WHICH NEURAL RESPONSES CAN INTERVENTION DATA VALIDATE?

Anonymous authors
Paper under double-blind review

ABSTRACT

Which variation should a neural response model explain to make a finite intervention experiment more informative? We study state-conditioned mean stimulation effects under randomized assignment and bounded responses, with finite-sample error control during sequential monitoring. A variance decomposition separates predictable common activity from effect heterogeneity and residual response noise. Matching null laws establish an inverse-margin validation lower bound and an exact information frontier with full, partial, and zero-adjustment regimes. In an eight-state model with asymmetric noise and heterogeneous effects, learned adjustment with accurate history at margin 0.01 reduces restricted correct-decision cost from 7,262 trials for a no-history oracle to 1,818, a 75.0% saving. All 256 learned runs answer correctly by 16,384 trials, compared with 244 for the oracle. With 70% history replacement, the advantage reverses: learned cost is 9,306 trials versus 7,778 without history. Across 256 population-history questions, improving observation increases optimal information by a median factor of 1.53; common-mean adjustment with the finer observation exceeds the coarse observation's information optimum in 253 questions. In a matched beta–binomial condition, a parametric likelihood-ratio test costs 60 trials, compared with 124 for the bounded-response oracle. These results identify when auxiliary prediction changes supported response conclusions and distinguish gains from fitting, observing state, and restricting the response class.

1 INTRODUCTION

An accurate neural response predictor can remove much of the variation in an intervention experiment. How much does that shorten the experiment? The question matters as population models generalize across animals and stimulus domains (Wang et al., 2025), and optical perturbations reach specific circuit elements (Packer et al., 2015; Daie et al., 2021). For a receiver population in a given behavioral state, the experimental question is often whether stimulation changes activity by more than a meaningful threshold. A predictor earns its experimental value by helping resolve that question with fewer interventions.

Two obstacles separate prediction accuracy from a supported decision. First, recorded histories reveal only part of the background activity. Population fluctuations depend on internal state and behavior (Goris et al., 2014; Stringer et al., 2019; Musall et al., 2019); state also modifies responses to perturbation (Mateo et al., 2011; Russell et al., 2024). Better fitting can recover more of the information already recorded. Better observation can reveal information that is currently missing. Second, even a perfectly adjusted response requires validation. Uniform finite-sample error control must allow rare outcomes that change the population effect, including when the responses observed so far are almost deterministic.

We connect these obstacles through the rate at which interventions accumulate evidence. Our first result gives a universal lower bound proportional to $1/(p_G\Delta)$ near the threshold, where p_G is the frequency of the state group and Δ the effect margin. A competing null law preserves history and randomization while introducing rare opposing responses. This construction explains why removing background variance cannot remove all validation cost.

Our central result solves the optimal evidence rate in a noisy history model. It reveals three regimes. Full background adjustment is optimal when unpredictable response noise dominates the margin.

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.

As noise falls, partial adjustment permits a larger valid evidence increment, despite retaining more variance. For sufficiently large margins, no adjustment is needed. Matching null laws establish sharpness for every valid procedure on these observations. Observing history can change the margin dependence from quadratic to linear; optimizing how it is used provides a further, generally smaller improvement.

We then connect this exact solution to richer experiments. A concave finite-history program computes the same information objective under skewed noise, unequal state frequencies, and heterogeneous effects. Predictable learning converts useful history into substantial trial savings, with an initial fitting cost. The information oracle also makes observation comparisons precise: in the population-history study, a common-mean predictor with a better observation often beats the best possible evidence factor with a poorer one. Boundary calibration and parametric comparisons explain which assumptions pay for faster validation. These results organize experimental investment around the question to be answered, the history available, and the response laws the inference procedure must accommodate.

2 Related work

Prediction and effect precision. Covariate adjustment improves randomized comparisons by explaining outcome variation (Lin, 2013; Wager et al., 2016). Held-out and predictable predictions permit flexible models while preserving unbiased effect scores (Wu & Gagnon-Bartsch, 2018; Hadad et al., 2021; Cook et al., 2024). Prediction-powered inference similarly combines predictions with labeled observations (Angelopoulos et al., 2023). These results motivate common-activity adjustment. We follow its reduction in variance through to the information required for a sequential threshold decision.

Prediction-adjusted sequential evidence. Off-policy confidence sequences use affine betting factors, prediction control variates, and empirical log-wealth optimization (Karampatziakis et al., 2021). Predictable truncation and selection by previous confidence-sequence width provide related adjustment rules (Waudby-Smith et al., 2024b). Expected log growth also guides prediction, sampling, and betting in risk control (Xu et al., 2024). We use this established machinery (Kelly, 1956; Waudby-Smith & Ramdas, 2024), and compare it with empirical-Bernstein (Howard et al., 2021) and asymptotic confidence sequences (Waudby-Smith et al., 2024a). Regression-adjusted anytime inference also extends to linear-model assumptions (Lindon et al., 2026).

Information limits and optimal evidence. Sequential change of measure relates decision probabilities to collected information (Kaufmann et al., 2016). Bounded-mean information projections admit a dual representation as optimal log betting growth (Honda & Takemura, 2015; Shekhar & Agrawal, 2026). Affine and complete-class representations characterize optimal evidence under linear and conditional constraints (Clerico, 2025a;b; Larsson et al., 2026; Clerico, 2026). Our finite-history program specializes this theory to randomized interventions. The contribution is its explicit noisy-history solution, the joint noise–margin limit, and the objective-matched comparison of improving observations and improving adjustment. The general duality provides the bridge between the closed-form model and richer response distributions.

3 Prediction, response definition, and validation

A trial supplies a pre-intervention history H, an independently randomized assignment $A \sim \mathrm{Bernoulli}(1/2)$, and a response $U = U(A) \in [0, 1]$. Trials are independent and identically distributed; activity within a trial can be temporally dependent. A question specifies a receiver, response window, history group G, and threshold δ. For $\delta \in (-1, 1)$, its target and margin are

$$T = \mathbb{E}[U(1) - U(0) \mid H \in G], \qquad \Delta = |T - \delta|.$$

The group occurs with probability $p_G > 0$, so acquiring one group trial costs $1/p_G$ calendar trials on average. We develop the above-threshold direction $T > \delta$; exchanging assignment labels gives the other one.

The response bound defines the scientific question. For a neural population, U can be the fraction of prespecified neuron–time bins containing an event, or an indicator of a response in a fixed window. Dependence among those bins changes trial variance; it does not change the bound or create additional interventions. Raw spike counts and $\Delta F/F$ require a different measurement choice. For example, a predeclared count cap C defines $U_C = \min(N, C)/C$ and an effect on capped counts. Its threshold must be expressed on that same scale. Changing C generally changes the estimand. Using a larger admissible response range also admits more competing null laws and can increase validation cost. A change of units that transforms both the support and threshold leaves the statistical problem unchanged. These distinctions matter when population averaging reduces response noise: a low-noise regime is a property of the chosen measurement, not a claim that individual cortical neurons have little trial variation (Goris et al., 2014; Stringer et al., 2019).

3.1 What prediction removes, and what still requires evidence

Within the group, write $c(H) = \{\mu_1(H) + \mu_0(H)\}/2$ for common activity and $\tau(H) = \mu_1(H) - \mu_0(H)$ for the conditional contrast. A prediction function g fitted on completed trials gives the unbiased score

$$\Gamma^g = 2(2A-1)\{U - g(H)\}, \qquad \mathbb{E}\Gamma^g = T, \qquad \mathrm{Var}(\Gamma^g) = V_0 + 4\mathbb{E}(c-g)^2, \tag{1}$$

where $V_0 = 2\mathbb{E}\{\mathrm{Var}(U(1) \mid H) + \mathrm{Var}(U(0) \mid H)\} + \mathrm{Var}\{\tau(H)\}$. Thus common-mean prediction minimizes score variance (Wager et al., 2016; Wu & Gagnon-Bartsch, 2018). Appendix A separates unobserved background variation from fitting error.

Sequential evidence adds a second consideration. A nonnegative factor f with $\mathbb{E}_Q f \le 1$ under every below-threshold law can be multiplied across trials. Reaching wealth $1/a$ controls the probability of a false directional answer by a, including under continuous monitoring. The expected log factor $I = \mathbb{E}_P \log f$ measures the rate of evidence accumulation; $\log(1/a)/I$ is its leading time scale. We require validity throughout the bounded randomized model. A supplied prediction remains the same function under a competing null, without external certification of its accuracy.

Proposition 1 (A validation floor) *Let $s = \mathrm{sign}(T - \delta)$ and suppose the false-answer probability is at most a under every null in direction s. Correct-answer probability at least $1 - \beta > a$ by budget B requires*

$$Bp_G \log \frac{1 + sT}{1 + s\delta} \ge \mathrm{kl}(1 - \beta, a). \tag{2}$$

If the correct-answer probability at a stopping time N is at least $1 - \beta$ and $\mathbb{E}_P N < \infty$, the same bound holds with B replaced by $\mathbb{E}_P N$.

A boundary null proves the result by occasionally replacing a trial's potential responses with maximally opposing endpoints, while preserving its history and randomization. Before such an outcome appears, observed scores can be perfectly stable. Sequential change of measure (Kaufmann et al., 2016) then gives an inverse-margin cost as $\Delta \to 0$. The bound uses the error allocated to the particular direction; a Bonferroni procedure for F questions sets $a = \alpha/(2F)$.

3.2 A computable frontier for general finite histories

Let $W = 2A - 1$ and let H take K values. Consider

$$f_{\lambda,b}(H, A, U) = 1 + \lambda(2WU - \delta) + b_H W, \quad \lambda \ge 0, \quad -1 + \lambda\delta \le b_h \le 1 - \lambda(2 + \delta). \tag{3}$$

The shared coefficient λ multiplies the effect; the randomization terms $b_H W$ have zero expectation because assignment is fair. The inequalities ensure nonnegativity even at the opposing response endpoints. For $\lambda > 0$, $g(h) = -b_h/(2\lambda)$ recovers a prediction-adjusted factor. We can therefore optimize prediction and its valid range using the evidence objective itself.

