State complexity of rare-switching tracking: universal bounds and sharp counter tradeoffs

Samuil Petkov
Department of Physics, École normale supérieure, Université PSL,
Paris, France
samuil.petkov@phys.ens.psl.eu

1 November 2026

hexagon:2611.00656v1

Abstract

We study the number of internal states required to track a symmetric binary Markov signal observed through a fixed binary symmetric channel. For homogeneous recursive controllers with independent private randomization, a posterior-covering converse gives the sharp leading state coefficient in the iterated nearly-Bayes-optimal limit, uniformly over controller kernels and stationary recurrent mixtures. At a fixed relative excess error, we also determine the exact leading coefficient within symmetric observation-driven birth–death counters with a prescribed sign decoder. Their acceptance probabilities may vary across states and tend to zero with the switching probability. A stationary flow argument links a confidence barrier to actual switch-time mass, and a discounted passage estimate controls repeated switching. For general MAP-decoded controllers, a point certificate for the same posterior covers shows that this counter coefficient is optimal at every fixed relative excess above an explicit threshold depending only on the channel; below that threshold the general frontier remains open, and the gap is narrowed. State cardinality is the resource; no thermodynamic work claim is made.

Mathematics Subject Classification (2020). Primary 62M20; Secondary 94A15, 60J10.

Keywords. finite-memory inference; hidden Markov models; rare transitions; causal filtering; state complexity; birth–death chains.

1 Introduction and results

Limited memory constrains both the confidence accumulated from noisy observations and the speed of adapting after a change. We quantify this tradeoff for a binary signal in the rare-switching regime. The model counts internal controller states rather than an observation window or an uncounted real posterior. Independent private coins are allowed, but an external clock and access to the past observation record are not.

Finite-state randomized inference has classical foundations in Hellman–Cover [1]. Bounded-memory decisions in changing environments were studied by Monte–Said [8] and Chatterjee–Guryev–Hu [9]. The rare-transition causal Bayes error scale is classical [6, 7]. We retain those results as antecedents and state our conclusions with their precise controller and limit quantifiers. The proof below also explains the stationary reset translation of the classical static bound.

Related work on finite-memory testing and estimation. Cover [2] studied hypothesis testing with statistics taking finitely many values. Hellman and Cover [1] analysed hypothesis testing by a finite-state, possibly randomized, automaton for a static hypothesis; their bound is the source of the reset converse in Appendix A. Hellman and Cover [3] quantified the memory saved by randomization, and Hellman [4] studied the effect of randomization on finite-memory decision schemes further. Wilson [10] studied optimal finite-state memory in a two-state decision problem as a source of biases in information processing. The survey of Berg, Ordentlich and Shayevitz [5] reviews inference under memory constraints, covering hypothesis testing, parameter estimation, and distribution property testing and estimation. Those

works concern a fixed hypothesis or parameter, with memory measured against asymptotic error. Here the hidden signal switches at a vanishing rate, and the comparison is with the causal Bayes error.

1.1 Model and conventions

A stationary fair hidden bit switches at each discrete update with probability $r \in (0, 1/2)$. Observations are conditionally independent through a binary symmetric channel of error $\eta \in (0, 1/2)$. A homogeneous recursive device updates from its stored state and the current observation, using independent private randomization. Its decision is made after that update. Let $e_m(r, \eta)$ be the infimum of stationary MAP error over all such kernels with at most m states and all their stationary joint laws, including recurrent mixtures. Let e_B be full causal Bayes error, and put

$$m_\epsilon(r, \eta) = \min\{m : e_m(r, \eta) \le (1 + \epsilon)e_B(r, \eta)\}.$$

$$p = 1 - \eta, \quad q = \eta, \quad R = p/q, \quad a = \log R,$$
$$D = (p - q)a, \qquad L = \log(1/r).$$

An empty minimum is infinite. All logarithms are natural, and time is counted in observation updates. Unless explicitly varied, η is fixed.

Theorem 1 (Universal nearly-Bayes state coefficient)**.** *For the controller class just defined,*

$$\lim_{\epsilon \downarrow 0} \liminf_{r \downarrow 0} \frac{m_\epsilon}{L} = \lim_{\epsilon \downarrow 0} \limsup_{r \downarrow 0} \frac{m_\epsilon}{L} = \frac{2}{a}.$$

No inner limit at every fixed ϵ, diagonal $\epsilon(r)$, or uniform limit as $\eta \to 1/2$ is asserted.

Our second result concerns the symmetric nearest-neighbor counters defined in Section 10. Their stored state is an integer in $\{-n, \dots, n\}$. Every move is driven by the current observation and accepted with a probability in $(0, 1]$ that depends on the state and on the direction of the move. These probabilities may depend on r and approach zero. Nonzero states are decoded by sign, with a fixed tie convention at zero. This decoder need not coincide with stationary MAP decoding. Write m_ϵ^{BD} for the least cardinality in this prescribed class attaining error at most $(1 + \epsilon)e_B$.

Theorem 2 (Sharp fixed-tolerance frontier for sign-decoded counters)**.** *For each fixed $\epsilon > 0$,*

$$\lim_{r \downarrow 0} \frac{m_\epsilon^{\mathrm{BD}}}{L} = \frac{2}{b_\epsilon}, \qquad b_\epsilon\big(p(p/q)e^{-b_\epsilon} - q\big) = \frac{D}{1 + \epsilon}, \qquad a < b_\epsilon < 2a.$$

This theorem optimizes over inhomogeneous acceptance profiles in the specified architecture. It does not identify the general MAP frontier.

Our third result returns to the general controller class. Put $\Gamma(b) = b(pRe^{-b} - q)$, so that b_ϵ above is the root of $\Gamma(b_\epsilon) = D/(1 + \epsilon)$ in $(a, 2a)$.

Theorem 3 (General fixed-tolerance frontier above a channel threshold)**.** *Let $\zeta_\eta \in (0, a)$ be the unique root of $\zeta + e^{\zeta - a} = 1$, and put*

$$1 + \epsilon_\eta^* = \frac{D(1 - \zeta_\eta)}{q\zeta_\eta(a + \zeta_\eta)}, \qquad \varkappa_\eta = \frac{D(1 - \zeta_\eta)}{q\zeta_\eta^2}.$$

Then $\epsilon_\eta^* > 0$, and for the general controller class:

(i) for every fixed $\epsilon \ge \epsilon_\eta^*$,

$$\lim_{r \downarrow 0} \frac{m_\epsilon}{L} = \frac{2}{b_\epsilon};$$

(ii) for every fixed $0 < \epsilon < \epsilon_\eta^*$,

$$\liminf_{r \downarrow 0} \frac{m_\epsilon}{L} \ge \frac{2(\varkappa_\eta - 1 - \epsilon)}{a\varkappa_\eta}.$$

For instance $\epsilon_\eta^* \approx 0.86, 0.67, 0.53, 0.43$ for $\eta = 0.1, 0.2, 0.3, 0.4$. The lower bound in (ii) is combined with the earlier ones in Section 7; a gap remains only for $0 < \epsilon < \epsilon_\eta^*$.

The universal converse charges coverage of two posterior-confidence regions. The restricted converse additionally uses stationary edge flow to locate a large likelihood barrier before almost all the switch-time mass. Sharp passage estimates then match the constant-profile counter construction. For Theorem 3, the same two covers are tested at a single point in each wide gap between centers instead of being integrated; each center is charged by at most one test point, and the resulting cost per gap is exactly that of the constant-acceptance counter. Full proofs follow; the classical reset and fixed-parameter geometric-horizon translations are included in an appendix.

2 Posterior-covering converse

Write $p = 1 - \eta$, $q = \eta$, $a = \log(p/q)$, $D = (p - q)a$ and $\rho = 1 - 2r$. Let $A < B < -a$, $H = B - A$, and $t_* = (1 + e^{-A})^{-1} > r$.

Lemma 4 (Posterior-covering inequality)**.** *For a stationary device of error e, suppose*

$$c = 1 - 2e - \frac{2e}{t_* - r} > 0.$$

Then every device with at most m states satisfies

$$\frac{2D}{r}(r + \rho e) \geq 2cH + c(m + 2)\psi\left(\left[\frac{2H}{m + 2} - a\right]_+\right), \qquad \psi(x) = e^x - 1 - x.$$

This includes arbitrary randomization probabilities and stationary mixtures.

