

Toward a Theory of Phase Transitions in Quantum Control Landscapes

Nicolò Beato^{1,*}, Pranay Patil^{2,3,†} and Marin Bukov^{1,‡}

¹*Max Planck Institute for the Physics of Complex Systems, Nöthnitzer Straße 38, 01187 Dresden, Germany*

²*Department of Physics, Indian Institute of Technology Madras, Chennai 600036, India*

³*Theory of Quantum Matter Unit, Okinawa Institute of Science and Technology Graduate University, Onna-son, Okinawa 904-0412, Japan*



(Received 22 August 2024; revised 2 July 2025; accepted 15 August 2025; published 29 October 2025)

Control landscape phase transitions (CLPTs) occur as abrupt changes in the cost function landscape upon varying a control parameter and can be revealed by nonanalytic points in statistical order parameters. A prime example are quantum speed limits, which mark the onset of controllability as the protocol duration is increased. Here, we lay the foundations of an analytical theory for CLPTs by developing Dyson, Magnus, and cumulant expansions for the cost function that capture the behavior of CLPTs with a controlled precision. Using linear and quadratic stability analysis, we reveal that CLPTs can be associated with different types of instabilities of the optimal protocol. This allows us to explicitly relate CLPTs to critical structural rearrangements in the extrema of the control landscape: Utilizing path-integral methods from statistical field theory, we trace back the critical scaling of the order parameter at the quantum speed limit to the topological and geometric properties of the set of optimal protocols, such as the number of connected components and its dimensionality. We verify our predictions by introducing a numerical sampling algorithm designed to explore this optimal set via a homotopic stochastic update rule. We apply this new toolbox explicitly to analyze CLPTs in the single- and two-qubit control problems whose landscapes are analytically tractable and compare the landscapes for bang-bang and piecewise continuous protocols. Our work provides the first steps toward a systematic theory of CLPTs and paves the way for utilizing statistical field theory techniques for generic complex control landscapes.

DOI: [10.1103/27nn-cb3t](https://doi.org/10.1103/27nn-cb3t)

Subject Areas: Condensed Matter Physics,
Quantum Physics, Statistical Physics

I. INTRODUCTION

Ever since quantum systems have become an object of investigation, the question of how to manipulate them optimally has been central to both verifying and understanding the theory behind quantum mechanics, as well as to developing new experimental techniques. Nowadays, control theory has become an indispensable tool in emerging quantum technologies [1–4]: Efficient quantum state manipulation is an essential prerequisite for the experimental realization of simulations on modern intermediate-noise-scale quantum devices [5–7], for the improvement of performance in quantum measurement devices [8–12], or quantum information processing [13–17]. Quantum control

theory aims to find protocols that optimize the values of physical quantities of interest: For example, in state preparation or gate synthesis problems, one is typically interested in minimizing infidelity or expectation values of a certain physical observable such as the energy or the total magnetization [18,19].

In quantum control problems, one distinguishes between two types of physical degrees of freedom: the *controlled system* and the *control variables* (see Fig. 1). The controlled system is the quantum system whose state we aim to control. The control variables correspond to the driving field applied to the controlled system. Time-dependent control protocols have an extensive number of degrees of freedom—one for each time step when time is discretized or a single continuous control *field* (in the sense of field theory) composed of infinitely many degrees of freedom. Consequently, the problem of finding optimal control protocols poses the formidable challenge of searching through high- or infinite-dimensional spaces.

The difficulty in finding analytical solutions arises from the lack of explicit expressions for the solution of the time-dependent Schrödinger equation, even for the simplest of quantum systems [3–5]. This motivated the development of an array of numerical optimal control algorithms. Among

*Contact author: nbeato@pks.mpg.de

†Contact author: pranay.patil@iitm.ac.in

‡Contact author: mgbukov@pks.mpg.de

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI. Open access publication funded by the Max Planck Society.




FIG. 1. In quantum control problems, there are two types of physical degrees of freedom: the *controlled system* (e.g., qubits) and the *control variables* defining the control protocol (e.g., modulation of electric or magnetic fields, or electromagnetic radiation, external forces, etc.). The figure of merit that measures the quality of control protocols defines the control landscape. Control phase transitions refer to critical changes in the structure of the control landscape, as a model parameter (e.g., the protocol duration T) is varied. Note that the control landscape can present a many-body problem in the control variables, even when the controlled system consists of as few as a single degree of freedom; see Sec. II.

them, two classes can be distinguished, namely, gradient-based methods (e.g., GRAPE [20] and Krotov [21]) and gradient-free algorithms (e.g., CRAB [22,23]). The former class uses the differentiability of the cost function; the latter relies on a clever sampling of the search space and identifies an optimal protocol through direct search methods [24,25].

The optimization procedure depends strongly on the structure of the underlying quantum control landscape; the latter is defined by an objective that measures the performance of solving the optimization task and is a functional of the protocol configuration [26]. Structural properties of the landscape, such as the number of local minima or saddle points, depend on both the controlled quantum system and the control variables [27,28]. For instance, the landscape associated with ground state preparation in a chain of qubits exhibits an exponential number of almost-optimal control protocols as local landscape minima, and a connection between quantum control landscapes and glassy systems was established [29]. However, despite the large body of previous work in quantum control, surprisingly little is known about the properties of control landscapes associated with experimentally relevant quantum control problems [27,30–34].

The bridge between quantum control landscapes and the physics of glassy systems becomes apparent once the cost function (defining the landscape) is interpreted as an effective energy function, with control variables playing the role of classical degrees of freedom [35–38]: From this point of view, optimal protocols correspond to ground states of the effective energy function. In particular, once the cost function is expanded in terms of the control variables, the n th-order expansion coefficient specifies the strength of the n -body effective interaction among control degrees of freedom. So far, the calculation of different terms in the expansion has required the (brute-force) exhaustive exploration of the entire control

space [29], a procedure computationally feasible only for a small number of control parameters. Nonetheless, this approach captures two important features of these effective interactions: They are long-range and multibody, directly reflecting the nonlocality and nonlinearity of the underlying quantum control problem.

Curiously, as the protocol duration is increased, quantum control landscapes have been shown to exhibit critical structural changes, called “control landscape phase transitions” (CLPTs) [29,39–41], in both few- and many-body controlled systems. The name derives from sharp nonanalytic behavior arising at the transition point in an order parameter that quantifies the correlations between minima in the optimization landscape. In light of the connection mentioned above between optimization problems and classical statistical physics, CLPTs are naturally associated with phase transitions occurring in the classical effective model. Moreover, since CLPTs are related to structural changes in the optimization landscape, they can affect the complexity of the underlying control problem. Therefore, an accurate analysis of CLPTs is an important step toward developing a better understanding of quantum control landscapes and shall provide new insights for improving existing optimization algorithms. Despite the clear numerical evidence for their existence, no analytical characterization of CLPTs has hitherto been developed.

In this work, we combine analytical and numerical methods to investigate different CLPTs occurring in prototypical quantum control setups.

First, we propose different controlled approximations for the quantum control landscape and analyze to what extent they reproduce the CLPTs present in the exact landscape. In particular, the central object is the evolution operator, which we expand using different analytical expansions (Dyson, Magnus, and cumulants); the cost function landscape (i.e., the infidelity in the case of state preparation or a physical observable in the case of observable optimization) is then obtained via projection onto the initial and target quantum states. Remarkably, the analytical expansions decouple the controlled quantum degrees of freedom from the control fields; hence, the framework remains general and valid for arbitrary quantum systems. From this result, we conjecture that CLPTs are ubiquitous in optimal control and can arise in both quantum and classical dynamical systems. Finally, our method treats piecewise continuous and bang-bang protocols on equal footing: The coefficients in the landscape expansions do not depend on the family of control protocols, but the landscapes do.

Second, we use the analytical expansion of the landscape to characterize CLPTs in single- and two-qubit setups. By working in a continuous protocol space, we make a precise connection between nonanalytic behaviors of order parameters and structural rearrangements in the control landscape. In particular, we keep track of optimal protocols while increasing adiabatically the duration of the quantum

evolution; this reveals CLPTs to be associated with different kinds of instabilities of the optimal protocols present in the landscape. The method is independent of the controlled quantum system and provides additional means to detect and characterize CLPTs.

Third, the landscape expansions allow us to estimate analytically the critical scaling behavior of order parameters at the quantum speed limit phase transition. To this end, we first frame the problem in the context of statistical field theory by defining an appropriate partition function. Then, we evaluate the partition function, defined using a path integral, within a Gaussian approximation.

Our work is a first attempt to apply methods from statistical field theory, designed to investigate generic complex optimization landscapes [38,42–53], to the study of quantum control landscapes [5,27,28]. We directly address the many-body character of the time-dependent *control fields*. Although we focus on one- and two-body models, our methods are formally independent of the properties of the controlled quantum system, and we discuss potential applications to many-body controlled systems. For this reason, the framework we present provides a starting point for analyzing more complex systems.

This paper is organized as follows. In Sec. II, we introduce the control problems we later use to benchmark our theory and discuss the corresponding control phase diagrams. In Sec. III, we derive the analytical expansions for the quantum control landscape and discuss their interpretation from the point of view of classical spin models in statistical physics. In Sec. IV, we discuss how CLPTs are connected to structural changes in the optimal level set, i.e., the set of control protocols that minimize the cost function. In particular, using linear and quadratic stability analysis, we detect and classify different CLPTs. In Sec. V, we introduce a sampling algorithm based on the Metropolis-adjusted Langevin dynamics, revisit the quantum speed limit (QSL) phase transition for piecewise continuous protocols, and study the properties of the optimal level set beyond the QSL. In Sec. VI, we estimate the critical scaling of the order parameters at the QSL phase transition, in cases of both continuous and bang-bang protocols. Finally, in Sec. VII, we comment on the challenges in extending the methods we develop to controlled many-body quantum systems. We summarize the main results and discuss their physical significance in Sec. VIII.

