

Stabilizing two-qubit entanglement with dynamically decoupled active feedback

Sacha Greenfield^{1,2,3}, Leigh Martin^{4,5}, Felix Motzoi^{4,6,7}, K. Birgitta Whaley^{4,6}, Justin Dressel^{3,8}, and Eli M. Levenson-Falk^{1,2,9,*}

¹*Department of Physics and Astronomy, University of Southern California, Los Angeles, California 90089, USA*

²*Center for Quantum Information Science and Technology, University of Southern California, Los Angeles, California 90089, USA*

³*Institute for Quantum Studies, Chapman University, Orange, California 92866, USA*

⁴*Berkeley Center for Quantum Information and Computation, Berkeley, California 94720, USA*

⁵*Department of Physics, University of California, Berkeley, California 94720, USA*

⁶*Department of Chemistry, University of California, Berkeley, California 94720, USA*

⁷*Peter Grünberg Institute of Quantum Control (PGI-8), Forschungszentrum Jülich GmbH, Wilhelm-Johnen-Straße, Jülich 52428, Germany*

⁸*Schmid College of Science and Technology, Chapman University, Orange, California 92866, USA*

⁹*Ming Hsieh Department of Electrical and Computer Engineering, University of Southern California, Los Angeles, California 90089, USA*

 (Received 7 August 2023; revised 11 December 2023; accepted 24 January 2024; published 12 February 2024)

We propose and analyze a protocol for stabilizing a maximally entangled state of two noninteracting qubits using active state-dependent feedback from a continuous two-qubit half-parity measurement in coordination with a concurrent, noncommuting dynamical decoupling drive. We demonstrate that such a drive can be simultaneous with the measurement and feedback, while also playing a key part in the feedback protocol itself. We show that robust stabilization with near-unit fidelity can be achieved even in the presence of realistic nonidealities, such as time delay in the feedback loop, imperfect state-tracking, inefficient measurements, dephasing from $1/f$ -distributed qubit-frequency noise, and relaxation. We mitigate feedback-delay error by introducing a forward-state-estimation strategy in the feedback controller that tracks the effects of control signals already in transit. More generally, the steady state is globally attractive without the need for ancillas, regardless of the error state, in contrast to most known feedback and error-correction schemes.

DOI: [10.1103/PhysRevApplied.21.024022](https://doi.org/10.1103/PhysRevApplied.21.024022)

I. INTRODUCTION

Maximally entangled multiqubit states are important resources for quantum information processing tasks. However, entangled states decohere more rapidly than single-qubit states, and so must either be generated on demand when needed or stored using some form of stabilization that prevents degradation. Circuit-based quantum error-correction protocols are examples of such active entanglement stabilization [1], but are resource intensive since they require additional ancillary qubits, periodic entangling unitary gates, and single-qubit measurements. Bath engineering protocols can stabilize entanglement [2–7], but typically require custom device designs to select the state.

A natural nonunitary tool to consider for such an entanglement stabilization task is the direct measurement of a

two-qubit observable, such as parity [8,9]. Periodic or continuous measurements of such an observable can prepare [10] and then stabilize [11–13] the entangled eigenspaces of the measurement and can perform multiqubit operations [14,15], even in cases where the qubits do not directly interact [16–23]. Stabilizing a specific state, however, generally requires bath engineering or additional feedback control [24–29] to break the symmetry of the measurement toward the desired outcome.

Implementing robust entanglement stabilization using measurement feedback is challenging in practice. The stabilized state should not only be a steady state under ideal conditions for the measurement and other engineered decoupling drives [30–36], but remain stabilized in the presence of realistic nonidealities [37]. Some dynamical states such as single-qubit Rabi oscillations [38] are less sensitive to environmental noise and so can be more easily stabilized, but entangled states are more susceptible to decoherence, which can drastically lower the achievable

*elevenso@usc.edu

steady-state entanglement [16–18]. Dynamical decoupling can reduce dephasing and prolong coherence [39–41], but does not produce true stabilization [42] and is typically separated from the measurement and feedback. Time delays in the feedback loop [43,44] and measurement inefficiency are also unavoidable in the laboratory [45], and can be disastrous to the stabilization fidelity if not accounted for in the protocol [33,46,47].

In this work, we show that it is possible to combine measurement, a feedback drive, and a dynamical decoupling drive simultaneously to achieve robust entanglement generation and stabilization. As a case study, we revisit the simple measurement-feedback protocol introduced in Ref. [33], which stabilizes a maximally entangled two-qubit Bell state with feedback from a half-parity measurement. This protocol was vulnerable to dephasing that would mix the target state with a degenerate measurement eigenstate, thus eliminating stability. We expand the protocol to include a concurrent drive that provides both dynamical decoupling [48–51] and acts to remove an undesired fixed point of the dynamics. We analyze both the optimum settings and resulting performance in the presence of experimental nonidealities, including the time delay of the feedback loop, measurement inefficiency, and environmentally induced relaxation and dephasing. Notably, we model environmental dephasing as more realistic non-Markovian qubit-frequency noise from environmental fluctuators, with either white or $1/f$ noise spectra [52], which we compare to the standard (Lindblad) Markovian techniques for modeling dephasing.

Our modified stabilization protocol benefits from two key innovations. First, our dynamical decoupling drive notably differs from standard approaches since it neither acts only during idle times nor acts only on auxiliary subspaces. Instead, the decoupling drive is always on as a second drive that can optionally be adjusted by the feedback controller. This drive both reduces dephasing and provides an active pathway for stabilization out of erroneous states. Second, we develop a forward-state-estimation procedure for the feedback controller that accounts for the time delay in the feedback loop by anticipating the evolution occurring from signals already in transit. We find through numerical simulations that the inclusion of these features significantly improves the performance of the stabilization protocol. For dephasing from $1/f$ fluctuators noise, in particular, the modified protocol can robustly achieve near-unit asymptotic fidelity to the target state in a wide range of conditions.

II. STATE STABILIZATION PROTOCOL

As in Refs. [16,33], our system is composed of two qubits that are not interacting, described by a joint (system) density operator $\rho_{\text{sys}}(t)$. Each qubit experiences environmental dephasing at an average rate Γ_2 , detailed later. We

assume a rotating frame, independent control drives, and coupling to a common half-parity measuring device (e.g., Refs. [10,13,16]).

The half-parity observable $\hat{N} = (\hat{\sigma}_z^{(1)} + \hat{\sigma}_z^{(2)})/2$ has three eigenvalues: -1 , 0 , and 1 . The eigenvalues -1 and 1 have the even-parity eigenvectors $|11\rangle$ and $|00\rangle$, respectively, while the eigenvalue 0 has a degenerate odd-parity subspace spanned by $|01\rangle$ and $|10\rangle$. The two odd-parity entangled Bell states, $|\psi_{\pm}\rangle \equiv (|01\rangle \pm |10\rangle)/\sqrt{2}$ are thus preserved by the measurement, while the even-parity Bell states $|\phi_{\pm}\rangle \equiv (|00\rangle \pm |11\rangle)/\sqrt{2}$ get disturbed.

The goal of the protocol is to use the two single-qubit control drives and the joint half-parity measurement to stabilize one maximally entangled Bell state: $|\psi_{+}\rangle$.

To use feedback to accomplish this goal, we apply control drives that depend upon the collected measurement record $r(t)$, which has the approximate form

$$r(t)dt = \text{Tr}(\rho_{\text{sys}}(t)\hat{N})dt + \sqrt{\tau} dW, \quad (1)$$

where $dW^2 = dt$ is a stochastic Wiener increment [24] and $\tau = 1/2\eta\Gamma$ is a timescale set by the measurement rate $\eta\Gamma$, given efficiency $\eta \in [0, 1]$ and measurement-dephasing rate Γ . This stochastic process models, e.g., a homodyne measurement of $\hat{N}$ after phase-sensitive amplification of the informational quadrature, then shifting and rescaling to normalize the mean signal to the eigenvalues of $\hat{N}$ [53].

We synchronize the applied single-qubit drives to be either corotating or counter-rotating, $\hat{\sigma}_y^{(1)} \pm \hat{\sigma}_y^{(2)} \equiv \hat{Y}_{\pm}$, yielding the effective two-qubit drive Hamiltonians

$$\hat{H}_{\Omega}(t) = \hbar\Omega\hat{Y}_{+}/2, \quad \hat{H}_{\Delta}(t) = \hbar\Delta\hat{Y}_{-}/2, \quad (2)$$

with amplitudes Ω and Δ that are generally functions of time t , the collected record $r(t - \tau_d)$ after a delay τ_d from the feedback loop, and/or an estimated state $\rho_{\text{est}}(t)$. Importantly, both $\hat{N}$ and $\hat{H}_{\Delta}$ act trivially on the stabilized $|\psi_{+}\rangle$, meeting the condition for global convergence [25,37]. The corotating drive $\hat{H}_{\Omega}$ with matched drive phases generates rotations in the $\{|\psi_{+}\rangle, |\phi_{-}\rangle\}$ subspace that contains the target state $|\psi_{+}\rangle$, making it the *feedback drive* [33]. The counter-rotating drive $\hat{H}_{\Delta}$ with relative π -phase shift generates rotations in the $\{|\phi_{+}\rangle, |\psi_{-}\rangle\}$ subspace orthogonal to $|\psi_{+}\rangle$; these rotations also suppress slowly varying dephasing noise on $|\psi_{+}\rangle$, making this the *dynamical decoupling drive*.

Figure 1 shows a conceptual schematic of how the two drives and the measurement both prepare and stabilize the entangled state $|\psi_{+}\rangle$. After preparing the product state $|+, +\rangle = (|\psi_{+}\rangle + |\phi_{+}\rangle)/\sqrt{2}$, the measurement stabilizes $|\psi_{+}\rangle$ at dephasing rate Γ , but stochastically transfers within the subspace $|\phi_{+}\rangle \leftrightarrow |\phi_{-}\rangle$. The feedback drive unitarily rotates between $|\phi_{-}\rangle \leftrightarrow |\psi_{+}\rangle$ at rate Ω , but the measurement feedback law biases this rotation to accumulate population in $|\psi_{+}\rangle$. Meanwhile, dephasing induces




FIG. 1. Schematic for stabilizing the entangled Bell state $|\psi_+\rangle$ (red). Two noninteracting qubits $\rho_{\text{sys}}(t)$ undergo a half-parity measurement with dephasing rate Γ (orange), yielding a noisy record r and stochastic transfer $\hat{M}_r$ between $|\phi_+\rangle \leftrightarrow |\phi_-\rangle$ while stabilizing $|\psi_\pm\rangle$. Environmental dephasing at rate Γ_2 (pink) also causes stochastic transfer $\hat{H}_{\Gamma_2}$ between $|\phi_+\rangle \leftrightarrow |\phi_-\rangle$ and $|\psi_+\rangle \leftrightarrow |\psi_-\rangle$. A controller uses the measurement record to internally track an estimated state $\rho_{\text{est}}(t)$, then applies a feedback drive $\hat{H}_\Omega$ (blue) after a delay τ_d at rate $\Omega[\rho_{\text{est}}(t)]$ that induces a biased transfer $|\phi_-\rangle \rightarrow |\psi_+\rangle$, as well as a dynamical decoupling drive $\hat{H}_\Delta$ (purple) at rate $\Delta[\rho_{\text{est}}(t)]$ that stabilizes $|\psi_+\rangle$ and $|\phi_-\rangle$ and also transfers $|\psi_-\rangle \leftrightarrow |\phi_+\rangle$.

transfers between $|\psi_+\rangle \leftrightarrow |\psi_-\rangle$ and $|\phi_+\rangle \leftrightarrow |\phi_-\rangle$ at rate Γ_2 , with $|\psi_-\rangle$ invisible to both Ω and the measurement. Thus, $|\psi_-\rangle$ is an undesired fixed point of the protocol so far. However, the decoupling drive unitarily rotates between $|\psi_-\rangle \leftrightarrow |\phi_+\rangle$ at a rate Δ ; the measurement (at rate Γ) and corotating drive (at rate Ω) then provide a path from $|\phi_+\rangle \rightarrow |\psi_+\rangle$ as described above. Neither the dephasing Γ_2 nor the decoupling Δ were considered in Ref. [33], nor in similar ancilla-free protocols in Refs. [10,22,23,31].

A. Modeling the feedback controller

Simulations of real-time measurement feedback must distinguish between the true (but inaccessible) state of the system $\rho_{\text{sys}}(t)$ and the estimated state $\rho_{\text{est}}(t)$ being internally tracked by the controller, both indicated in Fig. 1. Simulations must evolve both $\rho_{\text{sys}}(t)$ to determine the measurement record in Eq. (1) and $\rho_{\text{est}}(t)$ to determine the feedback drives in Eqs. (C1) in Appendix C [47,54].

The evolution of the internal state $\rho_{\text{est}}(t)$ has three key differences. First, the controller does not know the specific environmental fluctuations causing ensemble decoherence, so can only use an average decoherence model. Second, the controller only has the collected record $r(t)$ up to a delay τ_d before the output feedback signal will reach the system state $\rho_{\text{sys}}(t + \tau_d)$. Third, whereas the system state

evolves with simulation time step dt , the controller estimate $\rho_{\text{est}}(t)$ evolves with time step $\delta t = n dt$, where n is some integer. This time step is called a “control cycle” and represents the time interval in which the controller performs the Bayesian update, estimates the state, and updates the drive tones. Differentiating the two time steps is important because, whereas dt must be much smaller than any other simulation timescale to maintain accuracy of the discrete state update, δt will be experimentally constrained by field-programmable gate array (FPGA) controller latency and may be comparable to other timescales (such as a Rabi period).

B. Evolution with decoherence and measurement

We model dephasing as stochastic dispersive energy shifts from weakly coupling to a noisy environment [52]:

$$\hat{H}_w(t) = \hbar\omega(\chi_1(t)\hat{\sigma}_z^{(1)} + \chi_2(t)\hat{\sigma}_z^{(2)})/2. \quad (3)$$

Here $\chi_{1,2}(t)$ are independent zero-mean stochastic processes, which we generate with either white or $1/f$ power spectral densities to encompass the two extremes of commonly observed frequency noise. From here on we refer to these as “white noise” and “fluctuator noise.” We calibrate the amplitude $\hbar\omega$ so that ensemble averaging the fluctuations yields the dephasing rate Γ_{fl} (Appendix A).

After measuring $\hat{N}$, the observed record $r(t)$ in Eq. (1) integrated over a small time step dt is Gaussian distributed with mean $\text{Tr}(\rho_{\text{sys}}(t)\hat{N})$ and variance $1/2\eta\Gamma dt$. The induced partial collapse is thus encapsulated by the Kraus operator update [47,53]

$$\hat{M}_{r(t)dt} = \exp[\eta\Gamma dt(r(t)\hat{N} - \hat{N}^2/2)], \quad (4)$$

$$\mathcal{M}_{r(t)}[\rho(t)] = \hat{M}_{r(t)}\rho(t)\hat{M}_{r(t)}^\dagger / \text{Tr}(\hat{M}_{r(t)}^\dagger \hat{M}_{r(t)}\rho(t)). \quad (5)$$

For inefficient measurements, $\eta < 1$, the uncollected part of the signal must be averaged over as unknown, causing residual measurement dephasing. This dephasing is described with an additional jump-no-jump Lindblad map. We use the same procedure to model relaxation (T_1) processes, and group them together into the single map $\mathcal{L}_{dt}[\rho(t)]$ [55]:

$$\mathcal{L}_{dt}[\rho(t)] = \hat{J}_0\rho(t)\hat{J}_0^\dagger + \sum_{k=1}^3 (\gamma_k dt) \hat{J}_k\rho(t)\hat{J}_k^\dagger \quad (6)$$

with

$$\begin{aligned} \hat{J}_0 &= \sqrt{1 - \sum_{k=1}^3 (\gamma_k dt) \hat{J}_k^\dagger \hat{J}_k}, & \hat{J}_{1,2} &= \hat{\sigma}_-^{(1,2)} / \sqrt{2}, \\ \hat{J}_3 &= \hat{N} / \sqrt{2} & \gamma_{1,2} &= 1/T_1, & \gamma_3 &= (1 - \eta)\Gamma. \end{aligned}$$

The composite update map for the system state is thus

$$\rho_{\text{sys}}(t + dt) = [\mathcal{U}_{dt} \circ \mathcal{L}_{dt} \circ \mathcal{M}_{r(t)dt}][\rho_{\text{sys}}(t)] \quad (7)$$

with the unitary update, $\mathcal{U}_{dt}[\rho(t)] = \hat{U}_{dt,\text{tot}}(t)\rho(t)\hat{U}_{dt,\text{tot}}^\dagger$, where

$$\hat{U}_{dt,\text{tot}} = \exp(dt[\hat{H}_\Delta(t) + \hat{H}_\Omega(t) + \hat{H}_\omega(t)]/i\hbar), \quad (8)$$

recalling that Ω and Δ generally depend upon $\rho_{\text{est}}(t)$. This completely positive map maintains state integrity on iteration and is accurate to order dt , which is chosen to be much smaller than any other period in the system evolution.

