# Finite-copy risk for polynomial energies of thermal Gaussian states

Scott Schubr
Luminary Oddities

October 30, 2026

**Abstract**

We study how many independent copies of a centered gauge-invariant thermal Gaussian state are needed to estimate a homogeneous normally ordered polynomial energy. General quantum symmetry averaging provides an unbiased observable; here we determine its exact finite-copy risk, including all contraction orders and the dependence on the polynomial's coefficient tensor. The construction is valid even when the number of copies is smaller than the degree, and the estimator is a spectral measurement of the averaged observable. A lower spectral bound on the occupation covariance gives a sharp dimension-independent worst-case variance, attained by a single-mode monomial, and the exact worst-polynomial minimax risk among globally unbiased estimators. Allowing bias, matching universal bounds identify the joint degree, copy-count and occupation scale for consistent estimation; in the opposite regime the minimax relative risk tends to one. Worked examples show how multimode coefficient structure reduces variance and how the uniform bounds translate into sufficient copy budgets. These results provide finite-resource benchmarks for high-order thermal-state moment estimation.

**Keywords.** Quantum estimation; thermal Gaussian states; polynomial moments; minimax risk; symmetry averaging.
**MSC 2020.** 62F12, 62F10, 81P50.

## 1 Introduction

A thermal optical field can be interrogated through normally ordered moments. For example, a single-mode factorial photon moment has the form $\langle (b^\dagger)^n b^n \rangle$, while a two-mode intensity correlation is $\langle b_1^\dagger b_2^\dagger b_1 b_2 \rangle$. Both belong to the family $\operatorname{Tr} \rho f(b)^\dagger f(b)$, where $f$ is a known homogeneous holomorphic polynomial. We ask how accurately such an energy can be estimated from a finite number of independent copies of an unknown thermal Gaussian state, allowing collective measurements.

The positive coherent-state representation makes the target a classical complex Gaussian moment. It does not give the observer samples of the latent coherent amplitudes: the data consist of quantum states. This distinction changes the fluctuation calculation. Moreover, high-order moments require control beyond the first term of a fixed-degree large-sample expansion. When the degree grows with the number of copies, the omitted contraction orders can contribute on the same scale as the leading term.

Our answer has two complementary parts. An exact variance identity describes the dependence on the energy observable through reduced matrices of its covariance-weighted coefficient tensor. A sharp uniform bound then describes the hardest polynomial, independently of the number of modes. With degree $n$, copy count $N$ and a positive lower occupation bound $a$, the governing scale is $(1 + a^{-1})n^2/N$. The minimax relative squared error vanishes exactly when this scale tends to zero, and tends to one when the scale diverges. Thus the results supply a resource benchmark for

increasing-order moment estimation, as well as a way to distinguish particular observables that are easier than the worst case.

**Relation to existing estimation results.** The ingredients have substantial precedents. Guţă and Butucea introduced quantum U-statistics and their Hoeffding analysis [2]. Tsang's quantum Rao–Blackwell theorem gives invariant-state averaging as an estimator improvement [5, Section 5.3]. Dasgupta, Warsi and Chatterjee relate marginal kernels to functional gradients and prove uniqueness, efficiency and concentration results for finite-dimensional quantum U-statistics [1]. Normally ordered moment estimation for positive coherent mixtures is developed within quantum semiparametric estimation [6]; scalar thermal photon-count information is analyzed by Nair and Tsang [3]. The exact classical Gaussian-energy experiment with direct amplitude access is treated in [4].

Three distinctions locate the present result. First, the classical experiment in [4] observes Gaussian amplitudes, whereas here only their thermal-state mixture is available; bosonic contractions introduce the occupation-dependent weights in the risk. Second, the general averaging theorem in [5] supplies the construction principle, while the calculation here evaluates every contraction order for this polynomial family. Third, polynomial degree here refers to $f$ in the annihilation operators. The restriction to Gaussian states permits repeated copy labels and an unbiased observable for every $N \geq 1$, including $N < n$. This differs from extending a degree-$k$ polynomial functional over all finite-dimensional density matrices by a distinct-copy U-statistic [2, 1].