II. TOY MODELS FOR QUANTUM CLPTS

Consider the generic control problem

$$\hat{H}(t) = \hat{H}_0 + s(t)\hat{H}_1, \quad (1)$$

where the Hamiltonian is divided between a constant “drift” term $\hat{H}_0$ and a time-varying “control” term

$s(t)\hat{H}_1$. The control parameter s , regarded as a function of time t , is called control “protocol” (or “schedule”). In the qubit systems we consider below, the protocol modulates the external magnetic field in time, but more generic controls are also frequently used in experiments. To stay away from the trivial adiabatic limit, we consider a fixed protocol duration T of the order of the inverse energy scale in the drift term, i.e., $T \sim \|\hat{H}_0^{-1}\|$, and track evolution times $t \in [0, T]$. Consistently with typical limitations in present-day experiments [54], we assume a bounded control term $\hat{H}_1$ and restrict the problem to the class of bounded functions $|s(t)| \leq 1$, $\forall t \in [0, T]$.

We are interested in CLPTs emerging in state-preparation problems. Let us denote the initial state of the system by $|\psi_0\rangle$ and the target state (i.e., the state we want to prepare) by $|\psi_*\rangle$. Notice that the control term $s(t)\hat{H}_1$ modifies the bare drift evolution of the system. The optimal control problem then consists in finding those protocols $s_*(t)$ that achieve a maximum overlap between the target $|\psi_*\rangle$ and the time-evolved state $|\psi(t)\rangle = \hat{U}_s(t, 0)|\psi_0\rangle$ at the final time $t = T$; in other words, given T , we want to find (all) the protocols $s(t)$ that minimize the so-called “infidelity”

$$I(T)[s] = 1 - \left| \langle \psi_* | \hat{U}_s(T, 0) | \psi_0 \rangle \right|^2, \quad (2)$$

where $\hat{U}_s(T, 0) = \mathcal{T}e^{-i \int_0^T dt \hat{H}(t)}$ is the time-ordered evolution operator, dependent on the time-varying protocol s .

The infidelity defines a functional over protocol space (i.e., the space of all control functions), $s \mapsto I(T)[s]$, henceforth referred to as the “quantum control landscape”; the latter depends on both the specifics of the Hamiltonian $\hat{H}(t)$ and the details of the controlled system. Each point s of the infidelity landscape corresponds to a unique protocol $s(t)$; optimal protocols by definition reside in the (possibly many) global infidelity minima. By introducing both analytical and numerical techniques, we analyze changes in the structure of minima in the control landscape in few-qubit systems, as we vary the protocol duration T .

The quantum state preparation problem associated with few-qubit Hamiltonians of the form (1) was first studied in Refs. [29,40]. In particular, the search for optimal protocols was restricted to the so-called “bang-bang” class, where $|s(t)|$ takes its maximum allowed value. This restriction is formally motivated within the framework of Pontryagin’s maximum principle, which, under certain conditions, guarantees the existence of at least one optimal protocol within the bang-bang class [55]; that said, some control problems have optimal protocols mixing bang-bang and continuous parts [56]. Concretely, the time axis $[0, T]$ is divided in N equidistant time steps $\{t_1, \dots, t_N\}$ and, hence, $s(t_i) = s_i = \pm 1$ at each time step t_i . Optimal control algorithms, such as stochastic descent, can then be used

to collect a set of M (locally) optimal protocols $\mathcal{S} = \{s^{(1)}, \dots, s^{(M)}\}$ as a probe for the structure of minima in the quantum control landscape.

The piecewise-constant bang-bang protocol $s = \{s_i\}$ can be regarded as classical Ising spin degrees of freedom, while the infidelity landscape $I(T)[s]$ defines an effective Ising energy function. The corresponding spin model

$$I(T)[s] = c(T) + \frac{1}{N} \sum_{j=1}^N b_j(T) s_j + \frac{1}{2N^2} \sum_{i,j} J_{ij}(T) s_i s_j + \dots \quad (3)$$

can exhibit CLPTs, numerically investigated in Ref. [39]. Note that, despite the discretization of time, there is no underlying regular lattice structure; i.e., the control landscape is effectively an infinite-dimensional system (although discretized time induces a natural ordering of the lattice sites).

CLPTs can be detected using the order parameter

$$q(T) = \frac{1}{N} \sum_{i=1}^N [\langle s_i^2 \rangle_{\mathcal{S}} - \langle s_i \rangle_{\mathcal{S}}^2], \quad (4)$$

which is closely related to the Edwards-Anderson order parameter used to measure spin-glass order [38,53]. Here, $\langle s_i \rangle_{\mathcal{S}}$ and $\langle s_i^2 \rangle_{\mathcal{S}}$ represent the average on site i over the set of (locally) optimal protocols $\mathcal{S}$. By definition, $q(T)$ quantifies fluctuations among the set $\mathcal{S}$ with respect to the average protocol $\langle s \rangle_{\mathcal{S}}$. Within the bang-bang family, we can further simplify $q(T)$ to

$$q_{\text{BB}}(T) = 1 - \frac{1}{N} \sum_{i=1}^N \langle s_i \rangle_{\mathcal{S}}^2. \quad (5)$$

For convex landscapes, the set $\mathcal{S}$ contains a single protocol, and we have $q_{\text{BB}}(T) \equiv 0$. On the other hand, whenever protocols in $\mathcal{S}$ are uncorrelated, we have $\langle s_i \rangle_{\mathcal{S}}^2 = 0$, and $q_{\text{BB}}(T) \equiv 1$. Intermediate values of $q(T)$ correspond to correlations between protocols in $\mathcal{S}$.

Little is known about the nature and origin of control phase transitions, in general. Despite their formal resemblance to k -SAT problems and spin glasses [38], their analytical study is hampered by the lack of techniques to determine the form of the coupling constants c , b_j , and J_{ij} in the energy function $I(T)[s]$. A primary reason for this is the infinite character of the energy function in Eq. (3), comprising both long-range and many-body terms, which pose a daunting challenge for analytical studies. In turn, although numerical methods allow us to sample the control landscape and map out the corresponding phase diagram, they prove insufficient when it comes to determining and

understanding the detailed mechanisms that drive these phenomena.

In this paper, we introduce methods formally independent of the specific quantum system and focus our analysis on two few-qubit toy models. Before diving into the details, we briefly revisit their control phase diagrams and summarize the properties of the existing control phases.

A. Single-qubit control landscape

Consider the single-qubit system

$$\hat{H}(t) = h_z \hat{S}^z + s(t) h_x \hat{S}^x, \quad (6)$$

where $\hat{S}^\alpha$, $\alpha = x, y, z$, are the spin-1/2 operators. The control problem we investigate aims to prepare the target state $|\psi_*\rangle = |\text{GS}(h_x/h_z = -2)\rangle$, starting from the initial state $|\psi_0\rangle = |\text{GS}(h_x/h_z = 2)\rangle$; $|\text{GS}(h_x/h_z)\rangle$ denotes the ground state of the Hamiltonian specified by the ratio h_x/h_z (with $s = 1$). In addition, we fix the protocol duration T and set, throughout the evolution, $h_z = -1$, $h_x = -\sqrt{5}$ [i.e., $s(t)$ is the only time-dependent parameter]. Even though precise values of the control problem parameters affect the CLPTs' locations, the qualitative picture remains unaffected.

In Fig. 2(a), we show the corresponding bang-bang control phase diagram as a function of the parameter T . For $T < T_c$, the order parameter vanishes ($q_{\text{BB}} = 0$), and, hence, there exists a single optimal protocol. At $T = T_c$, the bang-bang infidelity landscape suddenly acquires additional local minima, associated with a cusp in q_{BB} . While a simple physical picture shows that at the critical duration T_c a *singular arc* appears in the optimal protocol [3,4,39], little is known about how the underlying microscopic spin model $I(T)[s]$ determines the properties of $q_{\text{BB}}(T)$, including the critical exponent at the transition. For $T_c \leq T \leq T_{\text{QSL}}$, the bang-bang infidelity landscape exhibits, in addition to a unique global minimum, multiple almost-optimal local minima; the correlations among these minima give rise to a finite value $0 < q_{\text{BB}} < 1$. At the quantum speed limit $T = T_{\text{QSL}}$, the landscape undergoes a second phase transition where the number of global minima proliferates; a second nonanalytic point in $q_{\text{BB}}(T)$ emerges. It is currently unclear how the presence of a finite quantum speed limit modifies the structure of the energy function $I(T)[s]$ nor how these changes affect the behavior of the order parameter $q_{\text{BB}}(T)$.

