Sharp calibration costs for high-order thermal Gaussian energies

Scott Schubr
Luminary Oddities

November 1, 2026

Abstract

We determine the joint cost of source measurements and loss calibration when estimating high-order energies of thermal Gaussian states. The source has unknown occupation covariance bounded below by a positive constant, and a common unknown transmissivity lies in a fixed interior interval. We observe independent source copies and single-photon calibration trials. As the polynomial degree grows with the square roots of both sample sizes, the worst-polynomial relative minimax risk converges to one minus the exponential of a sum of source and calibration costs. The result allows arbitrary biased decisions and collective source measurements, and is uniform in mode dimension and polynomial coefficients. Finite-sample bounds also identify the source and calibration budgets uniformly over the occupation floor, with a matched comparison to standard heterodyne. A reciprocal-binomial correction gives an attaining estimator whose moments remain controlled even at zero calibration counts. A supported product prior gives the matching quantum lower bound. The exact law yields the optimal deterministic allocation under specified sample costs. We also determine the complementary limit when calibration is supplied only as an interval: finite-copy uncertainty combines with an exact identification floor through a hyperbolic-secant factor. These results quantify the calibration precision needed to estimate growing-order source energies.

Keywords. Quantum estimation; thermal Gaussian states; loss calibration; minimax risk; high-order moments.
MSC 2020. 62C20, 62F12, 81P50.

1 Introduction

Estimating a source energy from attenuated light requires knowledge of both the observed state and the attenuation. For a homogeneous energy of degree n, an error in log transmissivity is multiplied by n. Calibration precision that suffices for a first moment may therefore be inadequate for a high-order energy, even when many source copies are available.

We study this effect in a specified two-resource experiment. An unknown thermal Gaussian source has occupation covariance $A \succeq aI$, where $a > 0$ is known. Pure loss with unknown transmissivity $\eta \in [l, h] \subset (0, 1)$ gives observed occupation covariance ηA. There are N source copies and M independently prepared single-photon calibration probes. Counting transmitted calibration photons produces a binomial observation. The target is a known homogeneous polynomial energy of the source, and loss is squared relative error.

Our principal result is an exact global minimax limit. If $n^2/N \to \kappa_s$ and $n^2/M \to \kappa_c$, the risk tends to

$$1 - \exp\left\{ -\left(1 + \frac{1}{la}\right)\kappa_s - \frac{1-l}{l}\kappa_c \right\}.$$

The upper bound holds for every dimension and polynomial. The lower bound uses scalar thermal sources and applies to all collective source measurements and arbitrary biased decisions. Thus the formula is a finite-risk conclusion at the critical sample scale, rather than only a local variance expansion. Its coefficients give the optimal deterministic division of a specified source/calibration budget.

The construction uses an existing optimal unbiased output-energy observable and a nonsingular reciprocal-binomial correction. Its first scaled calibration moment is exactly a binomial-tail probability. Uniform second-moment analysis pays the rare zero-count event, where a naive inverse empirical efficiency would be singular. The lower bound uses independent supported priors in observed occupation and transmissivity. Their support is chosen to retain the source floor after division by transmissivity; this makes the posterior factorization compatible with the original physical parameters.

Relation to previous work. The loss of identifiability between thermal brightness and detector efficiency is established in self-calibration studies [3]. Bayesian propagation of uncertain attenuation and detector parameters appears in [1]. Pure-loss calibration with number-state probes and its local precision theory are developed in [4]. We use independent one-photon trials as the specified calibration experiment.

Reciprocal-binomial regularization and inverse factorial moments have established statistical antecedents [7, 5]. Resource allocation for Bernoulli estimation, including transformed targets, is studied in [2]; the square-root allocation algebra is inherited. These ingredients explain the form of the calibration cost. The contribution here is the matched global high-degree source-energy risk, including its uniform quantum upper, feasible boundary prior lower and attaining correction.

The author's thermal-energy predecessor [6] supplies the output observable, its sharp variance envelope and the source-only critical law. The present paper treats the additional unknown source-to-output scale and pays for its calibration. Section 7 gives a complementary experiment with interval information and no calibration observations, deriving its exact identification floor and joint critical limit. The main proofs are in Sections 3–4; allocation and worked comparisons follow in Section 6.

