Sharp estimation costs for thermal Gaussian energies with known background

Scott Schubr
Luminary Oddities

November 1, 2026

Abstract

We determine the sampling cost of estimating high-order source energies from thermal Gaussian states with a known additive background. The source has unknown occupation covariance above a positive spectral floor, and the known background covariance may fail to commute with it. A finite inverse-heat correction produces a signed observable whose exact coefficient-dependent variance is the minimum among all globally unbiased quantum procedures. Its sharp worst-polynomial envelope gives matching arbitrary-bias minimax bounds. As degree grows, the quantum and standard-heterodyne risks converge to one minus an exponential at their respective resource scales, uniformly along arbitrary known sequences of source floors, attenuation, backgrounds and finite dimensions. The lower proof uses a normalized scalar likelihood whose relative curvature remains controlled through vanishing and diverging occupation regimes. The exact ratio of the two sample-cost coefficients is one plus the inverse output occupation floor. Consequently quantum relative consistency together with heterodyne complete collapse is possible precisely when that floor tends to zero. Known-background subtraction corrects expectations while retaining a quantified fluctuation cost. Scalar and noncommuting two-mode examples illustrate the exact risks and measurement budgets.

Keywords. Thermal Gaussian states; background noise; quantum estimation; minimax risk; heterodyne measurement.
MSC 2020. 62C20, 62F12, 81P50.

1 Introduction

A known thermal background can be removed from an expectation, but its random fluctuations still affect estimation. For high-order source energies, those fluctuations are amplified by the polynomial degree. The resulting sampling cost depends on the observed occupation as well as on the source contribution to it. We determine that cost and compare collective quantum measurements with standard heterodyne.

The observation is a centered thermal Gaussian state with occupation covariance

$$B = \eta A + B_0, \qquad A \succeq aI_d,$$

where $a > 0$, $0 < \eta \leq 1$ and the Hermitian background $B_0 \succeq 0$ are known. The source covariance A is unknown. For a given nonzero homogeneous holomorphic polynomial f of degree n, the target is its normally ordered source energy. We allow arbitrary biased real decisions and measure performance by relative squared error.

Write $c = \eta a$, $b = \|B_0\|$ and $s_0 = c + b$. The sharp quantum and heterodyne sample-cost coefficients are

$$\Gamma_Q = \frac{s_0(1+s_0)}{c^2}, \qquad \Gamma_H = \frac{(1+s_0)^2}{c^2}.$$

Our principal theorem gives the exact limiting minimax risk $1 - e^{-\lambda}$ whenever the corresponding scale $\Gamma_Q n^2/N$ or $\Gamma_H n^2/N$ tends to λ. The theorem holds along arbitrary known sequences of these experiments, including vanishing or diverging occupation and changing finite dimension. The exact coefficient ratio $\Gamma_H/\Gamma_Q = 1 + 1/s_0$ then identifies the measurement-separation boundary.

The upper construction retains more information than these worst-case coefficients: its exact all-order variance depends on the polynomial's reduced tensors and on two noncommuting covariance weights. We prove pointwise minimum variance among procedures unbiased over the full source-covariance model. For the arbitrary-bias lower, a maximal-background eigenmode gives a scalar geometric or exponential likelihood. A normalized source parameter makes both likelihood scores take the same form. Uniform control of their relative curvature pays the critical asymptotics even at the low-brightness boundary.

Prior methods and the increment. Calibrated amplifier-noise convolution and recovery of signal moments are established experimentally and mathematically [2]. Inverse-adjoint correction of expectations, with an inverse that need not be a physical channel, is developed for qubits and multiqubit noise in [3, 4]. Quantum Rao–Blackwell improvement and thermal-light block reduction appear in [6]. We use these established principles together with finite Gaussian moment calculus.

The author's thermal-energy predecessor [5] gives the no-background copy-invariant estimator, coefficient-dependent risk and source-only minimax law. Here finite inverse heat corrects an arbitrary known background; cross contractions retain two distinct matrix weights. The new normalized lower-bound argument treats the shifted source target and gives matching quantum and heterodyne critical laws throughout changing occupation regimes. Gaussian thermometry already exhibits high-temperature near-optimality of Gaussian measurements [1]; the comparison here concerns high-order source polynomial energies and global arbitrary-bias relative risk.