Proof. **Posterior contraction and observation gain.** For occupied memory states put $\pi_i = \Pr(Y = i)$, $\theta_i = \Pr(S = 1 \mid Y = i)$, $v_i = r + \rho\theta_i$, and $z_i = \operatorname{logit} v_i$. Let U be the posterior after the new observation but before memory compression. Independent private randomness implies

$$\Pr(S_t = 1 \mid Y_t) = \mathbb{E}[U \mid Y_t].$$

Jensen and stationarity apply to every convex call function $(x - t)_+$. For one binary observation its exact gain is

$$\mathbb{E}[(U - t)_+ \mid v] - (v - t)_+ = v(1 - t)k(\operatorname{logit} v - \operatorname{logit} t),$$

where

$$k(s) = \begin{cases} p - qe^{-s}, & -a \leq s \leq 0, \\ pe^{-s} - q, & 0 \leq s \leq a, \\ 0, & |s| > a. \end{cases} \qquad \int_{\mathbb{R}} k = D.$$

For $r < t < 1/2$, states with $\theta_i \geq 1/2$ contribute in total $r\sum_{\theta_i \geq 1/2} \pi_i(2\theta_i - 1) = r(1/2 - e)$ to the switching loss of the call function, since $\sum_i \pi_i\theta_i = 1/2$. Any adverse contribution requires $\theta_i > t - r$ and has magnitude at most r per unit mass. Their mass is at most $e/(t - r)$. The total switching loss is therefore at least

$$r\left(\frac{1}{2} - e - \frac{e}{t - r}\right).$$

Two covers and their area budget. Define globally truncated functions

$$F_-(x) = \frac{2}{r}\sum_{z_i < 0} \pi_i v_i k(z_i - x), \qquad F_+(x) = \frac{2}{r}\sum_{z_i > 0} \pi_i(1 - v_i)k(-z_i - x).$$

For $x \in [A, B]$, the omitted centers cannot contribute because $B < -a$. Combining posterior contraction, gain and switching loss gives $F_\pm(x) \geq c$. Exchanging the hidden labels proves the plus statement. Their total area obeys

$$\int_{\mathbb{R}} (F_- + F_+) \leq \frac{2D}{r}\sum_i \pi_i \min(v_i, 1 - v_i) = \frac{2D}{r}(r + \rho e).$$

Gap penalty. For a positive kernel sum $F \geq c$ on $[A, B]$, insert the endpoints as zero-weight fictitious centers and split at actual internal centers. In a center-free gap (ℓ, u) of length $h > a$, every left center has expired once $x > \ell + a$. All active terms then have the form $w_j(pe^{x-z_j} - q)$, so $F' \geq F$ almost everywhere. Continuity at support entrances gives

$$F(x) \geq ce^{x-(\ell+a)}, \qquad \ell + a \leq x \leq u,$$

and hence the excess area over this gap is at least $c\psi((h-a)_+)$. For $h \leq a$ the excess is nonnegative.

Each occupied memory-state index contributes to at most one cover. Thus the combined number of internal centers, counted separately within each cover, is at most m; reflected coordinate values may coincide. Their combined gap count g is at most $m + 2$, and their combined interval lengths are $2H$. Convexity of ψ yields excess area at least

$$cg\,\psi\left(\left[\frac{2H}{g} - a\right]_+\right).$$

This expression decreases with g: where its argument is positive, its derivative after removing c is $e^x(1 - x - a) + a - 1 \leq 0$ for $x \geq 0$. Substituting $g \leq m + 2$ and the total-area budget proves the lemma. □

Corollary 5. *Under the hypotheses of Lemma 4,*

$$\frac{2D}{r}(r + \rho e) \geq 2cH + \frac{c}{2(m+2)}[2H - a(m+2)]_+^2. \tag{1}$$

Proof. Since $\psi(x) \geq x^2/2$ for $x \geq 0$, the claim follows from Lemma 4. □

3 A causal counter witness with controlled adaptation

Put $d = p - q > 0$ and $L = \log(1/r)$. Use the ordinary saturating counter on $\{-n, \ldots, n\}$: add the current observation's sign and suppress moves beyond the endpoints. Define

$$h_+(i) = \mathbf{1}_{\{i<0\}} + \tfrac{1}{2}\mathbf{1}_{\{i=0\}}, \qquad h_- = 1 - h_+,$$

where $+$ denotes hidden state 1. Its stationary sign-rule error $e_{\mathrm{s}} = \mathbb{E}h_{S_t}(Y_t)$ bounds the optimal error $e(K)$ from above; the half-loss at zero adds no memory state.

Proposition 6. *With $n = \lceil L/a \rceil$ the saturating counter has $m = 2\lceil L/a \rceil + 1$ states and, as $r \downarrow 0$ with η fixed, $e(K) \leq rL/D + O_\eta(r)$. Moreover it has a stationary realization that is a function of past observations alone, so that $e_B \leq e(K)$.*

Proof. Hold the hidden state at $+$. The resulting kernel P_+ has invariant weights proportional to R^i. Its invariant sign error therefore satisfies

$$e_0 \leq C_0 R^{-n}, \qquad C_0 = \frac{1}{2} + \frac{1}{R-1}.$$

For the first hitting time T_n of n, define

$$G_+(y) = \mathbb{E}_y^+ \sum_{k<T_n} (h_+(Y_k) - e_0), \qquad G_+(n) = 0.$$

First-step conditioning gives $(I - P_+)G_+ = h_+ - e_0$ away from n. The invariant mean of both sides is zero, so the residual also vanishes at n, whose invariant mass is positive. Define G_- by reflection.

Before hitting n, the counter's drift is at least d. Stopping at $T_n \wedge N$ gives $d\mathbb{E}_y(T_n \wedge N) \leq n - y \leq 2n$. Monotone convergence yields $\mathbb{E}_y T_n \leq 2n/d$ and hence

$$G_s(y) \geq -2ne_0/d.$$

From a negative state, reaching zero takes expected time at most n/d. Starting nonnegative, couple the counter before T_n above a walk W started at zero with increments $+1, -1$ of probabilities p, q. The lower reflecting boundary only increases its increments. Thus its subsequent wrong-side occupation is bounded by

$$\sum_{k\geq 0} \Pr(W_k \leq 0) \leq \sum_{k\geq 0} \mathbb{E}e^{-aW_k/2} = \frac{1}{1-2\sqrt{pq}} =: C.$$

The denominator is positive because $p \neq q$. Dropping the nonpositive centering reward gives the uniform upper bound

$$-2ne_0/d \leq G_s(y) \leq n/d + C. \tag{2}$$

The actual update law and joint stationarity imply exactly

$$\begin{aligned}\mathbb{E}h_{S_t}(Y_{t-1}) - e_0 &= \mathbb{E}[G_{S_t}(Y_{t-1}) - G_{S_t}(Y_t)] \\ &= r\mathbb{E}[G_{-S}(Y) - G_S(Y)].\end{aligned}$$

Indeed, stationarity replaces the second expectation by $\mathbb{E}G_{S_{t-1}}(Y_{t-1})$, and the switch probability is r conditional on the preceding joint state. As $h_- = 1 - h_+$, the left error is $r + \rho e_{\mathrm{s}}$. Applying (2) gives

$$e(K) \leq e_{\mathrm{s}} \leq \frac{e_0 + r(n/d + C - 1 + 2ne_0/d)}{\rho}.$$

Choose $n = \lceil L/a \rceil$, so $e_0 \leq C_0 r$. Then

$$m = 2\lceil L/a \rceil + 1, \qquad e(K) \leq \frac{rL}{D} + O_\eta(r). \tag{3}$$

No initial side channel. A block of $2n$ positive observations sends every initial counter state to n. Conditional on the complete hidden path, disjoint such blocks are independent and each has probability at least $q^{2n} > 0$. Almost surely a coalescing block exists in the infinite past. Pulling the counter forward from that block defines a stationary counter state measurable from past observations alone. Consequently $e_B \leq e(K)$ for this witness; no correlated initialization is used as additional information. □

4 Stationary causal posterior and the Bayes scale

On the two-sided stationary hidden Markov model set

$$\mathcal{F}_t = \sigma(O_s : s \leq t), \qquad \Theta_t = \Pr(S_t = 1 \mid \mathcal{F}_t), \qquad e_B = \mathbb{E}\min(\Theta_t, 1 - \Theta_t).$$

Proposition 7 (Bayes scale)**.** *As $r \downarrow 0$ with η fixed, $e_B(r, \eta) \sim rL/D$.*

Proof. Finite-history conditional probabilities converge to Θ_t by martingale convergence, and their risks converge since the loss is bounded and continuous. The law of Θ_t is stationary, has mean $1/2$, and is invariant under prediction $V_t = r + \rho\Theta_{t-1}$ followed by the current BSC update.

The convex-call stationarity comparison therefore holds with equality before compression. The exact gain identity and switching-deficit estimate extend to integrals against the posterior law. Define the two globally truncated covers by replacing the finite sums above by these integrals. Nonnegative kernels permit Tonelli in their total-area bound. No finite-center count is used. Thus, for every band satisfying the earlier hypotheses,

$$\frac{2D}{r}(r + \rho e_B) \geq 2cH, \tag{4}$$

where c is computed using e_B.

Fix $0 < \kappa < 1/2$ and take

$$A = -(1-\kappa)L, \qquad B = -\kappa L, \qquad H = (1-2\kappa)L.$$

By (3), $e_B = O_\eta(rL)$. Moreover

$$t_* = \frac{r^{1-\kappa}}{1 + r^{1-\kappa}}, \qquad \frac{e_B}{t_* - r} = O_\eta(r^\kappa L) \longrightarrow 0.$$

For fixed κ, eventually $A < B < -a$, $t_* > r$ and $c > 0$, with $c \to 1$. Divide (4) by $2L$, let $r \downarrow 0$ first and then $\kappa \downarrow 0$, and combine with the causal upper witness. This proves the classical benchmark directly:

$$e_B(r, \eta) \sim \frac{r \log(1/r)}{D}. \tag{5}$$

□

This Bayes scale is classical [6, 7].