The additional conclusions are therefore an exact multimode risk identity, sharp uniform unbiased risk for every copy count, and a uniform degree/copy/occupation frontier in the declared thermal-state experiment. Sharp unbiased optimality is a worst-polynomial statement; the arbitrary-bias result is a uniform minimax order and limiting-phase statement. The exact formula separately retains the observable-specific variance. Section 6 illustrates these different guarantees with photon-count and two-mode product examples and sufficient copy budgets.

**Measurement model and proof route.** The resource counted is $N$ state copies. The allowed measurement is the spectral measurement of a positive collective observable obtained by averaging single-output polynomial energies over passive copy mixing. The averaging defines the observable being measured; random selection of a mixing direction gives a different experiment. The finite-photon-sector construction below specifies the measurement mathematically. This is a benchmark under unrestricted collective spectral measurements.

Copy-mixing invariance first gives unbiasedness. Bosonic Gaussian contractions and complex-sphere moments then give the exact variance. A scalar thermal subfamily, reduced to its complete negative-binomial photon-count statistic, gives sharp unbiased optimality. Finally scalar two-point testing supplies the biased minimax bounds and the consistency and collapse laws.

## 2 Experiment and main results

Let $d, n, N$ be positive integers. On $d$ bosonic modes write $b_1, \ldots, b_d$ for annihilation operators, satisfying $[b_i, b_j^\dagger] = \delta_{ij}\mathbb{I}$. For a positive definite Hermitian matrix $A$, define

$$\rho_A = \int_{\mathbb{C}^d} |z\rangle\langle z| \, d\,\mathrm{CN}(0, A)(z), \qquad A_{ij} = \mathbb{E}[Z_i \overline{Z_j}]. \tag{1}$$

Thus $A$ is the occupation covariance, with vacuum noise accounted for by the commutation relations. All states in this paper have the representation (1).

Let $f \neq 0$ be homogeneous and holomorphic of degree $n$, with symmetric coefficients

$$f(z) = \sum_{I \in [d]^n} c_I z_{i_1} \cdots z_{i_n}. \tag{2}$$

The target and loss are

$$F_f(A) = \operatorname{Tr} \rho_A f(b)^\dagger f(b) = \mathbb{E}|f(Z)|^2, \qquad L_A(y) = \left( \frac{y}{F_f(A)} - 1 \right)^2. \tag{3}$$

The equality follows from $b_i|z\rangle = z_i|z\rangle$ and positivity. We observe $\rho_A^{\otimes N}$, allowing any collective POVM and real-valued decision. The known parameter restriction is

$$A \succeq aI_d, \qquad a > 0, \qquad \gamma = 1 + a^{-1}. \tag{4}$$

This is a spectral floor in every mode. The estimator may depend on $f, d, N, a$.

For copy $r$, denote its annihilation operator by $b_{i,r}$. For an ordered tuple $\boldsymbol{r} \in [N]^n$ set

$$C_{\boldsymbol{r}} = \sum_I c_I \prod_{j=1}^n b_{i_j, r_j}, \qquad T_f = \frac{n!}{(N)^{\overline{n}}} \sum_{\boldsymbol{r} \in [N]^n} C_{\boldsymbol{r}}^\dagger C_{\boldsymbol{r}}. \tag{5}$$

Here $(N)^{\overline{k}} = N(N+1)\cdots(N+k-1)$, and $(N)^{\overline{0}} = 1$. Repeated indices $r_j$ are included.

For stating the exact formula, choose a mode basis diagonalizing $A$, with eigenvalues $\lambda_i$, and transform $c$ into that basis. Define

$$t_I = c_I \prod_{j=1}^n \sqrt{\lambda_{i_j}}, \qquad \psi = \frac{t}{\|t\|} \in \operatorname{Sym}^n(\mathbb{C}^d), \qquad R_k = \operatorname{Tr}_{n-k} |\psi\rangle\langle\psi|, \qquad D_A = \operatorname{diag}(1 + \lambda_i^{-1}). \tag{6}$$

The coordinate change in (6) is used in the analysis; the observable (5) is independent of the unknown $A$.