B. Two-qubit control landscape

While the control landscape of the single-qubit problem in Eq. (6) is not fully understood, finding the optimal protocol is rather straightforward due to the simplicity of the qubit system [57]. Therefore, we also investigate a two-qubit model




FIG. 2. CLPTs in the single- (a) and two-qubit (b) control problems. The curves shown are obtained from the stochastic descent (SD) algorithm in the exact infidelity landscape, searching in the space of bang-bang protocols. The minimum infidelity $\min_s I(T)[s]$ and the order parameter $q_{\text{BB}}(T)$ [cf. Eq. (5)] are shown as a function of the protocol duration T for a fixed number of bang-bang steps. Nonanalytic points in the order parameter $q_{\text{BB}}(T)$ mark control landscape phase transitions (dashed vertical lines). Similar transitions appear in both control problems: at $T = T_c$ and at $T = T_{\text{QSL}}$ (where the infidelity vanishes). The two-qubit problem possesses an additional symmetry-breaking transition at $T = T_{\text{sb}}$ related to the separation of the set of optimal protocols into two separated subsets which break the $s(t) \leftrightarrow -s(T-t)$ symmetry of the control problem [40].

$$\hat{H}(t) = J\hat{S}_1^z\hat{S}_2^z + h_z(\hat{S}_1^z + \hat{S}_2^z) + s(t)h_x(\hat{S}_1^x + \hat{S}_2^x) \quad (7)$$

with h_z and h_x as in the single-qubit problem and $J = -2$. The two qubits are coupled via an Ising interaction J and are exposed to a constant global z field and a time-varying global x field. The initial and target states read as $|\psi_0\rangle = |\text{GS}(h_x/h_z = 2, J/h_z = 1)\rangle$ and $|\psi_*\rangle = |\text{GS}(h_x/h_z = -2, J/h_z = 1)\rangle$, respectively. To the best of our knowledge, the exact solution to the corresponding constrained optimal control problem in Eq. (7) is unknown. This two-qubit model will allow us to study the wider applicability of the techniques we develop.

For $J = 0$, we recover the single-qubit problem discussed above; in particular, the control phase diagram exhibits an overconstrained phase ($T < T_c$) and a fully controllable phase $T > T_{\text{QSL}}$, separated by a correlated phase $T_c \leq T \leq T_{\text{QSL}}$ [cf. Fig. 2(a) and Ref. [40]]. However, for finite values of J , additional phases appear in the control landscape.

In Ref. [40], it was shown that the infidelity function features an Ising $\mathbb{Z}_2$ symmetry:

$$I_T[s(t)] = I_T[-s(T-t)], \quad (8)$$

corresponding to the simultaneous application of a protocol time-reversal symmetry $s(t) \rightarrow s(T-t)$ and global flip of the x/y projection $\hat{S}_j^{x/y} \rightarrow -\hat{S}_j^{x/y}$. On the level of the effective classical spin model, the same $\mathbb{Z}_2$ symmetry manifests itself as a global spin flip, combined with a reflection about the center of the lattice: $s_j \rightarrow -s_{N-j}$. The presence of this symmetry hinges on the choice of initial and target states.

Interestingly, symmetry breaking occurs within the correlated phase of the two-qubit bang-bang control phase diagram [cf. Fig. 2(b)], at a critical protocol duration T_{sb} . This leads to a doubly degenerate optimal level set that features two global minima, whose protocols can be transformed into one another by the symmetry operation. Moreover, the order parameter q_{BB} exhibits a discontinuity at the symmetry-breaking transition T_{sb} . The same $\mathbb{Z}_2$ symmetry is present in the single-qubit control problem (6), as well as in the corresponding multiqubit generalization [39]; nevertheless, rather mysteriously, the corresponding landscapes do not feature a symmetry-broken phase [29]. While this phenomenon was discovered and studied numerically, at present, it is not known what the necessary and sufficient conditions for it to occur are; it is also unclear what its relation to the other control phase transitions is.

In the following, we develop and discuss techniques to derive analytical expressions for the quantum control landscape. This allows us to go beyond bang-bang protocols and investigate the corresponding landscapes defined over the piecewise continuous protocol space.

III. PERTURBATIVE EXPANSIONS FOR THE CONTROL LANDSCAPE

We now proceed to present a systematic derivation of the analytical expansion for the infidelity landscape [cf. Eq. (3)] using three different controlled techniques: the Dyson, Magnus, and cumulant expansions. We test their accuracy when truncated to the leading few orders and compare and contrast their ability to capture the transitions in the control phase diagrams of the single- and two-qubit toy models.

A. Analytical expansions for the infidelity

Consider the general case of piecewise continuous protocols $s(t)$. We can write down a general expansion of the infidelity $I(T)[s]$ functional as

$$I(T)[s] = c(T) + \int_0^T dt b_t(T) s_t + \frac{1}{2} \int_0^T dt^2 J_{t_1 t_2}(T) s_{t_1} s_{t_2} + \dots \quad (9)$$

with some *a priori* unknown coefficients $c(T)$, $b_t(T)$, $J_{t_1 t_2}(T)$, ... that depend parametrically on the protocol

duration T . At the level of the underlying effective spin model, as the notation suggests, they play the role of a constant shift, an external field, and a two-body interaction strength, respectively.

To determine these unknown coefficients, we rewrite Eq. (2) in terms of the density matrix associated with the quantum state $\hat{U}_s(T, 0)|\psi\rangle$. In an N -qubit system, a generic density matrix can be expanded as $\rho = |\psi\rangle\langle\psi| = d^{-1} + \hat{S} \cdot \vec{n}$, where $d = 2^N$ is the Hilbert space dimension and the vector $\hat{S}$ contains all traceless spin operators acting on the Hilbert space (e.g., Pauli or Gell-Mann matrices for a qubit or qutrit system, respectively; Pauli strings for a qubit chain) and we choose the normalization convention $\text{Tr}\{\hat{S}^i \hat{S}^j\} = \delta_{ij}/2$ with $i, j = 1, \dots, d^2 - 1$. Notice that to each quantum state $|\psi\rangle \in \mathbb{C}^d$ corresponds a real “dual” vector $\vec{n} \in \mathbb{R}^{d^2-1}$; in terms of the dual vector $\vec{n}$, the infidelity becomes

$$\begin{aligned} I(T)[s] &= 1 - \left| \langle \psi_* | \hat{U}_s(T, 0) | \psi_0 \rangle \right|^2 \\ &= 1 - \langle \psi_* | \hat{U}_s(T, 0) \rho_0 \hat{U}_s(T, 0)^\dagger | \psi_* \rangle \\ &= 1 - \langle \psi_* | \left(d^{-1} + \hat{U}_s(T, 0) \hat{S} \hat{U}_s(T, 0)^\dagger \cdot \vec{n}_0 \right) | \psi_* \rangle \\ &= 1 - d^{-1} - \langle \psi_* | \hat{S} | \psi_* \rangle M_s(T, 0) \vec{n}_0 \\ &= 1 - d^{-1} - (1/2)(\vec{n}_* \cdot M_s(T, 0) \vec{n}_0). \end{aligned} \quad (10)$$

The rotation matrix $M_s(T, 0)$, which we refer to as control propagator, defines the representation of the unitary evolution group on the dual space $\mathbb{R}^{d^2-1}$; in other words, the quantum evolution operator is replaced by the orthogonal evolution operator over the space of real $(d^2 - 1)$ -dimensional vectors $\vec{n}$. Explicitly, $[M_s(T, 0)]_{ij} = 2\text{Tr}(\hat{S}^i \hat{U}_s(T, 0) \hat{S}^j \hat{U}_s(T, 0)^\dagger)$.

A few remarks are in order. (i) Note that, by going to the dual description, we manage to swap the original quantum expectation values for an “expectation” in the vectors $\vec{n}_{0,*}$; in a sense, the quantum nature of the original problem is now hidden in the dual space dimension which grows faster than the qubit Hilbert space size. This indicates that any transitions in the control landscape are not quantum by nature, but merely classical. (ii) On a technical level, the above procedure allows us to eliminate the absolute value square in the definition of the infidelity; this is useful, since it will allow us to more easily find the expression for the landscape from Eq. (9). (iii) Notice that this procedure is generic and applies to control problems where the cost function is defined in terms of observables (e.g., energy minimization) that possess easy-to-identify expansions in operator space: One simply has to replace ρ_0 by the corresponding observable. Last but not least, (iv) we have managed to reduce the problem of computing the infidelity landscape to the problem of expanding the control propagator $M_s(T, 0)$ in terms of the protocol $s(t)$.

To make further progress analytically, let us write the control propagator $M_s(T, 0)$ as

$$M_s(T, 0) = \mathcal{T} e^{\int_0^T m_s(t) dt}, \quad (11)$$

where $m_s(t)$ is the antisymmetric matrix generating the rotation $M_s(T)$; assuming the relation $[\hat{S}^i, \hat{S}^j] = i f_{ijk} \hat{S}^k$ with structure constants f_{ijk} , $m_s(t)$ can be computed directly from the Hamiltonian as $[m(t)]_{ij} = \sum_k f_{ijk} 2\text{Tr}(\hat{S}^k \hat{H}(t))$.

For a general Hamiltonian of the form in Eq. (1), we can separate the time dependence in the generator as $m_s(t) = m_0 + s(t)m_1$. However, this additive structure of the generator does not yet allow for a straightforward expansion in terms of the protocol s (see Appendix A). To circumvent this problem, we now perform a change-of-frame transformation, which leads to an effective resummation of all relevant subseries [58]. To this end, we use a rotating frame transformation on the level of the generator, resulting in

$$m_s(t) \mapsto m'_s(t) = M_0(t, 0)^t [m_s(t) - \partial_t] M_0(t, 0), \quad (12)$$

with $M_0(t, 0) = \exp(tm_0)$ and $(\cdot)^t$ denotes the matrix transpose. Note that physical quantities, such as the infidelity, remain invariant. Thus, in the primed reference frame, the infidelity can be written as

$$I_s(T) = 1 - d^{-1} - (1/2)(\vec{n}_* \cdot M_0(T, 0) M'_s(T, 0) \vec{n}_0), \quad (13)$$

where $M'_s(T, 0)$ is the transformed evolution operator associated with the generator $m'_s(t)$ in the new reference frame. We can also absorb the s -independent rotation in the target state: $\vec{n}'_*(T) = \vec{n}_* M_0(T, 0)$, to simplify notation. The only remaining protocol-dependent quantity is now the control propagator $M'_s(T, 0)$; hence, in the following, we seek a matrix-valued expansion in powers of s .