Corollary 8. *For every fixed $\epsilon > 0$, $m_\epsilon(r, \eta)$ is finite for all sufficiently small r and*

$$\limsup_{r \downarrow 0} \frac{m_\epsilon(r, \eta)}{L} \le \frac{2}{a}.$$

Proof. This follows from (3) and (5). □

5 The matched outer cardinality coefficient

Theorem 1 (restated). *For every fixed $\eta \in (0, 1/2)$, over all devices and stationary laws in the declared model,*

$$\boxed{\lim_{\epsilon \downarrow 0} \liminf_{r \downarrow 0} \frac{m_\epsilon(r, \eta)}{\log(1/r)} = \lim_{\epsilon \downarrow 0} \limsup_{r \downarrow 0} \frac{m_\epsilon(r, \eta)}{\log(1/r)} = \frac{2}{\log((1-\eta)/\eta)}.}$$

Lemma 9 (Attainment)**.** *For every m, the infimum $e_m(r, \eta)$ over the permitted stationary pairs is attained.*

Proof. Pad a device to exactly m states. Its two stochastic update matrices range over a compact product of simplices. Joint laws μ range over another compact simplex, and $\mu P_K = \mu$ is closed because P_K depends continuously on K. This feasible set is nonempty. The error $\sum_i \min\{\mu(0, i), \mu(1, i)\}$ is continuous, including zero-mass states. Thus the infimum e_m over the permitted stationary pairs is attained. This gives no continuity assertion as $r \downarrow 0$ and excludes no reducible kernel or stationary mixture. □

Proof of Theorem 1. Fix $\epsilon > 0$ and $0 < \kappa < 1/2$. For sufficiently small r, choose an attained optimizer at $m = m_\epsilon$. Its error is at most $(1 + \epsilon)e_B = O_\eta(rL)$. The same bands as above therefore have $c \to 1$. Put $M_r = (m_\epsilon + 2)/L$ and $h = 1 - 2\kappa$. Dividing (1) by L yields

$$\frac{2D(r + \rho e)}{rL} \ge 2ch + \frac{c}{2M_r}(2h - aM_r)_+^2.$$

The counter upper bound makes M_r bounded. A subsequence with $M_r \to 0$ is impossible: its right side diverges while its left side has limsup at most $2(1 + \epsilon)$. For any positive subsequential limit M, first let $r \downarrow 0$ along that subsequence and then $\kappa \downarrow 0$. The result is

$$(2 - aM)_+^2 \le 4\epsilon M.$$

Solving gives

$$\frac{2}{a + \epsilon + \sqrt{\epsilon^2 + 2a\epsilon}} \le \liminf_{r \downarrow 0} \frac{m_\epsilon}{L} \le \limsup_{r \downarrow 0} \frac{m_\epsilon}{L} \le \frac{2}{a}.$$

Only now send $\epsilon \downarrow 0$; the bounds squeeze both outer envelopes to $2/a$, proving the theorem. □

Remark (Exact boundary of the conclusion). The limit order is η fixed, then ϵ, κ fixed, then $r \downarrow 0$, then $\kappa \downarrow 0$, and finally $\epsilon \downarrow 0$. Neither an inner limit for every fixed $\epsilon > 0$ nor a result on arbitrary diagonals $\epsilon = \epsilon(r) \downarrow 0$ is asserted. There is no uniform conclusion as $\eta \to 1/2$. At fixed excess error, Theorem 3 identifies the frontier for $\epsilon \ge \epsilon_\eta^*$; it remains open for $0 < \epsilon < \epsilon_\eta^*$.

6 An attainable fixed-excess-error upper curve

The following construction gives an upper state coefficient at every fixed $\epsilon > 0$. Section 10 proves its optimality within the prescribed sign-decoded counter class, and Theorem 3 proves its optimality among all controllers for $\epsilon \geq \epsilon_\eta^*$; for smaller ϵ the general MAP frontier remains open.

Fix $\lambda \in (q/p, 1]$ and use states $-n, \dots, n$. Follow each observation by a unit step in its sign, accepting with probability λ if the current state has the opposite sign and with probability one otherwise. At zero always follow the observation; suppress steps beyond the endpoints. Acceptance coins are independent and private. Decode by sign, with an independent fair tie at zero. The stored integer is the entire internal state. Put

$$\Lambda_+ = R/\lambda, \quad \Lambda_- = \lambda R > 1, \quad g = \lambda p - q > 0, \quad b = \log \Lambda_+.$$

Proposition 10. *For fixed η and $\lambda \in (q/p, 1]$, with $n = \lceil L/b \rceil$, the counter has $m = 2n+1$ states, admits a causal stationary realization (a function of past observations and independent private coins), and*

$$\frac{m}{L} \longrightarrow \frac{2}{b}, \qquad e(K) \leq \frac{rL}{bg} + O_{\eta,\lambda}(r), \qquad \limsup_{r \downarrow 0} \frac{e(K)}{e_B} \leq \frac{D}{bg}.$$

Proof. For the birth–death chain with the hidden state held at plus, the stationary weights satisfy $\pi_i/\pi_0 = \Lambda_+^i$ and $\pi_{-i}/\pi_0 = \Lambda_-^{-i}$ for $1 \leq i \leq n$. Thus its sign error obeys

$$e_0 \leq \left(\frac{1}{2} + \frac{1}{\Lambda_- - 1}\right) \Lambda_+^{-n}.$$

Write $h_+(i) = \mathbf{1}_{\{i<0\}} + \frac{1}{2}\mathbf{1}_{\{i=0\}}$ and $h_- = 1 - h_+$. With T_n the first hitting time of n under the plus-fixed kernel P_+, define

$$G_+(y) = \mathbb{E}_y^+ \sum_{k<T_n} (h_+(Y_k) - e_0), \qquad G_+(n) = 0,$$

and define G_- by reflection. Before T_n the drift is at least g; stopping the position at $T_n \wedge N$ and then using monotone convergence gives $\mathbb{E}_y T_n \leq 2n/g$. The displayed expectation is therefore finite. First-step conditioning gives $(I - P_s)G_s = h_s - e_0$ away from the appropriate endpoint; averaging the residual under the fixed-sign invariant law proves the equation there also.

The hitting-time bound implies $G_s \geq -2ne_0/g$. From a negative state the expected time to zero is at most n/g. From a nonnegative start, until T_n, couple the plus-fixed device above an unrestricted walk started at zero with increment probabilities

$$\Pr(\Delta W = 1) = \lambda p, \qquad \Pr(\Delta W = -1) = q, \qquad \Pr(\Delta W = 0) = 1 - \lambda p - q.$$

Every device increment stochastically dominates this increment, including at zero and the lower endpoint; a common-uniform coupling preserves position domination. The upper reflecting boundary is excluded by stopping. For $z = \frac{1}{2}\log(\lambda p/q) > 0$,

$$\sum_{k \geq 0} \Pr(W_k \leq 0) \leq \sum_{k \geq 0} \mathbb{E} e^{-zW_k} = \frac{1}{(\sqrt{\lambda p} - \sqrt{q})^2} =: C_\lambda.$$

Dropping the nonpositive centering reward and then reflecting yields

$$-2ne_0/g \leq G_s(y) \leq n/g + C_\lambda.$$

For the actual post-observation stationary pair (S, Y), the Poisson equation and stationarity give exactly

$$r + \rho e_{\mathrm{s}} - e_0 = r\mathbb{E}[G_{-S}(Y) - G_S(Y)],$$

by the same pre-/post-observation identity used above. Consequently

$$e(K) \leq e_{\mathrm{s}} \leq \frac{e_0 + r(n/g + C_\lambda - 1 + 2ne_0/g)}{\rho}. \tag{6}$$

Choosing $n = \lceil L/b \rceil$ gives, for fixed η, λ,

$$\frac{m}{L} \longrightarrow \frac{2}{b}, \qquad e(K) \leq \frac{rL}{bg} + O_{\eta,\lambda}(r), \qquad \limsup_{r \downarrow 0} \frac{e(K)}{e_B} \leq \frac{D}{bg}. \tag{7}$$

Causal stationary realization. A block of $2n$ positive observations whose independent acceptance uniforms are all at most λ sends every initial state to n. Conditional on the complete hidden path, disjoint such blocks are independent, each with probability at least $(\eta\lambda)^{2n} > 0$. Almost surely a coalescing block exists in the infinite past. Pulling forward from it defines the stationary device as a function of past observations and independent private coins. These coins convey no hidden-state information beyond observations, so the full causal Bayes risk lower-bounds this randomized device's error. No initial side channel or external clock is used. □

Proposition 11. *For every fixed $\epsilon > 0$ there is a unique $b_\epsilon \in (a, 2a)$ with $b_\epsilon(pRe^{-b_\epsilon} - q) = D/(1+\epsilon)$, and $\limsup_{r\downarrow 0} m_\epsilon(r,\eta)/L \le U_\epsilon := 2/b_\epsilon < 2/a$.*