The controller update omits the noisy $\hat{H}_\omega$ as unknown and either replaces it with an average Lindblad dephasing map (Appendix A 5) or omits it as negligible over δt . It likewise uses a Lindblad map to model the effects of imperfect measurement efficiency and relaxation. It also performs a Bayesian update based on the measurement record averaged over the control cycle duration δt , i.e., $\bar{r}(t) \equiv (1/n) \sum_{m=0}^{n-1} r(t - m\delta t)$, where the full update is given by Eqs. (7) and (8) with $\rho_{\text{sys}}(t + dt) \rightarrow \rho_{\text{est}}(t + \delta t)$, $r(t) \rightarrow \bar{r}(t)$, and $dt \rightarrow \delta t$.

C. Forward estimation

Because of the feedback loop delay, the measurement record does not reflect the effects of the drives output over the past τ_d . This would cause inaccurate control if $\tau_d > \delta t$, as the system would overcorrect by repeatedly trying to rotate the state before the effects of those rotations could be measured. We therefore implement a novel *forward-estimation* protocol to correct for the delay. After the measurement update described above, the controller then makes a separate forward estimation by the delay time τ_d . This forward estimation predicts the effects of the control signals already output over the preceding τ_d , but of course does not include any measurement information that the controller would not have access to. For details of the forward-estimation procedure, see Appendix D. This $\rho_{\text{est}}(t + \tau_d)$ is used to determine the new control signals $\Omega(t + \delta t)$ and $\Delta(t + \delta t)$ to output, which are then added to the buffer of emitted control signals used for forward estimation in future control cycles. We show the effectiveness of this forward-estimation procedure in Sec. III C below.

D. Optimized feedback

The choice of feedback drives $\Omega(t)$ and $\Delta(t)$ in Eqs. (C1) in Appendix C affect the stabilization effectiveness, so we choose them to extremize the fidelity to $|\psi_+\rangle$,

$$\mathcal{F}[\rho_{\text{est}}(t)] = \langle \psi_+ | \rho_{\text{est}}(t) | \psi_+ \rangle, \quad (9)$$

after one time step δt . Imposing the conditions $\partial \mathcal{F}[\rho_{\text{est}}(t + \delta t)]/\partial \Delta = 0$ and $\partial \mathcal{F}[\rho_{\text{est}}(t + \delta t)]/\partial \Omega = 0$ yields

$$2\Delta_{\text{opt}}[\rho]\delta t = \tan^{-1}[2\text{Re}(\rho_{\phi_+, \psi_-})/(\rho_{\phi_+} - \rho_{\psi_-})], \quad (10)$$

$$2\Omega_{\text{opt}}[\rho]\delta t = \tan^{-1}[2\text{Re}(\rho_{\psi_+, \phi_-})/(\rho_{\psi_+} - \rho_{\phi_-})], \quad (11)$$

using the notation $\rho_{\psi,\phi} = \langle \psi | \rho | \phi \rangle$ and $\rho_\psi = \rho_{\psi,\psi}$. Intuitively, feedback coefficients are weighted *coherences* for the subspaces in which the drives and measurement act. In the limit of measurement strength much larger than other noise scales, this strategy produces a steady state that is self-healing for arbitrary error operators.

We also consider a dynamical decoupling (DD) strategy with an always-on constant drive $\Delta_{\text{opt}} \rightarrow \Delta_{\text{DD}}$. This simplification gives less-than-optimal fidelity in the next time step, but keeps the drive agnostic to state estimation errors, while increasing decoupling from dephasing (particularly for fluctuator noise) and still enabling the crucial transfer $|\psi_-\rangle \rightarrow |\phi_+\rangle$. The feedback drive $\hat{Y}_+$, however, disturbs the target state $|\psi_+\rangle$, so the optimal drive Ω_{opt} is necessarily state dependent even in this simplified protocol. As shown below, this trade-off of local optimality for enhanced decoupling is beneficial in some circumstances.

III. NUMERICAL SIMULATIONS

We first compare protocols with different combinations of feedback and counter-rotating drives across three different noise environments: pure fluctuator noise, fluctuator noise combined with T_1 decay, and pure white noise. To compare in a standard way across noise environments, we define the dephasing rate Γ_2 as follows: In the case of pure fluctuator or pure white noise, the total dephasing rate is simply $\Gamma_2 = \Gamma_{\text{fl}}$, i.e., all nonmeasurement-induced dephasing is due to averaged fluctuations, with Γ_{fl} calibrated separately (Appendix A 4). We assume that $T_1 \rightarrow \infty$ for both. In the case of fluctuator noise combined with T_1 decay, the dephasing is split evenly between the fluctuation-induced and T_1 -induced dephasing, such that $\Gamma_2 = \Gamma_{\text{fl}} + 1/2T_1$ and $\Gamma_{\text{fl}} = 1/2T_1 = \Gamma_2/2$, with T_1 decay described by Eq. (6).

A. Optimal feedback strategy

Figure 2 shows the time-dependent target-state fidelity from Eq. (9) after ensemble-averaging $N = 1000$ simulated trajectories per curve, with time $t \in [0, 150] \mu\text{s}$ discretized into bins $dt = 1 \text{ ns}$, and parameters $\Gamma_2 = 1/50 \mu\text{s}^{-1}$, $\Gamma = 1 \mu\text{s}^{-1}$, $\eta = 0.5$, and $\tau_d = 500 \text{ ns}$. We note that these values are well within the state of the art in the superconducting qubit community, with recently achieved efficiencies of $\eta > 0.7$ [56,57] and dephasing rates of $1/300 \mu\text{s}^{-1}$ [58]. The feedback delay time is constrained by the travel time for signals to enter and leave the cryostat and by processing time, leading to typical delay times of 200–400 ns [38]. Drives Ω and Δ are updated every $\delta t = 10\text{-ns}$ control cycle. In Figs. 2(a)–2(c), we compare three strategies using Ω_{opt} in Eq. (11):

(1) $\Omega_{\text{opt}}(\rho)$ and $\Delta = 0$, a protocol originally presented in Ref. [33];

- (2) $\Omega_{\text{opt}}(\rho)$ and Δ_{DD} , where a constant counter-rotating drive acts as simultaneous dynamical decoupling and feedback;
- (3) $\Omega_{\text{opt}}(\rho)$ and $\Delta_{\text{opt}}(\rho)$, where the counter-rotating drive choice is locally optimal [Eq. (10)].

We comment on these in turn.

(1) *No counter-rotating drive.*—In the presence of environmental noise, the corotating feedback drive Ω is inadequate for robust stabilization, regardless of the noise spectrum ($1/f$ versus white) or type (pure dephasing or dephasing and T_1). Decoherence populates $|\psi_{-}\rangle$, which is indistinguishable from $|\psi_{+}\rangle$ under the measurement and inaccessible to the corotating drive.

(2) *Simultaneous DD and feedback.*—For $1/f$ fluctuator noise [Fig. 2(a)], which causes a nonexponential short-time dephasing (Appendix A), a *state-agnostic constant decoupling drive simultaneous with the corotating feedback is asymptotically the best strategy, achieving near-unit fidelity*. The same approach leads to suboptimal saturating fidelities for fluctuator noise plus T_1 decay [Fig. 2(b)] or pure white noise dephasing [Fig. 2(c)]. These results make sense in light of known decoupling properties of $1/f$ noise and white noise [59] (Appendix A). With white noise dephasing or relaxation, the DD drive makes the $|\psi_{-}\rangle$ population accessible to the rest of the feedback loop (enabling some stabilization), but true dynamical decoupling does not occur.

(3) *Locally optimal feedback.*—For noise environments with a white component [Figs. 2(b) and 2(c)], the *locally optimal feedback on both corotating and counter-rotating drives is asymptotically the best strategy*. Furthermore, the qualitative similarities between fidelity time series for the two noise environments implies that the T_1 decay is mainly disruptive via its contribution of $1/2T_1$ to the overall dephasing rate Γ_2 , which makes sense given that decay of individual qubit populations from $|1\rangle \mapsto |0\rangle$ is a non-excitation-conserving map and can therefore be detected by the measurement and subsequently corrected by the feedback. Furthermore, since we model T_1 using Lindblad maps, the spectrum of the effective dephasing contribution is white, which further explains the similarities with pure white noise. The protocol performs well, particularly at short times, for pure $1/f$ fluctuators, but is asymptotically suboptimal if there is no broadband noise.

Figures 2(d)–2(f) show how decreasing measurement efficiency η and increasing time delay τ_d affects stabilization in the three noise environments, using the best feedback strategy for each environment, as determined above. Notably, for fluctuator noise, decreasing η only slows down the stabilization (due to the resulting decrease in measurement rate $\eta\Gamma$), but does not decrease the asymptotic fidelity. For white noise, as well as for fluctuator noise plus T_1 decay, the asymptotic fidelity is a steady state that balances the measurement rate $\eta\Gamma$ and dephasing rate Γ_2 ,




FIG. 2. Fidelity $\mathcal{F}(t)$ of the ensemble-averaged state ($N = 1000$ trajectories) to target $|\psi_{+}\rangle$, compared across different feedback protocols (a)–(c) and apparatus parameters (d)–(f). The ensemble is initialized in state $|+, +\rangle$ and evolves under optimal corotating drive $\Omega_{\text{opt},0}(\rho)$ and specified counter-rotating drive Δ . Except where specified, simulations use measurement-dephasing rate $\Gamma = 1 \mu\text{s}^{-1}$, environmental-dephasing rate $\Gamma_2 = 1/50 \mu\text{s}^{-1}$, feedback loop delay $\tau_d = 0.5 \mu\text{s}$, and measurement efficiency $\eta = 0.5$. (a)–(c) Comparison of three decoupling strategies, where Δ is either (1) off, (2) a constant $\Delta_{\text{DD}}/2\pi = 25 \text{ MHz}$ simultaneous with $\Omega_{\text{opt},0}$, or (3) an optimized state-dependent drive $\Delta_{\text{opt}}(\rho)$ simultaneous with $\Omega_{\text{opt},0}$. For dephasing from pure $1/f$ fluctuator noise in (a), simultaneous feedback and the constant counter-rotating (decoupling) drive achieve near-unit fidelity asymptotically. For fluctuator noise plus T_1 decay in (b), the state-dependent counter-rotating drive is most effective but saturates to imperfect stabilization. Pure white noise (c) mirrors this result with lower fidelities. (d)–(f) Comparison with different measurement efficiencies η and feedback loop delays τ_d , using the best control strategy for each noise type. For $1/f$ noise in (d), strategy (2) (constant Δ_{DD}) is applied, and nonidealities are observed to slow down stabilization without degrading the asymptotic near-unit fidelity. For fluctuator noise plus T_1 decay in (e) and white noise in (f), asymptotic fidelities depend on η and τ_d .

so it directly depends on η . These results are in agreement with other work showing that DD is most effective against nonwhite (time-correlated) noise [48].

We note that the locally optimal protocol attempts to optimize fidelity in the next time step, but assumes that

the decoherence rate is constant. However, a strong constant DD drive greatly reduces the effective dephasing rate in the case of slow ($1/f$) noise. We include this effect in our estimate of the state's dephasing (see Appendix A 5 for details). When $1/f$ dephasing is dominant, this DD approach is globally optimal; when white noise or relaxation is significant, the DD is less effective and so the locally optimal protocol is also globally optimal. Future work could explore whether hybrid strategies could be optimal for different ratios of slow and fast decoherence; for instance, a locally optimal protocol with periodic DD.

B. Optimizing control parameters

Figure 3 shows asymptotic fidelities at $t = 150 \mu\text{s}$, again applying the best strategy for each noise environment. The left panels, Figs. 3(a)–3(c), sweep the best strategy's main control parameter [Δ for Δ_{DD} and Γ for $\Omega_{\text{opt}}(\rho)$] versus the overall dephasing rate Γ_2 . As noted previously, for pure fluctuator or white noise [Figs. 3(a) and 3(c)], dephasing is entirely due to fluctuations ($\Gamma_2 = \Gamma_{\text{fl}}$), whereas for fluctuator noise plus T_1 decay [Fig. 3(b)], dephasing is evenly split between fluctuation and T_1 contributions ($1/2T_1 = \Gamma_{\text{fl}} = \Gamma_2/2$). Fixed parameters are $\tau_d = 0.5 \mu\text{s}$ and $\Gamma = 1 \mu\text{s}^{-1}$ for fluctuators. The right panels, Figs. 3(d)–3(f), sweep the measurement rate Γ versus the feedback loop delay τ_d for each noise environment, with fixed $\Gamma_2 = 1/50.0 \mu\text{s}^{-1}$, $\eta = 0.5$, and $\Delta_0/2\pi = 25 \text{ MHz}$ for fluctuators.

Similarly to Fig. 2, we observe that the pure fluctuator noise environment shows the highest stabilization fidelities, with $\mathcal{F}(t_f) > 0.98$ for any $\Gamma_2 \in [1/100.0, 1/10.0] \mu\text{s}^{-1}$ after optimizing the counter-rotating drive strength Δ_0 . Similarly, high fidelities are possible for all loop delays $\tau_d \in [0.1, 5.0] \mu\text{s}$ after optimizing over Γ . In particular, we highlight an optimal fidelity of $\mathcal{F} = 0.9999$ with control parameters $\Gamma = 5.76 \mu\text{s}^{-1}$ and $\Delta_0/2\pi = 25.0 \text{ MHz}$, operating at loop delay $\tau_d = 0.52 \mu\text{s}$ and dephasing rate $\Gamma_2 = 1/50.0 \mu\text{s}^{-1}$. While fluctuator noise plus T_1 decay and white noise environments demonstrate lower fidelities, we note that optimizing over measurement strength Γ still leads to high fidelities for realistic experimental parameters: using measurement strength $\Gamma = 10.0 \mu\text{s}^{-1}$ while operating at $\Gamma_2 = 1/50.0 \mu\text{s}^{-1}$ and $\tau_d = 0.52 \mu\text{s}$, optimal fidelities achieved are $\mathcal{F} = 0.95$ (fluctuator noise plus T_1 decay) and $\mathcal{F} = 0.92$ (white noise). This feedback delay is achievable with modern FPGA-based controllers, which we have tested and verified with placeholder calculations on a Quantum Machines' OPX controller; all other parameters are well within the current state of the art.