**Theorem 2.1** (Exact variance and sharp uniform bound)**.** *For all $d, n, N \geq 1$, the positive observable $T_f$ is unbiased for $F_f(A)$ and has finite second moment. Its relative variance is*

$$\frac{\operatorname{Var}_A(T_f)}{F_f(A)^2} = \sum_{k=1}^n \binom{n}{k}^2 \frac{k!}{(N)^{\overline{k}}} \operatorname{Tr}\left( R_k^2 D_A^{\otimes k} \right). \tag{7}$$

*Under* (4)*, put*

$$S(N, n, a) = \sum_{k=1}^n \binom{n}{k}^2 \frac{k! \gamma^k}{(N)^{\overline{k}}}. \tag{8}$$

*Then the relative variance is at most $S(N, n, a)$. Equality holds for $d = 1$, $f(z) = z^n$, $A = a$, so the bound is sharp simultaneously in all parameters.*

Define the worst-polynomial minimax risk

$$\mathcal{R}(N, n, a) = \sup_{d \geq 1} \sup_{0 \neq f \in \operatorname{Sym}^n(\mathbb{C}^d)} \inf_{\mathsf{M}, Y} \sup_{A \succeq aI_d} \mathbb{E}_A L_A(Y). \tag{9}$$

The notation for $f$ in (9) identifies polynomials with their coefficient tensors. Define $\mathcal{R}_{\mathrm{ub}}$ by restricting the infimum to estimators unbiased at every allowed $A$.

**Theorem 2.2** (Unbiased optimum and biased minimax order)**.** *For all $n, N \geq 1$ and $a > 0$,*

$$\mathcal{R}_{\mathrm{ub}}(N, n, a) = S(N, n, a). \tag{10}$$

*With $x = \gamma n^2/N$,*

$$\frac{7}{2304}\min(1, x) \;\leq\; \mathcal{R}(N, n, a) \;\leq\; \frac{S}{1+S} \;\leq\; 1 - e^{-x} \;\leq\; \min(1, x). \tag{11}$$

*The upper bound $S/(1+S)$ is attained as a risk bound by a spectral measurement of $T_f/(1+S)$.*

**Theorem 2.3** (Consistency and collapse)**.** *For any sequence of allowed triples $(N_j, n_j, a_j)$, write $x_j = (1 + a_j^{-1})n_j^2/N_j$. Then*

$$\mathcal{R}(N_j, n_j, a_j) \longrightarrow 0 \quad\Longleftrightarrow\quad x_j \longrightarrow 0, \tag{12}$$

$$x_j \longrightarrow \infty \quad\Longrightarrow\quad \mathcal{R}(N_j, n_j, a_j) \longrightarrow 1. \tag{13}$$

These results quantify the resources of collective quantum estimation in the declared state model. The exact formula also retains dependence on the polynomial through the reduced-tensor purities in (7).

# 3 Symmetry averaging and the exact risk formula

## 3.1 Construction, domains and first moment

Let $v$ be uniform on the complex unit sphere in $\mathbb{C}^N$ and write

$$b_i(v) = \sum_{r=1}^{N} v_r b_{i,r}, \qquad H_v = f(b(v))^\dagger f(b(v)). \tag{14}$$

The combined modes have canonical commutators, and their state in $\rho_A^{\otimes N}$ is again $\rho_A$. The sphere identity

$$\mathbb{E}_v[v_{r_1} \cdots v_{r_n} \overline{v_{s_1} \cdots v_{s_n}}] = \frac{1}{(N)^{\overline{n}}} \sum_{\pi \in S_n} \prod_{j=1}^{n} \delta_{r_j, s_{\pi(j)}} \tag{15}$$

follows by writing a standard complex Gaussian vector as its independent radius and direction. Its squared radius has $n$th moment $(N)^{\overline{n}}$, while the Gaussian numerator is the Wick sum in (15). Symmetry of $c$ and commutation of annihilators give

$$T_f = \mathbb{E}_v H_v. \tag{16}$$

In particular $\mathbb{E}_A T_f = F_f(A)$, since this is true for every $H_v$.