2 Experiment and principal results

For $A \succ 0$, let

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

be the centered gauge-invariant thermal Gaussian state. Thus $Z \sim \mathrm{CN}(0, A)$ has $\mathbb{E} Z Z^* = A$. Write a nonzero homogeneous holomorphic degree-n polynomial using a symmetric coefficient tensor as $f(z) = \sum_{I \in [d]^n} c_I z_{i_1} \cdots z_{i_n}$. The source target is

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

In particular $F_f(\eta A) = \eta^n F_f(A)$. The operators b_i are canonical annihilators with $[b_i, b_j^\dagger] = \delta_{ij}$.

Fix $a > 0$ and $0 < l < h < 1$. The data consist of N copies of $\rho_{\eta A}$, $A \succeq a I_d$, and an independent count $L \sim \mathrm{Bin}(M, \eta)$ from M one-photon calibration probes through the same pure-loss channel. A procedure may choose any collective source POVM conditional on L and any real decision Y.

Define

$$R_{N,M,n} = \sup_{d\geq 1} \sup_{0\neq f\in \mathrm{Sym}^n(\mathbb{C}^d)} \inf_{\mathsf{M},Y} \sup_{\substack{A\succeq aI_d \\ \eta\in[l,h]}} \mathbb{E}_{A,\eta}\left(\frac{Y}{F_f(A)} - 1\right)^2. \tag{2}$$

The estimator knows a, l, h, f, d, N, M, n. Source copies and calibration probes are separate resources. Put

$$\gamma_l = 1 + \frac{1}{la}, \qquad d_l = \frac{1-l}{l}. \tag{3}$$

Theorem 2.1 (Exact source/calibration critical risk)**.** *Let $n \to \infty$ and let N, M be positive integers with*

$$\frac{n^2}{N} \to \kappa_s \in [0,\infty), \qquad \frac{n^2}{M} \to \kappa_c \in [0,\infty).$$

Then the minimax risk in (2) *satisfies*

$$R_{N,M,n} \longrightarrow 1 - e^{-\gamma_l\kappa_s - d_l\kappa_c}. \tag{4}$$

If either $n^2/N \to \infty$ or $n^2/M \to \infty$, then $R_{N,M,n} \to 1$. Consequently, for arbitrary positive integer sequences N, M with $n \to \infty$,

$$R_{N,M,n} \to 0 \quad \Longleftrightarrow \quad \frac{n^2}{N} \to 0 \text{ and } \frac{n^2}{M} \to 0. \tag{5}$$

Theorem 2.2 (Finite-sample resource bounds)**.** *For fixed $0 < l < h < 1$, there are constants $c_{l,h}, C_l > 0$ such that, for every $a > 0$ and positive integers n, N, M,*

$$c_{l,h}\min\left(1, \gamma_l\frac{n^2}{N} + \frac{n^2}{M}\right) \leq R_{N,M,n} \leq C_l\min\left(1, \gamma_l\frac{n^2}{N} + \frac{n^2}{M}\right). \tag{6}$$

For standard heterodyne source measurements with the same calibration experiment, the corresponding risk $R^{\mathrm{H}}_{N,M,n}$ obeys the same two-sided order with γ_l replaced by γ_l^2, with constants depending only on l, h.

Here l, h are fixed and have nonzero separation. The exact dependence in (4) identifies the hardest source/calibration corner: the source floor and the lower transmissivity. The arbitrary-bias class is essential even at fixed sample size.

Proposition 2.3 (Finite calibration and exact unbiasedness)**.** *For every $n \geq 1$, $N \geq 1$ and finite $M \geq 0$, no integrable procedure is unbiased for the source target throughout the experiment* (2)*.*

Proof. Use $d = 1$, $f(z) = z^n$, and fix an observed occupation $s \geq ha$. Set the source occupation to s/η. The source state is then $\rho_s^{\otimes N}$ for every $\eta \in [l, h]$, and every source occupation remains at least a. Conditional on each $L = j$, an integrable procedure has a finite mean m_j independent of η. Indeed all binomial weights are positive on the interval, so integrability implies each conditional absolute mean is finite. The unconditional mean $\sum_{j=0}^{M} m_j\binom{M}{j}\eta^j(1-\eta)^{M-j}$ is a polynomial of degree at most M. Unbiasedness would make it equal to $n!s^n\eta^{-n}$ on an interval. Multiplication by η^n would give a polynomial with zero constant term identically equal to a positive constant, a contradiction. □

3 An attaining estimator and its uniform moments

3.1 The inherited output-energy statistic

For copy indices $r \in [N]$, let $b_{i,r}$ be the annihilator in mode i of copy r. Set

$$C_{\boldsymbol{r}} = \sum_{I} c_I \prod_{j=1}^{n} b_{i_j,r_j}, \qquad T_f = \frac{n!}{(N)^{\overline{n}}} \sum_{\boldsymbol{r}\in[N]^n} C_{\boldsymbol{r}}^{\dagger} C_{\boldsymbol{r}}, \tag{7}$$

where $(N)^{\overline{k}} = N(N+1)\cdots(N+k-1)$ and $(N)^{\overline{0}} = 1$. This positive observable preserves total photon number and is defined by its finite-dimensional Hermitian blocks on each total-number sector. Its spectral measurement supplies the source statistic used below.

