Robustness of quantum-enhanced adaptive phase estimation

Pantita Palittapongarnpim[1] and Barry C. Sanders[1,2]

[1]*Institute for Quantum Science and Technology, University of Calgary, Alberta, Canada T2N 1N4*

[2]*Program in Quantum Information Science, Canadian Institute for Advanced Research, Toronto, Ontario, Canada M5G 1M1*

(Received 26 November 2018; published 8 July 2019)

We aim to construct tests for evaluating whether policies for adaptive quantum-enhanced metrology are robust against unknown phase noise and are tractable to construct and to execute. Specifically, one of our tests determines scaling of phase-estimate precision with respect to photon number; the other two tests concern resource complexity with respect to the training time required to construct the policy and the execution time for the policy so constructed. The robustness test is performed on quantum-enhanced adaptive phase estimation by simulating the scheme under four phase-noise models corresponding to normal-distribution noise, random-telegraph noise, skew-normal-distribution noise, and log-normal-distribution noise. Control policies are devised either by an evolutionary algorithm under the same noisy conditions, albeit ignorant of its properties, or a Bayesian-based feedback method that assumes no noise. We have introduced an approach to evaluating quantum-control policies for metrology that relies on testing against unusual phase-noise models and accepting that the policies are robust only if scaling of phase-estimate precision beats the standard quantum limit and policy-design resource complexity, in the presence of phase noise, is polynomial in the number of photons.

DOI: 10.1103/PhysRevA.100.012106

I. INTRODUCTION

Quantum-enhanced metrology (QEM) employs a quantum state of N particles as a resource to estimate an unknown parameter ϕ with the goal of attaining imprecision,

$$\Delta\tilde{\phi} \in O(N^{-\wp}), \tag{1}$$

that asymptotically surpasses the standard quantum limit (SQL) $\wp = 1/2$ [1,2] but saturates at or below the Heisenberg limit (HL) $\wp = 1$ [3,4]. Our expression (1) differs from usual proportionality expressions in the literature, e.g., $\Delta\tilde{\phi} \propto N^{-\wp}$, by explicitly recognizing the relevance of lower-order terms [4] through the use of the big-O notation [5].

QEM is vital for high-precision applications, such as gravitational wave detection [6–8], atomic clocks [9,10], and magnetometry [4,11] whose systems are operating at the limit of their power tolerance. Some schemes consider ideal measurements that typically involve measuring multiple particles simultaneously [12,13], whereas adaptive QEM (AQEM) focuses on single-particle measurements augmented by feedback such that the SQL is beaten and the HL is approached [14–16].

AQEM performance critically depends on the choice of policy ϱ [17], which can be obtained by optimizing a known mathematical model [18–20] or by a trial-and-error approach [21–23]. Whereas policies from these methods are resistant to known noise models [18], whether they are robust against unknown noise is yet unstudied but critical property of a QEM scheme as noise can destroy the entanglement advantage and restore the SQL [24,25].

Our aim is to test the robustness of AQEM policies in the presence of noise with unknown properties. To this end, we start with an adaptive procedure that is known to attain quantum-enhanced precision without noise, then include Gaussian phase noise as well as skewed, long-tail, and discrete phase distributions to observe scaling of phase-estimate precision as the noise level increases. When we refer to the scaling of phase-estimate precision, it is with respect to photon number, in accordance with the usual practice, and we employ numerical simulation up to some maximum photon number $N_{\max}$ to determine this scaling. We use the term "trend" to refer to the apparent extrapolation of the plot beyond $N_{\max}$. In AQEM, proofs of asymptotic power-law or otherwise behavior are not now known.

Our test focuses on quantum-enhanced adaptive phase estimation (QEAPE), whose policies have been devised using Bayesian techniques [16,26] and by searching the policy space using an evolutionary algorithm (EA) [21–23]. Bayesian technique computes feedback based on a trusted, noiseless quantum model, whereas evolutionary algorithms [27] devise policies for feedback based on fitness of policies obtained through trial and error. This evolution is, by design, ignorant of the quantum-dynamical nature but employs heuristics to shrink the search space.

Here both Bayesian techniques and evolutionary algorithms are applied to QEAPE, including phase noise, which could arise from interferometer path-length fluctuations [28,29]. Sources for these noise processes can be grouped into mechanical (e.g., vibrations of optical components or mounts), acoustic (e.g., sound or air ventilation), thermal (e.g., heat sources and cooling), optomechanical, and other categories of fluctuations [30]. As modeling all these sources *ab initio* can be computationally expensive, which is especially daunting for feedback loops due to the simulation being executed multiple times, we employ simple phase-noise models to test robustness, time to design, and time to execute for candidate policies.

Typically, noise is assumed to be normal as a result of the central-limit theorem [31]. Periodicity of phase makes the normal distribution problematic unless the noise is small compared to 2π radians, which we assume here; technically, we would use the wrapped-up normal distribution [32]. As our aim is to test robustness for unknown noise, we consider three other noise distributions for our test: random telegraph [33], skew-normal [34], and log-normal [35] noise. The random-telegraph noise simulates a discrete noise process. Skew-normal and log-normal distributions represent asymmetric noise, which serve as distinct generalizations of the normal distribution. Both distributions are used to simulate noise in detectors and electronics [36–38].

For AQEM, we seek an efficient procedure that beats the SQL, and we choose ϱ that requires the least resource to run. We assess the policy-generating procedure according to the complexity of its time cost [39], which is evaluated by the scaling in the number of operations with the number of particles N. Here we consider two policy-design procedures, namely, a policy-search method based on differential evolution (DE) [40], which is a variant of evolutionary algorithms, and a method based on Bayesian inference, resulting in one policy designed by each method. To determine which policy is superior, we compare the complexity in space and time cost [5]. Thus, we are able to assess and compare the costs for generating policies and determine the best policy.

Through our analysis, we find that both Bayesian feedback and DE-designed policies are robust in the face of unknown phase noise. Specifically, the Bayesian method yields imprecision that approaches the HL and outperforms DE-designed policies for most noise models. This performance superiority is due to the Bayesian method effectively memorizing the measurement history through the agent's complete knowledge of the quantum state. Storing the entire model in the agent yields better imprecision scaling of phase-estimate precision but incurs higher space and operational time costs compared to the DE-derived policy.

In summary, our approach to assessing QEAPE policies is based on studying scaling of phase-estimate precision with respect to photon number and, furthermore, analyzing resource complexity given by the training time required to construct the policy and also the execution time for the policy so constructed. We have introduced an approach to evaluating quantum-control policies for metrology that relies on testing against unusual phase-noise models and accepting that the policies are robust only if scaling of phase-estimate precision beats the standard quantum limit and policy-design resource complexity, in the presence of phase noise, is poly$N$.

II. BACKGROUND

In this section, we present essential background knowledge for assessing robustness of adaptive quantum-enhanced phase estimation. We cover four key notions as subsections. In Sec. II A we discuss QEM including AQEM. Subsequently, in Sec. II B, we address a particular case of AQEM, namely, QEAPE. As noise is present for any practical QEM procedures, we briefly review the effect of noise on the imprecision scaling of phase-estimate precision in Sec. II C before we discuss phase noise models we use to evaluate policy robustness in Sec. II D.

A. Quantum-enhanced metrology

In this subsection, we explain QEM strategies, mainly collective nonadaptive [2] and adaptive measurements [41], and lower imprecision bounds attained by using the classical and quantum resources. Specifically, we focus on strategies that employ finite N input states. Whereas the collective nonadaptive strategy is useful to calculate lower bounds for imprecision, adaptive strategies offer the simpler alternative of individual-particle measurements that can achieve imprecision scaling of phase-estimate precision close to these theoretical bounds. Here we focus on single-parameter estimation; multiparameters QEM [42–44] would be the subject of future study.

1. Nonadaptive and adaptive strategies

Here we describe collective nonadaptive and adaptive QEM strategies, which are two of many types of QEM schemes. Whereas we focus on these two techniques, there are others, such as the sequential technique [45,46] and the ancilla-assisted techniques [47,48], which we do not cover.

a. Collective strategy. A collective nonadaptive scheme utilizes N d-level (typically, $d = 2$ for standard two-level atoms or two-path interferometry) particles prepared in a collective state

$$\rho \in \mathcal{S}\big(\mathscr{H}_d^{\otimes N}\big), \tag{2}$$

which is the space of positive-definite, trace-class, self-adjoint linear operators acting on a tensor product of N copies of a d-dimensional Hilbert space $\mathscr{H}_d$ [49]. The system, upon which metrology is performed, is represented as a quantum channel (completely positive trace-preserving map) $\mathcal{I}(\phi)$, acting on ρ, with ϕ being the single-unknown parameter of the channel [50]. In the special case of an isolated system without noise, decoherence or loss, the channel is represented by a unitary transformation [51]

$$\mathcal{I}(\phi)\rho = U(\phi)\rho U^\dagger(\phi) \tag{3}$$

for U being a unitary operator acting on $\mathscr{H}_d^{\otimes N}$.

After the particles exit the system, they are measured and this measurement is described as a positive-operator-valued measure [52,53], which comprises positive semidefinite operators

$$\hat{X}_x : \mathscr{H}_d^{\otimes N} \to \mathscr{H}_d^{\otimes N}, \quad \sum_x \hat{X}_x = \mathbb{1}, \tag{4}$$

assuming the measurement outcomes $\{x\}$ is a finite set. We use the notation $\mathbb{1}$ to denote an identity operator or identity matrix. This outcome (4) is random with probability

$$P_x = \mathrm{tr}(\hat{X}_x \mathcal{I}(\phi)\rho). \tag{5}$$

The measurement $\hat{X}_x$ is repeated multiple times to sample the distribution (5) sufficiently well to get good estimates, and then ϕ is inferred from these samples.

b. Bundle and individual-particle measurement. Instead of collective measurement, we can consider measuring subsets of particles, which we call bundles, and, at the extreme limit,

which is of interest here, measuring a single particle at a time. Mathematically, we split the particles into M bundles of L particles where $N = ML$ [41] so the Hilbert space can be expressed as

$$\left(\underbrace{\mathscr{H}_d^{\otimes L}}_{\text{bundle}}\right)^{\otimes M}. \tag{6}$$

In this case, both $\mathcal{I}(\phi)$ and $\hat{X}_x$ act on $\mathscr{H}_d^{\otimes L}$. For localized measurements on each bundle, the positive operator-valued measure (POVM) is

$$\bigotimes_{m=0}^{M-1} \hat{X}_{x_m}^{(m)}, \quad \hat{X}_{x_m}^{(m)} : \mathscr{H}_d^{\otimes L} \to \mathscr{H}_d^{\otimes L} \tag{7}$$

with outcomes from this tensor-product POVM being concatenations of M length L strings of d-dimensional digits,

$$\boldsymbol{x}_M = x_0 x_1 \ldots x_{M-1} \in \mathbb{N}_{d^L}^{\otimes M}, \tag{8}$$

where

$$x_m \in \mathbb{N}_{d^L} := \{0, 1, 2, \ldots, d^L - 1\} \tag{9}$$

measured from the mth bundle.

In one extreme case, each bundle contains only one particle, which leads to $M = N$ and $L = 1$. The string of outcomes becomes

$$\boldsymbol{x}_N = x_0 x_1 \ldots x_{N-1} \in \mathbb{N}_d^{\otimes N}. \tag{10}$$

The POVM is

$$\bigotimes_{m=0}^{N-1} \hat{X}_{x_m}^{(m)}, \quad \hat{X}_{x_m}^{(m)} : \mathscr{H}_d \to \mathscr{H}_d, \tag{11}$$

which is a tensor product of N qudit POVMs.

