

This is the accepted manuscript made available via CHORUS. The article has been published as:

Rapid estimation of drifting parameters in continuously measured quantum systems

Luis Cortez, Areeya Chantasri, Luis Pedro García-Pintos, Justin Dressel, and Andrew N. Jordan

Phys. Rev. A **95**, 012314 — Published 12 January 2017

DOI: [10.1103/PhysRevA.95.012314](https://doi.org/10.1103/PhysRevA.95.012314)

Rapid estimation of drifting parameters in continuously measured quantum systems

Luis Cortez*,^{1,2,3} Areeya Chantasri*,^{1,2} Luis Pedro García-Pintos,⁴ Justin Dressel,^{4,5} and Andrew N. Jordan^{1,2,4}

¹*Department of Physics and Astronomy, University of Rochester, Rochester, New York 14627, USA*

²*Center for Coherence and Quantum Optics, University of Rochester, Rochester, New York 14627, USA*

³*Facultad de Ciencias Físico Matemáticas, Universidad Autónoma de Nuevo León, Nuevo León 66455, Mexico*

⁴*Institute for Quantum Studies, Chapman University, Orange, California 92866, USA*

⁵*Schmid College of Science and Technology, Chapman University, Orange, California 92866, USA*

(Dated: December 19, 2016)

We investigate the determination of a Hamiltonian parameter in a quantum system undergoing continuous measurement. We demonstrate a computationally rapid method to estimate an unknown and possibly time-dependent parameter, where we maximize the likelihood of the observed stochastic readout. By dealing directly with the raw measurement record rather than the quantum state trajectories, the estimation can be performed while the data is being acquired, permitting continuous tracking of the parameter during slow drifts in real time. Furthermore, we incorporate realistic nonidealities, such as decoherence processes and measurement inefficiency. As an example, we focus on estimating the value of the Rabi frequency of a continuously measured qubit, and compare maximum likelihood estimation to a simpler fast Fourier transform. Using this example, we discuss how the quality of the estimation depends on both the strength and duration of the measurement; we also discuss the trade-off between the accuracy of the estimate and the sensitivity to drift as the estimation duration is varied.

I. INTRODUCTION

The problem of accurately measuring unknown parameters in an experimental system is of both fundamental and practical importance. The process of determining the parameters of the experiment serves as a valuable calibration of the experiment, and also determines the limitations of experimental accuracy. The understanding of how to minimize parameter uncertainties given a variety of possible measurement strategies has been developed into the science of quantum metrology over the past several decades [1, 2].

The usual approach taken in the laboratory is to repeatedly perform the following sequence of operations: prepare a quantum state, let it evolve unitarily in a way that depends on the unknown parameter, and then perform a (potentially unsharp) measurement to extract information about the parameter of interest. The concept of the quantum Fisher information [3] is fundamental to this approach, since it sets the bound on the minimum variance for all possible unbiased estimation strategies based on this final measurement, the so-called *Cramér-Rao bound*. However, there are other methods for estimating such an unknown parameter that go beyond the prepare-evolve-measure paradigm, which may be beneficial under certain circumstances. One such scenario is when the parameter changes slowly in time such that this variation cannot be predicted in advance, as is common in experimental laboratories (e.g., from thermal fluctuations). In this case, it is beneficial to be able to *continuously track* the changing parameter as it evolves in time, to sense when and how it is changing. Given knowledge of how the parameters are changing, introducing feedback control to stabilize the parameter then becomes possible [6, 7]. Such a situation brings into play the physics of open quantum systems and how they relate to metrology

[8–12].

To continuously track the changes of a parameter in time, it is natural to consider measurements that are also continuous in time [3–5]. In order to increase the speed of the estimation, the technique builds upon prior parameter information obtained from an initial broad system characterization. Assuming relatively slow drift of the parameter, the initial characterization narrows the search region of subsequent repeated estimations in real-time as a single noisy measurement record is monitored.

Previously, the physics of parameter estimation using continuous measurements has been analyzed by several groups. Ralph, Jacobs, and Hill [13] studied a weak continuous measurement of a quantum system, focusing on the quality of quantum state estimation affected by the estimate error in an unknown Hamiltonian parameter (an oscillation frequency of a qubit). Our work is closely related to theirs; however, we instead focus on the unknown parameter estimation and how to speed up the estimation using propagators and Fourier methods, avoiding numerically integrating a stochastic master equation used in [13]. Klaus Mølmer and collaborators developed the techniques in classical estimation theory for quantum measurement, showing how the parameter estimation can be carried out by Bayesian estimation in solving stochastic master equations, presented with specific examples of fluorescence two-level systems [14, 15]. They also showed how Fisher information is degraded in the quantum Zeno regime [16]. Mankei Tsang and collaborators have investigated the generalization of optimal Wiener and Kalman classical signal smoothing to quantum parameter estimation [17–22], and have more recently proposed using suboptimal Volterra filters for computational efficiency [23]. The potential use of continuous measurements for quantum state tomography is also starting to be explored [24–26].

In this work, we consider a quantum system undergoing such a continuous measurement, which gives rise to a stochastic measurement record that can be monitored in time. We wish to analyze this data to extract the value of an unknown parameter in a way that is both computationally rapid and statistically efficient to permit estimation on a short enough time scale where feedback control becomes possible to correct drift. For specificity, we focus on the determination of an unknown and drifting Rabi drive for the qubit. There are a number of open problems in this area which we now consider: How can one minimize the complexity of maximum likelihood algorithms so they are computationally fast? Is it possible to work only with raw measurement data, so numerical implementations of quantum filters that estimate the quantum state dynamically are not needed? Is it possible to incorporate numerically more efficient methods to narrow down the parameter search space? Can existing methods be generalized to account for experimental nonidealities, such as additional dephasing, detector inefficiency and energy relaxation? In this paper we work toward solutions to these problems, and outline how they can be experimentally implemented in superconducting circuits, among other possibilities. Continuous measurement with superconducting circuits have a proven ability to accurately track the quantum state in time [27] with excellent agreement with predicted statistics [28].

Our basic insight is to speed up the estimation protocol by avoiding the numerical integration of the stochastic master equation. Rather, we construct effective propagators directly from the observed measurement record that can be used in the maximum likelihood algorithm. These effective propagators use measurement operators with the sequence of digitized measurement results from, e.g., a homodyne measurement, together with unitary matrices with unknown Rabi drive frequency. These can all be evaluated numerically using the particular realization of the stochastic measurement results. By simple multiplication of the composite matrices in the effective propagator, a suitable likelihood function is straightforwardly constructed with the initial state, which can then be maximized. Such maximum likelihood estimation (MLE) saturates the Cramér-Rao bound for large data sets. Combining this MLE technique with an initial fast Fourier transform (FFT) technique, (which identifies a range of prior Rabi frequencies) provides a significant speed up by decreasing the number of trial frequencies needed for MLE. We further generalize this method to incorporate realistic non-idealities to prepare this method for experimental implementation in superconducting circuit architectures, where continuous homodyne and heterodyne measurement are now routinely carried out. We believe this method will be suitable to be directly programmed into a field-programmable gate array (FPGA) for rapid near-real-time implementation. Finally, we illustrate how this method can be applied to a time-varying, unknown Rabi drive, and show that we can accurately track even irregular motion in time. Notably, a *projective* measure-

ment version of this task has been accomplished by Shulman, Harvey, Nichol, *et al.* using such an FPGA in a triplet/singlet spin qubit in order to detect how the surrounding nuclear magnetic field was changing, and incorporate feedback to prolong the qubit dephasing time [29]. We thus present our own analysis of the projective equivalent with fixed spacing between measurements in the Appendix as a comparison to the present work.