Section 2 states the experiment and main conclusions. Sections 3–4 construct the exact optimum and its sharp envelope. Section 5 proves the common critical lower; Section 6 derives measurement consequences and worked examples.

2 Experiment and main conclusions

For $A \succ 0$, define the centered gauge-invariant Gaussian state

$$\rho_A = \int_{\mathbb{C}^d} |z\rangle\langle z| \frac{e^{-z^* A^{-1} z}}{\pi^d \det A}\, d^{2d}z.$$

Thus $Z \sim \mathrm{CN}(0, A)$ has $\mathbb{E}ZZ^* = A$ and the annihilators obey $\mathbb{E}b_i^\dagger b_j = A_{ji}$. For symmetric coefficients c_I, let $f(z) = \sum_{I\in[d]^n} c_I \prod_{j=1}^n z_{i_j} \neq 0$ and put

$$F_f(A) = \mathbb{E}|f(Z)|^2 = \mathrm{Tr}\,\rho_A f(b)^\dagger f(b) > 0. \tag{1}$$

Observe N independent copies of ρ_B, where

$$B = \eta A + B_0, \quad A \succeq aI_d, \quad a > 0, \quad 0 < \eta \leq 1, \quad B_0 \succeq 0. \tag{2}$$

All channel parameters and f are known. Define

$$R_Q = \sup_{0\neq f\in \mathrm{Sym}^n(\mathbb{C}^d)} \inf_{\mathsf{M},Y} \sup_{A\succeq aI_d} \mathbb{E}_A \left(\frac{Y}{F_f(A)} - 1\right)^2, \tag{3}$$

where the infimum is over collective POVMs on the N source copies and real decisions. Define R_H by restricting the measurement to standard heterodyne of each copy and allowing arbitrary subsequent decisions. The optimization order permits the procedure to depend on the known polynomial and experiment. Put

$$c = \eta a, \quad b = \|B_0\|, \quad s_0 = c + b, \qquad \Gamma_Q = \frac{s_0(1+s_0)}{c^2}, \quad \Gamma_H = \frac{(1+s_0)^2}{c^2}. \tag{4}$$

For $r > 0$ write

$$S_r(N,n) = \sum_{k=1}^{n} \binom{n}{k}^2 \frac{k! r^k}{(N)^{\overline{k}}}, \qquad (N)^{\overline{k}} = \prod_{j=0}^{k-1}(N+j), \quad (N)^{\overline{0}} = 1. \tag{5}$$

Theorem 2.1 (Sharp envelopes and minimax order)**.** *The exact worst-polynomial minimum globally unbiased relative variances are S_{Γ_Q} and S_{Γ_H} in the two experiments. For $\nu \in \{Q, H\}$ and $x_\nu = \Gamma_\nu n^2/N$,*

$$\frac{1}{512}\min(1, x_\nu) \le R_\nu \le \frac{S_{\Gamma_\nu}}{1+S_{\Gamma_\nu}} \le 1 - e^{-x_\nu} \le \min(1, x_\nu). \tag{6}$$

The unbiased variance envelope is attained by a scalar monomial on a maximal-background eigenmode with $A = aI_d$.

Theorem 2.2 (Exact critical laws along changing experiments)**.** *For each degree n, let $d_n \ge 1$, $a_n > 0$, $0 < \eta_n \le 1$ and a known finite Hermitian $B_{0,n} \succeq 0$ be arbitrary. Let $N_n \ge 1$ and $n \to \infty$. For either $\nu = Q$ or H, if*

$$x_{\nu,n} = \Gamma_{\nu,n}\frac{n^2}{N_n} \to \lambda \in [0, \infty), \tag{7}$$

then

$$R_{\nu,n} \to 1 - e^{-\lambda}. \tag{8}$$

If $x_{\nu,n} \to \infty$, then $R_{\nu,n} \to 1$. For arbitrary such sequences, $R_{\nu,n} \to 0$ if and only if $x_{\nu,n} \to 0$.

