Exact low-brightness minimax risk for quartic Gaussian energy

Scott Schubr
Luminary Oddities

November 1, 2026

Abstract

We determine the exact critical minimax relative risk for estimating the normally ordered quartic energy of an unknown centered single-mode Gaussian state. The occupation and complex anomalous covariance are both unknown, and arbitrary biased collective quantum procedures are allowed. If the occupation floor a tends to zero and the copy budget satisfies $Na^2 \to \lambda \in (0, \infty)$, the minimax risk tends to $(1 + 2\sqrt{\lambda})^{-2}$. A vanishing fraction of photon-counting observations estimates occupation; an affine correction of a known unbiased collective observable attains the limit uniformly over the unbounded state family. For the lower bound, a harmonic decomposition and common-phase averaging reduce a compact boundary experiment to a Poisson count with a positive shifted target. A proper Bayes equalizer solves that decision problem. We also prove complete collapse when $Na^2 \to 0$ and sharp high-information limits: the quantum risk is asymptotic to $1/(4Na^2)$, whereas standard heterodyne has risk $4/(Na^2)$. Thus the previously known equality of copy-cost orders coexists with a sharp factor-sixteen asymptotic efficiency advantage.

Keywords. Gaussian quantum states; quartic energy; minimax estimation; relative loss; squeezing.
MSC 2020. 62C20, 62F12, 81P50.

1 Introduction

For a centered Gaussian optical mode, the quartic correlation

$$F = \langle (b^\dagger)^2 b^2 \rangle = 2t^2 + |c|^2, \qquad t = \langle b^\dagger b \rangle, \quad c = \langle b^2 \rangle,$$

depends on occupation and squeezing. Estimating it uniformly is different from estimating the thermal expression $2t^2$, where c is fixed to zero. In the full Gaussian model with $t \geq a$, both collective quantum measurements and standard heterodyne have squared relative minimax risk of order $\min\{1, (Na^2)^{-1}\}$ [6]. That order theorem leaves two questions unanswered: what accuracy is possible at the critical budget $N \asymp a^{-2}$, and does equality of orders mean equality of efficiency?

This paper answers the first question exactly and the second with sharp constants. The quantum critical risk is

$$\rho(\lambda) = \frac{1}{(1 + 2\sqrt{\lambda})^2}, \qquad \lambda = \lim Na^2. \tag{1}$$

An estimator attaining this limit must balance bias and variance. Multiplying the optimal unbiased observable by a single constant does not suffice: its best critical risk is $1/(1 + 4\lambda)$. An additional occupation-dependent offset reaches (1). Occupation can be learned from a negligible pilot because its relevant sample scale is a^{-1}, while resolving the quartic contribution of weak squeezing costs a^{-2}. At $\lambda = 1/4$, the offset reduces limiting risk from $1/2$ to $1/4$.

The proof has complementary statistical parts. A collective harmonic label separates thermal occupation from radial pair excitations. Near the low-occupation boundary, the latter converge to a Poisson count. Estimating a positive offset plus its Poisson mean has a proper Bayes equalizer, which supplies the minimax lower bound. For attainability, we do not assume that the state lies in this compact boundary regime. Exact variance inequalities and an independent occupation pilot control the entire unbounded Gaussian family.

Relation to earlier work. The exact unbiased observable and its variance, as well as the quartic order theorem, are established in [6]. They are inputs here, not new claims. The present contribution is the exact arbitrary-bias critical decision law, its uniformly valid adaptive attainer, and the sharp high-information quantum/heterodyne comparison. Thermal Gaussian polynomial-energy estimation [7] concerns the smaller model with no unknown anomalous covariance.

Local information formulas for single-mode Gaussian states [5] and efficient moment-based Gaussian characterization [8] provide important parameter-estimation benchmarks. They do not by themselves determine a global relative-risk limit at a moving vacuum boundary. Recent cold/warm Gaussian tomography results [2] concern trace-distance reconstruction of a full state; our objective is a specified unbounded functional with relative loss. Local pure-state Gaussian-shift approximations [1] give related statistical ancestry, but do not automatically justify the mixed triangular experiment used below.