1. Dyson expansion

The Dyson series associated with $M'_s(T, 0)$ reads as

$$\begin{aligned} M'_s(T, 0) &= \mathbb{1} + \int_0^T dt \partial_s m'_s(t) s(t) \\ &\quad + \frac{1}{2} \int_0^T d^2 t \mathcal{T} [\partial_s m'_s(t_1) \partial_s m'_s(t_2)] s(t_1) s(t_2) \\ &\quad + \dots \end{aligned} \quad (14)$$

We can now easily read off the infidelity expansion coefficients in Eq. (9):

$$\begin{aligned} c(T) &= 1 - d^{-1} - (1/2)(\vec{n}'_*(T) \cdot \vec{n}_0), \\ b_t(T) &= -(1/2)(\vec{n}'_*(T) \cdot \partial_s m'_s(t) \vec{n}_0), \\ J_{t_1 t_2}(T) &= -(1/2)(\vec{n}'_*(T) \cdot \mathcal{T} [\partial_s m'_s(t_1) \partial_s m'_s(t_2)] \vec{n}_0), \end{aligned} \quad (15)$$

and similarly for higher-order terms. We have, thus, obtained explicit closed-form expressions for all coefficients in the infidelity expansion, which define the control landscape. Remarkably, the whole series expansion can be explicitly written once the quantity $\partial_s m'_s(t)$ is known.

As a side note, notice that, when restricted to bang-bang protocols, a second-order truncation of the series expansion yields an effective classical spin model that bears resemblance to the well-known Hopfield network [38].

It now becomes clear that the (matrix-valued) control propagator $M'_s(T, 0)$ generates the control landscape. In particular, sandwiching it between $\vec{n}'_*$ and $\vec{n}_0$ projects to a specific landscape for *any* choice of initial and target states and ultimately produces the corresponding infidelity expansion coefficients in Eq. (9). Moreover, by choosing $\vec{n}'_*$ appropriately, we can obtain the corresponding control landscape also for observables, e.g., when minimizing the energy or maximizing the expectation of some spin component, etc. Therefore, the control propagator is a fundamental object of control landscape theory.

As we see below, the Dyson series is one example of a functional expansion for the infidelity. In particular, Eq. (15) corresponds to the infidelity functional-Taylor expansion *centered* around the protocol $s = 0$ [written in terms of the orthogonal generator $m'_s(t)$]. In Appendix A, we generalize this result to an arbitrary center $s \neq 0$ and a generic reference frame; moreover, we derive the expansion for both unitary and orthogonal evolution operators. In Sec. IV, we use these general expressions to study the near-optimal region of the control landscape.

A well-known deficiency of the *truncated* Dyson series for the evolution operator is the violation of orthogonality of the resulting approximate operator; as a result, the corresponding Dyson-expanded infidelity landscape may not be restricted to the interval $[0, 1]$.

2. Magnus expansion

A possible remedy to the unitarity or orthogonality loss is given by the Magnus expansion [59–61] for the control propagator $M'_s(T, 0)$:

$$M'_s(T, 0) = \exp \left[\sum_{n=1}^{\infty} \Omega_{n,s}(T, 0) \right], \quad (16)$$

where the first few terms read as [62]

$$\begin{aligned} \Omega_{1,s}(T, 0) &= \int_0^T dt \partial_s m'_s(t), \\ \Omega_{2,s}(T, 0) &= \frac{1}{2} \int_0^T \int_0^{t_1} dt^2 [\partial_s m'_{t_1}, \partial_s m'_{t_2}] s_{t_1} s_{t_2}, \\ \Omega_{3,s}(T, 0) &= \frac{1}{6} \int_0^T \int_0^{t_1} \int_0^{t_2} dt^3 \left([\partial_s m'_{t_1}, [\partial_s m'_{t_2}, \partial_s m'_{t_3}]] \right. \\ &\quad \left. + [[\partial_s m'_{t_1}, \partial_s m'_{t_2}], \partial_s m'_{t_3}] s_{t_1} s_{t_2} s_{t_3} \right). \end{aligned}$$

Owing to its structure, truncating the Magnus series at finite order yields an orthogonal operator, and the resulting fidelity respects the $[0, 1]$ interval bound. Notice that the exponential map acting on the truncated expansion generates infinitely many terms that are absent in the Dyson expansion. However, the Magnus expansion does not allow for a direct scalar effective energy interpretation, since its terms are matrix valued [i.e., generators of the rotation group $\text{SO}(N^2 - 1)$].

3. Cumulant expansion

Substituting the Magnus expansion in Eq. (10), we find

$$I(T)[s] = 1 - d^{-1} - (1/2) \vec{n}'_* \cdot \exp(\Sigma) \vec{n}_0,$$

where Σ compactly denotes the Magnus expansion. Rewriting $I(T)[s] = 1 - e^{\log\{1 - I(T)[s]\}}$, we can expand the logarithm using the so-called cumulant expansion as

$$\begin{aligned} \log\{1 - I(T)[s]\} &= \log \frac{2d^{-1} + \vec{n}'_* \cdot \exp(\Sigma) \vec{n}_0}{2} \\ &= \sum_{m=0}^{\infty} \frac{\kappa_m}{m!}. \end{aligned} \quad (17)$$

This expansion is common in statistical mechanics, where the free energy F , linked to the partition function Z through the identity $-\beta F = \log Z$, is typically expanded in cumulants. In the present case, the fidelity $1 - I(T)[s]$ plays the role of the partition function. The first few terms of the cumulants expansion read as

$$\begin{aligned} \kappa_0 &= \log \frac{2d^{-1} + \vec{n}'_* \cdot \vec{n}_0}{2}, \\ \kappa_1 &= \frac{\vec{n}'_* \cdot \Sigma \vec{n}_0}{2d^{-1} + \vec{n}'_* \cdot \vec{n}_0}, \\ \kappa_2 &= \frac{\vec{n}'_* \cdot \Sigma^2 \vec{n}_0}{2d^{-1} + \vec{n}'_* \cdot \vec{n}_0} - \left(\frac{\vec{n}'_* \cdot \Sigma \vec{n}_0}{2d^{-1} + \vec{n}'_* \cdot \vec{n}_0} \right)^2. \end{aligned}$$

We now turn to the quality of the three infidelity expansions presented. Since we want to develop a technique to better understand the near-optimal region of the quantum control landscape, it is important to quantify the extent to which the infidelity expansions reproduce the behavior of already known quantities, such as the nonanalytic points in the order parameter $q_{\text{BB}}(T)$ across control phase transitions. For this reason, in the following, we compare the three expansions using the two toy models introduced in Sec. II.

B. Single-qubit problem

So far, the discussion has been carried out for a generic quantum system. In the case of Hamiltonian (6), one obtains

$$m'_s(t) = h_x s(t) [T_x \cos(h_z t) + T_y \sin(h_z t)], \quad (18)$$

where $\{T_i\}_1^3$ are the generators of the group $\text{SO}(3)$ defined by $(T_i)_{jk} = \epsilon_{ijk}$, with ϵ_{ijk} the $\text{SO}(3)$ structure constants. We emphasize that having a closed-form expression for $m'_s(t)$ is sufficient to generate *any* (multibody) coupling term in the Dyson expansion for the infidelity landscape.

In Fig. 3, we evaluate the performance of the three infidelity expansions using stochastic descent (see Appendix B). Since our motivation to introduce the expansions is the study of the quantum landscape transitions, we benchmark the three methods using the order parameter $q_{\text{BB}}(T)$, introduced in Eq. (5). In our simulations, we truncate Dyson and Magnus expansions up to third order.

The first transition around T_c is visible in all three expansions within the truncation orders considered. While the Dyson expansion is not able to capture the T_{QSL}




FIG. 3. Comparison of the order parameter $q_{\text{BB}}(T)$ [cf. Eq. (5)] obtained from SD in the landscape defined by the infidelity expansions introduced in Sec. III, in the single- and two-qubit control problem. Control protocols are here restricted to the bang-bang class. The black curve refers to the exact infidelity landscape in Eq. (2). The letters D, M, and C in the legend stand for Dyson, Magnus, and cumulant expansion, respectively, while the adjacent integer specifies the order of truncation. Single-qubit: the $T = T_c \simeq 0.98$ transition is approximately captured by all three methods. The second transition at $T = T_{\text{QSL}} \simeq 2.51$ is visible in only the Magnus and cumulant results. Two-qubit: the $T = T_c \simeq 0.56$ and $T = T_{\text{sb}} \simeq 1.57$ transitions are approximately captured by all three methods. The nonanalytic point at $T = T_{\text{QSL}} \simeq 2.95$ is less sharp than the corresponding point in the single-qubit problem, and only the Magnus expansion truncated at third order is able to reproduce the qualitative behavior of the exact curve.

transition at third order, the Magnus and cumulant expansion results follow remarkably well the exact curve $q_{\text{BB}}(T)$ (at third and fifth order, respectively). Thus, we conclude that higher than third-order terms, generated in the Magnus expansion from the exponential map but completely absent in the other two expansions, are essential to reproduce CLPTs in the exact infidelity landscape.