Proof. For an explicit attainable curve, parameterize $\lambda = Re^{-b}$, $a \le b < 2a$, and let

$$\Gamma(b) = b(pRe^{-b} - q).$$

Here $\Gamma(a) = D$, $\Gamma(2a) = 0$, and Γ is strictly decreasing on $[a, 2a)$. Indeed $\Gamma'(a) = d - pa < 0$, because $\log R > 1 - 1/R = d/p$; for $b < 1$, $\Gamma''(b) = pRe^{-b}(b-2) < 0$, while for $b \ge 1$ the formula $\Gamma'(b) = pRe^{-b}(1-b) - q$ is strictly negative directly. Thus each $\epsilon > 0$ determines a unique $b_\epsilon \in (a, 2a)$ by

$$b_\epsilon(pRe^{-b_\epsilon} - q) = \frac{D}{1+\epsilon}.$$

For every fixed $a \le b < b_\epsilon$, (7) has strict asymptotic slack relative to $1+\epsilon$, hence supplies a feasible device for all sufficiently small r. Taking the infimum over these fixed choices after taking the r-limsup proves

$$\limsup_{r\downarrow 0} \frac{m_\epsilon(r,\eta)}{L} \le U_\epsilon := \frac{2}{b_\epsilon} < \frac{2}{a}. \tag{8}$$

Feasibility at the equality boundary is not assumed. □

Remark. Implicit differentiation gives

$$b_\epsilon = a + \frac{D}{pa - d}\epsilon + O_\eta(\epsilon^2), \qquad U_\epsilon = \frac{2}{a} - \frac{2d}{a(pa-d)}\epsilon + O_\eta(\epsilon^2).$$

The universal covering inequality implies only the necessary condition

$$\epsilon \ge (M/2)\psi([2/M - a]_+)$$

for a finite positive subsequential coefficient M. Its quadratic consequence is

$$M \ge \frac{2}{a + \epsilon + \sqrt{\epsilon^2 + 2a\epsilon}} = \frac{2}{a} - \frac{2\sqrt{2}}{a^{3/2}}\sqrt{\epsilon} + O_a(\epsilon).$$

The necessary square-root allowance and the attainable linear improvement do not match, so this integrated covering bound alone implies neither optimality of this randomized family nor existence of the fixed-ϵ inner limit. The point certificate of Section 7 removes the mismatch for $\epsilon \ge \epsilon^*_\eta$ (Theorem 3). Constants above are for fixed η, λ and are not uniform as $\lambda \downarrow q/p$.

7 The fixed-tolerance frontier and its second outer endpoint

The bounds in this section concern the general controller class. The exact restricted counter frontier is proved in Section 10. Set $a = \log R$, $p = 1-\eta$, $q = \eta$, $d = p - q$, $D = da$ and $L = \log(1/r)$.

Proposition 12. *For every fixed finite $\epsilon > 0$ there is a constant $C_{\epsilon,\eta}$ such that, for all sufficiently small r, every feasible device satisfies $r + \rho e \le C_{\epsilon,\eta} rL$ and*

$$m \ge 1 + \frac{L - \log L - \log(2C_{\epsilon,\eta})}{a};$$

in particular $\liminf_{r\downarrow 0} m_\epsilon(r,\eta)/L \ge 1/a$.

Proof. For every fixed finite $\epsilon > 0$, feasibility implies $r + \rho e \le C_{\epsilon,\eta} r L$ eventually. The Hellman–Cover reset corollary therefore gives

$$R^{m-1} \ge \frac{1}{C_{\epsilon,\eta} r L} - 1 \ge \frac{1}{2C_{\epsilon,\eta} r L}$$

once $C_{\epsilon,\eta} r L \le 1/2$. Consequently

$$m \ge 1 + \frac{L - \log L - \log(2C_{\epsilon,\eta})}{a}, \qquad \liminf_{r \downarrow 0} \frac{m_\epsilon(r,\eta)}{L} \ge \frac{1}{a}.$$

This is a corollary of the inherited converse, rather than a new mechanism. □

Proposition 13. *Let $\psi(z) = e^z - 1 - z$. For each $\epsilon > 0$ the equation $\psi(z)/(a+z) = \epsilon$ has a unique root $z_\epsilon > 0$, and, with $b_\epsilon \in (a, 2a)$ as in Proposition 11,*

$$\max\left\{\frac{1}{a}, \frac{2}{a + z_\epsilon}\right\} \le \liminf_{r \downarrow 0} \frac{m_\epsilon}{L} \le \limsup_{r \downarrow 0} \frac{m_\epsilon}{L} \le \frac{2}{b_\epsilon}. \tag{9}$$

Proof. The function $z \mapsto \psi(z)/(a+z)$ vanishes at $z = 0$ and tends to infinity. Its derivative has numerator $(e^z - 1)(a+z) - \psi(z)$, which vanishes at zero and has derivative $e^z(a+z) > 0$. Thus it is strictly increasing for $z > 0$. Applying the covering inequality (Lemma 4) to every convergent subsequence of m_ϵ/L yields the sandwich. Boundedness needed for subsequences follows from the counter upper bound. The b_ϵ on the right is the root in $(a, 2a)$ defined in Proposition 11. □

A point certificate for the posterior covers. The lower bound in (9) integrates the covers and charges only the growth of a cover across a gap. Testing the covers at one point per wide gap gives a sharper bound. For $0 \le h < 2a$ put

$$\Upsilon(h) = \frac{D}{pRe^{-h} - q} = \frac{Dh}{\Gamma(h)} \quad (h > 0),$$

with Γ as in Proposition 11; thus $\Upsilon(a) = a$, and $\Upsilon(h)/h = D/\Gamma(h)$ is strictly increasing on $[a, 2a)$ and tends to infinity as $h \uparrow 2a$. With ζ_η and $\varkappa_\eta$ as in Theorem 3, set $h_\eta = a + \zeta_\eta$ and

$$\Upsilon_*(h) = \begin{cases} 0, & 0 \le h \le a, \\ \varkappa_\eta(h - a), & a \le h \le h_\eta, \\ \Upsilon(h), & h_\eta \le h < 2a. \end{cases}$$

Lemma 14. *Υ is strictly convex on $[0, 2a)$. The number $\zeta_\eta \in (0, a)$ is well defined, $\Upsilon(h_\eta) = \varkappa_\eta \zeta_\eta$ and $\Upsilon'(h_\eta) = \varkappa_\eta$. Consequently Υ_* is convex and continuous on $[0, 2a)$, $\Upsilon_*(0) = 0$, $\Upsilon_* \le \Upsilon$ on $[a, 2a)$, and $\Upsilon_*(h)/h$ vanishes on $(0, a]$ and is continuous and strictly increasing on $[a, 2a)$, tending to infinity as $h \uparrow 2a$.*

Proof. Write $y = pRe^{-h}$, so $y > q$ for $h < 2a$ and $dy/dh = -y$. Then $\Upsilon' = Dy/(y-q)^2$ and $\Upsilon'' = Dy(y+q)/(y-q)^3 > 0$. The function $\zeta \mapsto \zeta + e^{\zeta - a} - 1$ is increasing, equals $e^{-a} - 1 < 0$ at $\zeta = 0$ and $e^0 + a - 1 = a > 0$ at $\zeta = a$; hence $\zeta_\eta \in (0, a)$ is unique. At h_η, since $e^{-a} = q/p$, the defining equation gives $e^{-\zeta_\eta} = q/(p(1 - \zeta_\eta))$, so $y = q/(1 - \zeta_\eta)$ and $y - q = q\zeta_\eta/(1 - \zeta_\eta)$. Therefore

$$\Upsilon(h_\eta) = \frac{D(1 - \zeta_\eta)}{q\zeta_\eta} = \varkappa_\eta \zeta_\eta, \qquad \Upsilon'(h_\eta) = \frac{D(1 - \zeta_\eta)}{q\zeta_\eta^2} = \varkappa_\eta.$$

Thus the middle piece of Υ_* is the tangent of Υ at h_η, it passes through $(a, 0)$, and the three pieces agree at a and h_η. Their slopes 0, $\varkappa_\eta$ and $\Upsilon'(h) \ge \varkappa_\eta$ $(h \ge h_\eta)$ are nondecreasing, so Υ_* is convex; it lies below Υ on $[a, h_\eta]$ because a convex function lies above its tangents. Finally $\Upsilon_*(h)/h = \varkappa_\eta(1 - a/h)$ is strictly increasing on $[a, h_\eta]$, and $\Upsilon(h)/h = D/\Gamma(h)$ is strictly increasing on $[h_\eta, 2a)$ and tends to infinity, since Γ is strictly decreasing on $[a, 2a)$ with $\Gamma(2a) = 0$ (Proposition 11). □

Lemma 15 (Point-certificate covering inequality)**.** *Under the hypotheses of Lemma 4, for each of the two covers split $[A, B]$ at A, B and the distinct positions of its centers lying in $[A, B]$. Every resulting gap has length $h < 2a$, in particular $2H/(m+2) < 2a$, and*

$$\frac{2D}{r}(r + \rho e) \ge c \sum_{\text{gaps with } h > a} \Upsilon(h) \ge c\,(m+2)\,\Upsilon_*\left(\frac{2H}{m+2}\right),$$

where the sum runs over the gaps of both covers.

Proof. **Covers.** The proof of Lemma 4 shows that each cover has the form $F(x) = \sum_j w_j k(z_j - x)$ with $w_j \geq 0$, that $F \geq c$ on $[A, B]$, and that $D \sum_j w_j = \int F$, the two areas summing to at most $(2D/r)(r + \rho e)$. Recall $k(z - x) = 0$ when $|x - z| \geq a$, and $k(z - x) = pe^{x-z} - q$ when $z - a \leq x \leq z$. Fix one cover and one gap (y, y'). No center lies in (y, y'): centers in $[A, B]$ are split points, and the other centers lie outside $[A, B] \supseteq (y, y')$.