The oscillator harmonic decomposition is classical representation theory; its form for the metaplectic and orthogonal groups, and its relevance to identical Gaussian copies, are described in [3]. Common-phase averaging and proper Bayes equalizers are standard decision-theoretic tools. The Bayesian information inequality used for the high-information constants is the van Trees inequality [4]. We supply the particular probability approximation, global adaptation bounds and matching minimax deductions needed for this problem.

2 Experiment and main results

Let $[b, b^\dagger] = 1$. A centered single-mode Gaussian state is specified by

$$t = \operatorname{Tr} \varrho b^\dagger b \geq 0, \qquad c = \operatorname{Tr} \varrho b^2, \qquad u = |c|^2 \leq t(t+1). \tag{2}$$

Every pair satisfying (2) is allowed. In particular, the phase of c is unknown. We observe N independent copies, and the known floor is $t \geq a > 0$. An estimator is an arbitrary collective POVM with a measurable real decision Y. Define

$$R_Q(N, a) = \inf_Y \sup_{t \geq a,\, |c|^2 \leq t(t+1)} \mathbb{E}_{t,c} \left(\frac{Y}{2t^2 + |c|^2} - 1 \right)^2. \tag{3}$$

Bias is unrestricted. Infinite risks are permitted in the infimum. Define $R_H(N, a)$ by restricting the measurement to standard heterodyne on every copy, followed by arbitrary processing. Its record Z is centered complex Gaussian with $\mathbb{E}|Z|^2 = t + 1$ and $\mathbb{E}Z^2 = c$. Thus $R_Q \leq R_H \leq 1$, the last inequality following from the zero decision.

Theorem 2.1 (Exact critical risk)**.** *Along every sequence $a \to 0$ and positive integers N with $Na^2 \to \lambda \in (0, \infty)$,*

$$R_Q(N, a) \longrightarrow \rho(\lambda) = \frac{1}{(1 + 2\sqrt{\lambda})^2}. \tag{4}$$

If $Na^2 \to 0$, both $R_Q(N, a)$ and $R_H(N, a)$ tend to one.

Theorem 2.2 (Sharp high-information comparison)**.** *If $a \to 0$ and $Na^2 \to \infty$, then*

$$Na^2 R_Q(N,a) \longrightarrow \frac{1}{4}, \qquad Na^2 R_H(N,a) \longrightarrow 4. \tag{5}$$

These are simultaneous triangular limits, not an interchange of limits in (4)*.*

At vanishing relative mean-squared error, (5) gives a factor-sixteen asymptotic copy-cost ratio between the unrestricted quantum class and standard heterodyne. The attaining quantum procedure uses the spectral measurement described below.

2.1 The attaining procedure

For L copies define

$$K_L = \sum_{r=1}^{L} b_r^\dagger b_r, \qquad Q_L = \sum_{r=1}^{L} b_r^2, \qquad T_L = \frac{Q_L^\dagger Q_L + 2K_L(K_L-1)}{L(L+2)}.$$

The observable T_L is the established positive unbiased quartic estimator [6]. It is a self-adjoint direct sum on finite photon-number sectors, and all its moments used here exist at finite Gaussian covariance. Its spectral measurement is allowed in (3).

For sufficiently small a in the positive critical regime, let

$$m = \lceil N\sqrt{a} \rceil, \quad L = N - m, \quad \tau = \frac{1}{m}\sum_{r=1}^{m} n_r, \qquad \widehat{Y} = \alpha_L(\tau) T_L + B_L(\tau), \tag{6}$$

where n_r are individual photon counts on the first m copies and T_L is measured on the remaining copies, independently of the pilot. The known functions are

$$\alpha_L(x) = \frac{2\sqrt{L}x}{1+2\sqrt{L}x}, \qquad B_L(x) = \frac{4x^2}{1+2\sqrt{L}x} \quad (x \geq 0). \tag{7}$$

The procedure uses the observed pilot in place of the unknown t and attains both (4) and the quantum constant in (5). Its main measurement is the spectral measurement of T_L itself.