C. Two-qubit problem

Let us apply our theory to the two-qubit problem from Eq. (7). The rotating frame is defined by removing the drift term $\hat{H}_0$ from the Hamiltonian. In this case, the drift term has two contributions $\hat{H}_0 = \hat{H}_{0,1} + \hat{H}_{0,2}$:

$$\hat{H}_{0,1} = h_z \hat{S}_{\text{tot}}^z, \quad \hat{H}_{0,2} = JS_1^z S_2^z.$$

Since $[\hat{H}_{0,1}, \hat{H}_{0,2}] = 0$, the reference frame transformation factorizes as $\hat{U}_0 = \hat{U}_{0,1} \hat{U}_{0,2}$. A straightforward calculation shows that

$$\hat{H}'(t) = s(t) h_x \hat{U}_{0,2}^\dagger \left[\cos(h_z t) \hat{S}_{\text{tot}}^x - \sin(h_z t) \hat{S}_{\text{tot}}^y \right] \hat{U}_{0,2},$$

where $\hat{U}_{0,2}$ is a diagonal operator in the computational z basis.

In this case, the two-qubit quantum control problem allows for an orthogonal decomposition of the Hamiltonian operator into the triplet and singlet subspaces; moreover, the quantum states $|\psi_{0,*}\rangle$ have no component in the singlet subspace. Hence, we may restrict $\vec{S} = \vec{\lambda}/2$ to the $\text{SU}(3)$ basis defined by Gell-Mann matrices $\vec{\lambda}$. Explicitly, in the dual description, we have

$$m'_s(t) = s(t) \sqrt{2} h_x [T_1 \cos(h_z t + Jt/2) - T_2 \sin(h_z t + Jt/2) + T_6 \cos(h_z t - Jt/2) - T_7 \sin(h_z t - Jt/2)], \quad (19)$$

where the generators $\{T_i\}_1^8$ are given by $(T_i)_{jk} = f_{ijk}$ with f_{ijk} the $\text{SU}(3)$ structure constants.

Stochastic descent results for the two-qubit problem are reported in Fig. 3. Both the T_c and T_{sb} transitions are captured by the three expansions. In the latter case, the jump discontinuity of $q_{\text{BB}}(T)$ is smoothened out by the approximated landscape. On the contrary, the T_{QSL} transition is partially captured only by the Magnus expansion truncated at third order.

A final remark is in order. In this section, we used the infidelity expansions, centered around $s_0 = 0$, to detect CLPTs of the exact control landscape. The choice of the center $s_0 = 0$ is convenient, since it allows us to make no extra assumption regarding the shape of the optimal protocol(s) of the landscape. Nonetheless, as we see in the next section, better results can be achieved by centering the expansion around carefully chosen protocols ($s_0 \neq 0$).

This can be done whenever additional information regarding the properties of the optimal protocol(s) is known.

D. Landscape expansions in generic quantum control problems

Let us now comment on the potential applicability of landscape expansions to study CLPTs in other quantum control problems.

In Appendix A, we discuss the convergence properties of the landscape expansions introduced in this section and show that they depend on the norm of the generator integrated over the total evolution, namely,

$$\int_0^T dt \|H[s(t)]\| \quad \text{or} \quad \int_0^T dt \|m[s(t)]\|. \quad (20)$$

In the case of bounded protocols, the integral can be bounded by the product $T \max_t \{\|H[s(t)]\|\}$. The applicability of the techniques developed in this section for generic quantum control problems depends on this quantity. Physically, $\|\hat{H}(t)\|$ depends on the size of the quantum system, the interactions between its components, and the amplitude of the external control fields.

In particular, if one is interested in approximating the exact landscape within a given error (uniformly in the space of control protocols), the order of truncation of the expansion grows with the duration of evolution and the norm of the Hamiltonian $\|\hat{H}(t)\|$.

Nevertheless, when one is interested in studying CLPTs, it is not necessary to work with a uniform approximating functional. To see this, we observe that (i) CLPTs are determined by a subset of all the allowed protocols, namely, minima of the infidelity, and (ii) the numerical approximation of the infidelity is not required, as we are interested in the relative location of minima to capture the relevant physics of CLPTs, not necessarily in their infidelity values. Therefore, the ultimate limit for the applicability of landscape expansions in the study of CLPTs depends on the time-integrated norms in Eq. (20), restricted to the near-optimal region of the control landscape.

IV. STABILITY ANALYSIS AND CONTROL LANDSCAPE PHASE TRANSITIONS

In the previous section, we introduced three different analytical expansions for the control landscape, centered around the protocol $s = 0$. In Appendix A, we discuss the relation with the functional Taylor expansion. This section uses the infidelity Taylor expansion to analyze CLPTs in the single- and two-qubit control problem introduced in Sec. II. Our purpose is twofold. On the one hand, critical points in the control phase diagram are identified with different types of instability relative to the minimum of the quantum control landscape. On the other, we consider the Taylor expansion as a function of the parameter T to

identify and follow the landscape minimum. In particular, we first use the $T \rightarrow 0$ limit to analyze the landscape in a linear approximation and to identify the structure of the optimal protocol in this limiting case. Subsequently, we track how the optimal protocol changes as we increase T by enforcing “linear stability.”

The idea can be understood by using a one-dimensional analogy (see Fig. 4). Let $I = [-1, 1] \subset \mathbb{R}$ and $f: I \rightarrow \mathbb{R}$ be a convex function with minimum $x_0 \in I$. The Taylor expansion centered in $x = x_0$ and truncated at second order reads $f(x) = f(x_0) + ([df]/[dx])|_{x_0}(x - x_0) + (1/2!)([d^2f]/[dx^2])|_{x_0}(x - x_0)^2$. If x_0 does not belong to the boundary $\{-1, 1\}$, then the linear term vanishes $([df]/[dx]) = 0$ and the second-order term is positive $([d^2f]/[dx^2])(x_0) > 0$. Otherwise, if x_0 belongs to the boundary $\{-1, 1\}$, the linear term may be finite, but then either $f'(x) > 0$ if $x_0 = -1$ or $f'(x) < 0$ if $x_0 = 1$; notice that the two cases are summarized by the condition $\text{sgn}[f'(x_0)] = -\text{sgn}(x_0)$, with sgn the sign function. Finally, if we let $f(x) = f_T(x)$ to also continuously depend on a parameter T , we may then study the behavior of its minimum x_0 as a function of T .

The control landscape [represented by the function $f_T(x)$ in the above analogy] is a functional of the protocol s and depends on the duration of the quantum evolution $T \in [0, \infty)$. Our strategy is based on the continuous dependence of infidelity on the parameter T and relies on two ideas. First, once an optimal protocol at a given T is found, it can be used as an *Ansatz* to search for optimal protocol(s) at $T' = T + \Delta T$. We exploit this idea in two independent ways: (i) analytically, by formulating a variational *Ansatz* for the optimal shape of the protocol at T' using the optimal shape at T , and (ii) numerically, by using the numerical optimum at T as a starting point for the new




FIG. 4. Two different one-dimensional functions (blue and red curves) illustrating the stability conditions for optimal points in Eq. (21). The minimum of the red curve at $x = +0.5$ (red dot), lying inside the allowed domain $[-1, 1]$, satisfies $f'(x) = 0$ and $f''(x) > 0$. The minimum of the blue curve at $x = -1$ (blue dot), lying at the edge of the allowed domain $[-1, 1]$, satisfies $f'(x) > 0$.

run of the optimization algorithm at T' . Second, at a fixed T , we use the infidelity expansion around a given optimal protocol s_0 to study its stability under generic perturbations $s_0 + \delta s$. In this case, the stability (instability) of the protocol s_0 depends on whether the infidelity increases (decreases) after the perturbation $s + \delta s$ is applied.

We demonstrate that this procedure (which we dub “adiabatic tracing” below) is sufficient to map out the control phase diagram up to T_{QSL} in the one- and two-qubit problems, respectively (cf. Fig. 2). In addition, notice that the control phase diagram was originally obtained by searching for optimal configurations in the space of bang-bang protocols [39,40]. Here, we go beyond this restriction and study the landscape over the space of bounded piecewise continuous protocols, $s(t) \in [-1, 1]$. Interestingly, the phenomenology observed in the bang-bang case extends over to this larger space, but with a few caveats: In Sec. VI B, we comment on the relation between the control problem defined over bang-bang and bounded piecewise continuous protocol spaces.

A. Adiabatic tracing in quantum control landscapes

Observe that, when expanding the infidelity around a generic protocol $s_0(t)$, the coefficients $c(T)$, $b_t(T)$, and $J_{t_1 t_2}(T)$ in Eq. (9) depend on s_0 itself: In this section, we make this dependence explicit but also keep track of the parametric dependence on the protocol duration T :

$$c[s](T), \quad b_t[s](T), \quad J_{t_1 t_2}[s](T), \quad \text{etc.}$$

We define the linear stability of the protocol s_0 with respect to the infidelity landscape $I[s](T)$ analogously to the one-dimensional example $f_T(x)$ discussed above. The protocol s_0 is “linearly unstable” if

$$\begin{aligned} b_t[s_0](T) &\neq 0, \quad |s_{0,t}| < 1, \\ -\text{sgn}\{b_t[s_0](T)\} &\neq \text{sgn}(s_{0,t}), \quad |s_{0,t}| = 1, \end{aligned} \quad (21)$$

for some $t \in [0, T]$. The two cases distinguish the situations where $s(t)$ is inside $[-1, 1]$ or at its boundary (see Fig. 4).