The thermal theorem [6, Theorems 2.1–2.4] gives, at output covariance $B \succeq laI$,

$$\mathbb{E}_B T_f = F_f(B), \qquad \frac{\mathrm{Var}_B(T_f)}{F_f(B)^2} \le S_N, \quad S_N = \sum_{k=1}^{n} \binom{n}{k}^2 \frac{k!\gamma_l^k}{(N)^{\overline{k}}}. \tag{8}$$

The displayed bound is uniform in d, f, B and is sharp for a scalar monomial at $B = la$. It also gives $S_N \le e^{\gamma_l n^2/N} - 1$. If $n^2/N \to \kappa_s < \infty$, then

$$1 + S_N \to e^{\gamma_l \kappa_s}. \tag{9}$$

To see the limit directly, each fixed-k summand tends to $(\gamma_l\kappa_s)^k/k!$, and the positive summands are bounded by $C^k/k!$ when $\gamma_l n^2/N \le C$. This proves dominated summation, including $\kappa_s = 0$.

3.2 Reciprocal-binomial correction

Define the finite correction

$$V_{M,n}(L) = \prod_{j=1}^{n} \frac{M+j}{L+j}. \tag{10}$$

Its denominator stays positive even at $L = 0$. The following critical-scale version of inverse-factorial moment calculus [7, 5] is the estimate needed for source calibration.

Lemma 3.1 (Uniform calibration moments)**.** *Let $W = \eta^n V_{M,n}(L)$ with $L \sim \mathrm{Bin}(M, \eta)$. Exactly,*

$$\mathbb{E}_\eta W = \Pr\{\mathrm{Bin}(M+n, \eta) \ge n\}. \tag{11}$$

If $n \to \infty$ and n^2/M is bounded, then uniformly in $\eta \in [l, h]$,

$$\mathbb{E}_\eta W = 1 + o(1), \qquad \mathbb{E}_\eta W^2 = \exp\left(\frac{n^2}{M}\frac{1-\eta}{\eta}\right) + o(1). \tag{12}$$

Proof. The beta integral gives

$$\frac{1}{(L+1)\cdots(L+n)} = \frac{1}{(n-1)!}\int_0^1 x^L (1-x)^{n-1}\, dx.$$

Take expectation, use $\mathbb{E}x^L = (1-\eta+\eta x)^M$, and substitute $y = \eta(1-x)$. The result is the regularized incomplete beta integral with parameters $(n, M+1)$, which equals the binomial tail in (11). Since $\eta \ge l$ and $n = O(\sqrt{M})$, its missing probability is at most $e^{-c_l M}$ for all sufficiently large M.

For the second moment put $D = L - M\eta$. On $|D| \le M^{5/8}$, expand the finite logarithmic product to get

$$\log W = -\frac{nD}{M\eta} - \frac{n(n+1)}{2M}\frac{1-\eta}{\eta} + o(1). \tag{13}$$

The remainder is uniform: it is bounded by a constant depending only on l, h times

$$\frac{nD^2}{M^2} + \frac{n^2|D|}{M^2} + \frac{n^3}{M^2},$$

which tends to zero. For $k = 1, 2$, the binomial moment generating function yields

$$\log \mathbb{E}\exp\left(-\frac{knD}{M\eta}\right) = \frac{k^2n^2}{2M}\frac{1-\eta}{\eta} + O(n^3/M^2), \tag{14}$$

uniformly on the efficiency interval.