C. Optimality of forward estimation

We note that accounting for time delay in our state estimate is essential to the success of these protocols. If the




FIG. 3. Asymptotic fidelity $\mathcal{F}$ of the ensemble-averaged state ($N = 1000$ trajectories) to target $|\psi_+\rangle$ at $t_f = 150 \mu\text{s}$, with color displayed on a log scale. The optimal protocol for each noise type is used. Except where specified as a sweep variable, parameters default to $\eta = 0.5$, $\tau_d = 0.5 \mu\text{s}$, $\Gamma = 1 \mu\text{s}^{-1}$, $\Delta_0/2\pi = 25 \text{ MHz}$, $\Gamma_2 = 1/50.0 \mu\text{s}^{-1}$. (a) Fluctuator noise is mitigated with Δ_{DD} , sweeping the counter-rotating drive strength Δ_0 , for different dephasing rates $\Gamma_2 \in [0.01, 0.1] \mu\text{s}^{-1}$. Near-unit fidelity (> 0.99) is achieved for any dephasing rate with a sufficiently fast decoupling drive Δ_0 . (b) Fluctuator noise plus T_1 decay and (c) white noise are mitigated using Δ_{opt} , sweeping the measurement rate Γ for different Γ_2 . Fidelities are improved with stronger measurements (larger Γ), but lag behind the fidelities observed for fluctuators and Δ_{DD} . (d)–(f) Measurement rate Γ is swept for different feedback loop delay times τ_d for (d) fluctuator noise, (e) fluctuator noise plus T_1 decay, and (f) white noise. All demonstrate stabilization fidelities that are remarkably robust to loop delay, with short delays showing large optimal Γ (up to $10 \mu\text{s}^{-1}$ in the range swept) and long delays necessitating smaller Γ . For the realistically achievable $\Gamma_2 = 1/50.0 \mu\text{s}^{-1}$ and $\tau_d = 0.52 \mu\text{s}$, optimal fidelities achieved are (d) $\mathcal{F} = 0.9999$ at $\Gamma = 5.76 \mu\text{s}^{-1}$, (e) $\mathcal{F} = 0.95$ at $\Gamma = 10.0 \mu\text{s}^{-1}$, and (f) $\mathcal{F} = 0.92$ at $10.0 \mu\text{s}^{-1}$.

raw time-delayed state is used in Eqs. (10) and (11), we see only instability as the controller “overcorrects,” driving the system when previously emitted tones would already rotate to the target state. Our forward-estimation protocol addresses this issue. However, the behavior is reminiscent of a classical feedback controller with a feedback gain that is too large, which suggests a simpler protocol: turn down the strength of the feedback drives. We tested this deliberate weakening of the feedback drives such that $\Omega = \Omega_{\text{opt}}(\delta t/\Delta t)$. In this simpler protocol, the

non-forward-estimated (time-delayed) state is used as the state estimate, and Δt is a free parameter whose optimal value depends on other system timescales.

Three estimation strategies are compared in Fig. 4 for different loop delays τ_d and measurement rates Γ : no forward estimation with $\Delta t = \delta t$ (no attenuation), no forward estimation with the attenuation $\Delta t/\delta t$ optimized for each (τ_d, Γ) configuration, and forward estimation with $\Delta t = \delta t$. The first case leads to complete mixing of the state, resulting in a time-averaged fidelity of 0.25. This shows that some correction for feedback delay is essential. The second case—no forward estimation, but with attenuated drive—provides reliable stabilization, but the asymptotic fidelity achievable (averaged over the interval $[100, 150] \mu\text{s}$ to reduce noise) is generally lower than for the forward estimation. In particular, we note that tuning over the controllable parameter Γ indicates forward estimation as optimal for all τ_d tested, reducing infidelity by about 20% compared to the attenuated drive protocol. Thus, due to the low computational overhead of forward estimation, we selected forward estimation as the optimal estimation strategy for the data presented in Figs. 2 and 3, where, for simplicity, $\Delta t = \delta t$ was used for all forward-estimation simulations; optimizing over the feedback attenuation while still using forward estimation could further improve fidelity.

The attenuation-without-forward-estimation strategy is slightly simpler and so may be preferred in an experiment where computational overhead is the limiting factor, at the cost of lower fidelity. In a similar attempt at simplification, we also explored a version of the protocol with quasilinear feedback, where the feedback signal at time t is linearly proportional to the measurement signal $r(t)$ (see Appendix C 4). The proportionality constant derived depends on the prior state estimate, so it is not true linear feedback. However, while the linear feedback is qualitatively similar to the optimal feedback for stable trajectories, on average, it does not produce robust stabilization, which we again attribute to the finite time delay (see Appendix C 5).

IV. CONCLUSION AND DISCUSSION

We have presented and simulated a protocol that reliably generates and stabilizes an entangled state of two noninteracting qubits under realistic experimental conditions with dephasing from various realistic noise environments. More generally, in the limit of measurement strength much larger than other noise scales, our protocol produces a steady state that is self-healing for arbitrary error operators and without the use of ancillas. The success hinges on two key innovations of our protocol. First, we use a drive that both provides dynamical decoupling and transfers population out of an undesired fixed point of the dynamics, working concurrently with the feedback drive. Second, we develop a forward-estimation protocol for the feedback




FIG. 4. Ensemble fidelities averaged from 100 to 150 μs are compared on a log scale for different measurement rates Γ (x axis), loop delay times τ_d (color), and state estimation strategies (opaque versus translucent) for the fluctuator plus T_1 noise environment with $\Gamma_2 = 1/50.0 \mu\text{s}^{-1}$, $\eta = 0.5$, $\delta t = 10 \text{ ns}$. When control does not take time delay into account, fidelity is close to the purely mixed case $\mathcal{F} = 0.25$ (dashed lines). Forward estimation results in higher fidelities for all feedback loop delays τ_d after optimizing over the control parameter Γ , with typical infidelity reduction of about 20%, demonstrating the optimality of the protocol.

controller that compensates for loop delay, which can be used to improve other feedback control protocols. Together these provide evidence that continuous feedback control can achieve capabilities beyond what can be replicated with discrete (projective) measurement and gates, even in the presence of realistic nonidealities.

We note that, while our simulations take into account many experimental nonidealities, there are still approximations we have made. We treat the qubits as true two-level systems, while superconducting qubits are often weakly anharmonic. We note that the bandwidth of our drives ($2/\delta t = 50 \text{ MHz}$) is small compared to typical anharmonicity of a transmon (approximately 200 MHz), and so this assumption appears justified. We ignore stray qubit-qubit coupling, which can be eliminated in many architectures via multipath coupling (e.g., Ref. [60]). We ignore classical control crosstalk, which can be eliminated

through proper calibration. We also ignore the effects of finite readout cavity bandwidth (i.e., we make the “bad-cavity” approximation of Markovianity) and the effects of an improperly tuned measurement that can distinguish $|01\rangle$ from $|10\rangle$ (i.e., we assume a perfect half-parity measurement). These last assumptions are the most difficult to achieve experimentally. Future work could examine protocol performance when they are relaxed, for example by modeling cavity evolution directly as a truncated Fock space and including asymmetric measurement rates for $|01\rangle$ from $|10\rangle$.

Extensions of our work could optimize trade-offs between locally optimal feedback and dynamical decoupling based on noise spectra, optimize the forward-estimation protocol, and explore the impact of non-Markovian models of relaxation. Furthermore, our work can be generalized to protocols stabilizing other entangled states, including states with higher numbers of qubits. The generalization requires three components: (1) a joint measurement that has the target entangled state as an eigenstate; (2) combined single-qubit drives that rotate toward this target state from an entangled state that is not a measurement eigenstate (in our case the corotating drive); (3) combined single-qubit drives that have the target state as an eigenstate, but rotate all other degenerate measurement eigenstates to noneigenstate entangled states (in our case the counter-rotating drive). We leave the rigorous general analysis of this class of protocols to future work, but note that it would be fruitful to compare protocols obtained in this way to protocols obtained via more general reinforcement learning or gradient descent methods [61,62]. Applying machine-learning methods to finding optimal entanglement-generation protocols could guide the choice of generalized measurements, while identifying specific features of the protocols (e.g., the three components previously mentioned) could restrict the optimization space to reduce computational overhead of machine learning. We anticipate that our results will guide future studies of analog continuous measurement feedback protocols and provide an experimental roadmap for implementing fast digital processing while compensating for feedback delay.

ACKNOWLEDGMENTS

E.M.L.-F. thanks the members of QNL for facilitating this collaboration. S.G. thanks Arjendu Pattanayak for helpful conversations about the system-estimate distinction in quantum feedback control and Vinay Tripathi for discussions of simulating $1/f$ noise using fluctuators. S.G. also acknowledges funding from the Graduate Fellowships for STEM Diversity. This work was supported by the ONR under Grant No. N00014-21-1-2688, by Research Corp under Cottrell Scholarship 27550, by NSF-BSF under Grant Award No. 1915015, and by ARO/LPS under Grant Award No. W911NF-22-1-0258.

APPENDIX A: MODELING DEPHASING

We used three different simulation methods to model dephasing.

(1) *Lindblad*. We use time-discrete and completely positive Kraus maps to implement Markovian dephasing or decay for individual trajectories when more detailed modeling is not needed.

(2) *Fluctuator noise*. As a more physically realistic dephasing model, we simulate a bath of uncorrelated two-level fluctuators that are dispersively coupled to each qubit. The random switches of the fluctuators make the qubit frequencies stochastic processes with $1/f$ noise spectra, leading to nonexponential decoherence when ensemble averaged.

(3) *Markovian white noise*. For comparison to the fluctuator model, we also simulate a white noise fluctuation model where the qubit frequencies become Markovian stochastic processes, leading to exponential decoherence of Lindblad type when ensemble averaged.

We now summarize the theoretical background and implementation for each model and show numerically simulated examples.

1. Lindblad decoherence

We simulate average-sense decoherence using the “jump-no-jump” method for minimally unraveling continuous Lindblad evolution into discrete-time, completely positive maps (see, e.g., Eqs. (9) and (10) of Ref. [55]):

$$\rho(t) \mapsto \rho(t+dt) = \hat{J}_0 \rho(t) \hat{J}_0^\dagger + \sum_{k=1}^n (\gamma_k dt) \hat{J}_k \rho(t) \hat{J}_k^\dagger \quad (\text{A1})$$

with

$$\hat{J}_0 = \sqrt{\mathbb{1} - \sum_{k=1}^n (\gamma_k dt) \hat{J}_k^\dagger \hat{J}_k}. \quad (\text{A2})$$

This unraveling models rare events that occur through one of n available channels, modeled by Kraus operators $\{\hat{J}_k\}_{k=1}^n$ with rates γ_k , such that the probability of one event in a short duration dt approximates $\gamma_k dt$. The corresponding nonevent Kraus operator $\hat{J}_0$ follows from the positive operator-valued measure (POVM) normalization condition $\mathbb{1} = \hat{J}_0^2 + \sum_{k=1}^n (\gamma_k dt) \hat{J}_k^\dagger \hat{J}_k$. The total update then averages over all $n+1$ event possibilities, assuming a lack of knowledge regarding which event occurred. Using these time-discrete Kraus maps guarantees that the evolution of each time step is completely positive, so is less error prone than naively applying generic numerical integration methods to the Lindblad differential equation. The usual continuous Lindblad equation can be recovered by taking the limit

$dt \rightarrow 0$ of this discrete update, with the anticommutator term arising from nonevents.

While we emphasize fluctuator and white noise models of dephasing as the more physically realistic models for each sampled evolution trajectory, this preaveraged Lindblad decoherence model is an efficient way to model the expected ensemble-averaged dephasing evolution by applying it to each individual trajectory instead, acknowledging the absence of more detailed knowledge of the fluctuation mechanism for that trajectory. In particular, we use this crude ensemble-averaged model of the dephasing in projected forward-time estimations of the system state to compensate for signal propagation delay and maximize real-time feedback performance. For this purpose, single-qubit environmental dephasing using Eq. (A1) uses the phase-jump operator $\hat{J} = \hat{\sigma}_z/\sqrt{2}$ with rate Γ_2 . Similarly, residual ensemble dephasing due to inefficient measurement (see Appendix B2 below) uses the jump operator $\hat{J} = \hat{N}/\sqrt{2}$ with observable $\hat{N} = (\hat{\sigma}_z^{(1)} + \hat{\sigma}_z^{(2)})/2$ and rate $(1-\eta)\Gamma$, where Γ is the ensemble-averaged measurement-dephasing rate, $\eta \in [0, 1]$ is the measurement efficiency, and $\hat{N}$ is the half-parity operator being measured.

2. Non-Markovian fluctuator noise

The fluctuator approach to simulating dephasing postulates a set of N two-level systems, each of which dispersively interacts with the qubit to create classical telegraph noise. The two-level systems randomly switch between their two states in a Poisson process with characteristic rates $\{\gamma_i\}_{i=1}^N$. The rates are distributed log uniformly between minimum and maximum frequencies (γ_1, γ_N), i.e.,

$$\log \frac{\gamma_i}{\gamma_1} = \frac{i-1}{N-1} \log \frac{\gamma_N}{\gamma_1} \implies \gamma_i = \gamma_1 \left(\frac{\gamma_N}{\gamma_1} \right)^{(i-1)/(N-1)}. \quad (\text{A3})$$

Because this distribution is exponentially weighted towards low frequencies, the resulting fluctuations have nontrivial temporal correlations, leading to non-Markovian frequency noise that goes beyond simple Markovian (e.g., Lindblad) models of dephasing. This behavior better represents the experimentally observed correlations caused by two-level systems on-chip, which are a common source of dephasing in superconducting qubit systems. Simulating dephasing with fluctuator noise in this way generates a $(1/f)^\alpha$ power spectral density (PSD) [63–66], which is in qualitative agreement with experimentally measured PSD, particularly in cases where the device is charge- or flux-noise sensitive [63,67–69]. To implement the entanglement-generation experiment, in the main text we consider a flux-tunable qubit design may be necessary to engineer the desired joint measurement—as in Ref. [13]—in which case we should expect $1/f$ flux noise to be relevant.

For the simulations in the main text, we generate each stochastic telegraph noise realization with the following procedure.

(1) We define $N = 20$ frequencies $\{\gamma_i = 2\pi f_i\}$ distributed log uniformly between $(f_1, f_N) = (5 \text{ kHz}, 50 \text{ MHz})$ according to Eq. (A3).

(2) For each γ_i , we define a fluctuator to be a classical stochastic process $s_i(t)$ that takes the binary values -1 or 1 for all $t \in [0, t_f]$ up to the simulation duration t_f . Each fluctuator randomly switches between its two values as a Poisson process with characteristic switching rate γ_i , starting from an equiprobably random initial state $s_i(0) = \pm 1$.

For efficient simulation, we use the fact that the durations between random switches of a Poisson process are exponentially distributed, $T_j^{(i)} \sim \text{Exp}(\gamma_i)$, i.e., with probability density function $p(T_j^{(i)}) = \gamma_i \exp(-\gamma_i T_j^{(i)})$. Thus, we first sample the random durations between switches until the total duration, $\sum_{j=1}^N T_j^{(i)}$, exceeds the desired simulation duration t_f . We then construct the stochastic process $s_i(t)$ as a piecewise-constant function that flips sign after each duration $T_j^{(i)}$ in the sampled sequence, starting from $s_i(0)$.

(3) Finally, we construct the total stochastic process that affects each qubit frequency by averaging the effect of all $N = 20$ fluctuators. The resulting stochastic process $\chi(t) = (1/N) \sum_{i=1}^N s_i(t)$ has zero mean and a PSD with an approximate $1/f$ dependence, as verified in Fig. 5.

The simulated fluctuator envelopes modulate the qubit frequencies in the effective Hamiltonian

$$\hat{H}_{\text{fluct}}(t) = \frac{\omega}{2} \chi_1(t) \hat{\sigma}_z^1 + \frac{\omega}{2} \chi_2(t) \hat{\sigma}_z^2, \quad (\text{A4})$$

where $\chi_{1,2}(t)$ are independent fluctuator processes generated with the above procedure, while ω is a (constant) characteristic frequency that sets the strength of the dispersive frequency fluctuations.