For two-level particles, the state (2) is simplified to

$$\rho \in \mathcal{S}\big(\mathscr{H}_2^{\otimes N}\big), \tag{12}$$

and the POVM simplifies from (11) to

$$\bigotimes_{m=0}^{N-1} \hat{X}_{x_m}^{(m)}, \quad \hat{X}_{x_m}^{(m)} : \mathscr{H}_2 \to \mathscr{H}_2. \tag{13}$$

The outcome (10) is simplified to

$$\boldsymbol{x}_N \in \{0, 1\}^{\otimes N}, \tag{14}$$

which is an N-bit string. Henceforth, we restrict to the $d = 2$ (two-level system), $L = 1$ (single-particle-per-bundle case) for simplicity and without loss of generality.

c. Adaptive strategy. The adaptive strategy involves incorporating quantum feedback control [54] such that the system operation depends on both the unknown parameter ϕ and a control parameter Φ_m, for some degree of freedom, on the mth bundle. We assume that incorporating a control preserves the system acting as a channel and thus we write the channel acting on the mth bundle as $\mathcal{I}(\phi; \Phi_m)$. Measurement of the mth bundle leads to an update of the control parameter to Φ_{m+1} for the next bundle.

In control theory a policy ϱ is a procedure the controller uses to modify the plant (comprises a system, actuators and sensors) based on feedback from previous actions [55]. The control-parameter update is determined by a policy

$$\varrho : (\boldsymbol{x}_m, \Phi_m) \mapsto \Phi_{m+1}, \quad \boldsymbol{x}_m = x_0 x_1 \ldots x_{m-1}, \tag{15}$$

which uses the string of outcomes $\boldsymbol{x}_m$ and the most recent control parameter Φ_m to obtain the next control parameter Φ_{m+1}. This procedure continues until reaching the final particle, i.e., the Nth particle, at which point the estimate of the unknown parameter is

$$\tilde{\phi} := \Phi_N. \tag{16}$$

Therefore, the adaptive strategy can be used for single-shot measurement, i.e., inferring ϕ from one instance of the measurement procedure.

2. Imprecision limits

Imprecision of the estimate (16) is denoted $\Delta\tilde{\phi}$ (1). Assuming the measurement is optimal and that the quantum channel is noiseless [7,56], imprecision lower bounds are calculated for classical and quantum resources. These lower bounds are the SQL and HL, respectively. In a noisy system, which pertains in practice, a QEM scheme is unlikely to saturate the bound. Despite the presence of noise, the SQL is still the bound that must be surpassed to claim QEM, which requires using quantum resources. Here we review these limits as benchmarks for both the robustness test and to compare AQEM policies.

a. Standard quantum limit. The SQL is the imprecision lower bound if classical resources are used, which means that the input state (12) is separable, i.e., unentangled [57]. The simplest case of such a separable state is a tensor product of N independent particles [1],

$$\rho = \rho_1^{\otimes N}. \tag{17}$$

Interacting this state to the quantum channel leads to the output state

$$[\mathcal{I}(\phi)\rho_1]^{\otimes N}. \tag{18}$$

Measuring this state (18) according to the POVM (13) leads to output governed by the probability distribution,

$$P(x_m) = \mathrm{tr}(\hat{X}_{x_m} \mathcal{I}(\phi)\rho_1), \tag{19}$$

which is independent and identically distributed (iid) for $m \in \{0, 1, \ldots, M-1\}$. Calculating the imprecision, such as through the central limit theorem, leads to $\wp = 1/2$ because of this iid condition [58]. The scaling also holds for the imprecision lower bound, which is calculated from (18) using the Cramér-Rao lower bound [59], and is irrespective of the quantum channel.

b. Heisenberg limit. If quantum resources are employed, e.g., squeezing [60] or entanglement [61], the SQL can be surpassed [4,8]. The lower bound to using the quantum resource can be computed from the quantum version of the Cramér-Rao lower bound [42], which depends on the input state and the quantum channel [62]. Therefore, unlike the SQL, the HL is specific to the QEM scheme [63]. In the case of interferometric phase estimation, the HL is $\wp = 1$ [64], although this limit can only be attained through the use of optimal measurement. As an optimal POVM could be infeasible, QEAPE schemes

provides an attractive alternative to achieving close to this lower bound [65].

B. Adaptive phase estimation

Phase estimation underlies many QEM applications [1,8,51] and thus is widely used for devising quantum-enhanced techniques, including several AQEM schemes [18,21,65]. Here we explain QEAPE controlled by Bayesian feedback [16,26] or DE-designed policies [22,23,66], which we compare in terms of robustness and resource consumption for control.

1. QEAPE

One method of estimating phase is to use an interferometer, which infers phase shifts from the interference between two or more modes [67]. In particular, we use QEAPE based on a Mach-Zehnder interferometer, which has two modes and therefore we are looking at the case of $d = 2$ representing the modes. The mathematics of Mach-Zehnder interferometry applies to other forms of SU(2) interferometry, such as Ramsey, Sagnac, and Michelson interferometry [1,51,68]. Here we present the input state, adaptive channel, detection, feedback, inference, and imprecision.

a. Input state. For nonadaptive quantum interferometry with collective measurement, the unitary interferometric transformation is in the Lie group SU(2) with irrep ("irreducible representation," which is the Casimir-invariant label) $j = N/2$. For adaptive quantum interferometry or individual measurements, the interferometric unitary transformation is $\mathrm{SU}(2^N)$ for N particles and two paths. However, the two descriptions converge if the input state is permutationally symmetric; technically, Schur-Weyl duality dictates that the applicable transformation is SU(2) with irrep $j = N/2$ [69].

Notationally, modes are labeled by

$$\epsilon_m \in \{0, 1\}, \tag{20}$$

which conveys which of the two paths, such as input or output port or intrainterferometric path, pertains. Thus, the state $|\epsilon_m\rangle$ refers to mth photon being in path ϵ_m. The multiphoton basis is the tensor-product state

$$|\boldsymbol{\epsilon}_N\rangle = \bigotimes_{m=0}^{N-1} |\epsilon_m\rangle\,. \tag{21}$$

For ham $\boldsymbol{\epsilon}$ the Hamming weight, i.e., sum of bits, of $\boldsymbol{\epsilon}$, the permutationally symmetric basis is

$$|n, N_a - n\rangle = \binom{N_a}{n}^{-1/2} \sum_{\mathrm{ham}\,\boldsymbol{\epsilon}_{N_a}} \left|\boldsymbol{\epsilon}_{N_a}\right\rangle \tag{22}$$

for N_a the total number of particles in mode a.

The sine state serves as a symmetric state that minimizes phase-estimation imprecision [19,20,70], and is expressed as [16,26]

$$|\psi\rangle_N = \left(\frac{N}{2}+1\right)^{-1/2} \sum_{n,k=0}^{N} \sin\left(\frac{k+1}{N+2}\pi\right) e^{i\pi(k-n)/2} \times d^{N/2}_{n-N/2,\,k-N/2}\left(\frac{\pi}{2}\right) |n,\ N-n\rangle\,, \tag{23}$$

for $d^j_{m,m'}(\beta)$ the Wigner d function [71].

This state provides several advantages both from the metrological and computational perspective. First, the state can be viewed as an extension of the equal-number-photon state, which has been shown to deliver quantum-enhanced precision close to the Heisenberg scaling [51,72,73]. Second, this state also has the advantage of being robust against photon loss [69], which is a desirable property for practical QEM and so is used in the QEAPE procedures. This property is due to the state's ability to remain approximately invariant when few photons are lost. Third, the permutational symmetric structure of the state allows us to lower the computational cost of simulating the state from a matrix of size $2^N \times 2^N$ to $N + 1$.

b. Adaptive channel. The particles in the sine state (23) are divided into single-particle bundles ($L = 1$ case), each of which passes through the Mach-Zehnder interferometer. For noiseless interferometry, the quantum channel for one photon is

$$U_1(\phi; \Phi_m) = \exp[i(\phi - \Phi_m)\hat{\sigma}_y], \quad \phi, \Phi_m \in [0, 2\pi) \tag{24}$$

for $\hat{\sigma}_y$ a Pauli matrix [73]. Therefore, the channel is

$$U(\phi; \Phi_m) = U_1(\phi; \Phi_m) \otimes \cdots \otimes \mathbb{1}^{(N)} \tag{25}$$

acting on the state space (12).

In a physical implementation of an interferometer, mechanical disturbances, air-pressure changes, and the thermal fluctuations induce optical-path fluctuations [30]. These effects randomize the phase difference

$$\phi - \Phi_m, \tag{26}$$

which we assume to be contributed solely by ϕ according to the prior distribution $p(\phi)$, as ϕ is determined by a mechanism that is not completely specified, and the controllable phase shift Φ_m is always exact. The quantum channel is thus [66]

$$\mathcal{I}(\phi; \Phi_m) : \mathcal{S}\big(\mathscr{H}_2^{\otimes(N-m)}\big) \to \mathcal{S}\big(\mathscr{H}_2^{\otimes(N-m)}\big) : \rho_m \mapsto \int_{\phi=0}^{2\pi} d\phi\, p(\phi) U^\dagger(\phi; \Phi_m) \rho_m U(\phi; \Phi_m), \tag{27}$$

where ρ_m is the state after the $(m-1)$th photon is measured.

c. Detection, feedback, and inference. After the mth photon passes through the interferometer, the photon is detected by one of the single-photon detectors positioned outside the output ports. The information about the exit port is $x_m \in \{0, 1\}$, which is given to a controller. The controller then uses this information to compute Φ_{m+1} from ϱ before the next photon arrives. The procedure is repeated for subsequent photons until all photons are consumed, and the estimate is inferred from $\tilde{\phi} = \Phi_N$, assuming no loss of photons.

d. Imprecision. Imprecision of the estimate (1) is related to the Holevo variance [19],

$$(\Delta\tilde{\phi})^2 = V_{\mathrm{H}} := S^{-2} - 1, \tag{28}$$

by the sharpness function

$$S = \frac{1}{K}\left|\sum_{k=1}^{K} \exp\left[i(\phi_0^{(k)} - \tilde{\phi}^{(k)})\right]\right|, \tag{29}$$

which quantifies the width of a distribution over a periodic variable. The sharpness (29), hence the Holevo variance (28), is estimated by repeatedly simulating QEAPE $K = 10N^2$

times [21], each with an unknown (noiseless) interferometric phase shift $\phi_0 \in [0, 2\pi)$ randomly chosen from a uniform prior.

In the simulation, ϕ_0 is the mode (most frequent value) of the unimodal prior distribution $p(\phi)$. The sharpness S (29) is a random variable, which is effectively sampled from a sharpness distribution $P_{\text{sharp}}(S)$ with mean

$$\bar{S} = \sum SP_{\text{sharp}}(S). \tag{30}$$

If the interferometer is free of phase noise, the sharpness distribution $P_{\text{sharp}}(S)$ is sufficiently narrow that S computed from $K = 10N^2$ samples can be used to estimate the Holevo variance. Once the phase noise is included, $P_{\text{sharp}}(S)$ is not only shifted but widened and consequently $\bar{S}$ is used to estimate imprecision of the QEAPE procedure.

