Pontryagin-optimal control of a non-Hermitian qubit

Philippe Lewalle and K. Birgitta Whaley

Pontryagin–Optimal Control of a non-Hermitian Qubit

Philippe Lewalle[*] and K. Birgitta Whaley

Department of Chemistry, University of California, Berkeley, CA 94720, USA and
Berkeley Center for Quantum Information and Computation, Berkeley, CA 94720, USA
(Dated: January 10, 2023)

Open–system quantum dynamics described by non-Hermitian effective Hamiltonians have become a subject of considerable interest. Studies of non-Hermitian physics have revealed general principles, including relationships between the topology of the complex eigenvalue space and the breakdown of adiabatic control strategies. We study here the control of a single non-Hermitian qubit, similar to recently realized experimental systems in which the non-Hermiticity arises from an open spontaneous emission channel. We review the topological features of the resulting non-Hermitian Hamiltonian and then present two distinct results. First, we illustrate how to realize any continuous and differentiable pure–state trajectory in the dynamics of a qubit that are conditioned on no emission. This result implicitly provides a workaround for the breakdown of standard adiabatic following in such non-Hermitian systems. Second, we use Pontryagin's maximum principle to derive optimal trajectories connecting boundary states on the Bloch sphere, using a cost function which balances the desired dynamics against the controller energy used to realize them. We demonstrate that the latter approach can effectively find trajectories which maintain high state purity even in the case of inefficient detection.

I. INTRODUCTION

The dynamics of non-Hermitian (NH) systems have been an area of growing theoretical and experimental interest over the past two decades [1–8]. NH Hamiltonians arise naturally in the context of many open (non-conservative) systems. In contrast with Hermitian Hamiltonians, NH Hamiltonians may generically have complex eigenvalues and bi-orthogonal left and right eigenvectors. A number of unique effects can be understood in terms of the Riemann sheet topology of complex eigenvalues in parameter space; this topology is defined foremost by the presence of exceptional points (EPs), where there is a convergence of both the complex eigenvalues and eigenvectors of the NH Hamiltonian [9, 10]. In particular, parameter loops that encircle EPs in NH systems may lead to dynamics exhibiting gain/loss effects that can break principles of adiabatic following [11–17] and manifest chiral behavior [10, 18–24]. There is now considerable literature about non-Hermitian systems in general, emphasizing both topological aspects of complex non-Hermitian spectra that may be understood statically [10, 23], and investigations of the corresponding dynamics [25–29].

One particular realization of NH physics arises in the conditional dynamics of open quantum systems [30–36]. Explorations of such systems have illustrated the applicability of many of the concepts from the classical context, and raised fundamentally new issues and potential applications. In this quantum context, two approaches to studying NH systems are apparent: Topological properties of the NH Hamiltonian or Liouvillian [32, 33, 37] have been the subject of considerable interest, while the open and conditional quantum state dynamics underly-

ing those NH properties may be directly understood via quantum trajectory theory [38–42]. In the limit of slowly varying control parameters, the adiabatic principle suggests a clear connection between the system dynamics and NH topology. However, since the adiabatic principle is of limited applicability in NH systems [12], this slow limit does not offer a complete correspondence between the two, and connection of an NH system's dynamics to the underlying topology is not straightforward. For this fundamental reason and for practical considerations as well, the development of NH quantum control protocols that act effectively on shorter timescales is an important challenge for the field. To this end, we note there have been recent proposals employing perturbative approaches [43], NH shortcuts to adiabaticity [44], as well as investigation of experimentally–realizable periodic controls [36].

We propose here an alternative approach, based on use of the Pontryagin maximum principle, which is ubiquitous both the classical quantum control theory literature [45–48]. We show that this can be usefully applied to the control of NH quantum systems, investigating in particular the behavior of Pontryagin control for a relatively simple example of a single qubit with an effective NH Hamiltonian arising from its spontaneous emission channel (see Fig. 1). This NH qubit system is similar to that studied in recent experiments [31, 34–36], and as we show in this work, it proves to be a good testing ground for coherent control in the context of a NH quantum system with EPs.

The rest of the paper is organized as follows. We review the system dynamics, define the non-Hermitian Hamiltonian, and review its topological properties in Sec. II. In Sec. III, we take advantage of the relative simplicity of our chosen system, and demonstrate that it is straightforward to derive the controller paths needed to force the system along any continuous and differentiable pure state trajectory. In Sec. IV we formulate a cost func-

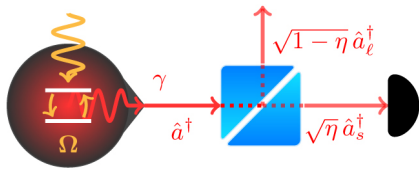FIG. 1. We sketch a possible realization of the scheme we consider. Unitary operations characterized by the Rabi rates $\boldsymbol{\Omega}$ are applied to a qubit. An effective non-Hermitian Hamiltonian arises due to the qubit's spontaneous emission. The dynamics conditioned on emission or non-emission of a photon are accessible if that emission is captured by a cavity or transmission line whose output is monitored by a photo–counter. Detection with efficiency η can be modeled by imagining that a beamsplitter between the system and detector causes some photons to be lost.

tion, and then apply Pontryagin's principle to derive the corresponding optimal trajectories the NH qubit can follow to connect boundary states over a desired evolution time. In particular, we illustrate that by using Rabi rotations as a control knob, arbitrary qubit state manipulations can be implemented under the NH dynamics, under both ideal (unit efficiency) detection of, or post–selection on, photon emission. Furthermore, decoherence under partially–conditional dynamics (i.e., those arising from inefficient detection of emitted photons) can be substantially mitigated. in Sec. V we provide a discussion of our results and potential future research based on the ideas we develop in this work.

II. NON-HERMITIAN JUMPLESS QUBIT DYNAMICS

We now consider the dynamics of a qubit that can spontaneously emit a photon into its environment with some probability, but *does not actually emit*. These dynamics are accessible, e.g., if the qubit emits into a read-out cavity and transmission line that route any photons to a photodetector. The dynamics of interest are then those which occur between any detector clicks, i.e. we focus on those dynamics which are inferred with knowledge that no photon was emitted through the open decay channel. These dynamics are thus conditioned on the outcomes of the photon counting measurements. This can also be realized in practice through post–selection using a third level [31]. The required continuous monitoring of the decay channel by photon counting [42] is illustrated in Fig. 1, and described in greater detail in Appendix A. We include the measurement efficiency η, which allows us to consider imperfect photon counters, as well as to interpolate between the case where the detector is present and when it is absent.

The conditional evolution of the qubit density matrix ρ is given by a modified Lindblad master equation

$$\dot{\rho} = i\,\rho\,\hat{\mathcal{H}}^\dagger - i\,\hat{\mathcal{H}}\,\rho + (1-\eta)\hat{L}\,\rho\,\hat{L}^\dagger + \eta\,\rho\,\mathrm{tr}\left(\hat{L}\,\rho\,\hat{L}^\dagger\right),\ (1)$$

where $\hat{L} = \sqrt{\gamma}\,\hat{\sigma}_-$ represents decay of the qubit into its environment at rate γ (or, equivalently, with characteristic time $T_1 = \gamma^{-1}$), $\eta \in [0,1]$ is the photon detection efficiency, and

$$\begin{aligned}\hat{\mathcal{H}} &= \tfrac{1}{2}\left(\Omega_x\,\hat{\sigma}_x + \Omega_y\,\hat{\sigma}_y + \Omega_z\,\hat{\sigma}_z\right) - \tfrac{i}{2}\,\gamma\,\hat{\sigma}^+\hat{\sigma}_- \\ &= \frac{1}{2}\begin{pmatrix} \Omega_z - i\,\gamma & \Omega_x - i\,\Omega_y \\ \Omega_x + i\,\Omega_y & -\Omega_z \end{pmatrix}\end{aligned} \quad (2)$$

is a non-Hermitian effective Hamiltonian that includes both unitary controls $\hat{H} = \tfrac{1}{2}\boldsymbol{\Omega}\cdot\hat{\sigma}_{\mathbf{q}}$ and NH decay dynamics. Eq. (1) gives the standard Lindblad master equation for $\eta = 0$ (no detection). We shall be most focused on the no–click decay dynamics accessible with ideal detection ($\eta = 1$), since the NH effective Hamiltonian $\hat{\mathcal{H}}$ is most closely associated with this ideal case. The dynamics (1) can equivalently be written as a dynamical system in the Bloch coordinates $\mathbf{q} = (x, y, z)^\top$ according to $\dot{\mathbf{q}} = \mathrm{tr}\,(\hat{\sigma}_{\mathbf{q}}\,\dot{\rho})$, where $\hat{\sigma}_{\mathbf{q}}$ are Pauli matrices. We will often refer to these dynamics using the shorthand notation $\dot{\mathbf{q}} = \boldsymbol{\mathcal{F}}(\mathbf{q}, \boldsymbol{\Omega})$, noting that we may decompose the total dynamics into a sum of the jumpless NH dynamics ($\boldsymbol{\Omega} = 0$), plus the unitary part ($\boldsymbol{\Omega} \neq 0$), according to

$$\dot{\mathbf{q}} = \underbrace{\boldsymbol{\mathcal{F}}(\mathbf{q}, \boldsymbol{\Omega} = 0)}_{\boldsymbol{\mathcal{F}}_0} + \boldsymbol{\Omega} \times \mathbf{q} = \boldsymbol{\mathcal{F}}(\mathbf{q}, \boldsymbol{\Omega}). \quad (3)$$

The right eigenvalues of $\hat{\mathcal{H}}$ read

$$\lambda_\pm = -\tfrac{i}{4}\gamma \pm \tfrac{1}{4}\sqrt{4\boldsymbol{\Omega}^2 - \gamma^2 - 4i\,\gamma\,\Omega_z} = -\tfrac{i}{4}\gamma \pm \tfrac{1}{4}\sqrt{\mathcal{J}}, \ (4)$$

with corresponding right eigenvectors

$$|\lambda_\pm\rangle = \mathcal{N}\left\{\left(2\Omega_z - i\gamma \pm \sqrt{\mathcal{J}}\right)|e\rangle + 2(\Omega_x + i\Omega_y)|g\rangle\right\} \quad (5)$$

(where $\mathcal{N}$ is a normalization factor). The left eigenvalues are the complex conjugates of the right eigenvalues. While $|\lambda_+\rangle$ and $|\lambda_-\rangle$ are no longer necessarily orthogonal to one another, the left and right eigenvectors do obey a bi-orthogonality relation [8]. A pair of Exceptional Points (EPs) appear as the solutions with $\mathcal{J} = 0$, satisfying $\Omega_z = 0$ and $4(\Omega_x^2 + \Omega_y^2) = \gamma^2$. The EPs are located on the equator of the Bloch sphere, and manifest as fixed points there when the Rabi drive $\boldsymbol{\Omega}$ tries to excite the qubit at a rate that exactly balances the conditional decay dynamics γ (e.g. at $z = 0$ and $x = 1$, $\dot{z} = 0$ is achieved by $\Omega_y = -\gamma/2$, which lies on the ring of EPs defined by $\Omega_x^2 + \Omega_y^2 = \gamma^2/4$). In the typical quantum–optical or circuit QED situations that are well–suited to realizing our physical situation, it is much more straightforward to modify $\boldsymbol{\Omega}$ precisely and quickly than to change γ. We will consequently proceed by assuming γ to be fixed, while treating $\boldsymbol{\Omega}$ as a dynamic control knob. Notice that the eigenvalues $\lambda_\pm$ have a rotational symmetry in Ω_x and Ω_y; we may consequently define $\tilde{\Omega}$ to be a drive vector in the equatorial xy Bloch–plane, and then understand the complex eigenspectrum by visualizing it as a function of $\tilde{\Omega}$ and Ω_z. Fig. 2 shows the real and imaginary parts of $\lambda_\pm$ in this parameter space.
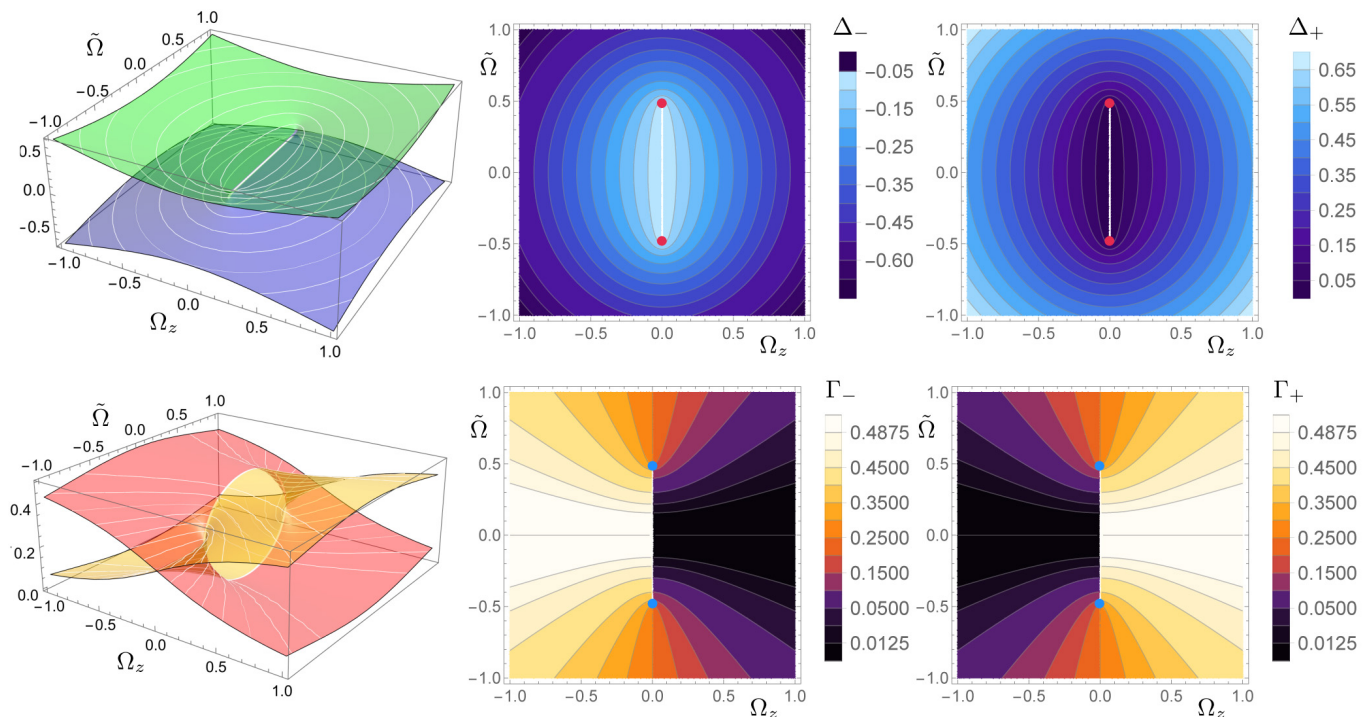
FIG. 2. We plot the real parts of the right eigenvalues $\lambda_\pm$ of $\hat{\mathcal{H}}$ (i.e. $\Delta_\pm$, top row), and the imaginary parts of $\lambda_\pm$ (i.e. $\Gamma_\pm$, bottom row). We fix γ, and all other rates (Ω, Δ, and Γ) are here expressed in units of γ. The resulting pair of Riemann sheets appear together on the left, and then the real and imaginary parts of each individual sheet are shown on the right. We mark the pairs of EPs with markers ●, and choose to place our branch cut between the two EPs along the $\Omega_z = 0$ line. Notice that a closed parameter loop enclosing one or both EPs ($n = 1, 2$) will force n exchange(s) of which gain/loss Γ value is greater; this is a topological property of the complex spectrum, that is associated with the failure of adiabatic control protocols when encircling EPs. This system differs from that of Ref. [31] only in minor details.