We justify removing the central event in these exponential moments. On $L \ge Ml/2$, the derivative of $\log W$ as a function of D has magnitude at most $2n/(Ml)$. At $D = 0$, every factor $\eta(M+j)/(M\eta+j)$ is at most one. Therefore

$$W^k \le \exp\left(\frac{2kn|D|}{Ml}\right).$$

For $|D| > M^{5/8}$, Hoeffding's bound $\Pr(|D| \ge u) \le 2e^{-2u^2/M}$ absorbs this linear exponential weight. Dividing $|D|$ into integer bands bounds its weighted tail by $CMe^{-cM^{1/4}} \to 0$. The same argument controls the exponential in (14). On $L < Ml/2$, the probability is at most $e^{-c_l M}$, whereas

$$V_{M,n}(L)^k \le \binom{M+n}{n}^k = \exp\{O(n\log M)\}.$$

Its contribution also vanishes. Equations (13)–(14) now give the kth moment $\exp\{(k^2-k)n^2(1-\eta)/(2M\eta)\} + o(1)$. This proves (12). □

Put $v_c = d_l n^2/M$, $D_{N,M} = (1+S_N)e^{v_c}$ and measure

$$Y = \frac{T_f V_{M,n}(L)}{D_{N,M}}. \tag{15}$$

Let $Q = T_f/F_f(\eta A)$. Independence and homogeneity give

$$\mathbb{E}\frac{Y}{F_f(A)} = \frac{\mathbb{E}W}{D_{N,M}}, \qquad \mathbb{E}\left(\frac{Y}{F_f(A)}\right)^2 = \frac{\mathbb{E}Q^2\mathbb{E}W^2}{D_{N,M}^2}.$$

By (8), $\mathbb{E}Q^2 \le 1 + S_N$; by Lemma 3.1, $\mathbb{E}W^2 \le e^{v_c} + o(1)$ and $\mathbb{E}W = 1 + o(1)$ uniformly. When both critical ratios are bounded, $D_{N,M}$ stays bounded and is at least one. Thus

$$\sup_{d,f,A,\eta} \mathbb{E}\left(\frac{Y}{F_f(A)} - 1\right)^2 \le 1 - \frac{1}{D_{N,M}} + o(1). \tag{16}$$

Together with (9), this proves the upper half of Theorem 2.1.

4 A feasible boundary prior gives the matching lower

The lower experiment is scalar: $f(z) = z^n$, output occupation $s = \eta t$, source occupation $t \geq a$. The known factor $n!$ cancels in relative loss, leaving target $F = (s/\eta)^n$. Scalar thermal states commute. Every collective source measurement is a parameter-independent randomization of the photon counts, whose total K has negative-binomial law

$$\Pr_s(K = k) = \binom{N+k-1}{k} \frac{s^k}{(1+s)^{N+k}}, \quad k \geq 0.$$

Conditional on K, the count vector is ancillary. Hence it suffices for a lower bound to give the statistician the full pair (K, L).

4.1 Supported independent coordinates

Put $\delta = n^{-1/4}$ and $m = n^{3/2}$. Let η have a beta prior with shapes $me_n, m(1-e_n)$, $e_n = l + \delta$, conditioned on

$$I_n = [l, l + 2\delta].$$

Independently, give $q = s/(1+s)$ a beta prior with shapes $mq_n, m(1-q_n)$, where $q_n = q(s_n)$, $s_n = a(l + 4\delta)$, conditioned on

$$E_n = \{q \geq q(a(l + 2\delta))\}.$$

For large n, $I_n \subset [l, h]$. On this rectangular prior support, $s \geq a(l + 2\delta) \geq a\eta$, so $t = s/\eta \geq a$. The support therefore lies in the actual source model. The unconditioned prior standard deviations are $O(n^{-3/4})$, smaller than the δ clearance from the conditioning boundaries. Both conditioning probabilities tend to one.

The likelihood factors in (q, η), so the posteriors are independent. The posterior minimizer of relative squared loss has conditional risk

$$1 - \frac{[\mathbb{E}(F^{-1} \mid K, L)]^2}{\mathbb{E}(F^{-2} \mid K, L)} = 1 - r_s(K) r_\eta(L), \tag{17}$$

where

$$r_s = \frac{[\mathbb{E}(s^{-n} \mid K)]^2}{\mathbb{E}(s^{-2n} \mid K)}, \qquad r_\eta = \frac{[\mathbb{E}(\eta^n \mid L)]^2}{\mathbb{E}(\eta^{2n} \mid L)}.$$

Cauchy–Schwarz puts both ratios in $[0, 1]$. We compute their limits on these supported priors, so bounded convergence can be applied to the ratios rather than to an unbounded loss.