3 A Poisson decision problem

Lemma 3.1 (Positive shifted Poisson target)**.** *For known $A > 0$, observe $J \sim \mathrm{Pois}(x)$, $x \geq 0$, and estimate $A + x$ under squared relative loss. The minimax value is $\rho(A)$. A minimax decision is*

$$d_A J + A(2 - d_A), \qquad d_A = \frac{2\sqrt{A}}{1+2\sqrt{A}}. \tag{8}$$

Proof. Its risk is

$$\frac{d_A^2 x + (1-d_A)^2 (A-x)^2}{(A+x)^2} = (1-d_A)^2,$$

because $d_A^2 = 4A(1-d_A)^2$. For the matching lower, take the proper prior

$$\pi_A(x) = C_A (A+x)^2 x^{r-1} e^{-sx}, \qquad s = (2\sqrt{A})^{-1}, \quad r = As = \sqrt{A}/2. \tag{9}$$

The relative-loss Bayes decision is

$$\frac{\mathbb{E}[(A+x)^{-1} \mid J]}{\mathbb{E}[(A+x)^{-2} \mid J]} = A + \frac{J+r}{1+s},$$

by gamma integration. This equals (8). Its constant risk is therefore also its Bayes risk, proving minimaxity. □

We will use the problem restricted to $0 \le x \le X$, then let $X \to \infty$. Its minimax values increase to $\rho(A)$. Here is one justification of this exhaustion. Nonnegative decisions suffice. Near-minimizers on expanding compact intervals have risk at most one plus a vanishing error. Once the interval contains $x = 1$, that risk bounds each coordinate of the decision sequence, since $\mathrm{Pois}(1)\{j\} > 0$. A diagonal subsequence converges coordinatewise; Fatou's lemma for each fixed x bounds the risk of the limiting global decision by the limit of the compact values. The opposite inequality follows by restriction. Thus no unbounded-loss limit is being inferred merely from compact weak convergence.

4 The collective boundary experiment

We give the probability reduction needed for the lower bound. Its purpose is to retain the pair coherence shared by the copies. Independently dephasing every copy would be a different experiment.

4.1 Harmonic labels and common-phase averaging

Write the state as squeezed thermal, with

$$\nu = \sqrt{(t+1/2)^2 - |c|^2} - 1/2, \qquad q = \frac{\nu}{1+\nu}, \qquad z = \frac{c}{1+t+\nu}, \qquad p = |z|^2. \tag{10}$$

An irrelevant sign of z depends on the squeeze convention. In these coordinates

$$t = \nu + (2\nu+1)\frac{p}{1-p}, \qquad c = (2\nu+1)\frac{z}{1-p}.$$

For $N \ge 2$ the N-oscillator Fock space decomposes into harmonic degree ℓ, multiplicity σ, and radial excitation $j \ge 0$. If $\mathcal{H}_\ell$ is the kernel of Q_N on the ℓ-photon sector, then

$$d_{\ell,N} = \dim \mathcal{H}_\ell = \binom{N+\ell-1}{\ell} - \binom{N+\ell-3}{\ell-2}, \qquad k_\ell = N/4 + \ell/2, \tag{11}$$

with the second binomial zero for $\ell < 2$. For an orthonormal harmonic basis $h_{\ell,\sigma}$, the normalized descendants are

$$|\ell, \sigma, j\rangle = \frac{(Q_N^\dagger)^j h_{\ell,\sigma}}{2^j \sqrt{j!(\ell+N/2)^{\bar{j}}}}.$$

On these vectors, $K_N = \ell + 2j$ and

$$Q_N^\dagger Q_N = 4j(\ell + N/2 + j - 1). \tag{12}$$

For completeness, in the Bargmann polynomial model multiplication by $\sum z_r^2$ is injective and its adjoint is $\sum \partial_r^2$. Each homogeneous sector splits orthogonally into its harmonic kernel and the

image from degree two below. Iteration gives the decomposition and dimension in (11). The commutators for

$$K_0 = (K_N + N/2)/2, \quad K_+ = Q_N^\dagger/2, \quad K_- = Q_N/2$$

are $[K_0, K_\pm] = \pm K_\pm$ and $[K_-, K_+] = 2K_0$; their ladder recurrence gives the normalization and (12). This is the oscillator harmonic decomposition underlying the duality in [3].