The dimension may change because the upper bound is uniform and the lower experiment uses one informative mode. The theorem applies separately at each measurement's scale. Its sharp coefficients give

$$\frac{\Gamma_H}{\Gamma_Q} = 1 + \frac{1}{s_0}. \tag{9}$$

In particular there exists a copy-budget sequence with quantum relative consistency and heterodyne risk tending to one precisely when $s_{0,n} \to 0$. Section 6 makes the quantifiers and critical-risk comparisons explicit.

3 A corrected observable and its exact risk

We first explain the correction at degree one. In a scalar mode with output occupation $s = \eta t + b$, let K be the total photon count. Then

$$\widehat{t} = \frac{K/N - b}{\eta}, \qquad \mathbb{E}\widehat{t} = t, \qquad \frac{\operatorname{Var} \widehat{t}}{t^2} = \frac{s(1+s)}{N(\eta t)^2}. \tag{10}$$

The corrected outcome can be negative. Known subtraction removes its mean contamination while the output occupation controls its fluctuations.

3.1 Finite inverse heat for all degrees

For a formal amplitude matrix $Z = (z_{i,r}) \in \mathbb{C}^{d\times N}$, define

$$P_f(Z) = \frac{n!}{(N)^{\overline{n}}} \sum_{\boldsymbol{r}\in[N]^n} \left| \sum_I c_I \prod_{j=1}^n z_{i_j,r_j} \right|^2. \tag{11}$$

This is also the complex-unit-sphere average of $|f(Zv)|^2$, by the sphere moment identity used in [5]. For a known covariance $D \succeq 0$ let

$$\Delta_D = \sum_{r=1}^N \sum_{i,j=1}^d D_{ji} \frac{\partial^2}{\partial \bar{z}_{i,r} \partial z_{j,r}}. \tag{12}$$

Its exponential on P_f is a finite polynomial. Normal-order quantization replaces $\bar{z}^p z^q$ by $(b^\dagger)^p b^q$. Put

$$K_f = \eta^{-n} : e^{-\Delta_{B_0}} P_f : . \tag{13}$$

It depends only on known quantities. Every term has equal total creator and annihilator degrees, so K_f preserves total photon number. Its finite Hermitian number-sector blocks define a self-adjoint direct sum with finite-particle core. Polynomial Gaussian moments give finite second moments at every allowed covariance.

For the exact variance choose a mode basis diagonalizing $C = \eta A$, with eigenvalues $c_i > 0$, and transform f into that basis. Define

$$\phi_I = \overline{c_I} \prod_{j=1}^n \sqrt{c_{i_j}}, \quad \psi = \phi/\|\phi\|, \quad R_k = \operatorname{Tr}_{n-k} |\psi\rangle\langle\psi|, \tag{14}$$

so $F_f(C) = n!\|\phi\|^2$. Also put

$$H = C^{-1/2} B C^{-1/2}, \qquad J = C^{-1/2}(B+I)C^{-1/2}, \qquad H_k = H^{\otimes k}, \quad J_k = J^{\otimes k}. \tag{15}$$

These quantities specify the variance, not the measurement construction. The conjugation in (14) matches $\mathbb{E} b_i^\dagger b_j = B_{ji}$; with unconjugated coefficient reductions both matrix weights would instead be transposed.

Theorem 3.1 (Exact corrected variance)**.** *The observable K_f is unbiased for $F_f(A)$ and has relative variance*

$$V_Q(A,f) = \sum_{k=1}^n \binom{n}{k}^2 \frac{k!}{(N)^{\overline{k}}} \operatorname{Tr}(R_k H_k R_k J_k). \tag{16}$$

Its spectral measurement has minimum variance at each $A \succeq aI_d$ among all real-decision POVMs integrable and unbiased throughout that source family. For heterodyne amplitude records, the statistic

$$U_f(Z) = \eta^{-n} e^{-\Delta_{B_0+I}} P_f(Z) \tag{17}$$

is the pointwise minimum-variance globally unbiased statistic, with exact relative variance

$$V_H(A, f) = \sum_{k=1}^{n} \binom{n}{k}^2 \frac{k!}{(N)_{\overline{k}}} \operatorname{Tr}(R_k J_k R_k J_k). \tag{18}$$