The paper is organized as follows. In Sec. II, we outline our strategy for determining the value of a static Hamiltonian parameter by maximizing the likelihood of observing a particular stochastic measurement record. We focus on the example of a driven qubit, where the parameter to be estimated is an initially static Rabi frequency. We then compare the maximum likelihood approach to a simpler FFT, which can be used to help identify a suitable frequency range for subsequent maximum likelihood estimation, and comment on the relative computational efficiency of each method. In Sec. III we generalize the static estimation method to a dynamic estimation method that is able to track arbitrary time-dependent parameters. Setting a desired estimation precision then specifies the time resolution for the tracking of drift. We demonstrate that we are able to accurately track dynamical parameters using this method. We conclude in Sec. IV. We also provide an Appendix that includes an analytic treatment of the maximum likelihood frequency estimation using periodic projective measurements, for handy comparison to the continuous case.

II. STATIC PARAMETER ESTIMATION

We start by describing the problem. Generally speaking, we consider the estimation of an unknown fixed parameter in the system Hamiltonian, given the output of some quantum measurement device. Suppose, for definiteness, that we are interested in measuring the Rabi oscillation rate of a driven qubit from the measurement output data. The unmeasured qubit is then described by a Hamiltonian

$$H = \frac{\hbar\Omega}{2} \sigma_y, \quad (1)$$

that rotates the qubit state in the x - z plane of the Bloch sphere at an angular frequency Ω , where we notate the usual Pauli operators as $\{\sigma_x, \sigma_y, \sigma_z\}$. In the absence of measurement, the quantum system unitarily evolves for a duration δt , which can be simply represented as a rotation in the configuration basis,

$$U = \begin{pmatrix} \cos \frac{\Omega \delta t}{2} & -\sin \frac{\Omega \delta t}{2} \\ \sin \frac{\Omega \delta t}{2} & \cos \frac{\Omega \delta t}{2} \end{pmatrix}. \quad (2)$$

We now wish to estimate the oscillation frequency Ω .

We begin our analysis by assuming a time-independent oscillation frequency Ω that is unknown beforehand, and

we examine different methods for estimating the parameter. In Sec. II A, we consider a maximum likelihood-based matrix multiplication method to estimate the parameter of interest, showing both idealized and more realistic cases that include experimental nonidealities. In Sec. II B, we compare this method to a simpler method based on FFT, and show that the FFT permits a quick-but-crude estimation of a prior range of frequencies that can be used to help accelerate the convergence of the maximum likelihood procedure.

A. Maximum Likelihood

Maximum likelihood methods presuppose a model that describes the stochastic physics with fixed parameters, then varies each unknown parameter to find the best estimate matching a target data set according to a suitable likelihood measure. In our case, the model is given by the quantum Bayesian update corresponding to repeated unsharp (generalized) measurements applied to a single qubit undergoing Hamiltonian evolution. The likelihood measure is simply the probability for obtaining the observed sequence of measurement results.

1. Ideal Continuous Measurement

The physical set-up of the Hamiltonian of the system is described by Eq. (1), to which we now add a (Markovian) continuous measurement of the σ_z operator. A thorough analysis of such a system may be found in Ref. [30], which considers a solid-state qubit formed by a double-quantum-dot and measured by the current flowing through a nearby quantum point contact. Notably, the analysis found therein applies mostly unchanged to more recent superconducting qubit measurements [31] that use circuit quantum electrodynamics [32, 33]. We consider an idealized such model here, to which we will later add additional experimental imperfections.

Such a continuous measurement weakly probes information about the qubit state coordinate $z(t) \equiv \langle \sigma_z \rangle(t)$ from the system, producing a (suitably renormalized) noisy record $r(t) \approx z(t) + \sqrt{\tau_m} \xi(t)$ that is approximately centered around $z(t)$ and masked with Gaussian white noise $\xi(t)$ (satisfying the correlation $\langle \xi(0)\xi(t) \rangle = \delta(t)$). The *characteristic measurement time-scale* τ_m in the noise-power determines the amount of time needed to distinguish the eigenstates $z = \pm 1$ with unit signal-to-noise ratio [33].

In practice, such continuous readouts $r(t)$ are digitized by the hardware into time bins $t_j \equiv j \delta t$ of duration δt . The output reported by the detector is thus a time-sliced picture of discrete outputs $r_j = \int_{t_j}^{t_{j+1}} r(t) dt / \delta t$, that are temporal averages of the continuous signal $r(t)$ over each time bin t_j . After a total duration $T \equiv N \delta t$, a measurement read-out $\{r_j\}$ is thus produced, consisting of real outputs r_j at each time step $j = 1, \dots, N$. Our goal is

to directly use such a readout $\{r_j\}$ to estimate the Rabi oscillation frequency Ω in Eq. (1).

We use an effective (*quantum Bayesian*) measurement model for such a time-sliced σ_z measurement (see, e.g., Ref. [33, 34]), that models the readout r for each time slice t_j as an independent random variable sampled from a Gaussian mixture distribution,

$$P(r) = \rho_{11} P(r|1) + \rho_{00} P(r|0), \quad (3)$$

$$P(r|0, 1) = \sqrt{\frac{\delta t}{2\pi\tau_m}} \exp \left[-\frac{\delta t(r \pm 1)^2}{2\tau_m} \right], \quad (4)$$

with $P(r|0)$ centered on $r = -1$ and $P(r|1)$ centered on $r = +1$, and where the probabilities ρ_{11} and ρ_{00} are density-matrix elements that correspond to the qubit z -populations at the beginning of the time-slice t_j . For sufficiently short time-slices δt , this Gaussian mixture approximates a single broad Gaussian of variance $\tau_m / \delta t$ that is centered at the qubit coordinate $z \equiv \rho_{11} - \rho_{00}$, thus recovering the Gaussian white noise picture $r(t) \approx z(t) + \sqrt{\tau_m} \xi(t)$ in the continuum limit.

For each time-step t_j with duration $\delta t \ll \tau_m$, the qubit is only weakly perturbed by the measurement while it evolves with the Hamiltonian in Eq. (1). The state back-action between t_j and $t_j + \delta t$ thus combines evolution with a partial-collapse of the prior system state ρ at t_j , and is described by the update rule

$$\rho' = \frac{M_r \rho M_r^\dagger}{\text{Tr}[M_r^\dagger M_r \rho]}, \quad (5)$$

in terms of *measurement operators* M_r that depend upon the readout $r = r_j$ observed in the interval $[t_j, t_{j+1}]$. For time steps $\delta t \ll (2\pi/\Omega)$ much smaller than a Rabi period, each measurement operator $M_r \approx U E_r^{1/2}$ can be approximately decomposed into a unitary part U and a positive operator E_r , such that U is given by the (r -independent, Ω -dependent) unitary evolution of Eq. (2) over the elapsed time δt , and the positive operators E_r are (r -dependent, Ω -independent) elements of a *positive operator-valued measure* (POVM) that is fully determined by the Gaussian probabilities in Eq. (4)

$$E_r = \begin{pmatrix} P(r|0) & 0 \\ 0 & P(r|1) \end{pmatrix}. \quad (6)$$

Note that full collapse of the wavefunction would be obtained by taking $\tau_m \rightarrow 0$ (Zeno measurement regime) [35], in which case E_r would converge to projection operators for each definite z -state of the qubit. Importantly, the probability of having obtained the readout r given the prior state ρ is then $P(r|\rho) = \text{Tr}[M_r \rho M_r^\dagger] = \text{Tr}[E_r \rho]$, which reproduces Eq. (3).