Dynamics with a complex eigenspectrum differ qualitatively from those of Hermitian systems with real eigenvalues. Even a naïve writing of the (un-normalized and non-trace–preserving) evolution $e^{-i\hat{\mathcal{H}}t} \sim \exp[-i\lambda_\pm t] = \exp[-\Gamma_\pm t - i\Delta_\pm t]$ reveals that while the real parts of $\lambda_\pm$ lead to effective eigenenergies $\Delta\pm \equiv \Im[i\lambda_\pm]$, the imaginary parts $\Gamma_\pm \equiv \Re[i\lambda_\pm]$ behave as effective gain/loss parameters. Their significance becomes clear when considering the possibility of adiabatic control (e.g., by fixing γ and then modifying the NH eigenstates by varying $\boldsymbol{\Omega}(t)$ slowly compared to γ). One finds that quasi-adiabatic following is possible, but only of the stable right eigenstate (where the more stable one at any given time is defined by having the smaller loss parameter) [12]. In our current example we generically will have one right eigenstate closer to $|e\rangle$ and the other closer to $|g\rangle$ (this is true except at the EP); the right eigenstate closer to $|g\rangle$ will be the stable one, at the expense of the other. Adiabatic following will be broken along any control trajectory that exchanges the stability of the two eigenstates. Such a stability exchange is topologically guaranteed when performing closed–loop encircling of an EP [10]. As such, the difference between the relative losses $\Gamma_\pm$ along a parameter trajectory concisely explain many of the NH phenomena reported in the lit-

erature, including breaking adiabaticity [12], and chiral state exchange [10, 18–21, 23]. In our present example this means that adiabatic following works only in the lower half of the Bloch sphere (closer to $|g\rangle$), and will fail on a timescale T_1 if we try to switch hemispheres.

The system (1) we have introduced is one of the simplest possible NH quantum systems one could devise. It is nevertheless surprisingly rich, containing spectral and dynamical features unique to NH systems, as well as being experimentally feasible [31]. This system is consequently an ideal sandbox for investigating NH physics and quantum control strategies [11, 44].

III. BASIC CONTROL: FORCING THE QUBIT STATE ALONG A PRESET TRAJECTORY

Above we have described two distinct pictures of the same system: On the one hand, we have a qubit with a complex eigenvalue spectrum containing exception points, and exhibiting a variety of behaviors that can be explained using the bi-orthogonal NH eigenstates. On the other hand, quantum trajectory theory tells us that the system dynamics may be straightforwardly expressed by the dynamical system (1) or (3) across timescales; this

picture implicitly contains features of the NH topology, but can be applied without careful analysis of that complementary topological NH structure [42].

In order to illustrate the power of the latter viewpoint, we show that we may straightforwardly force our qubit to follow any pure state trajectory we wish, assuming perfect conditional evolution (obtained when the measurement efficiency $\eta = 1$). In this case, our state $\mathbf{q}$ will be a unit vector pointing from the origin towards a point on the surface of the Bloch sphere, and $\mathbf{\Omega}$ will necessarily be an orthogonal vector generating rotations on the surface via the cross product in (3). We have dynamics $\dot{\mathbf{q}} = \mathcal{F}(\mathbf{q}, \mathbf{\Omega} = 0) + \mathbf{\Omega} \times \mathbf{q}$; suppose we now solve for the drive $\mathbf{\Omega}$ that forces the qubit to follow a desired pure state target trajectory $\mathbf{Q}(t)$ on the surface. We assume that $\mathbf{Q}(t)$ is continuous and differentiable, such that $\dot{\mathbf{Q}}$ exists over the entire time domain of interest. For such a pure state trajectory, $\mathbf{q}$ will point to the surface of the Bloch sphere, while $\dot{\mathbf{Q}}$ will be a tangent vector along the surface of the sphere; it follows that $\mathbf{\Omega} = \Omega\bar{\mathbf{\Omega}}$ must be a mutually–orthogonal vector pointing along the direction

$$\bar{\mathbf{\Omega}} = \left.\frac{\mathbf{q} \times (\dot{\mathbf{Q}} - \mathcal{F}_0)}{|\mathbf{q} \times (\dot{\mathbf{Q}} - \mathcal{F}_0)|}\right|_{\mathbf{q} \to \mathbf{Q}}, \qquad (6)$$

where we use the overbar to denote a unit vector, $\Omega = |\mathbf{\Omega}|$ is the magnitude, and $\mathcal{F}_0$ is a shorthand for the conditional dynamics due to decay only. This expression suggests that we may implement a controlled unitary characterized by $\mathbf{\Omega}$ that "makes up the difference" between the conditional decay dynamics and the target trajectory. The magnitude can be obtained by rearranging $\dot{\mathbf{q}} = \mathcal{F}_0 + \mathbf{\Omega} \times \mathbf{q} = \dot{\mathbf{Q}}$, and using the fact that $\mathbf{q}$, $\dot{\mathbf{Q}} - \mathcal{F}_0$, and $\mathbf{\Omega}$ are all mutually orthogonal by construction. In summary then, we may drive the system along an arbitrary pure–state trajectory $\mathbf{Q}(t)$ by applying the drive

$$\mathbf{\Omega} = \mathbf{Q} \times (\dot{\mathbf{Q}} - \mathcal{F}_0(\mathbf{Q})) \qquad (7)$$

for the ideal $\eta = 1$ case. An example of these dynamics is shown in Fig. 3. Note that (7) generalizes to $\mathbf{\Omega} = \mathbf{q} \times (\dot{\mathbf{Q}} - \mathcal{F}_0)/|\mathbf{q}|^2$ for the more general mixed state dynamics where $\mathbf{q}(t)$ is not confined to the surface of the Bloch sphere and thus does not necessarily follow the target trajectory $\mathbf{Q}(t)$ at all times. Note also that pinning the qubit to a desired pure state emerges naturally from this analysis as a simple sub-case of the dynamics under perfect measurement efficiency $\eta = 1$.

Even though the example in Fig. 3 includes a switch between hemispheres of the Bloch sphere, this trajectory does not necessarily imply a parameter loop that encloses an EP. We illustrate this in Fig. 3(d), with a set of trajectories corresponding to dynamics at differing speeds along the specific spatial trajectory shown in panel (a). It is evident that EP enclosure depends here on the speed at which we follow the path, with encircling occurring only for the slower loops. Nevertheless, these different parameter trajectories all result in identical dynamics on the Bloch sphere over their respective time intervals (i.e.,
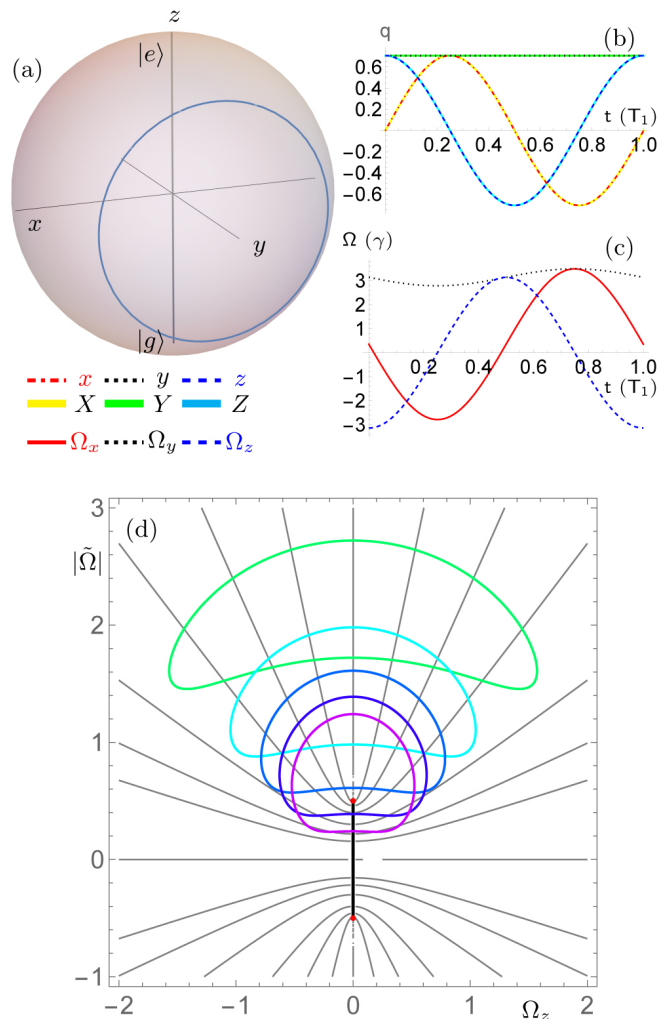

FIG. 3. We illustrate a sample target trajectory $\mathbf{Q} = (X, Y, Z)^\top = \frac{1}{\sqrt{2}}(\sin(2\pi\,t/T), 1, \cos(2\pi\,t/T))^\top$ on the surface of the Bloch sphere in (a). We are able to exactly follow this contrived target trajectory (b) via a smooth controller trajectory (c), both of which are evaluated here for final time $T = 1\,[T_1]$. In (d) we plot the contours of Γ (gray), with the branch cut in black and EPs in red. The $\mathbf{\Omega}(t)$ trajectories for $T = 2\,[T_1]$ (green), $T = 3\,[T_1]$ (cyan), $T = 4\,[T_1]$ (blue), $T = 5\,[T_1]$ (purple), and $T = 6\,[T_1]$ (magenta), are superposed over the contour plot. Note that to keep the figure two–dimensional, we summarize azimuthal unitary motion by the magnitude $|\tilde{\Omega}| = \sqrt{\Omega_x^2 + \Omega_y^2}$. We may understand the oscillation of Ω_y in (c) as correctly compensating for the changes in the gain loss ratio required to maintain the desired trajectory.

the path shown in panel (a)). While the EP and its attendant topology characterize the dynamics in the adiabatic limit (variations in $\mathbf{\Omega}$ much slower than T_1), other dynamical effects obscure those topological features when there is not an adiabatic–like separation of timescales.

Finally, we remark that the solutions (7) include the capability to realize something like a "shortcut to adiabaticity", even in the case where actual adiabatic follow-

ing fails, in this sense: First, we may define any parameter trajectory $\boldsymbol{\Omega}_0$ and then find the Bloch trajectory $\mathbf{Q}_0(t)$ which corresponds to the motion of a NH eigenstate under $\boldsymbol{\Omega}_0$. We can then use (7) to compute a new parameter trajectory $\boldsymbol{\Omega}_1$ which drives the qubit along $\mathbf{Q}_0(t)$ over any timescale.[1] Similarly, one could also use (7) to accelerate the NH dynamics that the system exhibits in response to a slowly–varying $\boldsymbol{\Omega}$ (slow enough to have adiabatic motion plus modifications primarily due to the NH topology described in the previous section).