2. Policy design

Here explain how the policy ϱ (15) is obtained. Specifically, we consider two approaches. First we consider the Bayesian approach, based on Bayes theorem [74], which is applied if the controller has complete knowledge about the quantum system and updates this knowledge based on measurement outcomes [75]. The second approach is to execute a direct policy search for an optimal or feasible ϱ [76], which has the advantage over the Bayesian approach in that a trusted model is not required.

a. Bayesian feedback. Early examples of AQEM employ the Bayesian approach [16,26,75], which involves a fully trusted and complete model of the quantum dynamics. The model is used for computationally intensive decision making as measurement data arrive, leading to updates of the prior. For each particle exiting the interferometer, the probability for the particle being detected in output port given by outcome x_m is computed by assuming a perfectly known input state and known quantum dynamics such as unitary interferometric evolution (25). The prior for ϕ_0 is then updated using Bayes's theorem. The width of this prior is quantified by $\sqrt{V_{\text{H}}}$ (28) at each measurement step; the optimal Φ_m that minimizes V_{H} for the next particle is then computed from this trusted model. Although the Bayesian method approaches the HL [26], ϱ might not be robust to input-state variability or to system noise, which is a problem when designing controllers [77,78].

b. Policy search using differential evolution. Policy search provides an attractive alternative to the Bayesian approach when a trusted model is lacking or if the computational overhead associated with the Bayesian approach is excessive. In this approach, the policy is given a task-appropriate structure and the performance is optimized by searching the space of the policy's parameters [76,79]. A policy-search method is called model free if the optimization process does not include learning a model of the dynamic but directly updates the policy parameters [76], which can be achieved by employing a black-box optimization algorithm [80] using only policy performance. Evolutionary algorithms form a class of black-box optimization algorithms whose policy update procedure is inspired by models of biological evolution [27], and DE (discussed in Appendix A) is one such algorithm that is capable of finding a feasible policy for QEAPE using many particles [22].

For the task of QEAPE, we assume a logarithmic-search heuristic Markovian update rule [21]

$$\Phi_{m-1} \mapsto \Phi_{m-1} - (-1)^{x_m}\Delta_m, \tag{31}$$

with phase-adjustment vector

$$\mathbf{\Delta} := (\Delta_1, \Delta_2, \ldots, \Delta_N), \tag{32}$$

optimized during the training stage. This update rule (31) corresponds to turning the "phase knob" up or down by a fixed amount Δ_m after the mth photon, subject only to the previous outcome and ignoring the full measurement history.

The performance of the policy is evaluated using the average sharpness (30) [66], as the interferometer is assumed to contain noise and this performance measure averages out noise. In practice, due to the high cost of sampling $P_{\text{sharp}}(S)$ (30), few runs are performed to obtain $\bar{S}$. A DE algorithm then uses this performance to search iteratively for $\mathbf{\Delta}$ such that $\bar{S}$ is maximized [23].

The use of DE incurs a time cost for generating ϱ, which is a one-time cost quantified by a loop analysis of the DE algorithm. Once a feasible policy is generated, the policy can be used in the QEAPE procedure as long as the dynamics of the system does not significantly change. The scaling of time cost with respect to N determines the complexity, which has been shown to be poly$N$ [22] so the degree of the polynomial conveys the complexity for generating ϱ.

C. Noisy quantum-enhanced metrology

Quantum-enhanced precision is usually diminished by experimental imperfections whether inherent in the input-state generation, loss and decoherence in the quantum channel, limitations of the measurement devices, or perhaps any combination of all these imperfections. [25,81]. Determining the imprecision lower bounds for different QEM strategies [82,83] in the presence of imperfection is therefore important for devising practical QEM procedures. In this subsection, we focus our discussion on summarizing previous work that model imperfection as noise in the quantum channel (as that is the main focus of work in this field), assuming that the optimal input state for these channels can be generated. Furthermore we discuss QEM schemes that have been shown to deliver SQL scaling and discuss schemes capable of breaking this limit. We do not mean to give a comprehensive review of the field but rather to discuss factors that could lead to robust QEM procedures.

One important aspect of these studies concerns so-called "no-go" results, for which noise in the quantum channel leads to an asymptotic phase-estimate precision corresponding to the SQL, has been observed by numerical investigation for both collective [25] and adaptive QEM schemes [84,85]. These studies suggest that the quantum enhancement would disappear in the limit of many photons in interferometric measurement. Theoretical studies show that noisy QEM channels described by uncorrelated semigroups lead to SQL scaling asymptotically as well [86,87]. On the other hand, noise models that are not described by uncorrelated semigroups, such as non-Markovian noise [88], is proven to yield a quantum-enhanced phase-estimate precision that exceeds the SQL in the asymptotic large-photon-number limit.

One strategy to surpass the SQL is based on applying time-dependent control, which is not time-homogeneous and hence not compliant with the requirements for the no-go result. One example of controlled QEM uses an ancillary system comprising many particles that are all entangled to the input state. These cases exploit quantum error correction [85,89,90], quantum teleportation [91], and dynamic decoupling [92]. In these systems, measurement can be performed on ancillae with the resultant information used to correct the input state system with respect to any error that might have occurred due to quantum-channel noise. Similarly, adaptive-measurement protocols have been applied to noisy QEM procedures without ancillary systems that show scaling exceeding SQL [10,66]. These quantum-control schemes suggest that adaptive control is a promising method for attaining quantum-enhanced precision, i.e., surpass the SQL, provided that appropriate feedback strategy can be devised.

D. Models of phase noise

In this subsection, we explain the choices of the phase-noise model for the robustness test. This noise is simulated by turning ϕ into a random variable that has a unimodal probability distribution with the peak at ϕ_0. The mode ϕ_0 is assumed to be the unknown parameter to be estimated. For the test, ϕ follows one of these four distributions: normal, three-stage random telegraph, skew-normal, or log-normal distribution. We summarize the relationship between the noise parameters and the variance and skewness [93] as we use both in selecting the parameters for the robustness test and the variance in particular to quantify the noise level.

1. Normal-distribution noise

Normal-distribution noise is important for testing robustness of the search algorithm because the normal distribution is especially prominent due to the central-limit theorem, which states that the arithmetic average of a random variable assumes a normal distribution as the number v of data points increases asymptotically, $v \to \infty$ [31]. Due to the ubiquity of the normal distribution in nature [94], we assume that typical mechanical vibrations, acoustic, and thermal fluctuations contribute to a normal-distribution noise model.

The normal distribution,

$$p(\phi) = \frac{e^{(\phi-\mu)^2/\sigma^2}}{\sqrt{2\pi}\sigma}, \tag{33}$$

is parametrized by the mean μ and standard deviation σ. As this distribution is symmetric, skewness γ is zero and thus the mode is at μ, and the variance is $V = \sigma^2$. In our simulations, we set $\mu \equiv \phi_0$ so the only free parameter is σ, which is bounded above by $\sigma < \pi$ as otherwise the width would exceed the domain of ϕ. We restrict $\sigma \leqslant 2$ as higher values of σ would be uninteresting as they would correspond to almost completely noisy measurements.

2. Random-telegraph noise

Random-telegraph noise [33] is a discrete distribution that, for each time step, randomly switches between two values, one being the correct and the other an erroneous value. Whereas this noise is most relevant to digital electronics, as it simulates a bit-flip error, this noise simulates other forms of digitized noise, such as salt-and-pepper noise in image processing [95].

In an adaptive interferometer, random-telegraph noise can arise from the control signal or from the stepping motor used to control the phase shifter. Although, in principle, the control signal is smooth, computerized control approximates the ideal signal using small stepwise changes, thereby introducing the possibility that the controllable phase shifter Φ_m is close to but not exactly where the ideal Φ_m should be. Similar effects can also be a result of the stepping-motor error, whose change is limited to discretized steps [96]. Here we exaggerate the size of the error in order to study effects of random-telegraph noise on a QEM scheme.

We modify two-stage random-telegraph noise to have three stages,

$$p(\phi) = \begin{cases} 1 - p_s, & \phi = \phi_0, \\ \frac{p_s}{2}, & \phi = \phi_0 \pm \delta. \end{cases} \tag{34}$$

The probability of switching to an erroneous value is p_s, and δ is the distance between the true and erroneous values leading to

$$V = p_s\delta^2, \quad \gamma \equiv 0 \tag{35}$$

with the last relation following from the symmetry of the distribution.

Unimodality of the distribution implies that $p_s < 2/3$. Furthermore, we restrict $\delta < \pi$ so that the distance between the two side peaks is less than 2π. To comply with both constraints and be able to raise the noise level to at least $V = 3$ so the result is comparable to other distributions, we fix $p_s = 1/2$ for the test and then vary only δ.

3. Skew-normal-distribution noise

The skew-normal distribution [34] is modified from a normal distribution by multiplying with a function whose skewness parameter is α. Typically, this noise model is used in simulations of filters and detectors [36,38] but not in interferometry where noise is approximated to be normally distributed. However, as the skew-normal distribution automatically includes the normal distribution as a limiting case (skewness goes to zero), this noise model can be viewed as a convenient generalization of the normal distribution when some skewness is suspected or should be treated.

Skewness of the distribution is described by the distribution

$$p(\phi) = \frac{e^{-(\phi-\mu)^2/2\sigma^2}}{\sqrt{2\pi}\sigma}\left[1 + \mathrm{erf}\left(\frac{\alpha}{\sqrt{2}\sigma}(\phi-\mu)\right)\right] \tag{36}$$

for erf the error function [97]. Skewness of the distribution is

$$\gamma = \frac{4-\pi}{2}\frac{2\beta}{\pi - 2\beta}, \quad \beta = \frac{\alpha^2}{1+\alpha^2}, \tag{37}$$

and the variance is

$$V = \sigma^2\left(1 - \frac{2\beta}{\pi}\right). \tag{38}$$

The mode, however, does not have a closed form although it remains close to μ as α/σ increases. For our simulations of

phase noise, we treat the mode as being at μ, so μ is always set to ϕ_0. The parameters α and σ are set to the desired skewness and variance for testing robustness.

4. Log-normal-distribution noise

Log-normal [35] noise has a heavy-tailed skewed distribution that provides another approach to generalizing the normal distribution (without including the normal distribution itself in the class). The log-normal distribution appears prominently in the study of biological [98] and electrical [99] networks, of electrical noises [37], and of many other fields of sciences [35] including astronomy [100]. The prevalence of the distribution in finite, positive data is often attributed to the "multiplicative central limit theorem," which states that the distribution of the geometric mean tends to a log-normal distribution for data size $v \to \infty$ [35,100]. As such, random data generated through nonlinear random processes tends to have a log-normal distribution.

Although no obvious multiplicative mechanism exists in an optical interferometer [30], physical phenomena that create ϕ_0 may play a role in adding noise and that noise, depending on the physics of the phenomena have a likely chance of being log-normally distributed if nonlinear multiplicative random processes are involved. As such, we include this noise distribution as one of the models we use for the robustness test.

In this case, the logarithm of the random variable has a normal distribution, leading to the distribution

$$p(\phi) = \frac{e^{-(\ln\phi-\mu')^2/2\sigma'^2}}{\sqrt{2\pi}\sigma'\phi} \tag{39}$$

with mode and variance

$$\phi_0 = e^{\mu'-\sigma'^2}, \quad V = (e^{\sigma'^2}-1)e^{2\mu'+\sigma'^2}, \tag{40}$$

respectively, and skewness

$$\gamma = (e^{\sigma'^2}+2)\sqrt{e^{\sigma'^2}-1}.$$

As this distribution is defined for $\phi \in (0, \infty)$, we first generate a random number within the compact phase domain given μ' and σ' and then apply the shift

$$\phi \mapsto \phi + \phi_0 - e^{\mu'-\sigma'^2} \tag{41}$$

so that the mode of the distribution is centered at ϕ_0 (40).

