test. It provides no consistently supported subtraction direction under the stated descriptive rule, while retaining all negative estimates. Positive directions do not establish AUROC gains or validate the path-conditional model, and five-fold sign agreement is not a confidence statement. Central numerical differentiation independently checks the sample-SNR derivative, with maximum absolute discrepancy 8.71×10^{-11}.

Setting	Current	Colored	Finite	Tuned EWMA
MATH Gen. 7B	0.836/0.853/0.722	0.983/0.995/0.751	0.978/0.990/0.750	0.973/1.012/0.754
MATH Math 1.5B	0.816/0.801/0.710	0.941/0.933/0.737	0.936/0.930/0.736	0.923/0.956/0.739
MATH PRM 7B	1.119/1.024/0.762	1.308/1.154/0.781	1.303/1.153/0.781	1.299/1.146/0.780
GSM8K Gen. 7B	0.916/0.747/0.702	0.958/0.816/0.719	0.956/0.816/0.719	0.945/0.830/0.722
GSM8K Math 1.5B	0.857/0.708/0.695	0.891/0.765/0.711	0.889/0.762/0.710	0.868/0.783/0.718
GSM8K PRM 7B	1.732/1.903/0.889	1.801/2.042/0.902	1.799/2.050/0.903	1.773/2.086/0.907

Table 8: Each cell reports stationary AR(1) working-model SNR / empirical SNR / empirical AUROC on the same training-eligible positions. The actual finite coefficient prefix is used for the model calculation. EWMA is selected by training within-position SNR. Prediction and empirical values use the same position pair weights and outer-fold averaging.

	Residual correlation: observed / AR(1)			Mean gap: observed / model
Setting	Lag 1	Lag 2	Lag 5	Lag 3
MATH Gen. 7B	0.469/0.469	0.403/0.220	0.330/0.023	0.856/0.707
MATH Math 1.5B	0.535/0.508	0.464/0.259	0.366/0.034	0.839/0.707
MATH PRM 7B	0.461/0.460	0.295/0.211	0.158/0.021	0.608/0.707
GSM8K Gen. 7B	0.301/0.295	0.203/0.087	0.207/0.002	0.613/0.280
GSM8K Math 1.5B	0.384/0.325	0.286/0.106	0.158/0.004	0.539/0.280
GSM8K PRM 7B	0.285/0.270	0.127/0.073	0.039/0.001	0.559/0.280

Table 9: Held-out observation diagnostics, averaged across folds with pair-count weights. The historical observation mean gap conditions on the current label and is divided by the fitted current gap; its model prediction is $\widehat{\nu}^{j}$. These are descriptive moment comparisons, not causal effects. Lag-three mean-gap cells contain at least 22 examples per anchor class in every fold; sparse GSM8K lag-four and lag-five gaps are omitted from this table.

	MATH			GSM8K		
Method	Gen. 7B	Math 1.5B	PRM 7B	Gen. 7B	Math 1.5B	PRM 7B
Per-step	0.9110	0.8476	1.1193	0.8997	0.8760	1.7344
EWMA, fixed 0.5	1.0113	0.9284	1.2553	0.9380	0.8933	1.8428
White exponential, balanced	0.9083	0.8217	1.1921	0.9245	0.8760	1.8521
White exponential, stationary	0.8459	0.7625	1.0674	0.9155	0.8674	1.8351
EWMA, pooled-AUROC-tuned	1.0083	0.9270	1.2533	0.9471	0.9074	1.8428
Finite LMMSE	1.0434	0.9766	1.3431	1.0781	1.0310	2.1835
Gaussian HMM	1.0510	0.9670	1.3256	1.0796	1.0301	2.2687
Position prior	0.5711	0.5711	0.5711	0.7239	0.7239	0.7239
Current score + prior	0.9329	0.8900	1.1547	1.0151	0.9829	2.0025
Fixed Wiener gain	0.5954	0.5931	0.6067	1.0231	0.9753	2.1461
AR(1) exponential, balanced	0.9191	0.8375	1.1963	0.9455	0.9036	1.8534
AR(1) exponential, stationary	0.8807	0.8023	1.1146	0.9424	0.9005	1.8543
White kernel, balanced	0.9160	0.8286	1.2011	0.9251	0.8769	1.8529
White kernel, stationary	0.8459	0.7625	1.0674	0.9155	0.8674	1.8351
Lag logistic	0.9984	0.9185	1.1807	0.9535	0.9021	1.7825
Lag logistic + position	1.0790	1.0044	1.3428	1.0494	0.9910	2.1741
Per-step logistic + position	0.9696	0.9190	1.1791	0.9989	0.9572	2.0121
Continuous CUSUM	0.7262	0.7307	1.0559	0.7792	0.7266	1.6854
Signed two-tap	0.9917	0.9143	1.2039	0.9459	0.9103	1.8373
EWMA, position-AUROC-tuned	0.9568	0.8456	1.2433	0.9097	0.8146	1.8019
Tuned fixed gain	1.0389	0.9503	1.3421	1.0657	1.0161	2.2055
EWMA logistic + position	1.0747	1.0046	1.3340	1.0560	1.0002	2.1684
Colored kernel, balanced	0.9559	0.8701	1.2373	0.9416	0.9011	1.8472
Colored kernel, stationary	0.9289	0.8473	1.1924	0.9405	0.8997	1.8529
EWMA, pooled-SNR-tuned	1.0098	0.9276	1.2544	0.9456	0.9075	1.8546
EWMA, position-SNR-tuned	0.8755	0.7227	1.2465	0.9039	0.8154	1.8226
Finite colored Fisher	0.8979	0.8152	1.2332	0.9366	0.8903	1.8625