4.2 The two beta moment ratios

For the source posterior write $U = \alpha + K$, $V = \beta + N$, where $\alpha = mq_n$, $\beta = m(1-q_n)$. Let p_j be the E_n probability under $\mathrm{Beta}(U - jn, V + jn)$, $j = 0, 1, 2$. Beta integration gives

$$r_s = \frac{\mathrm{B}(U-n, V+n)^2}{\mathrm{B}(U,V)\mathrm{B}(U-2n, V+2n)} \frac{p_1^2}{p_0 p_2}. \tag{18}$$

The shapes are positive for large n, since $\alpha \asymp n^{3/2} > 2n$. The source prior puts q a distance comparable to δ above its support boundary. Conditional photon variance is $Ns(1+s)$. Prior concentration and Chebyshev's inequality give $K/N = s_n + o_P(\delta)$, first on a fixed compact neighborhood of la, whose prior probability tends to one. Since $m/N = O(n^{-1/2}) = o(\delta)$ and $n/N = O(n^{-1})$,

the three tilted posterior means remain that distance above the boundary; their variances are $O_P(N^{-1}) = o_P(\delta^2)$. Consequently

$$p_j \to 1, \qquad U/N \to la, \qquad V/N \to 1 \quad \text{in probability.}$$

These statements also hold when $n^2/N \to 0$.

Writing ψ_1 for the trigamma function, the exact second difference is

$$\begin{aligned} &2\log \mathrm{B}(U-n, V+n) - \log \mathrm{B}(U,V) - \log \mathrm{B}(U-2n, V+2n) \\ &\quad = -\int_0^n \int_0^n [\psi_1(U-r-s) + \psi_1(V+r+s)]\, dr\, ds. \end{aligned} \tag{19}$$

Use $\psi_1(z) = z^{-1} + O(z^{-2})$ on the high-probability compact-ratio set. Since $n/N \to 0$, (19) tends in probability to $-\gamma_l \kappa_s$. Thus

$$r_s \to e^{-\gamma_l \kappa_s}. \tag{20}$$

For calibration put $A_\eta = \alpha_\eta + L$, $B_\eta = \beta_\eta + M - L$. Let u_j be the I_n probability under $\mathrm{Beta}(A_\eta + jn, B_\eta)$. Then

$$r_\eta = \frac{\mathrm{B}(A_\eta + n, B_\eta)^2}{\mathrm{B}(A_\eta, B_\eta)\mathrm{B}(A_\eta + 2n, B_\eta)} \frac{u_1^2}{u_0 u_2}. \tag{21}$$

Prior concentration and binomial variance give $L/M = l + \delta + o_P(\delta)$. Because $m/M = o(\delta)$ and $n/M = o(\delta)$, the three tilted posterior means stay $\delta(1 + o_P(1))$ from both endpoints of I_n, and their variances are $O_P(M^{-1}) = o_P(\delta^2)$. Hence $u_j \to 1$, $A_\eta/M \to l$ and $(A_\eta + B_\eta)/M \to 1$. With $G(x) = \log \mathrm{B}(A_\eta + x, B_\eta)$,

$$2G(n) - G(0) - G(2n) = -\int_0^n \int_0^n [\psi_1(A_\eta + r + s) - \psi_1(A_\eta + B_\eta + r + s)]\, dr\, ds, \tag{22}$$

which tends to $-d_l \kappa_c$. Therefore $r_\eta \to e^{-d_l \kappa_c}$. The moment ratios are bounded, so (17) converges in expectation to (4). Every supported-prior Bayes risk is a lower bound for (2); the scalar monomial is among its polynomials. This proves the matching lower, including either zero critical coefficient.

4.3 Diverging budgets and consistency

If $n^2/N \to \infty$, restrict $\eta = l$ and reveal it. Rescaling the target gives the thermal source-only experiment at occupation floor la; its minimax risk tends to one [6, Theorem 2.4].

For calibration collapse, keep an observed occupation $s \geq ha$ fixed. Take an interior efficiency η_0 and $\eta_1 = \eta_0 + \varepsilon_n$, where

$$\varepsilon_n = (n\sqrt{M})^{-1/2}.$$

When $n^2/M \to \infty$, we have $\varepsilon_n \to 0$, $M\varepsilon_n^2 \to 0$ and $n\varepsilon_n \to \infty$. The two binomial laws have total variation tending to zero by their $O(M\varepsilon_n^2)$ relative entropy. Source laws are identical, and the ratio of the two targets tends to infinity.

For completeness, if two positive targets are θ_0, θ_1 and their observation laws have densities p_0, p_1, their minimax relative risk is at least

$$\left(\frac{\theta_1 - \theta_0}{\theta_1 + \theta_0}\right)^2 (1 - \mathrm{TV}(p_0, p_1)). \tag{23}$$

Give target θ_i prior weight $\theta_i/(\theta_0 + \theta_1)$; for a common observation density the pointwise minimum weighted loss is the squared separation in (23). Multiply by $\min(p_0, p_1)$ and integrate. This proves

(23), which tends to one in the calibration construction. The zero decision has risk one, proving collapse.