2. Maximum Likelihood Estimation

For the estimation of the unknown value of Ω from the record taken over a total duration $T = N \delta t$, we require

the joint probability distribution for a long sequence of results $\{r_1, r_2, \dots, r_N\}$, which according to Eq. (5) has the simple form

$$P(r_1, \dots, r_N | \Omega) = \text{Tr}[M_N \rho M_N^\dagger], \quad (7)$$

where ρ is the known initial state and where we have defined an effective, multi-index measurement operator,

$$M_N(r_1, \dots, r_N) = M_{r_N} \dots M_{r_2} M_{r_1}, \quad (8)$$

as the simple product of the measurement operators M_{r_j} for each result r_j .

Given no prior information about the value of Ω , we use the log of the distribution in Eq. (7) as our *log-likelihood function* for an observed readout

$$\mathcal{L}(\Omega) = \ln P(r_1, \dots, r_N | \Omega) = \ln \text{Tr}[M_N^\dagger M_N \rho]. \quad (9)$$

That is, given an observed data set (r_1, \dots, r_N) from an experiment, we are interested in finding the *maximum likelihood estimator* Ω_{ML} , such that the distribution in Eq. (7), and therefore Eq. (9), is maximized for $\Omega = \Omega_{ML}$. Importantly, it is *not* necessary to track the quantum state (e.g., by using Eq. (5) or solving a stochastic master equation), since we are only interested in estimating the parameter Ω .

Consequently, only the functional form of the measurement operators are needed, and only the state-dependent part of each M_{r_j} will be relevant for the MLE optimization. As such, the state-independent Gaussian normalization may be conveniently discarded as uninformative. In what follows we rescale the measurement operator $M_{r_j} \mapsto U \tilde{E}_{r_j}^{1/2}$ by using an unnormalized equivalent of Eq. (6) that isolates the state-dependent part

$$\tilde{E}_j = \begin{pmatrix} \exp(-r_j \delta t / \tau_m) & 0 \\ 0 & \exp(r_j \delta t / \tau_m) \end{pmatrix}. \quad (10)$$

This replacement will later help numerical algorithms avoid products of very small numbers.

The Fisher information about the parameter Ω to be estimated is computed directly from the log-likelihood for all possible readouts

$$\mathcal{I}(\Omega) = \int \mathcal{D}r P(r_1, \dots, r_N) (\partial_\Omega \ln P(r_1, \dots, r_N | \Omega))^2, \quad (11)$$

where the measure $\mathcal{D}r$ indicates integration over all possible values of the N measurement results r_j . The Fisher information determines the minimum variance of the parameter Ω for unbiased estimators, the Cramér-Rao bound,

$$\text{Var } \Omega \geq \mathcal{I}^{-1}. \quad (12)$$

As we show in the Appendix for a related special case using periodic projective measurements, this bound on the variance can be saturated for large data sets by choosing the MLE procedure.

The uncertainty σ of the estimate, for large data sets, is approximately with width of a Gaussian fit to the distribution,

$$\sigma^{-2} = -\partial_\Omega^2 \ln P|_{\Omega=\Omega_{ML}}. \quad (13)$$

which can be conveniently found as the width of a parabolic fit to the computed log-likelihood function around the maximum Ω_{ML} , $\mathcal{L}(\Omega) \approx -(\Omega - \Omega_{ML})^2 / 2\sigma^2$. For large data sets, the variance of the maximum likelihood estimator saturates the Cramér-Rao bound.

From this point, there are two ways to proceed. The first semi-analytic way is to examine the maximum likelihood condition, $\partial_\Omega \mathcal{L}(\Omega) = 0$, given the (measured) data set $\{r_j\}$, which then yields the equation

$$\text{Tr}[(\partial_\Omega M_N^\dagger) M_N + M_N^\dagger \partial_\Omega M_N] \rho = 0. \quad (14)$$

This equation must be solved numerically in general. An iterative procedure is also possible since the measurement operator is a product of all previous operators. Therefore, we have the recursion relation

$$\partial_\Omega M_N = (\partial_\Omega M_{r_N}) M_{N-1} + M_{r_N} (\partial_\Omega M_{N-1}), \quad (15)$$

where $M_1 \equiv M_{r_1}$, and where $\partial_\Omega M_{r_j} = (\partial_\Omega U) \tilde{E}_{r_j}^{1/2}$ depend on Ω only through the unitary U . That is, to compute an updated likelihood from the likelihood at the previous step, we can use the already computed M_{N-1} and its derivative, as well as the newly computed measurement matrix M_{r_N} and its derivative to proceed to the next time step. This recursive simplification permits efficient parallel calculation of likelihood values for a range of Ω values during data collection. In order to confirm that the algorithm does not find a local maximum of the likelihood function, sampling in different regions of the parameter space can check the overall value of the likelihood function.

The second (simpler) way is to entirely numerically implement the MLE method. For each result r_j in the string of incoming data (r_1, r_2, \dots, r_N) we numerically compute the relevant matrix $M_{r_j} = U \tilde{E}_{r_j}^{1/2}$ that includes one of a variety of (precomputed) unitary matrices U that assume possible Ω values, as well as the rescaled POVM element \tilde{E}_{r_j} . For each chosen Ω value, we then compute the measurement operator M_N as a product of the string of N matrices M_{r_j} , and compute its associated log-likelihood function. This procedure results in a discretized function $\mathcal{L}(\Omega)$ of log-likelihoods over all sampled Ω . The best estimate Ω_{ML} maximizes this function, while the uncertainty σ is the best quadratic fit around this maximum according to Eq. (13). Adaptive searches for the maximum can be implemented to increase the efficiency of the search.

3. Numerical Simulations

Although frequency is a continuous parameter, the numerical maximum likelihood algorithm outlined above




FIG. 1. Log-likelihood $\mathcal{L}(\Omega)$ as a function of possible Rabi frequencies Ω . A long noisy z -measurement record $\{r_j\}$ of duration $T = 1\text{ms}$ was simulated with $N = 10^5$ discrete time steps $\delta t = 10\text{ns}$, a characteristic measurement time $\tau_m = 1\mu\text{s}$, and a true Rabi oscillation frequency $\Omega_T/2\pi = 1\text{MHz}$ (green solid). The maximum likelihood estimator (red dotted) is the peak at $\Omega_{ML}/2\pi = 1.0025\text{MHz}$ with 0.25% error. The quadratic fit $\mathcal{L} \approx -(\Omega - \Omega_{ML})^2/2\sigma^2$ to the peak yields the precision $\sigma/2\pi = 0.0026\text{MHz}$.

must search over a discrete set of trial values for the frequency. An initial search grid coarse-grains a frequency range of interest in units of $\delta\Omega$, and computes the maximum likelihood over this grid. Both the precision of the estimation and the computational efficiency of the procedure thus depend directly on the size of the search grid. A coarse frequency step size defines a big grid, which makes the search faster but with lower precision, and vice versa. In general, an adaptive mesh size is useful to find the desired precision, where the spacing between sampled Ω is iteratively refined to increase the resolution around the maximum.

