CONTRAST-LEVEL POSTERIOR REGRET CERTIFICATES FOR QUADRATIC COMPOSITES WITH UNCERTAIN GP HYPERPARAMETERS

Anonymous authors
Paper under double-blind review

ABSTRACT

Learning Gaussian-process hyperparameters changes uncertainty in nonlinear composite objectives. We study posterior regret certificates for finite-design squared-norm losses over continuous hyperparameter families. A common posterior coupling yields event-transfer bounds from signed-quadratic contrast coefficients without assuming sub-Gaussian composite tails. Under a correctly specified model, predictable observations and valid retained-family coverage, the bounds control joint false certification at a fixed terminal time. Structured two-design posteriors separate intrinsic linear ambiguity from the logarithmic penalty of a specified optimized concentration formula. An explicit event-specific remainder then gives a numerically validated finite separation: at the same posterior and parameter box, our certificate is below a fixed risk threshold while an optimally tilted, sign-conditioned information bound remains inconclusive. Informative observations reduce disagreement across a fixed family, with a sharp boundary rate, eventual almost-sure narrowing and a failure case for an unobserved relevant direction. Direct robust probability remains tighter in the finite example, and the information bound wins in a high-information case. The results characterize uncertainty and certificate conservatism; they do not establish faster Bayesian optimization.

1 INTRODUCTION

A composite objective maps latent responses to the loss used to select a design. Gaussian posterior distributions for those responses do not generally yield a Gaussian distribution for the loss. For example, if Y is standard normal, Y^2 has variance two although linearization at the posterior mean predicts no variation. A product of independent standard normals likewise has heavier tails than a sub-Gaussian variable. Gaussian component uncertainty and smooth composition therefore do not, by themselves, justify a reliable bound on the loss.

A second difficulty arises when hyperparameters are estimated from the same observations used for optimization. A regret probability computed with a fitted kernel describes the selected model. Other models compatible with the same data can assign a different probability to the regret event. Existing posterior-regret stopping methods already address simulation of the regret event and Monte Carlo error (Wilson, 2024). We ask how to bound the change in that event probability across a continuous family of GP hyperparameters without comparing the entire latent distribution.

Figure 1 shows the resulting decision problem: a recommendation can pass the risk budget under one posterior and fail under another. A reliable family-wide decision must account for this change in the regret event.

For squared-norm composites, the difference between two objective values is a signed quadratic form. Under a shared posterior coupling, its change across hyperparameters has a mean, a linear Gaussian component and a centered quadratic component. Bounding these coefficients retains cancellations that separate bounds on the two losses can discard. The operator norm controls the large-deviation term, while the Hilbert–Schmidt and linear norms control the square-root term. This distinction makes the effect of nonlinear uncertainty explicit.

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

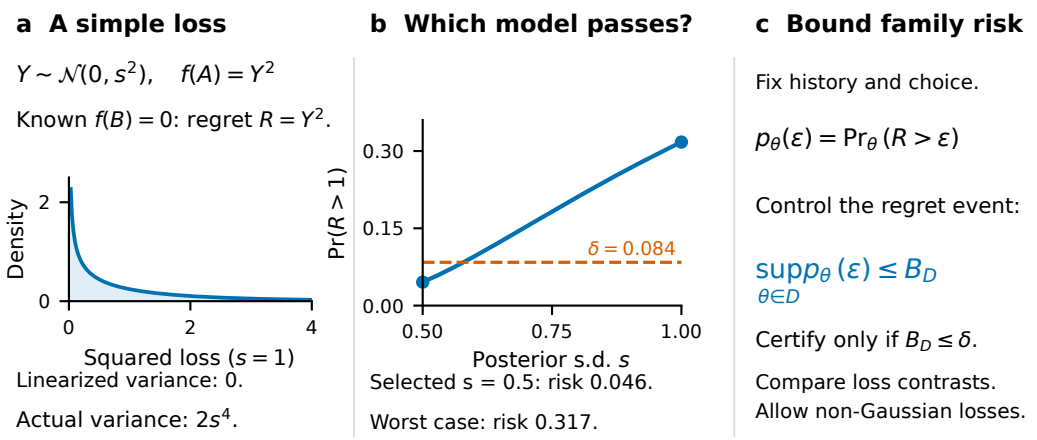

Figure 1: **A Gaussian response does not make a decision robust to model choice.** (a) In an elementary illustration, recommend A with loss Y^2, while B has known loss zero. Regret is $R = Y^2$ and linearization misses its variance. (b) For $Y \sim \mathcal{N}(0, s^2)$, the exact risk $\Pr(R > 1) = 2\Phi(-1/s)$ crosses the dashed risk budget as s ranges over $[0.5, 1]$. Selecting $s = 0.5$ alone would pass; the family would not. (c) The paper bounds the regret event across a hyperparameter family using loss contrasts. Panels (a–b) are an elementary analytical motivation; the RBF certificate comparisons appear later.

The analysis has three parts. First, we derive posterior regret bounds across continuous hyperparameter cells from the RBF contrast coefficients. Combining these bounds with family retention and fresh posterior simulation gives a joint false-certification guarantee at a fixed observation budget. Second, we distinguish intrinsic hyperparameter ambiguity from certificate conservatism: a structured posterior has linear local ambiguity, while a specified optimized concentration formula adds a logarithmic penalty. An event-specific remainder then establishes a finite decision advantage over an optimally tilted, sign-conditioned information bound at the same posterior. Third, we characterize how informative observations reduce disagreement across a fixed family, including a sharp boundary rate and a failure case when a relevant direction remains unobserved.

The finite separation concerns certificate tightness within the same GP family. Direct robust probability remains tighter, and a high-information case favors the information bound. Gaussian conditioning, quadratic concentration and probability sensitivity are established tools; our contribution is their posterior-specific analysis and integration into an explicit coverage contract.

The statistical scope is explicit: finitely many designs, independent latent GP output channels, a squared-norm objective, known positive Gaussian observation noise and a correctly specified fixed hyperparameter in a compact family. Policies may adapt to observations and fit hyperparameters. The results do not cover arbitrary nonlinear networks, learned noise, misspecified priors or an anytime stopping rule.

2 Related Work

Composite optimization. Astudillo & Frazier (2019) model component outputs and use their composition in expected improvement. Non-Gaussian composite uncertainty is therefore established motivation, not our novelty claim. BONSAI (Kudva & Paulson, 2025) studies robust optimization of function networks; its nominal-case Appendix B includes uniformly sub-Gaussian standardized objective marginals (Assumption A5). Our quadratic analysis avoids that premise, but addresses a more restricted objective class and a different certification target.

Posterior-regret certification. Wilson (2024) estimates posterior regret-event probabilities and controls simulation error for probabilistic stopping. Its general simulation formulation is not restricted to Gaussian objective marginals. Wang et al. (2026) develop GP-UCB stopping criteria from high-probability instantaneous-regret bounds. We do not claim to introduce posterior-regret certificates,

simulation-based decisions or risk splitting. Our proposed distinction is a continuous-hyperparameter transfer bound expressed through the coefficients of coupled quadratic contrasts. The terminal result below is a sufficient integration of this bound, rather than a stronger sequential stopping theorem.

Fully Bayesian treatment. Snoek et al. (2012, Section 3.1) integrate acquisition functions over the GP hyperparameter posterior. Such averaging accounts for model uncertainty but targets a hyperposterior average. Our certificate controls the regret probability over a retained family and then bounds joint false certification for each fixed true hyperparameter. These are different contracts; neither implies universal dominance.

Hyperparameter uncertainty and coupling. Capone et al. (2022) develop GP error bounds with unknown hyperparameters. Their prediction-error target differs from our regret-event target, while the need to account for hyperparameter uncertainty strongly overlaps. Neiswanger & Ramdas (2021) use martingales for uncertainty quantification with misspecified GP priors; our model-correct theorem makes no comparable misspecification claim. Pathwise Gaussian conditioning (Wilson et al., 2021) supplies the established posterior construction used here. The concentration argument follows Gaussian quadratic-form methods; related sub-Gaussian quadratic inequalities include Hsu et al. (2011). We give a proof for the signed noncentral form actually needed. Neither pathwise conditioning nor the underlying exponential-moment technique is claimed as new.

Comparison of statistical contracts. Wilson's Sections 3–4 include repeated-test budgeting and changing hyperparameters; Section 5.3 also discusses marginalization (Wilson, 2024); our distinction is not that fitting hyperparameters was absent. Capone et al.'s Theorems 3.5 and 3.7 use box-uniform prediction bounds and hyperposterior region probability (Capone et al., 2022). The contracts differ as follows.

Work	Controlled quantity	Parameter treatment	Time statement
Wilson PRB	Simulated posterior regret event	Current model; fitting allowed	Repeated tests with error budgets
Capone et al.	Whole-function prediction error	Lengthscale box and hyperposterior mass	Stated prediction-band contract
This paper	Quadratic regret event and joint false certificate	Continuous cells; retention covers each fixed truth under its GP law	Fixed terminal observation count

These distinctions identify the proposed addition; they establish no dominance across different probability contracts.

Probability sensitivity. Likelihood-ratio/score gradient estimation is established (Glynn, 1990). For a fixed event E and a smoothly parameterized positive density q_θ, differentiating under an integrable envelope gives $\partial_j P_\theta(E) = \mathbb{E}_\theta[\mathbf{1}_E \partial_j \log q_\theta]$. Indicator nonsmoothness does not itself prevent this probability derivative. Gaussian-product theory also already includes the independent one-zero-mean family used in our specialization (Gaunt & Ye, 2024). Our local contribution is the actual posterior sign and remainder calculation together with the sharp optimized-transfer comparison; probability differentiation and the Taylor support-function identity are not independent novelty claims.

Sensitivity and informative observations. Observable sensitivity bounds involving the product of observable standard deviation and square-root directional Fisher information are established (Dupuis et al., 2015, Theorem 2.13, Eq. (2.36)); their indicator specialization underlies our uniform derivative estimate. Posterior-expectation sensitivity to hyperparameters and its covariance representation are also classical robust-Bayes tools, discussed by Giordano et al. (2018). Our specific analysis derives the observation-noise dependence and positive-boundary coefficient for the actual quadratic-regret posterior. Gotoh et al. (2026) study posterior and likelihood uncertainty sets in Bayesian distributionally robust optimization, including zero posterior sensitivity for a point-mass posterior. Our uncertainty set instead indexes GP prior hyperparameters; the latent posterior event can stabilize across that fixed set without hyperparameter consistency. These distinctions do not imply that data-driven reduction of prior sensitivity is itself new.

Goal-oriented information bounds. The variational information bound of Dupuis et al. (Dupuis et al., 2015, Eq. (2.11), Definition 2.6), applied to an indicator, gives an optimally tilted event-probability bound. Section 7 uses that exact indicator construction and additionally exploits the centered posterior's parameter-independent signs. The finite advantage comes from an explicit event-specific remainder within the restricted GP family; it is not a new entropy inequality or superiority over exact robust probability.

3 CONTRAST-LEVEL EVENT TRANSFER

Fix $N \geq 2$ designs, $m \geq 1$ independent latent output channels, and $d_h = Nm$. For $\theta = (\alpha, \lambda)$ let $a = e^{\alpha}$, $\ell = e^{\lambda}$, and

$$C_\theta[(i,k),(j,l)] = \mathbf{1}_{k=l} a^2 \exp[-\|x_i - x_j\|^2/(2\ell^2)].$$

The loss is $f_h(i) = \|h(i)\|^2/m$; this paper's explicit instance is a quadratic composite, not every smooth nonlinear composition. At a fixed history, let $H : \mathbb{R}^{d_h} \to \mathbb{R}^n$ be the observation map and y the observed vector. Fix a parameter-independent mean m_0, known noise standard deviation $\sigma > 0$, and a recommendation $i_\star$. Let F_θ be the exact Gaussian Fourier feature factor, with $F_\theta F_\theta^* = C_\theta$. Define

$$V_\theta = HC_\theta H^* + \sigma^2 I, \qquad G_\theta = C_\theta H^* V_\theta^{-1},$$
$$\mu_\theta = m_0 + G_\theta(y - Hm_0), \qquad A_\theta = [(I - G_\theta H)F_\theta, -\sigma G_\theta].$$

On one auxiliary probability space assume an isonormal process W over the feature-plus-noise Hilbert space: its finite collections are jointly centered Gaussian, $\mathbb{E}W(u)W(v) = \langle u, v\rangle$, and it is linear in L^2. The coordinatewise sampler $h_\theta = \mu_\theta + A_\theta W$ has the exact selected Gaussian posterior law. Pathwise Gaussian conditioning gives this construction. The coupling compares posterior laws across parameters; the actual latent-data experiment and the final Monte Carlo simulation are specified separately.

For a competitor $x \neq i_\star$, let S_x select the recommendation's m coordinates first and the competitor's m coordinates second. Temporarily suppress θ in the coefficients. Write $J = \mathrm{diag}(I_m, -I_m)/m$, $A_x = S_x A$, $\mu_x = S_x \mu$, and

$$T_x = A_x^* J A_x, \qquad b_x = 2A_x^* J \mu_x, \qquad q_x = \langle \mu_x, J\mu_x\rangle + \mathrm{tr}(T_x).$$

The loss contrast therefore decomposes into its mean, a linear Gaussian term and a centered finite-rank Gaussian quadratic form. For $T = \sum_k \lambda_k e_k \otimes e_k$ in an orthonormal eigenbasis, write $I_2(T) = \sum_k \lambda_k(W(e_k)^2 - 1)$; the linear term $W(b)$ has variance $\|b\|^2$. Differences of these three coefficients, rather than distance between the entire latent laws, control regret changes.