No gap of length $2a$ or more. If $y' - y \geq 2a$, the point $x = y + a$ lies in $[A, B]$ and every center satisfies $|x - z| \geq a$, so $F(x) = 0 < c$, a contradiction.

One test point per wide gap. Let $h = y' - y \in (a, 2a)$ and $x_h = y + a \in (y, y') \subseteq [A, B]$. A center with $k(z - x_h) > 0$ satisfies $|z - x_h| < a$; as $(x_h - a, x_h] \subset (y, y')$, it lies in the window $[y', y + 2a)$. There $z - a < x_h \leq z$, and since $x_h - z \leq a - h$,

$$k(z - x_h) = pe^{x_h - z} - q \leq pe^{a-h} - q = pRe^{-h} - q = \frac{D}{\Upsilon(h)}.$$

Windows of distinct wide gaps of the same cover are disjoint: if (y_1, y_1') precedes (y_2, y_2'), then $y_2' > y_2 + a \geq y_1' + a > y_1 + 2a$. Hence every center satisfies $\sum_{h>a} \Upsilon(h) k(z_j - x_h) \leq D$, the sum running over the wide gaps of its cover. Multiplying by w_j and summing over j,

$$\int F = D \sum_j w_j \geq \sum_{h>a} \Upsilon(h) F(x_h) \geq c \sum_{h>a} \Upsilon(h).$$

Adding the two covers proves the first inequality.

Convexification. By Lemma 14, $\Upsilon \geq \Upsilon_*$ on the wide gaps and $\Upsilon_* = 0$ on the others, so the sum is at least $\sum_{\text{gaps}} \Upsilon_*(h)$. As in Lemma 4, the two covers have at most $m + 2$ gaps in total, of total length $2H$. If there are g gaps, Jensen's inequality gives at least $g\Upsilon_*(2H/g)$, with $2H/g < 2a$. Since Υ_* is convex with $\Upsilon_*(0) = 0$, $\Upsilon_*(h)/h$ is nondecreasing, so $g\Upsilon_*(2H/g) = 2H\, \Upsilon_*(2H/g)/(2H/g)$ does not increase when g is replaced by $m + 2 \geq g$. □

For each $\epsilon > 0$ let $\hat{b}_\epsilon \in (a, 2a)$ be the unique root of $\Upsilon_*(b) = (1 + \epsilon)b$, which exists by Lemma 14. Since $\Upsilon(h_\eta)/h_\eta = \varkappa_\eta \zeta_\eta/(a + \zeta_\eta) = 1 + \epsilon_\eta^*$,

$$\hat{b}_\epsilon = b_\epsilon \quad (\epsilon \geq \epsilon_\eta^*), \qquad \hat{b}_\epsilon = \frac{a\varkappa_\eta}{\varkappa_\eta - 1 - \epsilon} \in (a, h_\eta) \quad (0 < \epsilon < \epsilon_\eta^*).$$

Indeed, for $\epsilon \geq \epsilon_\eta^*$ the root lies in $[h_\eta, 2a)$, where $\Upsilon_*(b) = (1 + \epsilon)b$ reads $\Gamma(b) = D/(1 + \epsilon)$; otherwise it lies on the linear piece. Also $\epsilon_\eta^* > 0$, because $\Upsilon(h)/h > \Upsilon(a)/a = 1$ for $h > a$. In particular $\varkappa_\eta = (1 + \epsilon_\eta^*)h_\eta/\zeta_\eta > 1 + \epsilon$ in the second case.

Proof of Theorem 3. Fix $\epsilon > 0$ and $0 < \kappa < 1/2$, and use the bands of Proposition 7. As in the proof of Theorem 1, for sufficiently small r an attained optimizer at $m = m_\epsilon$ has error $e \leq (1 + \epsilon)e_B = O_\eta(rL)$, so $c > 0$ eventually and $c \to 1$. Lemma 15 gives

$$\frac{2D(r + \rho e)}{rL} \geq c\, M_r\, \Upsilon_*\!\left(\frac{2(1 - 2\kappa)}{M_r}\right), \qquad M_r = \frac{m_\epsilon + 2}{L},$$

and in particular $2(1 - 2\kappa)/M_r < 2a$. By Proposition 7 the left side has limsup at most $2(1 + \epsilon)$. Let $M = \liminf_{r \downarrow 0} m_\epsilon/L$, finite by Proposition 11, and fix a sequence $r_k \downarrow 0$ along which $M_r \to M$; neither depends on κ. Then $M \geq (1 - 2\kappa)/a > 0$. If $2(1 - 2\kappa)/M = 2a$, the right side tends to infinity along r_k by Lemma 14, which is impossible; hence $2(1 - 2\kappa)/M < 2a$, and continuity of Υ_* gives

$$M\, \Upsilon_*\!\left(\frac{2(1 - 2\kappa)}{M}\right) \leq 2(1 + \epsilon) \qquad \text{for every } \kappa \in (0, 1/2).$$

Letting $\kappa \downarrow 0$ excludes $2/M \geq 2a$ (the left side would be unbounded) and then yields $\Upsilon_*(b)/b \leq 1 + \epsilon$ for $b = 2/M$. By Lemma 14 and the definition of $\hat{b}_\epsilon$, $b \leq \hat{b}_\epsilon$, that is,

$$\liminf_{r \downarrow 0} \frac{m_\epsilon}{L} \geq \frac{2}{\hat{b}_\epsilon}.$$

For $0 < \epsilon < \epsilon_\eta^*$ this is (ii). For $\epsilon \geq \epsilon_\eta^*$ it reads $2/b_\epsilon$, and Proposition 11 supplies the matching limsup, proving (i). □

Remark. Combining with Propositions 12 and 13, for every fixed $\epsilon > 0$

$$\max\left\{\frac{2}{\hat{b}_\epsilon}, \frac{2}{a + z_\epsilon}\right\} \leq \liminf_{r\downarrow 0} \frac{m_\epsilon}{L} \leq \limsup_{r\downarrow 0} \frac{m_\epsilon}{L} \leq \frac{2}{b_\epsilon},$$

and $2/\hat{b}_\epsilon > 1/a$. The certificate gives no credit to gaps of length at most a, and this is the source of the remaining gap for $\epsilon < \epsilon_\eta^*$: if every such gap could be credited with $a + \Upsilon'(a)(h - a)$, the same proof, with Υ_* replaced by the convex function equal to this line on $[0, a]$ and to Υ on $[a, 2a)$, would give $2/b_\epsilon$ for every $\epsilon > 0$. No such certificate is claimed here. For $\eta = 0.3$ and $\epsilon = 1/4$ the proved interval is approximately $[1.77, 1.83]$, compared with the lower value 1.22 of (9).

Corollary 16 (Large-tolerance outer endpoint)**.**

$$\boxed{\lim_{\epsilon\to\infty} \liminf_{r\downarrow 0} \frac{m_\epsilon}{L} = \lim_{\epsilon\to\infty} \limsup_{r\downarrow 0} \frac{m_\epsilon}{L} = \frac{1}{a}.}$$

Proof. Since $b_\epsilon \uparrow 2a$ as $\epsilon \to \infty$, (9) gives the claim. □

The competitive factor is fixed before the rare-switching limit. This does not allow a diverging diagonal $\epsilon(r)$ or imply a bounded competitive ratio at coefficient exactly $1/a$.

Proposition 17. *(i) If a device sequence satisfies* $\limsup m(r)/L < 1/a$*, then* $e/e_B \to \infty$*. (ii) For every budget coefficient* $M > 1/a$ *there is a causal counter with* $m \leq ML$ *eventually whose competitive ratio* e/e_B *has finite limsup.*

Proof. If a device sequence satisfies $\limsup m(r)/L < 1/a$, choose $M < 1/a$ with $m(r) \leq ML$ eventually. The reset bound gives $r + \rho e \geq \xi r^{aM}$ for some $\xi > 0$; subtracting $r = o(r^{aM})$ and dividing by $e_B \sim rL/D$ proves $e/e_B \to \infty$. Conversely, every budget coefficient $M > 1/a$ permits a causal counter with finite limsup competitive ratio: choose $b \in (a, 2a)$ with $2/b < M$ when $M \leq 2/a$, or use the unmodified counter when $M > 2/a$. □

The endpoint $M = 1/a$ is not decided by these statements.

Lemma 18. *For* $\chi = 1/b \in (1/(2a), 1/a]$*, put* $z = pRe^{-1/\chi}$ *and*

$$J(\chi) = \frac{\chi}{pRe^{-1/\chi} - q}.$$

Then

$$J''(\chi) = \frac{z(z+q)}{\chi^3(z-q)^3} > 0.$$

Proof. Direct differentiation. □

Remark (A scalar averaging limitation). Thus any budget calculation averaging costs as χ_i and $J(\chi_i)$ satisfies $\sum_i w_i J(\chi_i) \geq J(\sum_i w_i \chi_i)$. This does not prove that an arbitrary state-dependent counter has those averaged costs; nor does comparison of upper bounds establish optimal actual risk.