IV. PONTRYAGIN OPTIMAL CONTROL: COST FUNCTIONS AND EXTREMIZATION

We have demonstrated above that we can easily force our system to follow particular dynamics. While we may then turn around and inquire about the topological properties of those solutions relative to the NH Hamiltonian, nothing in the method requires this. We now change emphasis, and consider controls that follow optimal dynamics between arbitrary boundary states with respect to some cost function, which we derive via the Pontryagin action extremization method. In other words, instead of specifying a target trajectory $\mathbf{Q}(t)$, we will now only specify an initial $\mathbf{q}_i$ and a final target state $\mathbf{Q}_f$, and derive optimal dynamics which connect them.

Pontryagin optimal control [47, 48] assumes that one has 1) some dynamical equations $\dot{\mathbf{q}} = \boldsymbol{\mathcal{F}}(\mathbf{q}, \boldsymbol{\Omega})$ encoding the system dynamics and their relationship to any available control knobs (here the Rabi drive $\boldsymbol{\Omega}$ takes on the latter role), and 2) some cost function $\mathscr{S}$ resembling an action functional. The aim is then to derive controller trajectories that extremize the cost $\mathscr{S}$ while also satisfying the dynamical constraint $\boldsymbol{\mathcal{F}}$. As we have a dynamical equation (3), we need only specify the cost against which we want to optimize. Generically, we wish to find controller trajectories $\boldsymbol{\Omega}(t)$ that take us rapidly from some initial state $\mathbf{q}_i$ to a chosen final state $\mathbf{Q}_f$, while using our drive power $\boldsymbol{\Omega}^2$ efficiently. We may consequently write a cost function of the form

$$\mathscr{S} = \int_0^T dt \, (\boldsymbol{\Lambda} \cdot \dot{\mathbf{q}} - \mathscr{H}), \tag{8a}$$

where the Hamiltonian

$$\mathscr{H} = \boldsymbol{\Lambda} \cdot \boldsymbol{\mathcal{F}}(\mathbf{q}, \boldsymbol{\Omega}) - \tfrac{\alpha}{2}(\mathbf{q} - \mathbf{Q}_f)^2 - \tfrac{\beta}{2}\boldsymbol{\Omega}^2 \tag{8b}$$

includes the Lagrange multipliers (co-states) $\boldsymbol{\Lambda}$ that constrain solutions to the dynamics of interest, as well as cost conditions for being far away from the target final

state and for using high drive power. The co-states are not physical; they enter only as a computational tool to derive the Pontryagin–optimized dynamics. We remark that under no circumstances should the object $\mathscr{H}$ used for optimization be confused with the quantum operators $\hat{H} = \hat{H}^\dagger$ or $\hat{\mathcal{H}} \neq \hat{\mathcal{H}}^\dagger$. The coefficients α and β are introduced in (8b) so that we may tune the relative weights of the different cost terms that we have chosen to include.

Pontryagin's principle then manifests as action extremization, i.e., taking $\delta\mathscr{S} = 0$ leads to

$$\partial_{\boldsymbol{\Lambda}}\mathscr{H} = \dot{\mathbf{q}} = \boldsymbol{\mathcal{F}}, \tag{9a}$$

$$\partial_{\mathbf{q}}\mathscr{H} = -\dot{\boldsymbol{\Lambda}}, \tag{9b}$$

$$\partial_{\boldsymbol{\Omega}}\mathscr{H}|_{\boldsymbol{\Omega}=\boldsymbol{\Omega}^\star} = 0, \tag{9c}$$

which are necessary conditions for optimality. We have Hamilton's equations of motion, plus a condition for deriving optimal controller trajectories $\boldsymbol{\Omega}^\star(t)$. Details of the optimal dynamics appear in Appendix B. It is straightforward to solve this last equation and find that optimal controller trajectories obey $\boldsymbol{\Omega}^\star = \frac{1}{\beta}\mathbf{q} \times \boldsymbol{\Lambda}$. Plugging in this optimal Rabi drive $\boldsymbol{\Omega}$ then yields the Hamiltonian

$$\mathscr{H}^\star = \tfrac{1}{2\beta}\left\{\mathbf{q}^2\,\boldsymbol{\Lambda}^2 - (\mathbf{q}\cdot\boldsymbol{\Lambda})^2\right\} + \boldsymbol{\Lambda}\cdot\boldsymbol{\mathcal{F}}_0 - \tfrac{\alpha}{2}(\mathbf{q}-\mathbf{Q}_f)^2. \tag{10}$$

The solutions of Hamilton's dynamical equations generated by $\mathscr{H}^\star$, initialized for a particular initial $\mathbf{q}_i$ and all possible $\boldsymbol{\Lambda}_i$, form a Lagrangian Manifold (LM) of candidate solutions [49]; if a globally optimal solution for given final boundary conditions exists, it can be found in this solution manifold. Note that for constant α, β, and γ, an optimal solution necessarily conserves its Pontryagin energy $\mathscr{E} = \mathscr{H}$. The main challenge in solving the optimal dynamics eq. (9) is in finding solutions that match particular boundary conditions $\mathbf{q}_i$ and $\mathbf{q}(T)$. That boundary value problem in the coordinate space can equivalently be formulated as an initial value problem in terms of the initial values of the states $\mathbf{q}_i$ and the co-states $\boldsymbol{\Lambda}_i$. We shall describe solutions to (9) on a Lagrangian Manifold (LM) that is defined from a subset of initial co-states $\boldsymbol{\Lambda}_i$ for a given initial $\mathbf{q}_i$. Specifically, the LM containing all $\boldsymbol{\Lambda}_i$ contains all possible optimal solutions originating at the given $\mathbf{q}_i$. This LM can be understood as a tool to translate between the initial value formulation of the control problem and the boundary value formulation, because there will always be at least one initial co-state $\boldsymbol{\Lambda}_i$ that generates a path on the LM leading to any attainable $\mathbf{q}(T)$. The LM thus provides a way of understanding the mapping between the initial co-states and attainable final states at later times. This tool has previously been used in a similar way to describe quantum trajectories following an optimal readout [50–52], and is discussed further in Appendix B.

The values of the coefficients α and β do impact the types of solutions that we can get. In Appendix B we

---

[1] Since the NH eigenstate depends on $\boldsymbol{\Omega}$, we will not necessarily be following the real–time NH eigenstate of $\boldsymbol{\Omega}_1$; we will just be following the NH eigenstate path that was defined by the other parameter path $\boldsymbol{\Omega}_0$. This will be the case for any control additions which modify the eigenstates of the Hamiltonian.

show that α effectively determines how fast the co-states $\mathbf{\Lambda}$ change as a function of the distance from the target state. The form of the optimal Rabi drive clearly indicates that β effectively scales how the co-states $\mathbf{\Lambda}$ translate to a controller trajectory, i.e., β just re-scales the strength of the controlled unitary dynamics relative to the decay dynamics of $\mathcal{F}_0$. In less precise but more suggestive terms, a larger value of α allows the controlled state trajectory to accelerate faster, while a large β tempers the impact of erratic co-state dynamics on the actual controller dynamics. We shall consequently occasionally refer to α and β as the controller's "agility" and "temperance", respectively, with an understanding that these parameters can partially counterbalance each other in practice. For simplicity we will typically set the temperance $\beta = 1$ in the examples below. This is because a constraint on the maximum available $\mathscr{E}$ plays a similar role and we necessarily already have to limit that energy in numerical simulations by choosing a finite volume of initial co-states $\mathbf{\Lambda}$.

IV.1. Pure State Control

We demonstrate here that it is possible to solve the control problem framed above numerically in the case $\eta = 1$, and thereby obtain controller and dynamical solutions which map arbitrary initial qubit states $\mathbf{q}_i$ to arbitrary final qubit states $\mathbf{Q}_f$ on a desired timeframe, (i.e., in a time T) within the context of the NH dynamics (1). Fig. 4 illustrates two pertinent examples. Panel (a) shows that we can drive a transition from $|g\rangle$ to $|e\rangle$ against decay, and panel (b) shows that we can drive a transition between two orthogonal states on the equator of the Bloch sphere. Both transitions can be driven on timescales faster than the average spontaneous emission time γ^{-1} with modest drive power. We furthermore find that for both initial conditions the associated LM covers the entire Bloch sphere "on its way" to the target state, indicating that for sufficiently large $\mathbf{\Omega}^2$ or $|\mathbf{\Omega}|$, i.e., for sufficiently fast Rabi driving, there are no forbidden pure states that our control scheme cannot access.

Assessing the examples of Fig. 3 and Fig. 4, and the underlying methods, we arrive at the following conclusions. Despite the apparent complexity of the present system when viewed through the lens of a non-Hermitian Hamiltonian, there are actually no substantial barriers impeding pure–state control of the ideal ($\eta = 1$) system state if we approach it instead as a dynamical system. Particularly, we have in this case been able to solve for the drive that forces the system along an exact desired pure state trajectory, and further shown that we may formulate reasonable cost functions against which to derive optimized state dynamics and the controls that generate them. We have done this with a relatively simple cost function, but nothing prevents us from investigating other cost functions that are motivated by more specific tasks.
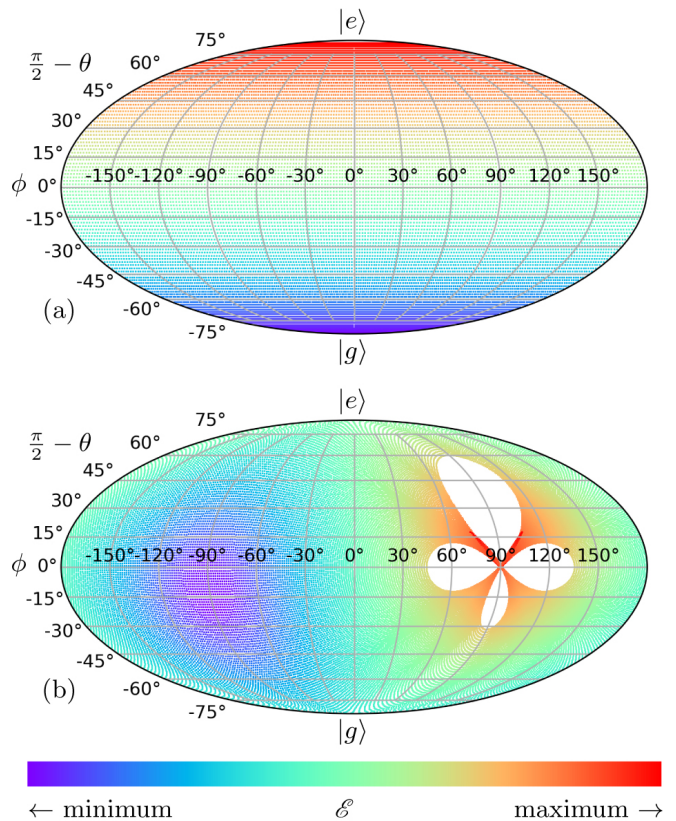