Theorem 1 (Continuous-cell posterior event transfer)**.** *Let D be a nonempty compact convex log-parameter cell, $c \in D$, and $\|\theta - c\|_\infty \leq r$ for all $\theta \in D$. Use the derived, norm-specific constants in Appendix B, or any smaller rigorously valid constants, satisfying for directions $\|u\|_\infty \leq 1$*

$$|D_u q_x(c)| \leq M_{1x}, \quad \|D_u b_x(c)\| \leq L_{1x}, \quad \|D_u T_x(c)\|_{\mathrm{HS}} \leq F_{1x}, \quad \|D_u T_x(c)\|_{\mathrm{op}} \leq O_{1x},$$

and, for the same direction norm, uniform second-directional bounds $M_{2x}, L_{2x}, F_{2x}, O_{2x}$ on D. For each coefficient type $C \in \{M, L, F, O\}$, set $\bar{C}_x = rC_{1x} + r^2 C_{2x}/2$. For $0 < \rho < 1$, define

$$t = \log((N-1)/\rho), \qquad \eta_D = \max_{x \neq i_\star} \left\{ \bar{M}_x + 2\sqrt{(\bar{F}_x^2 + \bar{L}_x^2/2)t} + 2\bar{O}_x t \right\}.$$

Let $R_\theta = f_{h_\theta}(i_\star) - \min_i f_{h_\theta}(i)$ and $p_\theta(s) = \Pr(R_\theta > s)$ under the fixed-history posterior. For every $\theta \in D$,

$$\Pr(R_\theta - R_c > \eta_D) \leq \rho, \qquad p_\theta(\varepsilon) \leq \min\{1, p_c(\varepsilon - \eta_D) + \rho\}.$$

Using $t = \log(2(N-1)/\rho)$ instead gives

$$|p_\theta(\varepsilon) - p_c(\varepsilon)| \leq \min\{1, \Pr(|R_c - \varepsilon| \leq \eta_D) + \rho\}.$$

Proof. Appendix B proves actual posterior C^2 regularity and all four controls. Taylor's theorem on the segment from c to θ gives the four coefficient displacement bounds. Write each contrast

increment as $\Delta q_x + I_2(\Delta T_x) + W(\Delta b_x)$, where Δ denotes its value at θ minus its value at c. Apply Lemma 6 to $\Delta T_x, \Delta b_x$ on the finite span of the selected rows from both endpoints, using the same W. This gives variance proxy $\|\Delta T_x\|_{\mathrm{HS}}^2 + \|\Delta b_x\|^2/2$ and operator coefficient $\|\Delta T_x\|_{\mathrm{op}}$; no single-endpoint rank bound is assumed. The finite-rank quadratic Gaussian concentration step uses the HS and operator bounds separately; no sub-Gaussian tail premise for the loss is needed. For each competitor it bounds the contrast displacement with failure $\rho/(N-1)$. A finite union bound and the Lipschitz property of the zero-augmented maximum give the regret displacement. The event inclusion $\{R_\theta > \varepsilon\} \subseteq \{R_c > \varepsilon - \eta_D\} \cup \{R_\theta - R_c > \eta_D\}$ proves transfer. The two-sided version uses two tails and the threshold boundary event. The posterior-law bridge identifies these probabilities with the actual fixed-history posterior probabilities. □

There is no probability union over the continuous parameter cell: this is a uniform deterministic bound on the individual probabilities for every parameter. It is not a simultaneous pathwise bound over all θ. If $r = 0$, the posteriors coincide and the discrepancy is exactly zero. If $N = 1$, regret is identically zero and the logarithmic formula is unnecessary.

What the integration establishes. Appendix C derives a fixed-terminal false-certificate bound $\min\{1, \alpha + \delta_{\mathrm{ret}} + \beta\}$ for a continuous retained family, with a fully specified adaptive observation and simulation contract. Section 7 establishes a finite threshold separation from a strong goal-oriented information bound using an event-specific remainder. Neither result asserts a BO optimization rate or monotonic narrowing under arbitrary observations.

4 Certification after adaptive observations

To use the transfer bound after adaptive observations, we specify the data-generating experiment and the final simulation. A fixed truth $\theta_* \in \Theta$ generates one latent vector $h \sim N(m_0, C_{\theta_*})$. At each of a fixed number T of steps, a measurable policy chooses H_t from an independent seed and past observations, then observes $Y_t = H_t h + \sigma \xi_t$. The noise vectors are independent standard Gaussian, and the policy rule is the same under every candidate parameter. A fitted hyperparameter may guide this policy.

For the final history Z, let $L_\theta(Z)$ be its normalized sequential likelihood. Fix a probability measure π independently of the observations, let $M = \int L_\theta \, d\pi$, and retain

$$\mathcal{R}(Z) = \{\theta : L_\theta(Z) \geq \delta_{\mathrm{ret}} M(Z)\}.$$

Cover this set by a finite measurable ordered family of compact convex cells with rigorous coefficient envelopes. Choose a recommendation, tolerance, cell anchors, transfer budgets $\rho_j < \alpha$ and sample counts K_j before the final simulation. For each anchor, estimate $q_j = p_{c_j}(\varepsilon - \eta_j \mid Z)$ from K_j conditionally independent exact posterior draws. Let V be the complete final simulation seed, conditionally independent of h given Z; all simulated draws are measurable functions of (Z, V). Independence is required within each cell, but shared random seeds across cells are allowed. Put

$$U_j = \min\{1, \widehat{q}_j + \sqrt{\log(1/\beta)/(2K_j)}\}.$$

Issue a certificate only if $U_j + \rho_j \leq \alpha$ for every cell. Any history-conditional one-sided binomial bound with error at most β may replace U_j; Appendix C gives an exact alternative.

Theorem 2 (Terminal false certification)**.** *Under this contract and the measurable, valid-cover conditions of Appendix C, for every fixed $\theta_* \in \Theta$,*

$$\Pr_{\theta_*}\{\textit{certificate issued and } f_h(i_\star) - \min_i f_h(i) > \varepsilon\} \leq \min\{1, \alpha + \delta_{\mathrm{ret}} + \beta\}.$$

The proof uses the first covering cell containing the fixed truth as a history-measurable witness. Only that cell's Monte Carlo error is needed, so no cell-count factor multiplies β. This does not provide simultaneous confidence bounds for all cells. The conclusion is a joint false-certificate bound, not a conditional error bound given certification. A certified outer cover of the entire parameter rectangle is always available; data-dependent removal requires rigorous likelihood bounds. We prove neither efficient cover construction nor retention contraction.

Using a sharper valid cell bound. Only a valid upper bound on the truth-containing cell is needed in the proof. A deterministic, history-measurable bound $b_j \geq \sup_{\theta \in D_j} p_\theta(\varepsilon \mid Z)$ can therefore replace the shifted-threshold bound. Certifying when every $b_j \leq \alpha$ gives joint false-certificate probability at most $\alpha + \delta_{\mathrm{ret}}$. A hybrid taking the minimum of b_j and the original Monte Carlo bound preserves the original $\alpha + \delta_{\mathrm{ret}} + \beta$ guarantee. Appendix C.4 proves both statements. The specialized event bound below is eligible only when its posterior assumptions and certified quantities hold; otherwise the general bound remains available.

5 Intrinsic ambiguity and certificate conservatism

Validity alone does not tell us how much conservatism the certificate adds. To measure that gap, consider two distinct designs $x_1 = -d/2$, $x_2 = d/2$, one zero-mean RBF output and one noisy normalized-sum observation $Y = (h_1 + h_2)/\sqrt{2} + \sigma\xi = y$. Fix $d, \sigma > 0$, the observed $y \in \mathbb{R}$, a finite center $c = (\log a, \log \ell)$ and a physical tolerance $\varepsilon > 0$. Both log hyperparameters vary in $D_r = c + [-r, r]^2$. Let

$$p_\theta = \Pr_\theta(R > \varepsilon \mid y), \qquad \mathcal{A}_r^+ = \sup_{\theta \in D_r} p_\theta - p_c.$$

Theorem 3 (Sharpness in the structured posterior)**.** *For sufficiently small $r > 0$, with $g = \nabla_\theta p_c$, the true ambiguity satisfies*

$$\mathcal{A}_r^+ = r\|g\|_1 + O(r^2), \qquad \|g\|_1 > 0.$$

The same expansion holds for $\sup_{\theta \in D_r} |p_\theta - p_c|$. If H_p bounds the directional Hessian on D_{r_0}, then for $0 < r \leq r_0$,

$$B_{\mathrm{event}} = \min\{1, p_c + r\|g\|_1 + H_p r^2/2\}$$

upper bounds the whole cell with excess at most $H_p r^2$ above its exact worst-case probability.

Let B_{conc} be the preceding one-reference signed-quadratic transfer formula, using the same common coupling, exact continuum coefficient envelopes and its optimal failure budget. With $f_c(\varepsilon) > 0$ the contrast density and $o_c = \sup_{\|u\|_\infty \leq 1} \|DT_c[u]\|_{\mathrm{op}} > 0$,

$$B_{\mathrm{conc}} - p_c \sim 2 f_c(\varepsilon) o_c r \log(1/r).$$

Theorem 12 gives the exact posterior integral, derivatives and remainder; Theorem 14 defines the comparator and proves the matching leading constant. The result separates the unavoidable linear ambiguity from extra conservatism in this formula. It concerns ambiguity penalties: both total probabilities converge to the same positive p_c. At a fixed smooth nonsingular Gaussian center with $p_c \in (0, 1)$, the optimized goal-oriented bound also has an $O(r)$ local penalty: Gaussian KL is $O(r^2)$ and its indicator-bound expansion is of square-root order in KL (Dupuis et al., 2015, Theorem 2.11). The logarithmic separation concerns the specified concentration formula; the later information-bound comparison concerns finite event-specific tightness.

The history and tolerance remain fixed as r decreases. Smaller cells may result from refining a numerical cover; this does not establish that the statistically retained family contracts. In fact, this experiment's likelihood depends only on $a^2(1 + k)$, and distinct parameters with that same value can have different regret probabilities (Appendix E.4.2). The sharp theorem is therefore a local analysis within the general uncertainty framework. It does not assert a learning rate, superiority over direct probability optimization, or an improved optimization trajectory.

6 When informative observations reduce ambiguity

The preceding result varies the cell radius while holding the data fixed. We now ask whether informative observations reduce disagreement when the hyperparameter family itself remains fixed. Retain the two-design RBF model, and now observe both $X_+ = (h_1 + h_2)/\sqrt{2}$ and $X_- = (h_1 - h_2)/\sqrt{2}$ with independent noise variance $t > 0$. Repeating measurements n times at each original design with known noise variance σ^2 gives exactly this model through transformed sample means, with $t = \sigma^2/n$. Conditional on their observed values $y_\pm$, the two coordinates are independent with

$$v_\pm = a^2(1 \pm k), \quad w_\pm = \frac{v_\pm}{v_\pm + t}, \qquad m_\pm = \frac{v_\pm y_\pm}{v_\pm + t}, \qquad V_\pm = \frac{t v_\pm}{v_\pm + t}.$$

Let $p_\theta(t, y) = \Pr_\theta(2X_+X_- > \varepsilon \mid y)$ for a fixed recommendation of design 1 and tolerance $\varepsilon > 0$. For a fixed compact convex log-parameter family K, define $\operatorname{osc}_K p = \sup_K p_\theta - \inf_K p_\theta$.

Theorem 4 (Observation-dependent ambiguity)**.** *Assume $v_\pm \geq v_* > 0$ on K. Uniformly over bounded observations, $\operatorname{osc}_K p = O(\sqrt{t})$. If also $|2y_+y_- - \varepsilon| \geq \gamma > 0$, then for sufficiently small t,*

$$\operatorname{osc}_K p \leq Ce^{-c/t},$$

with positive constants uniform over this bounded-data, fixed-margin class. This bounds disagreement, including on the unsafe side of the threshold.

The square-root rate is attained at a positive boundary. For fixed $y_+, y_- > 0$ and $\varepsilon = 2y_+y_-$, put $R_y = (y_+^2 + y_-^2)^{1/2}$. Then

$$\operatorname{osc}_K p = \frac{\varphi(0)y_+y_-}{R_y} \operatorname{osc}_K(v_+^{-1} + v_-^{-1})\sqrt{t} + O(t).$$

The leading coefficient is positive when K contains distinct amplitudes at fixed length scale. Under actual repeated observations at both designs, almost surely $\operatorname{osc}_K p(t_n, y_n)$ eventually decays exponentially in n, with realization-dependent constants. No contraction of K is assumed.

Appendix E gives the exact event derivatives, an explicit Hessian bound $H_t = O(\sqrt{t})$ and a local certificate with at most $H_t r^2$ excess over the exact cell supremum. It also supplies the uniform differentiated boundary remainder. These bounds can enter the deterministic/hybrid terminal interface of Appendix C.4. The almost-sure history result does not extend that interface to first-passing stopping.

The observation condition is substantive. With only the sum measured, equal-v_+ parameters can retain different limiting probabilities $\overline{\Phi}(\varepsilon/(2|y_+|\sqrt{v_-}))$ for $y_+ \neq 0$, where $\overline{\Phi} = 1 - \Phi$ is the standard normal upper tail (the limit is zero if $y_+ = 0$). Conversely, a direct Gaussian margin bound already attains the exponential rate when both coordinates are measured. The contribution here is the quantitative observation-regime characterization and its boundary and failure cases, rather than a universal advantage over direct event certificates.

7 A FINITE ADVANTAGE OVER A STRONG EVENT BOUND

A smaller asymptotic penalty need not change a decision at a finite cell radius. We now compare the same posterior history, parameter box, recommendation and regret threshold against both direct robust probability and a goal-oriented information bound with optimal exponential tilt (Dupuis et al., 2015).