Unbiasedness and variance. Gaussian covariance addition acts on polynomials by the heat semigroup. Thus

$$\mathbb{E}_B e^{-\Delta_{B_0}} P_f = \mathbb{E}_{B-B_0} P_f = F_f(C) = \eta^n F_f(A).$$

The coherent lower symbol of K_f is the corrected polynomial, so the positive Gaussian coherent mixture proves quantum unbiasedness. For heterodyne, records have covariance $B + I$; the same identity proves (17) unbiased.

For the quantum second moment use two sphere directions v, w, with copy overlap $s = v^* w$. Normal ordering in the product of corrected observables leaves internal contractions C in each factor, while cross contractions are B in the creator-to-annihilator direction and $B + I$ in the other direction. This can be checked directly by four formal source vectors: before coefficient extraction their generating exponential has internal forms $x^T C^T y$ and $z^T C^T w$, and cross forms $x^T B^T w$ and $y^T (B + I) z$, with the conjugate overlap factors. Inverse heat subtracts B_0 from the internal forms only.

Balance gives k cross contractions in each direction. Their multiplicity divided by $(n!)^2$ is $\binom{n}{k}^2$. Internal C contractions produce R_k, and whitening yields

$$\operatorname{Tr}(R_k H_k R_k J_k) = \|H_k^{1/2} R_k J_k^{1/2}\|_{\mathrm{HS}}^2 \geq 0.$$

The sphere overlap moment is $\mathbb{E}|v^* w|^{2k} = k!/(N)_{\overline{k}}$. The zero-cross term is the squared mean. Summing and subtracting it proves (16). The ordered contractions and the conjugation convention fix the transpose positions also for complex noncommuting matrices. All expansions are finite, and thermal polynomial moments justify the integrations.

For heterodyne, the inverse heat again changes internal contractions to C. Both cross directions have covariance $B + I$, giving (18). The optimality assertions are proved next. □

4 Optimality and sharp uniform envelopes

4.1 Completeness on a translated covariance cone

We give the statistical step beyond symmetry averaging [6]. It is the translated-cone version of the thermal completeness argument in [5].

Lemma 4.1. *Let $D \succ 0$ be fixed. If a self-adjoint, total-number-block-diagonal observable K is invariant under passive copy $U(N)$ mixing, integrable in $\rho_B^{\otimes N}$ for every $B \succeq D$, and has zero expectation throughout that family, then $K = 0$.*

Proof. The family contains $B = tI_d$ for all sufficiently large t. With $q = t/(1 + t)$, isotropic integrability gives $\sum_m q^m \operatorname{Tr}|K_m| < \infty$ on total-number blocks; hence $\|K_m\| \leq C_q q^{-m}$. For coherent

amplitude matrices, this bound makes $e^{\|Z\|_F^2}\langle Z|K|Z\rangle$ entire in the independent variables $\bar{Z}, Z$, since each total-degree-m bound has the factorial denominator $m!$.

Copy invariance makes the continuous coherent symbol a measurable function $g(ZZ^*)$, including deficient rank. Positive coherent mixtures give $\mathbb{E}_B g(ZZ^*) = 0$ and $\mathbb{E}_B|g(ZZ^*)| < \infty$. The scatter laws for N proper complex Gaussian columns form a complete exponential family when the positive Hermitian covariance varies over any open set. Explicitly, relative to an interior covariance the likelihood is a constant times $\exp[-\operatorname{Tr}((B^{-1} - B_*^{-1})ZZ^*)]$. Zero expectations in a neighborhood make the Laplace transform of the corresponding finite signed measure vanish there; uniqueness gives the zero measure. This argument remains valid for singular scatter when $N < d$, because the likelihoods have common support. Therefore the symbol is zero almost surely for a Gaussian law with positive density on amplitude space. Continuity makes it zero everywhere; uniqueness of the entire coefficient expansion identifies every Fock matrix entry as zero. □

Optimality in Theorem 3.1. Average an integrable competitor POVM over passive copy $U(N)$. Every thermal product is invariant, so its outcome law is unchanged at each allowed covariance. The central phase removes cross-total-number first-moment blocks. Isotropic thermal integrability controls each finite block; positivity and Cauchy–Schwarz control off-diagonal first moments. The averaged first blocks define a self-adjoint direct sum with geometrically bounded blocks as in Lemma 4.1.