As an example of such a purely numerical maximum likelihood method, we simulate a single measurement readout with a chosen true value of the Rabi oscillation frequency $\Omega_T = 2\pi f$ with $f = 1\text{MHz}$, while monitoring with a characteristic measurement time $\tau_m = 1\mu\text{s}$ at discrete time steps $\delta t = 10\text{ns}$ for a duration $T = 1\text{ms}$ (~ 1000 oscillations), and then compute \mathcal{L} for a range of frequencies around Ω_T . We plot the (unnormalized) log-likelihood function in Figure 1, where a dominant peak is clearly seen around the correct value. An estimation of the precision is given by fitting the log-likelihood with a polynomial function (dashed-blue), $-(\Omega - \Omega_{ML})^2/2\sigma^2$, around the peak $\Omega_{ML}/2\pi$. The estimated value of the frequency is $\Omega_{ML}/2\pi = 1.0025\text{MHz}$, which shows 0.0025 MHz (0.25%) error with a precision of $\sigma/2\pi = 0.0026\text{MHz}$, which is close to the minimum anticipated frequency resolution of $1/T = 0.001\text{MHz}$.

The proof-of-principle simulation shown above demonstrates that MLE allows us to determine the oscillation frequency Ω quite precisely using a measurement record about 1000 times the length of the characteristic mea-




FIG. 2. RMS error of maximum likelihood estimation (MLE) of frequency versus measurement time τ_m and total signal duration T , with $N_E = 600$ trajectory realizations. (a) Error of MLE, indicating that longer signals yield more accurate estimates, for a range of optimal measurement times τ_m . (b) Slice of (a) with constant $T = 40\mu\text{s}$, showing the “sweet spot” where the RMS error does not strongly depend on τ_m . (c) Slice of (a) with constant $\tau_m = 0.65\mu\text{s}$, showing the reduction of error with increased collection time.

surement time τ_m , over a span of roughly 1000 oscillations. However, realistic experiments have other characteristic timescales (like energy relaxation and dephasing) that will practically bound the duration T over which a measurement may be taken. Moreover, the frequency Ω itself may exhibit slow drift over longer timescales, which again practically bounds the duration T . We will consider both these cases later in this paper.

Keeping these practical limitations in mind for now, it is advantageous to optimize the estimation to use shorter measurement records. The estimation error will depend on two free timescale parameters, the duration T and the measurement time τ_m , in addition to a fixed timescale, the true Rabi period $2\pi/\Omega_T$. Hence, for a given target estimation error we wish to minimize the duration T by optimizing the measurement time τ_m .

To determine how the estimation error behaves as both T and τ_m are varied for a fixed Ω_T , we numerically simulated the error for a range of parameters. The duration T of each simulation was varied between $T = 1\text{--}50\mu\text{s}$ in increments of $1\mu\text{s}$. For each T , the measurement time τ_m was also varied between $\tau_m = 0.05\text{--}0.8\mu\text{s}$ in increments of $0.05\mu\text{s}$. The root-mean-square (RMS) deviation from

the true frequency Ω_T

$$\Omega_{RMS} = \sqrt{\frac{\sum_i (\Omega_i - \Omega_T)^2}{N_E}}, \quad (16)$$

was computed over an ensemble of $N_E = 600$ realizations for each parameter choice (T, τ_m) to quantify the estimation error. Results for the statistics are shown in Fig. 2(a) as a contour plot. Fig. 2(c) shows the horizontal slice through the dashed line corresponding to the optimal measurement time of $\tau_m = 0.65 \mu\text{s}$, and shows the improvement of the estimation error ($\sim 1/\sqrt{T}$) with an increase in T . Fig. 2(b) shows the vertical slice through the dashed line corresponding to $T = 40 \mu\text{s}$, showing that there is an optimum τ_m that minimizes the error when T is held fixed. The existence of the optimum τ_m is also discussed in [13] in terms of an optimal value for the measurement strength; however, we here investigate the measurement process in a much shorter range, i.e., the total measurement time is less than 50 oscillating cycles.

4. Including Non-idealities

The above discussion assumed that the stochastic evolution from the continuous measurement preserved the purity of the state. However, more realistic evolution must include additional experimental nonidealities that decrease the state purity. These nonidealities include qubit energy-relaxation with a characteristic timescale T_1 , environmental dephasing of the qubit to its energy basis with a characteristic timescale T_2 , and collection loss within the readout chain that leads to a net collection efficiency $\eta \in [0, 1]$. Other unexpected environmental effects over longer timescales that cause frequency drift (such as thermal fluctuations) will be considered later.

Including these effects requires an extended model of the state dynamics (and thus the maximum likelihood method) from pure states to mixed states. A particularly useful representation that accommodates mixed states is the Bloch paravector picture: given an unnormalized density operator ρ , such a paravector has the 4 real state coordinates $\vec{\rho} \equiv (x, y, z, p)$, such that $x = \text{Tr}[\sigma_x \rho]$, $y = \text{Tr}[\sigma_y \rho]$, $z = \text{Tr}[\sigma_z \rho]$, and $p = \text{Tr}[\rho]$. Here the state normalization p physically indicates the probability that the state was prepared. Dividing by this p renormalizes the state $\vec{\rho} \mapsto (x/p, y/p, z/p, 1)$ so that its first three coordinates are the usual Bloch coordinates (expectation values) conditioned on definite (successful) state preparation. By including the normalization of the state explicitly, we can linearize the state evolution due to measurement into a matrix product that generalizes the pure state case of Eq. (5).

Specifically, the ideal case of $\rho \mapsto M_r \rho M_r^\dagger$ with $M_r = U \tilde{E}_r^{1/2}$ can be written equivalently as the 4×4 matrix

product $\vec{\rho} \mapsto \mathbf{M}_r \vec{\rho}$, where $\mathbf{M}_r = \mathbf{V} \mathbf{F}_r$ and

$$\mathbf{V} \equiv \begin{bmatrix} \cos(\Omega \delta t) & 0 & \sin(\Omega \delta t) & 0 \\ 0 & 1 & 0 & 0 \\ -\sin(\Omega \delta t) & 0 & \cos(\Omega \delta t) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad (17)$$

$$\mathbf{F}_r \equiv \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & \cosh(r \delta t / \tau_m) & \sinh(r \delta t / \tau_m) \\ 0 & 0 & \sinh(r \delta t / \tau_m) & \cosh(r \delta t / \tau_m) \end{bmatrix}. \quad (18)$$

Notably, over the duration δt the z -measurement acts as a hyperbolic rotation in the z - p plane, while the Hamiltonian evolution acts as an elliptic rotation in the x - z plane. Interestingly, this behavior is completely analogous to the boosts and rotations of spacetime coordinates in Lorentz transformations [36]. To add nonidealities, it is convenient to re-express these rotations in terms of their infinitesimal generators, which are identical to the Lorentz transformation generators

$$\mathbf{V} \equiv \exp \left(\delta t \begin{bmatrix} 0 & 0 & \Omega & 0 \\ 0 & 0 & 0 & 0 \\ -\Omega & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \right), \quad (19)$$

$$\mathbf{F}_r \equiv \exp \left(\delta t \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & r/\tau_m \\ 0 & 0 & r/\tau_m & 0 \end{bmatrix} \right). \quad (20)$$

Properly, these generators should be summed before the exponentiation over the interval δt , but we keep them separated here for conceptual clarity and simplicity in the maximum likelihood calculation. Note that these expressions make it clear that both Ω and the averaged stochastic result r are assumed constant over the discretization interval δt .