FIG. 4. We plot samplings of the Lagrangian Manifold (LM) in two different situations. Each point in a sampled LM represents the state at time T that is generated by a different initial co-state $\mathbf{\Lambda}_i$, such that the resulting plot of the sampled LM shows the extent of the possible optimal solutions on their way to the target state $\mathbf{Q}_f$. (a) Plot of the LM after time $T = 2/3\,[\gamma]$, initialized at $|g\rangle$, with $\eta = 1$, $\alpha = 2$, $\beta = 1$, and target state $|e\rangle$. Panel (a) implies that if an unwanted jump to $|g\rangle$ were detected, there exist optimal solutions to reset the state in the next interval of no-jump evolution over a relatively short time and using modest control drive parameters. (b) Plot of the LM after $T = 1/2\,[\gamma]$, from a state initialized at $y = -1$, with dynamics targeting the orthogonal state $y = +1$. In both cases we are able to drive the transition of interest in reasonable times and find complete coverage of the Bloch sphere soon after the time shown, indicating that there is a controller path reaching every pure final state under the conditions shown (even on the way to a different target state). Smaller values of $\mathscr{E}$ (bluer regions) generically correspond to control trajectories with smaller $|\mathbf{\Lambda}|$, i.e., trajectories with low $\mathscr{E}$ tend to be the slower ones that will reach a target state over longer time intervals. Conversely, larger $\mathscr{E}$ values (redder regions) correspond to the fastest control paths that we allow, sitting on the leading edge of the manifold. The LM is bounded by a range of initial conjugate momenta (co-states) expressed in spherical coordinates, specifically $\Lambda_\theta(0) \in [-5, 5]$ and $\Lambda_\phi(0) \in [-5, 5]$ for both panels above. Expansion of this initial $\mathbf{\Lambda}$–volume would reveal faster coverage of the Bloch sphere and solutions over shorter time intervals (at the expense of higher Rabi power). Animations of the evolving LMs appear in the supplemental materials, and are described in Appendix C. The temporal context of the snapshots above is best understood via such animations.
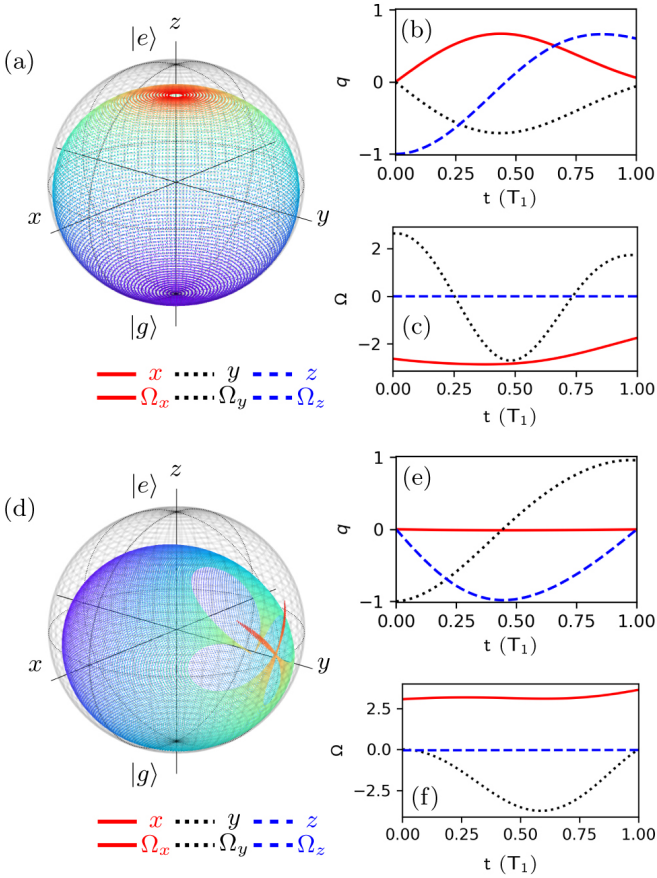
FIG. 5. Panels (a) and (d) repeat the manifold integrations of Fig. 4(a,b), respectively, using all the same parameters *except* that we now have inefficient conditioning, $\eta = 0.5$. Mixing due to the now partially–monitored decay is evident primarily close to $|e\rangle$, making pure states in the vicinity of $|e\rangle$ the most difficult to reach with modest values of $|\Omega|$. Panels (b) and (c) show the optimally–controlled trajectory which comes closest to the target state $|e\rangle$ at $T = 1$ $[T_1]$, and the controller trajectory $\Omega(t)$ generating it. It is evident that this trajectory struggles to terminate at $z \gtrsim 0.5$, highlighting the difficulty in reaching the excited state when inefficient photodetection is used. We will show in Fig. 6 below that this effect can be mitigated by allowing for higher Pontryagin energy and greater controller agility α. Panels (e) and (f) show the state trajectory and controller trajectory moving from $|-y\rangle$ to $|y\rangle$ over a duration $T = T_1$. Panel (e) shows a relatively high–fidelity operation, in stark contrast with panel (b). This is because we select here boundary conditions which can be reached by a trajectory which spends much of its time near $|g\rangle$, where the purity loss due to decay is relatively small. Further details appear in Appendix B. Animations of the evolving LMs appear in the supplemental materials, and are described in Appendix C.

IV.2. Mixed State Control and Inefficient Conditioning

We continue by considering the less straightforward case with inefficient photon detection, $\eta < 1$. Inef-
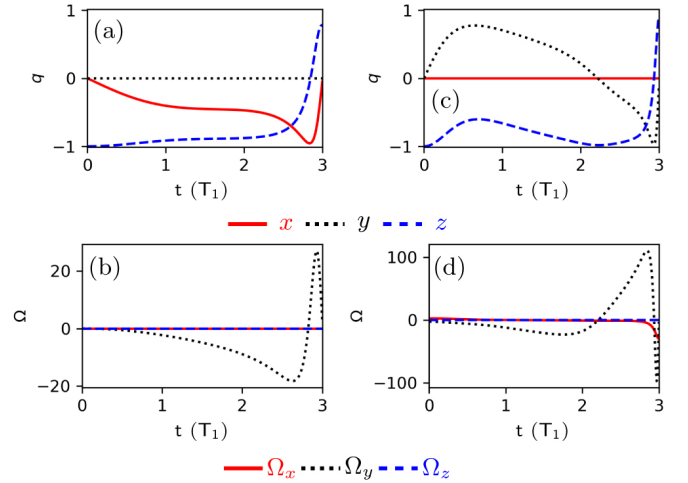

FIG. 6. We repeat Fig. 5(b,c), this time with an expanded LM $\Lambda_\theta \in [-10, 10]$ and $\Lambda_\phi \in [-10, 10]$, completely unconditional dynamics $\eta = 0$, and a longer duration $t = 3\,T_1$. Panels (a,b) use $\alpha = 2$ in (a,b), and panels (c,d) increase the agility to $\alpha = 10$. (The temperance is $\beta = 1$ in all cases). It is evident that despite the significantly worsened (now non-existent) detection, and longer evolution time compared with Fig. 5, we can actually get closer to $|e\rangle$. By allowing the controller use of greater power and the latitude to make more rapid adjustments, the optimization scheme is then able to keep a pure state near $|g\rangle$ until accelerating rapidly towards $|e\rangle$ at the end. Thus, the larger initial manifold and α values allow our controller to find trajectories in which decoherence is meaningfully mitigated. Further details appear in Appendix B, and in Fig. 7.

ficient detection or post–selection implies a loss of information that dissipates into environmental degrees of freedom without being detected and generically leads to mixed quantum states [42]. This presents a clear challenge for control [53, 54]. In particular, many high–purity qubit states may not be accessible to moderate–time control trajectories, because unitary controls cannot directly "undo" decoherence caused by inefficient detection (i.e., an observer's control abilities are limited if they possess only incomplete knowledge about the system of interest). This is apparent in Fig. 5, where we repeat the examples of Fig. 4 but with an imperfect detector efficiency $\eta = 0.5$. A decay channel in particular will tend to mix the qubit state over short times if we are near $|e\rangle$, but will then re-purify the qubit over timescales much longer than T_1 without a drive (i.e., $|g\rangle$ is a pure state, and the system will either jump there or asymptotically approach it if left alone). We can consequently begin to qualitatively understand the route towards reaching high–purity target states with poor detection efficiency. Here, in addition to elements of the ideal analysis above, there ought to be an implicit benefit to control trajectories which go through the ground state (or near it), so as to minimize decoherence over the course of the evolution.

To that end, an ideal trajectory to reach $|e\rangle$, which is the most difficult state to reach, is one that stays near

$|g\rangle$ to retain purity for a relatively long time during its evolution, and then rapidly accelerates towards $|e\rangle$ at the last possible moment. We elaborate on this point in Appendix B. It is natural to ask whether our optimization scheme can actually come up with a solution for the trajectory just described. The answer is that it can, and such a solution is explicitly shown in Fig. 6. The coefficient α turns out to play a meaningful role at this point: Increasing the "agility" allows the controller to change direction faster (see Appendix B), which is helpful in minimizing decoherence. In particular, Fig. 6 shows that the difficulties in performing the operation $|g\rangle \to |e\rangle$ illustrated in Fig. 5(a,b,c) can largely be overcome by increasing α.

V. DISCUSSION AND OUTLOOK

We have demonstrated here an instance where optimal control theory may be successfully applied to a non-Hermitian quantum system. We have focused on a relatively simple system (a decaying qubit) whose dynamics can be completely understood via quantum trajectory theory (see [42] and references therein), and whose pure–state dynamics are straightforwardly controllable. This has allowed us to analyze the conceptual aspects of non-Hermitian physics relative to our control schemes from a strong foundation. We have used this system to review a number of phenomena that are generic to non-Hermitian Hamiltonians in quantum mechanics (exceptional points, bi-orthogonal eigenstates, and so on), and then demonstrated that complementary and powerful control methods arise straightforwardly when the same system is approached from a quantum trajectory perspective instead.

Our results here constitute a particular type of feedback control [55]. The construction of a non-Hermitian Hamiltonian is conditioned on the readout (i.e., on the absence of photon emission events), but retains many of the desirable features of open–loop control between emission events. In particular, controls for extended times between jumps may be pre-computed. We might consequently characterize our work above as an exploration of "open–loop conditional quantum control". Despite the importance of the measurement record, we do not require feedback to be computed instantaneously, as occurs in many feedback problems requiring non-differentiable controller trajectories that respond to diffusive conditional dynamics at each time interval [40, 41, 54, 56].

We further note that conditional evolution of an open quantum system often admits a description in terms of a NH Hamiltonian, such that general features of our approach may apply to a wide class of related problems. In particular we have developed a pathway towards mitigating the impacts of inefficient detection in this context, with demonstrated efficacy in a simple example. Our present work further suggests novel research avenues to consider for control of more complex non-Hermitian quantum systems, by probing the intersection between dynamical and topological pictures of the system.

Much of the present interest in NH physics revolves around topological aspects of the phase space [23, 25, 57]. Such properties may be evaluated with respect to NH eigenvalue Riemann sheets as parameters are varied (as in, for example, Fig. 2), but connecting these ideas to the actual dynamical response of the system across arbitrary timescales and in general remains a challenge for the field. We view our application of Pontryagin's principle to a NH quantum system as one of many emerging strategies for the control of NH physics, which seek to avoid the limitations of adiabatic following in this context. The approach presented here is complementary to other efforts addressing non-Hermitian quantum control, including experimental work [36], and theoretical studies based on perturbation theories [43] or shortcuts to adiabaticity [44, 58–61]. Further development of a Pontryagin control approach for NH systems, and unification of topological NH concepts with a dynamical picture due to quantum trajectory theory, are promising avenues for continued research. Ultimately, we see this work as one element of a toolbox which may, together with other methods, lead to comprehensive control of non-Hermitian quantum systems that takes advantage of their unique physics for tasks of interest to quantum information science.

Appendix A: No–Jump Qubit Dynamics with a Decay Channel

We derive here our non-Hermitian Hamiltonian and the associated conditional dynamics, using a Bayesian description of continuously–monitored sponataneous emission (see Refs. [42, 62] and references therein). Physically, we assume that a qubit emits into a readout cavity (i.e., a fast decaying cavity) and/or transmission line, where the values of the qubit transition frequency and the density of environmental states result in emission at a rate γ (or timescale $T_1 = 1/\gamma$), with the characteristic exponential decay of the excited state population (i.e., we assume that the types of approximations first made by Weisskopf and Wigner [63] apply). In this setting, emitted photons can be captured with high efficiency and then be routed to a detection device, such that the conditional evolution is accessible. See Fig. 1.

It is then phenomenologically appropriate to consider a separable initial state of a qubit and output line leading to a detector, namely $(\zeta |e\rangle + \upsilon |g\rangle) \otimes |0\rangle$, evolving to

$$\zeta \left(\sqrt{e^{-\gamma \, \Delta t}} |e, 0\rangle + \sqrt{1 - e^{-\gamma \, \Delta t}} |g, 1\rangle \right) + \upsilon |g, 0\rangle \quad \text{(A1a)}$$

after a time Δt. This can equivalently be written as

$$
\begin{pmatrix} \sqrt{e^{-\gamma\,\Delta t}} & 0 \\ \sqrt{1 - e^{-\gamma\,\Delta t}}\,\hat{a}^\dagger & 1 \end{pmatrix} \begin{pmatrix} \zeta \\ \upsilon \end{pmatrix} \otimes |0\rangle, \qquad (A1b)
$$

where $\hat{a}^\dagger |0\rangle = |1\rangle$ represents the emission of a photon, and the matrix form highlights the action on the qubit state. Kraus operators can be obtained by selecting a final state of the line, which we will associate with a measurement outcome, leaving behind an operation on the qubit state only. We introduce a notion of loss between qubit and detector via insertion of a beamsplitter relation $\hat{a}^\dagger \to \sqrt{\eta}\,\hat{a}_s^\dagger + \sqrt{1-\eta}\,\hat{a}_\ell^\dagger$, which states that an emitted photon may go to a monitored signal port $\hat{a}_s^\dagger$ with probability η, or may be lost in transit with probability $1 - \eta$. We will henceforth refer to η as the measurement efficiency. This process generalizes (A1b) to

$$
\underbrace{\begin{pmatrix} \sqrt{e^{-\gamma\,\Delta t}} & 0 \\ \sqrt{1 - e^{-\gamma\,\Delta t}}\,(\sqrt{\eta}\,\hat{a}_s^\dagger + \sqrt{1-\eta}\,\hat{a}_\ell^\dagger) & 1 \end{pmatrix}}_{\mathcal{M}} \begin{pmatrix} \zeta \\ \upsilon \end{pmatrix} \otimes |00\rangle.
$$

$$(A2)$$

Note that the existing similar experiments emphasizing NH dynamics use a third qudit level and post–selection [31, 34–36] instead of continuous photo–counting, for practical reasons. This is however conceptually the same as the picture laid out here, since post–selection on no jumps having occurred to a third level at some time is equivalent to conditioning on an emission event not having occurred in any of the timesteps leading up to that moment.