Proposition 2 (Finite-history information oracle) *For any finite-support randomized law P with $T_P > \delta$, let $\mathcal{Q}_\delta$ be all bounded randomized laws with effect at most δ, allowing the history marginal to vary. Then*

$$I_H(P) = \max_{(\lambda,b)\ \textit{feasible in}\ (3)} \mathbb{E}_P \log f_{\lambda,b} = \min_{Q \in \mathcal{Q}_\delta} D(P\|Q). \tag{4}$$

This is the largest expected log evidence over all null-valid factors. A matching null has $T_Q = \delta$ and $f^* = dP/dQ^*$, yielding Equation (2) *with its logarithm replaced by $I_H(P)$.*

The program is concave with $K+1$ parameters. It permits unequal history frequencies, asymmetric noise, and heterogeneous effects. Replacing a history by a less informative observation cannot increase I_H, by data processing. This compares observation channels using the same objective on both sides. The representation specializes established constrained-null duality (Clerico, 2025b; Larsson et al., 2026; Shekhar & Agrawal, 2026). Its role here is to connect the explicit solution below to richer response models. Appendix A.2 gives the matching null construction and the extension to known unequal assignment probabilities. The contrast coefficient remains shared across histories; letting it vary freely would change the group-average target.

3.3 A CLOSED-FORM SOLUTION REVEALS THREE ADJUSTMENT REGIMES

Take independent fair signs H, R and the bounded family

$$U(a) = \frac{1 + hH + \sigma R + (2a-1)T}{2}, \qquad h, \sigma \geq 0, \quad h + \sigma + |T| < 1. \tag{5}$$

History reveals the background term hH; the residual R is unobserved before assignment. Its contribution to adjusted-score variance is σ^2. Write $d = 1 + \delta$, $\Delta = T - \delta > 0$, and $L = d + h$. For $g_r(H) = (1 + rH)/2$, expected log growth is

$$J(\lambda, r) = \frac{1}{4} \sum_{x,y \in \{-1,1\}} \log\{1 + \lambda[\Delta + (h-r)x + \sigma y]\}, \quad 0 \leq r \leq h, \quad 0 \leq \lambda \leq \frac{1}{d+r}. \tag{6}$$

Proposition 3 (Sharp noisy frontier) *The information rate in this family is $I_H = J(\lambda^*, r^*)$, where*

$$(\lambda^*, r^*) = \begin{cases} (\Delta/(\sigma^2 - \Delta^2), h), & \sigma^2 \geq \Delta(L + \Delta), \\ (1/d, 0), & \Delta \geq h, \quad \sigma^2 \leq (\Delta - h)(d + \Delta - h), \\ (\lambda_p, \lambda_p^{-1} - d), & \textit{otherwise}, \end{cases} \tag{7}$$

with $\lambda_p = 4/\{3(L-\Delta) + \sqrt{(L-\Delta)^2 + 8\sigma^2}\}$. The cases meet continuously. For the oracle's directional crossing time,

$$\frac{\ell}{I_H} \leq \mathbb{E}\tau_\ell \leq \frac{\ell + u^*}{I_H}, \qquad \ell = \log(1/a), \tag{8}$$

where u^* is the largest log increment and time is measured in group trials. The matching null also preserves the history marginal.

The regimes are full, zero, and partial adjustment. Full adjustment removes all predictable background variation. Partial adjustment retains some of it to permit a larger stake; at large margins, no adjustment is needed. For example, at $(h, \sigma, \delta, \Delta) = (0.30, 0, 0.05, 0.10)$, full adjustment gives a constant factor $1 + 0.10/1.35$ and log growth 0.07146. Partial adjustment uses $r = 0.20$ and stake 0.8, producing factors 1 and 1.16 with mean log growth 0.07421, despite score variance 0.01. This is the range–variance tradeoff in a concrete form.

For $h > 0$, full adjustment ends at $\Delta_c = \{\sqrt{L^2 + 4\sigma^2} - L\}/2$. At fixed positive noise, $I_H \sim \Delta^2/(2\sigma^2)$ near the threshold. At zero noise, $I_H \sim \Delta/L$. *Common-mean adjustment has the same leading zero-noise rate:* $I_{\text{common}} = \log(1+\Delta/L)$, whereas the optimum is $\frac{1}{2}\log\{(L+\Delta)/(L-\Delta)\}$ for $\Delta < h$. The change of margin exponent comes from observing history; partial adjustment improves its conversion into evidence at higher order. The joint limit makes the crossover precise:

$$\frac{I_H}{\Delta} \longrightarrow \begin{cases} L^{-1} - c/(2L^2), & c \leq L, \\ (2c)^{-1}, & c \geq L, \end{cases} \quad \text{when } \Delta \downarrow 0, \quad \sigma^2/\Delta \to c. \tag{9}$$

Appendix A gives the hidden-history rate and proofs.

How much can optimizing adjustment add beyond the common mean? At the experimental threshold $\delta = 0.05$, a numerical search over the admissible binary family finds a largest rate ratio of approximately 1.209, approached near zero noise and the response-support boundary. This numerical envelope puts the refinement on a smaller scale than the benefit of observing history. Its size depends on the threshold: over unrestricted $\delta \in (-1, 1)$ the ratio is unbounded as the threshold approaches -1 (Appendix A).

3.4 Learning costs observations as well as computation

In the binary family, least squares of $2U - 1$ on $(H, 2A - 1)$ estimates (h, T); residual variance estimates σ^2. Fit on completed trials, project onto the physical parameter set, and use Equation (7). A stake cap that approaches its valid bound keeps early wealth positive. Consistency implies $n^{-1} \log K_n \to I_H$ and $\tau_\ell/\ell \to 1/I_H$ almost surely. This learning result concerns increasing evidence thresholds; the finite oracle expectation is the separate guarantee in Equation (8).

For finite histories, we fit the history and conditional response distributions from completed observations and solve Equation (4). Feasibility preserves null validity even when these estimates are inaccurate. Two signed wealth processes use threshold $2/\alpha$ per question, or $2F/\alpha$ for F questions. The first direction to cross is the answer; simultaneous crossings remain unresolved. Every fitting observation counts toward acquisition cost. Appendix B specifies the estimators, updates, regularization, and stopping endpoints.

Finite-history learner for one threshold question

1. Initialize two directional wealths at one. Collect 16 group observations before fitting; these trials count toward the budget.
2. At each doubling prefix, estimate history frequencies and the response distribution in each history–assignment cell. Endpoint pseudo-counts stabilize cells with few observations.
3. For each direction, solve the concave information program using the fitted law. Multiply its coefficients by $1 - (n + 1)^{-1/2}$ to keep every factor strictly positive, where n is the completed group count.
4. Hold these coefficients fixed over the next block. Evaluate the factor on each newly observed history, assignment, and response, then multiply the corresponding wealth.
5. Stop at the first directional wealth reaching $2/\alpha$. For F questions use $2F/\alpha$; conflicting crossings remain unresolved.

The learner estimates a response distribution to choose evidence coefficients, but its validity comes from randomization and the response bound. This distinction permits flexible fitting without treating its predictions as certified. An independently trained predictor can shorten the initial learning phase, while the information oracle measures what the supplied history could support with unlimited fitting data. In an experimental planning calculation, its rate per calendar trial is $p_G I_H$: a rarer state directly increases the acquisition time for that question. A comparison of recording strategies should therefore hold the response definition, state group, and error allocation fixed.

4 From information rates to intervention budgets

We test three consequences of the theory: the noise–margin crossover, the value of informative history beyond the binary model, and the cost of accommodating a broad response class. Within each condition, methods use paired trial streams and repetitions are independent. Reported $\pm$ values are Monte Carlo standard errors (MCSEs). For single-question sequential comparisons, both signed processes cross at wealth 40, giving total error level 0.05. The primary endpoint is restricted correct-decision cost: wrong, conflicting, and unresolved first decisions retain the full horizon. Every fitting observation counts.

4.1 The frontier predicts scaling and the finite cost of learning

We cross six margins $\Delta \in \{0.01, 0.02, 0.04, 0.08, 0.16, 0.32\}$ and four residual amplitudes $\sigma \in \{0, 0.05, 0.15, 0.30\}$ in Equation (5), with $h = 0.30$ and $\delta = 0.05$. The 24 cells cover all three adjustment regimes. Each contains 1,024 repetitions to a horizon of 32,768 trials. The no-history, common-mean, and frontier oracles know the population parameters; the OLS learner starts fitting after 16 observations. Appendix B.1 specifies the updates and gives every cell's cost, MCSE, and attainment.

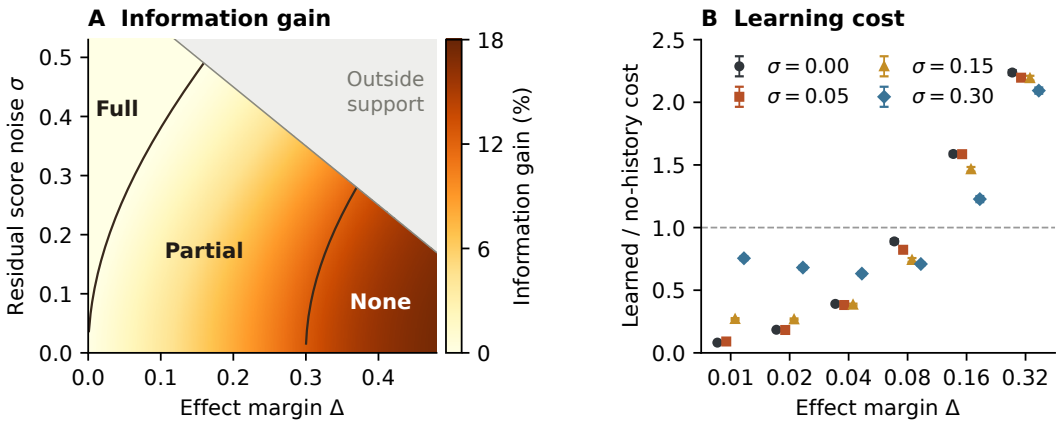

Figure 1: **Optimal adjustment and the value of learning depend on the margin.** A: exact percentage information gain over common-mean adjustment; boundaries separate full, partial, and zero adjustment. Gray lies outside response support. B: learned/no-history restricted-cost ratios, with paired MCSE bars across 1,024 streams per cell. Ratios below one favor learning. The points retain wrong and unresolved decisions at the horizon.