We can now readily add three types of nonideality. First, we include collection inefficiency $\eta \in [0, 1]$, such that the extra dephasing rate of the qubit due to measurement inefficiency is $\gamma_m \equiv (1 - \eta)/(2\eta\tau_m)$. This definition implies that the extra dephasing rate results from the averaged backaction of the signal that was never detected, and the rate vanishes when $\eta = 1$ [33]. Second, we include environmental dephasing with timescale T_2 that arises from sources other than measurement, which contributes an extra term $1/T_2$ to the total qubit dephasing rate. Third, we add energy-relaxation of the qubit at the timescale T_1 . This type of nonideality is more complicated; it preserves the probability p while shifting the population toward the ground state at an exponential rate $1/T_1$. In the z - p coordinates, this has the form, $z \mapsto z e^{-\delta t/T_1} - p(1 - e^{-\delta t/T_1})$. This shift in population also results in an additional dephasing term $(x, y) \mapsto (x, y) e^{-\delta t/2T_1}$. Writing all three of these non-

idealities using their infinitesimal generators, we obtain

$$\mathbf{F}_r \equiv \exp \left(\delta t \begin{bmatrix} -\gamma & 0 & 0 & 0 \\ 0 & -\gamma & 0 & 0 \\ 0 & 0 & -1/T_1 & r/\tau_m - 1/T_1 \\ 0 & 0 & r/\tau_m & 0 \end{bmatrix} \right), \quad (21)$$

in terms of the total dephasing rate

$$\gamma \equiv \gamma_m + \frac{1}{T_2} + \frac{1}{2T_1}, \quad \gamma_m \equiv \frac{1 - \eta}{2\eta\tau_m}. \quad (22)$$

This matrix exponential can easily be written in a closed form (and should be for numerical efficiency), but we omit it here for brevity. It is also worth recalling that this form of the measurement-evolution neglects normalization factors to simplify the linear evolution—the state may always be renormalized after each step during simulation if desired according to $\vec{\rho} \mapsto \mathbf{M}_r \vec{\rho} / (\vec{1} \cdot \mathbf{M}_r \vec{\rho})$, where $\vec{1} \equiv (0, 0, 0, 1)$ extracts the new state norm.

We can then write the log-likelihood function as

$$\mathcal{L}(\Omega) \equiv \ln \left(\vec{1} \cdot \mathbf{M}_N \vec{\rho} \right), \quad (23a)$$

$$\mathbf{M}_N = \mathbf{M}_{r_N} \cdots \mathbf{M}_{r_1}. \quad (23b)$$

The product of 2×2 matrices M_N from before has been simply replaced with the product of 4×4 matrices \mathbf{M}_N , while the pure state has been generalized to a mixed state vector. With this simple change, the rest of the maximum likelihood procedure outlined in the previous sections proceeds unaltered.

B. Fourier Methods

As we have shown, MLE methods can be used to accurately estimate the drive frequency. However, the necessary time to run the algorithm depends crucially on the frequency search range in the MLE algorithm. In order to speed up this search it is beneficial to have at least a rough estimate for the value of Ω_T .

We know the MLE method should saturate the Cramèr-Rao bound, so any other estimation method should perform less well. Nevertheless, other methods may have other desirable qualities, such as estimation speed. A simple and natural way to coarsely-yet-quickly estimate the drive frequency is by studying the power spectral density $S(\Omega)$ of the output signal $r(t)$, which can be quickly estimated from the fast Fourier transform (FFT) of the discretized signal $\{r_j\}$ at time points $\{t_j\}$ with spacing δt as $S(\Omega_j) = |\text{FFT}_{t_j \rightarrow \Omega_j} [\{r_j\}]|^2 \delta t$, where the resulting discrete frequencies $\{f_j \equiv \Omega_j/2\pi\}$ have spacing $\delta f = 1/T = 1/(N\delta t)$ and N is the number of time steps in the original signal $\{r_j\}$. Even a plain FFT method is enough to accurately estimate the frequency for some regime of parameters. However, better estimates can be achieved by simple filtering methods. Fig. 3 shows a simulated power spectral density

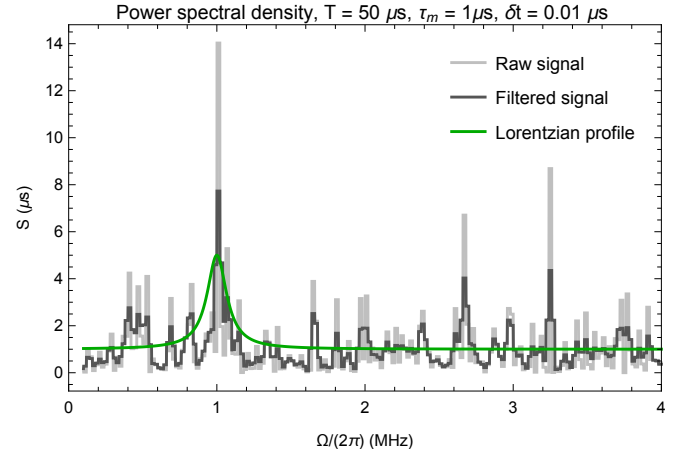


FIG. 3. Power spectral density for $T = 50 \mu\text{s}$, $\tau_m = 1 \mu\text{s}$ and $\delta t = 0.01 \mu\text{s}$, shown for a single realization of the measurement process (gray). After applying a triangular moving-average filter (black) with a width of roughly $T/(2\pi\tau_m)$ points of minimum frequency resolution $\delta\Omega/2\pi = 1/T$, the fluctuations in the power spectral density are reduced. With sufficient nearest-neighbor averaging, the filtered data will approach a Lorentzian profile (green) with increasing T .

of the output signal for $\tau_m = 1 \mu\text{s}$, a total runtime of $T = 50 \mu\text{s}$, and $\delta t = 10 \text{ ns}$, compared to a filtered power spectral density (using a triangular center-weighted moving average over 5 nearest bins of width δf). As the figure illustrates, in this weak measurement regime, the spectral density tends to a Lorentzian function as the total runtime increases [30]. The background noise is concentrated around τ_m , with a peak centered around the drive frequency of height $4\tau_m$. Moreover, the measurement time also characterizes the width of the peak, with a Full-Width-Half-Maximum of $1/(2\pi\tau_m)$. The total runtime T and time step δt determine the resolution and maximum frequency of the Fourier transform method, with $\delta f = 1/T$ and $f_{\text{max}} = 1/\delta t$, respectively. The resolution δf dictates the maximum size of the filtering window that one can apply, given that roughly $1/(2\pi\tau_m\delta f) = T/(2\pi\tau_m)$ points fit inside the peak of the Lorentzian.

In Fig. 4 we show the RMS error of the FFT estimated frequency after filtering for 200 realizations of the measurement, as a function of the measurement time τ_m , for a total runtime $T = 50 \mu\text{s}$. In this parameter regime the simple FFT method is fairly effective at estimating the frequency to within 10% error for τ_m between $0.3 \mu\text{s}$ to $0.8 \mu\text{s}$. Notice, however, that as τ_m decreases the measurement process tends to ‘pin’ the system into one of the eigenstates of the σ_z operator. Hence, in this “Zeno regime” the spectral density shows a peak at zero frequency [30], hindering the FFT frequency estimate and increasing the average error.