III. APPROACH

In this section, we devise a test to determine whether quantum-enhanced precision is feasible in the presence of unknown noise. We then assess whether power-law scaling of phase imprecision vs particle number N is valid asymptotically and establish a method to determine this power $\wp$. Finally, we define the resource for generating and executing the control policies in terms of how space and time costs scale with N. This scaling analysis is based, as we explain in the Introduction, on numerical analysis up to some maximum photon number $N_{\max}$, determining the trend of the log-log plot for Holevo variance V_{H} vs N based on statistical methods, and extrapolating to larger N to make scaling claims.

A. Robustness test

The robustness of QEAPE policies is determined by testing the policies in the presence of noise whose model is not recognized by the policies and the method that designs the policies, although DE-designed policies are learned in training that includes the noise. Here we define the test for QEAPE, including phase noise from Sec. II D. We specify the domain of N for simulating the phase estimation schemes to obtain V_{H} in noisy conditions. The noise parameters are variance V and skewness γ (Sec. II D), but here we fix γ for the asymptotic distributions, and we obtain the robustness-test threshold in terms of V, which is the maximum for each noise model such that the SQL is violated.

1. Varying N

To ascertain the asymptotic value for $\wp$, we simulate QEAPE for

$$N \in \{4, 5, \ldots, 100\}, \tag{42}$$

as V_{H} computed from this domain is sufficient to show a power-law relationship at high N (approaching $N_{\max}$). Furthermore, increasing N further requires changing double-precision arithmetic to quadruple-precision arithmetic to generate and manipulate the sine state without rounding error. Consequently, this increase in precision leads to a 15-fold increase in run time at $N = 100$, which is a large expense for generating a single data point. Therefore, we do not attempt to verify the robustness beyond this 100 particles.

2. Skewness

We fix skewness γ to a single value for all runs and only vary V because V is the dominant term in our noise models and γ has a small effect [101]. We fix the skewness for the asymmetric distribution to

$$\gamma = 0.8509, \tag{43}$$

which is sufficiently large to distinguish between the various noise models; otherwise all noise looks Gaussian. This value of γ (43) corresponds to $\alpha = 5$ for the skew-normal distribution where we are able to observe its effect on $\wp$ when compared to symmetric noise distributions. This same level of skewness corresponds to

$$\sigma' = 0.2715 \tag{44}$$

in the log-normal distribution.

3. Robustness threshold

Our ϱ is robust if the SQL-breaking condition $\wp > 1/2$ is satisfied for all four noise models in Sec. II D. As discussed in Sec. III A 2, we fix γ, and we ignore higher cumulants; thus, the ϱ robustness threshold is in terms of V, i.e., the maximum V such that $\wp > 1/2$ holds for all four noise models. This optimization problem is hard so we adopt a simpler characterization procedure instead to get insight into the robustness threshold. Our approach is to run the simulations for $V \in \{1, 2, 3\}$ for symmetric noise and $V \in \{1, 3, 5, 7\}$ for asymmetric noise, and we do not push beyond $V = 7$ to keep

below an imprecision width of 2π. We use these data to determine whether QEAPE policies pass the robustness test.

B. Determining asymptotic power-law scaling

To ascertain the robustness of QEAPE policies, the asymptotic $\wp$ is estimated from a subset of V_{H} at sufficiently high N, and determining this subset is done by fitting piecewise linear equations to a log-log plot of V_{H} vs N. In this subsection, we introduce five piecewise functions that are constructed from (numerical) observations regarding the trend of V_{H} vs N. We then explain the method of finding the break points between segments in the piecewise function and fitting the functions to the data. Using the criteria in Appendix B, we create a majority-vote method for selecting the function that best represents the data and thus $\wp$ from the last segment of the fit is used to estimate the asymptotic scaling.

1. Piecewise models

The power-law trend between V_{H} and N, which manifests as a linear relation between $\ln V_{\mathrm{H}}$ vs $\ln N$, differs under noisy conditions. Here we describe the trends we have observed that lead to piecewise linear functions. We construct five such functions, containing one to three segments that are then fitted to V_{H} vs N.

When the interferometer is noiseless, the relationship appears to be a power law captured in a linear equation, although the accept-reject criterion in the policy-search algorithm can lead to a different $\wp$ for $N > 93$. Once the noise level becomes high, typically $V > 1$, the relationship does not appear to be linear for low N. Therefore, we include segmented models that fit linear interpolation to the data in the first segment.

Combining these observations, we construct five piecewise-linear models that can potentially represent $\ln V_{\mathrm{H}}$ as a piecewise function of $\ln N$. The models have one to three segments, each segment connected at the break points determined by the fitting methods. Three of the models are one, two, and three linear models, whereas the other two are a two-segment model, where the first segment (low N) is a linear interpolation, and a three-segment model where the first segment is a linear interpolation and the second and third segments are linear.

2. Fitting method

Our method for fitting linear equations is based on the least-squares method [102]. However, because the functions in Sec. III B 1 are segmented, we include a step to optimize the break points depending on the specific function.

The full model for the regression analysis is the three-segment linear function, which is fitted using the linear-square method and the segments determined by a heuristic global optimization algorithm [103]. The two-segment linear function is also fitted using the same least-square method although a brute-force search is used to find the break point starting from $N = 4$.

The method for finding the break points for models with interpolation are different as the linear interpolation leads to a small residual. Thus, optimizing using the least-squares method can lead to a single segment of linear interpolation. For this reason, we first find the stop point for the first segment by fixing the latter segments to a single linear line and search for break point that results in a large decrease in sum square error. As for the single-segment linear model, we use a standard library to fit to the data.

3. Fitting figures of merit and model selection

After the functions are fitted to $V_{\mathrm{H}}(N)$, the criteria $\overline{R^2}$, $\mathrm{AIC_c}$, the F value, and Mallows's C_p (Appendix B 2) are calculated for each of the function and the fits are visually inspected. These criteria are used to select the function that best fits the data. Here we explain how the best function is chosen.

After the functions are fitted and the criteria are calculated, each fit is visually presented and inspected to ascertain that the segmentation fits the pattern. If corrections are unnecessary, the functions are then ranked for each of the criteria. Note that we do not perform the full F test as we discovered that reduced models typically fail the test even though there is no discernible difference when compared to the full model. However, the F value can still be used to quantify the difference between using the full and reduced model, so we use the F value to rank the functions instead of conducting a pass-fail test.

After the functions are ranked, the function that is voted as best by most of the criteria is chosen to represent the data. In the case where the full model, the three-segment linear function, is voted according to the criteria, the value of $\wp$ from the last segment and the subset of N where this value is computed is compared to the next alternative function to determine whether the function overfits the data. If $\wp$ from the two functions differs more than 0.001, then the full model is chosen; otherwise, the alternative function is chosen. The limit we use here is specified based on the precision used in this paper and can be changed based on the desired precision of $\wp$.

C. Resource complexity

To compare and select between policies and methods of generating policies, we determine the complexities for designing and for executing ϱ using the loop-analysis method [39] to determine both of these complexities. We begin this subsection by explaining the time complexity for designing policies. We then explain execution complexity by describing the controller's actions and concomitant resources, which are quantified by the space and time required for executing ϱ with N particles.

1. Design complexity

When an optimization algorithm is used to design ϱ, there is a time cost associated with the use of the algorithm. The scaling of the upper bound of this cost is called design complexity. We assume that the calculation is performed on a simulation of the AQEM task, as is common practice in policy design in quantum control [104], and, therefore, the time cost includes the cost of simulating the AQEM task. We assume that only a single processor is used for the purpose of comparing policy-design methods, although this cost can be reduced by parallelizing the optimization on multiple CPUs.

The design complexity for the policies in Sec. II B 2 is shown to be $O(N^6)$ through loop analysis [22], and this complexity does not change when noise is included. For policies that are devised through analytical optimization, such as Bayesian feedback, the cost is zero as no design algorithm is used.

2. Execution complexities

In this subsection, we explain the requisite resource complexity to implement AQEM policies, quantified by the scaling of the space and time costs with the number of particles N [5]. We begin by explaining the connection between an AQEM policy and an algorithm by viewing the controller as a computer, allowing us to use the method of algorithm analysis to calculate the complexities [39]. We then define the space and time cost for implementing policies and how these costs are calculated.

a. Controller. The controller holds ϱ (15) and implements ϱ to execute decisions based on feedback from detection of outgoing particles. The controller is essentially an agent who receives input from detectors and transmits a control signal to an actuator that shifts the interferometric phase. Thus, ϱ is represented as a computer algorithm expressed as a computer program. Computer memory is required to store ϱ and time to execute ϱ. The candidate ϱ can be designed by various means including Bayesian feedback (Sec. II B 2 a) and policy-search method (Sec. II B 2 b). Space and time costs are discussed in the next two paragraphs.

b. Space complexity. We determine the upper bound for space cost, which is the worst-case amount of memory used by an algorithm reported as a big-O function of the size of the problem [5]. As ϱ is executed by an algorithm, this worst-case, or maximum-size, memory corresponds to how much space is required to hold the critical information required to execute the feedback. For ϱ, the computer's space cost for memory depends on the type of ϱ. For Bayesian feedback, the size of the stored ϱ is $O(N^2)$, which is calculated from the size of the array used to store the quantum state transformation, and is shown in Sec. IV C and specifically in Table II. The size of the stored ϱ from the policy-search method is $O(N)$ [22]. This linear scaling of the policy size, with respect to the number N particles, is due to the generalized-logarithmic-search heuristic leading to the size N phase-adjustment vector (32).

c. Time complexity. Time complexity refers to the scaling of the upper bound for the time cost required to implement a single AQEM shot, i.e., the cost for using all N particles at once. This cost is calculated by assuming the time a particle takes to pass through an interferometer is constant, mimicking the physical implementation of the control procedure. We then use loop analysis, which counts the number of loops that perform operations, and we assume all take the same constant time [39]. For a shot of AQEM task using N particles, Φ_m is computed N times, corresponding to each particle passing through the interferometer and being detected. For each particle, nested loops for computing Φ_{m+1} exist according to ϱ. The complexity is reported as the scaling of this time cost with N.

We determine the execution cost for DE-designed policies by recognizing that the update of Φ_m according to Eq. (31) is constant in time. Therefore, QEAPE consists of only one loop over the number of particles, and the execution complexity is thus $O(N)$. Bayesian feedback, on the other hand, has nested loops for updating the quantum state, which is $O(N^2)$ in time complexity for every computation of Φ_m. Hence, the execution complexity is $O(N^3)$.

IV. RESULTS

In this section, we report results for the robustness test and the resources based on sampling from simulated QEAPE. We present and compare $V_{\mathrm{H}}(N)$ (28) obtained by a policy-search method discussed in Sec. II B 2 b and by Bayesian feedback discussed in Sec. II B 2 a. Our analysis considers all four types of noise discussed in Sec. II D. We report values of $\wp$ (1) from the regression procedure discussed in Appendix III B and resource complexities discussed in Sec. III C for both the DE-designed policies and the Bayesian feedback.

A. Variance vs number of particles

In this subsection, we present results for V_{H} as a function of number N of particles. Specifically, we present plots of V_{H} vs N from 4 to 100 particles, which is enough to determine scaling as discussed in Sec. III C. Both cases of using DE-designed policies and of using Bayesian feedback are presented as log-log plots (base 10) and compared to the SQL, with these plots obtained by computing from simulations of noiseless phase estimation using a product state $|0,1\rangle^{\otimes N}$ following the notation of Eq. (22). The HL is generated based on the intercept of the SQL data using the scaling of $1/N^2$ to provide a benchmark.