The thermal product $(1-q)^N q^{K_N}$ has harmonic weights

$$w_\ell(q) = d_{\ell,N} \frac{(1-q)^N q^\ell}{1-q^2}. \tag{13}$$

The multiplicity is uniform conditional on ℓ, and the radial state is

$$\tau_{q^2} = (1-q^2) \sum_{j \geq 0} q^{2j} |j\rangle\langle j|.$$

Normalization follows from $\sum_\ell d_{\ell,N} q^\ell = (1-q^2)/(1-q)^N$. Collective squeezing preserves ℓ, σ and acts on the radial ladder by

$$D_k(z) = e^{zK_+} (1-p)^{K_0} e^{-\overline{z} K_-}.$$

This identity follows by differentiating the product with the displayed commutators. In particular, its vacuum probabilities are

$$|\langle j | D_k(z) | 0 \rangle|^2 = (1-p)^{2k} \frac{(2k)^{\overline{j}}}{j!} p^j. \tag{14}$$

Thus measuring the harmonic and radial labels has exact law

$$P_{t,c}\{\ell, j\} = w_\ell(q)(1-q^2) \sum_{v \geq 0} q^{2v} |\langle j | D_{k_\ell}(z) | v \rangle|^2. \tag{15}$$

The labels are defined operationally by the joint spectral measurement of K_N and $Q_N^\dagger Q_N$ in the unrestricted measurement class.

For a phase-invariant target, average any procedure over the *common* phase $e^{i\phi K_N}$. Its maximum risk cannot increase. The corresponding averaged source is diagonal in the basis above: it already has no coherence between ℓ labels, is scalar on multiplicity, and common phase removes radial off-diagonal entries. Every POVM on these commuting states is a classical randomization of (ℓ, j). Conditional averaging of its real decision reduces squared loss. Therefore the minimax problem for phase-invariant targets equals the classical minimax problem with law (15).

4.2 A compact Poisson approximation

Lemma 4.1. *Suppose $a \to 0$, $Na^2 \to \lambda \in (0, \infty)$, and*

$$t = a\theta, \quad c = a\kappa, \qquad 1 \leq \theta \leq \Theta, \quad |\kappa| \leq C,$$

for fixed finite Θ, C. Uniformly on this parameter set, the law in (15) *has total variation distance tending to zero from*

$$w_\ell\left(\frac{a\theta}{1+a\theta}\right) \otimes \mathrm{Pois}\left(\lambda|\kappa|^2/2\right). \tag{16}$$

Proof. All the displayed parameters are physical for sufficiently small a. Rationalizing (10) gives $0 \le a\theta - \nu \le a^2C^2$, and $\nu \ge a\theta/2$ eventually. The relative entropy between geometric thermal distributions of means ν and $a\theta$ is $O(a^3)$ uniformly: its quadratic mean-parameter bound uses information $1/[s(1+s)]$ on the intervening interval. For N independent counts this is $O(Na^3) = O(a)$. The harmonic label is a channel from those thermal counts, so data processing and Pinsker replace $w_\ell(q)$ by the first factor in (16) with error $O(\sqrt{a})$.

For the radial comparisons, average conditional errors under the original weights $w_\ell(q)$ and apply the harmonic-weight replacement afterward. The radial state τ_{q^2} has trace distance $q^2 = O(a^2)$ from its vacuum, where trace distance includes the factor $1/2$. Squeezing and measurement do not increase that distance. Hence its radial law may be replaced by the negative binomial (14) with shape $r_\ell = N/2 + \ell$.