At zero residual noise, reducing the margin from 0.04 to 0.01 multiplies frontier cost by 3.97, from 125.5 ± 0.3 to 497.8 ± 0.7 trials. Without history, cost increases 16.60-fold, from 411 ± 10 to 6826 ± 170. At $\sigma = 0.30$, frontier cost instead increases 16.15-fold, from 402 ± 9 to 6497 ± 149. These patterns agree with the linear and quadratic margin scales. The transition from full to partial adjustment is at $\Delta_c = 0.0637$ for $\sigma = 0.30$. Common-mean and frontier factors coincide throughout the full-adjustment regime. Across this grid, optimizing adjustment adds at most 13.5% to their population information rate (Table 3).

The learner retains most of the practical history benefit near the threshold. At zero noise and margin 0.01, it costs 557.8 ± 0.7 trials, a 91.8% reduction relative to the no-history oracle, and answers correctly in every repetition. At margin 0.32, the no-adjustment oracle needs only 16.1 ± 0.1 trials, while the learner needs 36.0 ± 0.1: its initialization already consumes the oracle's typical budget. In the hardest cell, $(\Delta, \sigma) = (0.01, 0.30)$, learned cost is 9983 ± 217, with 1012 correct answers, two wrong answers, and ten unresolved paths out of 1024. The corresponding oracle has three unresolved paths and no wrong answers. This finite-budget gap is the cost of estimating a rate that is itself small.

4.2 Beyond binary histories: a common information objective

The next study combines eight unequally frequent states, asymmetric three-point noise, state- and arm-dependent response variances, and a heterogeneous treatment contrast. Arm means are $c_z + (a - 1/2)T_z$, with c_z spanning 0.22–0.78 and $T_z = 0.05 + \Delta + 0.04 b_z$, where b_z is centered over states. The observed history is the state label, independently replaced by a fresh population draw with probability ϵ. This channel preserves history frequencies while removing information. We cross four margins, three noise amplitudes, and three replacement probabilities, giving 36 exact information calculations. All responses remain in $[0.1014, 0.9706]$. Appendix C specifies the distributions and numerical oracle, including its concavity-based optimality certificate.

The information-optimal factor improves on common-mean adjustment by up to 13.2%, confirming that the range–variance tradeoff survives asymmetric, heterogeneous responses. We then compare five procedures on six conditions at noise amplitude 0.04: no-history, common-mean, and information oracles, plus empirical common-mean and empirical log-growth learners. Both learners fit the same conditional response distributions, with endpoint pseudo-counts and updates at doubling prefixes. They differ in whether the adjustment is fixed to the estimated common mean or jointly optimized with the stake. Each cell has 256 repetitions and horizon 16,384.

With an accurate history and margin 0.01, the log-growth learner needs 1818 ± 14 trials versus 7262 ± 255 without history, a 75.0% saving. All 256 learned paths answer correctly; the no-history

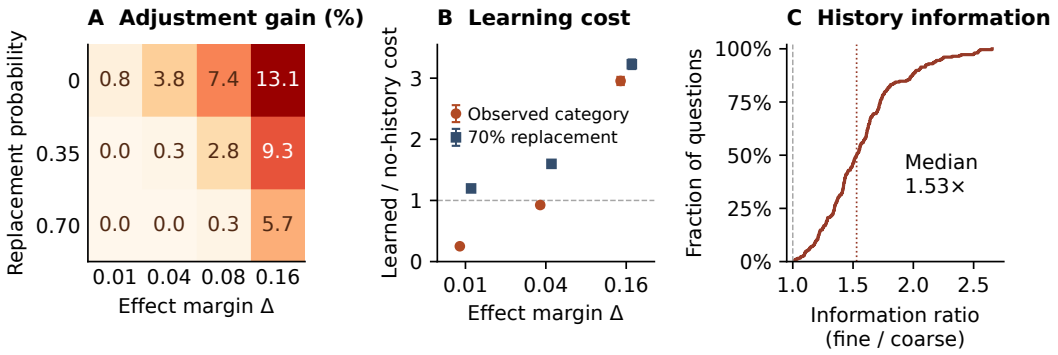

Figure 2: **Informative observations give larger gains than refining adjustment.** A: percentage information gain from optimizing adjustment in the eight-state model, at noise amplitude 0.04. B: learned/no-history restricted-cost ratios with paired MCSEs, 256 repetitions per condition. C: cumulative distribution of the exact fine/coarse information ratio across 256 population-history questions. Panel C describes variation across fixed questions; panels A and C involve population calculations, not Monte Carlo error.

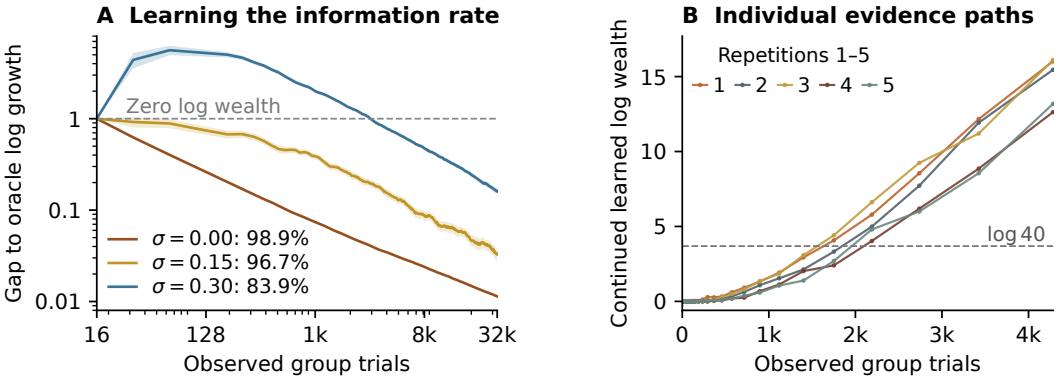

Figure 3: **Initial estimation losses persist beyond the first decision.** A: growth deficit $1 - \mathbb{E}[\log K_n]/(nI_H)$ in the binary model at $\Delta = 0.01$, shown on logarithmic axes with one-MCSE bands. The dashed line corresponds to zero accumulated log wealth; legend percentages give the final attained growth. B: continued learned log wealth in the first five eight-state repetitions at $\Delta = 0.01$ with accurate history. Points are saved checkpoints, connected for readability; the dashed line is the decision threshold. Continued wealth does not revise first decisions.

oracle answers 244. The information oracle needs 634 ± 8 trials, exposing a substantial remaining estimation cost. When replacement rises to 0.70, learned cost is 9306 ± 322, exceeding the no-history cost 7778 ± 291. At larger margins, short oracle experiments similarly leave little time to learn. The two learned objectives have nearly identical costs: in the favorable cell their paired difference is -0.6 ± 1.2 trials (common minus log). The positive result comes from recovering useful history, with only a small additional opportunity in the choice of adjustment objective.

Learning the rate takes time. Continued wealth makes the initial fitting cost visible (Figure 3). In the binary model at margin 0.01, mean log wealth per trial reaches 98.9%, 96.7%, and 83.9% of the oracle rate at 32,768 trials for noise amplitudes 0, 0.15, and 0.30. The high-noise learner initially loses wealth: a noisy estimate can choose a feasible but unprofitable adjustment or stake. Feasibility protects error control, while accurate estimation determines growth. These early losses take many observations to recover when the true information rate is small. In the eight-state favorable condition, mean normalized growth reaches 84.8% at 16,384 trials, despite the 75% stopping-cost saving already obtained. The first five repetitions cross at 1585, 1857, 1536, 2129, and 2023 trials, illustrating the remaining pathwise variation. Their continued paths describe learning after the operational decision; the first answers remain fixed.

Table 1: Boundary family error (%, mean $\pm$ binomial MCSE) with exact common-mean adjustment. Each law has 512 independent repetitions and eight questions at nominal family level 5%. Zero-event cells have a Wilson 95% upper limit of 0.75%. Appendix D gives counts and tuning sensitivity.

Response law	Log betting	Empirical Bernstein	AsympCS
Rare, base $t = 0.075$	1.8 ± 0.6	0.2 ± 0.2	31.4 ± 2.1
Rare, base $t = 0.025$	2.3 ± 0.7	0.0 ± 0.0	20.1 ± 1.8
Beta–binomial	1.0 ± 0.4	0.0 ± 0.0	2.7 ± 0.7

Improving observation can raise the attainable information. We next use 16 fixed banks containing 2,692 population histories, each an 8×128 array. Eight principal components and a deterministic eight-category partition summarize each bank within each history group. Four bounded smooth response functions define common activity; centered heterogeneous contrasts preserve each specified group effect. Conditional responses follow a beta–binomial law with concentration 60 and resolution 128. The resolution represents within-response aggregation, not 128 independent interventions. Sixteen banks, two groups, four receivers, and two margins yield 256 questions. Fine and coarse observations use replacement probabilities 0 and 0.65 on the same category; all response laws are held fixed. We sum the finite distributions exactly and solve the information program for each channel and question.

The fine/coarse information ratio has median 1.530 and interquartile range 1.346–1.725 (Figure 2C). Information increases for every question, as predicted by data processing. More strongly, the fine-channel common-mean factor exceeds the coarse channel's *globally optimal* information in 253 of 256 questions; the median ratio is 1.515. Within the fine channel, optimizing adjustment adds at most 5.5%. Thus observation quality can outweigh the entire remaining adjustment opportunity under a shared, optimal evidence objective. These ratios compare leading information time scales; the sequential study separately measures finite learning overhead.

4.3 The response class determines the price of validity

Rare outcomes distinguish finite-sample guarantees from asymptotic approximations. We compare log betting, empirical-Bernstein betting (Howard et al., 2021), and AsympCS (Waudby-Smith et al., 2024a) on three boundary-null laws with positive score variance. Two start from a constant adjusted contrast and introduce rare opposing endpoints using the same contamination mechanism as the validation floor. The third uses ordinary beta–binomial noise. Each experiment contains eight questions at $T = \delta = 0.05$, with an exact common-mean predictor, family error level 5%, and horizon 16,384. Each law has 512 independent repetitions.