Since, as we saw in the previous section, the MLE approach is optimal, we expect that it will outperform the Fourier transform in estimating the drive frequency with




FIG. 4. Frequency estimation via fast Fourier transform. The estimation of the peak frequency can be optimized by filtering the spectral density, as shown in Fig. 3. For a total runtime of $T = 50 \mu s$ and time step $\delta t = 0.01 \mu s$ the estimated frequency approaches the real frequency $\Omega_T/2\pi = 1$ MHz. The RMS error averaged over $N_E = 200$ realizations is of the order of 10% for measurement times $\tau_m \sim 0.3 \mu s$ to $\tau_m \sim 0.8 \mu s$. However, this error increases for shorter measurement times due to zero frequency Zeno pinning, or for longer measurement times, where the measurement is too weak for the output signal to accurately estimate the frequency.

the same duration T . In Fig. 5, we show how the RMS error scales with both runtime T and measurement time τ_m using the FFT method, which should be compared with the MLE scaling in Fig. 2 of the previous section. The best FFT error for short times ranges from 10–20% error, compared with 2–5% error obtained by MLE. The deterioration of performance for small τ_m due to Zeno pinning is also clearly visible. Note that for sake of comparison with MLE we have used a bandpass filter of width 2 MHz centered around $\Omega_T/2\pi = 1$ MHz in Fig. 5 to eliminate false positives at higher frequencies. For much longer durations T , the power spectral density produced by the FFT may be more aggressively window-averaged, so the bandpass filter may be removed.

C. Computational Efficiency

To see when FFT will provide an improvement over MLE in terms of computational efficiency, let us compute a rough estimate of how each method scales with the number of time points $N = T/\delta t$.

For the FFT method, the dominant contribution to the computation is performing the FFT itself, which scales as $N \log N$, with the other operations (squaring, filtering, maximizing) scaling linearly with N . For our improved MLE method, with a grid of n trial frequencies and matrices of dimension d , the dominant contribution scales as $d^3 n N$ due to matrix multiplications within a loop over trial frequencies, with other contributions (trace, maxi-

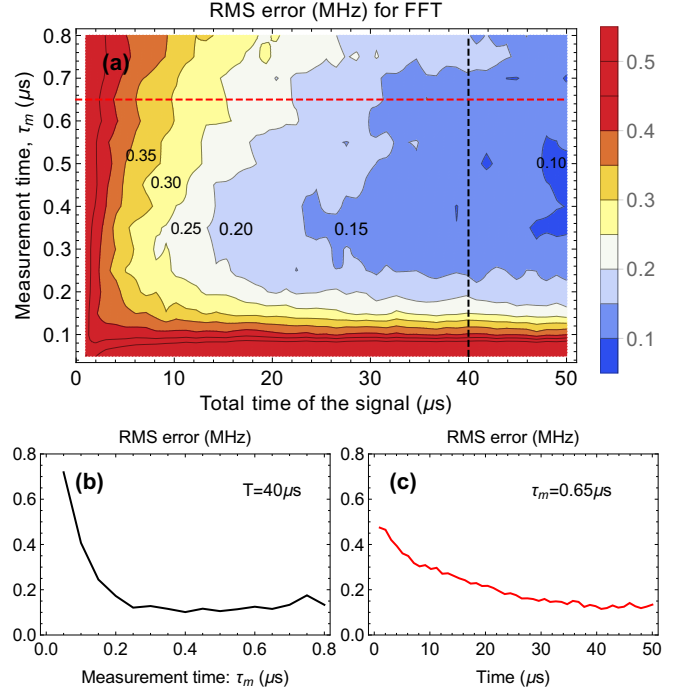


FIG. 5. RMS error of fast Fourier transform (FFT) frequency estimate versus measurement time τ_m and total signal duration T , with $N_E = 600$ trajectory realizations, and a bandpass filter to a window $\Omega \in [0, 2\Omega_T]$, as well as a triangular moving average over nearest spectral points of width $\delta\Omega/2\pi = 1/T$. (a) Error of FFT, showing uniformly worse performance than MLE. (b) Slice of (a) with the same constant $T = 40 \mu s$ as Fig. 2. (c) Slice of (a) with the same constant $\tau_m = 0.65 \mu s$. These crude estimates may be used to accelerate MLE.

mization) scaling at most linearly in N . For our method, d is either 2 or 4, thus giving an overall constant prefactor. The relevant comparison for the dominant scaling is thus $N \log N$ for FFT vs. $n N$ for MLE.

Using an initial long duration T , the computational efficiency of the FFT method may be used to accelerate the MLE estimation. The FFT-estimated frequency can then be used as starting point for the more precise MLE algorithm using shorter runtimes T . We note the unitaries U_j (depending on the trial Ω_j) can be pre-computed as numerical matrices, and simply called from a database when the given trial frequencies appear in order to optimize the algorithm. If one does not have any prior information about the possible value of the drive frequency, then one needs a large grid with $n \sim N$ points, so FFT will be faster than MLE by a crude factor of $\log N/N$. However, after using an initial FFT estimation to reduce the set of trial frequencies, then MLE may use a significantly smaller grid. If the grid is made sufficiently small, such that $n < \log N$, then MLE will be faster than FFT so that it can be appended with overhead scaling subdominantly in N .

As we will explore in detail in the next section, the scaling of MLE makes it particularly well suited for dy-

namical estimation problems. In such a problem, the frequency may be initially estimated via a combined FFT and MLE method, followed by periodic updates of this estimate using a faster MLE estimator scaling only linearly in N .

III. DYNAMIC PARAMETER ESTIMATION

The preceding analysis involves the estimation of a fixed Rabi frequency Ω . We now consider the following (more interesting) problem: what if the Rabi frequency Ω is not fixed, but instead changes in time due to experimental drift caused by changes in the environmental or control system at longer timescales (e.g., thermal variation)? We would then like to update our estimate $\Omega(t)$ continuously in time, which would allow us to monitor and compensate for such a drift *in situ*, in near-real-time. To accomplish this goal, we consider a moving window of a fixed duration T , such that the end-point of this window is the current time t (i.e., a fixed-delay scenario). Estimating the frequency as before within such a moving window then produces a time-evolving maximum-likelihood estimate $\Omega_{ML}(t)$. There is thus a trade-off between the maximum precision of the frequency estimation allowed by the chosen window ($\sigma > \delta f = 1/T$), and the resulting sensitivity to the timescale of a detectable drift $T_d > T$, which should be longer than the temporal averaging resulting from the estimation window.

As discussed in Sec. II C, by using an initial estimation (either FFT or MLE, or a combination) to narrow the frequency search range—assuming the drift is slow compared to the time scale of the estimation—the MLE algorithm scales only linearly in the number of time steps N , making real-time updates to the initial estimation computationally reasonable.

A. Including prior information in the estimate

In contrast to Sec. II A, where we assumed that we had no prior information about the parameter we are estimating, we now assume that we have some prior information from a previous estimation. Such a prior distribution $P_{\text{prior}}(\Omega)$ may be computed either from the RMS error, or from the uncertainty of the previous measurement, plus the typical expected drift uncertainty from the time-varying parameter. In these cases we can improve MLE by incorporating this information.

The maximum, *a posteriori*, probability of the parameter, taking into account the prior is given by Bayes rule, given the measurement record $r(t)$,

$$\begin{aligned} P(\Omega|r(t)) &= \frac{P(r(t)|\Omega)P_{\text{prior}}(\Omega)}{P(r(t))} \\ &\propto P(r(t)|\Omega)P_{\text{prior}}(\Omega). \end{aligned} \quad (24)$$

We may then modify our log-likelihood function as

$$\mathcal{L} \propto \ln P(r(t)|\Omega) + \ln P_{\text{prior}}(\Omega), \quad (25)$$

where the first term is our previous log-likelihood, and the second term takes into account the information from the prior experiments. Note that the denominator in Bayes rule can be dropped given that it is independent of Ω , the quantity over which we are maximizing.

