Single-shot quantum state estimation via a continuous measurement in the strong backaction regime

Robert L. Cook,[1,*] Carlos A. Riofrío,[2,1] and Ivan H. Deutsch[1]

[1]*Center for Quantum Information and Control, University of New Mexico, Albuquerque, New Mexico 87131, USA*
[2]*Dahlem Center for Complex Quantum Systems, Freie Universität Berlin, 14195 Berlin, Germany*

We study quantum tomography based on a stochastic continuous-time measurement record obtained from a probe field collectively interacting with an ensemble of identically prepared systems. In comparison to previous studies, we consider here the case in which the measurement-induced backaction has a non-negligible effect on the dynamical evolution of the ensemble. We formulate a maximum likelihood estimate for the initial quantum state given only a single instance of the continuous diffusive measurement record. We apply our estimator to the simplest problem: state tomography of a single pure qubit, which, during the course of the measurement, is also subjected to dynamical control. We identify a regime where the many-body system is well approximated at all times by a separable pure spin coherent state, whose Bloch vector undergoes a conditional stochastic evolution. We simulate the results of our estimator and show that we can achieve close to the upper bound of fidelity set by the optimal generalized measurement. This estimate is compared to, and significantly outperforms, an equivalent estimator that ignores measurement backaction.

PACS number(s): 03.65.Wj, 42.50.Dv, 05.10.Gg

I. INTRODUCTION

A fundamental task in quantum information processing is the ability both to reliably prepare an arbitrary quantum state and to experimentally verify its production. Traditional quantum state tomography (QST) relies on an exhaustive procedure where the target state is repeatedly prepared and then destructively measured in an informationally complete number of measurement settings. Such a procedure is often extremely time intensive, requiring a significant amount of both data and postprocessing time [1,2].

These inefficiencies can be significantly reduced when one can perform a weak continuous measurement, acting collectively on an identically prepared ensemble, in conjunction with a well-chosen dynamical control [3,4]. In particular, consider an ensemble of N systems prepared in an identical tensor product state $\rho_{\text{tot}} = \rho_0^{\otimes N}$, experiencing a known time-dependent control Hamiltonian while simultaneously coupled to a traveling wave probe. If the control drives the system such that a continuous measurement of the probe is informationally complete, then one can use this measurement record to obtain a high-fidelity estimate of the initial state of the system, ρ_0.

This protocol has been implemented in experiments [5,6] with quantum states encoded in the hyperfine spins of an ensemble of laser-cooled cesium atoms controlled with magneto-optical fields [6–8] and measured with polarization spectroscopy [9]. By applying an appropriate estimator to the measurement record, one can obtain high-fidelity reconstructions of arbitrary states in the 16-dimensional hyperfine ground-state manifold of cesium. However, these experiments were performed under far from idealized conditions. The reconstructions were ultimately limited by systematic errors and decoherence caused by spontaneous emission. While detrimental to the final fidelity, these limitations simplified the analysis, as the collective effects of quantum backaction were

completely negligible. Because of these facts, any fundamental limits of continuous-measurement-based QST have yet to be addressed.

Here, we extend this protocol to an idealized regime, free of technical imperfections and decoherence, where any limitations are solely due to the quantum backaction induced by the measurement itself and thus fundamental to the tomographic protocol. Recent works have studied how such continuous measurements in the presence of back action can be used for parameter estimation [10,11], hypothesis testing [12,13], and for the closely related problem of retrodicting a past system estimate given future measurements [14]. Such protocols have recently been implemented in continuous measurement of a superconducting qubit [15], a platform where quantum trajectories have been observed and predicted to give optimal routes between quantum states [16].

In the context of continuous collective measurement on an ensemble, the primary effects of measurement backaction are to introduce correlations between the atoms, i.e., spin squeezing [17], as well as to perturb the mean spin in a random and nonlinear way. Both effects greatly increase the complexity of the quantum state tomography, as the former necessitates a many-body description and the latter prevents the use of many standard tomographic techniques, e.g., convex optimization. This work addresses these issues by deriving a general likelihood function for a continuous-time diffusive measurement of a collective spin projection and derives an efficiently computable approximation in the case of pure qubits. We then use this function to numerically compute a maximum likelihood estimate (MLE) to reconstruct the initial state. We compare our results to the well-known bounds for the average fidelity [18], which is achieved by the optimal collective positive operator valued measure (POVM) [19].

The remainder of this paper is structured as follows. We first establish a general mathematical model for a continuous-time, collective-spin measurement via polarization spectroscopy, with particular emphasis on the conditions under which quantum backaction cannot be neglected. We then derive a maximum likelihood (ML) estimator for the initial state of

---

*rlcook.unm@gmail.com

the ensemble given a diffusive continuous-time measurement. From the general expression, we specialize to estimating the initial state of a pure qubit given an ensemble of identical copies. We derive an efficiently computable approximation to the exact expression under the condition that the entangling effects of the measurement backaction are negligible, while stochastic kicks to the Bloch vector induced by the measurement remain important. We then numerically test the performance of the approximate MLE for a moderate number of qubits and compare the results both to the optimal POVM for quantum tomography and an estimator that completely ignores the effect of measurement backaction. We conclude with a summary and outlook for future studies.

II. SPIN ESTIMATION THROUGH POLARIZATION SPECTROCOPY

We consider cold atomic spins measured via polarization spectroscopy as our model platform in which to examine the fundamental limits of QST based on continuous measurement and control [20]; a schematic is shown in Fig. 1. The measurement is made via the Faraday interaction, whereby the linear polarization of an off-resonant probe laser rotates in proportion to the collective magnetization of the atomic ensemble along the direction of propagation of the probe. For a system composed of N atoms identically coupled to the probe field, a measurement of this rotation results in a quantum nondemolition (QND) measurement of the collective angular momentum operator, $J_z = \sum_{i=1}^{N} j_z^{(i)}$, where $j_z^{(i)}$ is the z-axis projection of the ith atomic spin operator. This measurement
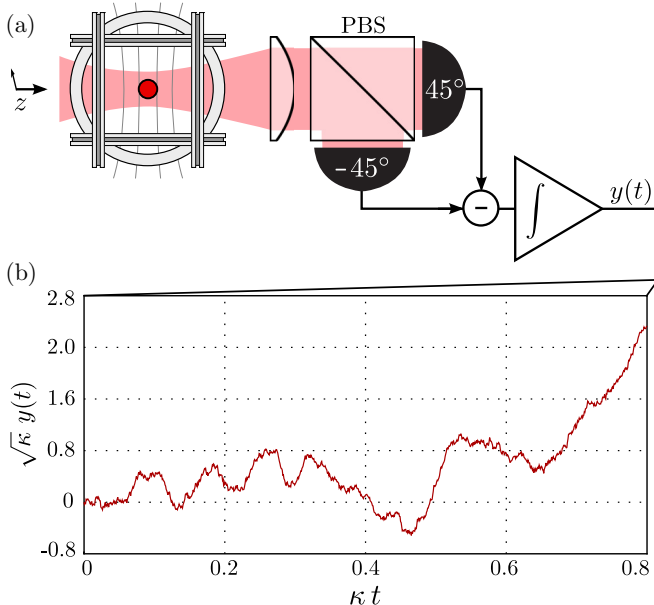

FIG. 1. (Color online) Schematic and sample measurement. (a) An atomic ensemble is probed by an off-resonant, linearly polarized laser, while simultaneously being subjected to external RF magnetic control fields. The outgoing laser is measured by a balanced polarimeter, whose integrated current generates a noisy measurement record $y(t)$. (b) A typical simulated measurement record for $N = 50$ symmetrically coupled qubits initialized in a spin coherent state along the x axis.

occurs at a rate κ, which is set by the input photon flux times the rate at which a single atom will scatter an incident photon into the orthogonal polarization mode.

A balanced polarimeter measuring in a basis $45°$ to the input polarization implements an effective homodyne measurement, where the probe field acts as the local oscillator [21]. In order to isolate the fundamental measurement statistics from any technical imperfections, we assume that the polarimeter is shot noise limited, i.e., introduces no additional noise, and is perfectly linear with unit quantum efficiency.