For a precise spectral construction, first work on finite total-photon sectors in the $dN$ modes. Each $C_{\boldsymbol{r}}^\dagger C_{\boldsymbol{r}}$ preserves the total number in each copy. Every sector with bounded total photon number is finite dimensional, and (5) defines a positive matrix there. The direct sum of the fixed-total-number matrices defines a positive self-adjoint operator with the finite-particle vectors as a core. Its spectral measurement is the estimator used here. Thermal states on finitely many modes have all polynomial number moments finite, which pays its second moment and justifies passage from the sector identities to the full state. Formula (16) refers to an averaged observable.

### 3.2 All contraction orders

We prove Theorem 2.1 by computing $\mathbb{E}_A(H_v H_w)$ for two unit vectors $v, w$. In an eigenbasis of $A$, the nonzero two-point contractions are

$$\mathbb{E}_A[b_i(v)^\dagger b_j(w)] = \delta_{ij}\lambda_i \sum_r \overline{v_r} w_r,$$
$$\mathbb{E}_A[b_i(v) b_j(w)^\dagger] = \delta_{ij}(1+\lambda_i) \sum_r v_r \overline{w_r}.$$

The product $H_v H_w$ has four groups, each of size $n$: creators at $v$, annihilators at $v$, creators at $w$, and annihilators at $w$. In a Wick pairing let $k$ be the number of contractions from the first annihilator group to the second creator group. Balance forces $k$ contractions in the reverse direction between the other two groups. The remaining $n-k$ contractions on each side are internal. The number of labeled pairings in this class is

$$\binom{n}{k}^4 (k!)^2((n-k)!)^2 = (n!)^2 \binom{n}{k}^2. \tag{17}$$

Symmetry of $c$ makes their coefficient contractions identical.

After whitening by (6), the internal contractions yield the $k$-body reduction $R_k$. More explicitly, for ordered multi-indices $I, J \in [d]^k$,

$$(R_k)_{I,J} = \sum_{L \in [d]^{n-k}} \psi_{I,L} \overline{\psi_{J,L}}.$$

The coefficient contraction for the $k$th class, divided by $\|t\|^4$, is

$$\sum_{I,J} \left( \prod_{\ell=1}^{k} (1 + \lambda_{i_\ell}^{-1}) \right) |(R_k)_{I,J}|^2 = \operatorname{Tr}(R_k^2 D_A^{\otimes k}). \tag{18}$$

The cross-copy contractions contribute $|\langle v, w\rangle|^{2k}$. Since $F_f(A) = n!\|t\|^2$, equations (17)–(18) give

$$\frac{\mathbb{E}_A(H_v H_w)}{F_f(A)^2} = \sum_{k=0}^{n} \binom{n}{k}^2 |\langle v, w\rangle|^{2k} \operatorname{Tr}(R_k^2 D_A^{\otimes k}), \tag{19}$$

where the $k = 0$ term is one. All terms are nonnegative real numbers. Finite thermal fourth moments of $f(b(v))$ also bound the ordered product by state Cauchy–Schwarz, uniformly over the compact sphere, so the integrals in (19) are legitimate.

For independent uniform $v, w$, unitary invariance and (15) imply

$$\mathbb{E}|\langle v, w\rangle|^{2k} = \frac{k!}{(N)^{\bar{k}}}.$$

This includes $N = 1$, when both sides equal one. Integrating (19) and subtracting its $k = 0$ term proves (7). Finally $D_A \preceq \gamma I$ and $\operatorname{Tr} R_k^2 \le 1$ prove (8). A single-mode coefficient tensor is pure at every reduction and $A = a$ makes every weight equal $\gamma$, proving simultaneous equality. This completes the proof of Theorem 2.1.