For necessity in (5), revealing $\eta = l$ gives the source-only lower bound. For calibration, use the same fixed-output submodel with two interior efficiencies separated by $c\min(n^{-1}, M^{-1/2})$. The target separation is bounded below by a positive multiple of $\min(1, n/\sqrt{M})$, while a sufficiently small fixed c keeps total variation below $1/2$. Equation (23) then bounds risk below by a positive constant times $\min(1, n^2/M)$. The upper already proves sufficiency. Theorem 2.1 follows.

5 Finite-resource bounds and the measurement comparison

We prove Theorem 2.2, retaining the dependence on a. Put $x = \gamma_l n^2/N$ and $y = n^2/M$. For $x + y > 1$ the zero decision gives risk one. Otherwise let $\widehat{\eta}$ be the projection of L/M onto $[l, h]$, and use $Y = T_f \widehat{\eta}^{-n}$. With $Z = L/M - \eta$ and $R = (\eta/\widehat{\eta})^n$, projection and the $1/l$ Lipschitz bound for the logarithm give

$$|R - 1|^2 \le \frac{n^2}{l^2} Z^2 e^{2n|Z|/l}.$$

The binomial central moment and Hoeffding moment-generating bounds give

$$\mathbb{E}Z^4 \le \frac{7}{16M^2}, \qquad \mathbb{E}e^{t|Z|} \le 2e^{t^2/(8M)}.$$

Cauchy–Schwarz therefore yields

$$\mathbb{E}(R-1)^2 \le \frac{\sqrt{14}}{4l^2}\frac{n^2}{M} e^{n^2/(l^2 M)} \le K_l y, \qquad K_l = \frac{\sqrt{14}}{4l^2} e^{1/l^2}. \tag{24}$$

Also $\mathbb{E}R^2 \le 2(1 + K_l)$. For $Q = T_f/F_f(\eta A)$, independence and $\mathbb{E}Q = 1$ give

$$\mathbb{E}(QR - 1)^2 = \mathrm{Var}(Q)\mathbb{E}R^2 + \mathbb{E}(R-1)^2.$$

Use $S_N \le (e-1)x$ when $x \le 1$ and (24). This proves the upper with, for example, $C_l = \max\{1, 2(e-1)(1 + K_l), K_l\}$.

For the source lower, restrict and reveal $\eta = l$. Calibration is then ancillary for A, and the thermal lower gives $R_{N,M,n} \ge (7/2304)\min(1, x)$ [6]. For the calibration lower keep scalar output $s \ge ha$ fixed, and choose

$$\eta_0 = (l + h)/2, \quad \eta_1 = \eta_0 + \varepsilon, \qquad \varepsilon = c_0 \min(n^{-1}, M^{-1/2}),$$

where

$$c_0 = \min\{(h - l)/4, l/2, \sqrt{l(1 - h)}/4\} > 0.$$

Both parameters are feasible. Their Bernoulli relative entropy is at most $\varepsilon^2/[l(1 - h)]$, so the M-sample total variation is below $1/4$. Their target harmonic separation is

$$\Delta = \tanh\left(\frac{n}{2}\log\frac{\eta_1}{\eta_0}\right) \ge \frac{n\varepsilon}{4h}.$$

Here $\log(\eta_1/\eta_0) \ge \varepsilon/h$, while its argument in tanh is at most $c_0/(2l) \le 1/4$, where $\tanh z \ge z/2$. Equation (23) now gives a positive constant depending only on l, h times $\min(1, y)$. Combining this with the source lower proves (6), since the maximum of the two lower bounds controls $\min(1, x + y)$ up to a constant.

For heterodyne, the thermal predecessor supplies an unbiased output-energy statistic with sharp envelope obtained by replacing γ_l by γ_l^2, and the matching scalar source minimax order [6, Section 6]. The same independent clipped-calibration upper and fixed-output calibration lower apply unchanged. This proves the heterodyne assertion.

The rate comparison identifies when a growing measurement advantage survives calibration costs. If n^2/M dominates both source terms, calibration sets both resource orders. If it is small compared with $\gamma_l n^2/N$, the thermal quantum/heterodyne source distinction remains visible. Thus calibrating the mean attenuation does not make its statistical cost free.