1. Policy search

Here we present the log-log plots of V_{H} vs N (base 10) from QEAPE using DE-designed policies, as shown in Fig. 1. Figs. 1(a)–1(d) present V_{H} for inclusion of normal-distribution noise, random-telegraph noise, skew-normal noise, log-normal noise, respectively.

Figure 1(a) also includes V_{H} from noiseless interferometry. This locus appears as a straight line in the plot, indicating a power-law relationship between V_{H} and N. As the noise variance V increases, this power-law relationship breaks into two parts, clearly visible in the V_{H} vs N plot from $V = 3$. This trend also appears at $V = 2$ as the model selection procedure Sec. III B 3 selects the two-segment model for this data set. The observation that the power-law relationship fails when noise is included is also evident in Figs. 1(b)–1(d). In these cases, the plots are fit to two- or three-segment linear equations as V_{H} appears to have a bump at low N as V increases.

The increase in phase noise V also results in an increase in the intercepts of V_{H} power-law lines; however, the rate of change appears to depend on the noise model. The difference is shown in Figs. 1(a) and 1(b); both include symmetric noise distributions but with different spacing of the intercepts. The same observation holds for Figs. 1(c) and 1(d), which are from asymmetric distributions. Comparing the four plots shows that the intercepts appear to increase more slowly for asymmetric

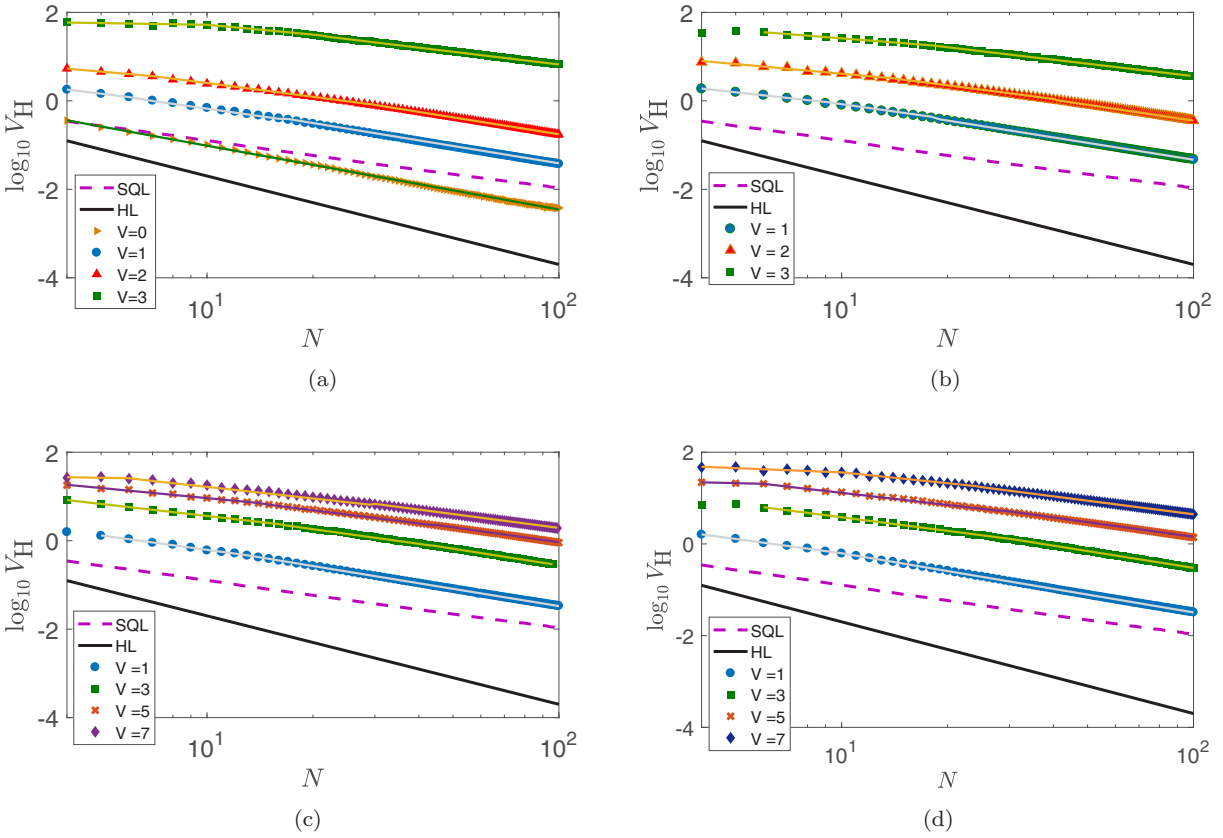

FIG. 1. Logarithmic plots (base 10) of Holevo variance from simulation of QEAPE. The policies are designed using policy-search method implemented in the specified noise condition, namely, (a) normal-distribution noise, (b) random-telegraph noise, (c) skew-normal-distribution noise, and (d) log-normal-distribution noise. The plot for the normal-distribution noise also includes the data from the noiseless simulation (brown side-facing triangle) and its linear fit (green solid). The blue circles are data when $V = 1$, the red triangles when $V = 2$, the green squares when $V = 3$, the brown plus when $V = 4$, the brown crosses when $V = 5$, and the purple diamonds when $V = 7$. The lines shown are the piecewise linear fits of the data whose scaling is reported. The solid black line is the HL and the dashed purple line in the SQL is generated from noiseless QEAPE.

distributions than for the symmetric distributions, being close to 1 for $V = 3$ in the former and $V = 2$ for the latter.

2. Bayesian feedback

Log-log plots of V_H (base 10) as a function of N, shown in Fig. 2, are computed from simulations of QEAPE controlled by Bayesian feedback. Figures 2(a)–2(d) present V_H in the presence of normal-distribution noise, random-telegraph noise, skew-normal noise, and log-normal noise respectively.

Similar to Fig. 1, the trend of V_H vs N in Fig. 2 shows that the power-law relationship also breaks into parts. Instead of a bump, V_H from Bayesian feedback exhibits noise for low N. For this reason, the model-selection procedure Sec. III B 3 favors the model with linear interpolation in the first segment. The subsequent segment appears straight in the log-log plots, although some, such as $V = 7$ in Fig. 2(c), shows a break into two linear segments.

The intercepts of the V_H vs N plots increase with the increase of V, and the observation of the changes are similar to when DE-designed policies are used (Sec. IV A 1). The asymmetric noise shows a slow increase in intercept when compared to symmetric noise, and the rate of change depends on the noise model.

B. Power-law scaling

In this subsection, we present values of $\wp$, summarized in Table I, that are estimated by fitting V_H plots in Sec. IV A. These $\wp$'s are from the last segment of the selected piecewise linear models (Sec. III B 1), which changes with the increase in V. We also include $\overline{R^2}$ (B2) to show the goodness of fit.

The power-law scaling for DE-designed policies, namely, $\wp_S$, shows a decrease as the noise level V increases, starting from the noiseless phase estimation at $2\wp_S = 1.459$. The DE-designed policies fail to deliver $\wp_S > 1/2$ when $V = 3$ for the symmetric noise distributions. This limit increases with asymmetric noise models to $V = 7$ in log-normal noise. Skew-normal noise only shows a scaling at approaches the SQL but does not breach it at all.

Similar trends are observed for Bayesian feedback. The scaling $\wp_B$ from noiseless interferometer closely approximates the HL at $2\wp_B = 1.957$ and approaches SQL when $V = 4$ for normal-distribution noise. This limit drops to $V = 3$ when random-telegraph noise is included. This limit also appears at $V = 7$ for log-normal noise, whereas the same noise level just leads to $\wp_B$ approaching SQL when skew-normal noise is present. This trend, aside from the case of normal-distributed noise, is the same as the trend for the DE-designed policies.

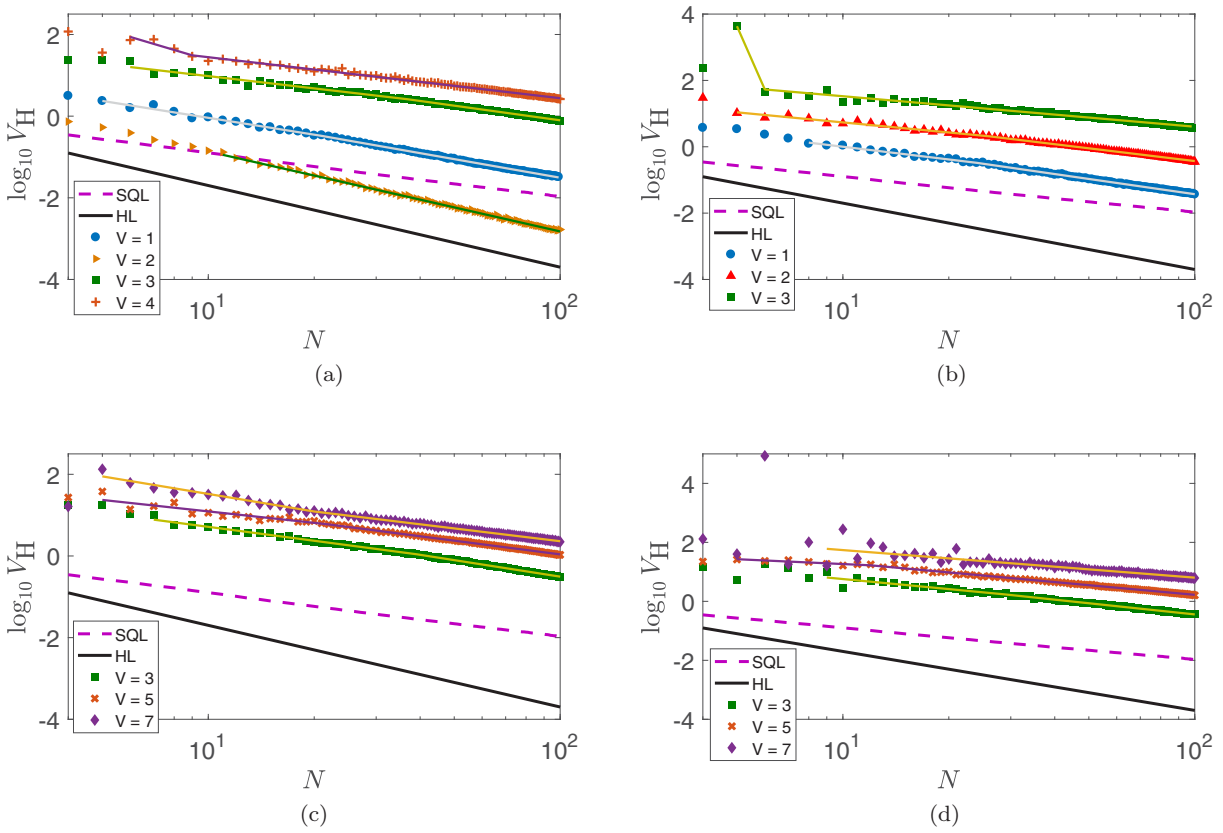

FIG. 2. Logarithmic plots (base 10) of Holevo variance from simulation of QEAPE using Bayesian feedback method. The simulation includes one of the four noise models, namely, (a) normal-distribution noise, (b) random-telegraph noise (c), skew-normal-distribution noise, and (d) log-normal-distribution noise. The plot for the normal-distribution noise also includes the data from the noiseless simulation (brown side-facing triangle) and its linear fit (green solid). The blue circles are data when $V = 1$, the red triangles when $V = 2$, the green squares when $V = 3$, the brown plus when $V = 4$, the brown crosses when $V = 5$, and the purple diamonds when $V = 7$. The lines shown are the piecewise linear fit of the data whose scaling is reported. The solid black line is the HL and the dashed purple line in the SQL generated from noiseless QEAPE.