Table 1 gives rare-response family errors of 31.4% and 20.1% for AsympCS, versus 1.8% and 2.3% for log betting. Its ordinary-response error is 2.7%. A seven-setting comparison covers sample-time and variance-time tuning at target times 128, 512, and 2048. The most conservative tested variance-time setting lowers the two rare-response errors to 15.4% and 12.9%. The behavior reflects finite monitoring before enough rare responses appear; it is compatible with the asymptotic guarantee. Appendix D gives the tuning units, formulas, and an inaccurate-predictor comparison.

Additional scientific assumptions can buy faster inference. In a matched two-state beta–binomial model, history levels are equally likely and have common means 0.3 and 0.7. We cross concentrations 20 and 60 with margins 0.025 and 0.05, using 256 repetitions per cell. A parametric likelihood-ratio test knows the common means, concentration, and constant-effect family. The bounded information oracle receives the same design alternatives but remains valid over the full bounded randomized class. At concentration 60 and margin 0.05, their costs are 60 ± 3 and 124 ± 3 trials. A regression-adjusted normal interval inspected once at 128 trials answers correctly in 96.1% of repetitions; its coverage is approximate and fixed-time.

Table 2 reports fixed-time power under the normal approximation alongside procedures with uniform sequential guarantees. The two sequential tests share the monitoring rule and error allocation, providing the direct comparison of response assumptions. Specifying the count family gives a larger advantage at lower response noise: at margin 0.05, its mean cost advantage grows from 20 trials at concentration 20 to 64 trials at concentration 60. All sequential paths in this comparison con-

Table 2: Sequential cost (mean $\pm$ MCSE) and correct-decision probability at a matched budget of 128 trials. The bounded procedure uses the finite-history information oracle. The fixed-time comparison is the common-mean adjusted normal interval. All sequential runs conclude correctly by 8192.

		Sequential cost		Correct at 128 (%)		
ν	Δ	Bounded	Parametric	Bounded	Parametric	Fixed-time
20	0.025	561 ± 23	545 ± 23	1.2	2.3	23.8
20	0.050	165 ± 5	145 ± 7	40.2	57.4	71.1
60	0.025	283 ± 8	220 ± 10	4.7	35.2	57.0
60	0.050	124 ± 3	60 ± 3	59.4	91.8	96.1

clude correctly by 8192. Appendix D.2 gives the likelihood-ratio monotonicity argument and the fixed-time interval. The sequential gap measures the benefit of restricting which response laws the experiment must rule out; fixed-time power supplies a separate reference for a precommitted acquisition budget.

5 Discussion

The experimental value of a predictor depends on the effect being validated. When residual noise dominates, common-mean prediction efficiently removes background variance. When noise is small, the range of admissible responses limits how much evidence each intervention can supply. Partial adjustment improves that conversion, while the largest gains can come from observing the state more clearly in the first place.

This distinction suggests a practical sequence for experimental design: define the bounded response and meaningful threshold, estimate the information available under candidate history measurements, and then assess how much of that information a feasible learner recovers within the acquisition budget. The population-history comparison shows why prediction error alone cannot answer the investment question. An improved measurement can raise the information ceiling beyond what further fitting with the existing measurement can achieve.

Two uses of state information lead to different design choices. Refining the observed history while keeping the group and response fixed cannot decrease oracle information. Defining finer groups instead creates new conditional-effect questions. Each subgroup has its own margin and information rate, and a smaller occurrence probability increases calendar cost for a given within-group rate. With F fixed questions and equal directional error allocation, the leading oracle time scale for one question is $\log(2F/\alpha)/(p_G I_H)$. A finer partition can increase some effect margins or reduce within-group variation, while also making groups rarer and allocating less error to each conclusion. Its value depends on which changes dominate at the available budget. The observation-channel comparison isolates one of these choices by holding the groups, responses, and question set fixed.

This distinction also guides model evaluation. Predicting which states have larger effects does not by itself remove the common-activity term in the adjusted score. A predictor of shared fluctuations can accelerate validation of a group-average effect without predicting its heterogeneity. An evaluation should therefore measure prediction of the relevant common activity alongside the evidence or decision cost for prespecified questions. The finite-budget comparison should include fitting trials, since a low population information cost can leave little time to learn.

The analysis assumes independent trials, fixed groups, and known randomization. The finite-history program accommodates unequal known assignment probabilities, whereas carryover or drift requires a corresponding sequential model. Choosing groups, thresholds, or response windows from outcomes would require separate design data or selection-aware error control. The evaluation budget should include any newly acquired pilot trials used to make those choices. Scientific restrictions on response laws can also lower the cost, as the parametric comparison illustrates. The appropriate frontier is therefore determined jointly by the measurement, the decision threshold, and the assumptions that the experiment can support.

AI USE STATEMENT

We used large language models to assist with drafting, structural revision, and language editing of the manuscript. The resulting text was reviewed and revised for consistency with the mathematical assumptions, algorithms, numerical results, and cited literature. We take responsibility for the final manuscript, including its claims, citations, and AI-assisted text.

REPRODUCIBILITY STATEMENT

Section 3 states the estimands, assumptions, and decision rules; Appendix A provides the proofs. Appendices B–D specify the response models, learning updates, comparison procedures, and calibration settings.

REFERENCES

Anastasios N. Angelopoulos, Stephen Bates, Clara Fannjiang, Michael I. Jordan, and Tijana Zrnic. Prediction-powered inference. *Science*, 382(6671):669–674, 2023. doi: 10.1126/science.adi6000.

Eugenio Clerico. On the optimality of coin-betting for mean estimation. *International Journal of Approximate Reasoning*, 187:109550, 2025a. doi: 10.1016/j.ijar.2025.109550.

Eugenio Clerico. Optimal e-value testing for properly constrained hypotheses. arXiv:2412.21125v3, 2025b. URL https://arxiv.org/abs/2412.21125v3.

Eugenio Clerico. Sequential testing of conditionally constrained hypotheses. arXiv:2606.06769v1, 2026. URL https://arxiv.org/abs/2606.06769v1.

Thomas Cook, Alan Mishler, and Aaditya Ramdas. Semiparametric efficient inference in adaptive experiments. In *Proceedings of the Third Conference on Causal Learning and Reasoning*, volume 236 of *Proceedings of Machine Learning Research*, pp. 1033–1064. PMLR, 2024. URL https://proceedings.mlr.press/v236/cook24a.html.

Kayvon Daie, Karel Svoboda, and Shaul Druckmann. Targeted photostimulation uncovers circuit motifs supporting short-term memory. *Nature Neuroscience*, 24(2):259–265, 2021. doi: 10.1038/s41593-020-00776-3.

Robbe L. T. Goris, J. Anthony Movshon, and Eero P. Simoncelli. Partitioning neuronal variability. *Nature Neuroscience*, 17(6):858–865, 2014. doi: 10.1038/nn.3711.

Vitor Hadad, David A. Hirshberg, Ruohan Zhan, Stefan Wager, and Susan Athey. Confidence intervals for policy evaluation in adaptive experiments. *Proceedings of the National Academy of Sciences*, 118(15):e2014602118, 2021. doi: 10.1073/pnas.2014602118.

Junya Honda and Akimichi Takemura. Non-asymptotic analysis of a new bandit algorithm for semi-bounded rewards. *Journal of Machine Learning Research*, 16(113):3721–3756, 2015. URL https://jmlr.org/papers/v16/honda15a.html.

Steven R. Howard, Aaditya Ramdas, Jon McAuliffe, and Jasjeet Sekhon. Time-uniform, nonparametric, nonasymptotic confidence sequences. *The Annals of Statistics*, 49(2):1055–1080, 2021. doi: 10.1214/20-AOS1991.

Nikos Karampatziakis, Paul Mineiro, and Aaditya Ramdas. Off-policy confidence sequences. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 5301–5310. PMLR, 2021. URL https://proceedings.mlr.press/v139/karampatziakis21a.html.

Emilie Kaufmann, Olivier Cappé, and Aurélien Garivier. On the complexity of best-arm identification in multi-armed bandit models. *Journal of Machine Learning Research*, 17(1):1–42, 2016. URL https://www.jmlr.org/papers/v17/kaufman16a.html.

J. L. Kelly, Jr. A new interpretation of information rate. *Bell System Technical Journal*, 35(4): 917–926, 1956. doi: 10.1002/j.1538-7305.1956.tb03809.x.

Martin Larsson, Aaditya Ramdas, and Johannes Ruf. Testing hypotheses generated by constraints. *Mathematics of Operations Research*, 2026. doi: 10.1287/moor.2025.1129. Published online May 12, 2026.

Winston Lin. Agnostic notes on regression adjustments to experimental data: Reexamining Freedman's critique. *The Annals of Applied Statistics*, 7(1):295–318, 2013. doi: 10.1214/12-AOAS583.

Michael Lindon, Dae Woong Ham, Martin Tingley, and Iavor Bojinov. Anytime-valid inference in linear models with applications to regression-adjusted causal inference. *Journal of the American Statistical Association*, pp. 1–12, 2026. doi: 10.1080/01621459.2026.2692052. Published online August 27, 2026.

Celine Mateo, Michael Avermann, Luc J. Gentet, Feng Zhang, Karl Deisseroth, and Carl C. H. Petersen. In vivo optogenetic stimulation of neocortical excitatory neurons drives brain-state-dependent inhibition. *Current Biology*, 21(19):1593–1602, 2011. doi: 10.1016/j.cub.2011.08.028.

Simon Musall, Matthew T. Kaufman, Ashley L. Juavinett, Steven Gluf, and Anne K. Churchland. Single-trial neural dynamics are dominated by richly varied movements. *Nature Neuroscience*, 22(10):1677–1686, 2019. doi: 10.1038/s41593-019-0502-4.

Adam M. Packer, Lloyd E. Russell, Henry W. P. Dalgleish, and Michael Häusser. Simultaneous all-optical manipulation and recording of neural circuit activity with cellular resolution in vivo. *Nature Methods*, 12(2):140–146, 2015. doi: 10.1038/nmeth.3217.