As in the one-dimensional example for nonconvex functions, one needs, in general, to define also the quadratic stability of the protocol s_0 . In this case, the second-order coefficient $J_{t_1 t_2}[s_0](T)$ plays the role of the Hessian matrix (i.e., the second-order term in the Taylor expansion) for multidimensional functions. In particular, the stability of s_0 depends on the set of eigenvalues $\{\lambda_n\}$, defined by [63]

$$\frac{1}{T} \int_0^T dt' J_{t't} f_t^{(n)} = \lambda_n f_t^{(n)}, \quad (22)$$

with $\{f_t^{(n)}\}_n$ the eigenfunctions (the dependence on s_0 and T is omitted for simplicity). Then, the protocol s_0 is “quadratically unstable” if there are negative eigenvalues

in the spectrum of the “Hessian operator” $J_{t_1 t_2}[s_0](T)$. In practice, we can perform the diagonalization numerically. Discretizing the time axis in L equal-size intervals, we find the matrix eigenvalue equation:

$$\frac{1}{L} \sum_{j=1}^L J_{ij} f_j^{(n)} = \lambda_n f_i^{(n)}, \quad n = 1, 2, \dots, L.$$

In this case, the number of eigenvalues is controlled by the number of discrete segments L we divide our protocol into and is not strictly physical. We report in Appendix C the finite- L scaling analysis.

We are now equipped and in a position to introduce each step of the method and discuss the corresponding results for the single- and two-qubit control problems. As a starting point, we consider the limit $T \rightarrow 0^+$ where the infidelity landscape can be approximated by a Taylor series truncated to first order (throughout our work, we use the $+$ and $-$ superscripts to distinguish the right or left limit, respectively). Expanding around the protocol $s = 0$,

$$I(T)[s] \approx I_T[0] + \int_0^T dt b_t[0](T) s(t),$$

one obtains the infidelity minimum $s_0(t) \equiv -\text{sgn}[b_t[0](T)]$. In fact, the above infidelity expansion can be interpreted as an effective classical energy $E(s) \equiv I(T)[s]$ of the field $s(t)$ in which $b_t[0](T)$ plays the role of an external field applied to the variable $s(t)$. Since this is a noninteracting problem, energy is trivially minimized when $s(t)$ has maximum absolute value and is antialigned with $b_t[0](T)$ [64]. In this way, in both the one- and two-qubit problems in Eqs. (6) and (7), for $T \rightarrow 0$ we find the optimal protocol

$$s_0(t) \equiv \begin{cases} +1 & t < T/2 \\ -1 & t > T/2. \end{cases} \quad (23)$$

B. Single-qubit problem

The starting point for our analysis is the protocol s_0 defined in Eq. (23). As a first step, we check its linear stability. In Fig. 5, we show the first-order term $b_t[s_0](T)$ of the infidelity Taylor series centered at s_0 , for $T = 0.25$ and $T = T_c$. In particular, we observe that s_0 is linearly stable for $T \in [0, T_c]$. Exactly at $T = T_c$, $b_t[s_0](T)$ (regarded as a function of the variable t) inverts its sign around the point $t = T/2$: s_0 has a linear instability for $T > T_c$ around the point $t = T/2$. In this way, we associate the T_c control landscape phase transition with linear instability of the protocol s_0 , optimal for $T \in [0, T_c]$.

Since the instability affects only a local neighborhood of $t = T/2$, it is suggestive to look for the new infidelity minimum in the region $T = T_c^+$ by locally (i.e., around $t = T/2$) modifying the previous optimal protocol $s_0(t)$. For $T = T_c$, in the neighborhood of $t = T/2$, the first-order




FIG. 5. Stability analysis in the single-qubit control problem. The plots show the linear (b_t , orange curve) and second-order terms ($J_{t_1 t_2}$, color map insets) of the infidelity Taylor expansion centered around optimal piecewise continuous protocols (s_t , blue curve) found by the adiabatic tracing method (cf. Sec. IV), for different values of the duration T of the quantum evolution (CLPTs in $T_c \simeq 0.98, T_{\text{QSL}} \simeq 2.51$). Linear term: b_t . For $T \in [0, T_c]$ the optimal protocol is s_0 in Eq. (23). (a) At $T = 0.25$, s_0 is antialigned with b_t and is linearly stable. (b) For $T = T_c$, s_0 becomes linearly unstable around $t = T/2$ as shown by the horizontal inflection of b_t around $t = T/2$. (c) For $T \in [T_c, T_{\text{QSL}}]$, the optimal protocol is $s_{\Delta_0(T)}$ in Eq. (24). At $T = 2.0$, s_{Δ} is stable provided Δ is chosen as Eq. (24) [cf. Fig. 6(a)]. (d) At $T = T_{\text{QSL}}$, the first-order term is exactly zero, and it flips sign for $T = T_{\text{QSL}}^+$: $s_{\Delta_0(T)}$ becomes linearly unstable. Quadratic term: $J_{t_1 t_2}$. In the language of statistical mechanics, the second-order term represents the effective two-body interactions of the field s_t . In this case, interactions are antiferromagnetic (AFM) and long-range: the system is frustrated. For $T = T_c$ (b), the protocol s_0 becomes linearly unstable around $t = T/2$, and the local AFM interactions around the instability point dictate the new optimal protocol structure $s_{\Delta(T)}$ [cf. Eq. (24)].

term $b_t[s_0](T)$ has a negligible contribution to the infidelity compared to the second-order term $J_{t_1 t_2}[s_0](T)$. Moreover, using the language of statistical physics, we can interpret $J_{([T_c/2], ([T_c/2])}[s_0](T_c) > 0$ as an antiferromagnetic interaction [cf. $T = T_c$ case in Fig. 5(b)]. This observation allows us to make an educated guess for the optimal protocol in the parameter region $T > T_c$:

$$s_{\Delta}(t) \equiv \begin{cases} 0 & t \in [\frac{T}{2} - \Delta, \frac{T}{2} + \Delta] \\ s_0(t) & \text{elsewhere,} \end{cases} \quad (24)$$

where $\Delta \in [0, T/2]$ is a free parameter in the *Ansatz* that may depend on the protocol duration T itself [65].

Past the critical point T_c , the family of protocols s_{Δ} parametrized by $\Delta \in [0, T/2]$ contains the optimal protocol. Searching for linearly stable protocols within the family s_{Δ} fixes Δ for each T . In fact, this constraint defines the curves $\Delta_0(T)$ and $\Delta_{\pm}(T)$ shown in Fig. 6(a) and selects the stable protocols among the family s_{Δ} . For $T \in [T_c, T_{\text{QSL}}]$, we find only one linearly stable protocol within this family. On the other hand, a bifurcation in the curve $\Delta(T)$ occurs which coincides exactly with the critical point at $T_{\text{QSL}} \simeq 2.51$, marking the end of the single-valued regime. More precisely, we find perfect agreement with the analytical behavior

$$\Delta(T) = \begin{cases} 0 & T \in [0, T_c], \\ \frac{1}{2}(T - T_c) & T \in [T_c, T_{\text{QSL}}], \\ \frac{1}{2}(T - T_c) + \delta_{\pm}(T) & T \gtrsim T_{\text{QSL}}, \end{cases} \quad (25)$$

where $\delta_{\pm}(T) \sim \pm(T - T_{\text{QSL}})^{1/2}$ as $T \rightarrow T_{\text{QSL}}^+$. Hence, for later convenience, let us define

$$\Delta_0(T) = (T - T_c)/2, \quad \Delta_{\pm}(T) = \Delta_0(T) + \delta_{\pm}(T). \quad (26)$$

The presence of the bifurcation at T_{QSL} implies that the infidelity landscape possesses two distinct linearly stable




FIG. 6. Stability analysis in the single-qubit control problem. (a) Δ parametrizes the family of protocols s_{Δ} in Eq. (24). $\Delta(T)$ is defined by linear stability of the corresponding protocol. For $T \in [0, T_{\text{QSL}}]$, $\Delta_0(T)$ is single valued: within the family s_{Δ} there is a single optimal protocol. The bifurcation $\Delta(T) = \Delta_{\pm}(T)$ occurs at $T = T_{\text{QSL}}$ so that for $T > T_{\text{QSL}}$ $\Delta(T)$ is double valued: within the family s_{Δ} there are two optimal protocols. The gray area corresponds to $\Delta > T/2$ and is inaccessible by definition. (b) Eigenvalues of the quadratic operator $J_{t_1 t_2}$ from the infidelity Taylor expansion centered around the optimal protocol $s_{\Delta(T)}$. For $T \rightarrow T_{\text{QSL}}$, $L - 2$ eigenvalues (blue solid curves) vanish: the protocol $s_{\Delta_0(T)}$ becomes quadratically unstable. The largest positive eigenvalue, approximately constant and with magnitude ~ 1 , is shown in Appendix C. For $T \geq T_{\text{QSL}}$, the two protocols $s_{\Delta_{\pm}(T)}$ are optimal ($I_T = 0$); the associated quadratic operator has $L - 2$ vanishing and two positive eigenvalues. We deduce the existence of $L - 2$ orthogonal deformations of the optimal protocols $s_{\Delta_{\pm}(T)}$ that preserve the vanishing infidelity.

protocols $s_{\Delta_{\pm}}$. For $T \geq T_{\text{QSL}}$, the infidelity associated with $s_{\Delta_{\pm}}$ vanishes (i.e., they are global optimal protocols) as it is visible in the inset in Fig. 6(a). Consistent with the global optimality condition, we also find $b_t[s_{\Delta_{\pm}}](T) = 0$, $\forall t \in [0, T]$. Therefore, for $T > T_{\text{QSL}}$, the infidelity landscape possesses at least two (and possibly more) globally optimal protocols $s_{\Delta_{\pm}}$. This indicates a nontrivial structure emerging in the near-optimal region in the control landscape.