TABLE I. Power-law scaling from QEAPE under noisy conditions using DE-designed policies $\wp_S$ and Bayesian feedback $\wp_B$.

	V	γ	$2\wp_S$	$\overline{R^2}_S$	$2\wp_B$	$\overline{R^2}_B$
SQL			1		1	
HL			2		2	
No noise			1.459	0.9998	1.957	0.9993
	1	0	1.302	0.9999	1.512	0.9985
Normal	2	0	1.267	0.9999		
	3	0	0.954	0.9992	1.190	0.9997
	4	0			1.004	0.9948
	1	0	1.266	0.9999	1.526	0.9991
Random telegraph	2	0	1.186	0.9997	1.277	0.9967
	3	0	0.935	0.9993	0.919	0.9892
	1	0.8509	1.296	0.9999		
Skew-normal	3	0.8509	1.246	0.9999	1.343	0.9987
	5	0.8509	1.118	0.9998	1.116	0.9927
	7	0.8509	1.039	0.9996	1.041	0.9964
	1	0.8509	1.290	0.9999		
Log-normal	3	0.8509	1.217	0.9998	1.258	0.9919
	5	0.8509	1.058	0.9997	1.086	0.9961
	7	0.8509	0.981	0.9994	0.9209	0.7965

The goodness-of-fit for these fits are reported in term of $\overline{R^2}$, where $\overline{R^2} = 1$ indicates a perfect fit. The values of the goodness $\overline{R^2}_S > 0.999$ for V_H delivered by DE-designed policies and $\overline{R^2}_B > 0.99$ for Bayesian feedback except for when a log-normal noise of $V = 7$ is present. Overall, the models chosen using the method in Sec. III B 3 provide good fits to the data, and the DE-designed policies always deliver fits with $\overline{R^2}_S > \overline{R^2}_B$.

C. Bounds on time and space costs

The result from calculating space and time complexities for both designing and implementing DE-designed policies and Bayesian feedback is shown in Table II. Here we compare these results.

TABLE II. Upper bound in policy space and time cost of the policy from DE algorithm (DE) and Bayesian feedback (BF).

Complexity	DE	BF
Design time	$O(N^6)$	
Policy space	$O(N)$	$O(N^2)$
Execution time	$O(N)$	$O(N^3)$

Time complexity for generating policies, named here design time, is poly$N$ but of a high polynomial degree when the DE algorithm is used. Bayesian feedback, which is designed through an analytical process, incurs no time cost for the design. When the execution time is compared, the time complexity goes from $O(N)$ for the DE-designed policy to $O(N^3)$ for Bayesian feedback as shown in Table II. Space complexity for the DE-designed policy is also linear with respect to the number of particles whereas Bayesian feedback requires a memory of $O(N^2)$, also shown in Table II.

V. DISCUSSION

In this section, we discuss policy robustness and why QEAPE policies are so strongly noise resistant. We explain the difference between the scaling parameter $\wp$ attained by DE-designed policies and Bayesian feedback. Our analysis enables a user to choose policy-design methods that minimize resource complexity for design and for execution as appropriate.

A. Robustness of QEAPE policies

In this subsection, we discuss robustness of QEAPE policies and the robustness threshold based on $\wp$ in Table I. Both DE-designed policies and Bayesian feedback are able to deliver $\wp > 1/2$ for all four noise models until $V = 3$ at which point scaling in the presence of random-telegraph noise fails to exceed the SQL. For noise models that are asymmetric, quantum-enhanced precision is observed up to $V = 7$ at skewness $\gamma = 0.8509$. This high level of robustness is surprising, especially so for Bayesian feedback, as the interferometer dynamic no longer agrees with the noiseless assumption.

The skewness in the phase-noise distribution appears to make QEAPE more robust at least up to $\gamma = 0.8509$ as observed in the robustness thresholds from the normal-distributed noise and the skew-normal noise. The normal distribution is a special case ($\gamma = 0$) of the skew-normal distribution, and when a normal-distributed noise is added, the robustness threshold is $V = 3$ whereas the presence of a skewed distribution increases the threshold to $V = 7$. The robustness threshold, of course, cannot increase with skewness indefinitely and would have to either reach an asymptote or decrease, although predicting the outcome is difficult as the phase shift is a periodic variable and the wrapping of the phase-shift distribution has to be taken into account.

The robustness of QEAPE appears to come only in small part from the feedback-control policy. The role of the feedback policy is highlighted by the threshold for normal-distribution noise, where the threshold for DE-designed policies is at $V = 3$ as opposed to Bayesian feedback at $V = 4$. However, for other noise distributions, the threshold are comparable, which indicates that other factors are of greater importance in determining the robustness of QEAPE.

One possible factor that creates robust QEAPE is the input state, which, in this work, is the sine state (23). We know that the permutational symmetry enables the state to be robust against loss [69], and we speculate that this symmetry maybe a necessary property for robustness against phase noise as well. However, permutational symmetry alone is insufficient to determine whether the quantum state is able to deliver quantum-enhanced precision. An example of this is the product state, which shows no quantum enhancement whatsoever [41], as expected intuitively when quantum resources such as entanglement or squeezing are not employed. The structure of the permutational symmetric state itself has to be studied to determine the key features that enables both quantum-enhanced precision and practical metrology.

B. Space cost and power-law scalings

Table I shows that power-law scaling delivered by Bayesian feedback is consistently superior to those scalings delivered by the DE-designed policies before the robustness threshold is reached. Here we use the space complexity of the policies in Table II to explain the reason behind this difference.

Table II shows that Bayesian feedback has a poly$N$ space cost that scales one polynomial degree higher than for DE-designed policies, which indicates that Bayesian feedback utilizes more information and hence is more complex than the DE-designed policies. By using a trusted quantum-state model, Bayesian feedback effectively uses the history of measurement outcomes $x_1 x_2 \ldots x_m$ to determine Φ_{m+1} instead of using just the current outcome x_m, which is the approach used by DE-designed policies in (31). As such, DE-designed policies are restricted by the generalized-logarithm-search strategy (Sec. II B 2) and so cannot deliver a value of $\wp_S$ that approaches the HL. Improvement of DE-designed policies is done by changing the update rule, i.e., the structure of the policy, so that ϱ uses a part of the measurement history.

C. Choosing a QEAPE policy

In this subsection, we explain how the space and execution time complexity (Sec. III C 2) is used to decide between competing policies and method of generating the policies. In particular, we discuss choosing between the DE-designed policies and Bayesian feedback.

The consideration of the space and time complexity of the policies comes after ascertaining that the candidate policies are able to deliver the target performance. In the case of QEAPE, the target is to attain $\wp > 1/2$, which both the DE-designed policies and Bayesian feedback are able to deliver. Both methods also have the same robustness threshold against phase noise of unknown distribution. Based on these comparisons, both policies appear to be equally suitable.

When comparing space and execution complexity (Sec. IV C), DE-designed policies show an advantage as scaling of both costs are bound to linear scaling with respect to particle number N, whereas Bayesian feedback is quadratic in space complexity and cubic in time complexity. For this reason, we favor DE-designed policies for robust QEAPE.

We do not consider design complexity as this cost can be offset by parallelizing the training. This cost may be of interest when learning occurs in a physical setup where parallelizing is not possible and one shot of the experiment is expensive. In this case, there could be an upper bound to the number of

experiments and hence time that can be invested in training ϱ. In that case, design complexity becomes relevant.

VI. CONCLUSION

We have tested quantum-enhanced adaptive-phase estimation (QEAPE) policy robustness based both on Bayesian feedback and on direct policy search. We compare performance with respect to resource complexities for designing and for implementing policies by numerical simulation and extrapolating beyond the maximum photon number $N_{\max}$ using the final function of our five fitted piecewise functions. By this method, we observe that both the DE-designed and the Bayesian-feedback policies are robust against phase noise up to the noise level characterized by $V = 3$. Although scaling of phase-estimate precision provided by Bayesian feedback is superior to the policies designed by direct policy search, direct policy search consumes far less time and space resource than does Bayesian feedback for policy execution.

Our methods for robustness testing and for comparing policies based on resource complexity could have applications to AQEM and quantum-control applications in general. Robustness is a necessary property of any QEM schemes so our tests are valuable and, moreover, can be adapted to quantify robustness of nonadaptive procedures and investigate the role of the input state to the robustness of QEM generally. Quantifying the resource used by control policies can be used to show efficacy of the policies not only in AQEM but in other quantum control tasks, and the comparison of the resource complexity can be used to select a policy that is most efficient in accomplishing a task.

In summary, we have introduced an approach to evaluating quantum-control policies for metrology that relies on testing against unusual phase-noise models and accepting that the policies are robust only if scaling of phase-estimate precision beats the standard quantum limit and policy-design resource complexity, in the presence of phase noise, is poly$N$. We have made explicit what is known in this field—that power-law behavior is observed by numerical simulation up to some $N_{\max}$ and extrapolated to claim asymptotic scaling—and introduced fitting the log-log plot of V_{H} vs N so that only the final (applicable for the largest values of N in the domain) of the piecewise-fitted functions is used to extrapolate large N behavior. We have also explained why time-homogenous analysis, which is used to prove no-go no-go results for QEM asymptotically for large N [86], does not evidently apply to our case.

ACKNOWLEDGMENTS

The computational work was enabled by the support of WestGrid [105] through Compute Canada Calcul Canada [106] and this project has been supported by NSERC and Alberta Innovates.

APPENDIX A: DIFFERENTIAL EVOLUTION

Differential evolution is an optimization algorithm based on a heuristic from a biological evolutionary process [40]. As such, differential evolution is an evolutionary algorithm [27]. This class of optimization algorithms operates by performing a global heuristic search in the solution space and selects or rejects solution candidates based on the candidate's fitness

Algorithm 1. Noise-resistant differential evolution.

procedure INITIALIZE CANDIDATES
  for all candidate in population of size N_{p} **do**
    Generate random vector $\boldsymbol{V}^{(p)}(G=1)$ of size N
    Calculate average fitness of $\boldsymbol{V}^{(p)}(G=1)$ from 2 samples of fitness
  end for
end procedure
procedure FIND A FEASIBLE SOLUTION
  while termination condition not yet met **do**
    for each candidate **do**
      Calculate fitness of $\boldsymbol{V}^{(p)}(G)$
      Update average fitness
    end for
    Procedure CREATING $(G+1)$ POPULATION
      for each candidate **do**
        Randomly select 3 members from population$(\boldsymbol{V}^{(p,1)}(G), \boldsymbol{V}^{(p,2)}(G), \boldsymbol{V}^{(p,3)}(G))$
        for each dimension i of $\boldsymbol{V}^{(p)}(G)$ **do**

$$\boldsymbol{D}_i^{(p)}(G) = \begin{cases} V_i^{(p,1)}(G) + \mathcal{F}[V_i^{(p,2)}(G) - V_i^{(p,3)}(G)] & \text{if rand} \leqslant Cr, \\ V_i(G) & \text{otherwise} \end{cases}$$

        end for
        Calculate average fitness of $\boldsymbol{D}^{(p)}(G)$ using 2 samples of fitness
        Select $\boldsymbol{V}^{(p)}(G)$ or $\boldsymbol{D}^{(p)}(G)$ to be $\boldsymbol{V}^{(p)}(G+1)$ whichever gives higher average fitness
      end for
    end procedure
  end while
  Select $\boldsymbol{V}^{(p)}$ with the highest average fitness as the solution
end procedure

value, which indicates how well the solution solves the optimization problem. Here we summarize a variant of noise-resistant DE [23] (Algorithm 1 that we devised based on one of the original variants of DE [40].