For integrability, first use an allowed isotropic state, whose weight is scalar on each number sector. The ordinary finite-dimensional trace-norm triangle inequality bounds $\operatorname{Tr}|K_m|$ by the competitor's absolute first moment traced on that sector. Summing with the scalar thermal weights gives integrability in every sufficiently hot isotropic state. For any given allowed B, choose such a state with geometric parameter $q \in (0,1)$ satisfying $q > \|B(I+B)^{-1}\|$. Then $\rho_B^{\otimes N}$ is dominated by a finite multiple of that isotropic state, which proves integrability of the direct sum at B. Equality of expectations follows from the actual averaged block-diagonal POVM, Fubini's theorem for its absolute first moment, and finite-number cutoffs. Subtract K_f. The difference has zero expectations on $B \succeq B_0 + cI_d$, hence is zero by Lemma 4.1.

On each finite number sector, matrix Cauchy–Schwarz for the averaged POVM gives its second-moment matrix at least the square of its first moment, which is the corresponding block of K_f. Multiply by the positive thermal block and sum increasingly over sectors. This proves the second-moment comparison even for infinite competitor variance. The means agree, so variance is minimized by K_f.

For heterodyne, condition any integrable unbiased decision on the complete scatter ZZ^*. The conditional distribution of amplitude records given scatter is covariance-independent, including deficient rank. The corrected polynomial (17) is copy-unitary invariant and thus a scatter function. The same open-parameter Laplace completeness argument identifies the conditioned first moment with it. Conditional Jensen gives the variance comparison. □

4.2 Sharp envelopes and finite minimax bounds

Since $C \succeq cI_d$, $B_0 \preceq bI_d$, put $\alpha = 1 + b/c$, $\beta = 1 + (b+1)/c$. Then

$$H \preceq \alpha I, \quad J \preceq \beta I, \quad \operatorname{Tr}(R_k H_k R_k J_k) \le \alpha^k \beta^k \operatorname{Tr} R_k^2 \le \Gamma_Q^k.$$

Here positive order is applied inside R_k and then paired with a positive matrix; commutation is unnecessary. The same argument bounds each heterodyne contraction by $\beta^{2k} = \Gamma_H^k$. Choose $A = aI_d$ and $f(z) = (u^*z)^n$, with u a unit top eigenvector of B_0. Its reductions are pure on that mode and attain every term. This proves the unbiased envelopes.

Measure $K_f/(1+S_{\Gamma_Q})$, or divide the heterodyne statistic by $1+S_{\Gamma_H}$. For an unbiased normalized statistic with variance at most S, this shrinkage has risk at most $S/(1+S)$. Also $S_r \le e^{rn^2/N}-1$, giving the upper in (6).

The lower needs only the same informative mode. Write its source occupation as ax, $x \ge 1$, and define

$$s(x)=b+cx, \qquad w(x)=1+b+cx, \qquad J_Q(x)=\frac{s(x)(1+s(x))}{c^2}, \quad J_H(x)=\frac{w(x)^2}{c^2}. \tag{19}$$

The target is a known positive constant times x^n. Other modes are held fixed and are independent of x. Quantum source states commute along this subfamily; total photon count is sufficient. For either scalar experiment,

$$\frac{d}{dh}D(P_1\|P_{1+h})=\frac{h}{J_\nu(1+h)}\le\frac{h}{J_\nu(1)}. \tag{20}$$

For geometric counts this follows by differentiating their relative entropy; for heterodyne the scalar record is complex normal of variance $w(x)$ and the same differentiation applies.