Use the two-coordinate posterior of Section 6 with $y_+ = y_- = 0$. Set $s_\theta = 2t\sqrt{w_+w_-}$ and $h_\theta = \varepsilon/s_\theta$. The exact product-Gaussian event probability is

$$p_\theta = \frac{1}{\pi}\int_0^{\pi/2} \exp(-h_\theta \sec\phi)\, d\phi.$$

This representation yields a substantially smaller remainder than bounding absolute posterior-density derivatives.

Proposition 5 (Event-specific remainder)**.** *On $D_r = c + [-r, r]^2$, suppose L_t and M_t uniformly bound $|D_u \log s_\theta|$ and $|D_u^2 \log s_\theta|$ for $\|u\|_\infty \leq 1$. Then*

$$H_t = \tfrac{1}{2}\{c_\sharp L_t^2 + e^{-1}M_t\}, \qquad c_\sharp = (2+\sqrt{5})e^{-(3+\sqrt{5})/2}$$

bounds $|D_u^2 p_\theta|$. Thus $B_{\mathrm{loc}} = \min\{1, p_c + r\|\nabla p_c\|_1 + H_t r^2/2\}$ upper bounds $B_* := \sup_{D_r} p_\theta$, with excess at most $H_t r^2$. For zero-valued sample means and any fixed compact positive-variance box, explicit RBF controls give $L_t, M_t, H_t = O(t)$ and certified ambiguity penalty $O(t)$.

Proof. Use polar coordinates for two independent standard normals. Their positive-product tail is $\pi^{-1}\int_0^{\pi/2} e^{-h\sec\phi}d\phi$. For $z = h\sec\phi$, $L_u = D_u \log s$ and $M_u = D_u^2 \log s$, differentiation gives

$$D_u p = \frac{1}{\pi}\int_0^{\pi/2} e^{-z} z L_u\, d\phi, \qquad D_u^2 p = \frac{1}{\pi}\int_0^{\pi/2} e^{-z}\{(z^2 - z)L_u^2 + zM_u\}\, d\phi.$$

n	Direct B_*	Local B_{loc}	B_{sign} enclosure	Local penalty
1	.08320605	.08325680	[.08472219,.08483404]	.00197431
2	.04380502	.04384427	[.04478719,.04482972]	.00107887
4	.01335810	.01338541	[.01383491,.01384632]	.00035071
8	.00143216	.00144921	[.00156326,.00156506]	.00005203
16	.00001998	.00002986	[.0000314401,.0000315538]	.00001037
32	4.953×10^{-9}	5.406×10^{-6}	$[2.8715, 2.9063]\times10^{-7}$	5.401×10^{-6}

Table 1: Matched certificates on the fixed hyperparameter box. The first five rows use the planned observation counts; $n = 32$ is the disclosed secondary check. Local penalty means $B_{\mathrm{loc}} - p_c = r\|\nabla p_c\|_1 + H_t r^2/2$ (the cap is inactive). It narrows with observations; direct control wins throughout and the information bound wins at $n = 32$. Point values are rounded; interval endpoints are outward-rounded.

Compactness and positive variances justify differentiation by domination. The stationary points of $(z^2 - z)e^{-z}$ are $(3 \pm \sqrt{5})/2$. Comparing their absolute values gives $\sup_{z\geq 0} |(z^2 - z)e^{-z}| = c_\sharp$; also $ze^{-z} \leq e^{-1}$. The normalized integration length $1/2$ gives H_t. Taylor's theorem gives validity and the $H_t r^2$ excess bound, by the same maximizing sign direction used in Corollary 16.

For explicit constants put $q_i = \partial_{\log \ell} \log v_i$, $q_i' = \partial^2_{\log \ell} \log v_i$ and $l_i = 2u_a + q_i u_\ell$. Their explicit RBF formulas are given in Appendix E. Direct differentiation yields

$$L_u = \tfrac{1}{2}\sum_i (1 - w_i) l_i, \qquad M_u = \tfrac{1}{2}\sum_i (1 - w_i)(q_i' u_\ell^2 - w_i l_i^2).$$

Writing $Q_i = \sup_D |q_i|$, $Q_i' = \sup_D |q_i'|$, $b_i = \sup_D (1 - w_i)$ and $W_i = \sup_D w_i$, valid choices are

$$L_t = \tfrac{1}{2}\sum_i b_i(2 + Q_i), \quad M_t = \tfrac{1}{2}\sum_i b_i\{Q_i' + W_i(2 + Q_i)^2\}.$$

If $v_i \geq v_* > 0$, then $b_i \leq t/v_*$, whereas Q_i, Q_i', W_i are bounded. This proves $L_t, M_t, H_t = O(t)$. Also $\|\nabla p_c\|_1 \leq L_t/(2e)$, so the fixed-radius penalty is $O(t)$. This is the Taylor penalty rate, not a sharp intrinsic ambiguity rate: zero-valued sample means also satisfy the safe-margin condition of Theorem 17, which gives exponential narrowing. No arbitrary-history monotonicity is asserted.

□

This specialization preserves the certificate structure and does not require sub-Gaussian product tails.

Let $K_r = \sup_{D_r} \mathrm{KL}(P_\theta \| P_c)$, and define the exact optimally tilted indicator bound

$$q_+(p, K) = \inf_{\lambda > 0} \frac{\log(1 - p + pe^\lambda) + K}{\lambda}.$$

It equals binary-KL inversion, saturated at one. The standard comparator is $B_{\mathrm{GI}} = q_+(p_c, K_r)$. We strengthen it by exploiting the common independent fair sign agreement: $B_{\mathrm{sign}} = \frac{1}{2} q_+(2p_c, K_r) \leq B_{\mathrm{GI}}$. Both comparators use the same Gaussian posterior family, exact indicator generating function and optimal tilt. Our certificate uses event-specific family structure beyond its KL summary; this is not a comparison of identical reduced summaries.

Fix $d = t = \varepsilon = 1$, $c = (0, 0)$, $r = .01$ and risk threshold $\alpha = .084$. Validated numerical enclosures establish

$$B_* < .083206055 < B_{\mathrm{loc}} < .083256800 < \alpha < .084722196 < B_{\mathrm{sign}}.$$

Thus our certificate passes while even the strengthened information bound remains inconclusive. The previous conservative Hessian certificate exceeds .09915 and fails. The event-specific bound gives $H_1 < .964$, versus the previous $H_1 > 318$. In this comparison, the sharper remainder turns an inconclusive certificate into one that passes.

Direct robust control is stronger here: monotonicity reduces B_* to one corner probability evaluation. We establish no computational saving over it. Table 1 shows the fixed-family narrowing sequence and the high-information reversal. Appendix F supplies the validation details. The setting was informed by earlier diagnostics, with the risk threshold fixed before this calculation. These stipulated histories with zero-valued sample means are a conditional construction, not evidence of typical adaptive BO performance.

Random-history check. A separate, predeclared check uses 100 independent two-design RBF trajectories, a fixed hyperparameter box, and $n = 1, 2, 4, 8, 16, 32$ observations per design (Appendix G). With the general nonzero-mean remainder, the median ambiguity penalty falls from 0.10234 to 0.00266. Our bound passes in $0/100$ histories at $n = 1$ and $74/100$ at $n = 32$; the ordinary goal-oriented comparator passes in $14/100$ and $74/100$, respectively, and never loses a paired decision to our bound. Thus the check supports small-scale evaluation and narrowing, but not comparative practical superiority. These are floating-point diagnostics, separate from the validated centered comparison.

8 Discussion and limitations

Our results connect a general contrast-level transfer bound to a fixed-terminal certification rule, then use structured posteriors to measure its conservatism. The finite comparison shows that an event-specific remainder can change a risk decision relative to a strong information bound. Informative observations can reduce disagreement even when the hyperparameter family does not contract.

The finite advantage uses a narrow two-parameter box and zero-valued sample means. Direct robust probability is tighter and simpler there; the information bound also wins at high information. On random histories, nonzero posterior means enter the general score and Hessian controls. Their uniform remainder is conservative: it accounts for a median 96.6% of the local penalty at $n = 1$, and the information comparator certifies more often. This diagnoses where the bound loses tightness, not an inevitable disadvantage of every event-specific certificate. The observation theorem allows nonzero histories but retains two designs and equal precision. Its eventual rate has realization-dependent constants, not a uniform sample-complexity guarantee.

Scaling and model assumptions. For fixed coefficient envelopes, Theorem 1 adds $\sqrt{\log((N-1)/\rho)}$ and $\log((N-1)/\rho)$ terms. The envelopes, posterior algebra and $N - 1$ contrast calculations can also grow with problem size; logarithmic factors alone do not establish scalability. The explicit RBF construction shares amplitude and length scale across channels. Separate channel parameters would require new derivative envelopes and a higher-dimensional cover; we do not instantiate that extension. Correct GP specification, known noise and valid retained-family coverage remain essential. The theorem does not cover systematic kernel misspecification or anytime stopping, and gives no efficient algorithm for constructing the cover.

Beyond quadratic losses. A possible extension is a quadratic approximation with a controlled remainder on a high-probability latent region. The remainder would shift the regret threshold, and the region's failure probability would enter the error budget. Smoothness alone supplies neither control, so this extension requires a separate theorem. Our written proofs and numerical enclosures also remain distinct from the partial Lean verification documented in the supplement.

Application pathway: hyperparameter tuning. Application configurations can form the finite design set, while θ indexes surrogate hyperparameters. For squared-norm losses satisfying the model assumptions, the certificate assesses a recommendation at a predetermined evaluation budget. It could complement search tools such as Optuna (Akiba et al., 2019); no integration, tuning advantage or large-scale benchmark is demonstrated here.

AI USE STATEMENT

Generative AI tools were used solely to polish the manuscript's language, improving grammar, clarity, and readability. The authors reviewed and approved all revisions and take full responsibility for the manuscript's content.

REPRODUCIBILITY STATEMENT

The main text states the model and statistical assumptions. The appendix provides the proofs, explicit posterior coefficient controls, comparator definitions, observation conditions and numerical enclosure procedure. The accompanying code includes pinned requirements, the exact finite-comparison settings, full ball-valued results and executable checks for both the reported advantage and retained baseline wins. The random-history diagnostic includes its frozen protocol, seed, simulated observations, all paired outcomes and plotting code; its floating-point quadrature is separate from ball validation. The motivation figure is generated from exact Gaussian identities. No empirical BO performance claim depends on an unprovided training run.

REFERENCES

Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. Optuna: A next-generation hyperparameter optimization framework. arXiv:1907.10902, 2019. URL https://arxiv.org/abs/1907.10902.

Raul Astudillo and Peter Frazier. Bayesian optimization of composite functions. In *ICML*, volume 97 of *PMLR*, pp. 354–363, 2019. URL https://proceedings.mlr.press/v97/astudillo19a.html.

Alexandre Capone, Armin Lederer, and Sandra Hirche. Gaussian process uniform error bounds with unknown hyperparameters for safety-critical applications. In *ICML*, volume 162 of *PMLR*, pp. 2609–2624, 2022. URL https://proceedings.mlr.press/v162/capone22a.html.

Paul Dupuis, Markos A. Katsoulakis, Yannis Pantazis, and Petr Plechac. Path-space information bounds for uncertainty quantification and sensitivity analysis of stochastic dynamics. arXiv:1503.05136, version 2, 2015. URL https://arxiv.org/pdf/1503.05136v2.

Robert E. Gaunt and Zixin Ye. Asymptotic approximations for the distribution of the product of correlated normal random variables. arXiv:2309.07734, version 2, 2024. URL https://arxiv.org/html/2309.07734v2. First submitted in 2023.

Ryan Giordano, Tamara Broderick, and Michael I. Jordan. Covariances, robustness, and variational bayes. arXiv:1709.02536v3, 2018. URL https://arxiv.org/pdf/1709.02536v3.

Peter W. Glynn. Likelihood ratio gradient estimation for stochastic systems. *Communications of the ACM*, 33:75–84, 1990. URL https://web.stanford.edu/~glynn/papers/1990/G90a.html.

Jun-ya Gotoh, Michael Jong Kim, and Andrew E. B. Lim. Posterior and likelihood sensitivity in Bayesian distributionally robust optimization. arXiv:2605.31306, version 1, 2026. URL https://arxiv.org/html/2605.31306v1.

Daniel Hsu, Sham M. Kakade, and Tong Zhang. A tail inequality for quadratic forms of subgaussian random vectors. arXiv:1110.2842, 2011. URL https://arxiv.org/abs/1110.2842.

Fredrik Johansson. Numerical integration in arbitrary-precision ball arithmetic. arXiv:1802.07942, 2018. URL https://arxiv.org/abs/1802.07942.

Akshay Kudva and Joel A. Paulson. BONSAI: Structure-exploiting robust Bayesian optimization for networked black-box systems under uncertainty. arXiv:2510.03893, version 1, 2025. URL https://arxiv.org/html/2510.03893v1.

Willie Neiswanger and Aaditya Ramdas. Uncertainty quantification using martingales for misspecified Gaussian processes. In *Proceedings of the 32nd International Conference on Algorithmic Learning Theory*, volume 132 of *Proceedings of Machine Learning Research*, pp. 963–982, 2021. URL https://proceedings.mlr.press/v132/neiswanger21a.html.

Jasper Snoek, Hugo Larochelle, and Ryan P. Adams. Practical bayesian optimization of machine learning algorithms. arXiv:1206.2944, version 2, 2012. URL https://arxiv.org/pdf/1206.2944v2.

Haowei Wang, Jingyi Wang, and Qiyu Wei. Regret-based (ϵ, δ)-optimal stopping criteria for Bayesian optimization. arXiv:2605.22561, version 1, 2026. URL https://arxiv.org/html/2605.22561v1.