In the absence of any dynamical control, the measurement outcomes of such a polarimeter averaged over a time T, x, occur with a probability $P(\mathsf{x}) = \mathrm{Tr}(E_\mathsf{x}\rho_{\mathrm{tot}})$, where

$$E_\mathsf{x} = \sqrt{\frac{\kappa T}{2\pi}} e^{-\frac{\kappa T}{2}(\mathsf{x}-J_z)^2} \qquad (1)$$

are the elements of a POVM [20,22,23]. As $T \to \infty$ this becomes a projective collective measurement of the observable J_z. At short times this describes a weak collective measurement of J_z with finite resolution set by the shot-noise variance of the measurement, $1/\kappa T$.

When the output homodyne current is monitored continuously in time, the integrated current describes a stochastic process, $\{y(t) : 0 \leqslant t \leqslant T\}$, where T is the fixed final time. For a system prepared in the definite initial condition $\rho_{\mathrm{tot}}(0)$, allowing for dynamical control and measurement backaction, this can be be written as [24]

$$y(t) = \sqrt{\kappa} \int_0^t \mathrm{Tr}[J_z \rho_{\mathrm{tot}}(s)] \, ds + w(t), \qquad (2)$$

where $\{w(t) : t \geqslant 0\}$ is a realization of the Wiener process and models the time integral over the uncorrelated shot noise introduced by the quantum limited measurements made at every time t [see Fig. 1(b)]. The goal of QST is to estimate $\rho_{\mathrm{tot}}(0)$ given an observation of $\{y(t) : 0 \leqslant t \leqslant T\}$ and a particular model for its statistics.

In a single run of the experiment, the evolution of $\rho_{\mathrm{tot}}(t)$ conditioned on this measurement record is governed by the well-known diffusive stochastic master equation (SME) [22,23] ($\hbar = 1$),

$$d\rho_{\mathrm{tot}}(t) = -i[H_c(t),\rho_{\mathrm{tot}}(t)]\, dt + \gamma_{\mathrm{diss}} \mathcal{D}[\rho_{\mathrm{tot}}(t)]\, dt$$
$$+ \frac{\kappa}{4} \mathcal{L}[\rho_{\mathrm{tot}}(t)]\, dt + \frac{\sqrt{\kappa}}{2} \mathcal{H}[\rho_{\mathrm{tot}}(t)]\, dv(t), \quad (3)$$
$$dv(t) \equiv dy(t) - \sqrt{\kappa}\, \mathrm{Tr}[\rho_{\mathrm{tot}}(t)J_z]\, dt,$$

where $H_c(t)$ is the externally applied control Hamiltonian and we have defined the maps

$$\mathcal{L}[\rho_{\mathrm{tot}}] \equiv J_z\, \rho_{\mathrm{tot}}\, J_z - \tfrac{1}{2} J_z^2\, \rho_{\mathrm{tot}} - \tfrac{1}{2} \rho_{\mathrm{tot}} J_z^2,$$
$$\mathcal{H}[\rho_{\mathrm{tot}}] \equiv J_z\, \rho_{\mathrm{tot}} + \rho_{\mathrm{tot}}\, J_z - 2\, \mathrm{Tr}[\rho_{\mathrm{tot}}(t)J_z]\rho_{\mathrm{tot}}. \qquad (4)$$

We have also included an additional general channel $\mathcal{D}[\rho_{\mathrm{tot}}]$ that accounts for any additional sources of decoherence occurring at a characteristic rate γ_{diss}. The stochastic process defined by the differential $dv(t)$ is known as the innovation process and determines the strength of the measurement backaction in a given interval. When the measurement record is consistent with the true state of the system, the innovation is a Wiener process, $dv(t) = dw(t)$, as follows from inverting

Eq. (2). However, the task of QST is to estimate an *unknown* quantum state given some data. To derive such an estimate, one must evolve a conditional state from an initial condition $\rho'_{\text{tot}}(0)$ not equal to the initial condition used to generate the data. Written in terms of the innovation, Eq. (3) is still valid, but in general, we cannot assume that the innovation is a Wiener process.

Previous experiments on QST via continuous measurement [5,6] operated in a regime where the control Hamiltonian $H_c(t)$ and decoherence rate γ_{diss} were much larger in magnitude than the measurement terms proportional to κ. The measurement duration, T, was chosen such that $\gamma_{\text{diss}}T < 1$. Thus, since $\kappa T \ll 1$, the stochastic measurement outcomes in the measurement record are completely dominated by the shot noise in the probe rather than the "projection noise" uncertainty of the state, $(\Delta J_z)^2_{\text{PN}} = \text{Tr}[\rho_{\text{tot}}(0)J_z^2] - \text{Tr}[\rho_{\text{tot}}(0)J_z]^2$. In that case, measurement backaction is negligible over the duration of the measurement and the system will remain unentangled. The expected value of the collective spin is then well approximated as $\text{Tr}[J_z\rho_{\text{tot}}(t)] \approx N\,\text{Tr}[j_z\tilde{\rho}(t)]$, where $\tilde{\rho}$ is the single-particle density operator that evolves solely under an unconditional master equation

$$\frac{d}{dt}\tilde{\rho}(t) = -i[h_c(t),\tilde{\rho}(t)] + \gamma_{\text{diss}}\mathcal{D}^{(1)}[\tilde{\rho}(t)], \tag{5}$$

where $h_c(t)$ is the single-atom control Hamiltonian, and $\mathcal{D}^{(1)}$ is the single-atom decoherence map. In this backaction-free approximation, the tomographic estimate for the initial state $\tilde{\rho}(0)$ reduces to a standard problem of constrained ML [4].

Here we consider the opposite regime, where $\gamma_{\text{diss}} = 0$ and κT is not necessarily small. This presents a formidable challenge due to the nonlinear nature of $\mathcal{H}[\rho_{\text{tot}}]$, as well as the fact that the future values of $y(t)$ depend on its past through the conditional nature of ρ_{tot}. For simplicity, we restrict our attention here to the case of pure-state, nondissipative dynamics. When $\gamma_{\text{diss}} = 0$ and assuming perfect measurement (i.e., unit quantum efficiency), the evolution of an initial pure state will remain pure. It is then sufficient to propagate a collective state vector, $|\Psi(t)\rangle$, which evolves according to a conditional Schrödinger equation (CSE) [22],

$$d|\Psi(t)\rangle = \left[-iH_c(t) - \tfrac{1}{8}\kappa(J_z - \langle J_z\rangle_{\Psi(t)})^2\right]|\Psi(t)\rangle\,dt$$
$$+ \tfrac{1}{2}\sqrt{\kappa}\,(J_z - \langle J_z\rangle_{\Psi(t)})\,|\Psi(t)\rangle\,dv(t), \tag{6}$$
$$dv(t) = dy(t) - \sqrt{\kappa}\langle J_z\rangle_{\Psi(t)}\,dt,$$

where $\langle J_z\rangle_{\Psi(t)} = \langle\Psi(t)|J_z|\Psi(t)\rangle$. Our goal is to deduce the initial state $|\psi(0)\rangle$ of one member of an identical ensemble, given an initial product state, $|\Psi(0)\rangle = |\psi(0)\rangle^{\otimes N}$, and a continuous measurement record of the form of Eq. (2), when the collective state evolves according to Eq. (6).

III. THE LIKELIHOOD FUNCTION

Estimating an initial quantum state from an observed measurement record is fundamentally a problem of statistical inference. Here we utilize an MLE given the measurement record $\{y(t) : 0 \leqslant t \leqslant T\}$, obtained over time from a collective measurement on a single ensemble. Our derivation hinges on the known form of the measurement record given in Eq. (2) and the fact that the quantum trajectory is consistent with our

model of homodyne detection. Given this, we are able to apply well-developed classical methods for analyzing continuous-time stochastic processes. Here we include a pedagogical derivation of the necessary expression, Eq. (12), in such a way as to minimize any references to measure theoretic probability theory. However, this expression can be independently derived in the fully rigorous setting [25, Theorem 7.19].