A better understanding of the T_{QSL} control phase transition is provided by the quadratic stability analysis and, in particular, by the eigenvalues of the Hessian operator. The results for the optimal protocol discussed so far $\{s_0, s_{\Delta_0}, s_{\Delta_{\pm}}\}$ are reported in Fig. 6(b) (see also Appendix C) and can be summarized as follows: (i) s_0 and s_{Δ_0} are quadratically stable in $[0, T_{\text{QSL}}]$. (ii) For s_{Δ_0} , all except two eigenvalues of $J_{t_1 t_2}[s_{\Delta_0}](T)$ change sign from positive to negative exactly at the quantum speed limit T_{QSL} . (iii) For $s_{\Delta_{\pm}}$, all except two eigenvalues of $J_{t_1 t_2}[s_{\Delta_{\pm}}](T)$ vanish as $T \rightarrow T_{\text{QSL}}^-$, and they remain zero for $T > T_{\text{QSL}}$; the remaining two nonvanishing eigenvalues are positive. Thus, the control phase transition at T_{QSL} is associated with linear and quadratic instability.

It is interesting to notice that the presence of infinitely many zero eigenvalues is associated with infinitely many deformations of the protocols $s_{\Delta_{\pm}}$ preserving zero infidelity. This fact has important implications. First, we deduce that at T_{QSL} the number of optimal protocols may grow from one to *infinite*. Second, quadratic stability analysis relative to the optimal protocols $s_{\Delta_{\pm}}$ suggests the existence of an infinite-dimensional “optimal level set,” defined implicitly by the equation $I(T)[s] = 0$, containing optimal and homotopically equivalent protocols; in this case, different optimal protocols (in the same connected component) could be continuously deformed into one another while preserving optimality. Such a property can be valuable in experimental applications: One can select a subset of optimal protocols based on desired criteria, such as robustness against experimental noise [28].

The local character of the linear and quadratic stability analysis does not allow for a direct study of the global properties of the optimal level set. For example, the constraint $|s(t)| \leq 1$ limits the allowed deformations around the two protocols $s_{\Delta_{\pm}}$, potentially preventing the exploration of the optimal level set. Hence, motivated by the stability analysis and keeping in mind the limitations of our local analytical analysis, in Sec. V, we investigate the existence of a level set of continuously connected optimal protocols using Metropolis-adjusted Langevin dynamics.

C. Two-qubit problem

The two-qubit control problem follows the same qualitative behavior of the single-qubit case up to $T = T_{\text{sb}}$. Following the same steps as in Sec. IV A, we correctly identify the optimal protocol s_0 for $T \rightarrow 0$ and detect, by

linear stability analysis, the critical point $T_c \simeq 1.57$. For $T \geq T_c$, only one protocol satisfies linear stability condition among the family (24) parametrized by $\Delta \in [0, T/2]$ (see Fig. 7).

The T_{sb} control phase transition presents new features. (i) As $T \rightarrow T_{\text{sb}}$, $\Delta(T) \rightarrow T/2$, so that $\Delta(T)$ reaches the boundary of its domain $[0, T/2]$ and $s_{\Delta(T)} \rightarrow 0$. (ii) For $T = T_{\text{sb}}$, the linear-order term $b_t[s_{\Delta(T)}](T)$ vanishes: $s = 0$ is a saddle point at $T = T_{\text{sb}}$. (iii) For $T = T_{\text{sb}}$, the lowest eigenvalue of the Hessian operator $J_{t_1 t_2}[s_{\Delta(T)}](T)$ vanishes (cf. Fig. 7). These properties are reminiscent of the behavior of control phase transitions in the one-qubit problem T_{QSL} , where the vanishing eigenvalues in the Hessian operator mark the presence of a branching phenomenon in the low-infidelity region of the landscape. However, for the two-qubit problem at $T = T_{\text{sb}}$, there is only one vanishing eigenvalue.

In Ref. [40], the T_{sb} control phase transition was related to a bifurcation process. It was shown that the unique optimal protocol (valid in $T \leq T_{\text{sb}}$) bifurcates into two optimal protocols (valid in $T \geq T_{\text{sb}}$) individually violating the symmetry $s(t) \leftrightarrow -s(T-t)$ of the quantum control problem. Interestingly, in the quadratic stability analysis, the eigenfunction associated with the vanishing eigenvalue has even parity; namely, it violates the symmetry $f(t) = -f(T-t)$ (cf. Fig. 7). This characteristic instability signals the presence of a bifurcation of a single optimal protocol into a pair of symmetry-violating optimal protocols.

For $T > T_{\text{sb}}$, piecewise continuous optimal protocols can be found using numerical optimization algorithms (for example, the LMC algorithm introduced in Sec. V).




FIG. 7. Stability analysis in the two-qubit control problem. (a) Eigenvalues of the quadratic operator $J_{t_1 t_2}$ from the infidelity Taylor expansion centered around the protocol $s_{\Delta(T)}$, optimal for $T \leq T_{\text{sb}}$. In the plot, we discretize the $[0, T]$ interval in $L = 64$ steps; as L grows, the smallest eigenvalue decreases but remains positive for $T < T_{\text{sb}} \simeq 1.57$. The CLPT in T_{sb} is associated with a single vanishing eigenvalue in the second-order term. (b) Eigenfunction $f(t)$ associated with the vanishing eigenvalue, plotted for $T \simeq T_{\text{sb}}$. As L grows, $f(t)$ approximates a sum of two delta distributions peaked at the interval boundary, $t = 0, T$. Notice the even parity $f(t) = f(T-t)$: this property is associated with the breaking of the symmetry of the quantum control problem, $s(t) \leftrightarrow -s(T-t)$, occurring at $T = T_{\text{sb}}$.

In Appendix C, we report the quadratic stability analysis at $T = T_{\text{QSL}}$ for the two-qubit problem, where the infidelity is expanded around optimal protocols found by the LMC algorithm. The results are analogous to the single-qubit case, except that there are now $L - 4$ vanishing eigenvalues, for $T \rightarrow T_{\text{QSL}}$.

V. STOCHASTIC DYNAMICS WITHIN THE OPTIMAL LEVEL SET BEYOND THE QUANTUM SPEED LIMIT

In Sec. IV, the control phase transition occurring at T_{QSL} has been associated with the formation of infinitely many global minima in the quantum control landscape, branching out from isolated optimal protocols. Past T_{QSL} , stability analysis suggests that the quantum control landscape possesses a level set of distinct protocols satisfying $I[s] = 0$ (i.e., “optimal”) and that can be transformed into each other via continuous deformations. In this section, we introduce a simple stochastic dynamics algorithm to explore the low-infidelity region of the landscape and probe the properties of the optimal level set $\{s(t) : I(T)[s] = 0\}$. Despite its simplicity, this approach can uncover important properties of the level set, such as connectivity, the local structure, and bulk size.

Let us introduce the algorithm we use to explore the optimal level set. First, we discretize the time domain $[0, T]$ in L steps $(t_1, \dots, t_L)$ and approximate the piecewise continuous protocol $s(t)$ with a piecewise-constant protocol $\mathbf{s} = [s(t_1), \dots, s(t_L)]$, where $s(t_i) \in [-1, 1]$. Then, we introduce a stochastic map $\mathbf{s}_n \mapsto \mathbf{s}_{n+1}$ by combining the gradient-free Langevin update rule

$$s_{n+1}(t_i) = s_n(t_i) + \xi(t_i),$$

$$\langle \xi(t_i) \rangle = 0, \quad \langle \xi(t_i) \xi(t_j) \rangle = \sigma^2 \delta_{ij}, \quad (27)$$

with the standard Metropolis acceptance rule

$$P_{\text{acc}}[s(t_i)_{n+1}] = \min[1, \exp(-\beta L \Delta I_{n+1})],$$

$$\Delta I_{n+1} = I_T[s_{n+1}] - I_T[s_n]. \quad (28)$$

Here, the subscript n labels the iteration of the stochastic dynamics, $\xi(t_i)$ is a Gaussian random variable of zero mean and variance σ^2 , and β has the intuitive meaning of (dimensionless) inverse temperature. This dynamics is a gradient-free version of the standard Langevin-Monte Carlo algorithm (also known as Metropolis-adjusted Langevin algorithm [66]); we thus refer to this algorithm as “LMC.” The reasons behind such a simple algorithm are twofold. On the one hand, we neglect the gradient term contribution (typically present in Langevin dynamics), because we use the algorithm to explore the optimal level set, characterized by a vanishing gradient term. On the other, we do not exploit information from the Hessian analysis in order to have an unbiased exploration algorithm.

LMC possesses two free parameters: $\beta, \sigma > 0$. In principle, they are independent of one another and need to be tuned carefully for the dynamics to explore efficiently the space of low-infidelity protocols. On one side, the dynamics is confined to a volume around the $I(T)[s] = 0$ hypersurface that shrinks as $\beta \rightarrow \infty$ and $\sigma \rightarrow 0$. On the other, as $\sigma \rightarrow 0$, the LMC dynamics slows down, and more time is needed for the optimal level set exploration. In what follows, we explore the optimal level set at $\beta = 10^4, 10^5, 10^6$ and choose the value of σ maximizing the acceptance ratio for each β . In these cases, the LMC dynamics possesses three distinct stages.