James T. Wilson. Stopping Bayesian optimization with probabilistic regret bounds. arXiv:2402.16811, version 2, 2024. URL https://arxiv.org/html/2402.16811v2.

James T. Wilson, Viacheslav Borovitskiy, Alexander Terenin, Peter Mostowsky, and Marc Peter Deisenroth. Pathwise conditioning of Gaussian processes. *Journal of Machine Learning Research*, 22(105):1–47, 2021. URL https://jmlr.org/papers/v22/20-1260.html.

A SIGNED QUADRATIC CONCENTRATION AND THE COMMON FACTOR

Lemma 6 (Finite-rank signed Gaussian tail)**.** *Let W be isonormal, T finite-rank self-adjoint and b a Hilbert-space vector. Define $I_2(T) = \sum_i \lambda_i(W(e_i)^2 - 1)$ for its nonzero eigenpairs (λ_i, e_i). For $Q = I_2(T) + W(b)$, set $v = \|T\|_{\mathrm{HS}}^2 + \|b\|^2/2$ and $o = \|T\|_{\mathrm{op}}$. For $t > 0$,*

$$\Pr\{Q > 2\sqrt{vt} + 2ot\} \le e^{-t}.$$

Proof. Choose an orthonormal basis of the finite-dimensional span of the range of T and b, diagonalizing T. The corresponding isonormal coordinates Z_i are independent standard normals and $Q = \sum_i [\lambda_i(Z_i^2 - 1) + b_i Z_i]$. For $0 < s < 1/(2o)$, Gaussian integration gives

$$\log \mathbb{E}e^{sQ} = \sum_i \left[-s\lambda_i - \frac{1}{2}\log(1 - 2s\lambda_i) + \frac{s^2 b_i^2}{2(1 - 2s\lambda_i)} \right] \le \frac{s^2 v}{1 - 2so}.$$

For the centered quadratic term use the power series of $-\log(1-z) - z$ and absolute values, bounding its terms of degree at least two by $z^2/[2(1-|z|)]$. Denominators in the linear term are at least $1 - 2so$. If $v > 0$, exponential Markov with $s = \sqrt{t}/(\sqrt{v} + 2o\sqrt{t})$ proves the result by substitution. If $v = 0$, then $T = b = 0$ and the strict upper-tail event is empty. The case $o = 0$ follows directly from the Gaussian mgf (or the same formula with no upper restriction on s). Applying the result to $-Q$ supplies the lower tail. □

An exact real Fourier factor uses

$$\phi_{\theta,i}(z) = a\big(\cos(z^\top x_i/\ell), \sin(z^\top x_i/\ell)\big), \qquad z \sim N(0, I).$$

Its Hilbert space is the direct sum of two real Gaussian L^2 spaces per independent output channel. The Gaussian characteristic function gives $\langle \phi_{\theta,i}, \phi_{\theta,j} \rangle = a^2 \exp(-\|x_i - x_j\|^2/(2\ell^2))$. Thus $F_\theta F_\theta^* = C_\theta$ exactly; Appendix B supplies regularity and bounds. Isonormal existence on this separable Hilbert space is a standard dependency, not a feature approximation.

For a fixed history, $A = [(I - GH)F, -\sigma G]$ satisfies

$$AA^* = (I - GH)C(I - GH)^* + \sigma^2 GG^* = C - CH^*(HCH^* + \sigma^2 I)^{-1}HC.$$

Together with the displayed posterior mean and Gaussianity this proves the exact posterior law, including singular latent covariances. Shared rows are essential for coefficient differences; equality of marginal covariance matrices alone does not identify a cross-parameter coupling.

B EXPLICIT RBF POSTERIOR COEFFICIENT CONTROLS

Output spaces are Euclidean; derivatives are along fixed parameter directions $\|u\|_\infty \le 1$. HS and operator norms are kept distinct.

B.1 PRIOR AND POSTERIOR DERIVATIVES

Let $a_+ = \max_D e^\alpha$, $\ell_- = \min_D e^\lambda$, $z_i = \|x_i\|/\ell_-$. The exact-factor bounds derived below are

$$p_0 = a_+\sqrt{Nm}, \quad p_1 = a_+\sqrt{m\sum_i (1 + z_i^2)}, \quad p_2 = a_+\sqrt{m\sum_i (1 + 7z_i^2 + 3z_i^4)}.$$

They bound F, F', F'' in HS norm. To see the constants directly, identify a cosine/sine row with ae^{it}, where $t = z^\top x_i/\ell$ and $s_i = \|x_i\|/\ell$. For a log-parameter direction (A, L) with $|A|, |L| \le 1$, its first and second directional derivatives are

$$ae^{it}(A - iLt), \qquad ae^{it}[A^2 - L^2t^2 + i(L^2 - 2AL)t].$$

Their squared L^2 norms are respectively

$$a^2(A^2 + L^2 s_i^2), \qquad a^2[A^4 + (2A^2L^2 - 4AL^3 + L^4)s_i^2 + 3L^4 s_i^4],$$

using the Gaussian second and fourth moments. Bounding the three latter coefficients by $1, 7, 3$ and summing the finite rows proves p_0, p_1, p_2. The real two-component row norm equals the complex norm used only as calculation notation. These constants depend on the chosen Fourier factor and fixed coordinate origin; they are not optimal over all couplings. Let $h = \|H\|_{\mathrm{op}}$, $z_y = \|y - Hm_0\|$, and define

$$\begin{aligned} c_0 &= p_0^2, & c_1 &= 2p_0p_1, & c_2 &= 2(p_1^2 + p_0p_2), \\ b_0 &= \sigma^{-2}, & b_1 &= b_0^2h^2c_1, & b_2 &= 2b_0^3h^4c_1^2 + b_0^2h^2c_2, \\ g_0 &= hc_0b_0, & g_1 &= h(c_1b_0 + c_0b_1), & g_2 &= h(c_2b_0 + 2c_1b_1 + c_0b_2). \end{aligned}$$

Lemma 7 (Actual posterior derivative bounds)**.** *The maps μ, A are C^2 on the entire log-parameter plane, with A in HS norm. The constants c_j, b_j, g_j bound $C^{(j)}, B^{(j)}, G^{(j)}$ in Euclidean operator norm, where $B = V^{-1}$ and $j = 0, 1, 2$.*

Proof. Since $V \succeq \sigma^2 I$, the spectral theorem gives $\|B\|_{\mathrm{op}} \le \sigma^{-2}$, including singular C. Differentiating the actual identities gives

$$\begin{aligned} C' &= F'F^* + F(F')^*, & C'' &= F''F^* + 2F'(F')^* + F(F'')^*, \\ V' &= HC'H^*, & V'' &= HC''H^*, \\ B' &= -BV'B, & B'' &= 2BV'BV'B - BV''B, \\ G' &= C'H^*B + CH^*B', & G'' &= C''H^*B + 2C'H^*B' + CH^*B''. \end{aligned}$$

Product norm inequalities give the constants. Smooth inversion of V and the HS-C^2 factor give the regularity, without inverting latent or posterior covariances. □

B.2 Selected-output bounds

Choose uniform nonnegative bounds p_{jx} for $\|(S_xF)^{(j)}\|_{\mathrm{HS}}$ and g_{jx} for $\|S_xG^{(j)}\|_{\mathrm{op}}$, $j = 0, 1, 2$. Always valid choices are $p_{jx} = p_j$, $g_{jx} = g_j$. Sharper p_{jx} follow by restricting the sums to the two selected designs, and sharper gain bounds may be supplied independently. Let $q_x^{\mathrm{rank}} = \min(2m, n)$ and $s_x = \sigma\sqrt{q_x^{\mathrm{rank}}}$. Define

$$\begin{aligned} a_{0x} &= p_{0x} + hg_{0x}p_0 + s_xg_{0x}, \\ a_{1x} &= p_{1x} + h(g_{1x}p_0 + g_{0x}p_1) + s_xg_{1x}, \\ a_{2x} &= p_{2x} + h(g_{2x}p_0 + 2g_{1x}p_1 + g_{0x}p_2) + s_xg_{2x}, \\ v_{0x} &= \|S_xm_0\| + g_{0x}z_y, & v_{1x} &= g_{1x}z_y, \qquad v_{2x} = g_{2x}z_y. \end{aligned}$$

Then a_{jx} bound $\|A_x^{(j)}\|_{\mathrm{HS}}$ and v_{jx} bound $\|\mu_x^{(j)}\|$. Also choose e_{jx} bounding $\|A_x^{(j)}\|_{\mathrm{op}}$; $e_{jx} = a_{jx}$ is always valid. Independently established smaller operator bounds can improve the dominant large-deviation term.

Proof. Differentiate $A_x = [S_xF - S_xGHF, -\sigma S_xG]$:

$$\begin{aligned} (S_xF - S_xGHF)' &= S_xF' - S_xG'HF - S_xGHF', \\ (S_xF - S_xGHF)'' &= S_xF'' - S_xG''HF - 2S_xG'HF' - S_xGHF''. \end{aligned}$$

Bound each product by operator times HS norm. The finite noise block satisfies $\|S_xG^{(j)}\|_{\mathrm{HS}} \le \sqrt{\min(2m, n)}g_{jx}$. The HS norm of a concatenation is at most the sum of the two block norms. The mean identities yield the v bounds. The common input space and coupling are unchanged. □

B.3 Four norm-specific coefficient controls

In the following formulas suppress x and put $j = 1/m$:

$$\begin{aligned} M_1 &= 2j(v_0v_1 + a_0a_1), & M_2 &= 2j(v_1^2 + v_0v_2 + a_1^2 + a_0a_2), \\ L_1 &= 2j(e_1v_0 + e_0v_1), & L_2 &= 2j(e_2v_0 + 2e_1v_1 + e_0v_2), \\ F_1 &= j(a_1e_0 + a_0e_1), & F_2 &= j(a_2e_0 + 2a_1e_1 + a_0e_2), \\ O_1 &= 2je_0e_1, & O_2 &= 2j(e_0e_2 + e_1^2). \end{aligned}$$

Lemma 8 (Derived coefficient derivative bounds)**.** *These constants satisfy the four first- and second-derivative controls in Theorem 1, with the first bounds valid throughout D.*

Proof. The actual coefficient derivatives are

$$\begin{aligned} T' &= (A')^* JA + A^* JA', \\ T'' &= (A'')^* JA + 2(A')^* JA' + A^* JA'', \\ b' &= 2[(A')^* J\mu + A^* J\mu'], \\ b'' &= 2[(A'')^* J\mu + 2(A')^* J\mu' + A^* J\mu'']. \end{aligned}$$

Use $\|J\|_{\mathrm{op}} = j$ and $\|U^* JV\|_{\mathrm{HS}} \le j\|U\|_{\mathrm{HS}}\|V\|_{\mathrm{op}}$ (or its reversed allocation) to obtain F_1, F_2. Operator submultiplicativity gives O_1, O_2 directly. The vector inequalities give L_1, L_2. For q, differentiation through its finite row sum gives

$$\begin{aligned} q' &= 2\langle \mu, J\mu' \rangle + 2\,\mathrm{tr}(A^* JA'), \\ q'' &= 2\langle \mu', J\mu' \rangle + 2\langle \mu, J\mu'' \rangle + 2\,\mathrm{tr}((A')^* JA') + 2\,\mathrm{tr}(A^* JA''). \end{aligned}$$

The trace product bound $|\,\mathrm{tr}(U^* JV)| \le j\|U\|_{\mathrm{HS}}\|V\|_{\mathrm{HS}}$ proves M_1, M_2. This differentiates a finite sum of scalar row inner products, not an unbounded trace functional on the ambient operator space. The row tensor representation also makes T C^2 in HS norm. □

Retaining cancellation at the center. The coarse first constants can be replaced by the exact quantities

$$C_{1x}^{\mathrm{exact}} = \max_{u \in \{-1,1\}^2} \|D_u g_x(c)\|_{\mathcal{Y}}, \quad (g_x, \mathcal{Y}) \in \{(q_x, \mathbb{R}), (b_x, \mathcal{K}), (T_x, \mathrm{HS}), (T_x, \mathrm{op})\}.$$

Linearity of the derivative and convexity of its norm make the vertex maximum exact. Finite row spans permit Gram-matrix evaluation, but numerical upper bounds must be certified. There is no analogous unrestricted corner rule for $D^2 g[u, u]$: keep a uniform proven second bound or derive a special structural bound.

C End-to-end probability statement and adaptive-law assumptions

C.1 Complete statistical contract

Let $\Theta \subset \mathbb{R}^2$ be a compact log-parameter rectangle and fix any $\theta_* \in \Theta$. Under $\mathbb{P}_{\theta_*}$ draw $h \sim N(m_0, C_{\theta_*})$ once. Independently draw a policy seed S and independent standard normal noise vectors $\xi_1, \ldots, \xi_T$, where $T < \infty$ and their dimension k are fixed. At step t, choose a finite real matrix $H_t = H_t(S, Y_{<t})$ by a measurable rule and observe $Y_t = H_t h + \sigma \xi_t$. The policy rule is the same under every candidate parameter. It may fit hyperparameters and optimize acquisitions using the recorded history, but cannot access θ_* or latent values other than through the observations. All seeds are on standard Borel spaces. Set $Z = (S, Y_{1:T})$, and stack the realized matrices and observations as $H(Z), y(Z)$.

Choose the recommendation, tolerance $\varepsilon(Z) \ge 0$, cells, bounds, and simulation sample counts as measurable functions of Z, before final certification simulation. The design set and prior mean are fixed; observation noise is known and positive. This contract is an exact GP model, averaged over the latent GP draw. It is not a frequentist guarantee uniformly over arbitrary fixed functions, and does not cover misspecified kernels or learned noise.