We begin by considering a general stochastic process, $\{x(t) : 0 \leqslant t \leqslant T\}$, defined by the integral

$$x(t) = \int_0^t m(\theta_0,s,x(s))\,ds + w(t), \tag{7}$$

where $m(\theta_0,t,x(t))$ is the instantaneous mean signal, which is assumed to be a time-dependent functional of the history of $\{x(t)\}$, and θ_0 is a vector of unknown parameters in the model. By assuming that $m(\theta_0,t,x(t))$ can only depend on $\{x(s) : 0 \leqslant s < t\}$, we are able to analyze the statistics of $\{x(t) : 0 \leqslant t \leqslant T\}$ by first considering the statistics of the Wiener process and then making a simple change of variables. The defining properties of the Wiener process are that (i) it has a continuous trajectory starting from 0 and (ii) its increments are independent, mean 0, Gaussian-distributed random variables, whose variance is equal to the increment's time duration. These criteria imply that the joint probability density for a sampling from the Wiener process is simply a product of nested Gaussians and therefore so will be the joint density for an equivalent sampling from $\{x(t)\}$.

To initially avoid some subtleties inherent to continuous-time stochastic processes, we first consider $\{x(t)\}$ sampled at a countably dense set of n times $\{t_i \in [0,\infty) : 0 = t_0 < t_1 < \cdots < t_n = T\}$ and then examine the continuous limit. To ease the notation, we define the quantities $m_i^{\theta_0} \equiv m(\theta_0,t_i,x(t_i))$, $\Delta t_i \equiv t_i - t_{i-1}$, and $\Delta x_i \equiv x(t_i) - x(t_{i-1})$. We obtain the continuous limit as $\Delta t_i \to 0$ or, equivalently, $n \to \infty$. For simplicity, we also assume that n is large enough that the approximation $\int_{t_{i-1}}^{t_i} m(\theta_0,s,x(s))\,ds \approx m_{i-1}^{\theta_0}\Delta t_i$ is valid. The joint probability that each sample x_i will be found in a corresponding interval $[a_i,b_i]$ is well approximated by the integrals

$$\mathbb{P}(\{x_i \in [a_i,b_i]\}) \approx \int_{a_1}^{b_1} dx_1 \ldots \int_{a_n}^{b_n} dx_n$$
$$\times \prod_{i=1}^n \frac{\exp\left[-\tfrac{1}{2\Delta t_i}\left(\Delta x_i - m_{i-1}^{\theta_0}\Delta t_i\right)^2\right]}{\sqrt{2\pi\,\Delta t_i}}. \tag{8}$$

A natural way to perform ML estimation would be to consider the integrand in Eq. (8) as the likelihood $\mathcal{L}_n(\theta)$, i.e., a function of the input parameter vector θ, given the sampled observation of $\{x(t) : 0 \leqslant t \leqslant T\}$ as determined by the unknown parameters θ_0:

$$\mathcal{L}_n(\theta) \equiv \prod_{i=1}^n \frac{\exp\left[-\tfrac{1}{2\Delta t_i}\left(\Delta x_i - m_{i-1}^{\theta}\Delta t_i\right)^2\right]}{\sqrt{2\pi\,\Delta t_i}}. \tag{9}$$

However, this likelihood fails to be of use in the continuous-time limit because as $\Delta t_i \to 0$, the short-time statistics are dominated by shot noise and are ultimately independent of θ. This can be seen by substituting the definition of Δx_i via

Eq. (7), which results in

$$\mathcal{L}_n(\theta) = \prod_{i=1}^{n} \frac{\exp\left\{-\frac{1}{2}\left[\frac{\Delta w_i}{\sqrt{\Delta t_i}} - \left(m_{i-1}^{\theta} - m_{i-1}^{\theta_0}\right)\sqrt{\Delta t_i}\right]^2\right\}}{\sqrt{2\pi\,\Delta t_i}}.$$

(10)

For any noise realization and $\Delta t_i > 0$, the random variables $\xi_i \equiv \Delta w_i / \sqrt{\Delta t_i}$ are mean 0 Gaussian random variables with *unit* variance. Therefore, at every time index, an estimator maximizing Eq. (9) would minimize the squared deviation of a number $\xi_i \sim O(1)$, from an expression proportional to $\sqrt{\Delta t_i}$. In the limit $\Delta t_i \to 0$, this is independent of θ and depends solely on the unwanted shot noise. While we could reduce this effect by coarse graining the measurement record over longer time intervals, such a procedure would also necessarily coarse grain over the time dependence in m, possibly resulting in a loss of information about θ_0.

Fortunately, we can make use of the full continuous measurement record by instead considering a *likelihood ratio* between a candidate parameter θ_1 and a reference parameter θ_2. Upon doing so the divergences represented by ξ_i cancel, leaving a useful expression in the continuous-time limit. Computing this ratio and simplifying gives

$$\frac{\mathcal{L}_n(\theta_1)}{\mathcal{L}_n(\theta_2)} = \exp\left\{\sum_{i=1}^{n}\left[m_{i-1}^{\theta_1} - m_{i-1}^{\theta_2}\right]\Delta x_i\right.$$
$$\left. - \frac{1}{2}\sum_{i=1}^{n}\left[\left(m_{i-1}^{\theta_1}\right)^2 - \left(m_{i-1}^{\theta_2}\right)^2\right]\Delta t_i\right\}. \quad (11)$$

The limit $\Delta t_i \to 0$ of this expression exists, is meaningful, and results in the exponentiated Itō integral,

$$\Lambda(\theta_1,\theta_2) \equiv \exp\left\{\int_0^T \left[m(\theta_1,t,x(t)) - m(\theta_2,t,x(t))\right]dx(t)\right.$$
$$\left. - \frac{1}{2}\int_0^T \left[m(\theta_1,t,x(t))^2 - m(\theta_2,t,x(t))^2\right]dt\right\}. \quad (12)$$

While this expression removes the short-time singularities present in the Wiener process, we now must also consider how the expected mean evolves not only given the "hypothesis" θ_1, but also given θ_2.

While it is most common to define a likelihood function as the probability density with respect to a uniform measure, e.g., dx, this is actually not strictly necessary. What is necessary is that, given a family of probability measures, $\{\mathbb{P}_\theta : \theta \in \Theta \subset \mathbb{R}^d\}$, there must exist a density that can transform the probability measure $\mathbb{P}_{\theta_2}$ into $\mathbb{P}_{\theta_1}$ without ever "dividing by 0"; in other words, they all must assign 0 probability to the same events and so are, in some sense, equivalent [26]. While we refer to Eq. (12) as a likelihood ratio, it is in fact a valid density for transforming the measure $\mathbb{P}_{\theta_2}$, which gives the proper probabilities for a diffusive trajectory with a drift parameterized by θ_2, into the measure $\mathbb{P}_{\theta_1}$. Thus, for every event A, $\mathbb{P}_{\theta_1}(A) = \int_A d\mathbb{P}_{\theta_1} = \int_A \Lambda(\theta_1,\theta_2)\,d\mathbb{P}_{\theta_2}$ [25, Theorem 7.19]. It is in this sense that, given a particular realization of $\{x(t)\}$, the maximum of Eq. (12) with respect to θ_1 truly is the MLE for θ.

To turn this general expression into the form we ultimately use, we first note that, given a measurement record, $\{y(t) : 0 \leqslant t \leqslant T\}$, and a valid initial condition, the expectation value $\langle J_z\rangle_{\Psi(t)}$ can be viewed as a time-dependent functional of the measurement record up to time t. We also note that a maximization of Λ with respect to its first argument is equivalent the maximization of a *log likelihood ratio* (LLR), $\lambda \equiv \ln\Lambda$. Under the replacements $x(t) \to y(t)$, $\theta_n \to \Psi_n(0)$, and $m(\theta_n,t,x(t)) \to \sqrt{\kappa}\,\langle J_z\rangle_{\Psi_n(t)}$, we have

$$\lambda(\Psi_1(0),\Psi_2(0)) \equiv \ln(\Lambda)$$
$$= \sqrt{\kappa}\int_0^T \left(\langle J_z\rangle_{\Psi_1(s)} - \langle J_z\rangle_{\Psi_2(s)}\right)dy(s)$$
$$- \frac{\kappa}{2}\int_0^T \left(\langle J_z\rangle_{\Psi_1(s)}^2 - \langle J_z\rangle_{\Psi_2(s)}^2\right)ds. \quad (13)$$

The MLE we use is then

$$|\Psi_{\mathrm{ML}}\rangle = \underset{\Psi_1 \in \mathbb{C}^{2^{\otimes N}}}{\arg\max}\left[\lambda(\Psi_1(0),\Psi_2(0))\right]. \quad (14)$$

In principle, the exact choice of $\Psi_2(0)$ is irrelevant for computing $|\Psi_{\mathrm{ML}}\rangle$, as the replacement $\Psi_2(0) \to \Psi_3(0)$ changes λ by a finite additive constant but does not affect where the maximum occurs. In practice, however, an initial condition that is radically different from the true one greatly reduces the numerical stability of Eq. (6). This fact impacts the choice of reference and the reconstruction algorithm we implement.