Lloyd E. Russell, Mehmet Fişek, Zidan Yang, Lynn Pei Tan, Adam M. Packer, Henry W. P. Dalgleish, Selmaan N. Chettih, Christopher D. Harvey, and Michael Häusser. The influence of cortical activity on perception depends on behavioral state and sensory context. *Nature Communications*, 15(1):2456, 2024. doi: 10.1038/s41467-024-46484-5.

Shubhanshu Shekhar and Shubhada Agrawal. Dual representation of minimum divergence under integral constraints. arXiv:2603.21027v1, 2026. URL https://arxiv.org/abs/2603.21027v1.

Carsen Stringer, Marius Pachitariu, Nicholas Steinmetz, Charu Bai Reddy, Matteo Carandini, and Kenneth D. Harris. Spontaneous behaviors drive multidimensional, brainwide activity. *Science*, 364(6437):eaav7893, 2019. doi: 10.1126/science.aav7893.

Stefan Wager, Wenfei Du, Jonathan Taylor, and Robert J. Tibshirani. High-dimensional regression adjustments in randomized experiments. *Proceedings of the National Academy of Sciences*, 113 (45):12673–12678, 2016. doi: 10.1073/pnas.1614732113.

Eric Y. Wang, Paul G. Fahey, Zhuokun Ding, Stelios Papadopoulos, Kayla Ponder, Marissa A. Weis, Andersen Chang, Taliah Muhammad, Saumil Patel, Zhiwei Ding, Dat Tran, Jiakun Fu, Casey M. Schneider-Mizell, MICrONS Consortium, R. Clay Reid, Forrest Collman, Nuno Maçarico da Costa, Katrin Franke, Alexander S. Ecker, Jacob Reimer, Xaq Pitkow, Fabian H. Sinz, and Andreas S. Tolias. Foundation model of neural activity predicts response to new stimulus types. *Nature*, 640(8058):470–477, 2025. doi: 10.1038/s41586-025-08829-y.

Ian Waudby-Smith and Aaditya Ramdas. Estimating means of bounded random variables by betting. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 86(1):1–27, 2024. doi: 10.1093/jrsssb/qkad009.

Ian Waudby-Smith, David Arbour, Ritwik Sinha, Edward H. Kennedy, and Aaditya Ramdas. Time-uniform central limit theory and asymptotic confidence sequences. *The Annals of Statistics*, 52 (6):2613–2640, 2024a. doi: 10.1214/24-AOS2408.

Ian Waudby-Smith, Lili Wu, Aaditya Ramdas, Nikos Karampatziakis, and Paul Mineiro. Anytime-valid off-policy inference for contextual bandits. *ACM / IMS Journal of Data Science*, 1(3):1–42, 2024b. doi: 10.1145/3643693.

Edward Wu and Johann A. Gagnon-Bartsch. The LOOP estimator: Adjusting for covariates in randomized experiments. *Evaluation Review*, 42(4):458–488, 2018. doi: 10.1177/0193841X18808003.

Ziyu Xu, Nikos Karampatziakis, and Paul Mineiro. Active, anytime-valid risk controlling prediction sets. In *Advances in Neural Information Processing Systems*, volume 37, pp. 60110–60132. Curran Associates, Inc., 2024. doi: 10.52202/079017-1920. URL https://proceedings.neurips.cc/paper_files/paper/2024/hash/6eb05d8bc6bd7bb6868c64b5802125bd-Abstract-Conference.html.

A Information bounds and optimal adjustment

A.1 Score variance and the rare-response null

All expectations below condition on one group. Write $\sigma_a^2(H) = \mathrm{Var}(U(a) \mid H)$. Fair independent assignment gives

$$\mathbb{E}[\Gamma^g \mid H] = \tau(H), \qquad \mathbb{E}[(\Gamma^g)^2 \mid H] = 2(\sigma_1^2 + \sigma_0^2) + \tau(H)^2 + 4(c-g)^2,$$

which proves Equation (1). To separate missing information from fitting error, let S be a latent response state and C a context included in H, with $U(a) \perp H \mid (S, C)$. Write $c_{S,C}$ and $\tau_{S,C}$ for their conditional common mean and contrast. Conditional projection gives

$$\mathrm{Var}(\Gamma^g) = V_{0,S} + 4\mathbb{E}\,\mathrm{Var}(c_{S,C} \mid H, G) + 4\mathbb{E}\{c(H) - g(H)\}^2, \tag{10}$$

where $V_{0,S} = 2\mathbb{E}(\sigma_{1,S,C}^2 + \sigma_{0,S,C}^2) + \mathrm{Var}(\tau_{S,C})$. The first two terms together equal V_0 in Equation (1).

For Proposition 1, let $\epsilon = \Delta/(1+sT) \in (0,1)$. Inside G, preserve history and assignment, but with probability ϵ replace the potential responses by $(U(1), U(0)) = (0,1)$ for $s=1$ and $(1,0)$ for $s=-1$. Keep the law outside G unchanged. The resulting null satisfies $T_Q = (1-\epsilon)T - \epsilon s = \delta$, and

$$D(P\|Q) = p_G D(P_G\|Q_G) \le p_G \log \frac{1}{1-\epsilon} = p_G \log \frac{1+sT}{1+s\delta}.$$

Data processing through the decision event proves the fixed-budget bound. For a stopping time N with finite P expectation, the stopped likelihood identity is $D(P^N\|Q^N) = \mathbb{E}_P N\, D(P\|Q)$ (Kaufmann et al., 2016). Its increments are integrable: the positive part is bounded by $\log(1/(1-\epsilon))$, and the negative part is integrable using $\log x \le x$. Applying the same decision-event argument proves the stopped bound. The supplied prediction function is unchanged under this competing law.

A.2 Finite histories: an exact computable oracle

We specialize the finite-constraint duality of Clerico (2025b); Larsson et al. (2026); Shekhar & Agrawal (2026). Let $z = 2WU - \delta$ and $c_h = W\mathbf{1}\{H = h\}$. The randomized null is described by $\mathbb{E}_Q z \le 0$ and $\mathbb{E}_Q c_h = 0$ for every h. Consequently every factor in Equation (3) has null expectation at most one. Because it is affine in the response, nonnegativity throughout $[0,1]$ is equivalent to the endpoint constraints in Equation (3).

Finite primal and matching null. Augment the distinct atoms of P by both response endpoints in each (H, A) branch, and denote the resulting finite support by $\mathcal{S}$. This loses no candidate for minimizing $D(P\|Q)$: any null mass outside the alternative support can be moved to the endpoints of its branch while preserving its mass and first response moment. If this places additional mass on an alternative atom, the divergence decreases. Writing $p_j = P\{j\}$, the primal is therefore

$$\min_{q \ge 0} \sum_{j: p_j > 0} p_j \log(p_j/q_j) \quad \text{subject to} \quad \sum_j q_j = 1, \quad \sum_j q_j c_{hj} = 0, \quad \sum_j q_j z_j \le 0.$$

A mixture of a small amount of P and a fully supported endpoint law with effect below δ is strictly feasible. Finite-dimensional convex duality applies. With normalization multiplier η, the dual is

$$\max_{\eta, \lambda, b:\, \lambda \ge 0} \left[1 - \eta + \sum_j p_j \log(\eta + \lambda z_j + b^\top c_j) \right], \qquad \eta + \lambda z_j + b^\top c_j \ge 0 \quad (j \in \mathcal{S}),$$

where only λ is constrained nonnegative; η is initially free. Feasibility implies $\eta \ge 0$ by taking expectation under a boundary null. When $\eta = 0$, the factor vanishes on a boundary null assigning positive mass to every endpoint, and hence on all responses by affinity. Thus any finite dual optimum has $\eta > 0$. Scaling the remaining coefficients by η leaves the term $1 - \eta + \log \eta$, maximized at $\eta = 1$. This gives Equation (4).

Stationarity gives $q_j^* = p_j / f_j^*$ whenever $p_j > 0$; complementarity permits extra mass only where $f_j^* = 0$. The alternative admits positive expected log growth for a sufficiently small positive stake, whereas $\lambda = 0$ gives at most zero by conditional Jensen. Hence $\lambda^* > 0$ and the

matching null has $T_{Q^*} = \delta$. For any null-valid factor E, Jensen's inequality under P gives $\mathbb{E}_P \log(E/f^*) \le \log \mathbb{E}_{Q^*} E \le 0$. This proves optimality over all factors. The finite alternative support also makes $\log f^*$ bounded, giving the same stopped change-of-measure bound as in Proposition 1. Data processing through a fixed history channel gives $I_H(P_{\text{coarse}}) \le I_H(P_{\text{fine}})$.

Known history probabilities and unequal randomization. The null above allows the history marginal to change. If its probabilities π_h are known externally, factors may additionally include offsets κ_h with $\sum_h \pi_h \kappa_h = 0$, and these can improve growth. The symmetric binary matching null already preserves its history marginal; the general finite-history matching null need not. For a known propensity $e_h \in (0,1)$, use $Z = AU/e_H - (1-A)U/(1-e_H)$ and $C_h = (A - e_h)\mathbf{1}\{H = h\}$. The same argument gives $f = 1 + \lambda(Z - \delta) + b_H(A - e_H)$, with

$$\frac{\lambda\delta - 1}{1 - e_h} \le b_h \le \frac{1 - \lambda\delta - \lambda/(1-e_h)}{e_h}.$$

This remains a convex program with $\lambda \le 1/(1+\delta)$.

The adjustment map may depend on the current history after its coefficients have been chosen from past trials. The contrast coefficient λ is shared across histories: replacing it by λ_h generally weights the conditional effects unequally. An unbiased reweighting that preserves the same average contrast for every response law must have a common net contrast coefficient. More general valid constructions are already covered by the all-factor optimum in Equation (4).