B. Time-dependent frequency tracking

We now illustrate the method of time-dependent parameter tracking of a slowly drifting Rabi frequency. We first generate a sample drifting frequency on a time scale of longer than $40 \mu\text{s}$, irregularly changing by around 40% of the value of the Rabi frequency in total. We then generate the measurement results, sampling from distributions using that slowly changing single qubit Hamiltonian. This provides a single realization, simulating a typical experiment of this type, where we have included the nonidealities discussed in Sec. II A 4 for realistic values of $T_1 = 50 \mu\text{s}$ and $T_2 = 30 \mu\text{s}$, and efficiency of $\eta = 0.5$.

From this data set, we then apply the two discussed estimation strategies, FFT and MLE, with a moving window duration of $T = 40 \mu\text{s}$, matched to the smallest time scale we wish to resolve. We choose the measurement time τ_m in order to optimize the statistical uncertainty for this time window choice. From Fig. 2, we make the choice $\tau_m = 0.65 \mu\text{s}$. As shown in Fig. 6, the methods are able to track this drifting frequency, to the expected accuracies discussed in the previous sections.

We have run the MLE methods for both the ideal model, and for the non-ideal model, incorporating the T_1 , T_2 , and η effects discussed in Sec. II A 4. The ideal simulation together with the MLE estimation method (red solid circles) and FFT (brown crosses) are shown in Fig. 6(a), while the simulation incorporating non-idealities mentioned above is shown in Fig. 6(b). In the later subfigure, two different MLE models are used, the first (blue open circles) assume completely ideal dynamics, while the second (red solid circles) incorporates the non-idealities into the estimation model. Brown crosses are again using the FFT method. As expected, the MLE methods performs significantly better than the FFT method. Generally, both MLE methods show good tracking fidelity, indicating there is not much difference in the estimation precision, despite the complexity increase in the non-ideal model. We explain this insensitivity as a combination of fast Rabi frequency compared to the relaxation rate $1/T_1$, with the rapid purification of the measurement compensating for the dephasing rate $1/T_2$ and inefficiency. The quality of the tracking fidelity varies from run to run, and can always be improved by increasing the duration of the time window.

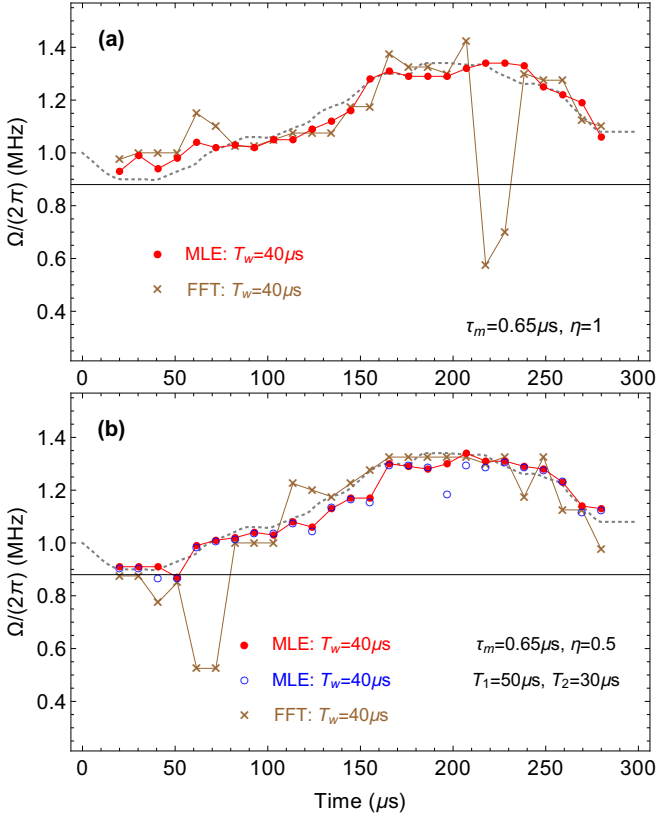


FIG. 6. The time-dependent frequency tracking, shown in two examples: (a) an ideal case with unit efficiency $\eta = 1$ and no extra dephasing, (b) a realistic case with 50% efficiency ($\eta = 0.5$), a phase relaxation time $T_2 = 30 \mu s$ and an energy relaxation time $T_1 = 50 \mu s$. The true drifting frequency in both plots is shown in dotted gray lines. In the top panel (a), brown crosses represent estimation from FFT and red solid circles represent estimation from MLE. In the bottom panel (b), brown crosses, red solid circles, and blue open circles represent estimation from FFT, MLE, and MLE ignoring non-idealities in the simulation, respectively. The moving window duration is $T_w = 40 \mu s$ for all estimation methods, stepped in $10 \mu s$ intervals. The estimated frequencies are plotted at the mid-point of each estimation window.

IV. CONCLUSION

We conclude that it is possible to track a drifting parameter with continuous monitoring, and have specified a method that improves the computational overhead to carry out the estimation analysis. There is a trade off between duration of the data window and of estimation accuracy. Given a target precision of the estimate, we have shown that this sets the temporal resolution on the drifting parameter, and illustrated that it can work by simulating measurement results. If the parameter is changing in an irregular fashion, a major conclusion is that weak continuous monitoring can faithfully track the parameter. We also showed that the method is computationally efficient for such tracking applications, scaling only lin-

early with the number of time points used in the estimation once an initial frequency range has been identified. This capability opens the door for real-time parameter tracking combined with adaptive feedback [6] for parameter stabilization.

The question of how to extend these methods to multiple quantum systems, and estimate parameters such as an interaction energy is an important open question that will be pursued in subsequent work.

Note added: After the public presentation of these results at the APS March meeting [37], but before the posting of our preprint, an independent work covering some of the same physics was posted on the arXiv, authored by Kiilerich and Mølmer [38].

ACKNOWLEDGMENTS

We thank Irfan Siddiqi, Alexander Korotkov and Shengshi Pang for helpful discussions. This work was supported by US Army Research Office Grants No. W911NF-15-1-0496, No. W911NF-13-1-0402, by National Science Foundation grant DMR- 1506081, by John Templeton Foundation grant ID 58558, by Development and Promotion of Science and Technology Talents Project Thailand, and by the National Council of Science and Technology (CONACyT) Mexico. We also acknowledge partial support by Perimeter Institute for Theoretical Physics. Research at Perimeter Institute is supported by the Government of Canada through Industry Canada and by the Province of Ontario through the Ministry of Economic Development & Innovation.

* L.C. and A.C. contributed equally to this work.

V. APPENDIX: MAXIMUM LIKELIHOOD METHOD FOR PERIODIC PROJECTIVE MEASUREMENTS

In this appendix we review how to find the maximum likelihood estimate for the qubit oscillation frequency, using only unitary evolution and periodic projective measurements. Starting in basis state $|0\rangle$, the unitary (2) followed by projective measurement will yield either result 0, or 1, so the system has probability $p_s = \cos(\Omega\tau/2)^2$ to be projected back into state $|0\rangle$ (given result 0), and probability $p_d = \sin(\Omega\tau/2)^2$ to be projected into state $|1\rangle$ (given result 1). Similarly, starting in the basis state $|1\rangle$, following the unitary operation and measurement, the system has probability $p_s = \cos(\Omega\tau/2)^2$ to be projected back into state $|1\rangle$ (given result 1), and probability $p_d = \sin(\Omega\tau/2)^2$ to be projected into state $|0\rangle$ (given result 0). Thus the system has probability p_s of staying the same as the previous step, and probability p_d of being different from the previous step after the measurement.