IV. ESTIMATING THE STATE OF A PURE QUBIT

As a first step towards understanding the fundamental limits of QST based on continuous-time measurement and control, we consider the simplest problem: reconstructing the state of a pure single qubit. We assume that we are initially given N qubits, each initialized in an unknown yet pure state $|\psi_0\rangle$. We also assume that the total evolution preserves the exchange symmetry of the system, thus allowing us to restrict our attention to states that are in the fully symmetric subspace of the many-body system. The evolution thus preserves the total collective angular momentum quantum number at its maximum value $J = N/2$. Therefore, instead of considering the entire 2^N-dimensional tensor-product Hilbert space, we are able to restrict our attention to the evolution to the $(d = N + 1)$-dimensional exchange-symmetric subspace.

A key ingredient of the protocol is to drive the system with a control Hamiltonian that ensures that the measurement record is informationally complete. Following the work of Riofrío *et al.* [4], we choose a control Hamiltonian, $H_c(t)$, that is randomized between a set of operators that rapidly generates the group of SU(2) rotations,

$$H_c(t) = \boldsymbol{b}(t) \cdot \boldsymbol{J} = \sum_i \boldsymbol{b}(t) \cdot \boldsymbol{\sigma}^{(i)}/2, \quad (15)$$

with

$$\boldsymbol{b}(t) = \frac{\pi}{2\tau}\sum_i \chi_{[i-1,i)}(t/\tau)\,\mathbf{e}_i. \quad (16)$$

Here $\{\mathbf{e}_i\}$ are uniformly sampled directions on the unit sphere, τ is the transition period, and the indicator function

$\chi_{[a,b)}(x) = 1$ for $x \in [a,b)$ and 0 otherwise. The choice of a Larmor frequency $\Omega_b = \pi/(2\tau)$ is an attempt to maximize the information gain; e.g., if $\mathbf{e}_i = \mathbf{e}_x$, then a $\pi/2$ rotation is needed to rotate the unobserved $\mathbf{e}_y$ component of the collective spin onto the measurement axis, $\mathbf{e}_z$.

For an arbitrary control law, the estimate $|\Psi_{\mathrm{ML}}\rangle$ in Eq. (14) does not have an analytic solution and therefore must be computed numerically. Taking the LLR as the cost function in the optimization, each evaluation of λ, Eq. (13), requires a comparison of two conditional states, as observed through their respective expectation values of J_z. This, in turn, requires an efficient method for integrating the CSE, Eq. (6), since a typical minimization algorithm will require many evaluations of λ. In general, this is a numerically intensive, as the dimension of the Hilbert space for the collective state in the symmetric Hilbert space grows as $N + 1$, and we seek to study the limits for large N. We can substantially reduce this numerical complexity by making an approximation on the measurement-induced dynamics.

To understand the appropriate approximation, let us consider how measurement backaction complicates the description of the dynamics. In general, the state of the symmetric ensemble of N particles is specified by all distinct symmetrized K-body correlation functions of Pauli products, $\langle \sigma_{\alpha_1}^{(1)} \sigma_{\alpha_2}^{(2)} \cdots \sigma_{\alpha_K}^{(K)} \rangle_{\mathrm{sym}}$, where $\sigma_{\alpha_i}^{(i)}$ acts on the ith spin with $\alpha_i \in \{x,y,z\}$ and $K = 1, \ldots, N$. For the special case of a spin coherent state (SCS), the state is completely specified only by the one-point correlation functions $n_\alpha = \langle \sigma_\alpha \rangle$: the Bloch vector of any of the identical qubits. The effect of measurement of the collective spin is twofold: (i) the Bloch vector is stochastically "kicked" when conditioned on the noisy measurement record; and (ii) higher order correlations (entanglement) are generated between the qubits. To lowest order, the measurement-induced correlations result in spin squeezing [17], specified by two-point correlations. For stronger measurements all correlations become important. In the absence of any control, the continuous measurement ultimately becomes projective, yielding a Dicke state (eigenstate of the collective J_z) as the steady state of a perfect QND measurement [27].

In the presence of strong randomized controls, the state evolution is dramatically different. In addition to causing precession of the mean spin, a transverse magnetic field will generally rotate the reduced uncertainty of the measured component into an orthogonal direction. The subsequent direction being measured will likely have increased uncertainty, i.e., is *antisqueezed*. This measurement will in turn reduce the previously increased uncertainty, resulting in at least a partial cancellation. The ultimate effect is that, with the application of rapid rotations about random directions, any spin squeezing produced at early times has a good chance of being undone at later times, and on average, the state remains close to an SCS.

An example of this effect is shown in Fig. 2, which contrasts the conditional evolution of a QND measurement of J_z without control to a system subjected to $10\,\pi/2$ rotations about random directions ($\Omega_b = 25\pi\kappa$) while being continuously measured. In both simulations we take $N = 75$ qubits ($J = 37.5$) initially prepared in an SCS along x. Figure 2(a) shows the trajectory that the mean spin takes on the Bloch sphere under the

influence of the controls, both in the absence of a continuous measurement and when conditioned on the measurement record in Fig. 2(b), *i* [solid (blue) line]. As a quantitative measure of the two-body correlations present in the system, in Fig. 2(b), *ii*, we plot the spin-squeezing parameter $\xi_T^2 \equiv \lambda_{\min}/J^2$ [28], where $\lambda_{\min}$ is the minimum eigenvalue of the matrix G with entries

$$G_{ij} = \frac{N}{2}\langle J_i J_j + J_j J_i \rangle_\Psi - (N-1)\langle J_i \rangle_\Psi \langle J_j \rangle_\Psi. \quad (17)$$

This particular parameter is qualitatively equivalent to the concurrence [29], a measure of pairwise entanglement between qubits. We see that, in the absence of the controls, squeezing grows monotonically, reaching its maximum value at a final time $\kappa t = 0.2$. However, in the presence of the controls the squeezing does not monotonically increase but, instead, reaches a maximum value at time $\kappa t \approx 0.1$ and then returns to a value near 0.

We can see how the controls average out the effect of squeezing and entanglement by plotting the spin-Husimi Q function at various sample times,

$$Q(t,\vartheta,\varphi) \equiv \frac{N+1}{4\pi} |\langle \vartheta,\varphi|\Psi(t)\rangle|^2, \quad (18)$$

where $|\vartheta,\varphi\rangle \equiv |\mathbf{n}(\vartheta,\varphi)\rangle^{\otimes N}$ is an SCS whose Bloch vector $\mathbf{n}$ is parameterized by the spherical coordinate angles ϑ and φ. For spin-squeezed states, the Q function takes the form of an approximately Gaussian distribution, centered at the mean spin position and with its minor axis oriented in the squeezing direction. Figure 2(c) shows contour plots of the Q function, both with and without controls, at times $\kappa t_0 = 0$, $\kappa t_1 = 0.03$, $\kappa t_2 = 0.1$, and $\kappa t_3 = 0.2$. The Q function in the presence of controls begins as an unsqueezed SCS and proceeds to rotate about the z axis, staying roughly near the equator. During this time it is also being squeezed, as its minor axis has strong overlap with the measurement axis. This continues until $\kappa t \sim 0.08$, when the controls rotate the mean spin to be near the $-z$ axis. As it does so, the minor and major axes are also rotated, so that by time $\kappa t = 0.1$ the antisqueezed major axis is almost aligned with the measurement axis. The remainder of the evolution returns the mean spin to near the equator, in such a way as to preserve this orientation and subsequently undoes the accumulated squeezing, as seen in the Q function at the final time, $\kappa t_3 = 0.2$. This is contrasted with the uncontrolled evolution, which shows a sequence of increasingly eccentric ellipses whose minor axes are always oriented along the z axis.