That law is a Poisson mixture with gamma intensity of mean $r_\ell p/(1-p)$ and variance $r_\ell p^2/(1-p)^2$. The coupling bound $\mathrm{TV}(\mathrm{Pois}(x), \mathrm{Pois}(y)) \le |x-y|$ yields error at most $\sqrt{r_\ell}p/(1-p)$ from the Poisson law of its mean. Here $p = O(a^2)$ and

$$\mathbb{E}\ell = N\nu - \frac{2q^2}{1-q^2} \le Na\Theta.$$

Averaging the error gives $O(a^2\sqrt{N}) = O(a)$. Also

$$\mathbb{E}\left| r_\ell \frac{p}{1-p} - \frac{Na^2|\kappa|^2}{2} \right| = O(a),$$

since $p/(1-p) = a^2|\kappa|^2(1+O(a))$ and $a^2\mathbb{E}\ell = O(a)$. A second Poisson coupling and $Na^2 \to \lambda$ prove the claim. □

The harmonic channel in (16) is noisy and depends on θ. It is not an exact observation of an unknown real parameter. We only need $\theta = 1$, when that channel is ancillary for κ.

4.3 The critical lower bound

Fix $\theta = 1$ and restrict $|\kappa| \le C$, with arbitrary phase. The target divided by a^2 is $2 + |\kappa|^2$. Project the normalized decision Y/a^2 to $[2, 2+C^2]$; this can only reduce risk, and makes the loss uniformly bounded. Lemma 4.1 therefore transfers the compact minimax problem to the Poisson experiment in (16). Multiplication of both normalized target and decision by $\lambda/2$ gives target $\lambda + x$, with $x = \lambda|\kappa|^2/2$. The ancillary harmonic channel cannot improve a convex-loss decision, since its randomization can be averaged out. Consequently

$$\liminf R_Q(N,a) \;\ge\; \inf_{\delta} \sup_{0 \le x \le \lambda C^2/2} \mathbb{E}_{\mathrm{Pois}(x)} \left(\frac{\delta(J)}{\lambda + x} - 1 \right)^2.$$

Let $C \to \infty$. The compact exhaustion after Lemma 3.1 proves

$$\liminf R_Q(N,a) \ge \rho(\lambda). \tag{17}$$

Uniform attainment over the entire original state family is established next.

5 Uniform attainment without known occupation

The exact inherited variance of T_L is [6]

$$\begin{aligned}\mathrm{Var}(T_L) &= \frac{A_1}{L} + \frac{3A_2}{L(L+2)}, \\ A_1 &= 2u + 16t^3(t+1) + (52t^2+20t)u + 4u^2, \\ A_2 &= 4t^2(t+1)^2 + 16t(t+1)u + 4u^2.\end{aligned} \tag{18}$$

Using $F^2 \geq 4t^4$, $F^2 \geq 8t^2u$ and $F^2 \geq u^2$ gives the all-domain bounds

$$\frac{A_1 - 2u}{F^2} \leq \frac{29}{2} + \frac{13}{2t}, \qquad \frac{A_2}{F^2} \leq 7 + \frac{4}{t} + \frac{1}{t^2}. \tag{19}$$

These hold for every $t > 0, u \geq 0$, hence over the unbounded physical domain.

5.1 Oracle identity

Temporarily suppose that t is known and set $Y_t = \alpha_L(t)T_L + B_L(t)$. Its bias is $(1-\alpha_L(t))(2t^2-u)$. Since

$$\alpha_L(t)^2\frac{2u}{L} + (1-\alpha_L(t))^2(2t^2-u)^2 = (1-\alpha_L(t))^2F^2,$$

its risk is exactly $\rho(Lt^2)$ plus the nonnegative variance remainder. In particular

$$\sup_{t\geq a,u} \mathbb{E}(Y_t/F - 1)^2 \leq \rho(La^2) + e(a,L), \qquad e(a,L) = \frac{29/2 + 13/(2a)}{L} + \frac{3(7 + 4/a + 1/a^2)}{L(L+2)}. \tag{20}$$

If $La^2 \to \lambda > 0$, then $e(a,L) \to 0$. This identity translates the Poisson equalizer to the original quantum observable, including states far outside the compact lower-bound family.