We now adapt these expressions to describe the conditional evolution of the qubit. Three distinct outcomes are possible for a measurement over a finite timestep Δt, each represented by a Kraus operator: The detector may click due to photon emission, as represented by the operator

$$
\begin{aligned}
\hat{\mathcal{M}}_{10} = \langle 1_s\,0_\ell|\,\mathcal{M}\,|0\,0\rangle &= \begin{pmatrix} 0 & 0 \\ \sqrt{\eta(1 - e^{-\gamma\,\Delta t})} & 0 \end{pmatrix} \\
&\approx \sqrt{\Delta t}\,\sqrt{\eta\,\gamma}\,\hat{\sigma}_- + \mathcal{O}\left(\Delta t^{\frac{3}{2}}\right),
\end{aligned} \quad (A3a)
$$

or the detector may not click, which is due to either no photon being emitted

$$
\begin{aligned}
\hat{\mathcal{M}}_{00} = \langle 0_s\,0_\ell|\,\mathcal{M}\,|0\,0\rangle &= \begin{pmatrix} \sqrt{e^{-\gamma\,\Delta t}} & 0 \\ 0 & 1 \end{pmatrix} \\
&\approx \hat{\mathbb{I}} + \Delta t \underbrace{\begin{pmatrix} -\gamma/2 & 0 \\ 0 & 0 \end{pmatrix}}_{\hat{\mathfrak{Z}}} + \mathcal{O}\left(\Delta t^2\right),
\end{aligned} \quad (A3b)
$$

or to a photon being emitted but going unobserved due to imperfect detection

$$
\begin{aligned}
\hat{\mathcal{M}}_{01} = \langle 0_s\,1_\ell|\,\mathcal{M}\,|0\,0\rangle &= \begin{pmatrix} 0 & 0 \\ \sqrt{(1-\eta)(1 - e^{-\gamma\,\Delta t})} & 0 \end{pmatrix}, \\
&\approx \sqrt{\Delta t}\,\sqrt{(1 - \eta)\,\gamma}\,\hat{\sigma}_- + \mathcal{O}\left(\Delta t^{\frac{3}{2}}\right).
\end{aligned}
$$

$$(A3c)$$

We define $\hat{L} \equiv \sqrt{\gamma}\,\sigma_-$, such that $\hat{\mathcal{M}}_{10} \approx \sqrt{\Delta t}\,\sqrt{\eta}\,\hat{L}$ and $\hat{\mathcal{M}}_{01} \approx \sqrt{\Delta t}\,\sqrt{1-\eta}\,\hat{L}$. The operators above represent a complete set of outcomes, in that together they complete a positive operator valued measure (POVM), satisfying

$$
\hat{\mathcal{M}}_{00}^\dagger \hat{\mathcal{M}}_{00} + \hat{\mathcal{M}}_{01}^\dagger \hat{\mathcal{M}}_{01} + \hat{\mathcal{M}}_{10}^\dagger \hat{\mathcal{M}}_{10} = \hat{\mathbb{I}}. \qquad (A4)
$$

Notice that (A4) holds both for the general forms of these matrices and for the operators approximated to $\mathcal{O}(\Delta t)$. The latter further implies that $\hat{\mathfrak{Z}} = -\frac{1}{2}\hat{L}^\dagger\hat{L}$, which is an expected and general property of such jump operators.

We are now in a position to describe the conditional qubit dynamics (i.e., the dynamics of the qubit that can be inferred from the "click" or "no–click" outcome of the photodetector monitoring the spontaneous emission channel with efficiency η). In the event of a click, we infer a jump of the qubit density matrix ρ to $|g\rangle\langle g|$, according to

$$
\rho(t + \Delta t) = \frac{\hat{\mathcal{M}}_{10}\,\rho(t)\,\hat{\mathcal{M}}_{10}^\dagger}{\mathrm{tr}\left(\hat{\mathcal{M}}_{10}\,\rho(t)\,\hat{\mathcal{M}}_{10}^\dagger\right)} = |g\rangle\langle g|. \qquad (A5)
$$

Note that after detecting a jump the final state $|g\rangle$ is obtained irrespective of our estimate of the state prior to the detector click. However, we are much more interested in the dynamics that arise when the detector does not click; these dynamics come from the state update

$$
\rho(t + \Delta t) = \frac{\hat{\mathcal{M}}_{00}\,\rho(t)\,\hat{\mathcal{M}}_{00}^\dagger + \hat{\mathcal{M}}_{01}\,\rho(t)\,\hat{\mathcal{M}}_{01}^\dagger}{\mathrm{tr}\left(\hat{\mathcal{M}}_{00}\,\rho(t)\,\hat{\mathcal{M}}_{00}^\dagger + \hat{\mathcal{M}}_{01}\,\rho(t)\,\hat{\mathcal{M}}_{01}^\dagger\right)}, \quad (A6)
$$

which includes a weighted average over the the two sub-processes that lead to the no-click outcome. We may take this expression, and re-write it as a dynamical equation for the density matrix by making expansions to $\mathcal{O}(\Delta t)$. In the notation above, we find

$$
\begin{aligned}
\dot{\rho} &= \hat{\mathfrak{Z}}\,\rho + \rho\,\hat{\mathfrak{Z}}^\dagger + (1-\eta)\hat{L}\,\rho\,\hat{L}^\dagger - \rho\,\mathrm{tr}\left(\hat{\mathfrak{Z}}\,\rho + \rho\,\hat{\mathfrak{Z}}^\dagger + (1-\eta)\hat{L}\,\rho\,\hat{L}^\dagger\right) \\
&= \underbrace{(1-\eta)\hat{L}\,\rho\,\hat{L}^\dagger - \tfrac{1}{2}\hat{L}^\dagger\hat{L}\,\rho - \tfrac{1}{2}\rho\,\hat{L}^\dagger\hat{L}}_{\text{Linear (may be written as a Liouvillian)}} + \underbrace{\eta\,\rho\,\mathrm{tr}\left(\hat{L}\,\rho\,\hat{L}^\dagger\right)}_{\text{Non-Linear}},
\end{aligned}
$$

$$(A7)$$

which is appropriate between any detector clicks. We have written a trace–preserving equation of motion, which consequently explicitly includes a non-linear term that enforces conservation of probability (i.e. a normalization). In the absence of the non-linear term, the equation can be written in a linear (Liouvillian) form [32, 33]. The non-linear trace–preserving term arises when $\eta > 0$ because the probability of the associated event of no detector click is not equal to one. We have identified $\hat{\mathfrak{Z}} = -\frac{1}{2}\hat{L}^\dagger \hat{L}$ for $\hat{L} = \sqrt{\gamma}\,\hat{\sigma}_-$, and see that the Lindblad master equation is recovered for the unmonitored case (i.e., for the unconditional dynamics that arise from $\eta = 0$).

We then arrive at the expressions used in the main text by adding unitary qubit rotations defined by a Rabi drive $\boldsymbol{\Omega}$ to the dynamics above (these unitaries function as a control Hamiltonian throughout the main text), resulting in $\dot{\rho} =$

$$i[\rho, \hat{H}] + (1-\eta)\hat{L}\,\rho\,\hat{L}^\dagger - \tfrac{1}{2}\hat{L}^\dagger\hat{L}\,\rho - \tfrac{1}{2}\rho\,\hat{L}^\dagger\hat{L} + \eta\,\rho\,\mathrm{tr}\left(\hat{L}\,\rho\,\hat{L}^\dagger\right) \tag{A8a}$$

$$\text{with} \quad \hat{H} = \tfrac{1}{2}\left(\Omega_x\,\hat{\sigma}_x + \Omega_y\,\hat{\sigma}_y + \Omega_z\,\hat{\sigma}_z\right). \tag{A8b}$$

Note that such an expression is intrinsically conceived in the limit of small timesteps. The lack of commutation between the control unitary $\hat{\mathcal{U}} = e^{-i\hat{H}\Delta t}$ with any of the Kraus operators $\hat{\mathcal{M}}$ may be neglected to $O(\Delta t)$, but will typically contribute non-trivially to $O(\Delta t^{\frac{3}{2}})$ and beyond.

With these assumptions in place, the equation of motion can be recast as a dynamical system in the qubit's Bloch coordinates ($\dot{\mathbf{q}} = \mathrm{tr}(\hat{\sigma}_{\mathbf{q}}\,\rho)$), specifically as

$$\dot{x} = \mathcal{F}_x(\mathbf{q}, \boldsymbol{\Omega}) = \tfrac{1}{2}\gamma\,x\,(\eta(z+1)-1) + z\,\Omega_y - y\,\Omega_z, \tag{A9a}$$

$$\dot{y} = \mathcal{F}_y(\mathbf{q}, \boldsymbol{\Omega}) = \tfrac{1}{2}\gamma\,y\,(\eta(z+1)-1) + x\,\Omega_z - z\,\Omega_x, \tag{A9b}$$

$$\dot{z} = \mathcal{F}_z(\mathbf{q}, \boldsymbol{\Omega}) = \tfrac{1}{2}\gamma\,(1+z)\,(\eta(z+1)-2) + y\,\Omega_x - x\,\Omega_y, \tag{A9c}$$

for $\mathbf{q} \equiv (x, y, z)^\top$ and $\boldsymbol{\Omega} \equiv (\Omega_x, \Omega_y, \Omega_z)^\top$. We may equivalently separate out the parts of the dynamics due to the measurement (leading terms) and the Rabi drive, according to

$$\dot{\mathbf{q}} = \boldsymbol{\mathcal{F}}(\mathbf{q}, \boldsymbol{\Omega}) = \gamma\,\tilde{\boldsymbol{\mathcal{F}}}(\mathbf{q}) - \mathbf{q} \times \boldsymbol{\Omega}, \tag{A9d}$$

where $\tilde{\boldsymbol{\mathcal{F}}}(\mathbf{q}) = \tfrac{1}{\gamma}\,\boldsymbol{\mathcal{F}}_0$, i.e., $\tilde{\boldsymbol{\mathcal{F}}}$ are the un-controlled conditional dynamics, with the decay timescale factored out.

We may then regroup some terms in (A8a) to identify the non–Hermitian Hamiltonian of primary interest. This yields

$$\hat{\mathcal{H}} = \hat{H} - \tfrac{i}{2}\hat{L}^\dagger\hat{L} = \frac{1}{2}\begin{pmatrix} \Omega_z - i\,\gamma & \Omega_x - i\,\Omega_y \\ \Omega_x + i\,\Omega_y & -\Omega_z \end{pmatrix}, \tag{A10a}$$

in the context of the conditioned dynamics

$$\dot{\rho} = i\,\rho\,\hat{\mathcal{H}}^\dagger - i\,\hat{\mathcal{H}}\,\rho + (1-\eta)\hat{L}\,\rho\,\hat{L}^\dagger + \eta\,\rho\,\mathrm{tr}\left(\hat{L}\,\rho\,\hat{L}^\dagger\right) \tag{A10b}$$

with $\hat{L} = \sqrt{\gamma}\,\hat{\sigma}_-$. Eq. (A10b) is simply a re-writing of (A8a) in terms of the NH Hamiltonian.

Appendix B: Pontryagin Optimal Control: Details and Extended Discussion

1. The Pontryagin Maximum Principle

Our use of Pontryagin's principle for optimal control revolves around an action that is generically of the form

$$\mathscr{S} = f(\mathbf{q}_T) + \int_0^T dt\,(\mathbf{p}\cdot\dot{\mathbf{q}} - \mathscr{H}(\mathbf{q}, \mathbf{p}, \mathbf{u})), \tag{B1}$$

where $\mathbf{q}$ are state coordinates as in the main text, $\mathbf{p}$ are the co-states (i.e., conjugate "momenta", or Lagrange multipliers, directly analogous to $\boldsymbol{\Lambda}$), and $\mathbf{u}$ are some controller variables. Optimization involves solving $\delta\mathscr{S} = 0$, where $\delta\mathscr{S}$

$$= \delta f(\mathbf{q}_T) + \int_0^T (\delta[\mathbf{p}\cdot\dot{\mathbf{q}}] - \delta\mathscr{H}(\mathbf{q}, \mathbf{p}, \mathbf{u})),$$

$$= \mathbf{p}\cdot\delta\mathbf{q}\big|_0^T + \frac{\partial f}{\partial\mathbf{q}_T}\delta\mathbf{q}_T + \int_0^T dt\,\{\delta\mathbf{p}\cdot\dot{\mathbf{q}} - \dot{\mathbf{p}}\cdot\delta\mathbf{q} - \delta\mathscr{H}\}, \tag{B2}$$

with

$$\delta\mathscr{H} = \frac{\partial\mathscr{H}}{\partial\mathbf{q}}\delta\mathbf{q} + \frac{\partial\mathscr{H}}{\partial\mathbf{p}}\delta\mathbf{p} + \frac{\partial\mathscr{H}}{\partial\mathbf{u}}\delta\mathbf{u}. \tag{B3}$$

The equations

$$\dot{\mathbf{q}} = \frac{\partial\mathscr{H}}{\partial\mathbf{p}}, \quad \dot{\mathbf{p}} = -\frac{\partial\mathscr{H}}{\partial\mathbf{q}}, \quad \frac{\partial\mathscr{H}}{\partial\mathbf{u}} = 0 \tag{B4}$$

arise straightforwardly from requiring that the integral contribution goes to zero, independent of the variations $\delta\mathbf{q}$, $\delta\mathbf{p}$, and $\delta\mathbf{u}$.