Notably, after ensemble averaging, these frequency fluctuations produce a characteristic decay envelope that is distinct from the purely exponential decay of Markovian (Lindblad) dephasing models, as can be seen in Fig. 5, with quadratic decay at small times rather than linear decay. However, at longer times the fluctuator coherence decay more closely approximates exponential decay, with the fluctuation strength ω indirectly setting a characteristic decay rate Γ_2 that can be compared to the exponential decay rate of Markovian dephasing models. We describe how to calibrate the scaling frequency ω to achieve a desired Γ_2 in Appendix A4 below.




FIG. 5. Comparison of the temporal and spectral characteristics of white noise and fluctuator noise. (a) The white noise process (top) is stationary and self-similar at different timescales, as shown in the inset comparison, while the fluctuator noise process (bottom) is temporally correlated and thus more structured. (b) The PSD of white noise is frequency independent, while the PSD of our sampled fluctuator noise has approximately $1/f$ frequency dependence (blue fit curve) between 0.1 and 25 MHz, with roll-off into a flat spectrum at both high and low frequencies. Note that the power spectra have been smoothed with a moving window average of width $50\Delta f = 1$ MHz for visual clarity.

3. Markovian white noise

For comparison with the non-Markovian fluctuator model, we also explicitly simulate Markovian (white noise) fluctuations that produce purely exponential decay

of the ensemble-averaged coherence. Using the same Hamiltonian as in Eq. (A4), each stochastic envelope χ_k is instead sampled at discrete times $t_\ell = t_0 + \ell dt$ from a zero-mean Gaussian distribution, $\chi_k(t_\ell) \sim \text{Norm}(0, \tau/dt)$ with variance τ/dt characterized by a timescale τ and the time discretization dt . The resulting envelope functions $\chi_k(t)$ are then Markovian random processes with a piecewise-constant structure that limit to white noise processes satisfying $\langle \chi_k(t) \chi_k(t') \rangle = \tau \delta(t - t')$ in the time-continuum limit $dt \rightarrow 0$. The exponential decay rate of the ensemble averaged Bloch radius in the x - y plane is then

$$\Gamma_2 = \int_{t_0}^t \langle [\omega \chi_k(t)] [\omega \chi_k(t')] \rangle dt' = \frac{1}{2} \omega^2 \tau \implies \omega = \sqrt{\frac{2\Gamma_2}{\tau}}. \quad (\text{A5})$$

We numerically confirm this scaling of the decay rate with ω in Fig. 7(b) in the next subsection. For simplicity in the simulations, we choose $\tau = dt$ to be the same as the time step duration, making the sampled χ_k have unit variance.

4. Calibrating the fluctuation strength for a target dephasing rate

To fairly compare dephasing timescales produced by the different noise models, we numerically calibrate the dependence between the noise strength parameter ω in Hamiltonian (A4) and a rate Γ_2 that characterizes the dephasing envelope. Importantly, the strength ω is not directly measurable experimentally, while the ensemble-averaged decay profile of the coherence can be observed in a Ramsey experiment. Calibrating ω as a function of the characteristic rate Γ_2 thus better connects our simulated noise models to experimental observation. Moreover, matching effective decay rates in the absence of feedback is critically important to meaningfully compare and assess the performance of each feedback protocol in the main text.

A subtle complication with establishing a meaningful rate comparison is that, while exponential decay is memoryless with a constant decay rate Γ_2 , fluctuator noise has a temporally structured decay that is quadratic for short times and exponential at late times, so any fit to purely exponential decay is problematic. This non-Markovian behavior arises from the exponential memory kernel of the fluctuator noise time correlations and is precisely what makes it possible to dynamically decouple from short-time low-frequency fluctuator noise experimentally (e.g., using a spin-echo midgate). Nevertheless, we can still identify a characteristic rate Γ_2 that sets the time scale of the total decay profile and reduces to the exponential decay rate in the Markovian limit, which is sufficiently close for the purposes of the main text. Most importantly, we confirm that the feedback protocol fidelity results in the main




FIG. 6. (a) Ramsey decay curves for fluctuators with different choices of frequency drift amplitude ω , fit to a two-parameter stretched exponential decay, $x(t; \Gamma_2, \beta) = \exp(-(\Gamma_2 t)^\beta)$. (b) Relation between drift amplitudes ω and the dephasing rate parameter Γ_2 , which is fit to a power law used to set the target dephasing rates in the main text. The frequency reference $\gamma_{\text{geom}} = \sqrt{\gamma_1 \gamma_N} = (2\pi)500$ kHz is the geometric mean of the $N = 20$ log-distributed fluctuator frequencies between $\gamma_1/2\pi = 5$ kHz and $\gamma_N/2\pi = 50$ MHz. (c) Relation between drift amplitudes ω and the dephasing power parameter β .

text are not artifacts of any miscalibration of our fluctuator models, but are instead fundamental to the choice of time-dependent feedback.

Since we can crudely think of the fluctuator noise spectrum as a mixture of $1/f$ and white noise, we can

postulate a mean decay profile along a “stretched exponential” that interpolates between the two. To calibrate ω , we simulate Ramsey experiments as follows: we initialize ensembles of single-qubit $|+x\rangle$ states undergoing fluctuator or white noise, sweep the parameter ω over the range $[0.25, 1]/\mu\text{s}$ for fluctuator noise and $[1, 5]/\mu\text{s}$ for white noise, extract the ensemble-averaged Bloch coordinate x as a function of time, and fit the resulting decay curves to a stretched-exponential function,

$$x(t) = \exp(-(\Gamma_2 t)^\beta), \quad (\text{A6})$$

characterized by both a rate Γ_2 and a dimensionless power β . The power $\beta = 1$ corresponds to purely exponential decay, while the short-time decay anticipated for $1/f$ noise should be quadratic with $\beta = 2$, suggesting that an average fit for β will be somewhere between 1 and 2 for the full fluctuator noise decay profile. Equation (A6) is a common experimental choice for fitting nonexponential decay with a single rate parameter and fits the broad structure of our simulated decay curves reasonably well. Moreover, the extracted rates Γ_2 for nonexponential decay are also very close to the rates obtained from naïve exponential fits to the same data, which supports the use of the obtained Γ_2 as a reasonable comparison to purely exponential decay rates.

For the fluctuator noise model, the stretched exponential fits shown in Fig. 6(a) determine the relationship between ω and Γ_2 shown in Fig. 6(b), which has the polynomial fit

$$\frac{\Gamma_2}{\gamma_{\text{geom}}} = 0.12 \left(\frac{\omega}{\gamma_{\text{geom}}} \right)^{1.55} \quad (\text{fluctuator noise}). \quad (\text{A7})$$

For this fit, we choose the reference rate to be the geometric mean of the log-uniformly-spaced fluctuator frequencies, $\gamma_{\text{geom}} = \sqrt{\gamma_1 \gamma_N} = \exp[(1/N) \sum_i \log \gamma_i] = (2\pi)(500 \text{ kHz}) = \pi/\mu\text{s}$, yielding the swept parameter range $\omega/\gamma_{\text{geom}} \in [1/4\pi, 1/\pi]$. Other choices of frequency reference will scale the prefactor of the fit, but will not change the dependence on ω . The numerical results for the prefactor (0.12) and power (1.55) generally depend on other parameters that we keep constant in our simulations, like the number of fluctuators N and the number of frequency decades $d = \log_{10}(\gamma_{\text{max}}/\gamma_{\text{min}})$ spanned about the mean γ_{geom} , but these additional dependencies are not important for the results of the main text. The relationship between ω and the power β is also shown in Fig. 6(c) and is approximately linear, $\beta(\omega) = 1.07 + 0.47(\omega/\gamma_{\text{geom}})$.

For the white noise fluctuation model, the stretched exponential fits shown in Fig. 7(a) yield the relationship between ω and Γ_2 shown in Fig. 7(b), with a polynomial fit

$$\Gamma_2 dt = 0.49(\omega dt)^{1.99} \quad (\text{white noise}), \quad (\text{A8})$$

which confirms the dependence expected from Eq. (A5) with the noise timescale choice of $\tau = dt$. Similarly, the




FIG. 7. (a) Ramsey decay curves for white noise with different choices of frequency drift amplitude ω , fit to a two-parameter stretched exponential decay, $x(t; \Gamma_2, \beta) = \exp(-(\Gamma_2 t)^\beta)$. (b) Relation between drift amplitudes ω and the dephasing rate parameter Γ_2 , which is fit to a power law used to set the target dephasing rates in the main text. The white noise characteristic time $\tau = dt = 10$ ns is set to the simulation time step duration for convenience. (c) Relation between drift amplitudes ω and the dephasing power parameter β . Exponential decay corresponds to the power $\beta = 1$, confirmed by the white noise fit up to small deviations due to fit imperfections.

power β is consistent with the exponential-decay power of 1, as shown in Fig. 7(c).

These fit equations can then be inverted to obtain the conversion from Γ_2 to ω used for the simulations in the

main text:

$$\omega(\Gamma_2) = 3.9 \gamma_{\text{geom}} \left(\frac{\Gamma_2}{\gamma_{\text{geom}}} \right)^{0.65} \quad (\text{fluctuators}), \quad (\text{A9a})$$

$$\omega(\Gamma_2) = \sqrt{\frac{2\Gamma_2}{dt}} \quad (\text{white noise}). \quad (\text{A9b})$$

5. Effective dephasing rate with constant dynamical decoupling Rabi drive

In addition to calibrating the relationship between the fluctuation amplitude ω in Hamiltonian (A4) and the resulting characteristic dephasing rate Γ_2 , we also calibrate the effective dephasing rate Γ_{eff} observed in the presence of a dynamical decoupling drive. This is applicable to the “simultaneous DD and feedback” protocol in the main text, where $\Delta(\rho) = \Delta_{\text{DD}}$ is constant and acts to decouple the system from slow fluctuations. In the case of a single qubit undergoing dephasing with no feedback, a constant Rabi drive $\Delta \hat{\sigma}_y/2$ around the y axis orthogonal to the frequency fluctuation z axis $\hat{\sigma}_z$ will mitigate the ensemble dephasing by reducing the net time that the qubit coherence is affected by frequency fluctuations. The effective dephasing rate $0 < \Gamma_{\text{eff}} < \Gamma_2$ will be reduced in proportion to the ratio Δ/Γ_2 of decoupling and bare dephasing timescales. The effectiveness of such a decoupling technique depends on the short-time behavior of the ensemble dephasing, which is linear for exponential decay but has slower quadratic dependence for fluctuator-based dephasing. Thus, this simple dynamical decoupling technique is significantly more effective at mitigating coherence decay for fluctuator noise than for white noise.

Knowing the expected effective dephasing Γ_{eff} in the presence of a decoupling Rabi drive is important for our feedback protocol in the main text. The forward state estimation used to compensate for feedback delay uses an effective dephasing Lindblad term in the filter (see Appendix D below), which must use an accurate effective dephasing rate that accounts for the drive. Including this effective ensemble dephasing in the forward estimation improves the state-tracking fidelity and, as a result, the feedback stabilization performance.

Figure 8 shows the effective dephasing rate Γ_{eff} for a sweep of decoupling Rabi amplitudes Δ and bare dephasing rates Γ_2 for a single qubit. Because the qubit state vector is rotating in the x - z plane on average, the purity [i.e., $P = \text{Tr} \rho^2 = (1 + r^2)/2$ with $r^2 = x^2 + y^2 + z^2$] is used to extract a time constant. Since the drive generally decouples from the slow noise components contributing to the nonexponential decay, a single parameter fit is sufficient:

$$P(t) = \frac{1 + \exp(-2\Gamma_{\text{eff}} t)}{2}. \quad (\text{A10})$$

For $\Delta = 0$, $\Gamma_{\text{eff}} = \Gamma_2$, as expected.




FIG. 8. Ratio of the effective qubit dephasing rate Γ_{eff} to the bare rate Γ_2 in the presence of a constant decoupling Rabi drive $\Delta\hat{\sigma}_y/2$. The qubit frequency fluctuations $\omega\chi(t)\hat{\sigma}_z/2$ have the amplitude ω calibrated to produce the ensemble-dephasing rate Γ_2 in the absence of the decoupling drive, based on Figs. 6(b) and 7(b). The small-time quadratic dependence of the fluctuator-based dephasing dramatically increases the effectiveness of the decoupling drive, with the effective dephasing rate decreasing as roughly $1/\Delta$. In contrast, for white noise, the effective dephasing drops to $\Gamma_{\text{eff}} = \Gamma_2/2$ for any nonzero Δ due to the system state spending approximately half the time close to the dephasing-immune z axis.

For white noise, $\Delta/\Gamma_2 > 1/2$ yields the asymptotic limit $\Gamma_{\text{eff}}/\Gamma_2 \rightarrow 1/2$, as shown in Fig. 8 and predicted from the white noise correlator $\langle \chi_k(t)\chi_k(t') \rangle = \tau\delta(t-t')$ [70]. The intuition behind this limit is that, under a constant Rabi drive, the qubit spends about half its time near the z axis, where it is immune to $\hat{\sigma}_z$ rotations, so dephases at half the rate.

In contrast, for fluctuators, there are much larger reductions in Γ_{eff} as Δ is increased. As shown in Fig. 8, for larger $\Delta/\Gamma_2 \gg 1$, the rate Γ_{eff} decays as roughly Γ_2/Δ :

$$\Gamma_{\text{eff}}/\Gamma_2 = 0.09(\gamma_{\text{geom}}/\Gamma_2)^{0.71}(\Gamma_2/\Delta)^{1.1} \quad (\text{A11})$$

with the $(\gamma_{\text{geom}}/\Gamma_2)$ prefactor accounting for the Γ_2 -dependent intercepts on the log-log plot.

This $1/\Delta$ behavior allows Γ_{eff} to become arbitrarily small for a large enough decoupling drive. For simplicity and accuracy, we use the numerical results shown in Fig. 8 to calibrate Γ_{eff} in the state-tracking filter used in the main text.

APPENDIX B: MEASUREMENT MODEL

We use a Bayesian update model of time-stepped sequential weak measurements to describe the backaction of the joint continuous measurement on the two-qubit system. For completeness, we briefly review the essentials of this method.

In the time-continuum limit, the monitoring of an observable $\hat{N}$ with a homodyne readout will produce a normalized signal $r(t)$ that approximates a stochastic Wiener process

$$r(t)dt = \text{Tr}(\rho_{\text{sys}}(t)\hat{N})dt + \sqrt{\tau}dW \quad (\text{B1})$$

with Wiener increment $dW^2 = dt$ (assuming the Itô picture) and a characteristic measurement timescale $\tau = 1/2\Gamma\eta$ related to the ensemble measurement-dephasing rate Γ through the measurement efficiency $\eta \in [0, 1]$ (see Appendix B 2 below). This approximation is the foundation of stochastic Schrödinger equation or master equation methods for describing continuous measurement. Averaging this readout for a duration dt yields a Gaussian random variate

$$r_t = \int_{t_i}^{t_i+dt} r(t')\frac{dt'}{dt} = \text{Tr}(\rho_t\hat{N}) + \sqrt{\frac{\tau}{dt}}\xi_t \quad (\text{B2})$$

with variance τ/dt , written here in terms of its mean with additive zero-mean Gaussian noise ξ_t with unit variance, i.e., $\xi_t \sim \mathcal{N}(0, 1)$, which can be efficiently sampled numerically. The observable expectation value with the state $\rho_t = \rho_{\text{sys}}(t_i)$ approximately determines the mean value of the signal over the time bin of width dt . This relation also makes it clear that the measurement timescale τ represents the integration time required to achieve a unit signal-to-noise ratio.

In the laboratory, the finite bandwidth of the signal analyzer discretizes the time axis into bins dt in this way, making the observed readout signal more accurately a discrete readout time series $\vec{r} = \{r_1, \dots, r_i, \dots, r_n\}$ with approximately Gaussian random variates. This random time series will determine the measurement backaction on the system. In other words, for a simulation with time vector $\vec{t} = \{t_1, t_2, \dots, t_n\}$, the state evolution will be completely determined by the randomly sampled Gaussian noise $\vec{\xi} = \{\xi_1, \xi_2, \dots, \xi_n\}$ combined with the initial condition ρ_0 and full knowledge of the unitary, dissipative, and measurement evolution.