Given these facts, we propose the ansatz that the exact conditional state is well approximated by a conditional SCS, a state that is always a separable product, $|\Psi(t)\rangle \approx |\mathbf{n}(t)\rangle^{\otimes N}$, where $\mathbf{n}(t)$ is a conditional single-qubit Bloch vector. This ansatz allows us to extend the continuous-measurement QST protocol to include the effect of measurement backaction, by returning an estimate that depends only on the evolution of a single-body density operator. The direction of the Bloch vector will evolve under the control Hamiltonian, Eq. (15), with a stochastic component arising from the measurement backaction. We make this approximation by deriving the conditional evolution $\langle \boldsymbol{J} \rangle_{\Psi(t)}$ under the assumption that all moments are computed under an SCS approximation. The equation of motion for $\langle \boldsymbol{J} \rangle_{\Psi(t)}$ follows from the exact CSE,
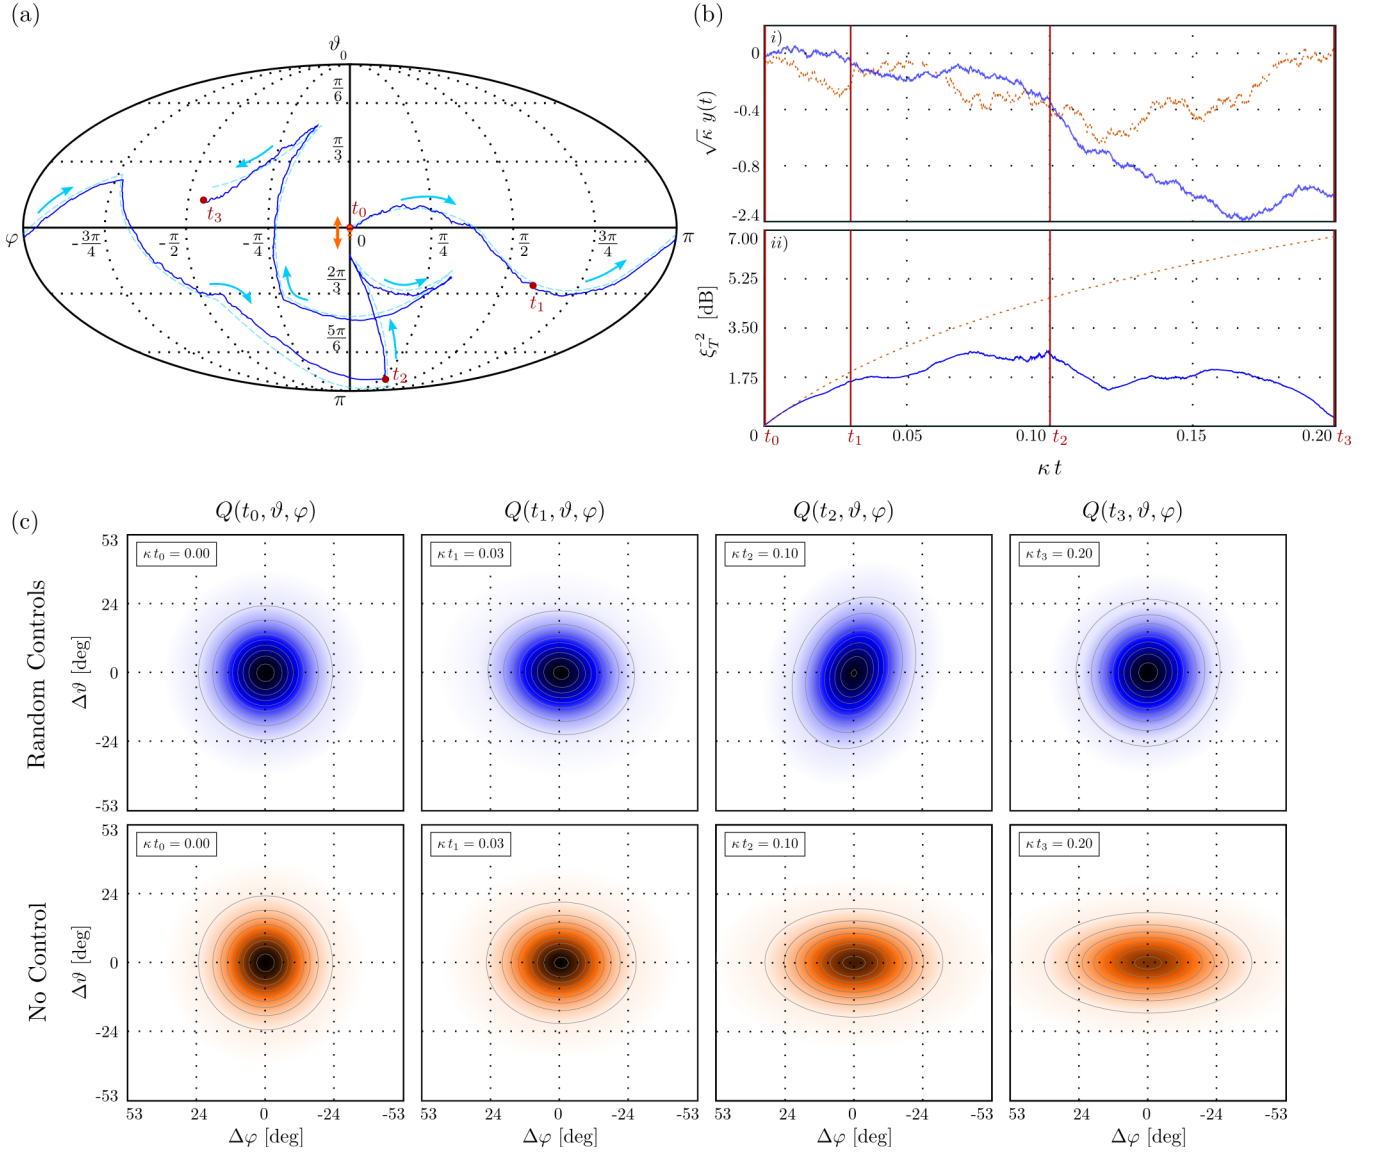
FIG. 2. (Color online) Simulations for $N = 75$ qubits, initially prepared in an SCS, polarized along x. (a) Trajectory of the mean spin on the Bloch sphere for the cases of a randomized control Hamiltonian with measurement [solid (dark-blue) line], a control Hamiltonian without measurement, i.e., $\kappa = 0$ [dashed (light-blue) line], and no controls but with measurement [dotted (orange) line]. (b) i: Simulated measurement records in the presence [solid (blue) line] and absence [dotted (orange) line] of a randomized control Hamiltonian. ii: The amount of spin squeezing (in dB) generated the presence [solid (blue) line] and absence [dotted (orange) line] of the controls. The squeezing is a measure of the correlations between qubits generated by the measurement backaction. (c) Conditional spin Q functions for both trajectories taken at sample times t_0, t_1, t_2, and t_3. The random rotations generated by the control Hamiltonian average out the effect of squeezing, leaving the collective state close to a product SCS.