Lemma 9 (Predictable sampling gives the actual posterior)**.** *Conditional on $S = s$, the normalized observation density is*

$$L_\theta(s, y) = \int \prod_{t=1}^{T} \varphi_{\sigma^2 I_k}(y_t - H_t(s, y_{<t})u)\, N(m_0, C_\theta)(du).$$

For each realized history it equals

$$\varphi_{H(Z) C_\theta H(Z)^* + \sigma^2 I_{Tk}}(y - H(Z)m_0) > 0.$$

The conditional law of h given Z, under parameter θ, is the exact posterior used in Theorem 1 with the realized $H(Z), y$.

Proof. Conditional on $h, S, Y_{<t}$, predictability and fresh independent noise give the tth density factor. Multiplication followed by integration over the Gaussian prior yields the first formula. Its normalization follows by iterated integration in time. At a fixed history, $H(Z)$ is a numerical matrix, so Gaussian integration gives the second formula and the Gaussian posterior. The denominator is positive because its covariance dominates $\sigma^2 I$. Bayes' formula therefore supplies the stated regular conditional version. The adaptive Y is generally not a jointly Gaussian vector; its normalized density is justified by the sequential product, not by treating $H(Z)$ as an independent random design. Matrix entries are measurable, inversion occurs only on positive matrices, and these formulas give measurable posterior kernels. □

C.2 A RETAINED CONTINUOUS FAMILY

Fix a reference probability measure π on Θ independently of the experiment, and let $0 < \delta_{\mathrm{ret}} < 1$. Define

$$M(Z) = \int_{\Theta} L_{\vartheta}(Z)\, \pi(d\vartheta), \qquad \mathcal{R}(Z) = \{\theta : L_{\theta}(Z) \geq \delta_{\mathrm{ret}} M(Z)\}.$$

For every fixed truth θ_*, $\mathbb{P}_{\theta_*}(\theta_* \notin \mathcal{R}(Z)) \leq \delta_{\mathrm{ret}}$. To prove this, condition on each seed and apply Tonelli to get $\mathbb{E}_{\theta_*}[M/L_{\theta_*} \mid S] = 1$; Markov's inequality gives the assertion and averaging over the seed preserves it. Continuity of the likelihood on compact Θ makes M finite and ensures a maximizer with $L_{\theta} \geq M$, so the retained set is nonempty. No hyperprior on the truth is assumed. A data-fitted reference π is outside this argument.

Choose a finite ordered list of nonempty compact convex cells $D_j(Z)$ with measurable membership, centers and radii, whose union contains $\mathcal{R}(Z)$. The list length $J(Z)$ is a finite measurable positive integer, and each cell has a measurable membership graph on its index domain. Require a certified outer cover, not only a test at grid points. A fixed rectangular partition covering all Θ is a simple valid choice. All constants from Appendix B must be valid on each entire cell. Choose $0 < \rho_j(Z) < \alpha$, with $0 < \alpha < 1$, and compute η_j by Theorem 1. Put $q_j(Z) = p_{c_j}(\varepsilon(Z) - \eta_j \mid Z)$.

C.3 EXACT SIMULATION AND DECISION

Let the fresh simulation object V satisfy $V \perp h \mid Z$. Conditional on Z, generate $K_j(Z) \geq 1$ independent exact posterior draws at anchor c_j, measurably from (Z, V). Measurable positive-semidefinite square roots and a fresh Gaussian array suffice, including for singular covariances. Independence across cells is unnecessary. Define

$$\widehat{q}_j = \frac{1}{K_j} \sum_{l=1}^{K_j} \mathbf{1}\{R(h^{jl}, Z) > \varepsilon(Z) - \eta_j\}, \qquad U_j = \min\left\{1, \widehat{q}_j + \sqrt{\frac{\log(1/\beta)}{2K_j}}\right\}, \quad 0 < \beta < 1.$$

The counts, thresholds, cells and recommendation must be fixed before these draws. The one-sided Bernoulli Hoeffding bound gives $\Pr(q_j > U_j \mid Z) \leq \beta$. For completeness the centered Bernoulli log-mgf has second derivative at most $1/4$, so its log-mgf is at most $s^2/8$; exponential Markov and optimization give $\Pr(q - \widehat{q} > b) \leq e^{-2Kb^2}$. Capping at one preserves coverage.

Exact binomial alternative. For the fixed, history-measurable count $K = K_j$, let $S = K\widehat{q}_j$. If $S < K$, define $U^{\mathrm{bin}}(S)$ as the unique $u \in (0, 1)$ satisfying

$$\sum_{k=0}^{S} \binom{K}{k} u^k (1-u)^{K-k} = \beta;$$

set $U^{\mathrm{bin}}(K) = 1$. For any true Bernoulli probability q, the left side is decreasing in u and increasing in S. Thus $\{q > U^{\mathrm{bin}}(S)\}$ is a lower tail of $S \sim \mathrm{Bin}(K, q)$ whose probability is at most β. Conditioning on the history fixes K, so this bound can replace U_j. At $S = 0$ it equals $1 - \beta^{1/K}$; it need not improve every count and does not allow stopping simulation adaptively.

Certify, event $\mathcal{C}$, only when every covering cell satisfies $U_j + \rho_j \leq \alpha$. If any required rigorous bound, cover or exact draw is unavailable, the specified mathematical procedure does not issue a certificate.

Theorem 10 (Fixed-terminal actual-model guarantee)**.** *Under the full contract above, for every fixed* $\theta_* \in \Theta$*,*

$$\mathbb{P}_{\theta_*}\big(\mathcal{C} \text{ and } f_h(i_\star(Z)) - \min_i f_h(i) > \varepsilon(Z)\big) \le \min\{1, \alpha + \delta_{\mathrm{ret}} + \beta\}.$$

There is no cell-count factor multiplying β*.*

Proof. Let $E = \{\theta_* \in \mathcal{R}(Z)\}$. On E select the first covering cell $j_*(Z)$ containing θ_*. Extend j_* to be one outside E. The events $E \cap \{j_* = j\}$ intersect measurable cell membership with complements of earlier membership, so j_* is history-measurable. With $F_j = \{q_j > U_j\}$, conditional coverage gives

$$\Pr(E \cap F_{j_*} \mid Z) = \mathbf{1}_E \sum_{j \le J(Z)} \mathbf{1}_{\{j_* = j\}} \Pr(F_j \mid Z) \le \beta \mathbf{1}_E.$$

Thus $\Pr(E \cap \{q_{j_*} > U_{j_*}\}) \le \beta$ without a union over cells. On $G = \mathcal{C} \cap E \cap \{q_{j_*} \le U_{j_*}\}$, Theorem 1 and the decision rule give

$$p_{\theta_*}(\varepsilon(Z) \mid Z) \le q_{j_*} + \rho_{j_*} \le U_{j_*} + \rho_{j_*} \le \alpha.$$

If V denotes final simulation randomness and B the true bad-regret event, Lemma 9 and conditional independence give the actual-law identity

$$\mathbb{E}_{\theta_*}[\mathbf{1}_B \mid Z, V] = p_{\theta_*}(\varepsilon(Z) \mid Z).$$

Integrating on the (Z, V)-measurable event G yields $\Pr(G \cap B) \le \alpha \Pr(G) \le \alpha$. Outside G, a false certificate entails retention failure or the witness's Monte Carlo failure. Their probabilities are at most δ_{ret} and β. Addition and the probability cap prove the theorem. □

C.4 Deterministic and hybrid certificate substitution

Corollary 11 (Deterministic and hybrid cell certificates)**.** *Under the correctly specified, fixed-terminal observation and retention contract of Appendix C, let* $D_j(Z)$ *be the same finite measurable outer cover. Suppose* $b_j(Z) \in [0, 1]$ *is history-measurable and satisfies*

$$b_j(Z) \ge \sup_{\theta \in D_j(Z)} p_\theta(\varepsilon(Z) \mid Z)$$

for almost every history, for every fixed truth in the model family. If a deterministic rule issues a certificate only when every $b_j \le \alpha$*, its joint false-certificate probability is at most* $\min\{1, \alpha + \delta_{\mathrm{ret}}\}$*. Alternatively, retain the original simulation contract and certify only if, for every cell,* $\min\{b_j, \min(1, U_j + \rho_j)\} \le \alpha$*. This hybrid rule has the original bound* $\min\{1, \alpha + \delta_{\mathrm{ret}} + \beta\}$*.*

Proof. On retention, let $j_*(Z)$ be the first cell containing the fixed true parameter. In the deterministic rule, certification gives $p_{\theta_*}(\varepsilon \mid Z) \le b_{j_*} \le \alpha$. Integrate the conditional probability over this history-measurable certification event. Outside retention, pay at most δ_{ret}. For the hybrid, on the original witness Monte Carlo good event, both b_{j_*} and $\min(1, U_{j_*} + \rho_{j_*})$ upper bound the true posterior regret probability. Their minimum also does. The original joint independence of the entire simulation object from the actual latent truth conditional on Z permits the same conditional integration. Retention and witness simulation failure add at most $\delta_{\mathrm{ret}} + \beta$. No union over cells is added. □

Corollary 13 supplies b_j on histories satisfying its posterior assumptions, positive tolerance and $0 < r \le r_0$. History-selected centers and radii preserve pointwise validity, not a history-uniform asymptotic rate. Use $b_j = 1$ on ineligible histories, including $\varepsilon = 0$; the hybrid simulation bound remains available. No extra transfer budget is added to a direct event bound. Analytic quantities or rigorous numerical enclosures are required.

D Local sharpness and concentration-transfer loss

D.1 Structured posterior

Take distinct designs $x_1 = -d/2$, $x_2 = d/2$ with $d > 0$, one zero-mean RBF output, known observation noise $\sigma > 0$, and one observation $Y = (h_1 + h_2)/\sqrt{2} + \sigma\xi = y$. Allow $y \in \mathbb{R}$ and

recommend design 1. Write $\theta = (\alpha, \lambda) = (\log a, \log \ell) \in \mathbb{R}^2$. Define

$$k = e^{-d^2/(2\ell^2)}, \quad b = 1 + k, \quad \delta = \sqrt{1-k}, \quad w = \frac{a^2 b}{a^2 b + \sigma^2}, \quad s = 2a\delta\sigma\sqrt{w}, \quad \nu = \frac{|y|\sqrt{w}}{\sigma}.$$

All parameters here are finite and $0 < k, w < 1$, $s > 0$. Hold d, σ, y and the physical tolerance $\varepsilon > 0$ fixed. The posterior has independent sum and difference coordinates

$$S \sim N(wy, \sigma^2 w), \qquad D \sim N(0, a^2\delta^2), \qquad C = h_1^2 - h_2^2 = 2SD.$$

Consequently C has the same marginal law as $s(Z + \nu)X$, for independent standard normals X, Z. Reflecting Z handles negative y. This marginal reduction is used only for event probabilities; it does not replace the original Fourier/Matheron coupling in the comparator below.

Let $R = \max(0, C)$ and $p_\theta(v) = \Pr_\theta(R > v \mid y)$. Fix a center $c \in \mathbb{R}^2$ and the full two-parameter box $D_r = c + [-r, r]^2$. Write

$$A_r^+ = \sup_{\theta \in D_r} p_\theta(\varepsilon) - p_c(\varepsilon), \qquad A_r = \sup_{\theta \in D_r} |p_\theta(\varepsilon) - p_c(\varepsilon)|.$$

The history, designs and tolerance remain fixed as $r \downarrow 0$.

D.2 Sharp event-level ambiguity

Let φ, Φ denote the standard Gaussian density and distribution function. Set $\overline{\Phi} = 1 - \Phi$ and $z = \varepsilon/s$.

Theorem 12 (Exact probability, positive sensitivity and sharp local cost)**.** *In the preceding model,*

$$p_\theta(\varepsilon) = \int_0^\infty \varphi(x) \left[\overline{\Phi}(z/x - \nu) + \overline{\Phi}(z/x + \nu)\right] dx. \tag{1}$$

This probability is C^2 in θ and belongs to $(0, 1/2)$. Put $e_j = \partial_j \log s$ and $n_j = \partial_j \nu$. Its gradient is

$$g_j(\theta) = \int_0^\infty \varphi(x) \left[e_j \frac{z}{x}\{\varphi(z/x - \nu) + \varphi(z/x + \nu)\} + n_j\{\varphi(z/x - \nu) - \varphi(z/x + \nu)\}\right] dx. \tag{2}$$

With $\dot{b} = kd^2/\ell^2$, the coefficients are explicitly

$$e_\alpha = 2 - w, \qquad n_\alpha = \nu(1 - w),$$
$$e_\lambda = \frac{\dot{b}}{2}\left[\frac{1-w}{b} - \frac{1}{2-b}\right], \qquad n_\lambda = \frac{\nu(1-w)\dot{b}}{2b}.$$

In particular $g_\alpha(\theta) > 0$ for every finite y.

For any fixed $r_0 > 0$, there is a finite, fully specified Hessian bound H on D_{r_0} such that, for $0 < r \le r_0$, both A_r^+ and A_r lie between

$$r\|g(c)\|_1 - \tfrac{1}{2}Hr^2 \quad \text{and} \quad r\|g(c)\|_1 + \tfrac{1}{2}Hr^2. \tag{3}$$

Thus $A_r^+ = r\|g(c)\|_1 + O(r^2)$ and $A_r = r\|g(c)\|_1 + O(r^2)$, with $\|g(c)\|_1 > 0$.