This mathematical approach to control, where optimal controls are derived via action extremization, variational calculus or geodesics, has a long history in both classical [45, 46] and quantum [47, 48] control. This includes work on related physical systems, where an action–extremization principle has similarly been used to find extremal–probability measurement records and paths (instead of unitary controllers) for continuously–monitored quantum systems [42, 50–52, 62, 64–68]. An application of a similar method towards simultaneous optimization of a diffusive measurement record and unitary controls has recently been carried out [69].

2. Optimal Equations of Motion

In the main text, we introduce a Hamiltonian for Pontryagin control of our qubit, which reads

$$\mathscr{H} = \boldsymbol{\Lambda} \cdot \boldsymbol{\mathcal{F}}(\mathbf{q}, \boldsymbol{\Omega}) - \underbrace{\left(\frac{\alpha}{2} \left\{ (x - x_f)^2 + (y - y_f)^2 + (z - z_f)^2 \right\} + \frac{\beta}{2} \left\{ \Omega_x^2 + \Omega_y^2 + \Omega_z^2 \right\} \right)}_{\mathscr{L}(\mathbf{q}, \boldsymbol{\Omega})}. \tag{B5}$$

Note that the actual cost function (i.e. the part of the Hamiltonian that is not directly implementing the constraint to the dynamical equations $\dot{\mathbf{q}} = \boldsymbol{\mathcal{F}}$) may be recognized as the corresponding Lagrangian $\mathscr{L}$. We have chosen a simple form of the cost function in this work; it is possible to follow the same procedure with a modified cost function that may lead to solutions better optimized for a particular task.

We now go through the optimization calculations based on $\mathscr{H}$ in some detail. We consider the drive optimization condition first, finding

$$\frac{\partial \mathscr{H}}{\partial \boldsymbol{\Omega}} = 0 \quad \rightarrow \quad \Omega_x^\star = \tfrac{1}{\beta} \left(y\, \Lambda_z - z\, \Lambda_y \right), \quad \Omega_y^\star = \tfrac{1}{\beta} \left(z\, \Lambda_x - x\, \Lambda_z \right), \quad \Omega_z^\star = \tfrac{1}{\beta} \left(x\, \Lambda_y - y\, \Lambda_x \right). \tag{B6}$$

This may remind the reader of expressions for the angular momentum, if the Lagrange multipliers $\boldsymbol{\Lambda}$ are understood as conceptually analogous to linear momenta (i.e., the choice of $\mathscr{L}$ above leads to $\boldsymbol{\Omega}^\star = \frac{1}{\beta} \mathbf{q} \times \boldsymbol{\Lambda}$). We emphasize again our remark in the main text that this form of the optimal drive indicates that the choice of β is somewhat redundant with the choice we must make in practice about where to bound the Lagranian Manifold, and hence also the Pontryagin energy $\mathscr{E} = \mathscr{H}$. For this reason we pay relatively little attention below to this "temperance" parameter.

Substituting in $\boldsymbol{\Omega} \rightarrow \boldsymbol{\Omega}^\star = \frac{1}{\beta} \mathbf{q} \times \boldsymbol{\Lambda}$ leads to

$$\begin{aligned}
\mathscr{H}^\star &= \boldsymbol{\Lambda} \cdot \left(\boldsymbol{\mathcal{F}}(\mathbf{q}, \boldsymbol{\Omega} = 0) - \tfrac{1}{\beta} \mathbf{q} \times (\mathbf{q} \times \boldsymbol{\Lambda}) \right) - \tfrac{\alpha}{2} |\mathbf{q} - \mathbf{Q}_f|^2 - \tfrac{1}{2\beta} |\mathbf{q} \times \boldsymbol{\Lambda}|^2 \\
&= \boldsymbol{\Lambda} \cdot \boldsymbol{\mathcal{F}}(\mathbf{q}, \boldsymbol{\Omega} = 0) + \tfrac{1}{2\beta} \mathbf{q}^2\, \boldsymbol{\Lambda}^2 - \tfrac{1}{2\beta} (\mathbf{q} \cdot \boldsymbol{\Lambda})^2 - \tfrac{\alpha}{2} |\mathbf{q} - \mathbf{Q}_f|^2,
\end{aligned} \tag{B7}$$

which is the Hamiltonian generating optimally–controlled dynamics. Note that after substituting in $\boldsymbol{\Omega} = \boldsymbol{\Omega}^\star$ everywhere, some terms from the uncontrolled part of the dynamical equations $\boldsymbol{\mathcal{F}}(\mathbf{q}, \boldsymbol{\Omega} = 0) = \boldsymbol{\mathcal{F}}_0$ will remain. The Hamiltonian above may equivalently be expanded to read

$$\begin{aligned}
\mathscr{H}^\star = &\tfrac{\gamma}{2} \left\{ [\eta(z+1) - 1] [x\, \Lambda_x + y\, \Lambda_y] + (1+z) \Lambda_z [\eta(z+1) - 2] \right\} \\
&+ \tfrac{1}{2\beta} \left\{ \Lambda_x^2 (y^2 + z^2) + \Lambda_y^2 (x^2 + z^2) + \Lambda_z^2 (x^2 + y^2) \right\} - \tfrac{1}{\beta} \left\{ xy\, \Lambda_x \Lambda_y + xz\, \Lambda_x \Lambda_z + yz\, \Lambda_y \Lambda_z \right\} \\
&- \tfrac{\alpha}{2} \left\{ (x - X_f)^2 + (y - Y_f)^2 + (z - Z_f)^2 \right\}.
\end{aligned} \tag{B8}$$

The subsequent equations of motion (Hamilton's dynamical equations) now read

$$\dot{x}^\star = \tfrac{1}{2} \gamma\, x \{ \eta(z+1) - 1 \} - \tfrac{1}{\beta}\, x (\Lambda_y y + \Lambda_z z) + \tfrac{1}{\beta}\, \Lambda_x \left(y^2 + z^2 \right), \tag{B9a}$$

$$\dot{y}^\star = \tfrac{1}{2} \gamma\, y \{ \eta(z+1) - 1 \} - \tfrac{1}{\beta}\, y (\Lambda_x x + \Lambda_z z) + \tfrac{1}{\beta}\, \Lambda_y \left(x^2 + z^2 \right), \tag{B9b}$$

$$\dot{z}^\star = \tfrac{1}{2} \gamma (z+1) \{ \eta(z+1) - 2 \} - \tfrac{1}{\beta}\, z (\Lambda_x x + \Lambda_y y) + \tfrac{1}{\beta}\, \Lambda_z \left(x^2 + y^2 \right), \tag{B9c}$$

$$\dot{\Lambda}_x^\star = \alpha(x - X_f) - \tfrac{1}{2} \gamma\, \Lambda_x \{ \eta(z+1) - 1 \} + \tfrac{1}{\beta}\, \Lambda_x (\Lambda_y y + \Lambda_z z) - \tfrac{1}{\beta}\, x \left(\Lambda_y^2 + \Lambda_z^2 \right), \tag{B9d}$$

$$\dot{\Lambda}_y^\star = \alpha(y - Y_f) - \tfrac{1}{2} \gamma\, \Lambda_y \{ \eta(z+1) - 1 \} + \tfrac{1}{\beta}\, \Lambda_y (\Lambda_x x + \Lambda_z z) - \tfrac{1}{\beta}\, y \left(\Lambda_x^2 + \Lambda_z^2 \right), \tag{B9e}$$

$$\dot{\Lambda}_z^\star = \alpha(z - Z_f) - \gamma\, \Lambda_z \{ \eta(z+1) - 1 \} - \tfrac{1}{2} \gamma\, \eta (\Lambda_y y + \Lambda_x x) + \tfrac{1}{\beta}\, \Lambda_z (\Lambda_x x + \Lambda_y y) - \tfrac{1}{\beta}\, z \left(\Lambda_x^2 + \Lambda_y^2 \right). \tag{B9f}$$

It is these equations which are explicitly integrated in the course of finding dynamical solutions. Some features of the dynamics are more apparent in this form: For instance, the leading terms on each $\dot{\boldsymbol{\Lambda}}$ equation (i.e., $\alpha(\mathbf{q} - \mathbf{Q}_f) - ...$) clearly illustrate how the magnitude and direction of the control drive is adjusted if we find ourselves far from the target state. More specifically, α essentially determines how rapidly the control drive can

change, such that paths derived with a larger α are capable of greater acceleration. As shown in Fig. 6, this property is highly useful to reach pure target states under inefficient detection.

In the main text we have worked with a Lagrangian Manifold (LM) of control parameter solutions defined by choosing a volume of initial $\boldsymbol{\Lambda}_i$, corresponding to a volume of initial conditions for optimal control trajectories. Solutions in this LM are unique within the full (six–dimensional) phase space of $\mathscr{H}$, but are not necessarily unique when projected into the smaller, physical, $\mathbf{q}$–space. Regions of the Bloch sphere where multiple solutions meet the given boundary conditions are directly analogous to the caustics that appear in optics and diverse other physical settings, that are described mathematically by catastrophe theory [49–51, 70–73].

3. On the Existence of Control Solutions Meeting Particular Boundary Conditions

We comment here on the simplest case (of pure states with $\eta = 1$), and then on the general mixed–state case with $\eta \leq 1$.

a. Pure State Solutions

It is easy to see from (A9d) that trivial purity–preserving solutions exist for $\eta = 1$ and in the limit $|\boldsymbol{\Omega}| \gg \gamma$. If the unitary control dynamics are fast, i.e., they occur on timescales where the decay is negligible, we may factor out the drive magnitude $\dot{\mathbf{q}} = |\boldsymbol{\Omega}| (\bar{\Omega} \times \mathbf{q}) + \gamma \boldsymbol{\mathcal{F}}$ (where $\bar{\Omega}$ is the unit vector setting the rotation axis), and immediately see that

$$\frac{\dot{\mathbf{q}}}{|\boldsymbol{\Omega}|} = \bar{\Omega} \times \mathbf{q} + \frac{\gamma}{|\boldsymbol{\Omega}|} \tilde{\boldsymbol{\mathcal{F}}} \approx \bar{\Omega} \times \mathbf{q}. \tag{B10}$$

In other words, in the limit of fast drive ($\gamma \ll |\boldsymbol{\Omega}|$), the decay dynamics can be treated as a perturbation, and the trivial linear control problem arising from Rabi drive alone is recovered in the limit where the controls are applied very strongly/quickly. We consequently conclude that our ability to map arbitrary pure states to other ar-

bitrary pure states in the ideal ($\eta = 1$) case is limited only by the controller power.

The mixed state case is less straightforward. We will find that it is useful to supplement the Cartesian representation of our dynamics above with a spherical coordinate representation before considering the mixed state solutions in detail.

b. Dynamics and Control in Spherical Coordinates

Let us convert (A9) to spherical coordinates $x = R \cos \phi \sin \theta$, $y = R \sin \phi \sin \theta$, and $z = R \cos \theta$, thereby representing the dynamics by

$$\dot{R} = \gamma \left\{ \tfrac{1}{2} \cos \theta \left(\eta(1 + R^2) - 2 \right) \right. \\ \left. + \tfrac{1}{4} R(\eta - 1)(\cos(2\theta) + 3) \right\}, \tag{B11a}$$

$$\dot{\theta} = -\frac{\gamma(\eta + (\eta - 1)R \cos \theta - 2) \sin \theta}{2R} \\ - \Omega_x \sin \phi + \Omega_y \cos \phi, \tag{B11b}$$

$$\dot{\phi} = \Omega_z - (\Omega_x \cos \phi + \Omega_y \sin \phi) \cot \theta. \tag{B11c}$$

A large part of our interest in the spherical form of the problem stems from the dynamics of state purity $\mathcal{P} = \mathrm{tr}\left(\rho^2\right)$. This expression is closely related to the radial equation of motion, i.e., $\dot{\mathcal{P}} = 2 \mathrm{tr}\left(\rho \dot{\rho}\right) = R \dot{R}$. This can equivalently be written as $\dot{\mathcal{P}} =$

$$\tfrac{1}{2} \gamma \left[(x^2 + y^2)\{\eta(1 + z) - 1\} + z(z + 1)\{\eta(1 + z) - 2\} \right]. \tag{B12}$$

Since these dynamics related to changes in state purity are independent of ϕ, they can be completely represented within a cross-sectional plane of the Bloch sphere; graphical representations appear in Fig. 7.