6 Allocation and worked resource comparisons

Let $c_s, c_c > 0$ be the specified costs of one source copy and one calibration probe. They may describe time or preparation cost; source occupation remains unknown, so these are prescribed costs rather than an assertion of equal photon energies. For a total budget B_n impose $c_s N + c_c M \le B_n$.

Corollary 6.1 (Optimal deterministic allocation)**.** *If $n^2/B_n \to \kappa \in (0,\infty)$, then the infimum of* (2) *over deterministic positive integer allocations satisfies*

$$\inf_{c_s N + c_c M \le B_n} R_{N,M,n} \longrightarrow 1 - \exp\{-\kappa(\sqrt{c_s\gamma_l} + \sqrt{c_c d_l})^2\}. \tag{25}$$

An attaining allocation obeys

$$\frac{N}{M} \to \sqrt{\frac{\gamma_l c_c}{d_l c_s}}, \qquad \frac{c_s N}{B_n} \to \frac{\sqrt{c_s\gamma_l}}{\sqrt{c_s\gamma_l} + \sqrt{c_c d_l}}. \tag{26}$$

Proof. For positive N, M, Cauchy–Schwarz gives

$$\left(\frac{\gamma_l}{N} + \frac{d_l}{M}\right)(c_s N + c_c M) \ge (\sqrt{c_s\gamma_l} + \sqrt{c_c d_l})^2,$$

with equality at (26). Rounding an interior allocation changes its scaled limits by $o(1)$, and Theorem 2.1 gives the upper. For the lower, first extend every admissible allocation by extra observations so that $c_s N + c_c M = B_n + O(1)$: the remaining budget can be made less than a fixed sample cost. Extra observations can be discarded, so this extension can only decrease risk. Now take a subsequence on which the source cost fraction converges; the two cost fractions sum to $1 + o(1)$. If it is zero or one, the corresponding critical ratio diverges and risk tends to one. An interior fraction gives the exact limit in Theorem 2.1, bounded below by (25). This proves the infimum limit. □

For equal costs, take $a = 1/10$, $l = 4/5$ and any fixed $h \in (4/5, 1)$. Then $\gamma_l = 27/2$ and $d_l = 1/4$. The optimal source fraction is

$$\frac{\sqrt{27/2}}{\sqrt{27/2} + 1/2} \simeq 0.8802.$$

At $B_n/n^2 \to 100$, the optimal exponent is approximately 0.17424 and the limiting risk is approximately 0.1599. Equal allocation gives exponent 0.275 and limiting risk approximately 0.2404. Calibration is therefore costly but need not receive half the observations. These values describe the growing-degree limit, not a finite-n numerical minimax computation.

Along sequences in the critical range, allocating an order-n^2 calibration budget produces a nonvanishing calibration penalty. Relative consistency requires both budgets to grow faster than n^2. The allocation conclusion is for the specified deterministic experiment and costs, as in the classical design logic of [2].

7 Calibration supplied only as an interval

A complementary experiment provides N source copies and the information $0 < l \le \eta \le h \le 1$, with no calibration trials. Here endpoints may vary with n. Let $R^{\mathrm{int}}_{N,n}$ be (2) with L absent. Write

$$u = l^n, \quad v = h^n, \quad \chi = \frac{4uv}{(u+v)^2} = \operatorname{sech}^2\left(\frac{n}{2}\log\frac{h}{l}\right), \quad x = \left(1 + \frac{1}{la}\right)\frac{n^2}{N}.$$

Let $S_- = S(N, n, la)$ be the sum in (8).

Theorem 7.1 (Interval calibration)**.** *The decision $Y = 2T_f/[(1+S_-)(u+v)]$ satisfies*

$$\max\left\{1-\chi, \frac{7}{2304}\min(1,x)\right\} \le R^{\mathrm{int}}_{N,n} \le 1 - \frac{\chi}{1+S_-} \le 1 - \chi e^{-x}. \tag{27}$$

The exact minimax risk with observed output covariance known is $1-\chi$. For arbitrary sequences, $R^{\mathrm{int}}_{N,n} \to 0$ precisely when $x \to 0$ and $n\log(h/l) \to 0$. If either budget diverges, the risk tends to one.

For fixed $a > 0$, suppose $n \to \infty$, $l, h \to \eta_0 > 0$, $x \to \lambda \in [0,\infty)$ and $n\log(h/l) \to \tau \in [0,\infty)$. Then

$$R^{\mathrm{int}}_{N,n} \longrightarrow 1 - e^{-\lambda}\operatorname{sech}^2(\tau/2). \tag{28}$$

Proof. Put $z = \eta^n \in [u, v]$. For $c = 2/[(1+S_-)(u+v)]$, relative risk is bounded by $(1+S_-)c^2z^2 - 2cz + 1$. This convex quadratic has equal endpoint values $1 - \chi/(1+S_-)$. The bound $S_- \le e^x - 1$ proves the upper.