As in other evolutionary algorithms, the solution candidate p in a population of $N_{\rm p}$ candidates for the Gth generation is represented by a vector

$$\boldsymbol{V}^{(p)}(G) = \left(V_1^{(p)}(G), V_2^{(p)}(G), \ldots, V_N^{(p)}(G)\right), \tag{A1}$$

where N is the size of the solution space, which in the case of QEAPE is the number N of photons. A candidate's ability to solve the optimization problem is quantified using a fitness function such that the maximum fitness value corresponds to the optimal solution. The heuristic of an evolutionary algorithm typically only uses these fitness values in the optimization process, and so the algorithm can be applied to control problems where the system's dynamic is not fully known but the fitness of the control procedure can be observed.

What distinguishes DE from other evolutionary algorithms is the procedure for creating the offspring,

$$\boldsymbol{D}^{(p)}(G) = \left(D_1^{(p)}(G), D_2^{(p)}(G), \ldots, D_N^{(p)}(G)\right) \tag{A2}$$

(line 16 of Algorithm 1), and the selection rule (line 19 of Algorithm 1). The term "differential" refers to the difference term $[V_i^{(p,2)}(G) - V_i^{(p,3)}(G)]$. Randomness introduced in the selection of

$$[\boldsymbol{V}^{(p,1)}(G), \boldsymbol{V}^{(p,2)}(G), \boldsymbol{V}^{(p,3)}(G)] \tag{A3}$$

enables the algorithm to execute the search of the solution space whereas the greedy-selection rule enables the algorithm to converge quickly to an area with good solutions [107]. This rapid convergence on good solutions makes DE an attractive algorithm for global optimization problems.

We note that DE can perform poorly when noise is included in the fitness function [108,109], which is important for our application. Thus, we devise a variant of the DE algorithm that is noise-resistant (Algorithm 1) based on the idea that averaging reduces the noise in the fitness landscape and thus enables the algorithm to determine the quality of the solution with accuracy. Hence, we use the average fitness value in the selection procedure, which enables us to find policies that are able to break the SQL in the presence of phase noise [23].

APPENDIX B: REGRESSION ANALYSIS

The imprecision $\Delta\tilde{\phi}$ and N are asymptotically power-law related (1). However, when the system is noisy, this relationship fails for low N, with the actual bound on N depending on the noise model. We employ regression analysis to select the subset of $V_{\rm H}$ at high N that scales as $N^{-\wp}$ and estimate the corresponding $\wp$ by building piecewise functions and selecting the best candidate to represent the data. In this section, we explain our regression-analysis procedure for fitting a model given a set of data.

1. Fitting the model

Regression analysis aims to determine the mathematical relationship between dependent ($V_{\rm H}$ here) and independent variables (here N) [110]. The process of building this mathematical model begins with selecting a function $f(N)$ based on the knowledge of the mechanism and observations of the trends [111]. The function is only a best guess as the discerned trend could be subjective.

After a function is selected, the function is then fitted to the data by finding the parameters that minimize the error between the predicted $V_{\rm H}$ and the data $V_{\rm H}$ [112]. The method we employ is the least-squares estimation [102], used in linear regression to calculate the variables by constraining the gradients to zero and solving the resulting system of linear equations. We choose this method as we fit linear and piecewise linear equations to a log-log plot of $V_{\rm H}$ and N.

2. Consistency of the fit

As the fitted function is only an educated guess, the fitting result must be examined for inconsistencies with respect to the model's assumptions [112]. An alternative function can then be proposed, fitted, and compared to the previous function in order to find one that best represents the data. Deciding on the best model from the set is done using statistical criteria that either estimate the goodness of the fit to the data or between two models fitted to the same data [113]. The model that is consistently shown to fit well according to each of the criteria is then selected to represent the data.

Common linear-regression criteria include

(1) Coefficient of determination:

$$R^2 = 1 - \frac{\sum_N \left(V_{\rm H}^{(N)} - f(N)\right)^2}{\sum_N \left(V_{\rm H}^{(N)} - \bar{V}_{\rm H}\right)}, \tag{B1}$$

adjusted to

$$\overline{R^2} = R^2 - \frac{b}{v - b - 1}(1 - R^2), \tag{B2}$$

with v the number of data points in the fit.

(2) Corrected Akaike information criteria: $\mathrm{AIC_c}$ quantifies information lost due to the discrepancy between the model function and the true function $g(N)$, and is "corrected" to avoid overfitting,

(3) F test assesses a full model (maximum number of parameters b), as the null hypothesis, vs a reduced model (reduction from the full model) as the alternative hypothesis [111,114], and

(4) Mallows's C_p [111,113] (but we use b rather than the traditional p for the number of parameters) estimates the mean-square prediction error [113] to compare a reduced model to the full model, where the reduced model with the smallest C_p close to b is chosen.

Each of these criteria is designed to penalize functions with many parameters b to avoid overfitting the data [111].

[1] V. Giovannetti, S. Lloyd, and L. Maccone, Science **306**, 1330 (2004).

[2] V. Giovannetti, S. Lloyd, and L. Maccone, Nat. Photon. **5**, 222 (2011).

[3] S. L. Braunstein, C. M. Caves, and G. J. Milburn, Ann. Phys. (NY) **247**, 135 (1996).

[4] G. Tóth and I. Apellaniz, J. Phys. A: Math. Theor. **47**, 424006 (2014).

[5] D. Vrajitoru and W. Knight, *Practical Analysis of Algorithms* (Springer, Cham, Switzerland, 2014), Chap. 1, pp. 1–7.
[6] J. N. Hollenhorst, Phys. Rev. D **19**, 1669 (1979).
[7] C. M. Caves, K. S. Thorne, R. W. P. Drever, V. D. Sandberg, and M. Zimmermann, Rev. Mod. Phys. **52**, 341 (1980).
[8] C. M. Caves, Phys. Rev. D **23**, 1693 (1981).
[9] J. J.. Bollinger, W. M. Itano, D. J. Wineland, and D. J. Heinzen, Phys. Rev. A **54**, R4649 (1996).
[10] J. Borregaard and A. S. Sørensen, Phys. Rev. Lett. **111**, 090801 (2013).
[11] S. Danilin, A. V. Lebedev, A. Vepsäläinen, G. B. Lesovik, G. Blatter, and G. S. Paraoanu, npj Quantum Inf. **4**, 29 (2018).
[12] Z. Zhang and L. M. Duan, New J. Phys. **16**, 103037 (2014).
[13] J. C. F. Matthews, X.-Q. Zhou, H. Cable, P. J. Shadbolt, D. J. Saunders, G. A. Durkin, G. J. Pryde, and J. L. O'Brien, Npj Quantum Inf. **2**, 16023 (2016).
[14] H. M. Wiseman, Phys. Rev. Lett. **75**, 4587 (1995).
[15] M. A. Armen, J. K. Au, J. K. Stockton, A. C. Doherty, and H. Mabuchi, Phys. Rev. Lett. **89**, 133602 (2002).
[16] D. W. Berry and H. M. Wiseman, Phys. Rev. Lett. **85**, 5098 (2000).
[17] U. Rosolia, X. Zhang, and F. Borrelli, Annu. Rev. Control, Robotics, Autonomous Systems **1**, 259 (2018).
[18] S. Roy, I. R. Petersen, and E. H. Huntington, New J. Phys. **17**, 063020 (2015).
[19] H. M. Wiseman and R. B. Killip, Phys. Rev. A **56**, 944 (1997).
[20] H. M. Wiseman and R. B. Killip, Phys. Rev. A **57**, 2169 (1998).
[21] A. Hentschel and B. C. Sanders, Phys. Rev. Lett. **104**, 063603 (2010).
[22] N. B. Lovett, C. Crosnier, M. Perarnau-Llobet, and B. C. Sanders, Phys. Rev. Lett. **110**, 220501 (2013).
[23] P. Palittapongarnpim, P. Wittek, E. Zahedinejad, S. Vedaie, and B. C. Sanders, Neurocomputing **268**, 116 (2017).
[24] B. M. Escher, R. L. de Matos Filho, and L. Davidovich, Nat. Phys. **7**, 406 (2011).
[25] R. Demkowicz-Dobrzanski, J. Kołodynski, and M. Guţă, Nat. Commun. **3**, 1063 (2012).
[26] D. W. Berry, H. M. Wiseman, and J. K. Breslin, Phys. Rev. A **63**, 053804 (2001).
[27] A. E. Eiben and J. E. Smith, *Introduction to Evolutionary Computing*, Natural Computing, 2nd ed. (Springer, Berlin, Germany, 2015), Chap. 6, pp. 99–116.
[28] N. Bobroff, Appl. Opt. **26**, 2676 (1987).
[29] L. C. Sinclair, F. R. Giorgetta, W. C. Swann, E. Baumann, I. Coddington, and N. R. Newbury, Phys. Rev. A **89**, 023805 (2014).
[30] I. Filinski and R. A. Gordon, Rev. Sci. Instrum. **65**, 575 (1994).
[31] T. A. Severini, *Elements of Distribution Theory*, Cambridge Series in Statistical and Probabilistic Mathematics (Cambridge University Press, Cambridge, UK, 2005), Chap. 12, pp. 365–399.
[32] E. Breitenberger, Biometrika **50**, 81 (1963).
[33] D. S. Newman, Ann. Math. Stat. **39**, 890 (1968).
[34] A. Azzalini and A. Capitanio, in *The Skew-Normal and Related Families*, edited by D. R. Cox, A. Agresti, B. Hambly, S. Holmes, and X.-L. Meng, Institute of Mathematical Statistics Monograph (Cambridge University Press, Cambridge, UK, 2014), Chap. 2, pp. 24–56.
[35] E. Limpert, W. A. Stahel, and M. Abbt, BioScience **51**, 341 (2001).
[36] M. Pourahmadi, Commun. Stat. Theory Methods **36**, 1803 (2007).
[37] W. Jouini, IEEE Signal Process. Lett. **18**, 423 (2011).
[38] J. Rezaie and J. Eidsvik, Comput. Stat. Data Anal. **75**, 1 (2014).
[39] D. Vrajitoru and W. Knight, *Practical Analysis of Algorithms* (Ref. [5]), Chap. 5, pp. 169–293.
[40] R. Storn and K. Price, J. Global Optim. **11**, 341 (1997).
[41] P. Palittpongarnpim, P. Wittek, and B. C. Sanders, in *Proceedings of the SPIE Quantum Communications and Quantum Imaging XIV* (SPIE, Bellingham, WA, 2016), Vol. 9980, pp. 99800H–99800H-11.
[42] C. W. Helstrom, J. Stat. Phys. **1**, 231 (1969).
[43] P. C. Humphreys, M. Barbieri, A. Datta, and I. A. Walmsley, Phys. Rev. Lett. **111**, 070403 (2013).
[44] J.-D. Yue, Y.-R. Zhang, and H. Fan, Sci. Rep. **4**, 5933 (2014).
[45] B. L. Higgins, D. W. Berry, S. D. Bartlett, H. M. Wiseman, and G. J. Pryde, Nature (London) **450**, 393 (2007).
[46] R. Demkowicz-Dobrzański and L. Maccone, Phys. Rev. Lett. **113**, 250801 (2014).
[47] M. Hotta, T. Karasawa, and M. Ozawa, Phys. Rev. A **72**, 052334 (2005).
[48] M. Sbroscia, I. Gianani, L. Mancino, E. Roccia, Z. Huang, L. Maccone, C. Macchiavello, and M. Barbieri, Phys. Rev. A **97**, 032305 (2018).
[49] K. Landsman, *Foundations of Quantum Theory: From Classical Concepts to Operator Algebras* (Springer International, Cham, Switzerland, 2017), Chap. 4, pp. 103–123.
[50] A. S. Holevo, *Quantum Systems, Channels, Information: A Mathematical Introduction*, De Gruyter Studies in Mathematical Physics Vol. 16 (Walter de Gruyter, Berlin, Germany, 2012), pp. 103–131.
[51] B. Yurke, S. L. McCall, and J. R. Klauder, Phys. Rev. A **33**, 4033 (1986).
[52] M. Hayashi, *Quantum Information: An Introduction* (Springer, Berlin, 2006), pp. 9–25.
[53] H. M. Wiseman and G. J. Milburn, *Quantum Measurement and Control* (Cambridge University Press, Cambridge, MA, 2009).
[54] M. R. James and H. I. Nurdin, in *2015 IEEE Conference on Control Applications (CCA)* (IEEE CCA, Sydney, NSW, 2015), pp. 1–12.
[55] P. H. Lewis and C. Yang, *Basic Control Systems Engineering* (Prentice-Hall, Upper Saddle River, NJ, 1997), Chap. 1, pp. 1–12.
[56] S. Schreppler, N. Spethmann, N. Brahms, T. Botter, M. Barrios, and D. M. Stamper-Kurn, Science **344**, 1486 (2014).
[57] J. Kołodyński and R. Demkowicz-Dobrzański, New J. Phys. **15**, 073043 (2013).
[58] L. Maccone, Phys. Rev. A **88**, 042109 (2013).
[59] S. M. Kay, in *Fundamentals of Statistical Signal Processing: Estimation Theory*, edited by A. V. Oppenheimer (Prentice-Hall, Upper Saddle River, NJ, 1993), Chap. 3.
[60] S. L. Braunstein, Phys. Rev. A **71**, 055801 (2005).
[61] W. K. Wootters, Philos. Trans. R. Soc. London A **356**, 1717 (1998).