The control problem as a whole can also be recast in spherical coordinates, where the change from Cartesian $\mathbf{q} = (x, y, z)$ and $\boldsymbol{\Lambda} = (\Lambda_x, \Lambda_y, \Lambda_z)$ to spherical $\mathbf{q} = (R, \theta, \phi)$ and $\boldsymbol{\Lambda} = (\Lambda_R, \Lambda_\theta, \Lambda_\phi)$ is performed via canonical transformation:

$$
\begin{aligned}
x &\rightarrow & R \cos \phi \sin \theta, \\
y &\rightarrow & R \sin \phi \sin \theta, \\
z &\rightarrow & R \cos \theta, \\
\Lambda_x &\rightarrow & \Lambda_R \sin \theta \cos \phi + (\Lambda_\theta/R) \cos \theta \cos \phi - (\Lambda_\phi/R) \csc \theta \sin \phi, \\
\Lambda_y &\rightarrow & \Lambda_R \sin \theta \sin \phi + (\Lambda_\theta/R) \cos \theta \sin \phi + (\Lambda_\phi/R) \csc \theta \cos \phi, \\
\Lambda_z &\rightarrow & \Lambda_R \cos \theta - (\Lambda_\theta/R) \sin \theta.
\end{aligned}
\tag{B13}
$$

Applying the above transformation to the Cartesian form of $\mathscr{H}$ preserves the Poisson bracket between all dynam-

ical variables, and is equivalent to re-assembling the

Hamiltonian in spherical coordinates as follows

$$\mathscr{H} = \Lambda_R \, \dot{R} + \Lambda_\theta \, \dot{\theta} + \Lambda_\phi \, \dot{\phi} - \mathscr{L}(\mathbf{q}, \mathbf{\Omega}). \qquad (B14)$$

Note that the equations for the optimal control (i.e., from solving $\partial_{\mathbf{\Omega}}\mathscr{H} = 0$) now read

$$\Omega_x^\star = -\Lambda_\theta \, \sin\phi - \Lambda_\phi \, \cot\theta \, \cos\phi, \qquad (B15a)$$

$$\Omega_y^\star = \Lambda_\theta \, \cos\phi - \Lambda_\phi \, \cot\theta \, \sin\phi, \qquad (B15b)$$

$$\Omega_z^\star = \Lambda_\phi. \qquad (B15c)$$

These are quite helpful in constructing simulations, in that they formalize an intuition that the co-state variable Λ_R plays no role in the controller trajectory (because Λ_R corresponds to changes in purity, which the unitary controller cannot implement directly). Consequently, we may work with a two–dimensional Lagrangian Manifold (LM), completely defined by an initial mesh of Λ_θ and Λ_ϕ, rather than a three–dimensional one in Λ_x, Λ_y, and Λ_z. This reduction in the dimensionality of the integrated manifold considerably simplifies the numerics.

c. Mixed State Solutions

The rate of change of the state purity $\dot{\mathcal{P}}$ is especially important to us in the cases of inefficient detection, $\eta < 1$, as well as an initially impure qubit state, because the amount of evolution time spent in regions of $\dot{\mathcal{P}} > 0$ will upper-bound the final state purity (and thereby constrain which final states are attainable). Recall that unitary dynamics (i.e., $\mathbf{\Omega}$) cannot directly change the state purity. Purity is here increased (decreased) purely through the acquisition (loss) of information the qubit has shared with its optical environment via spontaneous emission. What we can do, however, is to devise controller trajectories that drive the system to a state where the no-click (no-jump) measurement information leads to purity increase, i.e., a measurement induced change in purity that depends on the qubit state. In this manner we can use our controls to indirectly affect $\mathcal{P}$.

Qualitative aspects of the relevant dynamics may be immediately inferred from Fig. 7. For instance, the line $\dot{\mathcal{P}} = 0$ necessarily goes through $|g\rangle$ irrespective of η, and this point is the only state where $\dot{\mathcal{P}} = 0$ *and* $\mathcal{P} = 1$ for $\eta < 1$. It follows that the jumpless dynamics (including the Lindbladian dynamics without detection) can only asymptotically approach $\mathcal{P} = 1$. More specifically, the dynamics without drive ($\mathbf{\Omega} = 0$) decay as

$$z(t) = \frac{2\,u_0}{\eta\,u_0 - e^{\gamma\,t}(\eta\,u_0 - 2)} - 1 \qquad (B16)$$

for $u_0 = 1 + z_0 \in [0, 2]$. These solutions asymptotically approach $z = -1$ in the long–time limit for all parameters

except for $z_0 = 1$ (initial $|e\rangle$) and $\eta = 1$ [42, 74]. Therefore, if one needs to significantly increase state purity, there are two options to accomplish this (which are not mutually incompatible). The first is to drive the trajectory close to $|g\rangle$ to the greatest extent possible, then wait long enough for the jumpless dynamics to purify the state to the desired degree, and then quickly drive the qubit from near $|g\rangle$ to the target state. The second is to alternatively actually measure a click, which resets the qubit directly to $|g\rangle$, at which point the attainable state purity is limited solely by the speed of the controls. Note that a slower trajectory from $|g\rangle \to \mathbf{Q}_f$ will have more time to lose purity, and therefore have a more constrained range of accessible final states (recall Fig. 5).

We have demonstrated in Fig. 6 that our optimal equations of motion are, especially for sufficiently large α, capable of generating the type of solutions we have just described. A larger value of the "agility" parameter α helps with this, because it allows our controller to "accelerate" towards the target state a short time before we intend to reach it, and to stay near $|g\rangle$ to retain purity prior to that.

Appendix C: Supplementary Plots and Animations

We provide supplementary animations in the accompanying slide deck [75]. These animations are designed to offer clear visual illustrations of the following points:

1. (Slide 2) We illustrate bi–orthogonal states on the Bloch sphere (see [8]), in preparation for subsequent animations.

2. (Slides 3–4) We demonstrate how adiabatic following can succeed or fail, using closed loop parameter trajectories of duration $T = 100\,T_1$. Adiabatic following works just as in the Hermitian case if our entire trajectory stays on the stable eigensheet (slide 3), but fails when we encircle an EP and thereby force a stability exchange (slide 4) (see [12]).

3. (Slides 5–6) Stability exchange when encircling an EP can also be understood as resulting in chiral state exchange. We illustrate a situation where we perform a clockwise encircling that stays on the stable eigensheet (slide 5), in contrast with a counter–clockwise encircling that attempts to follow the lossy eigensheet [10].

4. (Slides 7–8) Pure state optimal manifold dynamics from Pontryagin control cover the entire surface of the Bloch sphere over moderate time intervals for moderate manifold energy boundaries (which limit the maximum Rabi rotation rates $|\mathbf{\Omega}|$ available to the controller). These are animated versions of each panel in Fig. 4.

5. (Slides 9–13) Pontryagin control under inefficient measurement ($\eta < 1$) leads to difficulties in the
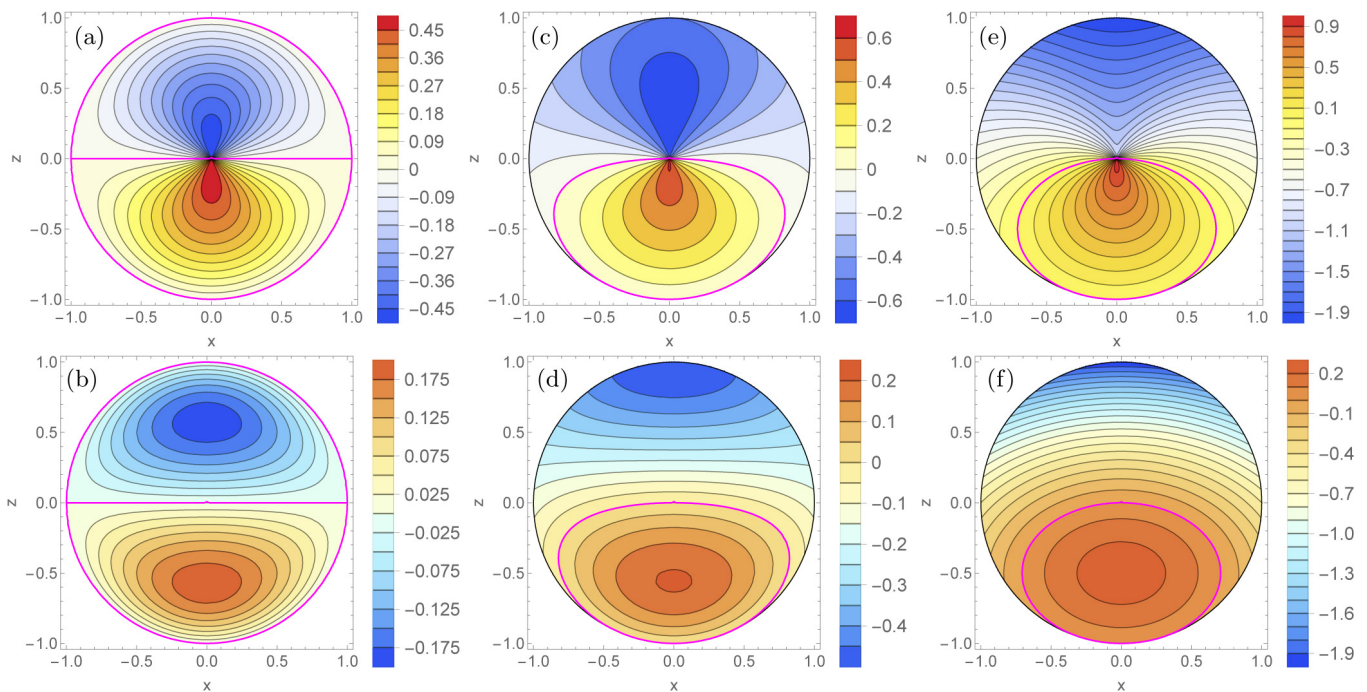
FIG. 7. Graphical representation of the change in purity for mixed state no–jump dynamics (A9). We plot the rate of change of the Bloch vector length, $\dot{R}$ (B11a), on the top (panels a,c,e) and the rate of change of the purity $\dot{\mathcal{P}}$ on the bottom (panels b, d, f). Results are shown for detection efficiencies $\eta = 1$ (panels a, b), $\eta = 0.75$ (panels c, d), and $\eta = 0$ (panels e, f). Contour plots show lines of constant $\dot{R}(x,z)$ or $\dot{\mathcal{P}}(x,z)$, such that the dynamical flow will run perpendicular to the contour lines, i.e. along the gradient. Both $\dot{R}$ and $\dot{\mathcal{P}}$ values are given in units of γ. The $\dot{R} = 0$ and $\dot{\mathcal{P}} = 0$ lines are shown in magenta; purity increases within the magenta boundary, and decreases outside of it. We can understand that purity decreases in the upper half of the Bloch sphere (i.e., closer to $|e\rangle$) in the inefficient conditional dynamics, and does so faster for worse efficiencies. On the other hand, the dynamical attraction towards $|g\rangle$ can be used to increase purity for states in a large region between the equator and $|g\rangle$, across a wide range of measurement efficiencies. These results generically imply that control protocols can reach states with $\mathcal{P}_f > \mathcal{P}_i$ if the target trajectory spends sufficient time traversing the region where $\dot{\mathcal{P}} > 0$.

exact targeting of pure final states. However, we illustrate here that attainment of nearly–pure final states is possible, and that larger values of α

can significantly improve the purity of final states reached by trajectories $t \gtrsim T_1$; our animations here supplement Figs. 5 and 6.

[1] C. M. Bender and S. Boettcher, Real spectra in non-Hermitian Hamiltonians having PT symmetry, Phys. Rev. Lett. **80**, 5243 (1998).

[2] A. Mostafazadeh, Pseudo-Hermiticity versus PT symmetry: The necessary condition for the reality of the spectrum of a non-Hermitian Hamiltonian, Journal of Mathematical Physics **43**, 205 (2002); Pseudo-Hermiticity versus PT-symmetry. II. A complete characterization of non-Hermitian Hamiltonians with a real spectrum, Journal of Mathematical Physics **43**, 2814 (2002); Pseudo-Hermiticity versus PT-symmetry III: Equivalence of pseudo-Hermiticity and the presence of antilinear symmetries, Journal of Mathematical Physics **43**, 3944 (2002).

[3] W. D. Heiss, Exceptional points of non-Hermitian operators, Journal of Physics A: Mathematical and General **37**, 2455 (2004).

[4] W. D. Heiss, The physics of exceptional points, Journal of Physics A: Mathematical and Theoretical **45**, 444016 (2012).

[5] R. El-Ganainy, K. Makris, M. Khajavikhan, Z. H. Musslimani, S. Rotter, and D. N. Christodoulides, Non-Hermitian physics and PT symmetry, Nat. Phys. **14**, 11–19 (2018).

[6] c. Özdemir, S. Rotter, F. Nori, and L. Yang, Parity–time symmetry and exceptional points in photonics, Nature Materials , 783–798 (2019).

[7] M.-A. Miri and A. Alù, Exceptional points in optics and photonics, Science **363**, eaar7709 (2019).

[8] Y. Ashida, Z. Gong, and M. Ueda, Non-Hermitian physics, Advances in Physics **69**, 249 (2020).

[9] C. Dembowski, B. Dietz, H.-D. Gräf, H. L. Harney, A. Heine, W. D. Heiss, and A. Richter, Encircling an exceptional point, Phys. Rev. E **69**, 056216 (2004).