yielding the Itō equation,

$$d\langle \boldsymbol{J} \rangle = \boldsymbol{b}(t) \times \langle \boldsymbol{J} \rangle dt - \tfrac{1}{8}\kappa \langle [J_z,[J_z,\boldsymbol{J}]] \rangle \, dt$$
$$+ \tfrac{1}{2}\sqrt{\kappa}(\langle J_z \boldsymbol{J} + \boldsymbol{J} J_z \rangle - 2\langle J_z \rangle \langle \boldsymbol{J} \rangle)\, dv(t)$$
$$= \boldsymbol{b}(t) \times \langle \boldsymbol{J} \rangle dt - \tfrac{1}{8}\kappa(\langle \boldsymbol{J} \rangle - \mathbf{e}_z \langle J_z \rangle)\, dt$$
$$+ \tfrac{1}{2}\sqrt{\kappa}(\langle J_z \boldsymbol{J} + \boldsymbol{J} J_z \rangle - 2\langle J_z \rangle \langle \boldsymbol{J} \rangle)\, dv(t), \quad (19)$$

where all expectation values are computed with state $|\Psi(t)\rangle$. Under the SCS approximation, $\langle \boldsymbol{J} \rangle \approx \frac{N}{2}\mathbf{n}(t)$ and $\langle J_z \boldsymbol{J} + \boldsymbol{J} J_z \rangle - 2\langle J_z \rangle \langle \boldsymbol{J} \rangle \approx \frac{N}{2}[\mathbf{e}_z - \langle \sigma_z \rangle \mathbf{n}(t)]$. The conditional evo-

lution of the Bloch vector $\mathbf{n}(t)$ thus obeys the stochastic differential equation (SDE),

$$d\mathbf{n}(t) = \left(\boldsymbol{b}(t) \times \mathbf{n}(t) - \tfrac{1}{8}\kappa(\mathbf{n}(t) - \mathsf{z}(t)\,\mathbf{e}_z) \right) dt$$
$$+ \tfrac{1}{2}\sqrt{\kappa}\,(\mathbf{e}_z - \mathsf{z}(t)\,\mathbf{n}(t))\, dv(t), \quad (20a)$$
$$dv(t) = dy(t) - \sqrt{\kappa}\tfrac{N}{2}\mathsf{z}(t)\, dt, \quad (20b)$$

where $\mathbf{n} \equiv (\mathsf{x},\mathsf{y},\mathsf{z}) \equiv (\langle \sigma_x \rangle, \langle \sigma_y \rangle, \langle \sigma_z \rangle)$. This is the same SDE we would derive for the conditional evolution of a single qubit, with the exception that the innovation $v(t)$ expects a signal scaled by the factor $J = N/2$. Note that this equation is valid

for both pure and mixed single-qubit states, a fact we exploit in our reconstruction algorithm.

To test the quality of this approximation, we compare the exact evolution of the collective state $|\Psi(t)\rangle$, governed by the CSE, Eq. (6), to that given by the SCS approximation, $|\Psi_{\mathrm{SCS}}(t)\rangle = |\mathbf{n}(t)\rangle^{\otimes N}$, governed by Eq. (20). Given the same SCS initial condition in both cases, we compare these states in two ways. First, we compute the fidelity $\mathcal{F} = |\langle \Psi_{\mathrm{SCS}}(t)|\Psi(t)\rangle|^2$, as a function of time. Second, we compute the root mean square (RMS) error between $\langle J_z \rangle_{\Psi(t)}$ and $\langle J_z \rangle_{\mathrm{SCS}(t)}$ as defined by the quantity

$$\Delta \mathsf{z}_{\mathrm{err}}(t) \equiv \sqrt{\left\langle \left(\tfrac{1}{J}\langle J_z\rangle_{\Psi(t)} - \mathsf{z}(t)\right)^2\right\rangle_\nu}. \tag{21}$$

The expectation values have been scaled by the total spin length J to allow for a comparison between different values of N. This quantity impacts the performance of the estimator, since any error in $\langle J_z \rangle$ directly impacts the LLR. The ensemble average is computed for $\nu = 100$ unit vectors uniformly sampled over the Bloch sphere and uses only a single noise realization per state.

Figure 3 shows this average fidelity, $\langle \mathcal{F} \rangle_\nu$, for a variety of numbers of qubits, N, both with and without 40 $\pi/2$ rotations about random directions, for a total measurement time $\kappa T = 0.8$. The SCS approximation performs poorly in the absence of the controls and for large N, showing a worst-case average fidelity of $\langle \mathcal{F} \rangle_\nu \sim 0.47$ for $N = 100$. In the presence of the controls the approximation performs well, maintaining the fidelity at a level $\langle \mathcal{F} \rangle_\nu > 0.80$ for all N tested. The nonmonotonic decrease in the average fidelity implies that the controls could be optimized to maximize this value, however, it is unclear if such an optimization would return an optimal tomographic estimate. Figures 3(c) and 3(d) show that for all of the N that we simulated, the SCS approximation tracks the mean spin with $\Delta \mathsf{z}_{\mathrm{err}}(t) < 0.1$ and that, in the presence of the controls, $\Delta \mathsf{z}_{\mathrm{err}}(t) \sim 0.025$. The case $N = 1$ shows that
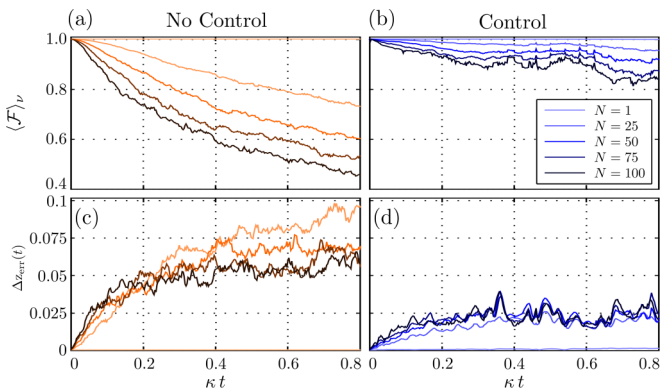

FIG. 3. (Color online) Performance of the separable SCS approximation. Average fidelity between the exact state $|\Psi(t)\rangle$ and the SCS $|\mathbf{n}(t)\rangle^{\otimes N}$ as a function of time (a) for no controls and (b) applying 40 $\pi/2$ rotations over a time $\kappa T = 0.8$. RMS error $\Delta \mathsf{z}_{\mathrm{err}}(t)$ as defined in Eq. (21) is plotted (c) for no controls and (d) with the same control law as in (b). The average is over $\nu = 100$ random initial unit vectors, uniformly sampled over the Bloch sphere, with a single noise realization per state. We make these comparisons for $N = 1$, 25, 50, 75, and 100 qubits, plotted with a correspondingly increasing contrast and decreasing fidelity.

Eq. (20) exactly reproduces the single-qubit evolution, up to numerical precision.

V. NUMERIC SIMULATIONS

Armed with the SCS approximation we are able to efficiently compute an approximate version of the LLR given in Eq. (13). We now test the performance of our tomographic procedure via a series of numerical simulations. In the absence of a closed-form solution to Eqs. (13) and (20), we must also find $|\Psi_{\mathrm{ML}}\rangle$ through a numerical search. While we may choose from any number of algorithms (e.g., gradient assent), we use a particularly simple procedure here. As the Bloch sphere is such a small search space, we simply sample a suitably dense set of initial conditions and then choose as our estimate the element that maximizes λ. We operate with a density of samples such that the average infidelity between nearest neighbors is $\sim 6 \times 10^{-4}$. This ensures that we will obtain an estimate that is sufficiently close to the true state. As an example, given $N = 100$ qubits, the optimum POVM bound sets an average infidelity of 0.01 [18], implying that any deficits in our procedure should not be attributed to the finite number of samples.

In practice, we need to consider an additional step in our protocol. A CSE with an informationally complete measurement record is, in principle, stable [30]. This means that, given a measurement record generated from an initial state $|\Psi(0)\rangle$, it is possible to integrate a CSE from any initial condition $|\Psi(0)'\rangle \neq |\Psi(0)\rangle$, such that $|\Psi(t)'\rangle \rightarrow |\Psi(t)\rangle$ as $t \rightarrow \infty$. In a sense, this means that the CSE is self-correcting for the initial misinformation. Unfortunately, we find that the numerical stability of both Eqs. (6) and (20) is quite poor when the initial condition is nearly orthogonal to the true state. This affects our reconstruction procedure because an instability in computing either the candidate or the reference state can result in λ reaching arbitrarily large or small values. To correct for this issue, we compute λ first for *mixed* initial conditions and then for a spread of pure states in the direction of the most likely mixed state. This two-step procedure greatly improves the numerical stability because a mixed state polarized in a direction orthogonal to the true state still has some overlap with that state.