[62] K. Jacobs, *Qauntum Measurement Theory and its Applications* (Cambridge University Press, Cambridge, UK, 2014), Chap. 6, pp. 303–322.

[63] M. Zwierz, C. A. Pérez-Delgado, and P. Kok, Phys. Rev. A **85**, 042112 (2012).

[64] R. S. Bondurant and J. H. Shapiro, Phys. Rev. D **30**, 2548 (1984).

[65] H. M. Wiseman, D. W. Berry, S. D. Bartlett, B. L. Higgins, and G. J. Pryde, IEEE J. Sel. Top. Quantum Electron. **15**, 1661 (2009).

[66] A. Hentschel and B. C. Sanders, Phys. Rev. Lett. **107**, 233601 (2011).

[67] J. A. Armstrong, J. Opt. Soc. Am. **56**, 1024 (1966).

[68] P. Hariharan and B. C. Sanders, in *Progress in Optics*, edited by E. Wolf (Elsevier, Amsterdam, 1996), Vol. 36, Chap. 2, pp. 49–128.

[69] A. Hentschel and B. C. Sanders, J. Phys. A: Math. Theor. **44**, 115301 (2011).

[70] G. Summy and D. Pegg, Opt. Commun. **77**, 75 (1990).

[71] M. I. Mishchenko, Wigner d-functions, in *Electromagnetic Scattering by Particles and Particle Groups: An Introduction* (Cambridge University Press, Cambridge, UK, 2014), Appendix F, pp. 385–389.

[72] M. J. Holland and K. Burnett, Phys. Rev. Lett. **71**, 1355 (1993).

[73] B. C. Sanders and G. J. Milburn, Phys. Rev. Lett. **75**, 2944 (1995).

[74] K. F. Riley, M. P. Hobson, and S. J. Bence, *Mathematical Methods for Physics and Engineering: A Comprehensive Guide*, 3rd ed. (Cambridge University Press, Cambridge, UK, 2006), Chap. 30, pp. 1119–1220.

[75] Z. Hradil, R. Myška, J. Peřina, M. Zawisky, Y. Hasegawa, and H. Rauch, Phys. Rev. Lett. **76**, 4295 (1996).

[76] M. P. Deisenroth, G. Neumann, and J. Peters, *A Survey on Policy Search for Robotics*, edited by H. Christensen and R. Siegwart, Foundations and Trends Vol. 2 (now publishers, Delft, 2013).

[77] J. R. Leigh, *Control Theory*, 2nd ed. (Institution of Engineering and Technology, London, UK, 2004), Chap. 6, pp. 61–81.

[78] S. Wang, J. M. Simkoff, M. Baldea, L. H. Chiang, I. Castillo, R. Bindlish, and D. B. Stanley, in *11th IFAC Symp. Dynamics and Control of Process Systems Including Biosystems 2016* (Elsevier, Amsterdam, 2016), Vol. 49, pp. 25–30.

[79] D. E. Moriarty, A. C. Schultz, and J. J. Grefenstette, J. Artif. Intell. Res. **11**, 241 (1999).

[80] T. Rückstieß, F. Sehnke, T. Schaul, D. Wierstra, Y. Sun, and J. Schmidhuber, Paladyn J. Behav. Robot. **1**, 14 (2010).

[81] A. Datta, L. Zhang, N. Thomas-Peter, U. Dorner, B. J. Smith, and I. A. Walmsley, Phys. Rev. A **83**, 063836 (2011).

[82] R. Nichols, T. R. Bromley, L. A. Correa, and G. Adesso, Phys. Rev. A **94**, 042101 (2016).

[83] Z. Huang, C. Macchiavello, and L. Maccone, Phys. Rev. A **97**, 032333 (2018).

[84] T. A. Wheatley, D. W. Berry, H. Yonezawa, D. Nakane, H. Arao, D. T. Pope, T. C. Ralph, H. M. Wiseman, A. Furusawa, and E. H. Huntington, Phys. Rev. Lett. **104**, 093601 (2010).

[85] R. Demkowicz-Dobrzański, J. Czajkowski, and P. Sekatski, Phys. Rev. X **7**, 041009 (2017).

[86] A. Smirne, J. Kołodyński, S. F. Huelga, and R. Demkowicz-Dobrzański, Phys. Rev. Lett. **116**, 120801 (2016).

[87] A. Fujiwara and H. Imai, J. Phys. A: Math. Theor. **41**, 255304 (2008).

[88] A. W. Chin, S. F. Huelga, and M. B. Plenio, Phys. Rev. Lett. **109**, 233601 (2012).

[89] E. M. Kessler, I. Lovchinsky, A. O. Sushkov, and M. D. Lukin, Phys. Rev. Lett. **112**, 150802 (2014).

[90] X.-M. Lu, S. Yu, and C. H. Oh, Nat. Commun. **6**, 7282 (2015).

[91] Y. Matsuzaki, S. Benjamin, S. Nakayama, S. Saito, and W. J. Munro, Phys. Rev. Lett. **120**, 140501 (2018).

[92] P. Sekatski, M. Skotiniotis, and W. Dür, New J. Phys. **18**, 073034 (2016).

[93] T. A. Severini, *Elements of Distribution Theory*, Cambridge Series in Statistical and Probabilistic Mathematics (Cambridge University Press, Cambridge, UK, 2005), Chap. 4, pp. 94–131.

[94] L. Kirkup, *Data Analysis for Physical Scientists: Featuring Excel®*, 2nd ed. (Cambridge University Press, Cambridge, UK, 2012), Chap. 3, pp. 90–145.

[95] C. Boncelet, *Handbook of Image and Video Processing*, edited by A. Bovik (Elsevier, Amsterdam, 2005), pp. 397–409.

[96] P. Acarnley, *Stepping Motors - A Guide to Theory and Practice*, 4th ed. (Institution of Engineering and Technology, London, 2002), Chap. 1, pp. 1–12.

[97] K. F. Riley, M. P. Hobson, and S. J. Bence, *Mathematical Methods for Physics and Engineering: A Comprehensive Guide*, 3rd ed. (Cambridge University Press, Cambridge, UK, 2006), Chap. 18, pp. 577–647.

[98] E. A. Kish, C.-G. Granqvist, A. Dér, and L. B. Kish, Cogn. Neurodyn. **9**, 459 (2015).

[99] S. Kai, S. Higaki, M. Imasaki, and H. Furukawa, Phys. Rev. A **35**, 374 (1987).

[100] H. Mouri, Phys. Rev. E **88**, 042124 (2013).

[101] P. Palittapongarnpim, P. Wittek, and B. C. Sanders, in *2017 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, IEEE (IEEE Systems, Man, and Cybernetics Society, Banff, AB, 2017), pp. 294–299.

[102] D. C. Montgomery, E. A. Peck, and G. G. Vining, *Introduction to Linear Regression Analysis*, Wiley Series in Probability and Statistics, 5th ed. (Wiley, Hoboken, 2012), Chap. 2, pp. 12–66.

[103] C. Jekel, Fitting a piecewise linear function to data, https://jekel.me/2017/Fit-a-piecewise-linear-function-to-data/ (2017).

[104] N. Khaneja, T. Reiss, C. Kehlet, T. Schulte-Herbrüggen, and S. J. Glaser, J. Magn. Reson. **172**, 296 (2005).

[105] www.westgrid.ca.

[106] www.computecanada.ca.

[107] S. Das and P. N. Suganthan, IEEE Trans. Evol. Comput. **15**, 4 (2011).

[108] T. Krink, B. Filipic, and G. B. Fogel, in *Proceedings of the 2004 Congress on Evolutionary Computation* (IEEE, Piscataway, NJ, 2004), Vol. 1, pp. 332–339.

[109] J. Vesterstrom and R. Thomsen, in *Proceedings of the Congress on Evolutionary Computation, 2004 (CEC2004)* (IEEE, Piscataway, NJ, 2004), Vol. 2, pp. 1980–1987.

[110] D. C. Montgomery, E. A. Peck, and G. G. Vining, *Introduction to Linear Regression Analysis*, Wiley Series in Probability and Statistics, 5th ed. (Wiley, Hoboken, NJ, 2012), Chap. 1, pp. 1–11.

[111] S. Chatterjee and J. S. Simonoff, *Handbook of Regression Analysis* (Wiley, Hoboken, NJ, 2013), Chap. 2, pp. 23–52.

[112] S. Chatterjee and A. S. Hadi, *Regression Analysis by Example*, Wiley Series in Probablity and Statistics, 5th ed. (Wiley, Hoboken, NJ, 2012), Chap. 1, pp. 1–24.

[113] D. C. Montgomery, E. A. Peck, and G. G. Vining, *Introduction to Linear Regression Analysis*, Wiley Series in Probability and Statistics, 5th ed. (Wiley, Hoboken, NJ, 2012), Chap. 10, pp. 327–371.

[114] S. Chatterjee and J. S. Simonoff, *Handbook of Regression Analysis* (Wiley, Hoboken, NJ, 2013), Chap. 1, pp. 3–22.