8 Sharp passage time for an inhomogeneous observation-driven arm

This is a dynamical theorem for a specified class of controller arms. The stationary bridge for the sign-decoded counter class is proved in Section 10; a reduction of general controller graphs remains open.

Theorem 19 (Passage time of an inhomogeneous arm)**.** *Fix $p > q > 0$, $p+q = 1$, $R = p/q$ and $a = \log R$. For each n, choose outward acceptances $u_0, \dots, u_{n-1}$ and inward acceptances $v_1, \dots, v_n$, all in $(0, 1]$. Under the old hidden sign the device on $0, \dots, n$ moves outward with probability pu_i and inward with probability qv_i, with unused probability holding and endpoints reflecting. Its stationary endpoint ratio is*

$$\frac{\pi_{\mathrm{old}}(n)}{\pi_{\mathrm{old}}(0)} = \prod_{i=1}^{n} \frac{pu_{i-1}}{qv_i} = e^{\Xi_n}.$$

Under the opposite sign, the same kernel moves inward with probability pv_i and outward with probability qu_i. Let T_n be its mean time from n to first hitting 0, with n reflecting. If $\Xi_n/n \to b \in \mathbb{R}$, then, for $b < 2a$,

$$\liminf_{n\to\infty} \frac{T_n}{n} \geq \frac{1}{pRe^{-b} - q}. \tag{10}$$

For $b \geq 2a$, $T_n/n \to \infty$. For every $b \in [a, 2a)$ the bound is attained asymptotically by $u_i = 1$, $v_i = Re^{-b}$. All probabilities may depend on n and approach zero; no uniform positive lower bound is imposed. Time is counted in observation updates.

Proof. Write t_i for the opposite-sign mean hitting time of zero, $t_0 = 0$, and $d_i = t_i - t_{i-1}$. Finite positive inward probabilities ensure finite hitting times. First-step conditioning gives

$$pv_i d_i - qu_i d_{i+1} = 1 \quad (i < n), \qquad pv_n d_n = 1.$$

Back substitution and summing the d_i yield the exact expansion

$$T_n = \sum_{1 \leq i \leq j \leq n} \frac{q^{j-i}}{p^{j-i+1}} \frac{\prod_{\ell=i}^{j-1} u_\ell}{\prod_{\ell=i}^{j} v_\ell}. \tag{11}$$

Put $\beta_i = a + \log u_{i-1} - \log v_i$, so $\Xi_n = \sum_i \beta_i$. An interval of length $k = j - i + 1$ contributes exactly

$$\frac{1}{p} \frac{\exp\{\sum_{\ell=i}^{j} \beta_\ell - a(2k-1)\}}{u_{i-1}} \geq \frac{1}{p} \exp\left\{ \sum_{\ell=i}^{j} \beta_\ell - a(2k-1) \right\}.$$

The length-one summands in (11) give $T_n \geq 1/(pv_i)$ for every i. Hence

$$M_n := \max_i (\beta_i)_+ \leq \max\{0, a + \log(pT_n)\}.$$

Among all length-k intervals, each index occurs at most k times; the total omitted multiplicity is $k(k-1)$. Negative β_i only improve the lower bound, so

$$\sum_{i=1}^{n-k+1} \sum_{\ell=i}^{i+k-1} \beta_\ell \geq k\Xi_n - k(k-1)M_n.$$

Jensen's inequality, separately at each length, therefore gives for $K \leq n$

$$T_n \geq \sum_{k=1}^{K} \frac{n-k+1}{p} \exp\left\{ -a(2k-1) + \frac{k\Xi_n - k(k-1)M_n}{n-k+1} \right\}. \tag{12}$$

If $\liminf T_n/n = \infty$ there is nothing to prove. Otherwise choose a subsequence attaining its finite liminf, with $T_n \leq Cn$. On it $M_n = O(\log n)$. For every fixed K, divide (12) by n and take this subsequence limit:

$$\liminf \frac{T_n}{n} \geq \frac{1}{p} \sum_{k=1}^{K} e^{kb - a(2k-1)}.$$

Only now let $K \to \infty$. For $b \geq 2a$ the sum diverges, ruling out a finite liminf. For $b < 2a$ its value is

$$\frac{e^{b-a}}{p(1 - e^{b-2a})} = \frac{1}{pRe^{-b} - q},$$

proving (10). For sharpness, let $\lambda = Re^{-b} \in (q/p, 1]$ and $g = p\lambda - q > 0$. The same recurrence gives

$$T_n = \frac{n}{g} - \frac{q}{g^2} \left[1 - \left(\frac{q}{p\lambda} \right)^n \right], \qquad \Xi_n = nb.$$

Its bounded correction proves the asserted sharpness. □

The endpoint theorem alone does not identify stationary error. The next section controls geometric interruption, and Section 10 supplies the barrier and switching-law bridge for the prescribed counter class. General controller graphs and jumps remain outside that reduction.

9 Geometric interruption and stationary wrong-arm occupation

The passage theorem also gives a lower bound when another hidden switch may interrupt recovery. We retain the actual conditional switch-time state distribution.

Geometrically discounted passage. Use the opposite-sign arm kernel of the preceding section. Write τ_0 for its first hitting time of zero, and, for $0 < r < 1$, put $\omega = 1 - r$ and

$$\mathcal{D}_i^{(r)} = \mathbb{E}_i \sum_{t=0}^{\tau_0 - 1} \omega^t, \qquad T_i = \mathbb{E}_i \tau_0.$$

Here ω is a geometric survival probability, distinct from the switching-correlation parameter $\rho = 1 - 2r$.

Lemma 20. *With this notation,*

$$\frac{T_n}{1 - r + rT_n} \le \mathcal{D}_n^{(r)} \le T_n. \tag{13}$$

Proof. Indeed, skip-free paths from $j > i$ must hit i before zero. If σ_i is that hitting time, the strong Markov property gives

$$\mathcal{D}_j^{(r)} - \mathcal{D}_i^{(r)} = (1 - \mathbb{E}_j \omega^{\sigma_i})(r^{-1} - \mathcal{D}_i^{(r)}) \ge 0.$$

For the substochastic kernel P on $1, \dots, n$, its fundamental matrix $(I - P)^{-1}$ is entrywise nonnegative. The recursions imply

$$(I - P)T = \mathbf{1}, \qquad (I - P)\mathcal{D}^{(r)} = \omega^{-1}(\mathbf{1} - r\mathcal{D}^{(r)}).$$

Since $\mathcal{D}_i^{(r)} \le \mathcal{D}_n^{(r)}$, multiplication by that fundamental matrix yields $\mathcal{D}_n^{(r)} \ge \omega^{-1}(1 - r\mathcal{D}_n^{(r)})T_n$, which rearranges to (13). The upper bound is pathwise. No second moment, mixing-time estimate or lower bound on acceptance probabilities is used. □

Corollary 21. *If $n \to \infty$, $rn \to 0$ and $\Xi_n/n \to b < 2a$, then*

$$\liminf \frac{\mathcal{D}_n^{(r)}}{n} \ge \frac{1}{pRe^{-b} - q}. \tag{14}$$

For $b \ge 2a$, $\mathcal{D}_n^{(r)}/n \to \infty$.

Proof. Theorem 19 bounds T_n/n from below and Lemma 20 compares $\mathcal{D}_n^{(r)}$ with T_n. To justify both assertions even when rT_n is large, fix any finite lower bound $T_n/n \ge \xi$ eventually. The increasing function $x \mapsto x/(1 - r + rx)$ gives $\mathcal{D}_n^{(r)}/n \ge \xi/(1 - r + rn\xi) \to \xi$. Only afterwards let ξ approach the passage lower bound, or tend to infinity. Thus rare trapping does not invalidate this lower bound through geometric truncation. □

An exact stationary identity. Return to the original symmetric switching model, with update kernels K_s, $s \in \{+, -\}$. For any stationary joint law let μ_s be the conditional law of the post-observation device state given hidden sign s. Fix a stationary MAP decoder, breaking ties deterministically, and let f_s indicate states decoded as $-s$.

Proposition 22. *For every stationary joint law, the actual error of this decoder satisfies the convergent exact identity*

$$e = \frac{r}{2} \sum_{s=\pm} \sum_{k \ge 0} (1 - r)^k \mu_{-s} K_s^{k+1} f_s. \tag{15}$$

Proof. The switch coin is independent of the preceding device state. Therefore

$$\mu_s = (1-r)\mu_s K_s + r\mu_{-s}K_s, \qquad \mu_s = r\mu_{-s}K_s[I-(1-r)K_s]^{-1}.$$

These formulas hold separately for each stationary recurrent mixture; irreducibility of the controller is unnecessary. The identity for e is a last-switch decomposition with the actual conditional initial laws, rather than equilibrium initial laws. □

Proposition 23. *Suppose, for sign s, the kernel K_s has a wrong-decision arm $1,\dots,N_s$, with skip-free transitions, holding and reflection at N_s until hitting the boundary 0. Every state of this arm is decoded as $-s$. Let $w_s = \mu_{-s}(\{n_s,\dots,N_s\})$ for some $n_s \ge 1$, and let $\mathcal{D}^{(r)}_{s,n_s}$ be the discounted time to zero starting at n_s under that kernel. Then*

$$e \ge \frac{r}{2(1-r)} \sum_{s=\pm} w_s\big(\mathcal{D}^{(r)}_{s,n_s} - 1\big). \tag{16}$$

Proof. Monotonicity of hitting times along the arm and (15) reduce the claim to a single start. For a start at $j \ge n_s$, the contribution before its first hit is $\mathbb{E}_j \sum_{t=1}^{\tau_0-1}(1-r)^{t-1} = (\mathcal{D}^{(r)}_{s,j}-1)/(1-r)$. All later errors and all other starts have nonnegative contributions, proving the inequality. □

To combine it directly with (14), take $n_s = N_s$. A reflected truncation at an interior threshold needs a separate discounted comparison; none is assumed here.