A.3 Solving the binary family and constructing its matching null

Write $W = 2A - 1$, $X = WH$, and $Y = WR$. Then X, Y are independent fair signs and $\Gamma^{g_r} = T + (h - r)X + \sigma Y$. The valid range is $\lambda(d + |r|) \le 1$. For fixed λ, expected log growth is even and concave in the residual amplitude $h - r$. Hence its best feasible adjustment is

$$r(\lambda) = \min\{h, 1/\lambda - d\}, \qquad 0 < \lambda \le 1/d.$$

For $\lambda \le 1/L$, full adjustment gives $J = \frac{1}{2}\log\{(1+\lambda\Delta)^2 - \lambda^2\sigma^2\}$. Its unconstrained optimum lies in this interval exactly when $\sigma^2 \ge \Delta(L + \Delta)$, yielding the first case of Equation (7).

On $1/L \le \lambda \le 1/d$, put $C = L + \Delta$, $D = L - \Delta$. The factors on $X = 1$ and $X = -1$ are $f_{+,y} = \lambda(C + \sigma y)$ and $f_{-,y} = 2 - \lambda(D - \sigma y)$, respectively. Thus

$$J = \frac{1}{4}\log\left\{\lambda^2(C^2 - \sigma^2)[(2 - \lambda D)^2 - \lambda^2\sigma^2]\right\}. \tag{11}$$

Stationarity reduces to $2 - 3\lambda D + \lambda^2(D^2 - \sigma^2) = 0$, with applicable root $4/(3D + \sqrt{D^2 + 8\sigma^2})$. The derivative at $1/d$ is nonnegative precisely when $\Delta \ge h$ and $\sigma^2 \le (\Delta - h)(d + \Delta - h)$. These conditions produce the remaining two cases; the partial-regime conditions ensure a positive denominator even when $D < 0$. Concavity in $(\lambda, \lambda r)$ establishes the optimum in this factor class. The following matching law establishes its unrestricted optimality and preserves the history frequencies as well.

Give every alternative atom its P mass divided by f; combine any coincident atoms. Let $S = \frac{1}{2}\sum_y 1/f_{+,y}$ and $W_- = \frac{1}{2}\sum_y 1/f_{-,y}$. Each (H, A) branch then has mass $S/4$ or $W_-/4$. In the full-adjustment regime, scalar stationarity gives $S = W_- = 1$. In the partial regime, differentiating Equation (11) gives $J'(\lambda) = (1 - W_-)/\lambda$, so $W_- = 1$. Since $h - r \ge 0$, paired factors obey $f_{+,y} \ge f_{-,y}$ and $S \le W_-$. Add $(1 - S)/4$ to the opposing endpoint in each $X = 1$ branch. In the no-adjustment regime, $J'(1/d) \ge 0$ gives $W_- \le 1$; add masses $(1 - S)/4$ and $(1 - W_-)/4$ in the respective branch types.

Every branch now has mass $1/4$. The added endpoint is $U = 0$ in treatment and $U = 1$ in control, where $f = 1 - \lambda(d + rX) = 0$ whenever added mass is positive. The same formulas cover coincident atoms, regime boundaries, $h = 0$, and $\sigma = 0$. Therefore

$$f = dP_G/dQ_G^*, \qquad 1 = \mathbb{E}_{Q_G^*} f = 1 + \lambda^*(T_{Q_G^*} - \delta), \qquad D(P_G \| Q_G^*) = I_H.$$

The null is randomized and has $T_{Q_G^*} = \delta$; its conditional arm laws can always be coupled into potential outcomes independent of assignment. The entropy argument used above gives optimality over all factors. Preserving the outside-group law gives the all-procedure lower bound. Under P, log increments are bounded and have positive mean I_H. A Chernoff bound gives finite expected crossing time; Wald's identity and bounded overshoot yield Equation (8).

A.4 Hidden history and the noise–margin limit

Without history, (A, U) is equivalently (A, Z) with $Z = 2(2A-1)(U-0.5) = T + hX + \sigma Y$, independent of fair A. Its exact rate is

$$I_{\varnothing} = \max_{0 \le \lambda \le 1/d} J(\lambda, 0). \tag{12}$$

For the maximizing stake, assign each distinct score atom null mass $P(Z = z)/\{1 + \lambda(z - \delta)\}$. At an interior optimum these masses sum to one and have mean δ. At $\lambda = 1/d$ their sum is at most one; place the remainder at $z = -1$, again giving mean δ. Take independent fair A and potential responses $U(1) = (1 + Z)/2$, $U(0) = (1 - Z)/2$. The resulting randomized null has the displayed scalar factor as its likelihood ratio, proving the rate for the complete reduced observation.

For fixed $\sigma > 0$, the full-adjustment formula gives $I_H \sim \Delta^2/(2\sigma^2)$, whereas expansion of Equation (12) gives $I_{\varnothing} \sim \Delta^2/\{2(h^2 + \sigma^2)\}$. For fixed $h > 0$, $d > 0$ and $\sigma^2/\Delta \to c < L$, the partial solution satisfies

$$\lambda^* = L^{-1} + O(\Delta), \qquad h - r^* = \Delta(1 - c/L) + o(\Delta), \qquad I_H = \Delta/L - \sigma^2/(2L^2) + o(\Delta).$$

For $c > L$, use the full-adjustment formula; continuity gives the common limit at $c = L$. This proves Equation (9). At zero noise, $r^* = (h - \Delta)_+$ and

$$I_H = \begin{cases} \frac{1}{2}\log\{(d + h + \Delta)/(d + h - \Delta)\}, & \Delta \le h, \\ \frac{1}{2}\log\{[(d + \Delta)^2 - h^2]/d^2\}, & \Delta \ge h. \end{cases}$$

For $\sigma = 0 < h$ and $\Delta \to 0$, hidden history gives $I_{\varnothing} \sim \Delta^2/(2h^2)$, while both observed-history adjustments have leading rate Δ/L.

Size of the adjustment gain. At fixed $\delta = 0.05$, numerical optimization over the physical binary family with three independent differential-evolution starts finds an information ratio $I_H/I_{\text{common}} \approx 1.2095$ near $(h, \Delta, \sigma) = (0.4004, 0.5496, 0)$, approaching the response-support boundary. This is a numerical envelope rather than a proved global upper bound. There is no finite envelope uniform over all thresholds. For example, fix $h = 0.3$, $\sigma = 0$, and $T = 0$. As $\delta \downarrow -1$, the no-adjustment regime gives

$$I_H = \log\frac{1}{1 + \delta} + \frac{1}{2}\log(1 - h^2) \longrightarrow \infty, \qquad I_{\text{common}} \longrightarrow \log(1 + 1/h).$$

Thus the information ratio is unbounded when the threshold approaches an endpoint. The finite-trial stopping comparison also includes overshoot and the minimum of one observation.

A.5 Predictable learning

In the binary family, $2U_i - 1 = hH_i + TW_i + \sigma R_i$. The empirical design matrix for (H_i, W_i) tends to the identity, so past-trial least squares estimates (h, T) consistently; its mean squared residual tends to σ^2. Projection onto $\{(h, T, \sigma) : h, \sigma \ge 0,\ h + \sigma + |T| \le 1\}$ preserves consistency at the interior alternative. For each sign, a nonpositive estimated margin gives zero stake; otherwise apply Equation (7), capped at $\{1 - (m + 1)^{-1/2}\}/(1 + r + s\delta)$ after m completed trials. Choices depend only on the past, and under every bounded randomized null the signed factor has conditional expectation at most one. This proves finite-sample validity throughout that observation model.

Inside the family, consistency and continuity of the optimizer make the learned factor converge to f^* on every alternative atom. The support is finite and $f^* > 0$ there, so the log increments converge uniformly on that support. Cesàro averaging and the iid strong law give $n^{-1}\log K_n \to I_H$; positive limiting slope then gives $\tau_\ell/\ell \to 1/I_H$ almost surely as $\ell \to \infty$. These are statements about continued directional wealth; they do not interchange a limit and expected stopping time. For the finite-history learner, precommitted feasible coefficients likewise ensure validity, independently of the quality of its fitted response distribution.

B Algorithms and complete numerical designs

Group-specific processes update on group visits and hold their wealth constant outside the group. For a signed question s, exchanging the assignment labels when $s = -1$ gives target sT and thresh-

old $s\delta$. All estimators below use completed observations; their feasible factors therefore retain the conditional validity established in Appendix A.

B.1 Binary-family design and complete cellwise results

The 24 cells in Table 3 use $h = 0.30$ and $\delta = 0.05$ and contain four full-adjustment, 18 partial-adjustment, and two no-adjustment settings. Each has 1,024 independent trial streams and a horizon of 32,768 group trials. Different cells have independent random streams; methods and thresholds within each cell share all observations. Responses follow Equation (5), so their conditional arm variance is $\sigma^2/4$.

The no-history, common-mean, and frontier oracles use their respective population-optimal stakes. Their negative-direction stakes are zero under these positive alternatives. The learned procedure fits least squares of $2U - 1$ on H and $2A - 1$ using only completed observations, with a diagonal ridge of 10^{-8} for initialization. Residual mean squared error estimates σ^2. The history coefficient is truncated below at zero, and the parameter vector is radially projected onto $h + \sigma + |T| \leq 1$. Updates occur after 16, 32, 64, and every subsequent 64 completed trials. Stakes are zero before the first update. Both signed directions use these estimates with separate optimizers, and their stakes are capped by $\{1 - (m+1)^{-1/2}\}/(1 + r + s\delta)$ after m observations. All fitting observations count toward the stopping cost.

The primary endpoint is the first signed crossing of 40. Wrong decisions, conflicting simultaneous crossings, and unresolved paths retain cost 32,768 in the restricted correct-decision mean. Two learned paths make a wrong decision, both in the $(\Delta, \sigma) = (0.01, 0.30)$ cell; there are no conflicting crossings. Continued wealth also records crossings of 10^3 and 10^6 on the same paths. These higher thresholds do not change the primary decisions.

Table 3 gives cellwise uncertainty, avoiding a pooled standard error across distinct parameter settings. Figure 1 uses the ratio of learned and no-history mean costs; its Monte Carlo standard error uses their paired pathwise covariance. Zero-noise common-mean factors are deterministic, whereas partial adjustment retains assignment-induced variation. The horizon-restricted costs in the table differ from the uncensored expectation bounded in Equation (8).