# 4 Sharp unbiased risk

In one mode the thermal state is diagonal in the photon basis, with geometric probabilities

$$p_t(k) = \frac{t^k}{(1+t)^{k+1}}, \qquad k \geq 0. \tag{20}$$

For $N$ copies, the total count $K$ has negative-binomial probabilities

$$\Pr_t(K = k) = \binom{N+k-1}{k} \frac{t^k}{(1+t)^{N+k}}.$$

Conditioned on $K$, the individual counts are uniform over compositions of $K$ into $N$ nonnegative parts. This distribution is independent of $t$. Since the quantum family is commuting, every quantum measurement is a classical randomization of the full count vector. Conditional averaging given $K$ therefore reduces the variance of any unbiased estimator.

For $f(z) = z^n$, formula (5) simplifies to

$$T_f = \frac{n!(K)_{\underline{n}}}{(N)^{\overline{n}}}, \qquad F_f(t) = n!t^n. \tag{21}$$

The identity follows by counting ordered choices of $n$ photons across the $N$ modes, or by the sphere average. The factorial moment $\mathbb{E}(K)_{\underline{n}} = (N)^{\overline{n}} t^n$ verifies unbiasedness directly.

The negative-binomial family is complete even after restriction to $t \geq a$. To see this, put $q = t/(1+t)$. If an integrable function $g(K)$ has expectation zero for every allowed $t$, its power series

$$\sum_{k \geq 0} g(k) \binom{N+k-1}{k} q^k$$

vanishes on $q \in (a/(1+a), 1)$. Integrability at every larger $q < 1$ gives convergence throughout the unit disk. Analytic uniqueness forces all coefficients to vanish. Consequently (21) is the unique unbiased function of $K$, and conditional averaging proves that it has minimum variance among all scalar unbiased measurements and decisions.

At $t = a$ its relative variance is $S(N, n, a)$ by Theorem 2.1. This proves the lower bound in (10); the estimator of Theorem 2.1 proves the upper bound for every $d, f$. Thus (10) holds with the displayed order of optimization.

For comparison, the factorial-product identity gives the scalar relative second moment in the alternative form

$$\mathbb{E}_t \left( \frac{T_f}{F_f(t)} \right)^2 = \sum_{j=0}^{n} \binom{n}{j}^2 j! \frac{(N)^{\overline{2n-j}}}{((N)^{\overline{n}})^2} t^{-j}. \tag{22}$$

Equating (22) with one plus (7) supplies a scalar identity between the falling-factorial and contraction expansions.

# 5 Biased minimax bounds and asymptotic phases

## 5.1 Upper bound

If $T$ is unbiased and has relative variance at most $S$, its rescaling $T/(1+S)$ has relative risk at most

$$\frac{S}{(1+S)^2} + \left( \frac{1}{1+S} - 1 \right)^2 = \frac{S}{1+S}.$$

Moreover $\binom{n}{k}^2 k! \le n^{2k}/k!$ and $(N)^{\overline{k}} \ge N^k$. Summing the exponential series gives $S \le e^x - 1$ and proves all upper bounds in (11).

### 5.2 A uniform two-point lower bound

It suffices to use the scalar family (20), since it is included in (9). Compare $t_0 = a$ with $t_1 = a(1+h)$. Direct computation gives

$$\mathrm{KL}(p_{t_0} \| p_{t_1}) = (1+a)\log\left(1 + \frac{ah}{1+a}\right) - a\log(1+h). \tag{23}$$

Its derivative in $h$ is

$$\frac{ah}{(1+h)(1+a+ah)} \le \frac{h}{\gamma},$$

so the one-copy divergence is at most $h^2/(2\gamma)$ for every $h \ge 0$. Set

$$h = \frac{1}{4}\min\left(\frac{1}{n}, \sqrt{\frac{\gamma}{N}}\right).$$

The $N$-copy divergence is at most $1/32$, and Pinsker's inequality bounds the total variation distance by $1/8$. The target ratio is $r = (1+h)^n$. Its relative separation

$$\delta = \frac{r-1}{r+1}$$

satisfies $\delta \ge nh/3$, because $nh \le 1/4$, $r \le e^{1/4}$ and $r - 1 \ge nh$.