The occupation bound retains the actual switch-time weights. Section 10 derives useful weights by stationary tail flow for the specified counter class. No corresponding reduction for arbitrary controller graphs is assumed.

10 Sharp tradeoff within symmetric sign-decoded birth–death counters

This section closes the stationary barrier and initial-law bridge for a specified recursive architecture. It does not reduce arbitrary controllers to that architecture. In particular the prescribed sign decoder need not be the actual MAP decoder of an inhomogeneous device.

Fix $p > q > 0$, $p+q=1$, $a = \log(p/q)$ and $L = \log(1/r)$, with $0 < r < 1/2$. On $-n,\dots,n$, the plus-sign kernel moves from $i > 0$ outward with probability pu_i when $i < n$, and inward with probability qv_i, reflecting at n. On $i < 0$ these probabilities are respectively $qu_{|i|}$ away from zero and $pv_{|i|}$ towards zero. From zero the probabilities of moving to $+1, -1$ are pu_0, qu_0. All remaining probability holds; $u_0,\dots,u_{n-1}, v_1,\dots,v_n \in (0,1]$. The minus kernel is the reflection of the plus kernel. All parameters can depend on r. Decode nonzero states by their sign and break the zero tie towards plus. Let E_r be this decoder's stationary error, after the current update.

Theorem 24 (Sharp tradeoff for sign-decoded counters)**.** *If $n/L \to \chi \in (1/(2a), 1/a]$, then*

$$\liminf_{r\downarrow 0} \frac{E_r}{rL} \ge J(\chi) := \frac{\chi}{p(p/q)e^{-1/\chi} - q}. \tag{17}$$

The constant is attained by the constant-acceptance counters already constructed, with $b = 1/\chi$ and $v_i = (p/q)e^{-b}$, $u_i = 1$. If $n/L \to \chi \le 1/(2a)$, then $E_r/(rL) \to \infty$. Thus this architecture has sharp competitive factor $DJ(\chi)$ for $\chi \in (1/(2a), 1/a]$, where $e_B \sim rL/D$. This is a restricted optimum, not the all-controller fixed-tolerance theorem. Here $r \downarrow 0$ may be taken along any sequence $r_k \downarrow 0$ with $n/L \to \chi$ along it; the proof uses only such a sequence.

Proof. **Stationary order and the barrier.** The finite joint chain is irreducible and aperiodic; all acceptances are positive and endpoints have positive holding. Write μ_i for the conditional device mass given hidden plus. Reflection gives the minus-conditional mass μ_{-i}, and

$$E_r = \sum_{i<0} \mu_i + \mu_0/2, \qquad \nu_i = (1-r)\mu_i + r\mu_{-i}.$$

Both kernels preserve stochastic order on the line: the adjacent up/down probabilities have sum at most $p+q=1$, including the two edges incident to zero. Moreover K_+ stochastically dominates K_- from each common starting state. Starting the two conditional laws equally, their update preserves their stochastic ordering, since their pre-update difference is $(1-2r)$ times the preceding difference. Convergence to the unique stationary joint law therefore proves $\mu_+ \succeq_{\mathrm{st}} \mu_-$. This is tail ordering, not pointwise posterior ordering or sign-MAP alignment.

Across the edge $(j-1, j)$, stationary tail balance yields

$$pu_{j-1}\nu_{j-1} - qv_j\nu_j = r\sum_{k=j}^{n}(\mu_k - \mu_{-k}) \geq 0.$$

Set $\Xi_i = \sum_{j=1}^{i} \log[pu_{j-1}/(qv_j)]$. It follows that $\nu_i \leq \nu_0 e^{\Xi_i}$ and $\nu_0 = \mu_0 \leq 2E_r$. Fix $\delta \in (0,1)$. On a subsequence with $E_r = O(rL)$ and $n = O(L)$, let k be the first positive index whose cumulative μ mass is at least δ. Such k exists eventually. Some $i \leq k$ has $\mu_i \geq \delta/n$, while

$$w := \sum_{j=i}^{n}\mu_j \geq 1 - 2E_r - \delta, \qquad \Xi_i \geq \log\frac{(1-r)\delta}{2nE_r} \geq L - O(\log L).$$

Thus the large barrier occurs before almost all the actual switch-time mass, without assuming endpoint concentration.

Interior truncation and passage. Under the new minus sign, let $\mathcal{D}_j$ be the discounted hitting time of zero in the full positive arm, and $\mathcal{D}_i^{\mathrm{ref}}, T_i^{\mathrm{ref}}$ the discounted and ordinary times when the arm is instead reflected at i. Monotonicity gives $\mathcal{D}_{i+1} \geq \mathcal{D}_i$. Restricting the full discounted equations to $1, \ldots, i$ leaves at its last coordinate the extra nonnegative term $(1-r)qu_i(\mathcal{D}_{i+1} - \mathcal{D}_i)$ on the right side of the reflected resolvent equation. Its inverse is nonnegative, so $\mathcal{D}_i \geq \mathcal{D}_i^{\mathrm{ref}}$. For $i = n$ the quantities coincide. The previous occupation inequality, valid also for a fixed decoder, gives

$$E_r \geq \frac{rw}{1-r}(\mathcal{D}_i - 1), \qquad \mathcal{D}_i^{\mathrm{ref}} \geq \frac{T_i^{\mathrm{ref}}}{1 - r + rT_i^{\mathrm{ref}}}.$$

Hence $\mathcal{D}_i^{\mathrm{ref}} = O(L)$ and, by rearranging the latter inequality, $T_i^{\mathrm{ref}} \leq (1-r)\mathcal{D}_i^{\mathrm{ref}}/(1 - r\mathcal{D}_i^{\mathrm{ref}}) = O(L)$.

Let $\beta_j = \log[pu_{j-1}/(qv_j)]$. The previous interval expansion gives $\max_{j \leq i}(\beta_j)_+ = O(\log L)$ and

$$\frac{T_i^{\mathrm{ref}}}{i} \geq \frac{1}{p}\exp(\Xi_i/i - a).$$

Together with $\Xi_i \geq L - O(\log L)$ this first implies $i \geq CL/\log L$ for some $C > 0$, and then rules out $i/L \to 0$: the lower bound for T_i^{ref}/L would diverge. On any subsequence attaining a finite error liminf, pass further to $i/L \to s \in (0, \chi]$, $\Xi_i/i \to b$; the displayed bound ensures b is bounded, and $b \geq 1/s$. The sharp arm passage theorem rules out $b \geq 2a$ and gives

$$\liminf \frac{E_r}{rL} \geq (1-\delta)\frac{s}{p(p/q)e^{-b} - q} \geq (1-\delta)J(s).$$

Here $rT_i^{\mathrm{ref}} \to 0$ follows from the proved $O(L)$ bound, rather than being assumed. On $(1/(2a), 1/a]$, J is strictly decreasing: the earlier convexity formula gives $J'' > 0$ and $J'(1/a) = (p-q-pa)/(p-q)^2 < 0$. Thus $J(s) \geq J(\chi)$. Let $\delta \downarrow 0$ after the rare-switching limit to obtain (17). If $\chi \leq 1/(2a)$, the forced relations $s \leq \chi$ and $b \geq 1/s \geq 2a$ contradict finite error liminf. The earlier constant-profile upper bound is exactly $J(\chi)$, proving attainment. None of these arguments upgrades sign-decoded error to the smaller MAP error when the two decoders disagree, or treats jumps and general controller graphs. □

Corollary 25 (Restricted fixed-tolerance state frontier)**.** *Let $m_\epsilon^{\mathrm{BD}}(r)$ be the least cardinality among precisely these positive-acceptance, symmetric sign-decoded counters whose error is at most $(1+\epsilon)e_B$. For every fixed $\epsilon > 0$,*

$$\lim_{r \downarrow 0}\frac{m_\epsilon^{\mathrm{BD}}(r)}{L} = \frac{2}{b_\epsilon}, \qquad b_\epsilon(p(p/q)e^{-b_\epsilon} - q) = \frac{D}{1+\epsilon}.$$

Proof. The earlier strict-slack construction gives the limsup; it belongs to this class with $u_i = 1$ and $v_i = \lambda$, and its fair tie at zero has the same stationary error $\sum_{i<0}\mu_i + \mu_0/2$ as the plus tie, by symmetry. For any smaller subsequential cardinality coefficient, the theorem either forces infinite competitive error or a factor $DJ(\chi) > 1 + \epsilon$, since J is strictly decreasing and $\chi < b_\epsilon^{-1}$. This contradicts feasibility and proves the liminf. No inner limit for unrestricted controllers is inferred. □

11 Scope and open questions

The general fixed-tolerance MAP frontier is identified for $\epsilon \geq \epsilon^*_\eta$ (Theorem 3) and remains open for $0 < \epsilon < \epsilon^*_\eta$. The sign-decoded counter theorem concerns the prescribed sign decoder and nearest-neighbor architecture. None of the theorems is uniform as the observation channel becomes uninformative. State cardinality is not memory entropy or thermodynamic work. Classical Bayes and static finite-memory results are credited above.