C Finite-history bridge: distributions and algorithms

Eight-state response family. For the latent category $Z \in \{1, \ldots, 8\}$, the probabilities and common responses are

$$\pi = (0.28, 0.20, 0.16, 0.12, 0.09, 0.07, 0.05, 0.03),$$
$$c = (0.22, 0.30, 0.37, 0.44, 0.55, 0.64, 0.72, 0.78).$$

Let $z_j = -1 + 2(j-1)/7$ and $\tau_j = 0.05 + \Delta + 0.04(z_j - \sum_k \pi_k z_k)$. For a random variable R taking values $(-1, 0.2, 0.8)$ with probabilities $(3/10, 13/30, 4/15)$, define

$$U(a) = c_Z + (a - 1/2)\tau_Z + \sigma q_Z k_a R,$$
$$q = (0.4, 0.55, 0.8, 1, 0.65, 0.9, 0.5, 0.75), \quad (k_0, k_1) = (0.8, 1.2).$$

Here R is independent of Z and fair assignment. Its mean is zero, so the target effect is exactly $0.05 + \Delta$; its asymmetry and the $q_Z k_a$ multipliers produce nonconstant, asymmetric response noise. All responses in the grid lie between 0.1014 and 0.9706. The observed category equals Z with probability $1 - \epsilon$; otherwise it is an independent draw from π. The replacement is unmarked. Its transition matrix is $C_\epsilon = (1 - \epsilon)I + \epsilon \mathbf{1}\pi^\top$. For $\epsilon_2 \geq \epsilon_1$, applying another replacement channel of strength $(\epsilon_2 - \epsilon_1)/(1 - \epsilon_1)$ to C_{ϵ_1} gives C_{ϵ_2}. Thus these observations are ordered by information loss.

The exact-rate grid crosses $\Delta \in \{0.01, 0.04, 0.08, 0.16\}$, $\sigma \in \{0.01, 0.04, 0.08\}$, and $\epsilon \in \{0, 0.35, 0.70\}$. The probability of every observed (H, A, U) atom is obtained by summing over Z and R. The information oracle optimizes the finite-history program in the main text. The common-mean comparator fixes $g(H) = \{\mathbb{E}[U(0) \mid H] + \mathbb{E}[U(1) \mid H]\}/2$ and optimizes its stake. The no-history oracle optimizes both its constant adjustment and stake. This gives each comparator its best population stake under the same bounded-response null model. Table 5 reports the complete grid of information gains over common-mean adjustment.

Table 3: Binary-family first decisions at evidence threshold 40. Costs are means ± Monte Carlo standard errors over 1,024 independent trial streams per cell; all methods within a stream share observations. Correct-attainment percentages list no-history/frontier/learned methods in that order. Wrong or unresolved decisions retain cost 32,768. I_H/I_{common} compares population growth rates.

σ	Δ	Regime	No history	Frontier	Learned	Correct (%)	I_H/I_{common}
0.00	0.01	Partial	6826 ± 170	498 ± 0.7	558 ± 0.7	99.3/100.0/100.0	1.004
0.00	0.02	Partial	1600 ± 34	249 ± 0.5	295 ± 0.5	100.0/100.0/100.0	1.007
0.00	0.04	Partial	411 ± 9.7	125 ± 0.3	161 ± 0.4	100.0/100.0/100.0	1.015
0.00	0.08	Partial	103 ± 2.3	64 ± 0.3	91 ± 0.3	100.0/100.0/100.0	1.031
0.00	0.16	Partial	34 ± 0.4	32 ± 0.2	55 ± 0.2	100.0/100.0/100.0	1.063
0.00	0.32	None	16 ± 0.1	16 ± 0.1	36 ± 0.1	100.0/100.0/100.0	1.135
0.05	0.01	Partial	6803 ± 156	549 ± 4.0	617 ± 4.2	99.8/100.0/100.0	1.003
0.05	0.02	Partial	1697 ± 39	263 ± 1.4	310 ± 1.4	100.0/100.0/100.0	1.006
0.05	0.04	Partial	430 ± 9.9	128 ± 0.6	164 ± 0.6	100.0/100.0/100.0	1.014
0.05	0.08	Partial	112 ± 2.4	64 ± 0.3	92 ± 0.3	100.0/100.0/100.0	1.029
0.05	0.16	Partial	35 ± 0.4	32 ± 0.2	55 ± 0.2	100.0/100.0/100.0	1.062
0.05	0.32	None	16 ± 0.1	16 ± 0.1	36 ± 0.1	100.0/100.0/100.0	1.134
0.15	0.01	Full	8022 ± 175	1692 ± 39	2168 ± 46	99.2/100.0/100.0	1.000
0.15	0.02	Partial	2036 ± 49	434 ± 8.6	545 ± 9.7	100.0/100.0/100.0	1.000
0.15	0.04	Partial	537 ± 12	163 ± 2.0	206 ± 2.2	100.0/100.0/100.0	1.007
0.15	0.08	Partial	135 ± 3.1	72 ± 0.6	100 ± 0.6	100.0/100.0/100.0	1.021
0.15	0.16	Partial	39 ± 0.6	34 ± 0.2	57 ± 0.3	100.0/100.0/100.0	1.053
0.15	0.32	Partial	17 ± 0.1	17 ± 0.1	36 ± 0.1	100.0/100.0/100.0	1.124
0.30	0.01	Full	13222 ± 260	6497 ± 149	9983 ± 217	94.5/99.7/98.8	1.000
0.30	0.02	Full	3345 ± 76	1634 ± 36	2278 ± 51	100.0/100.0/100.0	1.000
0.30	0.04	Full	838 ± 19	402 ± 9.1	530 ± 11	100.0/100.0/100.0	1.000
0.30	0.08	Partial	214 ± 4.8	109 ± 2.0	152 ± 2.3	100.0/100.0/100.0	1.002
0.30	0.16	Partial	53 ± 1.2	41 ± 0.5	65 ± 0.5	100.0/100.0/100.0	1.028
0.30	0.32	Partial	18 ± 0.2	18 ± 0.2	38 ± 0.2	100.0/100.0/100.0	1.092

Table 4: Finite-history restricted correct-decision cost, with MCSE in parentheses, across 256 independent repetitions per row. All methods share the trial stream within a row. The horizon is 16,384.

Δ	ϵ	No history	Common oracle	Information oracle	Empirical log	Empirical common
0.01	0	7262 (255)	637 (7)	634 (8)	1818 (14)	1817 (14)
0.01	0.70	7778 (291)	7184 (273)	7185 (273)	9306 (322)	9295 (322)
0.04	0	488 (22)	153 (1)	148 (2)	452 (3)	452 (3)
0.04	0.70	500 (21)	461 (20)	465 (20)	800 (21)	801 (21)
0.16	0	41 (1)	40 (0.1)	36 (0.5)	122 (1)	123 (1)
0.16	0.70	43 (1)	45 (1)	42 (1)	139 (2)	139 (2)

Numerical solution. Writing $d = 1 + \delta$, $\theta = d\lambda$, and $v_h = \lambda(2g(h) - 1)$ gives the factor

$$f = 1 + \frac{\theta}{d}\{W(2U - 1) - \delta\} - Wv_H, \qquad \theta \geq 0, \quad \theta + |v_h| \leq 1.$$

For fixed θ, each v_h solves a one-dimensional concave problem; the outer problem then has one variable. We find interior stationary points by derivative roots and include the constraint boundaries. To avoid numerical cancellation at extremely small beta–binomial endpoint probabilities, the implementation contracts the feasible set by 10^{-10}. If a is the objective gradient at the returned (θ, v), concavity bounds the gap to the uncontracted optimum by

$$\max\{a_\theta, \textstyle\sum_h |a_h|\} - a_\theta\theta - \sum_h a_h v_h.$$

This certificate is below 1.8×10^{-9} for every reported finite-family and population-history optimum.

Learning and stopping. Six sequential conditions use $\sigma = 0.04$, $\Delta \in \{0.01, 0.04, 0.16\}$, and $\epsilon \in \{0, 0.70\}$. Each contains 256 independent repetitions of at most 16,384 trials. Conditions use separate random-number streams, and methods within a condition share the same observations. Both signed decisions require wealth 40, allocating total error 0.05 across the two signs. The three population comparators know the alternative direction and place zero stake in the opposite direction. The two learned rules monitor both directions. First decisions are absorbing; the cost of a wrong, conflicting, or unresolved first decision is the horizon.

Both learners begin after 16 trials and update at completed sample sizes $16, 32, 64, \ldots, 8192$. If n_h and n_{ha} are the history and history–arm counts, their fitted law has

$$\widehat{P}(H = h) = \frac{n_h + 2}{n + 16}, \qquad \widehat{P}(A = a \mid H = h) = \frac{1}{2},$$
$$\widehat{P}(U \in \cdot \mid h, a) = \frac{\sum_{i:H_i=h,A_i=a} \delta_{U_i} + \frac{1}{2}\delta_0 + \frac{1}{2}\delta_1}{n_{ha} + 1}.$$

The endpoint pseudo-counts stabilize sparsely observed categories. The empirical-log learner optimizes the finite-history objective under this law. The empirical-common learner instead fixes its fitted common mean and optimizes only the stake. Before the next block, both multiply all fitted evidence coefficients by $1 - (n+1)^{-1/2}$. This mixes the proposed factor with the constant factor one, preserving validity. Neither the fitted law nor the coefficients use the next trial's history. Opposite signs use $W \mapsto -W$ and $\delta \mapsto -\delta$.

Table 4 gives the six conditions' costs and MCSEs. No first decision is wrong or conflicting. At margin 0.01 with accurate history, no history answers 244/256 repetitions and every other method answers 256/256. With replacement probability 0.70, the corresponding completion counts are 221, 234, 234, 205, and 205. Every method completes all repetitions in the remaining four conditions. The information oracle's average cost need not be below every competitor's finite-horizon cost in each Monte Carlo sample: its optimality concerns population log growth.

Population-history response law. Each of the 16 fixed banks contains population histories of shape 8×128, with 2,692 histories in total. Within each bank, we flatten the histories, center each coordinate, and compute the first eight principal-component scores. Each score is divided by its population standard deviation; its sign makes the loading of largest magnitude positive. Write the

Table 5: Exact information gains in the eight-state family. Entries give $100(I_H/I_{\text{common}} - 1)$ for the full grid; both rules use optimized population stakes.