Table 10: Pooled signed balanced SNR for all 27 methods; higher is better. Values average the five outer problem folds. The original 24 methods retain their original fitted scores and selection criteria. The last three rows are added by the SNR bridge. Position SNR uses training-eligible positions only, and is the pair-weighted mean of position-specific SNR values.

	MATH			GSM8K		
Method	Gen. 7B	Math 1.5B	PRM 7B	Gen. 7B	Math 1.5B	PRM 7B
Per-step	0.8532	0.8012	1.0244	0.7473	0.7079	1.9026
EWMA, fixed 0.5	0.9825	0.9206	1.1442	0.8272	0.7799	2.0848
White exponential, balanced	1.0109	0.9538	1.1424	0.8301	0.7837	2.0786
White exponential, stationary	1.0116	0.9554	1.1218	0.8310	0.7848	2.0857
EWMA, pooled-AUROC-tuned	0.9683	0.9074	1.1485	0.8120	0.7583	2.0848
Finite LMMSE	1.0070	0.9512	1.1474	0.8283	0.7801	2.0796
Gaussian HMM	0.9676	0.9287	1.1247	0.8115	0.7686	2.1515
Position prior	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
Current score + prior	0.8532	0.8012	1.0244	0.7473	0.7079	1.9026
Fixed Wiener gain	1.0116	0.9554	1.1218	0.8310	0.7848	2.0857
AR(1) exponential, balanced	1.0105	0.9527	1.1429	0.8233	0.7735	2.0583
AR(1) exponential, stationary	1.0115	0.9544	1.1289	0.8256	0.7757	2.0741
White kernel, balanced	1.0106	0.9531	1.1455	0.8300	0.7834	2.0776
White kernel, stationary	1.0116	0.9554	1.1218	0.8310	0.7848	2.0857
Lag logistic	0.9716	0.9099	1.1237	0.8161	0.7629	2.0744
Lag logistic + position	0.9947	0.9333	1.1528	0.8235	0.7725	2.0880
Per-step logistic + position	0.8532	0.8012	1.0244	0.7473	0.7079	1.9026
Continuous CUSUM	0.8537	0.8513	0.9641	0.6980	0.5872	1.7990
Signed two-tap	0.9511	0.8874	1.1047	0.8105	0.7631	2.0385
EWMA, position-AUROC-tuned	1.0090	0.9515	1.1481	0.8310	0.7833	2.0826
Tuned fixed gain	1.0106	0.9535	1.1493	0.8291	0.7831	2.0838
EWMA logistic + position	0.9928	0.9301	1.1434	0.8261	0.7784	2.0847
Colored kernel, balanced	0.9949	0.9333	1.1537	0.8161	0.7648	2.0416
Colored kernel, stationary	1.0007	0.9394	1.1538	0.8192	0.7674	2.0610
EWMA, pooled-SNR-tuned	0.9731	0.9094	1.1456	0.8068	0.7582	2.0705
EWMA, position-SNR-tuned	1.0124	0.9558	1.1465	0.8299	0.7829	2.0862
Finite colored Fisher	0.9898	0.9296	1.1530	0.8160	0.7617	2.0500

Table 11: Within-position signed balanced SNR for all 27 methods; higher is better. Values average the five outer problem folds. The original 24 methods retain their original fitted scores and selection criteria. The last three rows are added by the SNR bridge. Position SNR uses training-eligible positions only, and is the pair-weighted mean of position-specific SNR values.