To obtain the state trajectory $\vec{\rho} = \{\rho_1, \rho_2, \dots, \rho_n\}$ from $\vec{\xi}$, we use the discrete Bayesian filter [53] presented in the main text,

$$\rho_{t+1} = \mathcal{L}_t \frac{\hat{U}_t \hat{M}_{r_t} \rho_t \hat{M}_{r_t}^\dagger \hat{U}_t^\dagger}{\text{Tr}(\hat{M}_{r_t}^\dagger \hat{M}_{r_t} \rho_t)}. \quad (\text{B3})$$

Here, $\mathcal{L}_t$ is a Kraus map representing total average dephasing (due to, e.g., the imperfect measurement efficiency discussed in Appendix B 2 below), $\hat{U}_t$ is a unitary representing the coherent evolution between times t_i and $t_i + dt$, and $\hat{M}_{r_t}$ is the measurement Kraus operator that implements the measurement backaction and depends on the sampled r_t as described below. The set of $\hat{M}_{r_t}$ over all

possible r_i forms a Gaussian POVM with completeness condition,

$$\int_{-\infty}^{\infty} dr |c(r)|^2 \hat{M}_r^\dagger \hat{M}_r = \sqrt{\frac{dt}{2\pi\tau}} \int_{-\infty}^{\infty} dr e^{-dt(r-\hat{N})^2/2\tau} = \hat{I} \quad (\text{B4})$$

with an r -dependent normalization $|c(r)|^2$ factor that cancels in the state update of Eq. (B3).

1. Derivation of the measurement operator

To derive the measurement operator specific to our setup directly from the correspondence principle, we start from the standard form of a measurement Kraus operator for $\hat{N} = (\hat{\sigma}_z^{(1)} + \hat{\sigma}_z^{(2)})/2$ that associates the eigenspace projection operators for its three eigenvalues of $-1, 0, 1$ with their corresponding conditional readout likelihoods [47,53],

$$c(r)\hat{M}_r = \sqrt{P(r|00)}|00\rangle\langle 00| + \sqrt{P(r|11)}|11\rangle\langle 11| + \sqrt{P(r|01, 10)}(|01\rangle\langle 01| + |10\rangle\langle 10|), \quad (\text{B5})$$

anticipating a state-independent factor $c(r)$ that will cancel. By construction, the POVM condition $\int dr |c(r)|^2 \hat{M}_r^\dagger \hat{M}_r = \hat{I}$ then follows from the normalization of each likelihood distribution.

For a calibrated readout r averaged over a time bin dt , the readout likelihood distributions,

$$P(r|00) = \sqrt{\frac{dt}{2\pi\tau}} \exp\left(-\frac{dt}{2\tau}(r-1)^2\right), \quad (\text{B6a})$$

$$P(r|11) = \sqrt{\frac{dt}{2\pi\tau}} \exp\left(-\frac{dt}{2\tau}(r+1)^2\right), \quad (\text{B6b})$$

$$P(r|01, 10) = \sqrt{\frac{dt}{2\pi\tau}} \exp\left(-\frac{dt}{2\tau}r^2\right), \quad (\text{B6c})$$

are each Gaussian with variance τ/dt centered at an eigenvalue of $\hat{N}$ for each eigenspace, assuming the quantum computing convention $\hat{\sigma}_z = |0\rangle\langle 0| - |1\rangle\langle 1|$. Thus,

$$\hat{M}_r = e^{(dt/4\tau)(2r-1)}|00\rangle\langle 00| + e^{-(dt/4\tau)(2r+1)}|11\rangle\langle 11| + (|01\rangle\langle 01| + |10\rangle\langle 10|), \quad (\text{B7a})$$

$$c(r) = \left(\frac{dt}{2\pi\tau}\right)^{1/4} e^{-(dt/4\tau)r^2} \quad (\text{B7b})$$

are the forms of the state-dependent Kraus operator $\hat{M}_r$ and the state-independent normalization coefficient $c(r)$ that will cancel in the state update. Written in terms of the

Bell basis, this measurement operator has the useful closed form

$$\begin{aligned} \hat{M}_r = & |\psi_+\rangle\langle\psi_+| + |\psi_-\rangle\langle\psi_-| \\ & + e^{-dt/4\tau} \left[\cosh\left(\frac{r dt}{2\tau}\right) (|\phi_+\rangle\langle\phi_+| + |\phi_-\rangle\langle\phi_-| \right. \\ & \left. + \sinh\left(\frac{r dt}{2\tau}\right) (|\phi_+\rangle\langle\phi_-| + |\phi_-\rangle\langle\phi_+|) \right] \end{aligned}$$

that depends directly on the measured operator and its square,

$$\hat{N} = \frac{\hat{\sigma}_z^{(1)} + \hat{\sigma}_z^{(2)}}{2} = |\phi_+\rangle\langle\phi_-| + |\phi_-\rangle\langle\phi_+|, \quad (\text{B8a})$$

$$\hat{N}^2 = |\phi_+\rangle\langle\phi_+| + |\phi_-\rangle\langle\phi_-|, \quad (\text{B8b})$$

$$\hat{N}^3 = \hat{N}. \quad (\text{B8c})$$

Note that this state-dependent part can also be written compactly as

$$\hat{M}_r = \exp[(dt/2\tau)(r\hat{N} - \hat{N}^2/2)], \quad (\text{B9})$$

in agreement with the expected form of the Gaussian POVM in Eq. (B4). Unlike for single-qubit observable monitoring (e.g., Ref. [47]), the r -independent quadratic term with $\hat{N}^2$ is nontrivial for the state update since it picks out the relevant projection for the subspace in which the measurement acts.

2. Inefficient measurement

In a perfectly efficient measurement, the digitized measurement record is informationally complete, so after conditioning the state on the information in the collected record no information is lost and the state remains pure. However, a real measurement apparatus is not perfectly efficient due to signal loss during transit, as well as a degradation of the remaining information in the signal from additional noise due to thermal interactions, electronics noise, or the amplification chain.

The overall measurement efficiency can be characterized with a parameter $\eta \in [0, 1]$, such that $\eta = 1$ corresponds to perfectly efficient measurement with no information loss, $\eta = 0$ corresponds to total information loss, and $0 < \eta < 1$ is a partially efficient measurement. Any loss of information increases the measurement time $\tau = 1/2\Gamma\eta$ needed to integrate the signal to achieve a unit signal-to-noise ratio. The measurement operator in Eq. (B9) thus already accounts for this loss through the parameter τ .

However, any lost information will be effectively averaged over in trajectory validation tomography, which will produce additional ensemble-average dephasing. In the limit of $\eta = 0$ with no collected signal, the best estimate one can claim is the ensemble-averaged evolution,

which, for Markovian noise, should be in accordance with a Lindblad master equation. Thus, we include the additional ensemble dephasing from information loss in each trajectory via the Lindblad map in Eq. (A1), using $\hat{J} = \hat{N}/\sqrt{2}$ as the measured operator with a residual ensemble measurement-dephasing rate $(1 - \eta)\Gamma$,

$$\mathcal{L}[\rho] = (1 - \eta)(dt\Gamma/2)(\hat{N}\rho\hat{N}) + \hat{J}_0\rho\hat{J}_0^\dagger, \quad (\text{B10a})$$

$$\hat{J}_0 = \sqrt{1 - (1 - \eta)(dt\Gamma/2)\hat{N}^2}. \quad (\text{B10b})$$

APPENDIX C: DERIVATION OF FEEDBACK PROTOCOLS

The feedback task is to apply a control drive that maximizes the fidelity to the target state $|\psi_+\rangle$, while minimizing the fidelity with the invisible state $|\psi_-\rangle$. Such a control drive will be generated by the co- and counter-rotating drive Hamiltonians defined in the main text,

$$\hat{H}_\Omega(t) = \Omega \frac{\hat{Y}_+}{2}, \quad \hat{H}_\Delta(t) = \Delta \frac{\hat{Y}_-}{2}, \quad (\text{C1})$$

where $\hat{Y}_\pm \equiv \hat{\sigma}_y^1 \pm \hat{\sigma}_y^2$ and the frequencies Ω and Δ can generally depend upon the time t , the delayed and normalized measurement record $r(t - \tau_d)$, and/or the quantum state estimate $\rho_{\text{est}}(t)$.

We use the standard (Uhlmann) definition for the state fidelity,

$$\mathcal{F}(\sigma, \rho) \equiv \text{Tr} \sqrt{\sqrt{\sigma} \rho \sqrt{\sigma}}, \quad (\text{C2})$$

which, for a pure target state $\sigma = |\psi\rangle\langle\psi|$ and estimated state $\rho = \rho_{\text{est}}$, simplifies to

$$\mathcal{F}_\psi \equiv \mathcal{F}(|\psi\rangle\langle\psi|, \rho_{\text{est}}) = \langle\psi|\rho_{\text{est}}|\psi\rangle, \quad (\text{C3})$$

and can be interpreted as the population of ρ_{est} in the target state $|\psi\rangle$.

For constant Ω and Δ , the unitary evolution generated by the control Hamiltonian $\hat{H}_\Omega + \hat{H}_\Delta$ in isolation over a time interval dt is

$$\begin{aligned} \hat{U}_{\Omega, \Delta, dt} &= e^{-i(\Omega \hat{Y}_+ + \Delta \hat{Y}_-)dt/2} \\ &= e^{-i(\Omega + \Delta)\hat{\sigma}_y/2} \otimes e^{-i(\Omega - \Delta)\hat{\sigma}_y/2} \\ &= (|\psi_+\rangle\langle\psi_+| + |\phi_-\rangle\langle\phi_-|) \cos(\Omega dt) \\ &\quad + (|\psi_+\rangle\langle\phi_-| - |\phi_-\rangle\langle\psi_+|) \sin(\Omega dt) \\ &\quad + (|\phi_+\rangle\langle\phi_+| + |\psi_-\rangle\langle\psi_-|) \cos(\Delta dt) \\ &\quad + (|\phi_+\rangle\langle\psi_-| - |\psi_-\rangle\langle\phi_+|) \sin(\Delta dt), \end{aligned}$$

where we note the following useful Bell state identities:

$$\frac{\hat{Y}_+}{2i} = |\psi_+\rangle\langle\phi_-| - |\phi_-\rangle\langle\psi_+|, \quad (\text{C4a})$$

$$\left(\frac{\hat{Y}_+}{2}\right)^2 = |\psi_+\rangle\langle\psi_+| + |\phi_-\rangle\langle\phi_-|, \quad (\text{C4b})$$

$$\left(\frac{\hat{Y}_+}{2}\right)^3 = \frac{\hat{Y}_+}{2}, \quad (\text{C4c})$$

$$\frac{\hat{Y}_-}{2i} = |\phi_+\rangle\langle\psi_-| - |\psi_-\rangle\langle\phi_+|, \quad (\text{C4d})$$

$$\left(\frac{\hat{Y}_-}{2}\right)^2 = |\phi_+\rangle\langle\phi_+| + |\psi_-\rangle\langle\psi_-|, \quad (\text{C4e})$$

$$\left(\frac{\hat{Y}_-}{2}\right)^3 = \frac{\hat{Y}_-}{2}. \quad (\text{C4f})$$

These identities imply the partition of unity

$$\begin{aligned} \hat{I} &= \left(\frac{\hat{Y}_+}{2}\right)^2 + \left(\frac{\hat{Y}_-}{2}\right)^2 \\ &= |\phi_+\rangle\langle\phi_+| + |\phi_-\rangle\langle\phi_-| + |\psi_+\rangle\langle\psi_+| + |\psi_-\rangle\langle\psi_-|. \end{aligned} \quad (\text{C5})$$

For completeness, we also note that these operators are part of two natural sets of Pauli-like two-qubit operators. The other two in the $\{|\psi_+\rangle, |\phi_-\rangle\}$ subspace are

$$\frac{\hat{X}_+}{2} = \frac{\hat{\sigma}_z^{(1)}\hat{\sigma}_x^{(2)} + \hat{\sigma}_x^{(1)}\hat{\sigma}_z^{(2)}}{2} = |\psi_+\rangle\langle\phi_-| + |\phi_-\rangle\langle\psi_+|, \quad (\text{C6a})$$

$$\frac{\hat{Z}_+}{2} = \frac{\hat{\sigma}_z^{(1)}\hat{\sigma}_z^{(2)} - \hat{\sigma}_x^{(1)}\hat{\sigma}_x^{(2)}}{2} = |\phi_-\rangle\langle\phi_-| - |\psi_+\rangle\langle\psi_+|, \quad (\text{C6b})$$

$$\left[\frac{\hat{X}_+}{2}, \frac{\hat{Y}_+}{2}\right] = 2i\frac{\hat{Z}_+}{2} \quad (\text{and cyclic perm.}), \quad (\text{C6c})$$

while the other two in the orthogonal $\{|\psi_-\rangle, |\phi_+\rangle\}$ subspace are

$$\frac{\hat{X}_-}{2} = \frac{\hat{\sigma}_z^{(1)}\hat{\sigma}_x^{(2)} - \hat{\sigma}_x^{(1)}\hat{\sigma}_z^{(2)}}{2} = |\psi_-\rangle\langle\phi_+| + |\phi_+\rangle\langle\psi_-|, \quad (\text{C7a})$$

$$\frac{\hat{Z}_-}{2} = -\frac{\hat{\sigma}_z^{(1)}\hat{\sigma}_z^{(2)} + \hat{\sigma}_x^{(1)}\hat{\sigma}_x^{(2)}}{2} = |\psi_-\rangle\langle\psi_-| - |\phi_+\rangle\langle\phi_+|, \quad (\text{C7b})$$

$$\left[\frac{\hat{X}_-}{2}, \frac{\hat{Y}_-}{2}\right] = 2i\frac{\hat{Z}_-}{2} \quad (\text{and cyclic perm.}). \quad (\text{C7c})$$

We now briefly derive two strategies for setting the control frequencies Ω and Δ so this control unitary maximizes the target fidelity.

1. Optimal feedback without measurement or dephasing

We can derive a feedback protocol if we assume a known initial estimated state ρ_{est} from which an optimal drive can be inferred. The updated state estimate after applying a control drive is $\rho'_{\text{est}} = \hat{U}_{\Omega, \Delta, \delta t} \rho_{\text{est}} \hat{U}_{\Omega, \Delta, \delta t}^\dagger$, which should be compared to the states $|\psi_{\pm}\rangle$. For a pure target state $\sigma = |\psi\rangle\langle\psi|$ and the post-control-estimated state $\rho = \rho'_{\text{est}}$, the state fidelity has the form

$$\begin{aligned} \mathcal{F}_{\psi} &\equiv \mathcal{F}(|\psi\rangle\langle\psi|, \rho'_{\text{est}}) \\ &= \langle\psi|\rho'_{\text{est}}|\psi\rangle \\ &= \langle\psi|\hat{U}_{\Omega, \Delta, \delta t} \rho_{\text{est}} \hat{U}_{\Omega, \Delta, \delta t}^\dagger |\psi\rangle. \end{aligned} \quad (\text{C8})$$

The state fidelities to $|\psi_{\pm}\rangle$ following a control drive $\hat{U}_{\Omega, \Delta, \delta t}$ are thus

$$\begin{aligned} \mathcal{F}_{\psi+}(\Omega) &= \frac{1}{2}(\rho_{\psi+} + \rho_{\phi-} + (\rho_{\psi+} - \rho_{\phi-}) \cos(2\Omega\delta t) \\ &\quad + 2\text{Re}(\rho_{\psi+, \phi-}) \sin(2\Omega\delta t)), \end{aligned} \quad (\text{C9a})$$

$$\begin{aligned} \mathcal{F}_{\psi-}(\Delta) &= \frac{1}{2}(\rho_{\psi-} + \rho_{\phi+} + (\rho_{\psi-} - \rho_{\phi+}) \cos(2\Delta\delta t) \\ &\quad - 2\text{Re}(\rho_{\psi-, \phi+}) \sin(2\Delta\delta t)), \end{aligned} \quad (\text{C9b})$$

where $\rho_{\psi\pm} = \langle\psi_{\pm}|\rho_{\text{est}}|\psi_{\pm}\rangle$, $\rho_{\phi\pm} = \langle\phi_{\pm}|\rho_{\text{est}}|\phi_{\pm}\rangle$, $\rho_{\psi+, \phi-} = \langle\psi_{+}|\rho_{\text{est}}|\phi_{-}\rangle$, and $\rho_{\phi+, \psi-} = \langle\phi_{+}|\rho_{\text{est}}|\psi_{-}\rangle$ are the relevant populations and coherences of the pre-control-estimated state ρ_{est} . Because $|\psi_{+}\rangle$ is invariant under $\hat{H}_{\Delta}$, and $|\psi_{-}\rangle$ under $\hat{H}_{\Omega}$, their respective fidelities are only affected by one of the two drive parameters, Ω or Δ .