Take $h=\frac{1}{4}\min(n^{-1},\sqrt{\Gamma_\nu/N})$. Pinsker's inequality and (20) give total variation at most $1/8$. The target harmonic separation is

$$\Delta=\tanh\Big(\frac{n}{2}\log(1+h)\Big)\ge\frac{nh}{5},$$

because $h\le 1/4$, $\log(1+h)\ge 4h/5$ and $\tanh z\ge z/2$ for $0\le z\le 1/8$. For any two positive targets θ_0,θ_1 with observation laws p_0,p_1,

$$R\ge\left(\frac{\theta_1-\theta_0}{\theta_1+\theta_0}\right)^2(1-\mathrm{TV}(p_0,p_1)). \tag{21}$$

Indeed prior weights $\theta_i/(\theta_0+\theta_1)$ give the squared harmonic separation as the minimum common-density weighted loss. Multiply by $\min(p_0,p_1)$ and integrate. Our pair therefore gives $R_\nu \ge (7/3200)\min(1,x_\nu)$, which implies the stated $1/512$ bound. This proves Theorem 2.1.

5 The uniform critical lower bound

We prove Theorem 2.2 without fixing the limiting occupation. First, for $r_n=\Gamma_{\nu,n}$ and $x_n=r_nn^2/N_n$ bounded, the kth envelope summand is

$$\left[\left(\frac{n(n-1)\cdots(n-k+1)}{n^k}\right)^2\frac{N_n^k}{(N_n)^{\overline{k}}}\right]\frac{x_n^k}{k!}.$$

Both coefficients Γ_Q,Γ_H are at least one, so $N_n\to\infty$. The bracket tends to one for fixed k, and summands are dominated by $C^k/k!$. Thus $S_{r_n}\to e^\lambda-1$, proving the critical upper. At $\lambda=0$ this already proves the result.

5.1 A common normalized likelihood

Assume $x_n\to\lambda\in(0,\infty)$ and suppress n on known parameters. Use the hard eigenmode in (19). In the quantum experiment the photon total K has negative-binomial mean $Ns(x)$ and variance

$Ns(x)(1+s(x))$. In the heterodyne experiment $T = \sum_r |Z_r|^2$ has mean $Nw(x)$ and variance $Nw(x)^2$. The normalized unconstrained maximum-likelihood estimates are, respectively,

$$\widehat{x}_Q = \frac{K/N - b}{c}, \qquad \widehat{x}_H = \frac{T/N - 1 - b}{c}.$$

They obey $\mathbb{E}_x \widehat{x} = x$ and $\mathrm{Var}_x\, \widehat{x} = J_\nu(x)/N$. Ignoring data-only factors, the likelihoods are

$$L_Q(x) = s(x)^K (1+s(x))^{-(N+K)}, \qquad L_H(x) = w(x)^{-N} e^{-T/w(x)}.$$

In both cases the exact score is

$$\frac{d}{dx} \log L_\nu(x) = \frac{N(\widehat{x} - x)}{J_\nu(x)}. \tag{22}$$

On $x \in [1/2, 2]$, uniformly in all known channel parameters,

$$0 \le \frac{d}{dx} \log J_\nu(x) \le \frac{2}{x}, \qquad \frac{1}{4} \le \frac{J_\nu(x)}{J_\nu(1)} \le 4. \tag{23}$$

Each factor in (19) is a positive affine function of x, which proves these bounds directly. Thus the normalization controls both low and high brightness.

5.2 Supported prior and tilted likelihood integrals

Give x the uniform prior on $I_n = [1+\delta, 1+3\delta]$, $\delta = n^{-1/4}$. This is wholly inside the source-floor constraint. Put $r = n^{-4/5}$. The prior mass within $3r$ of an endpoint is $O(r/\delta) \to 0$. From (23), uniformly on the prior,

$$\Pr_x(|\widehat{x} - x| > r) \le \frac{C J_\nu(1)}{N r^2} = O((n^2 r^2)^{-1}) \to 0.$$

Hence with prior-predictive probability tending to one, $\widehat{x}$ is at least r inside I_n and $\widehat{x} = 1 + O(\delta)$. All following relative estimates hold uniformly on this event.