Remark 26 (Numerical evidence on the open frontier; not used in any proof). The following study was carried out before Theorem 3 was proved, when only the bounds of Proposition 13 were available; for $\epsilon \geq \epsilon^*_\eta$ its conclusion is now a theorem. It suggests that the upper value $2/b_\epsilon$ is the sharp one also for $\epsilon < \epsilon^*_\eta$. First, for $\eta \in \{0.1, 0.2, 0.3\}$, $r \in \{10^{-2}, 10^{-3}, 10^{-4}\}$ and $2 \leq m \leq 8$, multistart local optimization over all randomized m-state kernels with MAP decoding found no controller better than the best nearest-neighbour chain with MAP decoding, except in two near-Bayes cases ($\eta = 0.1$, $r = 10^{-2}$, $m = 7, 8$, excess below 0.3%), where a deterministic automaton with two-step jumps was better by less than 0.1% of e_B. Second, for $\eta = 0.3$ and $\epsilon \in \{1/4, 1\}$, the discrete slopes $2/(L_\epsilon(m+2) - L_\epsilon(m))$, where $L_\epsilon(m)$ is the value of $\log(1/r)$ at which the best m-state chain has error exactly $(1+\epsilon)e_B$, stabilize near $2/b_\epsilon$ (about 1.83 and 1.49) for $7 \leq m \leq 13$, as do those of the symmetric sign-decoded counters of Section 10, whose limit is proved in Corollary 25. Both are well above the lower value $\max\{1/a, 2/(a+z_\epsilon)\}$ (about 1.22 and 1.18). Here $\epsilon^*_{0.3} \approx 0.53$, so the value $2/b_1 \approx 1.49$ is now proved for all controllers, while for $\epsilon = 1/4$ Theorem 3 gives the lower value $2/\hat{b}_{1/4} \approx 1.77$. These computations are local optimizations at moderate m and r; they neither certify global optima nor reach the asymptotic regime, and they do not settle the open question.

AI assistance disclosure

OpenAI GPT-5.6 Sol, accessed through ChatGPT and Codex, was used to develop and audit the proofs and for drafting. Anthropic Claude Opus 5.5, used through Claude Code, was used for editing, an independent line-by-line check of the proofs, and the numerical study reported in Remark 26. Anthropic Claude Opus 5.5 (Claude Code) produced the point-certificate argument, namely Theorem 3, Lemmas 14 and 15, and the remark following them. The author reviewed and takes responsibility for every statement and proof. AI systems are not authors of this work, and the work has not been peer reviewed.

References

[1] M. E. Hellman and T. M. Cover, Learning with finite memory, *Ann. Math. Statist.* **41**, 765–782 (1970), doi:10.1214/aoms/1177696958. Author-hosted scan. The equal-prior error formula used here is Theorem 3, equation (33).

[2] T. M. Cover, Hypothesis testing with finite statistics, *Ann. Math. Statist.* **40**(3), 828–835 (1969), doi:10.1214/aoms/1177697590.

[3] M. E. Hellman and T. M. Cover, On memory saved by randomization, *Ann. Math. Statist.* **42**(3), 1075–1078 (1971), doi:10.1214/aoms/1177693334.

[4] M. E. Hellman, The effects of randomization on finite-memory decision schemes, *IEEE Trans. Inform. Theory* **18**(4), 499–502 (1972), doi:10.1109/TIT.1972.1054846.

[5] T. Berg, O. Ordentlich and O. Shayevitz, Statistical inference with limited memory: a survey, arXiv:2312.15225.

[6] R. Khasminskii and O. Zeitouni, Asymptotic filtering for finite state Markov chains, *Stochastic Process. Appl.* **63**, 1–10 (1996). Author version. The author-hosted version may differ from the final publication.

[7] C. Nair, E. Ordentlich and T. Weissman, Asymptotic filtering and entropy rate of a hidden Markov process in the rare transitions regime, *IEEE International Symposium on Information Theory* (2005). Author manuscript.

[8] D. Monte and M. Said, The value of (bounded) memory in a changing world, *Economic Theory* (2014), doi:10.1007/s00199-013-0771-1. The 2013 author manuscript is available at MPRA.

[9] K. Chatterjee, K. Guryev and T.-W. Hu, Bounded memory in a changing world: Biases in behaviour and belief, *J. Econ. Theory* **206**, 105556 (2022), doi:10.1016/j.jet.2022.105556.

[10] A. Wilson, Bounded memory and biases in information processing, *Econometrica* **82**(6), 2257–2294 (2014), doi:10.3982/ECTA12188.

A Classical bounds and model translations

A.1 The inherited converse follows from Hellman–Cover

Recall $\rho = 1 - 2r$, the switching-correlation parameter. The classical theorem of Hellman and Cover [1] allows homogeneous randomized automata. Its Theorem 2 bounds likelihood-ratio spread; Theorem 3, equation (33), gives the equal-prior error bound

$$P_{\text{error}} \geq (1 + \gamma^{(m-1)/2})^{-1}.$$

For iid BSC signals under a fixed binary hypothesis, $\gamma = R^2$. The following reduction applies that established bound without replacing a changing signal by a static one without justification.

Proposition 27. *Every stationary device with at most m states and stationary error e (including every stationary mixture) satisfies $r + \rho e \geq (1 + R^{m-1})^{-1}$.*

Proof. For any chosen physical stationary device law, let

$$\pi_s(y) = \Pr(Y_t = y \mid S_t = s), \qquad f_s(y) = \Pr(Y_{t-1} = y \mid S_t = s), \qquad \mu = (\pi_0 + \pi_1)/2,$$

and let Q_s be its observation-averaged memory transition matrix conditional on the new hidden state s. The stationary hidden chain is symmetric, so

$$\pi_s = f_s Q_s, \qquad f_s = (1 - r)\pi_s + r\pi_{1-s} = 2r\mu + \rho f_s Q_s.$$

In a static testing experiment with iid BSC signals, use the homogeneous randomized update

$$K'(z \mid y, o) = 2r\mu(z) + \rho K(z \mid y, o).$$

The vector μ is a fixed kernel parameter, not a hypothesis-dependent initialization or an online side channel. Its conditional transition matrix resets to μ with probability $2r$, and hence contracts total variation by at most $\rho < 1$. It has the unique stationary law f_s, reached from any fixed initial state independently of the hypothesis. Hellman–Cover therefore implies

$$\frac{1}{2}\sum_y \min\{f_0(y), f_1(y)\} \geq \frac{1}{1 + R^{m-1}}.$$

Because $f_0 - f_1 = \rho(\pi_0 - \pi_1)$, the error on the left is exactly $r + \rho e$. This proves the inherited switching bound for each stationary mixture, including reducible original kernels and zero transition probabilities. This is a corollary of Hellman–Cover. The posterior-covering penalty additionally charges gaps in likelihood confidence. □

A.2 A fixed-parameter geometric-horizon equivalence

We record a fixed-parameter equivalence for the same memory kernels, independent private randomization, and terminal decisions based only on the current memory state.

Fix r, η, m. Initialize the hidden state fairly and the memory independently, and take an independent geometric decision time N with $\Pr(N = n) = h(1 - h)^n$ for $n \geq 0$. For a kernel K, write P_K for the physical joint transition matrix and

$$f_h = h\sum_{n \geq 0}(1 - h)^n \nu_0 P_K^n, \qquad f_h = h\nu_0 + (1 - h) f_h P_K.$$

Let $E_m(h)$ minimize the terminal MAP error $\Phi(f_h) = \sum_y \min\{f_h(0,y), f_h(1,y)\}$ over the same kernels and allowed independent initial memory laws. The decision rule cannot use N as an additional uncounted state.

Proposition 28. *For fixed r, η, m, $\lim_{h\downarrow 0} E_m(h) = e_m(r,\eta)$.*

Proof. For the lower inequality, take approximate minimizing kernels and occupations along a sequence $h \downarrow 0$. Compactness gives a convergent subsequence $(K, f_h) \to (K_*, f_*)$. The occupation identity implies $f_* P_{K_*} = f_*$, and continuity of Φ gives a limiting error at least e_m.

For the reverse inequality, choose an attained stationary optimum. Concavity of Φ permits a recurrent component with error no larger than the stationary mixture. Full support of the switching and observation probabilities makes every recurrent joint class a product $\{0,1\} \times C$ with a recurrent memory class C. Initialize a fixed memory state in C, independently of the fair hidden state. The finite irreducible joint class has a unique stationary law, and its Abel occupation limit is that law even when the class is periodic. This supplies terminal errors converging to the chosen optimal stationary error.

□

Remark. The proof fixes r, η, m before sending $h \downarrow 0$. It gives no convergence rate uniform in controllers with tiny transition probabilities or in the regime $r \downarrow 0$, $m \asymp \log(1/r)$. Consequently a uniform growing-cardinality implication would require a separate argument.