Corollary 13 (A locally sharp robust certificate)**.** *For $0 < r \le r_0$, with H as in Theorem 12, the deterministic quantity $B_{\text{event}}(r) = \min\{1, p_c(\varepsilon) + r\|g(c)\|_1 + Hr^2/2\}$ is a uniform upper bound for $p_\theta(\varepsilon)$ throughout D_r. Its excess above the exact worst-case probability is at most Hr^2. For fixed length scale and any amplitude interval $[a_-, a_+]$, the exact worst-case probability is $p_{a_+,\ell}(\varepsilon)$.*

D.3 A matched separation from the previous transfer formula

Keep exactly the same center, two-parameter box, event and history. Let A_θ be the original common Fourier/Matheron residual factor and μ_θ its posterior mean, with $J = \operatorname{diag}(1, -1)$. Define on that common Hilbert space

$$T_\theta = A_\theta^* J A_\theta, \qquad v_\theta = 2A_\theta^* J \mu_\theta, \qquad q_\theta = \mu_\theta^* J \mu_\theta + \operatorname{tr} T_\theta = 0.$$

The last identity follows because the two means and the two posterior marginal variances are equal. Use the *exact* envelopes

$$O_r = \sup_{D_r} \|T_\theta - T_c\|_{\text{op}}, \quad F_r = \sup_{D_r} \|T_\theta - T_c\|_{\text{HS}}, \quad L_r = \sup_{D_r} \|v_\theta - v_c\|.$$

For $0 < \rho < 1$, set

$$\eta_r(\rho) = 2\sqrt{(F_r^2 + L_r^2/2)\log(1/\rho)} + 2O_r \log(1/\rho).$$

The best certificate within this fixed-reference transfer formula is

$$B_{\mathrm{conc}}(r) = \inf_{0<\rho<1} \min\{1, p_c(\varepsilon - \eta_r(\rho)) + \rho\}.$$

The comparator uses exact probabilities and envelopes, with an optimal risk split.

Theorem 14 (Sharp logarithmic loss of fixed-reference concentration transfer)**.** *Let $f_c(\varepsilon) > 0$ be the density of C_c at ε, and set $o_c = \sup_{\|u\|_\infty \le 1} \|DT_c[u]\|_{\mathrm{op}}$. Then $0 < o_c < \infty$, and*

$$B_{\mathrm{conc}}(r) - p_c(\varepsilon) = 2f_c(\varepsilon)o_c\, r\log(1/r)\,(1 + o(1)). \tag{4}$$

Consequently

$$\frac{B_{\mathrm{conc}}(r) - p_c(\varepsilon)}{A_r^+} \sim \frac{2f_c(\varepsilon)o_c}{\|g(c)\|_1}\log(1/r) \longrightarrow \infty, \qquad \frac{B_{\mathrm{event}}(r) - p_c(\varepsilon)}{A_r^+} \longrightarrow 1.$$

The ratio concerns ambiguity penalties*; the total probabilities have the same nonzero limiting reference probability. It is not a divergence of total certificate ratios.*

The separation concerns this fixed-reference, fixed-coupling formula only.

D.4 The original posterior coupling

Fix two distinct inputs x_1, x_2, a finite length scale $\ell > 0$, and the amplitude-one Gaussian-frequency feature rows ϕ_i in the prior Hilbert space. The exact kernel identity gives unit row norms and inner product $k = \exp[-\|x_1 - x_2\|^2/(2\ell^2)] \in (0, 1)$. The cosine and sine features are linear in amplitude in their actual L^2 spaces; hence $F_a = aF_1$ as an operator, with no feature truncation. Embed the prior rows in the feature-plus-one-noise Hilbert space, and let e_0 be the noise coordinate. Put $b = 1 + k$, $\delta = \sqrt{1-k}$ and define

$$e_s = \frac{\phi_1 + \phi_2}{\sqrt{2b}}, \qquad e_d = \frac{\phi_1 - \phi_2}{\sqrt{2}\,\delta}.$$

Both denominators are positive. Expanding inner products gives $\|e_s\| = \|e_d\| = 1$ and $\langle e_s, e_d\rangle = 0$. The noise coordinate is orthogonal to the prior space and has norm one. These three vectors are independent of amplitude.

Observe the normalized sum through $H = (1,1)/\sqrt{2}$, with known noise standard deviation $\sigma > 0$. Its prior signal row is $a\sqrt{b}\,e_s$. Since $HC_aH^* = a^2b$ and $C_aH^* = a^2b(1,1)^\top/\sqrt{2}$, the gain is $G_a = \frac{w_a}{\sqrt{2}}\binom{1}{1}$, $\qquad w_a = \frac{a^2b}{a^2b+\sigma^2}$. With zero prior mean, the posterior mean is $\mu_{a,i} = w_a y/\sqrt{2}$. The exact signed residual row implemented by the Matheron factor is

$$a\phi_i - G_{a,i}(a\sqrt{b}\,e_s + \sigma e_0) = \frac{a\sqrt{b}(1-w_a)e_s - \sigma w_a e_0 \;\pm\; a\delta e_d}{\sqrt{2}}.$$

The sign is positive for $i = 1$ and negative for $i = 2$. These are the exact rows underlying the single-sum structured posterior. Applying the same isonormal process W therefore gives the same draws almost surely, for every fixed amplitude comparison. Only finitely many linearity identities are required in each comparison.

Appendices A and C establish the posterior and adaptive-law identities.

D.5 Proofs

D.5.1 Posterior reduction and differentiation

Gaussian conditioning on the sum changes its variance from a^2b to $\sigma^2 w$ and its mean to wy. The difference coordinate is independent of the observed sum and keeps variance $a^2(1-k)$. This proves the stated independent product representation. Conditional on $X = x > 0$, the event $s(Z + \nu)X > \varepsilon$ has probability $\overline{\Phi}(z/x - \nu)$. Conditional on $X = -x$, its probability is $\overline{\Phi}(z/x + \nu)$.

Integrating proves (1). The contrast is symmetric because X is symmetric and independent; it has positive probability in every nonempty interval away from zero. Hence its positive-threshold survival probability lies strictly between zero and one half.

For $u = z/x$, let

$$G(u,\nu) = \overline{\Phi}(u-\nu) + \overline{\Phi}(u+\nu), \quad S = \varphi(u-\nu) + \varphi(u+\nu), \quad D = \varphi(u-\nu) - \varphi(u+\nu).$$

Then $G_u = -S$, $G_\nu = D$, $u_j = -e_j u$, yielding (2). Logarithmic differentiation of w gives $w_\alpha = 2w(1-w), \qquad w_\lambda = w(1-w)\dot{b}/b$; differentiating s, ν yields the displayed coefficients. Since $u > 0$, $\nu \geq 0$ and $(u-\nu)^2 \leq (u+\nu)^2$, we have $D \geq 0$. Also $e_\alpha = 2 - w > 1$ and $n_\alpha \geq 0$. The first integrand term for g_α is strictly positive for every $x > 0$, so $g_\alpha > 0$, including $y = 0$.

For domination and a computable Hessian bound, define on D_{r_0} $T = \sup \nu$, $E_1 = \max_j \sup |e_j|$, $E_2 = \max_{i,j} \sup |\partial_i e_j|$, $N_1 = \max_j \sup |n_j|$, $N_2 = \max_{i,j} \sup |\partial_i n_j|$. These compact extrema of explicit smooth functions may be replaced by verified upper bounds. Put

$$M_0 = \varphi(0), \qquad M_j = 2^{j-1}\left[(j/e)^{j/2}\varphi(0) + T^j\varphi(0)\right], \quad j = 1,2,3.$$

For every $u \geq 0$, $0 \leq \nu \leq T$, $u^j\varphi(u \pm \nu) \leq M_j$. Indeed $u \leq |u \pm \nu| + T$, use $(v+T)^j \leq 2^{j-1}(v^j + T^j)$ and maximize $v^j\varphi(v)$ at $v = \sqrt{j}$. The $j = 0$ bound is immediate.

The second derivatives satisfy

$$\begin{aligned} G_{uu} = G_{\nu\nu} &= (u-\nu)\varphi(u-\nu) + (u+\nu)\varphi(u+\nu), \\ G_{u\nu} &= -(u-\nu)\varphi(u-\nu) + (u+\nu)\varphi(u+\nu), \\ u_{ij} &= (e_i e_j - \partial_i e_j)u. \end{aligned}$$

The chain rule gives

$$\partial_{ij} G = G_{uu} u_i u_j + G_{u\nu}(u_i n_j + u_j n_i) + G_{\nu\nu} n_i n_j + G_u u_{ij} + G_\nu \partial_i n_j.$$

The terms are uniformly dominated by a constant times $\varphi(x)$, justifying twice-continuous differentiation under the integral. Using $\int_0^\infty \varphi(x)dx = 1/2$, each Hessian entry is bounded in absolute value by

$$\begin{aligned} H_0 = {} & E_1^2(M_3 + TM_2) + 2E_1N_1(M_2 + TM_1) \\ & + N_1^2(M_1 + TM_0) + (E_1^2 + E_2)M_1 + N_2M_0. \end{aligned}$$

Thus $H = 4H_0$ bounds $|u^T \nabla^2 p_\theta(\varepsilon) u|$ for $\|u\|_\infty \leq 1$.

The contrast density at a positive v is

$$f_\theta(v) = \int_0^\infty \frac{\varphi(x)}{sx}\left[\varphi\left(\frac{v}{sx} - \nu\right) + \varphi\left(\frac{v}{sx} + \nu\right)\right] dx > 0. \tag{5}$$

It is finite and continuous locally for $v > 0$: writing $u = v/(sx)$ converts $1/(sx)$ into u/v, so the preceding M_1 bound supplies local integrable domination. In particular $p_c(\varepsilon - h) - p_c(\varepsilon) = f_c(\varepsilon)h + o(h) \quad (h \downarrow 0)$. No density assertion at the atom of regret at zero is needed.

D.5.2 Sharp ambiguity and the event certificate

For $\theta = c + ru$, $\|u\|_\infty \leq 1$, Taylor's theorem on its line segment gives

$$p_\theta(\varepsilon) - p_c(\varepsilon) = rg(c)^T u + \mathcal{E}, \qquad |\mathcal{E}| \leq Hr^2/2.$$

The maximum of $g^T u$ on the square is $\|g\|_1$. Choosing a sign vertex realizes that linear value, while every point is upper bounded by it. This proves the two bounds for A_r^+. The upper bound for A_r follows from $|g^T u| \leq \|g\|_1$, and its lower bound follows from $A_r \geq A_r^+$. The strict amplitude derivative makes the leading coefficient positive. The event certificate upper bounds every parameter by the upper Taylor bound. Subtracting the lower bound on the exact supremum shows its excess is at most Hr^2; clipping at one cannot increase that excess. Strict amplitude monotonicity proves the amplitude-only endpoint statement.

D.5.3 Regularity and a nonzero operator slope in the original coupling

Appendix B establishes C^2 regularity of the original common Fourier/Matheron factor and its contrast coefficients. In particular, T_θ and v_θ are C^1 in the required norms.

Uniform first-order expansion over the unit parameter square gives $O_r/r \longrightarrow o_c, \qquad F_r = O(r), \qquad L_r = O(r)$. For example, $T_{c+ru} - T_c = rDT_c[u] + o(r)$ uniformly for $\|u\|_\infty \le 1$; taking operator norms and suprema proves the first assertion. Compactness and continuity make o_c finite.

To prove positivity, vary amplitude alone while fixing the center length scale. The amplitude-one Fourier sum/difference rows normalize to two fixed orthonormal vectors, independent of amplitude; adjoining the observation-noise vector gives the same three coordinates as before. In those coordinates

$$C_a = \delta\{2yU(a)Z_d + 2\sqrt{b}V(a)Z_sZ_d - 2\sigma U(a)Z_0Z_d\}, \quad U(a) = \frac{a^3b}{a^2b+\sigma^2}, \quad V(a) = \frac{a^2\sigma^2}{a^2b+\sigma^2}.$$

The derivative quadratic operator is a symmetric star matrix, so

$$\|\partial_\alpha T_c\|_{\mathrm{op}} = \delta\sqrt{b(\partial_\alpha V)^2 + \sigma^2(\partial_\alpha U)^2} > 0, \qquad \partial_\alpha U = aw(3-2w) > 0.$$

It follows that $o_c \ge \|\partial_\alpha T_c\|_{\mathrm{op}} > 0$.

D.5.4 Proof of the sharp optimized transfer asymptotic

Abbreviate $p_0 = p_c(\varepsilon)$, $f_0 = f_c(\varepsilon)$, $L = \log(1/r)$, and let $E_r = B_{\mathrm{conc}}(r) - p_0$. The transfer inequality guarantees $E_r \ge A_r^+ \ge 0$.

For the upper bound choose $\rho = r$ for $r < 1$. Since $O_r/r \to o_c$ and $F_r, L_r = O(r)$, $\eta_r(r) = 2o_c rL(1+o(1)) + O(r\sqrt{L})$. This tends to zero. The density expansion gives

$$\limsup_{r\downarrow 0} \frac{E_r}{rL} \le 2f_0o_c;$$

the added $\rho = r$ is negligible relative to rL.

For the lower bound take any sequence $r \downarrow 0$ and choose $\rho_r \in (0,1)$ whose clipped cost is within r^2 of the infimum; such a choice follows from the definition of infimum, without assuming attainment. The upper bound implies this cost is $p_0 + O(rL)$. Since $p_0 < 1$, clipping is inactive for small r. Therefore

$$0 \le \rho_r \le O(rL), \qquad 0 \le p_c(\varepsilon - \eta_r(\rho_r)) - p_0 \le O(rL).$$

The second assertion forces $\eta_r(\rho_r) \to 0$: otherwise a fixed positive shift of at most $\varepsilon/2$ and the strictly positive density on that threshold interval give a fixed positive probability increase. Consequently the density expansion applies to the selected shifts. The first assertion implies $\log(1/\rho_r) \ge L - \log L - O(1) = L(1-o(1))$. Because $\eta_r(\rho_r) \ge 2O_r\log(1/\rho_r)$, the selected excess is at least $(f_0 + o(1))\eta_r(\rho_r) \ge 2f_0o_c\,rL(1-o(1))$. The r^2 infimum approximation is negligible. Arbitrariness of the sequence proves the lower limit and (4). The ratios follow from Theorem 12 and Corollary 13.