In the first step, we use $M_1 = 250$ isotropically distributed mixed states, whose Bloch vectors form the set $\mathcal{N}_1 = \{\mathbf{n}_m \in \mathbb{R}^3 : \|\mathbf{n}_m\| = 3/4, m = 1, \ldots, M_1\}$. To identify an acceptable pure reference state, we find the mixed-state sample $\mathbf{n}^\star \in \mathcal{N}_1$ that maximizes the approximate LLR,

$$\lambda_{\mathrm{scs}}(\mathbf{n}_m, \mathbf{n}_r) = \frac{\sqrt{\kappa N}}{2} \int_0^T [\mathsf{z}_m(s) - \mathsf{z}_r(s)] \, dy(s)$$
$$- \frac{\kappa N^2}{8} \int_0^T [\mathsf{z}_m(s)^2 - \mathsf{z}_r(s)^2] \, ds, \tag{22}$$

where we choose the unbiased reference initial condition $\|\mathbf{n}_r(0)\| = 0$. From this mixed state, we then define the new reference vector $\mathbf{n}'_r \equiv \mathbf{n}^\star / \|\mathbf{n}^\star\|$ and uniformly sample $M_2 = 250$ pure states within a neighborhood of this vector. Specifically, we form the set $\mathcal{N}_2 = \{\mathbf{n}_m \in \mathbb{R}^3 : \|\mathbf{n}_m\| = 1, \mathbf{n}_m \cdot \mathbf{n}'_r \geqslant \cos(\pi/4), m = 1, \ldots, M_2\}$. We then report as an estimate the single-qubit state $|\mathbf{n}_{\mathrm{ML}}\rangle$ whose Bloch vector $\mathbf{n}_{\mathrm{ML}} \in \mathcal{N}_2$
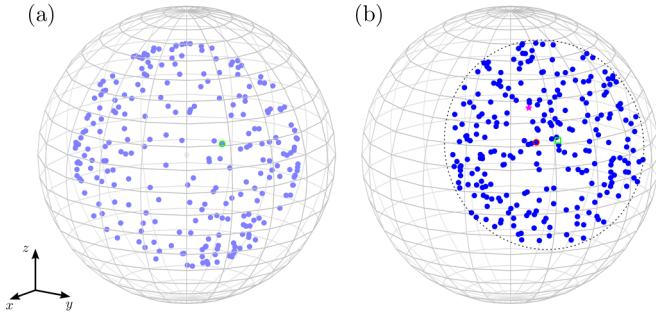
FIG. 4. (Color online) Typical sample distributions. (a) The set $\mathcal{N}_1$ shown inside the unit sphere for $M_1 = 250$ and $\|\mathbf{n}_m(0)\| = 3/4$. The point that maximized λ_{scs} relative to $\mathbf{n}_r(0) = 0$, given a measurement record made with $N = 75$ qubits, is circled (in green). (b) The resample set $\mathcal{N}_2$ for $M_2 = 250$ is plotted in blue on the Bloch sphere, with the maximum angular deviation of $\pi/4$ indicated by the dashed line. Also shown is the reference point $\mathbf{n}_r'$ [(green) square] and the true initial state [(magenta) star]. The resample point that maximizes λ_{scs} is circled (in red).

maximizes $\lambda_{\mathrm{scs}}(\mathbf{n}_m, \mathbf{n}_r')$. Figure 4 shows a typical realization of both sample sets, for a simulation over $N = 75$ qubits.

In order to characterize the performance of our protocol, we perform a series of numerical simulations for a variety of N. In each simulation we wish to compare the average infidelity between our estimate and the true input state, $1 - \langle \mathcal{F} \rangle_\nu$, averaged over uniformly sampled inputs and measurement realizations. We also compare our protocol to two additional measurement schemes. The first comparison is to the fundamental bound set by the optimum POVM, with $1 - \langle \mathcal{F} \rangle_{\mathrm{opt}} = 1/(N+2)$ [18]. The second comparison is to an alternative model of the continuous measurement, one that completely ignores measurement backaction. In other words, we wish to compare the above model to a model where the measurement record is approximated by

$$\tilde{y}(t) \approx w(t) + \frac{\sqrt{\kappa} N}{2} \int_0^t \langle \mathbf{n}(0) | \sigma_z(s) | \mathbf{n}(0) \rangle \, ds, \qquad (23)$$

where $\sigma_z(s)$ is the Heisenberg evolved Pauli z operator and $w(t)$ is a Wiener process. This model is equivalent to the $\gamma_{\mathrm{diss}} = 0$ limit of the single-atom density matrix $\tilde{\rho}$ defined in Eq. (5). While such a model is a good approximation when the total measurement time is very short compared to $1/\kappa$, we expect the effect of measurement backaction to have a significant impact on our estimator.

To make a fair comparison, we use a nearly identical algorithm in the backaction-free case as in the estimator described above. In this case we no longer have a problem with the numerical stability of our estimator because the Heisenberg equation of motion for σ_z is independent of the state. Therefore, we need not perform a two-step sampling procedure. We thus uniformly sample M pure Bloch vectors with a density equal to the final density of samples that we used in the procedure above, which requires $M = 1700$. We then choose the sampled state that maximizes a backaction-free version of the LLR, where the conditional expectation values $\langle J_z \rangle_{\Psi_i(t)}$ are replaced by $N \langle \mathbf{n}_i | \sigma_z(t) | \mathbf{n}_i \rangle / 2$, with the first sample $\mathbf{n}_1$ serving as the reference.
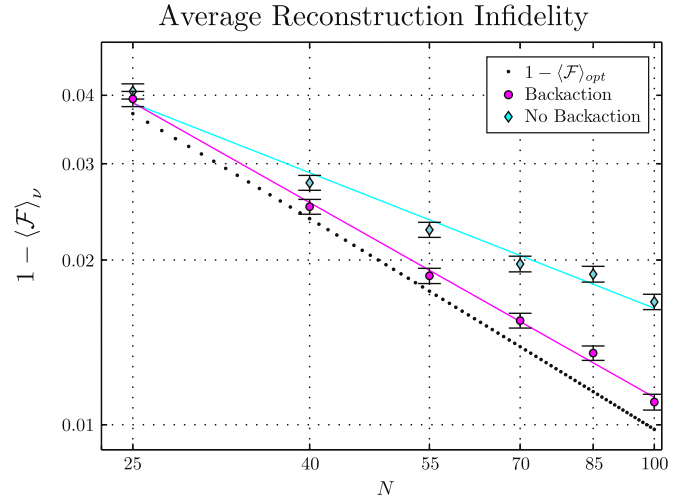
Average Reconstruction Infidelity

FIG. 5. (Color online) Average quantum state reconstruction infidelities (log-log axes) for different estimators. Circles show the performance of the MLE based on the LLR, Eq. (14), with the SCS approximation to the dynamical evolution, Eq. (22). A power-law fit to these data yields $1 - \langle \mathcal{F} \rangle_\nu \propto N^{-0.89}$, which is close to the bound set by the optimal POVM, $1 - \langle \mathcal{F} \rangle_\nu \approx N^{-1}$ (dotted line). Diamonds show the performance of an estimator that assumes a backaction-free measurement model given by Eq. (23) and achieves a power-law scaling, $1 - \langle \mathcal{F} \rangle_\nu \propto N^{-0.60}$. Error bars show a standard error of $\pm \sqrt{\mathrm{Var}[1 - \mathcal{F}]/\nu}$.

Figure 5 shows the results of numerical simulations for our reconstruction procedure with and without backaction. Plotted on a log-log scale is the average infidelity, $1 - \langle \mathcal{F} \rangle_\nu$, for $N = 25, 40, 55, 70, 85,$ and 100 qubits. For every N, we average over $\nu = 1000$ initial single-qubit states, with a single measurement realization per state. Every simulation used the same control law, with 40 randomized $\pi/2$ rotations, as well as a final time $\kappa T = 0.8$. Also shown are linear-least-squares fits to a power law, aN^b. With backaction, the best-fit parameters are $a = 0.69 \pm 0.08$ and $b = -0.89 \pm 0.03$, and without backaction, $a = 0.29 \pm 0.06$ and $b = -0.62 \pm 0.05$. These exponents are compared to the optimal scaling of $1/N$. By implementing the SCS approximation, we have introduced $\lesssim 5\%$ systematic errors in computing $\langle J_z \rangle$, which propagates into the approximate LLR, λ_{scs}, ultimately contributing to the suboptimal scaling.