Extremizing these fidelities using the conditions $d\mathcal{F}_{\psi+}/d\Omega = 0 = d\mathcal{F}_{\psi-}/d\Delta$ then determines the optimal state-estimate-dependent constant drives to apply over a single control cycle δt ,

$$\begin{aligned} \tan[2\Omega_{\text{opt},0}(\rho_{\text{est}})\delta t] &= 2\Omega_{\text{opt},0}(\rho_{\text{est}})\delta t + \mathcal{O}(\delta t^3) \\ &= \frac{2\text{Re}\rho_{\psi+, \phi-}}{\rho_{\psi+} - \rho_{\phi-}}, \end{aligned} \quad (\text{C10a})$$

$$\begin{aligned} \tan[2\Delta_{\text{opt}}(\rho_{\text{est}})\delta t] &= 2\Delta_{\text{opt}}(\rho_{\text{est}})\delta t + \mathcal{O}(\delta t^3) \\ &= \frac{2\text{Re}\rho_{\phi+, \psi-}}{\rho_{\phi+} - \rho_{\psi-}}. \end{aligned} \quad (\text{C10b})$$

Because the right-hand sides of both expressions are $\mathcal{O}(1)$ in δt , which itself is an integer multiple of dt , the arctangent is not expanded in calculations of $\Omega_{\text{opt},0}$, Δ_{opt} used in simulations. Note that, unlike state-estimate-independent linear feedback schemes (e.g., Ref. [47]), the drive $\Omega\hat{Y}_{+}$ inherently disrupts the target state in our protocol, which forces the dependence on a state estimate to achieve optimal feedback control. This optimal feedback drive applies when the measurement backaction is ignored or negligible. In particular, the basic rate $\Omega_{\text{opt},0}$ will acquire corrections

when accounting for the measurement properly, as detailed in the next section.

For completeness, we also note that these optima can be written in terms of the observables in Eqs. (C6) and (C7),

$$\Omega_{\text{opt},0}(\rho_{\text{est}})\delta t = \frac{1}{2} \tan^{-1}(C_{+}(\rho_{\text{est}})), \quad (\text{C11a})$$

$$\Delta_{\text{opt}}(\rho_{\text{est}})\delta t = \frac{1}{2} \tan^{-1}(C_{-}(\rho_{\text{est}})), \quad (\text{C11b})$$

where

$$C_{+}(\rho) = \frac{2\text{Re}\rho_{\psi+, \phi-}}{\rho_{\psi+} - \rho_{\phi-}} = -\frac{\langle\hat{X}_{+}\rangle}{\langle\hat{Z}_{+}\rangle}, \quad (\text{C12a})$$

$$C_{-}(\rho) = \frac{2\text{Re}\rho_{\phi+, \psi-}}{\rho_{\phi+} - \rho_{\psi-}} = -\frac{\langle\hat{X}_{-}\rangle}{\langle\hat{Z}_{-}\rangle}. \quad (\text{C12b})$$

2. Optimal feedback with measurement but no dephasing

When concurrently monitoring an observable $\hat{N} = (\hat{\sigma}_z^{(1)} + \hat{\sigma}_z^{(2)})/2$, the feedback control will generally depend upon the delayed and normalized measurement record $r(t - \tau_d)$. As a minimal feedback protocol, the corotating drive should depend on the record and a state estimate ρ_{est} ,

$$\Omega(t) = \Omega[r(t - \tau_d), \rho_{\text{est}}(t)], \quad (\text{C13})$$

while the counter-rotating drive can be kept constant or independently optimized to the $\Delta_{\text{opt}}(\rho_{\text{est}})$ in Eq. (C11b). Combining the control drive in Eq. (C4a) with the measurement operator associated with $r(t - \tau_d)$ yields a composite update operator for the measurement interaction and evolution with delay time τ_d plus a following control time step δt of $\hat{U}_{\Omega, \Delta, \delta t, r} \hat{U}_{\tau_d + \delta t} \hat{M}_r$.

The time delay τ_d complicates the analysis, since the state generally evolves between when the measurement is performed and when the feedback control is applied. Distinguishing the control cycle time step δt from the evolution time step dt results in a similar complication, since even a $\tau_d = 0$ system will undergo evolution that goes uncorrected for an interval δt . However, as an initial approximation, we can take the zero-delay limit $\tau_d \rightarrow 0$ by ignoring the intermediate evolution $\hat{U}_{\tau_d + \delta t}$ before the control, and take the controller latency to be negligible by setting $\delta t = dt$. In this zero-delay case, the fidelity to the target state $|\psi_{+}\rangle$ after a measurement-and-control time step dt has the form

$$\mathcal{F}_{\psi+} = \frac{\langle\psi_{+}|\hat{U}_{\Omega, \Delta, dt, r} \hat{M}_r \rho_{\text{est}} \hat{M}_r^\dagger \hat{U}_{\Omega, \Delta, dt, r}^\dagger |\psi_{+}\rangle}{\text{Tr}[\hat{M}_r^\dagger \hat{M}_r \rho_{\text{est}}]}. \quad (\text{C14})$$

Treating the parameter Ω as tunable and using the shorthand notation $\hat{M} = \hat{U}_{\Omega, \Delta, dt, r} \hat{M}_r$ for the composite operator, the extremization condition for this fidelity becomes

$$\frac{\partial \mathcal{F}_{\psi+}}{\partial \Omega} = 0 = 2\text{Re} \langle \phi_- | \hat{M} \rho_{\text{est}} \hat{M}^\dagger | \psi_+ \rangle. \quad (\text{C15})$$

Using the Bell state identities

$$\hat{M}^\dagger | \psi_+ \rangle = \cos(\Omega dt) | \psi_+ \rangle \sin(\Omega dt) e^{-dt/4\tau} \left[\cosh\left(\frac{rdt}{2\tau}\right) | \phi_- \rangle + \sinh\left(\frac{rdt}{2\tau}\right) | \phi_+ \rangle \right], \quad (\text{C16a})$$

$$\hat{M}^\dagger | \phi_- \rangle = -\sin(\Omega dt) | \psi_+ \rangle + \cos(\Omega dt) e^{-dt/4\tau} \left[\cosh\left(\frac{rdt}{2\tau}\right) | \phi_- \rangle + \sinh\left(\frac{rdt}{2\tau}\right) | \phi_+ \rangle \right], \quad (\text{C16b})$$

the fidelity condition evaluates to a combination of components of ρ_{est} :

$$\begin{aligned} 2\text{Re} \langle \phi_- | \hat{M} \rho_{\text{est}} \hat{M}^\dagger | \psi_+ \rangle &= -\sin(2\Omega dt) \rho_{\psi+} \\ &+ 2 \cos(2\Omega dt) e^{-dt/4\tau} \left[\cosh\left(\frac{rdt}{2\tau}\right) \text{Re} \rho_{\psi+, \phi-} + \sinh\left(\frac{rdt}{2\tau}\right) \text{Re} \rho_{\phi+, \psi+} \right] \\ &+ \sin(2\Omega dt) e^{-dt/2\tau} \left[\cosh^2\left(\frac{rdt}{2\tau}\right) \rho_{\phi-} + \sinh^2\left(\frac{rdt}{2\tau}\right) \rho_{\phi+} + \sinh\left(\frac{rdt}{\tau}\right) \text{Re} \rho_{\phi+, \phi-} \right]. \end{aligned} \quad (\text{C17})$$

Setting this to zero, we can then solve directly for the optimal drive:

$$\begin{aligned} \tan(2\Omega_{\text{opt}} dt) &= \frac{2e^{-dt/4\tau} [\cosh(rdt/2\tau) \text{Re} \rho_{\psi+, \phi-} + \sinh(rdt/2\tau) \text{Re} \rho_{\phi+, \psi+}]}{\rho_{\psi+} - e^{-dt/2\tau} [\cosh^2(rdt/2\tau) \rho_{\phi-} + \sinh^2(rdt/2\tau) \rho_{\phi+} + \sinh(rdt/\tau) \text{Re} \rho_{\phi+, \phi-}]} \\ &= \frac{2\text{Re} \rho_{\psi+, \phi-}}{\rho_{\psi+} - \rho_{\phi-}} + \frac{rdt}{2\tau} \left[\frac{2\text{Re} \rho_{\psi+, \phi+}}{\rho_{\psi+} - \rho_{\phi-}} + \frac{4\text{Re} \rho_{\psi+, \phi-} \text{Re} \rho_{\phi+, \phi-}}{(\rho_{\psi+} - \rho_{\phi-})^2} \right] \\ &\quad - \frac{dt}{4\tau} \left[\frac{\text{Re} \rho_{\psi+, \phi-}}{\rho_{\psi+} - \rho_{\phi-}} - \frac{4\text{Re} \rho_{\psi+, \phi+} \text{Re} \rho_{\phi+, \phi-} + 2\text{Re} \rho_{\psi+, \phi-} (\rho_{\phi+} - \rho_{\phi-})}{(\rho_{\psi+} - \rho_{\phi-})^2} - \frac{2\text{Re} \rho_{\psi+, \phi-} (2\text{Re} \rho_{\phi+, \phi-})^2}{(\rho_{\psi+} - \rho_{\phi-})^3} \right] \\ &\quad + \mathcal{O}(dt^{3/2}). \end{aligned} \quad (\text{C18})$$

Here the second equality uses $(rdt)^2 = \tau dt$.

Identifying the first term as the rescaled coherence $C_+(\rho)$ determining the measurement-free rate $\Omega_{\text{opt},0}$ in Eq. (C11a) and recalling that $\tau = 1/2\Gamma\eta$, this solution can be rewritten in the form

$$\Omega_{\text{opt}} dt = \frac{1}{2} \tan^{-1} \left(C_+(\rho_{\text{est}}) + 2\eta\Gamma P(\rho_{\text{est}}) rdt + \eta\Gamma Q(\rho_{\text{est}}) \frac{dt}{2} \right) \quad (\text{C19})$$

with the characteristic state-estimate-dependent factors

$$P(\rho_{\text{est}}) = \frac{\text{Re} \rho_{\psi+, \phi+}}{\rho_{\psi+} - \rho_{\phi-}} + C_+(\rho_{\text{est}}) \frac{\text{Re} \rho_{\phi+, \phi-}}{\rho_{\psi+} - \rho_{\phi-}}, \quad (\text{C20a})$$

$$Q(\rho_{\text{est}}) = -\frac{C_+(\rho_{\text{est}})}{2} + \frac{4\text{Re} \rho_{\psi+, \phi+} \text{Re} \rho_{\phi+, \phi-} + 2\text{Re} \rho_{\psi+, \phi-} (\rho_{\phi+} - \rho_{\phi-})}{(\rho_{\psi+} - \rho_{\phi-})^2} + \frac{2\text{Re} \rho_{\psi+, \phi-} (2\text{Re} \rho_{\phi+, \phi-})^2}{(\rho_{\psi+} - \rho_{\phi-})^3}. \quad (\text{C20b})$$

The dominant terms are the measurement-free rate $C_+(\rho)$, of order 1, followed by the white noise term of order $dt^{1/2}$ in $rdt = 2\text{Re} \rho_{\psi+, \phi-} dt + dW/\sqrt{2\Gamma\eta}$. Expanding the arctangent about the $\mathcal{O}(1)$ term $C_+(\rho_{\text{est}})$ and keeping terms of order $dt^{3/2}$ gives

$$\Omega_{\text{opt}} dt \approx \frac{1}{2} \tan^{-1}(C_+(\rho_{\text{est}})) + \frac{\eta\Gamma P(\rho_{\text{est}}) rdt}{1 + C_+(\rho_{\text{est}})^2} = \Omega_{\text{opt},0}(\rho_{\text{est}}) + \eta\Gamma \left(\frac{P(\rho_{\text{est}})}{1 + C_+(\rho_{\text{est}})^2} \right) rdt. \quad (\text{C21})$$

Thus, incorporating the most recent measurement information $r dt$ leads to a *linear* correction, multiplied by a state-dependent factor. In the case where the coherence ρ_{ψ_+, ϕ_-} is being consistently and successfully rotated into the ρ_{ψ_+} population, $C_+(\rho_{\text{est}}) \approx 0$ and the feedback is linear up to the state estimation for $P(\rho_{\text{est}})$:

$$\Omega_{\text{opt}} dt \approx \eta \Gamma P(\rho_{\text{est}}) r dt. \quad (\text{C22})$$

Note that $P(\rho_{\text{est}})$ can also be written in terms of the operators in Eqs. (C6), (C7), and (B8) as

$$P(\rho_{\text{est}}) = -\frac{\langle \hat{X}_+ \hat{N} + \hat{N} \hat{X}_+ \rangle}{4\langle \hat{Z}_+ \rangle} + \frac{2\langle \hat{X}_+ \rangle \langle \hat{N} \rangle}{\langle \hat{Z}_+ \rangle^2}. \quad (\text{C23})$$

This means that the two dominant terms depend mainly on the premeasurement coherence $\langle \hat{X}_+ \rangle$ and the measurement-disturbed coherence $\langle \hat{X}_+ \hat{N} + \hat{N} \hat{X}_+ \rangle / 2$, both weighted by population bias $\langle \hat{Z}_+ \rangle$ in the subspace. This further motivates thinking of Eq. (C22) as the optimal drive when the feedback is on track: given an initial condition ρ_0 with $\langle \hat{X}_+ \rangle = 0$, the drive following a measurement will just be $P(\rho_0) = -(\langle \hat{X}_+ \hat{N} + \hat{N} \hat{X}_+ \rangle) / 4\langle \hat{Z}_+ \rangle$. Since this drive is optimal, it should rotate all measurement-disturbed coherence into $|\psi_+\rangle$, thus restoring $\langle \hat{X}_+ \rangle = 0$ until the next measurement. However, the introduction of nonidealities (time delay and fluctuator or white noise) may cause the trajectory to deviate from this optimal feedback, making the $C_+(\rho_{\text{est}})$ terms critical to the feedback success. This is explored further in Appendix C 5 below.

3. Optimal feedback with measurement and dephasing

The preceding section neglected the average dephasing occurring during the time step dt prior to the control being applied. However, this dephasing can be readily included in principle. Since the specifics of the dephasing are unknown to the feedback controller, the fidelity optimization condition of Eq. (C15) should be modified to

$$\frac{\partial \mathcal{F}_{\psi_+}}{\partial \Omega} = 0 = 2\text{Re} \langle \phi_- | \mathcal{L}_{dt}(\hat{M} \rho_{\text{est}} \hat{M}^\dagger) | \psi_+ \rangle \quad (\text{C24})$$

with the average dephasing modeled by a Lindblad state update

$$\begin{aligned} \mathcal{L}_{dt}(\rho) = & \frac{dt\Gamma_2}{2} (\hat{\sigma}_z^{(1)} \rho \hat{\sigma}_z^{(1)} + \hat{\sigma}_z^{(2)} \rho \hat{\sigma}_z^{(2)}) \\ & + \frac{(1-\eta)dt\Gamma}{2} \hat{N} \rho \hat{N} + \hat{J}_0 \rho \hat{J}_0^\dagger, \end{aligned} \quad (\text{C25a})$$

$$\hat{J}_0^2 = 1 - \frac{dt\Gamma_2}{2} (2) - \frac{(1-\eta)dt\Gamma}{2} \hat{N}^2, \quad (\text{C25b})$$

which includes both the effective single-qubit dephasing and the residual measurement dephasing due to imperfect

efficiency $\eta < 1$. However, including this dephasing over the single time step dt will only add corrections of order dt to the deterministic feedback function $Q(\rho_{\text{est}})$ in Eq. (C19), which is already subdominant, so the dephasing corrections can also be safely neglected.