For a scalar observed occupation $s \ge ha$, the source occupations $s/h, s/l$ give identical source data. Their energies have ratio v/u. The optimal constant action is their harmonic midpoint and its maximum risk is $((v-u)/(v+u))^2 = 1-\chi$. Conversely, when the output covariance is known, $2F_f(B)/(u+v)$ has at most this risk throughout the interval. This proves the exact identification floor. Restricting $\eta = l$ gives the source-only minimax lower in (27). These lower bounds and the upper prove consistency; divergence uses the source-only collapse or the identification floor, respectively.

For the critical upper, the same dominated series argument as (9) gives $1 + S_- \to e^\lambda$; $\chi \to \operatorname{sech}^2(\tau/2)$. For the matching lower choose scalar output occupation s independently of η, with endpoint prior probabilities

$$\Pr(\eta = l) = \frac{v}{u+v}, \qquad \Pr(\eta = h) = \frac{u}{u+v}.$$

Support s above ha. Since the observation depends only on s, the efficiency posterior is unchanged and independent of the output posterior. Its moment ratio is exactly

$$\frac{(\mathbb{E}\eta^n)^2}{\mathbb{E}\eta^{2n}} = \chi.$$

To obtain the source factor when $\lambda > 0$, use $q = s/(1+s)$ with a beta prior of total shape $m = n^{3/2}$, centered $\delta = n^{-1/4}$ above $q(ha)$ and conditioned on $q \ge q(ha)$. Since $ha \to \eta_0 a > 0$, the center stays in a compact subinterval of $(0,1)$. The source posterior and its two negative-moment tilts remain $\delta(1+o_P(1))$ above the boundary by the argument of Section 4. Formula (18) and the trigamma identity yield a ratio tending to $e^{-\lambda}$: the prior's higher output floor ha and the upper's la have the same positive limit. Thus the posterior risk tends to $1 - \chi e^{-\lambda}$ in expectation, proving (28). When $\lambda = 0$, the identification floor and upper already give the result. □

For example, let $x \to 0.1$ and $n \log(h/l) \to 1$. The identification floor tends to $\tanh^2(1/2) \simeq 0.2136$, and the joint risk tends to approximately 0.2884. A known-loss source experiment at the same copy exponent has risk $1 - e^{-0.1} \simeq 0.0952$. This separates the persistent ambiguity from sampling noise. Shrinking the log-calibration width to $o(1/n)$ removes that ambiguity; collecting calibration data instead leads to the paid experiment of Theorem 2.1.

Reproducibility. The accompanying standard-library Python program checks the reciprocal-binomial first-moment identity, moment positivity, interval endpoint equalization and calibration posterior beta ratios by exact rational arithmetic. The analytic proofs establish the uniform limits. The source-energy construction and its own reproducibility evidence are supplied in the cited thermal predecessor.

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

References

[1] K. M. R. Audenaert and S. Scheel, *Statistical inference from imperfect photon detection*, New J. Phys. **11** (2009), 113052. https://doi.org/10.1088/1367-2630/11/11/113052.

[2] S. C. Medin, J. Murray-Bruce, D. Castañón and V. K. Goyal, *Beyond binomial and negative binomial: adaptation in Bernoulli parameter estimation*, IEEE Trans. Comput. Imaging **5** (2019), 570–584. https://doi.org/10.1109/TCI.2019.2913108.

[3] D. Mogilevtsev, J. Řeháček and Z. Hradil, *Self-calibration for self-consistent tomography*, New J. Phys. **14** (2012), 095001. https://doi.org/10.1088/1367-2630/14/9/095001.

[4] R. Nair, *Quantum-limited loss sensing: multiparameter estimation and Bures distance between loss channels*, Phys. Rev. Lett. **121** (2018), 230801. https://doi.org/10.1103/PhysRevLett.121.230801.

[5] T. R. L. Phillips and A. Zhigljavsky, *Approximation of inverse moments of discrete distributions*, Statist. Probab. Lett. **94** (2014), 135–143. https://doi.org/10.1016/j.spl.2014.07.007.

[6] 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.

[7] J. Wei, P. He and T. Tong, *Estimating the reciprocal of a binomial proportion*, Int. Statist. Rev. **92** (2024), 1–16. https://doi.org/10.1111/insr.12539.