Integrating (22) and using (23) gives, for $h = x - \widehat{x}$,

$$\log \frac{L(\widehat{x}+h)}{L(\widehat{x})} = -\frac{N h^2}{2 J_\nu(\widehat{x})} + O\left(\frac{N|h|^3}{J_\nu(\widehat{x})}\right). \tag{24}$$

The exact integral also bounds this log ratio above by $-c_* N h^2 / J_\nu(1)$ on the entire shrinking prior, with a universal positive c_*. This avoids separate bounds on large canceling derivatives.

Let $I_j = \int_{I_n} L(x) x^{-jn}\, dx$, $j = 0, 1, 2$. On $|h| \le r$, the error in (24) is $O(n^2 r^3) = O(n^{-2/5})$. Also

$$-jn \log(\widehat{x} + h) = -jn \log \widehat{x} - \frac{jnh}{\widehat{x}} + O(nh^2),$$

with error $O(nr^2) = O(n^{-3/5})$. Completing the square yields

$$I_j = L(\widehat{x}) \widehat{x}^{-jn} \sqrt{\frac{2\pi J_\nu(\widehat{x})}{N}} \exp\left(\frac{j^2 n^2 J_\nu(\widehat{x})}{2N\widehat{x}^2}\right)(1 + o(1)). \tag{25}$$

We justify the full integral in this formula. Outside $|h| \le r$, the likelihood and target are bounded by their common prefactor times

$$\exp\left(-c_* \frac{N}{J_\nu(1)} h^2 + Cn|h|\right).$$

Since $N/J_\nu(1) \asymp n^2$, the linear term is absorbed for $|h| \geq r$; the resulting integral is at most $Ce^{-c'n^2r^2}$ times the prefactor, negligible relative to the central factor of order $1/n$. The completed Gaussian's shift is $O(1/n) = o(r)$ and its standard deviation is $O(1/n)$, so its replacement by a full-line integral is valid too. Thus (25) is a relative asymptotic for the actual supported integrals.

For target x^n, the posterior minimum relative squared risk is

$$1 - \frac{I_1^2}{I_0 I_2}.$$

The explained-risk ratio lies in $[0, 1]$ by Cauchy–Schwarz. Formula (25) gives

$$\frac{I_1^2}{I_0 I_2} = \exp\left(-\frac{n^2 J_\nu(\widehat{x})}{N\widehat{x}^2}\right)(1 + o(1)) \longrightarrow e^{-\lambda}$$

in prior-predictive probability. Boundedness gives convergence in expectation, so the Bayes lower matches the uniform upper. For quantum observations, the commuting eigenmode states make photon counts sufficient against every collective POVM; for heterodyne, the scalar exponential sufficient statistic retains all record information. Other fixed modes are ancillary. This proves (8).

5.3 Collapse and consistency

If $x_n \to \infty$, use the hard pair $x = 1, 1 + h_n$ with $h_n = x_n^{1/4}/n$. Equation (20) holds for every $h_n \geq 0$, giving $ND(P_1 \| P_{1+h_n}) \leq 1/(2\sqrt{x_n}) \to 0$. The target ratio satisfies $(1 + h_n)^n \geq 1 + nh_n = 1 + x_n^{1/4} \to \infty$. Equation (21) therefore gives risk tending to one. The zero decision supplies the upper one. Finally, the two-sided finite bounds (6) prove the consistency equivalence. This completes Theorem 2.2.

6 Measurement boundary and worked examples

Equation (9) gives an exact comparison of the critical copy budgets. At fixed nonzero critical exponent, achieving the same limiting risk requires the heterodyne budget to be larger by $1 + 1/s_0$. At a common budget with $x_Q \to \lambda \in (0, \infty)$ and $s_0 \to s \in (0, \infty)$,

$$R_Q \to 1 - e^{-\lambda}, \qquad R_H \to 1 - e^{-(1+1/s)\lambda}.$$

If instead $s_0 \to \infty$, both risks tend to $1 - e^{-\lambda}$. If $s_0 \to 0$, heterodyne risk tends to one at that same quantum critical budget.