4. Linear feedback stochastic Schrödinger equation

Assuming ideally efficient linear feedback in the form of Eq. (C19), we can also derive an effective stochastic Schrödinger equation that describes the feedback dynamics. For simplicity, consider feedback drives of the minimal form

$$\Omega(t) = \Gamma P(\rho_{\text{est}}) r(t), \quad \Delta = \Gamma c, \quad (\text{C26})$$

with a common scaling constant Γ set to match the ensemble-average measurement-dephasing rate. For the corotating drive, the tunable state-estimate-dependent function $P(\rho_{\text{est}})$ scales the measurement record, again assuming zero time delay for convenience, while the strength of the counter-rotating drive can be relatively tuned with the dimensionless constant c .

Following the form of the update in Eq. (C14), the renormalized state after both measurement and the unitary feedback drive is

$$|\psi(t+dt)\rangle = \frac{\hat{U}_{\Omega, \Delta, dt, r} \hat{M}_r}{\sqrt{\langle \psi | \hat{M}_r^\dagger \hat{M}_r | \psi \rangle}} |\psi\rangle. \quad (\text{C27})$$

In this zero-delay case, using the white noise approximation for the readout in Eq. (B1), the form of the measurement operator in Eq. (B9), and the form of the control unitary in Eq. (C4a) yields the following composite update map to order dt , recalling that $(r dt)^2 = \tau dt + \mathcal{O}(dt^{3/2})$:

$$\begin{aligned} \frac{\hat{U}_{\Omega, \Delta, dt, r} \hat{M}_r}{\sqrt{\langle \hat{M}_r^\dagger \hat{M}_r \rangle}} \approx & \hat{I} + (r dt) \left[\frac{\hat{N} - \langle \hat{N} \rangle}{2\tau} + \Gamma P \frac{\hat{Y}_+}{2i} \right] \\ & + dt \left[\Gamma c \frac{\hat{Y}_-}{2i} + \frac{\Gamma P}{2} \frac{\hat{Y}_+}{2i} (\hat{N} - \langle \hat{N} \rangle) \right] \\ & - \frac{dt}{2} \left[\frac{\hat{N}^2 + 2\hat{N} \langle \hat{N} \rangle - 3\langle \hat{N} \rangle^2}{4\tau} \right. \\ & \left. - \tau (\Gamma P)^2 \left(\frac{\hat{Y}_+}{2i} \right)^2 \right]. \end{aligned} \quad (\text{C28})$$

After substituting the white noise relation $r dt = \langle \hat{N} \rangle dt + \sqrt{\tau} dW$ and writing the measurement time $\tau = 1/2\Gamma$ in terms of the dephasing rate Γ , assuming perfect efficiency $\eta = 1$, this yields a stochastic Schrödinger equation that

includes (instantaneous and perfectly efficient) feedback:

$$d|\psi\rangle = \sqrt{\frac{\Gamma}{2}} dW \left[(\hat{N} - \langle \hat{N} \rangle) - i \frac{P}{2} \hat{Y}_+ \right] |\psi\rangle - i \frac{\Gamma}{2} dt \left[c \hat{Y}_- + \frac{P}{2} \hat{Y}_+ (\hat{N} + \langle \hat{N} \rangle) \right] |\psi\rangle - \frac{\Gamma}{2} dt \left[\frac{(\hat{N} - \langle \hat{N} \rangle)^2}{2} + \frac{P^2}{8} \hat{Y}_+^2 \right] |\psi\rangle. \quad (\text{C29})$$

5. Correspondence of deterministic and record-dependent feedback

The optimal protocols derived in Appendices C 1 and C 2 are equivalent under a special choice of state estimate. In particular, Eq. (C14) can be rearranged as

$$\begin{aligned} \mathcal{F}_{\psi+} &= \langle \psi_+ | \hat{U}_{\Omega, \Delta, dt, r} \left(\frac{\hat{M}_r \rho_{\text{est}} \hat{M}_r^\dagger}{\text{Tr}[\hat{M}_r^\dagger \hat{M}_r \rho_{\text{est}}]} \right) \hat{U}_{\Omega, \Delta, dt, r}^\dagger | \psi_+ \rangle \\ &= \langle \psi_+ | \hat{U}_{\Omega, \Delta, dt, r} \hat{\rho}_{\text{est}}^\dagger \hat{U}_{\Omega, \Delta, dt, r} | \psi_+ \rangle, \end{aligned} \quad (\text{C30})$$

which is equivalent to using the Bayesian-updated state ρ_{est}' for the deterministic feedback in Eq. (C8). Thus, the optimal feedback that includes the linear feedback term [Eq. (C19)] is implicitly realized by applying deterministic feedback [Eq. (C11a)] on the Bayesian-updated state.

This has important practical consequences. The fact that the measurement-dependent feedback requires state tracking mitigates its potential advantage over the deterministic feedback. In a state-independent linear feedback protocol such as that of Ref. [38], feedback can be implemented in hardware, eliminating computational overhead. However, since state estimation is required with both the record-dependent and record-independent expressions we derived, it is optimal to compute the computationally simpler deterministic expression in Eq. (C8) after performing the Bayesian update. In this case, the record dependence enters the feedback implicitly through the Bayesian update.

For completeness, we compare simulations of two approaches in Fig. 9. The left column depicts deterministic feedback, which uses only the $\Omega_{\text{opt},0}$ term from Eq. (C11a). The right column depicts record-proportional feedback, which uses only the first term in Eq. (C20a), i.e.,

$$\begin{aligned} P(\rho_{\text{est}}) &= \frac{\text{Re} \rho_{\psi+, \phi+}}{\rho_{\psi+} - \rho_{\phi-}} + \frac{2 \text{Re} \rho_{\psi+, \phi-} \text{Re} \rho_{\phi+, \phi-}}{(\rho_{\psi+} - \rho_{\phi-})^2} \\ &\approx \frac{\text{Re} \rho_{\psi+, \phi+}}{\rho_{\psi+} - \rho_{\phi-}}. \end{aligned} \quad (\text{C31})$$

This is chosen on the basis that the coherence $\text{Re} \rho_{\psi+, \phi-}$ and the measured $\langle \hat{N} \rangle$ are proportional to $\sqrt{\rho_{\phi-}}$ and so become small near $|\psi_+\rangle$. Including r , the noise will dominate over deterministic terms in Eq. (C19), so $\Omega_{\text{opt}} dt \approx$




FIG. 9. Fidelity $\mathcal{F}(t)$ to target $|\psi_+\rangle$, comparing deterministic feedback $\Omega_{\text{opt},0}(\rho)$ (left) to record-proportional feedback $\Omega_{\text{opt},1}(\rho) = \eta \Gamma P(\rho) r$ (right), with $1/f$ noise $\Gamma_2 = 1/52 \mu\text{s}^{-1}$, $\Delta_{\text{DD}}/2\pi = 10 \text{ MHz}$, $\eta = 1/2$, $\tau_d = 500 \text{ ns}$, and $\Gamma = 1 \mu\text{s}^{-1}$. (a),(b) Overlaid ensemble of 100 trajectories (translucent purple) highlighting one trajectory (dark purple) and the ensemble average (black). Deterministic feedback in (a) stabilizes the ensemble mean over time, while proportional feedback in (b) with a delayed record $r(t - \tau_d)$ does not reliably stabilize all trajectories. (c),(d) Comparison of the deterministic drive $\Omega_{\text{opt},0}[\rho(t)]/\eta \Gamma r(t - \tau_d)$ scaled by the delayed record (black) and the envelope $P[\rho(t - \tau_d)]$ used in proportional feedback with no forward estimation (aqua). For deterministic feedback in (c), the drive $\Omega_{\text{opt},0}$ is computed and used to get the highlighted trajectory in (a) first, then P is computed for that observed trajectory and shows close agreement to the deterministic envelope (inset). For proportional feedback in (d), the P is computed and used to get the highlighted and stabilized trajectory in (b) first, then $\Omega_{\text{opt},0}$ is computed, but shows poor agreement.

$\eta \Gamma P[\rho] r dt$ so long as the feedback remains close to the optimal trajectory.

The deterministic feedback in Fig. 9(a) stabilizes more reliably than the proportional feedback in Fig. 9(b). This makes sense given that the deterministic feedback includes the zero-order term while also including record information implicitly, whereas the proportional feedback excludes the zero-order term. Interestingly, for the highlighted trajectory stabilized with deterministic feedback, the forward-estimated drive $\Omega_{\text{opt},0}[\rho_{\text{est}}(t)]/\eta \Gamma r(t - \tau_d)$ scaled by the delayed record in Fig. 9(c) (black) is in close agreement with the $P[\rho_{\text{est}}(t - \tau_d)]$ that would have been calculated for that trajectory *without forward estimation* (aqua). This supports the intuition that the deterministic feedback implicitly includes the record information, and also confirms that the proportional term is the dominant effect most of the time. However, for a trajectory stabilized with proportional feedback only, the used P in Fig. 9(d) (aqua) shows poor

agreement with the drive $\Omega_{\text{opt},0}/\eta\Gamma r$ (black) that would have been calculated. This suggests that the infrequent, large discrepancies between $\Omega_{\text{opt},0}$ and P in Fig. 9(c) represent nonlinear corrections that are critical to the feedback success. Without these large corrections in the proportional feedback case, the deterministic protocol always predicts a large correction to return to the more stable regime where proportional feedback becomes reliable.

APPENDIX D: IMPROVING FEEDBACK CONTROL WITH FORWARD STATE ESTIMATION

The feedback control derived in the previous section ignored the time delay τ_d between when the measurement record is collected and the control is applied. However, in practice, this time lag is necessary for information to travel between the system and measurement apparatus, so concurrent system evolution will take place during this delay that should not be neglected in the state estimate ρ_{est} being used to predict the feedback control that should be applied. Thus, the applied feedback control should anticipate what the *future state* will be, accounting for the evolution between when the measurement interaction occurred in the past to produce the collected record and when the feedback signal reaches the system.

While unitary control drives are known and may be readily accounted in such a *forward estimation* of the future state, other types of evolution like measurement and fluctuations are not known during the delay. As such, the best estimate one can make for the evolution during the delay is the *mean evolution* that averages over the unknowns and uses the information about what drives have already been output by the controller but have not yet arrived at the system.

The feedback loop time delay τ_d between when the measurement microwave field interacts with the system and when the corresponding feedback control field reaches the system can be broken into three conceptual parts,

$$\tau_d = \tau_{s \rightarrow c} + \tau_c + \tau_{c \rightarrow s}, \quad (\text{D1})$$

where $\tau_{s \rightarrow c}$ is the time of flight for the readout to make it through the measurement chain to the controller, τ_c is the time it takes for the controller to compute the feedback to apply, and $\tau_{c \rightarrow s}$ is the time of flight for the control signal to reach the system. Thus, the controller should understand the readout received at time t as $r(t - \tau_{s \rightarrow c})$, meaning it reflects information in the system state at a time $\tau_{s \rightarrow c}$ in the past, but it should issue a control signal that is appropriate for the future system state $\rho(t - \tau_{s \rightarrow c} + \tau_d) = \rho(t + \tau_c + \tau_{c \rightarrow s})$. However, this distinction is unimportant for the controller, which can operate with a shifted internal clock time $t \mapsto t + \tau_{s \rightarrow c}$ so that its internal t is matched to the time of the measurement interaction with the past system state, in which case it can assume that it receives a

signal $r(t)$ and must output a feedback signal for the future state $\rho(t + \tau_d)$ a full delay time τ_d in the future.

Other than this full delay time τ_d , the only important timescale for the feedback is the processing time of the controller τ_c . This time sets the minimum duration required to change the output feedback control signal. Whereas the readout of the system is modeled as a discrete process with time step dt according to Eq. (B2), the readout processed by the controller is buffered and averaged in larger bins of duration τ_c to produce the values $\tilde{r}(t)$. In this way, no signal is lost during the processing time, but the controller is nevertheless restricted to outputting a signal that is time stepped into bins no finer than its processing time τ_c . For the purposes of this work, we assume that this processing time is what sets the overall control time step $\delta t = \tau_c$ in the analysis.

With the shifted clock, the feedback controller should track an internal state estimate $\rho_{\text{est}}(t)$ corresponding to the actual system state $\rho_{\text{sys}}(t)$ that produced the received readout $\tilde{r}(t)$. It should then assume that the feedback signal will be applied to the appropriate *forward estimate* of the system state,

$$\rho_{\text{est}}(t + \tau_d) = \prod_{k=0}^{\lfloor \tau_d/\delta t \rfloor} \mathcal{E}_{\delta t}(t + k\delta t)[\rho_{\text{est}}(t)], \quad (\text{D2a})$$

$$\mathcal{E}_{\delta t}(t_k) = \mathcal{L}_{\delta t} \circ \mathcal{U}_{\delta t}(t_k), \quad (\text{D2b})$$

that has been evolved forward by $\lfloor \tau_d/\delta t \rfloor$ time steps δt , including both known unitary dynamics $\mathcal{U}_{\delta t}(t_k)[\cdot] = \hat{U}_{\delta t}(t_k)(\cdot)\hat{U}_{\delta t}^\dagger(t_k)$ for the drives that were already sent by the controller during the preceding τ_d (which have not yet arrived at the system) and the average decoherence $\mathcal{L}_{\delta t}$ that is anticipated during each time step δt over the delay time τ_d .

In practice, the average decoherence is relatively small if the delay τ_d is sufficiently short, so can be neglected as an initial approximation. In this case, the forward estimation is entirely determined by the sequence of control unitaries for the drives that the controller has already sent. Since the unitary generators $\hat{Y}_\pm$ commute, the controller only needs to save a running buffer sum of the emitted drives, $(\Omega(t), \Omega(t - \delta t), \dots, \Omega(t - \tau_d))$ and $(\Delta(t), \Delta(t - \delta t), \dots, \Delta(t - \tau_d))$, from which it can keep running sums Ω_{τ_d} and Δ_{τ_d} over each buffer by subtracting the oldest and adding the newest when the buffer is updated at each time step. The forward-estimation unitary operator used for simulations in the main text is then simply

$$\hat{U}_{\tau_d} = \exp(-i(\Omega_{\tau_d}\hat{Y}_+ + \Delta_{\tau_d}\hat{Y}_-)\delta t/2) \quad (\text{D3})$$

with the corresponding closed form in Eq. (C4a). This sort of buffering and calculation is straightforward and fast on an FPGA controller.