[10] Q. Zhong, M. Khajavikhan, D. Christodoulides, and R. El-Ganainy, Winding around non-Hermitian singularities, Nat. Comm. **9**, 4808 (2018).

[11] A. Kvitsinsky and S. Putterman, Adiabatic evolution of an irreversible two level system, Journal of Mathematical Physics **32**, 1403 (1991).

[12] G. Nenciu and G. Rasche, On the adiabatic theorem for nonself-adjoint Hamiltonians, Journal of Physics A: Mathematical and General **25**, 5741 (1992).

[13] M. V. Berry and R. Uzdin, Slow non-Hermitian cycling: exact solutions and the stokes phenomenon, Journal of Physics A: Mathematical and Theoretical **44**, 435303 (2011).

[14] M. V. Berry, Optical polarization evolution near a non-Hermitian degeneracy, Journal of Optics **13**, 115701 (2011).

[15] A. Leclerc, G. Jolicard, and J. P. Killingbeck, Discussion of the adiabatic hypothesis in control schemes using exceptional points, Journal of Physics B: Atomic, Molecular and Optical Physics **46**, 145503 (2013).

[16] D. Viennot, Adiabatic quantum control hampered by entanglement, Journal of Physics A: Mathematical and Theoretical **47**, 295301 (2014).

[17] T. J. Milburn, J. Doppler, C. A. Holmes, S. Portolan, S. Rotter, and P. Rabl, General description of quasia-diabatic dynamical phenomena near exceptional points, Phys. Rev. A **92**, 052124 (2015).

[18] J. Doppler, A. A. Mailybaev, J. Böhm, U. Kuhl, A. Girschik, F. Libisch, T. J. Milburn, P. Rabl, N. Moiseyev, and S. Rotter, Dynamically encircling an exceptional point for asymmetric mode switching, Nature **537**, 76–79 (2016).

[19] A. U. Hassan, B. Zhen, M. Soljačić, M. Khajavikhan, and D. N. Christodoulides, Dynamically encircling exceptional points: Exact evolution and polarization state conversion, Phys. Rev. Lett. **118**, 093002 (2017).

[20] A. Pick, S. Silberstein, N. Moiseyev, and N. Bar-Gill, Robust mode conversion in NV centers using exceptional points, Phys. Rev. Research **1**, 013015 (2019).

[21] H. Wang, S. Assawaworrarit, and S. Fan, Dynamics for encircling an exceptional point in a nonlinear non-Hermitian system, Opt. Lett. **44**, 638 (2019).

[22] Q. Zhong, S. Nelson, Ş. K. Özdemir, and R. El-Ganainy, Controlling directional absorption with chiral exceptional surfaces, Opt. Lett. **44**, 5242 (2019).

[23] J. Höller, N. Read, and J. G. E. Harris, Non-Hermitian adiabatic transport in spaces of exceptional points, Phys. Rev. A **102**, 032216 (2020).

[24] Q. Zhong, A. Hashemi, Ş. K. Özdemir, and R. El-Ganainy, Control of spontaneous emission dynamics in microcavities with chiral exceptional surfaces, Phys. Rev. Research **3**, 013220 (2021).

[25] H. Xu, D. Mason, L. Jiang, and J. G. E. Harris, Topological energy transfer in an optomechanical system with exceptional points, Nature **537**, 80–83 (2016).

[26] Yu-Kun, L. P. Peng, Q.-T. Cao, D. Xu, J. Wiersig, Q. Gong, and Y.-F. Xiao, Spontaneous t-symmetry breaking and exceptional points in cavity quantum electrodynamics systems, Science Bulletin **63**, 1096 (2018).

[27] Y.-X. Wang and A. A. Clerk, Non-hermitian dynamics without dissipation in quantum systems, Phys. Rev. A **99**, 063834 (2019).

[28] K. Wang, A. Dutt, K. Y. Yang, C. C. Wojcik,

[29] J. Vučković, and S. Fan, Generating arbitrary topological windings of a non-Hermitian band, Science **371**, 1240 (2021).

[29] K. Wang, A. Dutt, C. C. Wojcik, and S. Fan, Topological complex-energy braiding of non-Hermitian bands, Nature **598** (2021).

[30] Y. Wu, W. Liu, J. Geng, X. Song, X. Ye, C.-K. Duan, X. Rong, and J. Du, Observation of parity-time symmetry breaking in a single-spin system, Science **364**, 878 (2019).

[31] M. Naghiloo, M. Abbasi, Y. N. Joglekar, and K. W. Murch, Quantum state tomography across the exceptional point in a single dissipative qubit, Nat. Phys. **15**, 1232–1236 (2019).

[32] F. Minganti, A. Miranowicz, R. W. Chhajlany, and F. Nori, Quantum exceptional points of non-Hermitian Hamiltonians and Liouvillians: The effects of quantum jumps, Phys. Rev. A **100**, 062131 (2019).

[33] F. Minganti, A. Miranowicz, R. W. Chhajlany, I. I. Arkhipov, and F. Nori, Hybrid-Liouvillian formalism connecting exceptional points of non-Hermitian Hamiltonians and Liouvillians via postselection of quantum trajectories, Phys. Rev. A **101**, 062112 (2020).

[34] W. Chen, M. Abbasi, Y. N. Joglekar, and K. W. Murch, Quantum jumps in the non-hermitian dynamics of a superconducting qubit, Phys. Rev. Lett. **127**, 140504 (2021).

[35] W. Chen, M. Abbasi, B. Ha, S. Erdamar, Y. N. Joglekar, and K. W. Murch, Decoherence-induced exceptional points in a dissipative superconducting qubit, Phys. Rev. Lett. **128**, 110402 (2022).

[36] M. Abbasi, W. Chen, M. Naghiloo, Y. N. Joglekar, and K. W. Murch, Topological quantum state control through exceptional-point proximity, Phys. Rev. Lett. **128**, 160401 (2022).

[37] J. Huber, P. Kirton, S. Rotter, and P. Rabl, Emergence of PT-symmetry breaking in open quantum systems, SciPost Phys. **9**, 52 (2020).

[38] H. J. Carmichael, *An Open Systems Approach to Quantum Optics* (Springer, Berlin, 1993).

[39] A. Barchielli and M. Gregoratti, *Quantum trajectories and measurements in continuous time* (Springer-Verlag, Berlin, Heidelberg, 2009).

[40] H. M. Wiseman and G. J. Milburn, *Quantum Measurement and Control* (Cambridge University Press, 2009).

[41] K. Jacobs, *Quantum Measurement Theory and its Applications* (Cambridge University Press, 2014).

[42] P. Lewalle, S. K. Manikandan, C. Elouard, and A. N. Jordan, Measuring fluorescence to track a quantum emitter's state: a theory review, Contemporary Physics **61**, 26 (2020).

[43] H. Ribeiro and F. Marquardt, Accelerated non-reciprocal transfer of energy around an exceptional point (2021), arXiv:2111.11220 [quant-ph].

[44] Y. Chen, Q. Wu, B. Huang, J. Song, Y. Xia, and S. Zheng, Improving shortcuts to non-Hermitian adiabaticity for fast population transfer in open quantum systems, Annalen der Physik **530**, 1700247 (2018).

[45] D. Liberzon, *Calculus of Variations and Optimal Control Theory* (Princeton University Press, Princeton, NJ, 2012).

[46] H. Schättler and U. Ledzewicz, *Geometric Optimal Control* (Springer, New York, 2010).

[47] D. D'Alessandro, *Introduction to Quantum Control and*

Dynamics, 2nd ed. (Chapman and Hall/CRC, 2021).

[48] U. Boscain, M. Sigalotti, and D. Sugny, Introduction to the Pontryagin maximum principle for quantum optimal control, PRX Quantum **2**, 030203 (2021).

[49] V. I. Arnold, *Mathematical Methods of Classical Mechanics* (Springer, New York, 1989).

[50] P. Lewalle, A. Chantasri, and A. N. Jordan, Prediction and characterization of multiple extremal paths in continuously monitored qubits, Phys. Rev. A **95**, 042126 (2017).

[51] M. Naghiloo, D. Tan, P. M. Harrington, P. Lewalle, A. N. Jordan, and K. W. Murch, Quantum caustics in resonance-fluorescence trajectories, Phys. Rev. A **96**, 053807 (2017).

[52] P. Lewalle, J. Steinmetz, and A. N. Jordan, Chaos in continuously monitored quantum systems: An optimal-path approach, Phys. Rev. A **98**, 012141 (2018).

[53] C. P. Koch, Controlling open quantum systems: tools, achievements, and limitations, Journal of Physics: Condensed Matter **28**, 213001 (2016).

[54] T. L. Patti, A. Chantasri, L. P. García-Pintos, A. N. Jordan, and J. Dressel, Linear feedback stabilization of a dispersively monitored qubit, Phys. Rev. A **96**, 022311 (2017).

[55] J. Zhang, Y. xi Liu, R.-B. Wu, K. Jacobs, and F. Nori, Quantum feedback: Theory, experiments, and applications, Physics Reports **679**, 1 (2017).

[56] S. Zhang, L. S. Martin, and K. B. Whaley, Locally optimal measurement-based quantum feedback with application to multiqubit entanglement generation, Phys. Rev. A **102**, 062418 (2020).

[57] Y. S. S. Patil, J. Höller, P. A. Henry, C. Guria, Y. Zhang, L. Jiang, N. Kralj, N. Read, and J. G. E. Harris, Measuring the knot of degeneracies and the eigenvalue braids near a third-order exceptional point (2021).

[58] S. Ibáñez, S. Martínez-Garaot, X. Chen, E. Torrontegui, and J. G. Muga, Shortcuts to adiabaticity for non-Hermitian systems, Phys. Rev. A **84**, 023415 (2011); S. Ibáñez and J. G. Muga, Adiabaticity condition for non-Hermitian Hamiltonians, Phys. Rev. A **89**, 033403 (2014).

[59] B. T. Torosov, G. Della Valle, and S. Longhi, Non-Hermitian shortcut to adiabaticity, Phys. Rev. A **87**, 052502 (2013).

[60] H. Li, H. Z. Shen, S. L. Wu, and X. X. Yi, Shortcuts to adiabaticity in non-Hermitian quantum systems without rotating-wave approximation, Opt. Express **25**, 30135 (2017).

[61] D. Guéry-Odelin, A. Ruschhaupt, A. Kiely, E. Torrontegui, S. Martínez-Garaot, and J. G. Muga, Shortcuts to adiabaticity: Concepts, methods, and applications, Rev. Mod. Phys. **91**, 045001 (2019).

[62] A. Jordan, A. Chantasri, P. Rouchon, and B. Huard, Anatomy of fluorescence: quantum trajectory statistics from continuously measuring spontaneous emission, Quantum Stud.: Math. Found. **3**, 237–263 (2016).

[63] V. Weisskopf and E. Wigner, Berechnung der natürlichen linienbreite auf grund der diracschen lichttheorie, Zeitschrift für Physik **63**, 54 (1930).

[64] A. Chantasri, J. Dressel, and A. N. Jordan, Action principle for continuous quantum measurement, Phys. Rev. A **88**, 042110 (2013).

[65] S. J. Weber, A. Chantasri, J. Dressel, A. N. Jordan, K. W. Murch, and I. Siddiqi, Mapping the optimal route between two quantum states, Nature **511**, 570 (2014).

[66] A. Chantasri and A. N. Jordan, Stochastic path-integral formalism for continuous quantum measurement, Phys. Rev. A **92**, 032125 (2015).

[67] A. Chantasri, I. Guevara, K. T. Laverick, and H. M. Wiseman, Unifying theory of quantum state estimation using past and future information, Physics Reports **930**, 1 (2021).

[68] T. Karmakar, P. Lewalle, and A. N. Jordan, Stochastic path-integral analysis of the continuously monitored quantum harmonic oscillator, PRX Quantum **3**, 010327 (2022).

[69] W. Kokaew, T. Chotibut, and A. Chantasri, Quantum state-preparation control in noisy environment via most-likely paths (2022).

[70] M. A. Alonso, Rays and waves, in *Phase-space optics: fundamentals and applications*, edited by M. E. Testorf, B. M. Hennelly, and J. Ojeda-Castañeda (McGraw-Hill, New York, 2009).

[71] V. I. Arnold, *Catastrophy Theory* (Springer-Verlag, Berlin, 1992).

[72] V. P. Maslov and M. V. Fedoriuk, *Semi-Classical Approximations in Quantum Mechanics* (Reidel, Dordrecht, Holland, 1981).

[73] R. G. Littlejohn, The Van Vleck formula, Maslov theory, and phase space geometry, Journal of Statistical Physics **68**, 7 (1992).

[74] D. Tan, N. Foroozani, M. Naghiloo, A. H. Kiilerich, K. Mølmer, and K. W. Murch, Homodyne monitoring of postselected decay, Phys. Rev. A **96**, 022104 (2017).

[75] Pontryagin–optimal control of a non-Hermitian qubit: Supplementary animations, `https://drive.google.com/file/d/1Ahp9alxd3EkNtTj8IckOVpWzxZODsTU5/view?usp=sharing`.