The performance of the backaction-free estimator is best understood by considering not only the final-state reconstruction given the entire measurement record, but also the family of estimates generated by using data up a time $0 \leqslant t \leqslant T$. The stability of the CSE implies that for initial conditions $\Psi_1(0) \neq \Psi_2(0)$ we have the convergence $\langle J_z \rangle_{\Psi_1(t)} - \langle J_z \rangle_{\Psi_2(t)} \to 0$ as $t \to \infty$. The effect of this is that the LLR will either reach or asymptotically approach a steady-state value at long times. This convergence is necessarily implemented through the innovation, which occurs more rapidly for larger N, as follows from Eq. (20b). However, the unitary evolution in the backaction-free scheme is unable to implement such a convergence, and thus the LLR will never reach steady state. This ultimately biases the estimate away from the true state at

long times, which can been seen in the poor performance of the backaction-free estimate for large N.

VI. SUMMARY AND OUTLOOK

We have studied a protocol that performs QST using a single continuous measurement record of an ensemble of identical copies when the system is subjected to dynamical control and measurement backaction. We have considered the simplest case: estimation of the direction of the Bloch vector of a pure qubit in the absence of decoherence and systematic errors. This allowed us to focus on the effects of measurement backaction that complicate the estimator due to the nonlinearity of the conditional state evolution and the many-body nature of the dynamics induced by the entangling QND measurement.

We formulated an ML estimator and showed that it is possible to obtain a high-fidelity reconstruction of an initial SCS using only a single realization of a continuous collective measurement and dynamical control. Numerical simulations indicate that this estimate nearly reaches the bound set by the optimal POVM. If we fail to include the effect of measurement backaction in the conditional dynamics of the mean spin direction, an otherwise equivalent estimator becomes biased towards a poorer estimate at long times.

A key feature of our estimator was a simplification of the dynamical model in which the effects of measurement backaction act solely to induce random kicks in the direction of the Bloch vector while the entangling effects of the QND measure are small. This followed from the fact that the random rotations of the Bloch vector generated by the control Hamiltonian acted to average out the effects of squeezing as well as any higher-order correlations between the qubits. Such an approximation should also follow for ensembles of higher dimensional $d > 2$ qudits, and our protocol should be easily extended to reconstruction of such qudits. In this case we would require SU(d) controllability, as has been demonstrated for atomic spins [6]. The estimator would then be based on the stochastic evolution of a generalized Bloch vector, analogous to the qubit case considered here.

The next natural extension of our protocol is to include two-point atom-atom correlations in our estimator. This could allow us to improve the fidelity of our estimator and generalize the class of states we can reconstruct, including spin-squeezed states or other Gaussian states described by two-atom correlations. Going beyond two-body correlations, the matrix product state formalism provides a natural framework for studying many-body effects [31,32]. The application of the matrix product state ansatz to ML tomography has been studied [33], with good results. By translating this to the continuous measurement setting discussed here, one might be able to extract many-body correlations efficiently and robustly, of particular interest in the context of quantum simulators [34].

[1] H. Haffner, W. Hansel, C. F. Roos, J. Benhelm, D. Chek-al-kar, M. Chwalla, T. Korber, U. D. Rapol, M. Riebe, P. O. Schmidt *et al.*, Nature **438**, 643 (2005).

[2] D. Leibfried, E. Knill, S. Seidelin, J. Britton, R. B. Blakestad, J. Chiaverini, D. B. Hume, W. M. Itano, J. D. Jost, C. Langer *et al.*, Nature **438**, 639 (2005).

[3] A. Silberfarb, P. S. Jessen, and I. H. Deutsch, Phys. Rev. Lett. **95**, 030402 (2005).

[4] C. A. Riofrío, P. S. Jessen, and I. H. Deutsch, J. Phys. B **44**, 154007 (2011).

[5] G. A. Smith, A. Silberfarb, I. H. Deutsch, and P. S. Jessen, Phys. Rev. Lett. **97**, 180403 (2006).

[6] A. Smith, C. A. Riofrío, B. E. Anderson, H. Sosa-Martinez, I. H. Deutsch, and P. S. Jessen, Phys. Rev. A **87**, 030102 (2013).

[7] S. Chaudhury, S. Merkel, T. Herr, A. Silberfarb, I. H. Deutsch, and P. S. Jessen, Phys. Rev. Lett. **99**, 163002 (2007).

[8] S. T. Merkel, P. S. Jessen, and I. H. Deutsch, Phys. Rev. A **78**, 023404 (2008).

[9] G. A. Smith, S. Chaudhury, A. Silberfarb, I. H. Deutsch, and P. S. Jessen, Phys. Rev. Lett. **93**, 163602 (2004).

[10] A. H. Kiilerich and K. Mølmer, Phys. Rev. A **89**, 052110 (2014).

[11] S. Gammelmark and K. Mølmer, Phys. Rev. Lett. **112**, 170401 (2014).

[12] M. Tsang, Phys. Rev. Lett. **108**, 170502 (2012).

[13] K. Mølmer, arXiv:1408.4568 [quant-ph] (2014).

[14] S. Gammelmark, B. Julsgaard, and K. Mølmer, Phys. Rev. Lett. **111**, 160401 (2013).

[15] D. Tan, S. Weber, I. Siddiqi, K. Mølmer, and K. W. Murch, arXiv:1409.0510 [quant-ph] (2014).

[16] S. J. Weber, A. Chantasri, J. Dressel, A. N. Jordan, K. W. Murch, and I. Siddiqi, Nature **511**, 570 (2014).

[17] K. Hammerer, A. S. Sørensen, and E. S. Polzik, Rev. Mod. Phys. **82**, 1041 (2010).

[18] S. Massar and S. Popescu, Phys. Rev. Lett. **74**, 1259 (1995).

[19] E. Bagan, A. Monras, and R. Muñoz-Tapia, Phys. Rev. A **71**, 062318 (2005).

[20] I. H. Deutsch and P. S. Jessen, Opt. Commun. **283**, 681 (2010).

[21] B. Q. Baragiola, L. M. Norris, E. Montaño, P. G. Mickelson, P. S. Jessen, and I. H. Deutsch, Phys. Rev. A **89**, 033850 (2014).

[22] K. Jacobs and D. A. Steck, Contemp. Phys. **47**, 279 (2006).

[23] H. M. Wiseman and G. J. Milburn, *Quantum Measurement and Control* (Cambridge University Press, Cambridge, 2010).

[24] L. Bouten, R. van Handel, and M. R. James, SIAM J. Control Optim. **46**, 2199 (2007).

[25] R. S. Liptser and A. N. Shiriaev, *Statistics of Random Processes: I. General Theory* (Springer, Berlin, 2001).

[26] P. D. Feigin, Adv. Appl. Prob. **8**, 712 (1976).

[27] J. K. Stockton, R. van Handel, and H. Mabuchi, Phys. Rev. A **70**, 022106 (2004).

[28] G. Tóth, C. Knapp, O. Gühne, and H. J. Briegel, Phys. Rev. A **79**, 042334 (2009).

[29] X. Yin, X. Wang, J. Ma, and X. Wang, J. Phys. B **44**, 015501 (2011).

[30] R. Van Handel, Infin. Dimens. Anal. Quantum Prob. Relat. Topics **12**, 153 (2009).

[31] D. Perez-Garcia, F. Verstraete, M. M. Wolf, and J. I. Cirac, Quantum Inf. Comput. **7**, 401 (2007).

[32] F. Verstraete, V. Murg, and J. Cirac, Adv. Phys. **57**, 143 (2008).

[33] T. Baumgratz, A. Nüßeler, M. Cramer, and M. B. Plenio, New J. Phys. **15**, 125004 (2013).

[34] A. Wright, M. Chiao, D. Gevaux, A. Klopper, I. Georgescu, and B. Verberck (eds.), Nature Phys. Insight Quantum Simulat. 8, No. 4 (2012).