5.2 Removing the oracle

The photon pilot in (6) satisfies

$$\mathbb{E}\tau = t, \qquad \mathrm{Var}\,\tau = \frac{t(t+1)+u}{m} \leq \frac{2t(t+1)}{m}. \tag{21}$$

For any $x \geq 0$, direct subtraction and differentiation give

$$|\alpha_L(x) - \alpha_L(t)| \leq \frac{|x-t|}{t}, \qquad B_L'(x) = \frac{2}{\sqrt{L}}\left[1 - (1+2\sqrt{L}x)^{-2}\right] \leq \frac{2}{\sqrt{L}}. \tag{22}$$

Writing $\Delta\alpha = \alpha_L(\tau) - \alpha_L(t)$ and $\Delta B = B_L(\tau) - B_L(t)$, we obtain uniformly over $t \geq a$

$$\mathbb{E}(\Delta\alpha)^2 \leq \frac{2(1+a^{-1})}{m}, \qquad \mathbb{E}(\Delta B/F)^2 \leq \frac{2(a+1)}{mLa^3}. \tag{23}$$

The second bound uses $F \geq 2t^2$; the function $(t+1)/t^3$ decreases. Thus the pilot variance alone controls the adaptation error.

Independence of the pilot and main measurement now gives

$$\mathbb{E}\left(\frac{\widehat{Y} - Y_t}{F}\right)^2 \leq 2\mathbb{E}(\Delta\alpha)^2\mathbb{E}(T_L/F)^2 + 2\mathbb{E}(\Delta B/F)^2. \tag{24}$$

Equations (18)–(19), together with $2u/F^2 \leq 1/(4t^2)$, bound $\mathbb{E}(T_L/F)^2$ uniformly when La^2 stays bounded away from zero. Since $m/N \to 0$, $ma \to \infty$ and $La^2 \to \lambda$, the right side of (24) tends to zero. Minkowski's inequality and (20) prove the upper matching (17), completing the positive critical part of Theorem 2.1.

5.3 What the offset changes

The same exact variance shows

$$\sup_{t\ge a,u} \operatorname{Var}(T_N)/F^2 \longrightarrow \frac{1}{4\lambda}.$$

The upper follows from (19); $t = a, u = 2a^2$ attains its limit. A decision sT_N with one known constant s has maximum risk $(1-s)^2 + s^2 \sup \operatorname{Var}(T_N)/F^2$. Optimizing s gives critical value $1/(1+4\lambda)$ within this class. By contrast, (6) attains $\rho(\lambda)$. For example,

λ	best constant multiple of T_N	full quantum minimax
1/4	1/2	1/4
1	1/5	1/9

The offset uses information about the thermal part of the target, rather than suppressing every component by the same factor. The greatest improvement over that specified constant-multiple class is a factor two, at $\lambda = 1/4$.

6 Collapse and sharp high-information constants

6.1 Collapse below the critical budget

For the thermal state of occupation a, let $S(h)$ be real squeezing and put

$$h = \sqrt{a/\sqrt{N}}, \qquad G = \frac{b^2 - (b^\dagger)^2}{2i}, \qquad S(h) = e^{ihG}.$$

The squeezed state has $t_h = a + (2a+1)\sinh^2 h \ge a$ and $|c_h| = (2a+1)\sinh h \cosh h$. If $Na^2 \to 0$, then $h/a \to \infty$ and $\sqrt{N}h \to 0$. Its target is at least h^2, whereas the unsqueezed target is $2a^2$, so their ratio diverges.

Purifying the common thermal seed gives outcome total variation at most

$$|h|\sqrt{N \operatorname{Var}_{\varrho_a} G}, \qquad \operatorname{Var}_{\varrho_a} G = a^2 + a + 1/2. \tag{25}$$

Indeed the characteristic-function overlap obeys $|\operatorname{Tr} \varrho_a e^{ihG}|^2 \ge 1 - h^2 \operatorname{Var} G$; tensoring and trace-distance contraction give (25). For positive targets f_0, f_1, the two-point relative-loss lower bound is

$$\left(\frac{f_1 - f_0}{f_1 + f_0}\right)^2 (1 - \mathrm{TV}).$$

It follows by integrating the pointwise minimum with prior weights $f_0/(f_0+f_1)$ and $f_1/(f_0+f_1)$ against the common part of the two laws. Here it tends to one; the zero decision matches it. This proves collapse, including bounded N sequences.