E Informative observations: statements and proofs

E.1 Actual observation model and probability

Let $x_1 = -d/2, x_2 = d/2$, $d > 0$, and let (h_1, h_2) have zero-mean RBF covariance with amplitude $a > 0$ and length scale $\ell > 0$. Put $\theta = (\log a, \log \ell), \quad k = e^{-d^2/(2\ell^2)}$,

$$X_+ = (h_1 + h_2)/\sqrt{2}, \quad X_- = (h_1 - h_2)/\sqrt{2}, \quad v_\pm = a^2(1 \pm k).$$

Observe $Y_i = X_i + \sqrt{t}\xi_i = y_i$ for $i \in \{+,-\}$, independent standard Gaussian ξ_i and known $t > 0$. Repeated original-design measurements give these sufficient statistics with $t = \sigma^2/n$; the orthonormal transformation preserves independent equal-variance noise. For recommendation 1, regret exceeding $\varepsilon > 0$ is the event $E_\varepsilon = \{2X_+X_- > \varepsilon\}$. Gaussian conditioning gives independent posterior coordinates

$$X_i = m_i + \sqrt{V_i}Z_i, \qquad w_i = \frac{v_i}{v_i + t}, \quad m_i = w_iy_i, \quad V_i = tw_i, \qquad Z_+, Z_- \overset{\mathrm{iid}}{\sim} N(0,1). \quad (6)$$

Write $p_\theta(t, y) = \mathbb{P}_\theta(E_\varepsilon \mid y)$, holding t, y, ε fixed in parameter derivatives. Both means may be nonzero; no symmetry or amplitude monotonicity is assumed.

Fix a compact convex $K \subset \mathbb{R}^2$, contained in an open parameter neighborhood, and $Y < \infty$ with $|y_i| \le Y$. Choose constants

$$v_i \ge v_* > 0, \quad |D_u \log v_i| \le L, \quad |D_u^2 \log v_i| \le M \quad (\theta \in K,\ \|u\|_\infty \le 1). \tag{7}$$

These exist for the RBF variances. Specifically, with $b = kd^2/\ell^2$ and $b' = b(d^2/\ell^2 - 2)$,

$$\partial_\alpha \log v_i = 2, \quad \partial_\lambda \log v_\pm = \frac{\pm b}{1 \pm k}, \quad \partial_{\lambda\lambda} \log v_\pm = \frac{\pm b'}{1 \pm k} - \frac{b^2}{(1 \pm k)^2},$$

while the other second derivatives vanish. Compact suprema of these explicit functions specify L, M.

E.2 Uniform event sensitivity and a shrinking remainder

For direction u, define $l_i = D_u \log v_i$ and

$$A_i = \frac{y_i \sqrt{w_i}(1 - w_i)}{\sqrt{t}} l_i, \qquad B_i = \tfrac{1}{2}(1 - w_i) l_i, \qquad S_u = \sum_i \{A_i Z_i + B_i (Z_i^2 - 1)\}.$$

Define nonnegative bounds

$$\begin{aligned} a_t &= Y L \sqrt{t}/v_*, & b_t &= Lt/(2v_*), \\ a_t' &= Y(M + L^2)\sqrt{t}/v_*, & b_t' &= (M + L^2)t/(2v_*), \\ J_t &= 2(a_t^2 + 2b_t^2), & H_t &= 2\{a_t' + 2a_t^2 + 3a_t b_t + \sqrt{2} b_t' + 4b_t^2\}. \end{aligned}$$

Theorem 15 (Actual probability and uniform derivative control)**.** *For fixed $t > 0$, p_θ is C^2 in θ and*

$$D_u p_\theta = \mathbb{E}_\theta[\mathbf{1}_{E_\varepsilon} S_u], \qquad |D_u p_\theta| \le \sqrt{p_\theta(1 - p_\theta)} \sqrt{\sum_i (A_i^2 + 2B_i^2)} \le \tfrac{1}{2}\sqrt{J_t}. \tag{8}$$

Moreover $|D_u^2 p_\theta| \le H_t$ uniformly over K, bounded y, and $\|u\|_\infty \le 1$. Thus $J_t = O(t)$ and $H_t = O(\sqrt{t})$ as $t \downarrow 0$, with constants independent of the history in this class.

An exact one-dimensional implementation of the probability and first derivative is available. Put $s(z) = m_+ + \sqrt{V_+}\, z$, $a(z) = (\varepsilon/(2s(z)) - m_-)/\sqrt{V_-}$, and

$$q(z) = \begin{cases} \overline{\Phi}(a(z)), & s(z) > 0, \\ \Phi(a(z)), & s(z) < 0. \end{cases}$$

At the measure-zero point $s(z) = 0$, assign any bounded value. Then

$$p_\theta = \int_{\mathbb{R}} \varphi(z) q(z)\, dz, \tag{9}$$

$$D_u p_\theta = \int_{\mathbb{R}} \varphi(z) \left[q(z)\{A_+ z + B_+(z^2 - 1)\} + \operatorname{sgn}(s(z)) \varphi(a(z)) \{A_- + B_- a(z)\} \right] dz. \tag{10}$$

Corollary 16 (Local certificate and fixed-family narrowing)**.** *For a box $D_r = c + [-r, r]^2 \subset K$, put $g = \nabla p_c$ and $\mathcal{A}_r^+ = \sup_{D_r} p_\theta - p_c$. Then*

$$r\|g\|_1 - \tfrac{1}{2} H_t r^2 \le \mathcal{A}_r^+ \le r\|g\|_1 + \tfrac{1}{2} H_t r^2.$$

The same bounds hold for $\sup_{D_r} |p_\theta - p_c|$. The valid cell certificate

$$B_{\mathrm{loc}} = \min\{1, p_c + r\|g\|_1 + H_t r^2/2\}$$

has excess at most $H_t r^2 = O(r^2\sqrt{t})$ above the exact cell supremum. No nonzero-gradient assumption is needed; if $g = 0$ only the quadratic upper bound is implied. Over the entire fixed K,

$$\operatorname{osc}_K p := \sup_K p - \inf_K p \le \min\{1, \tfrac{1}{2} \operatorname{diam}_\infty(K) \sqrt{J_t}\} = O(\sqrt{t}). \tag{11}$$

E.3 Decision margin, a sharp boundary rate, and actual histories

Theorem 17 (Exponential narrowing away from the threshold)**.** *Suppose* $|2y_+y_- - \varepsilon| \geq \gamma > 0$, $|y_i| \leq Y$, *and choose* $0 < d_0 \leq \min\{1, \gamma/(8(Y+1))\}$. *If* $tY/v_* \leq d_0/2$, *then, uniformly in* $\theta \in K$,

$$\left|p_\theta(t,y) - \mathbf{1}_{\{2y_+y_- > \varepsilon\}}\right| \leq b(t) := \min\{1, 4e^{-d_0^2/(8t)}\}, \qquad \operatorname{osc}_K p \leq b(t). \tag{12}$$

In the safe case $2y_+y_- \leq \varepsilon - \gamma$, $b(t)$ *itself is a uniform regret certificate. In the unsafe case the result yields agreement near probability one, not safety.*

Theorem 18 (The square-root rate is attained on the boundary)**.** *Fix* $y_+ > 0, y_- > 0, \varepsilon = 2y_+y_- > 0$, *and write* $R_y = (y_+^2 + y_-^2)^{1/2}$. *Uniformly in* $\theta \in K$,

$$p_\theta(t,y) = \frac{1}{2} - \frac{\varphi(0)y_+y_-}{R_y}\left(\frac{1}{v_+} + \frac{1}{v_-} + \frac{1}{R_y^2}\right)\sqrt{t} + O(t). \tag{13}$$

The $O(t)$ *remainder remains* $O(t)$ *after up to two parameter derivatives on a compact neighborhood of* K. *In particular,*

$$\partial_\alpha p_\theta(t,y) = \frac{2\varphi(0)y_+y_-}{R_y}\left(\frac{1}{v_+} + \frac{1}{v_-}\right)\sqrt{t} + O(t) > 0$$

for sufficiently small t, *uniformly on* K. *If* K *contains two points with distinct* $v_+^{-1} + v_-^{-1}$ *(for example distinct amplitudes at fixed length scale), then*

$$\operatorname{osc}_K p = \frac{\varphi(0)y_+y_-}{R_y}\operatorname{osc}_K(v_+^{-1} + v_-^{-1})\sqrt{t} + O(t). \tag{14}$$

Hence (11) *has the optimal order uniformly over bounded histories.*

Corollary 19 (Repeated measurements without parameter contraction)**.** *Fix any true finite* θ_* *in the model and repeatedly observe both original design values with known independent Gaussian noise variance* σ^2. *Let* y_n *be the transformed sample means and* $t_n = \sigma^2/n$. *For any fixed compact* K *as above and* $\varepsilon > 0$, *almost surely there exist finite random* $N, C, c > 0$ *such that*

$$\operatorname{osc}_{\theta \in K} p_\theta(t_n, y_n) \leq Ce^{-cn} \qquad (n \geq N).$$

The limit is the actual event indicator. No contraction of K *is required; the random constants give neither uniform sample complexity nor anytime certification.*

Proposition 20 (A nonzero floor when only the sum is observed)**.** *Suppose only* X_+ *is observed with noise variance* $t \downarrow 0$ *and the limiting observed value is* $y_+ \neq 0$. *Along a fixed-prior-variance curve* $v_+ = v$ *with two distinct* v_- *values,*

$$p_\theta(t, y_+) \longrightarrow \overline{\Phi}\left(\frac{\varepsilon}{2|y_+|\sqrt{v_-}}\right).$$

The two limiting probabilities differ. The sum likelihoods are identical at every t; *likelihood retention cannot distinguish the two parameters.*

E.4 Proofs

E.4.1 Posterior, scores, and explicit Hessian bound

Equation (6) follows by completing the square separately in the two independent coordinates. Independence of the transformed sample-average noises follows from their equal variances and the orthonormal transformation. For each fixed positive t, on a compact parameter neighborhood and bounded data, posterior variances have positive lower and finite upper bounds. The density and its first two parameter derivatives are therefore dominated by an integrable polynomial times a Gaussian density in the latent coordinates. Differentiating the integral over the fixed measurable event is valid. This also proves C^2 regularity; no differentiation of an indicator is used.

For a directional derivative, $D_u w_i = w_i(1-w_i)l_i$, $D_u m_i = y_i w_i(1-w_i)l_i$ and $D_u \log V_i = (1-w_i)l_i$. Thus the Gaussian density score is exactly S_u. Its mean is zero and its variance is

$\sum_i (A_i^2 + 2B_i^2)$ by orthogonality of Gaussian Hermite polynomials of degrees one and two. Cauchy–Schwarz applied to $(\mathbf{1}_E - p)S_u$ proves (8). Because $1 - w_i \le t/v_*$, $|A_i| \le a_t$ and $|B_i| \le b_t$.

Differentiate at fixed latent X. Writing $l_i' = D_u^2 \log v_i$ gives

$$A_i' = \frac{y_i\sqrt{w_i}(1-w_i)}{\sqrt{t}}\{l_i' + (\tfrac{1}{2} - \tfrac{3}{2}w_i)l_i^2\},$$
$$B_i' = \tfrac{1}{2}(1-w_i)\{l_i' - w_i l_i^2\}, \qquad D_u Z_i = -A_i - B_i Z_i.$$

Consequently $|A_i'| \le a_t'$, $|B_i'| \le b_t'$ and

$$D_u S_u = \sum_i \{A_i' Z_i - A_i^2 - A_i B_i Z_i + B_i'(Z_i^2 - 1) - 2A_i B_i Z_i - 2B_i^2 Z_i^2\}.$$

The second derivative is $D_u^2 p = \mathbb{E}[\mathbf{1}_E(S_u^2 + D_u S_u)]$. Using $\mathbb{E}|Z| \le 1$, $\mathbb{E}|Z^2 - 1| \le \sqrt{2}$, $\mathbb{E}Z^2 = 1$, and the score variance gives

$$|D_u^2 p| \le \sum_i \{|A_i'| + 2A_i^2 + 3|A_i B_i| + \sqrt{2}|B_i'| + 4B_i^2\} \le H_t.$$

The expressions are valid for all $t > 0$; their small-t orders follow directly. Conditioning on $Z_+ = z$ proves (9). The identities $\mathbb{E}[Z\mathbf{1}_{Z>a}] = \varphi(a)$ and $\mathbb{E}[(Z^2-1)\mathbf{1}_{Z>a}] = a\varphi(a)$, with negative versions for $Z < a$, give (10).

Apply the Taylor sign-direction argument of Appendix D.5.2 with the uniform bound H_t. It gives Corollary 16, including the zero-gradient case and the $H_t r^2$ excess bound. Integrating (8) along segments in convex K gives (11).

E.4.2 Margin bound and the actual history limit

If $|X_i - y_i| \le d_0$ for both coordinates, then

$$|2X_+X_- - 2y_+y_-| \le 4Yd_0 + 2d_0^2 \le (4Y+2)d_0 < \gamma.$$

Thus the event indicator agrees with the observed-product indicator on this rectangle. The posterior bias obeys $|m_i - y_i| \le tY/v_* \le d_0/2$, and $V_i \le t$. The ordinary Gaussian tail bound and a union bound give

$$\mathbb{P}_\theta(\exists i : |X_i - y_i| > d_0 \mid y) \le 4\exp(-d_0^2/(8t)).$$

If the common indicator is zero, every probability lies in $[0, b(t)]$; if it is one, every probability lies in $[1 - b(t), 1]$. Hence the oscillation is bounded by $b(t)$, not twice $b(t)$. This proves Theorem 17.