Threshold any estimator at the harmonic target midpoint $2F_0F_1/(F_0+F_1)$. Each classification error entails relative squared loss at least $\delta^2$. The sum of the testing errors is at least one minus total variation, so the maximum of the two risks is at least $\delta^2(1-1/8)/2$. Since $n^2h^2 = \min(1,x)/16$, this proves the lower bound $7\min(1,x)/2304$. A quantum POVM cannot improve the distance of the commuting family. Theorem 2.2 follows.

### 5.3 Consistency and complete collapse

Equation (12) follows immediately from the upper and lower bounds in (11). To prove (13), take along the given sequence

$$h = \frac{x^{1/4}}{n}, \qquad t_0 = a, \quad t_1 = a(1+h).$$

The divergence bound above gives $N\,\mathrm{KL} \le 1/(2\sqrt{x})$, hence the total variation $\eta_x$ tends to zero. Meanwhile $r = (1+h)^n \ge 1 + x^{1/4}$ tends to infinity.

Let $P_0, P_1$ be the full photon-count laws and let their common subprobability measure have density $\min(dP_0, dP_1)$; its mass is $1 - \eta_x$. For every real $y$, direct minimization gives

$$\frac{1}{1+r}\left(\frac{y}{F_0} - 1\right)^2 + \frac{r}{1+r}\left(\frac{y}{F_1} - 1\right)^2 \ge \left(\frac{r-1}{r+1}\right)^2. \tag{24}$$

The minimizing $y$ is the harmonic midpoint. Integrating (24) over the common measure bounds every estimator's maximum risk from below by

$$(1-\eta_x)\left(\frac{r-1}{r+1}\right)^2 \longrightarrow 1.$$

Classical randomization, and therefore arbitrary quantum measurement on this scalar subfamily, preserves this lower bound. The identically zero decision has risk one, proving Theorem 2.3.

# 6 Worked examples and resource consequences

## 6.1 Photon counts and fewer copies than the degree

For $f(z) = z^n$ in one mode, the target is $n!t^n$ and the total photon count $K$ over $N$ copies has the negative-binomial law used above. The estimator is the falling-factorial statistic (21). At degree two it becomes

$$T_f = \frac{2K(K-1)}{N(N+1)}, \qquad F_f(t) = 2t^2,$$
$$\frac{\mathrm{Var}_t(T_f)}{F_f(t)^2} = \frac{4\gamma_t}{N} + \frac{2\gamma_t^2}{N(N+1)}, \qquad \gamma_t = 1 + t^{-1}.$$

For $N = 1$ this is the single-copy observable $K(K-1)$. Its expectation already gives the degree-two target because the thermal distribution relates all factorial moments to the same occupation parameter $t$. Thus $N < n$ is compatible with unbiasedness in this restricted state family. Multiple copies reduce its fluctuations, with both displayed orders contributing to the exact variance. At degree one the sample mean photon count has relative variance $\gamma_t/N$.

## 6.2 A two-mode product and the role of coefficient geometry

Take $f(z) = z_1 z_2$ and $A = tI_2$. The target is $F_f(tI_2) = t^2$. In the tensor convention (2), the only nonzero coefficients are $c_{12} = c_{21} = 1/2$, so

$$\psi = \frac{|1,2\rangle + |2,1\rangle}{\sqrt{2}}, \qquad R_1 = \frac{I_2}{2}, \qquad \mathrm{Tr}\, R_1^2 = \frac{1}{2}, \qquad \mathrm{Tr}\, R_2^2 = 1.$$

Substitution in (7) gives

$$\frac{\mathrm{Var}_{tI_2}(T_f)}{F_f(tI_2)^2} = \frac{2\gamma_t}{N} + \frac{2\gamma_t^2}{N(N+1)}. \tag{25}$$