σ	ϵ	$\Delta = 0.01$	0.04	0.08	0.16
0.01	0.00	0.93	3.91	7.52	13.17
0.01	0.35	0.01	0.28	2.93	9.45
0.01	0.70	0.00	0.02	0.39	5.80
0.04	0.00	0.80	3.78	7.40	13.07
0.04	0.35	0.01	0.26	2.82	9.34
0.04	0.70	0.00	0.02	0.33	5.70
0.08	0.00	0.44	3.36	7.04	12.76
0.08	0.35	0.01	0.22	2.47	9.00
0.08	0.70	0.00	0.02	0.19	5.40

resulting row as $x \in \mathbb{R}^8$ and define $G = \mathbf{1}\{x_1 \geq 0\}$. Sampling selects G fairly and then a uniform bank row within that group. No response enters these features.

For receiver $q \in \{1, \ldots, 4\}$, common activity is

$$c_q(x) = \frac{1}{2} + a_q \tanh\big(w_q^\top x + b_q \tanh(x_{2q-1}) \tanh(x_{2q})\big)\,,$$
$$a = (0.28, 0.30, 0.26, 0.32), \quad b = (0.30, -0.30, 0.30, -0.30),$$

where each w_q is the corresponding row of the following matrix, divided by its Euclidean norm:

$$\begin{pmatrix} 1 & 0.4 & 0.2 & 0 & 0 & 0.2 & 0 & 0 \\ -0.3 & 0.8 & 0 & 0.4 & 0 & 0 & 0.2 & 0 \\ 0.2 & 0 & 0.8 & -0.4 & 0.2 & 0 & 0 & 0 \\ 0 & 0.3 & -0.2 & 0.7 & 0 & 0.3 & 0 & 0.2 \end{pmatrix}.$$

Let m_{gq} be the within-group bank mean of $\tanh(x_{q+4})$. The contrast and arm means are

$$\tau_q(x) = 0.05 + D_{gq}\Delta + 0.025\{\tanh(x_{q+4}) - m_{gq}\},$$
$$D = \begin{pmatrix} -1 & -1 & 1 & 1 \\ -1 & 1 & -1 & 1 \end{pmatrix}, \quad \mu_{aq}(x) = c_q(x) + (a - 1/2)\tau_q(x).$$

Thus the group effect is exactly $0.05 + D_{gq}\Delta$. Conditional on the row and assignment, draw $P_q \sim \text{Beta}(60\mu_{aq}, 60(1 - \mu_{aq}))$ and $128U_q \sim \text{Binomial}(128, P_q)$. The resolution 128 defines the bounded response; one draw remains one intervention. Both channels below keep this response law fixed.

Observation channels. Within each group, we cluster the eight-dimensional coordinates into eight categories by Lloyd's algorithm. Initialization chooses the point farthest from the overall mean, then repeatedly chooses the point farthest from the previously chosen centers; ties use the lowest row index. Updates use cluster means, stop when labels are unchanged (at most 200 updates), and remove empty categories.

The fine channel observes the category. The coarse channel replaces it with probability 0.65 by an independent category drawn from its within-group marginal. The unobserved bank row and its response law are identical under both channels. We sum the beta–binomial probability of each response value $0, 1/128, \ldots, 1$ over all rows assigned to each observed category, giving an exact finite joint law. Two margins, 0.025 and 0.05, yield 256 questions across banks, groups, and receivers, each evaluated under both channels. These calculations introduce no Monte Carlo error.

The exact fine/coarse oracle information ratios range from 1.007 to 2.654. For the fine-common/coarse-optimum comparison, the three exceptions to an improvement have ratios below one, with a minimum of 0.954.

D CALIBRATION AND RESPONSE ASSUMPTIONS

D.1 RARE RESPONSES AND ASYMPTOTIC TUNING

All eight questions are at $T = \delta = 0.05$, with group probability $1/2$, family level 0.05, and 16,384 trials. Each law has 512 independent repetitions. Predictors $g = 0.5$ and 0.1 give score bounds 1 and 1.8. Each trial has one fair group and one fair assignment shared across four receivers. The rare-outcome laws start with $U(a) = \{1 + (2a - 1)t\}/2$, where $t = 0.075$ or 0.025. For $s = \mathrm{sign}(t - \delta)$, independently across trials and receivers, replace the potential response pair by the opposite endpoints with probability $\epsilon = |t - \delta|/(1 + st)$; then $(1 - \epsilon)t - \epsilon s = \delta$. The ordinary law is beta–binomial with 128 counts, arm means $0.475, 0.525$, and concentration 60. Every law has positive score variance.

Finite-sample evidence rules. For each question, let $z_i = 2(2A_i - 1)(U_i - g)$ on its ith group visit and $M = 2\max(g, 1 - g)$. Log betting uses $x_i^s = s(z_i - \delta)/(M + s\delta)$ and factors $1 + b_s x_i^s$. After calendar trials 64, 128, 256, and every 256 thereafter, choose $b_s \in [0, 0.99]$ to maximize $\sum_{i \le m} \log(1 + b_s x_i^s) + \log(1 - b_s^2)$ over completed visits. The two pseudo-observations represented by the final term stabilize tuning and do not contribute to wealth. The initial stake is 0.15.

For empirical Bernstein, define the past-score predictor $a_i = \mathrm{clip}_{[-M,M]}(\sum_{j<i} z_j/i)$ and $v_i = (z_i - a_i)^2/(4M^2)$. At the same checkpoints, use

$$D_s = \sum_{i \le m} \frac{s(z_i - \delta)}{2M}, \quad b_s = \min\left\{0.99, \frac{(D_s)_+}{(D_s)_+ + 1 + \sum_{i \le m} v_i}\right\}.$$

Its next log increment is $b_s s(z_i - \delta)/(2M) + \{\log(1 - b_s) + b_s\} v_i$; the initial stake is again 0.15. The predictable empirical-Bernstein inequality ensures null validity (Howard et al., 2021). Both rules hold their stake until the next checkpoint, monitor from the first visit, and cross at $320 = 2 \cdot 8/0.05$. A crossing in either direction at this boundary is a family error, including simultaneous crossings.

Asymptotic evidence and tuning. Let $V_m = \sum_{i \le m} z_i^2 - m\bar{z}_m^2$, $a = 0.05/8$, and $c_a = -W_{-1}(-a^2/e) - 1$. We compare the variance-time and sample-time half-widths from Theorems 2.5 and 2.2 of Waudby-Smith et al. (2024a):

$$h_m^{\mathrm{V}} = \frac{\sqrt{(V_m + \rho^{-2}) \log\{(1 + \rho^2 V_m)/a^2\}}}{m}, \qquad \rho^2 = c_a/(t_* v_*),$$

$$h_m^{\mathrm{T}} = \frac{\sqrt{V_m \{1 + (m\rho^2)^{-1}\} \log\{(1 + m\rho^2)/a^2\}}}{m}, \qquad \rho^2 = c_a/t_*.$$

The reference $t_* = 512$, $v_* = 0.05$ targets accumulated variance 25.6, not a sample size of 25.6. Sensitivity uses $t_* \in \{128, 512, 2048\}$; variance-time variants receive the population score variance as v_*. All variants share trials and start monitoring after 128 group observations.

Table 6: Boundary family errors out of 512 repetitions per law. Columns pair the accurate ($g = 0.5$) and deliberately offset ($g = 0.1$) predictors. Variance-time sensitivity uses each law's population score variance; the reference uses $v_* = 0.05$.

	Rare, $t = 0.075$		Rare, $t = 0.025$		Beta–binomial	
Procedure	0.5	0.1	0.5	0.1	0.5	0.1
Continuous log betting	9	7	12	6	5	5
Empirical Bernstein	1	2	0	1	0	2
Variance time, reference	161	4	103	4	14	4
Sample time, $t_* = 128$	290	7	237	6	12	7
Sample time, $t_* = 512$	295	11	241	6	17	8
Sample time, $t_* = 2048$	276	10	212	10	11	7
Variance time, $t_* = 128$	289	7	233	6	12	7
Variance time, $t_* = 512$	267	11	201	6	17	8
Variance time, $t_* = 2048$	79	10	66	10	11	7

Table 6 gives all error counts. Under the ordinary law, the seven settings give family errors of 2.1–3.3%. Adding assignment variation with $g = 0.1$ brings every setting below 5% in these repetitions.

D.2 What additional response assumptions buy

History has two equally likely levels, with known common means $c(H) \in \{0.3, 0.7\}$. Conditional on history and balanced assignment, $128U$ is beta–binomial with mean parameter $c(H) + (2A - 1)T/2$ and concentration ν. We cross $\nu \in \{20, 60\}$ with $T = 0.05 + \Delta$, $\Delta \in \{0.025, 0.05\}$, using 256 independent repetitions per condition. Both sequential tests monitor both signs against the design alternatives $T_s = 0.05 + s\Delta$, at threshold 40. The bounded test uses the finite-history information-optimal factor, computed separately at each T_s. The parametric test uses $L_n^s = \prod_{i \le n} p_{T_s}(K_i \mid H_i, A_i)/p_{0.05}(K_i \mid H_i, A_i)$ with known common means and concentration.

The latter's validity follows from the beta–binomial adjacent ratio,

$$\frac{p_\mu(k+1)}{p_\mu(k)} = \frac{128 - k}{k + 1} \frac{k + \nu\mu}{127 - k + \nu(1 - \mu)}.$$

It increases with μ, giving monotone likelihood ratios. Consequently, each signed likelihood ratio has conditional mean at most one on its corresponding constant-effect null. This guarantee uses the specified count family and its known nuisance parameters; bounded betting allows any randomized bounded-response law.

At each separately prescribed budget, the fixed-time comparator uses $\bar{z}_n \pm 1.96\, s_z/\sqrt{n}$, with $z_i = 2(2A_i - 1)\{U_i - c(H_i)\}$. This is a regression-adjusted normal interval, with approximate fixed-time coverage. It is inspected once at the designated budget. Table 2 compares decisions on the same trial streams. One fixed-time decision is wrong at 128 trials in the $(20, 0.025)$ condition.