For complete order separation, assume $s_{0,n} \to 0$ and choose

$$N_n = \left\lceil n^2 \sqrt{\Gamma_{Q,n}\Gamma_{H,n}} \right\rceil.$$

Then $x_Q \to 0$, $x_H \to \infty$: quantum relative consistency and heterodyne complete collapse coexist. Conversely, those two risk conclusions force $x_Q \to 0$ and $x_H \to \infty$ by Theorem 2.2 and subsequence extraction. Hence their ratio diverges, which is exactly $s_{0,n} \to 0$. A background bounded below prevents this growing order separation. A vanishing background can preserve it even when it dominates the source: $b_n \gg c_n$, $b_n \to 0$ gives coefficient ratio asymptotic to $1/b_n$.

6.1 Scalar check and critical budgets

Take $\eta = 1$, source floor $a = 1/2$ and scalar background $b = 1/2$. At the floor $s_0 = 1$, (10) gives quantum relative variance $8/N$. Heterodyne uses the mean of $|Z|^2 - 1 - b$, giving relative variance $16/N$. These are exact finite-sample degree-one variances. For growing degree, the equal-risk critical copy-budget ratio is two. More generally output floors $s_0 = 0.01, 1, 100$ give ratios $101, 2, 1.01$, respectively. These latter values compare asymptotic sample-cost coefficients rather than finite-degree risk ratios.

6.2 Complex noncommuting degree-two check

Let $\eta = 1$, $A = \mathrm{diag}(1, 4)$,

$$B_0 = \begin{pmatrix} 1 & i/3 \\ -i/3 & 2 \end{pmatrix}, \qquad f(z) = z_1^2 + (1+i)z_1 z_2 + (2-i)z_2^2.$$

The source and background covariances do not commute, and B_0 is positive definite. The source energy is $F_f(A) = 170$. Formula (16) gives relative variances

$$V_Q = \frac{360092}{21675} \quad (N = 1), \qquad V_Q = \frac{467248}{65025} \quad (N = 2).$$

Direct normal-ordered multiplication of the inverse-heat observable gives the same rational values. This example tests the complex conjugation and two-sided matrix contraction rather than replacing the background by its operator norm.

Reproducibility. The accompanying Python standard-library verifier checks the corrected operator moments against the trace formula in twenty exact real and complex noncommuting two-mode cases, and checks both normalized score identities in seventy-two rational parameter cases. A second verifier checks the normalized likelihood identities and relative-curvature bounds over a broader rational grid. The proofs establish the uniform phase theorem; finite algebraic controls support its formulas.

Acknowledgment. OpenAI Codex was used for mathematical, computational and editorial assistance in preparing this work.

References

[1] M. F. B. Cenni, L. Lami, A. Acín and M. Mehboudi, *Thermometry of Gaussian quantum systems using Gaussian measurements*, Quantum **6** (2022), 743. https://doi.org/10.22331/q-2022-06-23-743.

[2] C. Eichler, D. Bozyigit, C. Lang, L. Steffen, J. Fink and A. Wallraff, *Experimental state tomography of itinerant single microwave photons*, Phys. Rev. Lett. **106** (2011), 220503. https://doi.org/10.1103/PhysRevLett.106.220503.

[3] S. Mangini, L. Maccone and C. Macchiavello, *Qubit noise deconvolution*, EPJ Quantum Technol. **9** (2022), 29. https://doi.org/10.1140/epjqt/s40507-022-00151-0.

[4] S. Roncallo, L. Maccone and C. Macchiavello, *Multiqubit noise deconvolution and characterization*, Phys. Rev. A **107** (2023), 022419. https://doi.org/10.1103/PhysRevA.107.022419.

[5] S. Schubr, *Optimal unbiased estimation of polynomial energies in thermal Gaussian states*, author-published edition 3, Luminary Oddities (2026). https://research.luminaryoddities.com/finite-copy-risk-polynomial-energies-thermal-gaussian-states/v/3.

[6] M. Tsang, *Operational meanings of a generalized conditional expectation in quantum metrology*, Quantum **7** (2023), 1162. https://doi.org/10.22331/q-2023-11-03-1162.