We now are given a sequence of 0s and 1s resulting from a lengthy number of measurements. Suppose we also know the initial state of the system. How can we

extract the value of Ω if we do not know it in advance? We see from the discussion above that the probability of a sequence of results can be determined by the number of switches from the same result to a different result, n . Given N measurements, the total probability is

$$P(n, N|\Omega) = \binom{N}{n} (\sin^2 \Omega\tau/2)^n (\cos^2 \Omega\tau/2)^{N-n}, \quad (26)$$

where the prefactor normalizes the distribution. This is just a binomial probability distribution with probability p_d of switching and probability $p_s = 1 - p_d$ of staying the same. In this simple example the Fisher information about the parameter Ω may be calculated straightforwardly,

$$\mathcal{I} = \sum_{n=0}^N P(n, N|\Omega) [\partial_\Omega \ln P(n, N|\Omega)]^2 = N\tau^2, \quad (27)$$

so the standard deviation is bounded by the Cramér-Rao bound, $\sigma_\Omega \geq 1/(\tau\sqrt{N})$.

We may now use the maximum likelihood method (MLE) to find an estimate of Ω and the uncertainty in the estimate. We expand the log-likelihood as a function of Ω as

$$\begin{aligned} \ln P \sim \text{const} + \frac{\partial}{\partial \Omega} \ln P(n, N|\Omega_{ML})(\Omega - \Omega_{ML}) \\ + (1/2) \frac{\partial^2}{\partial \Omega^2} \ln P(n, N|\Omega_{ML})(\Omega - \Omega_{ML})^2 + \dots, \end{aligned} \quad (28)$$

and define Ω_{ML} as the frequency that maximizes it. In turn, the second derivative evaluated at $\Omega = \Omega_{ML}$ gives the (negative) inverse uncertainty variance in the estimate. This yields

$$\tau \Omega_{ML} = 2 \arcsin \sqrt{\frac{n}{N}}, \quad (29)$$

with a standard deviation uncertainty of $\sigma_\Omega = 1/(\tau\sqrt{N})$, indicating that the maximum likelihood method saturates the Cramér-Rao bound. We note that although the uncertainty is independent of n , the estimate has a divergent slope at $n = 0$ and $n = N$, with a minimum slope at $n = N/2$, indicating that the estimate is least sensitive to fluctuations in n at $\tau\Omega = \pi/2$. We have checked this method by simulating the sequence of outputs from such a measurement by first fixing a known value Ω , running the simulation and estimating the parameter from the simulation data. We then compare the estimate to the known value. This method works robustly for different values of n and N , and gives estimated values consistent with the expected uncertainty.

-
- [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] H. Wiseman and G. Milburn, *Quantum Measurement and Control* (Cambridge University Press, 2010).
 - [4] K. Jacobs, *Quantum Measurement Theory and its Applications* (Cambridge University Press, 2014).
 - [5] A. Barchielli and M. Gregoratti, *Quantum trajectories and measurements in continuous time* (Springer-Verlag Berlin Heidelberg, 2009).
 - [6] R. Vijay, C. Macklin, D. H. Slichter, S. J. Weber, K. W. Murch, R. Naik, A. N. Korotkov, and I. Siddiqi, *Nature* **490**, 77 (2012).
 - [7] W. Cui and F. Nori, *Phys. Rev. A* **88**, 063823 (2013).
 - [8] B. M. Escher, R. L. de Matos Filho, and L. Davidovich, *Nat. Phys.* **7**, 406 (2011).
 - [9] R. Demkowicz-Dobrzański, J. Kołodyński, and M. Guţă, *Nat. Commun.* **3**, 1063 (2012).
 - [10] M. Tsang, *New J. Phys.* **15**, 073005 (2013).
 - [11] J. Kołodyński and R. Demkowicz-Dobrzański, *New J. Phys.* **15**, 073043 (2013).
 - [12] S. Alipour, M. Mehboudi, and A. T. Rezakhani, *Phys. Rev. Lett.* **112**, 120405 (2014).
 - [13] J. F. Ralph, K. Jacobs, and C. D. Hill, *Physical Review A* **84**, 052119 (2011).
 - [14] S. Gammelmark and K. Mølmer, *Physical Review A* **87**, 032115 (2013).
 - [15] S. Gammelmark and K. Mølmer, *Physical review letters* **112**, 170401 (2014).
 - [16] A. H. Kiilerich and K. Mølmer, *Physical Review A* **92**, 032124 (2015).
 - [17] M. Tsang, *Phys. Rev. Lett.* **102**, 250403 (2009).
 - [18] M. Tsang, *Phys. Rev. A* **80**, 033840 (2009).
 - [19] M. Tsang, *Phys. Rev. A* **81**, 013824 (2010).
 - [20] 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).
 - [21] H. Yonezawa, D. Nakane, T. A. Wheatley, K. Iwasawa, S. Takeda, H. Arao, K. Ohki, K. Tsumura, D. W. Berry, T. C. Ralph, H. M. Wiseman, E. H. Huntington, and A. Furusawa, *Science* **337**, 1514 (2012).
 - [22] K. Iwasawa, K. Makino, H. Yonezawa, M. Tsang, A. Davidovic, E. Huntington, and A. Furusawa, *Phys. Rev. Lett.* **111**, 163602 (2013).
 - [23] M. Tsang, *Phys. Rev. A* **92**, 062119 (2015).
 - [24] G. A. Smith, A. Silberfarb, I. H. Deutsch, and P. S. Jessen, *Physical review letters* **97**, 180403 (2006).
 - [25] A. Smith, C. A. Riofrío, B. E. Anderson, H. Sosa-Martinez, I. H. Deutsch, and P. S. Jessen, *Physical Review A* **87**, 030102 (2013).
 - [26] P. Six, P. Campagne-Ibarcq, I. Dotsenko, A. Sarlette, B. Huard, and P. Rouchon, *Physical Review A* **93**,

- 012109 (2016).
- [27] K. W. Murch, S. J. Weber, C. Macklin, and I. Siddiqi, *Nature* **502**, 211 (2013).
 - [28] S. J. Weber, A. Chantasri, J. Dressel, A. N. Jordan, K. W. Murch, and I. Siddiqi, *Nature* **511**, 570 (2014).
 - [29] M. D. Shulman, S. P. Harvey, J. M. Nichol, S. D. Bartlett, A. C. Doherty, V. Umansky, and A. Yacoby, *Nature communications* **5**, 5156 (2014).
 - [30] A. N. Korotkov, *Phys. Rev. B* **63**, 085312 (2001).
 - [31] R. Vijay, D. H. Slichter, and I. Siddiqi, *Phys. Rev. Lett.* **106**, 110502 (2011).
 - [32] J. Gambetta, A. Blais, M. Boissonneault, A. A. Houck, D. I. Schuster, and S. M. Girvin, *Phys. Rev. A* **77**, 012112 (2008).
 - [33] A. N. Korotkov, in *Quantum machines, Lecture notes of the Les Houches Summer School (Session 96, July 2011)*, edited by M. Devoret et al. (Oxford University Press, Oxford, 2014) Chap. 17, p. 533, arXiv:1111.4016.
 - [34] A. N. Jordan and A. N. Korotkov, *Contemporary Physics* **51**, 125 (2010).
 - [35] M. A. Nielsen and I. L. Chuang, *Quantum computation and quantum information* (Cambridge university press, 2010).
 - [36] A. N. Jordan and A. N. Korotkov, *Physical Review B* **74**, 085307 (2006).
 - [37] L. Cortez-Gonzalez and A. N. Jordan, in *APS Meeting Abstracts, E48.013* (2016).
 - [38] A. H. Kiilerich and K. Mølmer, arXiv preprint arXiv:1605.00902 (2016); A. H. Kiilerich and K. Mølmer, *Phys. Rev. A* **94**, 032103 (2016).