For repeated observations, the strong law for each noise-average sequence gives $y_n \to (X_+, X_-)$ almost surely. The true Gaussian vector is nondegenerate; conditioning on X_+ shows $\mathbb{P}(2X_+X_- = \varepsilon) = 0$ for positive ε (the case $X_+ = 0$ also has probability zero). On a realization outside this null set, choose finite Y bounding the eventual averages and a positive γ smaller than half the limiting margin. Eventually Theorem 17 applies, with fixed positive d_0 and $t_n = \sigma^2/n$. It gives the stated exponential rate with realization-dependent constants. This proves Corollary 19 on the fixed family K.

Unobserved-direction ambiguity. Along $v_+ = v$, the sum likelihood is $N(0, v + t)$ for every observed value, so the retention rule treats all such parameters identically. Their posterior sum laws agree, but $v_- = v(1-k)/(1+k)$ varies strictly with k. At finite t, the product representation in (1) has unchanged ν and scale strictly increasing in v_-. Thus their positive-threshold regret probabilities differ. Refining a cover reduces within-cell variation without removing this between-cell ambiguity.

As $t \downarrow 0$, $X_+ \mid y_+$ converges in probability to y_+, while the independent X_- retains law $N(0, v_-)$. The limiting product has no atom at positive ε, so continuity-set convergence gives Proposition 20. Its tail probability is strictly increasing in v_-. This establishes both finite-data nonidentifiability and the limiting ambiguity floor.

E.4.3 BOUNDARY EXPANSION WITH A UNIFORM DIFFERENTIATED REMAINDER

Fix $a = y_+ > 0, b = y_- > 0$, write $h = \sqrt{t}$, and work on a compact parameter neighborhood K' of K. Throughout this proof a, b denote observed coordinates, not the GP amplitude. Let $w_i(h) = \frac{v_i}{v_i+h^2}, \quad q_i(h) = \sqrt{w_i(h)}$,

$$s(h,z) = aw_+(h) + hq_+(h)z, \quad D(h) = \frac{v_+ + v_- + h^2}{(v_+ + h^2)(v_- + h^2)}.$$

On $s > 0$, the conditional event probability is $\Phi(U)$, where

$$U(\theta, h, z) = \frac{bw_-(h) - ab/s(h,z)}{hq_-(h)} = \frac{b\{w_-(h)q_+(h)z - ahD(h)\}}{q_-(h)s(h,z)}. \tag{15}$$

The second equality follows from $w_+w_- - 1 = -h^2D(h)$ and removes the apparent singularity at $h = 0$.

Choose $h_0 > 0$ sufficiently small. If $0 < h \le h_0$, $|z| \le h^{-1/4}$ and $0 \le u \le h$, then $s(u,z) \ge a/2$ uniformly on K'. This follows from $w_+(u) \to 1$ uniformly and $u|z| \le h^{3/4}$. On this domain the right side of (15) is a smooth quotient with denominators bounded away from zero. Its derivatives $\partial_\theta^\beta \partial_h^j U$, $|\beta| \le 2, j \le 2$, are bounded by $C(1+|z|^m)$ for some fixed finite C, m independent of h, z, θ. Indeed the derivatives of w_i, q_i, D are uniformly bounded on $K' \times [0, h_0]$, while differentiating $1/s$ introduces only finitely many factors linear in z and powers of s^{-1}.

Direct differentiation at $h = 0$ gives

$$U(\theta, 0, z) = bz/a, \qquad U_h(\theta, 0, z) = -b\{v_+^{-1} + v_-^{-1} + z^2/a^2\}.$$

Let $F = \Phi(U)$. Gaussian density derivatives are bounded. Thus $\partial_\theta^\beta F_{hh}$ has a uniform polynomial bound in z for $|\beta| \le 2$. The *exact integral remainder*

$$F(\theta,h,z) - F(\theta,0,z) - hF_h(\theta,0,z) = h^2 \int_0^1 (1-u)F_{hh}(\theta, uh, z)\, du \tag{16}$$

therefore implies, for every such multi-index,

$$\left|\partial_\theta^\beta \{F(\theta,h,z) - F(\theta,0,z) - hF_h(\theta,0,z)\}\right| \le Ch^2(1+|z|^m).$$

Integrating against $\varphi(z)$ on $|z| \le h^{-1/4}$ gives an $O(h^2)$ remainder uniformly through two parameter derivatives.

For completeness, define the full conditional integrand $Q_\theta(h,z)$ by $\Phi(U)$ on $s > 0$ and $1 - \Phi(U)$ on $s < 0$, with the original first expression for U in (15). At $s = 0$ set it to zero. We now bound its parameter derivatives on the excluded z tails. Since $q_- \ge c_* > 0$ on $K' \times [0, h_0]$, differentiating U at most twice gives, after increasing C if necessary,

$$|\partial_\theta^\beta U| \le Ch^{-1}(1+|z|)^2(1+|s|^{-3}), \qquad 1 \le |\beta| \le 2,$$

where $h_0 \le 1$ can be imposed. On $|s| \ge a/2$, denominators are bounded, so boundedness of φ, φ' bounds derivatives of Q by a polynomial in $h^{-1}, |z|$. On $0 < |s| < a/2$, use $0 < w_- \le 1$ to obtain

$$|bw_- - ab/s| \ge \frac{ab}{2|s|}, \qquad |U| \ge \frac{ab}{2h|s|}.$$

For every fixed integer j, both $|s|^{-j}\varphi(U)$ and $|s|^{-j}|\varphi'(U)|$ are bounded uniformly for $h \le h_0$: Gaussian decay absorbs reciprocal powers of s. The chain rule then yields, with a conservative common bound,

$$|\partial_\theta^\beta Q_\theta(h,z)| \le Ch^{-4}(1+|z|)^4, \qquad |\beta| \le 2. \tag{17}$$

The derivatives extend continuously by zero at $s = 0$, as does Q; hence the change of sign creates no differentiation boundary term. From (17), the integral of every such derivative over $|z| > h^{-1/4}$ is smaller than any fixed power of h. The same is true for the Gaussian integrals of $F(0,z)$ and $hF_h(0,z)$ and their parameter derivatives over this tail. The cutoff is independent of θ, so it contributes no

parameter derivative. These facts justify both differentiation of the full integral and extension of the central Taylor integrals to all $\mathbb{R}$.

Consequently

$$p_\theta = \int \varphi(z)\Phi(bz/a)\,dz - hb\int \varphi(z)\varphi(bz/a)\left(v_+^{-1} + v_-^{-1} + z^2/a^2\right)dz + O(h^2),$$

where the remainder is uniform in the $C^2(K)$ norm. Symmetry makes the first integral $1/2$. Completing the Gaussian square gives

$$\int \varphi(z)\varphi(bz/a)dz = \varphi(0)\frac{a}{\sqrt{a^2+b^2}}, \qquad \int z^2\varphi(z)\varphi(bz/a)dz = \varphi(0)\frac{a^3}{(a^2+b^2)^{3/2}}.$$

This proves (13) and the claimed C^2 remainder. Differentiating v_i^{-1} with respect to log amplitude gives $-2v_i^{-1}$, proving the positive leading derivative. Its coefficient has a positive minimum on K. Finally a uniform $O(t)$ error changes each supremum or infimum by at most $O(t)$, which proves (14).

F Finite comparison: essential proof and checks

F.1 Comparator validity and direct control

The entropy variational inequality gives $P(E) \le q_+(Q(E), \mathrm{KL}(P\|Q))$. A density constant on E and E^c attains the binary-KL bound; hence the indicator tilt is optimized exactly, rather than replaced by Pinsker's inequality. In the centered product model, sign agreement is an independent fair coin for every parameter. Folding the Gaussian coordinates preserves KL because their sign laws are identical and independent of magnitudes. Applying the same inequality to $\{2|X_+X_-| > \varepsilon\}$ gives B_{sign}. A fair-coin gate is a Markov kernel, so binary-KL data processing implies $B_{\mathrm{sign}} \le B_{\mathrm{GI}}$.

With $A = a^2$, the product variance factor satisfies

$$w_+w_- = \frac{A^2(1-k^2)}{A^2(1-k^2) + 2At + t^2}.$$

It increases strictly in A and decreases strictly in k, which increases with length scale. The positive product tail increases with s, so $B_* = p_{c+(r,-r)}$. One corner integration suffices; no computational advantage for the local method is claimed.

F.2 Narrowing, failure case and numerical validation

Use the unchanged settings of Table 1: $d = \varepsilon = 1$, $D = [-.01, .01]^2$, $\alpha = .084$, zero-valued sample means and $t = 1/n$. The first five counts were fixed in advance; $n = 32$ is a secondary check. Ball comparisons verify strict penalty narrowing, direct-control wins throughout, and the information-bound win at $n = 32$. Full enclosures accompany the code.

Computations use 160-bit Arb ball arithmetic and validated integration (Johansson, 2018). Integrate on $[0, \pi/2 - \eta]$ with $\eta = .01$ and add a tail enclosure of radius $\eta e^{-h/\sin\eta}/\pi$. For the gradient integral the radius is $\eta z_0 e^{-z_0}/\pi$, $z_0 = h/\sin\eta > 1$ in every call. Coefficient bounds enclose the entire parameter box, using $w_i = (1 + t/v_i)^{-1}$ to avoid interval dependency. Gaussian KL equals $\frac{1}{2}\sum_i (w_i/w_{i,c} - 1 - \log(w_i/w_{i,c}))$. Corner values give a lower bound on its supremum; an upper bound sums scalar maxima over exact marginal variance ranges. Each scalar maximum occurs at an endpoint because the scalar KL is monotone away from its reference variance. Binary-KL bisection with strict ball comparisons then encloses the optimized comparator. Neither global KL corner optimality nor unvalidated quadrature is assumed.

The zero-data sequence is a conditional illustration of probability zero under continuous noise, not a test of retention coverage or adaptive BO performance. Actual-history narrowing is Corollary 19.

G Random-history robustness diagnostic

Fixed protocol. We generated 100 independent latent RBF draws at true log hyperparameters $(0, 0)$ on designs $\{-0.5, 0.5\}$, with independent unit-variance Gaussian observation noise and nested counts

$n \in \{1, 2, 4, 8, 16, 32\}$ per design (seed 20260911). Recommendation 1, regret threshold 1, risk budget 0.084, and $D = [-0.01, 0.01]^2$ were fixed before outcomes. All histories were retained: the 600 evaluations comprise 100 independent trajectories with paired counts. The box contains the truth by construction; learned retention and anytime stopping are not tested.

Matched calculations. All methods share the posterior history and target. We use the general nonzero-mean Hessian bound and ordinary optimized indicator KL comparator; the centered angular remainder and sign conditioning do not apply. Corner KL values and analytic scalar-range bounds bracket the KL radius. On a fixed 3×3 grid G, direct robust probability is bracketed by

$$L = \max_{z \in G} p_z, \qquad U = \min\left\{1, \max_{z \in G}\left[p_z + \frac{r}{2}\|\nabla p_z\|_1 + \frac{Hr^2}{8}\right]\right\},$$

where $r = 0.01$ and H is the same uniform second-directional bound. The grid covers D within infinity-distance $r/2$, giving $L \le \sup_D p_\theta \le U$ in exact arithmetic; this is a covering bound, not exact optimization.

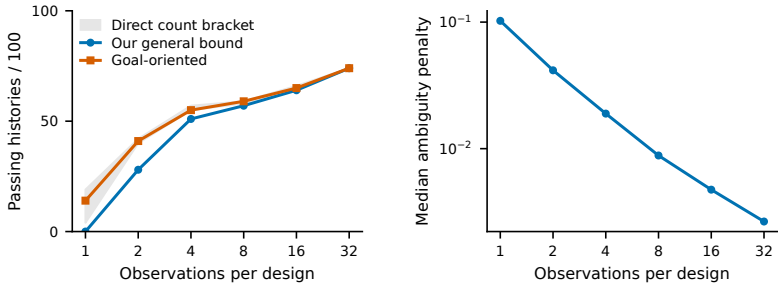

Figure 2: **Narrowing without a comparative advantage on random histories.** Left: passing counts for 100 independent trajectories, paired across observation counts. Shading brackets ideal direct-robust passing counts via the numerical covering bound; it is not a confidence interval. Right: median ambiguity penalty. Floating-point diagnostics; no histories excluded.

Outcomes and numerical scope. The comparator never loses a paired decision; its bracket endpoints agree on all 600 decisions. Direct-robust passing counts lie in $[4, 19]$ at $n = 1$ and $[74, 74]$ at $n = 32$. A lower passing count from a loose covering bound does not imply inferior ideal direct control. The median penalty drops by 97.4%; the Hessian remainder accounts for a median 96.6% at $n = 1$. Quadrature uses tolerance 10^{-8} and standardized limits ± 12. For the center parameter at all six counts in the fixed first ten trajectories (60 evaluations), tolerance 10^{-11} changes probability/gradient values by less than 5×10^{-11}. Tail and rounding errors are not interval-enclosed. Full paired outcomes accompany the code; the diagnostic supports numerical stability and narrowing, not certified calibration or scalable BO performance.

Verification boundary. Earlier supplied Lean audits check selected Gaussian, coefficient, concentration and conditional-probability components; they are not rebuilt here. The Fourier adapter, adaptive-law assembly, and new sharpness, observation and remainder results have written proofs only. Isonormal existence remains an unverified standard dependency. Numerical enclosures do not imply Lean verification of the full theorem or implementation.