- [1] Daniel A. Lidar and Todd A. Brun, *Quantum Error Correction* (Cambridge University Press, Cambridge, 2013).
- [2] Shyam Shankar, Michael Hatridge, Zaki Leghtas, K. M. Sliwa, Aniruth Narla, Uri Vool, Steven M. Girvin, Luigi Frunzio, Mazhar Mirrahimi, and Michel H. Devoret, Autonomously stabilized entanglement between two superconducting quantum bits, *Nature* **504**, 419 (2013).
- [3] M. E. Kimchi-Schwartz, L. Martin, E. Flurin, C. Aron, M. Kulkarni, H. E. Tureci, and I. Siddiqi, Stabilizing entanglement via symmetry-selective bath engineering in superconducting qubits, *Phys. Rev. Lett.* **116**, 240503 (2016).
- [4] F. Tacchino, A. Auffèves, M. F. Santos, and D. Gerace, Steady state entanglement beyond thermal limits, *Phys. Rev. Lett.* **120**, 063604 (2018).
- [5] Sheng-li Ma, Xin-ke Li, Xin-yu Liu, Ji-kun Xie, and Fu-li Li, Stabilizing Bell states of two separated superconducting qubits via quantum reservoir engineering, *Phys. Rev. A* **99**, 042336 (2019).
- [6] Sheng-li Ma, Jing Zhang, Xin-ke Li, Ya-long Ren, Ji-kun Xie, Ming-tao Cao, and Fu-li Li, Coupling-modulation-mediated generation of stable entanglement of superconducting qubits via dissipation, *Europhys. Lett.* **135**, 63001 (2021).
- [7] Alan C. Santos and R. Bachelard, Generation of maximally entangled long-lived states with giant atoms in a waveguide, *Phys. Rev. Lett.* **130**, 053601 (2023).
- [8] Rusko Ruskov and Alexander N. Korotkov, Entanglement of solid-state qubits by measurement, *Phys. Rev. B* **67**, 241305 (2003).
- [9] Nathan S. Williams and Andrew N. Jordan, Entanglement genesis under continuous parity measurement, *Phys. Rev. A* **78**, 062322 (2008).
- [10] D. Riste, M. Dukalski, C. A. Watson, G. de Lange, M. J. Tiggelman, Ya M. Blanter, K. W. Lehnert, R. N. Schouten, and L. DiCarlo, Deterministic entanglement of superconducting qubits by parity measurement and feedback, *Nature* **502**, 350 (2013).
- [11] Joseph Kerckhoff, Hendra I. Nurdin, Dmitri S. Pavlichin, and Hideo Mabuchi, Designing quantum memories with embedded control: Photonic circuits for autonomous quantum error correction, *Phys. Rev. Lett.* **105**, 040502 (2010).
- [12] Razieh Mohseninia, Jing Yang, Irfan Siddiqi, Andrew N. Jordan, and Justin Dressel, Always-on quantum error tracking with continuous parity measurements, *Quantum* **4**, 358 (2020).
- [13] William P. Livingston, Machiel S. Blok, Emmanuel Flurin, Justin Dressel, Andrew N. Jordan, and Irfan Siddiqi, Experimental demonstration of continuous quantum error correction, *Nat. Commun.* **13**, 2307 (2022).
- [14] S. Hacohe-Gourgy, L. P. Garcia-Pintos, L. S. Martin, J. Dressel, and I. Siddiqi, Incoherent qubit control using the quantum Zeno effect, *Phys. Rev. Lett.* **120**, 020505 (2018).
- [15] E. Blumenthal, C. Mor, A. A. Diring, L. S. Martin, P. Lewalle, D. Burgarth, K. B. Whaley, and S. Hacohe-Gourgy, Demonstration of universal control between non-interacting qubits using the quantum Zeno effect, *npj Quantum Inf.* **8**, 1 (2022).
- [16] Nicolas Roch, Mollie E. Schwartz, Felix Motzoi, Christopher Macklin, Rajamani Vijay, Andrew W. Eddins, Alexander N. Korotkov, K. Birgitta Whaley, Mohan Sarovar, and Irfan Siddiqi, Observation of measurement-induced entanglement and quantum trajectories of remote superconducting qubits, *Phys. Rev. Lett.* **112**, 170501 (2014).
- [17] Felix Motzoi, K. Birgitta Whaley, and Mohan Sarovar, Continuous joint measurement and entanglement of qubits in remote cavities, *Phys. Rev. A* **92**, 032308 (2015).
- [18] Christian Dickel, J. J. Wesdorp, N. K. Langford, S. Peiter, Ramiro Sagastizabal, Alessandro Bruno, Ben Criger, F. Motzoi, and L. DiCarlo, Chip-to-chip entanglement of transmon qubits using engineered measurement fields, *Phys. Rev. B* **97**, 064508 (2018).
- [19] Joseph Kerckhoff, Luc Bouten, Andrew Silberfarb, and Hideo Mabuchi, Physical model of continuous two-qubit parity measurement in a cavity-QED network, *Phys. Rev. A* **79**, 024305 (2009).
- [20] Areeya Chantasri, Mollie E. Kimchi-Schwartz, Nicolas Roch, Irfan Siddiqi, and Andrew N. Jordan, Quantum trajectories and their statistics for remotely entangled quantum bits, *Phys. Rev. X* **6**, 041052 (2016).
- [21] S. Hacohe-Gourgy and L. S. Martin, Continuous measurements for control of superconducting quantum circuits, *Adv. Phys.: X* **5**, 1813626 (2020).
- [22] Leigh S. Martin and K. Birgitta Whaley, Single-shot deterministic entanglement between non-interacting systems with linear optics, *ArXiv:1912.00067* (2019).
- [23] Philippe Lewalle, Cyril Elouard, and Andrew N. Jordan, Entanglement-preserving limit cycles from sequential quantum measurements and feedback, *Phys. Rev. A* **102**, 062219 (2020).
- [24] H. M. Wiseman and G. J. Milburn, *Quantum Measurement and Control* (Cambridge University Press, Cambridge, NY, USA, 2010).
- [25] S. G. Schirmer and Xiaoting Wang, Stabilizing open quantum systems by Markovian reservoir engineering, *Phys. Rev. A* **81**, 062306 (2010).
- [26] Constantin Brif, Raj Chakrabarti, and Herschel Rabitz, Control of quantum phenomena: Past, present and future, *New. J. Phys.* **12**, 075008 (2010).
- [27] Jing Zhang, Yu-xi Liu, Re-Bing Wu, Kurt Jacobs, and Franco Nori, Quantum feedback: Theory, experiments, and applications, *Phys. Rep.* **679**, 1 (2017).
- [28] G. de Lange, D. Riste, M. J. Tiggelman, C. Eichler, L. Tornberg, G. Johansson, A. Wallraff, R. N. Schouten, and L. DiCarlo, Reversing quantum trajectories with analog feedback, *Phys. Rev. Lett.* **112**, 080501 (2014).
- [29] Clément Sayrin, Igor Dotsenko, Xingxing Zhou, Bruno Peaudecerf, Théo Rybarczyk, Sébastien Gleyzes, Pierre Rouchon, Mazhar Mirrahimi, Hadis Amini, Michel Brune, Jean-Michel Raimond, and Serge Haroche, Real-time quantum feedback prepares and stabilizes photon number states, *Nature* **477**, 73 (2011).
- [30] Naoki Yamamoto, Koji Tsumura, and Shinji Hara, Feedback control of quantum entanglement in a two-spin system, *Automatica* **43**, 981 (2007).
- [31] Mazhar Mirrahimi and Ramon Van Handel, Stabilizing feedback controls for quantum systems, *SIAM J. Control Optim.* **46**, 445 (2007).
- [32] Naoki Yamamoto, Hendra I. Nurdin, Matthew R. James, and Ian R. Petersen, Avoiding entanglement sudden death

- via measurement feedback control in a quantum network, *Phys. Rev. A* **78**, 042339 (2008).
- [33] Leigh Martin, Felix Motzoi, Hanhan Li, Mohan Sarovar, and K. Birgitta Whaley, Deterministic generation of remote entanglement with active quantum feedback, *Phys. Rev. A* **92**, 062321 (2015).
- [34] L. C. G. Govia, A. Lingenfelter, and A. A. Clerk, Stabilizing two-qubit entanglement by mimicking a squeezed environment, *Phys. Rev. Res.* **4**, 023010 (2022).
- [35] Tristan Brown, Emery Doucet, Diego Ristè, Guilhem Ribeill, Katarina Cicak, Joe Aumentado, Ray Simmonds, Luke Govia, Archana Kamal, and Leonardo Ranzani, Trade off-free entanglement stabilization in a superconducting qutrit-qubit system, *Nat. Commun.* **13**, 3994 (2022).
- [36] E. Doucet, L. C. G. Govia, and A. Kamal, Scalable entanglement stabilization with modular reservoir engineering, Preprint [ArXiv:2301.05725](https://arxiv.org/abs/2301.05725) (2023).
- [37] Felix Motzoi, Eli Halperin, Xiaoting Wang, K. Birgitta Whaley, and Sophie Schirmer, Backaction-driven, robust, steady-state long-distance qubit entanglement over lossy channels, *Phys. Rev. A* **94**, 032313 (2016).
- [38] R. Vijay, C. Macklin, D. H. Slichter, S. J. Weber, K. W. Murch, R. Naik, A. N. Korotkov, and I. Siddiqi, Stabilizing Rabi oscillations in a superconducting qubit using quantum feedback, *Nature* **490**, 77 (2012).
- [39] Bibek Pokharel, Namit Anand, Benjamin Fortman, and Daniel A. Lidar, Demonstration of fidelity improvement using dynamical decoupling with superconducting qubits, *Phys. Rev. Lett.* **121**, 220502 (2018).
- [40] Vinay Tripathi, Huo Chen, Mostafa Khezri, Ka-Wa Yip, E. M. Levenson-Falk, and Daniel A. Lidar, Suppression of crosstalk in superconducting qubits using dynamical decoupling, *Phys. Rev. Appl.* **18**, 024068 (2022).
- [41] Bibek Pokharel and Daniel A. Lidar, Demonstration of algorithmic quantum speedup, *Phys. Rev. Lett.* **130**, 210602 (2023).
- [42] Kaveh Khodjasteh, Tamas Erdelyi, and Lorenza Viola, Limits on preserving quantum coherence using multipulse control, *Phys. Rev. A* **83**, 020305 (2011).
- [43] Kenji Kashima and Naoki Yamamoto, Control of quantum systems despite feedback delay, *IEEE Trans. Automat. Contr.* **54**, 876 (2009).
- [44] Yanan Liu, Daoyi Dong, Sen Kuang, Ian R. Petersen, and Hidehiro Yonezawa, Two-step feedback preparation of entanglement for qubit systems with time delay, *Automatica* **125**, 109174 (2021).
- [45] Hadis Amini, Ram A. Somaraju, Igor Dotsenko, Clément Sayrin, Mazyar Mirrahimi, and Pierre Rouchon, Feedback stabilization of discrete-time quantum systems subject to non-demolition measurements with imperfections and delays, *Automatica* **49**, 2683 (2013).
- [46] Charlene Ahn, Andrew C. Doherty, and Andrew J. Landahl, Continuous quantum error correction via quantum feedback control, *Phys. Rev. A* **65**, 042301 (2002).
- [47] Taylor Lee Patti, Areeya Chantasri, Luis Pedro García-Pintos, Andrew N. Jordan, and Justin Dressel, Linear feedback stabilization of a dispersively monitored qubit, *Phys. Rev. A* **96**, 022311 (2017).
- [48] Daniel A. Lidar, in *Quantum Information and Computation for Chemistry* (John Wiley & Sons, Ltd, Hoboken, NJ), p. 295.
- [49] Soumya Singha Roy, T. S. Mahesh, and G. S. Agarwal, Storing entanglement of nuclear spins via Uhrig dynamical decoupling, *Phys. Rev. A* **83**, 062326 (2011).
- [50] R. L. Franco, Antonio D'Arrigo, Giuseppe Falci, Giuseppe Compagno, and Elisabetta Paladino, Preserving entanglement and nonlocality in solid-state qubits by dynamical decoupling, *Phys. Rev. B* **90**, 054304 (2014).
- [51] Trevor McCourt, Charles Neill, Kenny Lee, Chris Quintana, Yu Chen, Julian Kelly, Jeffrey Marshall, V. N. Smelyanskiy, M. I. Dykman, Alexander Korotkov, and Et Al, Learning noise via dynamical decoupling of entangled qubits, *Phys. Rev. A* **107**, 052610 (2023).
- [52] E. Paladino, Y. M. Galperin, G. Falci, and B. L. Altshuler, $1/f$ noise: Implications for solid-state quantum information, *Rev. Mod. Phys.* **86**, 361 (2014).
- [53] Alexander N. Korotkov, Quantum Bayesian approach to circuit QED measurement, [ArXiv:1111.4016](https://arxiv.org/abs/1111.4016) (2011).
- [54] S. S. Szigeti, S. J. Adlong, M. R. Hush, A. R. R. Carvalho, and J. J. Hope, Robustness of system-filter separation for the feedback control of a quantum harmonic oscillator undergoing continuous position measurement, *Phys. Rev. A* **87**, 013626 (2013).
- [55] Mostafa Khezri, Justin Dressel, and Alexander N. Korotkov, Qubit measurement error from coupling with a detuned neighbor in circuit QED, *Phys. Rev. A* **92**, 052306 (2015).
- [56] A. Eddins, J. M. Kreikebaum, D. M. Toyli, E. M. Levenson-Falk, A. Dove, W. P. Livingston, B. A. Levitan, L. C. G. Govia, A. A. Clerk, and I. Siddiqi, High-efficiency measurement of an artificial atom embedded in a parametric amplifier, *Phys. Rev. X* **9**, 011004 (2019).
- [57] F. Lecocq, L. Ranzani, G. A. Peterson, K. Cicak, X. Y. Jin, R. W. Simmonds, J. D. Teufel, and J. Aumentado, Efficient qubit measurement with a nonreciprocal microwave amplifier, *Phys. Rev. Lett.* **126**, 020502 (2021).
- [58] Alexander P. M. Place, *et al.*, New material platform for superconducting transmon qubits with coherence times exceeding 0.3 milliseconds, *Nat. Commun.* **12**, 1779 (2021).
- [59] Goren Gordon, Gershon Kurizki, and Daniel A. Lidar, Optimal dynamical decoherence control of a qubit, *Phys. Rev. Lett.* **101**, 010403 (2008).
- [60] Fei Yan, Philip Krantz, Youngkyu Sung, Morten Kjaergaard, Daniel L. Campbell, Terry P. Orlando, Simon Gustavsson, and William D. Oliver, Tunable coupling scheme for implementing high-fidelity two-qubit gates, *Phys. Rev. Appl.* **10**, 054062 (2018).
- [61] Riccardo Porotti, Antoine Essig, Benjamin Huard, and Florian Marquardt, Deep reinforcement learning for quantum state preparation with weak nonlinear measurements, *Quantum* **6**, 747 (2022).
- [62] Riccardo Porotti, Vittorio Peano, and Florian Marquardt, Gradient-ascent pulse engineering with feedback, *PRX Quantum* **4**, 030305 (2023).

- [63] E. Paladino, Y. M. Galperin, G. Falci, and B. L. Altshuler, $1/f$ noise: Implications for solid-state quantum information, *Rev. Mod. Phys.* **86**, 361 (2014).
- [64] Huo Chen and Daniel A. Lidar, HOQST: Hamiltonian open quantum system toolkit, *Commun. Phys.* **5**, 112 (2022).
- [65] Vinay Tripathi, Huo Chen, Eli Levenson-Falk, and Daniel A. Lidar, Modeling low- and high-frequency noise in transmon qubits with resource-efficient measurement, *ArXiv:2303.00095* (2023).
- [66] Ka Wa Yip, Open-system modeling of quantum annealing: Theory and applications, *ArXiv:2107.07231* (2021).
- [67] Jonas Bylander, Simon Gustavsson, Fei Yan, Fumiki Yoshihara, Khalil Harrabi, George Fitch, David G. Cory, Yasunobu Nakamura, Jaw-Shen Tsai, and William D. Oliver, Noise spectroscopy through dynamical decoupling with a superconducting flux qubit, *Nat. Phys.* **7**, 565 (2011).
- [68] R. Trappen, *et al.*, Decoherence of a tunable capacitively shunted flux qubit, *ArXiv:2307.13961* (2023).
- [69] F. Yoshihara, K. Harrabi, A. O. Niskanen, Y. Nakamura, and J. S. Tsai, Decoherence of flux qubits due to $1/f$ flux noise, *Phys. Rev. Lett.* **97**, 167001 (2006).
- [70] Daniel Lidar, *Open Quantum Systems* (unpublished).