6.2 A Bayesian information bound for arbitrary bias

For a smooth scalar submodel with positive target $g(s)$ and compactly supported continuously differentiable prior π vanishing at its endpoints, integration by parts and Cauchy–Schwarz give

$$\sup_s \mathbb{E}_s(Y/g(s) - 1)^2 \ge \frac{(\mathbb{E}_\pi g')^2}{(\sup_{\operatorname{supp}\pi} g^2)(\mathbb{E}_\pi I_Y + I_\pi)}, \qquad I_\pi = \int \frac{(\pi')^2}{\pi}. \tag{26}$$

This is the van Trees argument [4], applied first to ordinary squared error and then using its upper bound by $\sup g^2$ times the maximum relative risk. Decisions may first be projected to the target range on this prior support, reducing every risk and making them bounded. Thus arbitrary bias and integrability cause no additional assumption. A convenient prior is $\pi_0(s) = \cos^2(\pi s/2)$ on $[-1, 1]$, with information π^2; half-width w gives $I_\pi = \pi^2/w^2$.

For the quantum lower use the squeezed thermal seed ϱ_a and parameter r, with

$$r_0 = \sqrt{2}a, \qquad w = a\delta, \qquad \delta = (Na^2)^{-1/4}.$$

Every prior point has occupation at least a. Uniformly on this shrinking interval, $F/a^2 \to 4$ and $F'/a \to 2\sqrt{2}$. The N-copy purification has classical measurement information at most

$$4N \operatorname{Var}_{\varrho_a} G = 4N(a^2 + a + 1/2) = 2N(1 + o(1)).$$

The bound follows from amplitude Cauchy–Schwarz for any POVM and the squared norm of the perpendicular pure-state derivative. The generator has finite second moment, so the purified curve is differentiable in Hilbert norm. Since $I_\pi/N = \pi^2/\sqrt{Na^2} \to 0$, (26) yields $\liminf Na^2 R_Q \geq 1/4$.

The unshrunk T_N supplies the matching upper directly from (19):

$$Na^2 \operatorname{Var}(T_N)/F^2 \leq \frac{1}{4} + \frac{13}{2}a + \frac{29}{2}a^2 + \frac{3(1 + 4a + 7a^2)}{N + 2}. \tag{27}$$

The adaptive procedure (6) also attains this constant: multiply (24) by Na^2. Its two terms are respectively $O(\sqrt{a})$ and $O((ma)^{-1})$, while $Na^2 e(a, L) \to 0$ and $Na^2 \rho(La^2) \to 1/4$. Minkowski proves the same upper for $\widehat{Y}$.

6.3 The heterodyne constant

For $N \geq 2$ define the unbiased pair statistic

$$U_N = \frac{1}{N(N-1)} \sum_{r \neq s} \left[2(|Z_r|^2 - 1)(|Z_s|^2 - 1) + \Re(Z_r^2 \overline{Z_s^2}) \right].$$

Write $v = t + 1$. Its exact variance is

$$\begin{aligned} \operatorname{Var} U_N &= \frac{B_1}{N} + \frac{B_2}{N(N-1)}, \\ B_1 &= 16t^2(v^2 + u) + 32tvu + 4uv^2 + 4u^2, \\ B_2 &= 12v^4 + 48v^2u + 12u^2. \end{aligned} \tag{28}$$

To verify the formula, put $X = |Z|^2 - v$, $P = Z^2 - c$. Wick moments give $\mathbb{E}X^2 = v^2 + u$, $\mathbb{E}|P|^2 = 2v^2$, $\mathbb{E}P^2 = 2c^2$, and $\mathbb{E}XP = 2vc$. The first Hoeffding projection is $2tX + \Re(\bar{c}P)$; the canonical second kernel is $2X_1X_2 + \Re(P_1\overline{P}_2)$. Their variances are $B_1/4$ and $B_2/2$, respectively, proving (28).