At the same occupation, the single-mode monomial $z^2$ has leading term $4\gamma_t/N$. The product energy has half that leading variance, while its second-order term is unchanged. For $t = 1$ the exact relative variances are

| $N$ | 1 | 2 | 3 |
|---|---|---|---|
| $z_1 z_2$ | 8 | 10/3 | 2 |
| $z^2$ | 12 | 16/3 | 10/3 |

These values compare the particular unbiased observables at the stated isotropic covariance. The reduced matrices belong to the polynomial coefficient tensor; their purities encode how its indices are distributed across modes. This is the observable-specific information retained by the exact formula and lost on replacing every purity by its worst-case value one. The uniform minimax result instead takes the supremum over polynomials and allowed covariances.

## 6.3 Sufficient copy budgets and growing degree

For a desired relative squared-error bound $0 < r < 1$, Theorem 2.2 gives the sufficient condition

$$N \geq \frac{(1 + a^{-1})n^2}{-\log(1-r)}. \tag{26}$$

Indeed this makes $x \le -\log(1-r)$, and the observable $T_f/(1+S)$ has risk at most $1-e^{-x} \le r$. Evaluating the finite sum $S$ can give a tighter sufficient budget. Formula (26) is a conservative design guarantee under the stated collective measurement model.

For example, at degree $n = 10$ and floor $a = 1$, choosing $N = 2000$ gives $x = 0.1$ and a uniform relative mean-square bound $1 - e^{-0.1} < 0.096$. At floor $a = 0.1$, $N = 11000$ gives the same bound. The comparison displays the occupation cost while keeping degree, loss and measurement class fixed. The rate lower bound and Theorem 2.3 establish the corresponding order and vanishing-risk criterion for the hardest polynomial.

For fixed $n, a$ and $N \to \infty$, the sharp unbiased worst-case variance has leading coefficient $\gamma n^2/N$. For growing degrees the exact finite sum is more informative. If $n \to \infty$ and $x = \gamma n^2/N \to \lambda \in (0, \infty)$, each fixed-$k$ summand in (8) tends to $\lambda^k/k!$. Exponential domination then gives $S \to e^\lambda - 1$, and the rescaled estimator has a worst-case risk bound tending to $1 - e^{-\lambda}$. This limit shows explicitly why retaining all contraction orders matters at the degree/copy transition. Theorem 2.3 separately characterizes vanishing risk and collapse for the full minimax problem.

**Computational verification.** The accompanying exact-arithmetic program checks the contraction formula against an independent normal-order expansion for nontrivial two-mode polynomials and copy counts, including the product example (25). It also checks the scalar negative-binomial identity (22). The proofs above establish the results for arbitrary finite dimension and degree.

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

## References

[1] A. Dasgupta, N. A. Warsi and P. Chatterjee, *Uniqueness, Cramér–Rao efficiency and concentration bounds for quantum U-statistics*, arXiv:2609.08745v2 (2026). https://arxiv.org/abs/2609.08745v2.

[2] M. Guţă and C. Butucea, *Quantum U-statistics*, J. Math. Phys. **51** (2010), 102202. https://doi.org/10.1063/1.3476776.

[3] R. Nair and M. Tsang, *Quantum optimality of photon counting for temperature measurement of thermal astronomical sources*, Astrophys. J. **808** (2015), 125. https://doi.org/10.1088/0004-637X/808/2/125.

[4] S. Schubr, *Exact minimax estimation of homogeneous complex Gaussian energies*, author-published edition 2, Luminary Oddities (2026). https://research.luminaryoddities.com/exact-minimax-homogeneous-complex-gaussian-energies/v/2.

[5] 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.

[6] M. Tsang, F. Albarelli and A. Datta, *Quantum semiparametric estimation*, Phys. Rev. X **10** (2020), 031023. https://doi.org/10.1103/PhysRevX.10.031023.