Separating $16t^2 + 4u$ in B_1 and using the inequalities for (19) gives

$$\frac{B_1}{F^2} \leq \frac{4}{t^2} + \frac{13}{t} + \frac{29}{2}, \qquad \frac{B_2}{F^2} \leq 3(1 + t^{-1})^4 + 6(1 + t^{-1})^2 + 12.$$

Consequently, uniformly over the entire state family,

$$Na^2 \operatorname{Var}(U_N)/F^2 \leq 4 + 13a + \frac{29}{2}a^2 + \frac{3(1 + a)^4 + 6a^2(1 + a)^2 + 12a^4}{(N - 1)a^2}, \tag{29}$$

which tends to four when $Na^2 \to \infty$.

For the lower, restrict to thermal states $c = 0$. Heterodyne information for t is $N/(1+t)^2$. Take the prior of half-width $a\delta$ centered at $a(1+2\delta)$, with the same δ as above. It lies entirely above a; on its support $g = 2t^2$, $g'/a \to 4$, $\sup g^2/a^4 \to 4$, and $I_\pi/N \to 0$. Equation (26) gives $\liminf Na^2 R_H \geq 4$, matching (29) and completing Theorem 2.2.

The lower-bound witnesses differ: the quantum path has $t/a \to 1$ and $|c|/a \to \sqrt{2}$, whereas the heterodyne path is thermal. These witnesses do not assert uniqueness of a least-favorable state. A local comparison at one fixed covariance would not establish the global minimax ratio.

7 Consequences and limits

The critical curve identifies how bias should be used when weak squeezing contributes an unresolved quartic energy. Its Poisson reduction retains collective pair information while thermal occupation is learned on a faster scale. The global pilot construction is essential: a compact boundary approximation alone would not prove minimax attainability over all occupations and squeezing strengths.

Unknown squeezing therefore has two distinct resource effects. It raises the quantum occupation exponent to the heterodyne order, as already known [6], but leaves a sharp factor-sixteen advantage at low brightness and vanishing risk. The exact critical curve and attaining procedure quantify the remaining measurement advantage for one centered Gaussian mode and its quartic energy.

Reproducibility. The accompanying standard-library program checks the quartic Wick coefficients, the exact heterodyne projection variances, the uniform algebraic bounds and the equalizer identities using exact rational arithmetic. The analytic arguments establish the asymptotic minimax statements.

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

References

[1] C. Butucea, M. Guţă and M. Nussbaum, *Local asymptotic equivalence of pure states ensembles and quantum Gaussian white noise*, Ann. Statist. **46** (2018), 3676–3706. `https://doi.org/10.1214/17-AOS1672`.

[2] S. Chen, A. A. Mele, F. A. Mele and J. Preskill, *When are bosonic Gaussian states classical to learn?*, arXiv:2609.26705v1 (2026). `https://arxiv.org/abs/2609.26705v1`.

[3] H. Dai and Y. Zhang, *Uncovering non-Gaussianity through multi-copy symmetries*, arXiv:2608.03755v1 (2026). `https://arxiv.org/abs/2608.03755v1`.

[4] R. D. Gill and B. Y. Levit, *Applications of the van Trees inequality: a Bayesian Cramér–Rao bound*, Bernoulli **1** (1995), 59–79. `https://doi.org/10.2307/3318681`.

[5] O. Pinel, P. Jian, N. Treps, C. Fabre and D. Braun, *Quantum parameter estimation using general single-mode Gaussian states*, Phys. Rev. A **88** (2013), 040102(R). `https://doi.org/10.1103/PhysRevA.88.040102`.

[6] S. Schubr, *Optimal unbiased Gaussian quantum estimation and the cost of unknown squeezing*, author-published edition 1, Luminary Oddities (2026). https://research.luminaryoddities.com/optimal-unbiased-gaussian-quantum-estimation-unknown-squeezing/v/1.

[7] 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.

[8] N. Tripier-Mondancin, I. Karuseichyk, M. Walschaers, V. Parigi and N. Treps, *Optimal moment-based characterization of a Gaussian state*, Phys. Rev. Lett. **135** (2025), 110805. https://doi.org/10.1103/1kln-n48s.