- (1) *Thermal relaxation stage*— $\mathbf{s}_n$ descends the infidelity landscape, from the initial random configuration to the low-infidelity region, reaching the bottom of one (of possibly many) landscape “valley” [Fig. 8(a)].
- (2) *Diffusion stage*— $\mathbf{s}_n$ moves stochastically along the bottom of the valley, since the Metropolis acceptance rule constrains the size of infidelity fluctuations $\Delta I \sim L\beta$ [Fig. 8(b)].
- (3) *Equilibrium sampling stage*— $\mathbf{s}_n$ explores the landscape valley and samples a statistically representative set of optimal protocols [Fig. 8(c)]. This happens on a timescale that is much larger than that of the diffusion stage and, thus, ensures that individual samples are uncorrelated.

Although the main motivation for the LMC algorithm is the study of the optimal level set, we find that the control phase transition at T_{QSL} affects each of the three stages. As a consequence, next we separately present the main results from the three stages of LMC dynamics.

A. Thermal relaxation stage

Upon initialization of the LMC algorithm, each component of the initial protocol $\mathbf{s}_0$ is drawn randomly from a uniform distribution in the interval $[-1, 1]$. Subsequently, the stochastic dynamics evolves the protocol $\mathbf{s}_n$ decreasing the corresponding infidelity $I[\mathbf{s}_n]$, thus “descending” the landscape. By studying how the infidelity changes as the protocol evolves from the initial random configuration, we obtain information about the structure of the quantum control landscape above the low-infidelity region.

For $\beta \gtrsim 10^3$, the behavior of the algorithm during thermal relaxation is largely independent of $\sigma \lesssim 10^{3/2}$. In general, during this stage, the nonequilibrium dynamics can be divided into two substages: One characterizes the motion at the beginning of thermal relaxation, while the other appears subsequently and persists until the system thermalizes. Interestingly, the whole behavior can be captured by an effective one-dimensional model that admits a closed-form solution and provides a better intuition on the structural changes occurring in the landscape across the different control phase transitions (cf. Appendix D 1). In particular, the CLPT affects the nonequilibrium dynamics in the thermal relaxation stage of LMC, and one may then




FIG. 8. Illustration of LMC dynamics in a two-dimensional toy-model system $\mathbf{s}_n = (x_n, y_n) \in \mathbb{R}^2$ subject to the landscape $V(\mathbf{s}_n) = -(1/8)x_n^2 + (1/4)x_n^4$ (“Mexican hat”) with $\beta = 10^4$ and $\sigma = 10^{-3/2}$ [cf. Eq. (27)]. The trajectory of $\mathbf{s}_n$ in a single LMC run is plotted up to three different iterations. The stochastic motion along the optimal level set (the landscape valley denoted by the red dashed curve) becomes visible. In general, LMC dynamics can be separated into three consecutive stages. (a) Thermal relaxation stage. The system $\mathbf{s}_n$ starts from the random configuration $\mathbf{s}_0$ and reaches the region surrounding the optimal level set $I_T = 0$. (b) Diffusion stage. By stochastic moves, $\mathbf{s}_n$ diffuses in the optimal level set; this motion cannot be suppressed by increasing β (provided the system has enough time to diffuse). (c) Equilibrium sampling stage. Asymptotically, $\mathbf{s}_n$ explores the optimal level set. Using LMC dynamics, we sample different configurations from the optimal level set and study their statistical properties.

detect the presence of the CLPT without sampling optimal protocols. We remark that, so far, CLPTs have been characterized exclusively in terms of statistical properties of the locally optimal protocols, whereas here we provide a dynamical characterization.

B. Diffusion stage

Once the protocol $\mathbf{s}_n$ reaches the bottom of a valley in the infidelity landscape, it moves stochastically along the flat directions of the landscape. If the infidelity minimum is unique (or the infidelity minima are isolated from each other), fluctuations of $\mathbf{s}_n$ are suppressed by the effective inverse temperature β ; cf. Eq. (28). Otherwise, if different infidelity minima are continuously connected with each other, protocol fluctuations within the flat directions are not suppressed by taking $\beta \rightarrow \infty$. Therefore, if the T_{QSL} control phase transition is associated with the creation of infinitely many flat directions in the infidelity valley, then the LMC dynamics after thermal relaxation should be strongly affected by this control phase transition.

For a quantitative description, it is useful to fix the parameter β (in the Metropolis acceptance rule) and vary the size of the stochastic jumps $\sigma \ll \sqrt{\beta}$. In this limit, we approximate the LMC dynamics with an L -dimensional diffusive motion subject to a constraining potential given by the infidelity. For short times, LMC explores the neighborhood of the initial point, and its motion can be decomposed in an approximately free motion along the valley (the tangent space) and fluctuations along the orthogonal directions [cf. Fig. 8(b)]. At long times, the LMC dynamics diffuses around the optimal level set, ideally exploring this set in a relatively dense way (i.e., such that the mean values of the observables we sample using LMC dynamics converge to constant values).

In more concrete terms, during the diffusive stage, we consider the covariance matrix related to the protocol fluctuations $\mathbf{s}_n$. Using this quantity, for each β we estimate the minimum time required to diffuse through the optimal level set (cf. Appendix D 2). In addition, eigenvalues and eigenvectors of the covariance matrix can be used to study the local structure (during early-diffusion time) or the bulk size (for asymptotic diffusion time) of the optimal level set.

C. Equilibrium sampling stage

Ultimately, the T_{QSL} control phase transition affects the statistical properties of protocols lying in the low-infidelity region of the landscape. Accordingly, we primarily use LMC to sample protocols continuously connected with each other through sufficiently small stochastic jumps; cf. Eq. (27). In particular, we use LMC’s diffusive motion to move along the infidelity valleys and store the protocol configuration $\mathbf{s}_n$ every Δn_{dec} iterations (see Appendix D 2) such that, between two samples, LMC explores a global portion of the optimal level set. For each set of LMC parameters, we collect R LMC runs acquiring M optimal protocols each. Notice that, in each run, LMC starts from a different random protocol $\mathbf{s}_0$ and the system first relaxes to the global optimal level set before the sampling of the M optimal protocols begins.

As a first observable, consider the order parameter $q(T)$ defined in Eq. (4) which measures the size of fluctuations around the mean protocol. The average $\langle \cdot \rangle$ is here performed over each LMC run separately, in order to separate contributions from different connected components of the optimal level set. The existence of an optimal level set composed of isolated points is related to a vanishing $q(T)$ in the limit $\beta \rightarrow \infty$. In Fig. 9, we show $q(T)$ for different values of the inverse temperature β and the system size L ,




FIG. 9. Single- and two-qubit control problems. Jump discontinuity of the order parameter $q(T)$ in Eq. (4) at the quantum speed limit; $q(T)$ quantifies fluctuations of the piecewise continuous protocol as it explores the optimal level set under LMC dynamics. A nonvanishing value of $q(T)$ (as $\beta \rightarrow \infty$) implies that the optimal level set possesses a family of optimal protocols connected to each other via continuous transformations (homotopy). The black dashed curve shows the analytical prediction discussed in Sec. VIA. In the two-qubit case, the errors associated with the mean value of $q(T)$ are relatively small for $T \lesssim 3.3$, but they show a sudden increase for $T \gtrsim 3.3$ for the same number of LMC steps. As discussed in Ref. [41], this is related to the presence of a new disconnected component in the optimal level set arising at $T \simeq 3.3$.

for the single- and two-qubit problems. In the single-qubit case, as β grows, we find $q(T) \rightarrow 0$ for $T \in [0, T_{\text{QSL}}]$: This confirms that for $T < T_{\text{QSL}}$ the infidelity minimum is unique for piecewise-constant protocols. By contrast, for $T > T_{\text{QSL}}$ the minimum is not unique, since $q(T)$ is independent of β . Remarkably, the scaling analysis with L reveals the T_{QSL} transition as a first-order transition: Namely, $q(T)$ has a jump discontinuity in $T = T_{\text{QSL}}$. In the two-qubit case, similarly, we find $q(T) \rightarrow 0$ (as β is increased) for $T \in [0, T_{\text{QSL}}]$. Notice that this does not contradict the presence of two isolated optimal protocols in the interval $[T_{\text{sb}}, T_{\text{QSL}}]$: Each LMC run is confined to one of the two valleys which are not continuously connected to each other. Therefore, the protocol fluctuations still depend on the inverse temperature β . For $T = T_{\text{QSL}}$, $q(T)$ acquires a finite value with a jump discontinuity (as in the single-qubit case). Interestingly, for $T \gtrsim 3.3$ in the two-qubit case the presence of multiple connected components in the optimal level set gives rise to a larger standard deviation of the mean value $q(T)$ for the same number of LMC samples [error bar in Figs. 9(c) and 9(d)] [41].




FIG. 10. Schematic visualization of different sets of protocols $R_1, R_2, \dots$ sampled by independent LMC runs. As $\beta \rightarrow \infty$, stochastic jumps between different connected components of the optimal level set, represented by the red sets C_1 and C_2 , become ever more unlikely during LMC dynamics. In this limit, the sets $R_1, R_2, \dots$ are confined in either one of the two components. The distance between two sets, defined in Eq. (D6) and represented by the arrows in the figure, can vanish, as the number of sampled sets increases, only when the two sets are located in the same connected component. Whenever two sets are located in different connected components, the distance reaches asymptotically (in LMC iterations) a nonvanishing value.

The number of connected components in the optimal level set can be estimated by computing the distance between pairs of LMC runs. The M protocols sampled within each LMC run are connected with each other by the LMC dynamics (see, e.g., Fig. 8 for a schematic cartoon). Assuming that LMC approximates homotopic transformations between optimal protocols, we regard the M protocols collected during each LMC run as continuously connected to one another. On the contrary, protocols in different runs may belong to different connected components of the optimal level set, since LMC thermalizes to different protocols in different runs. Therefore, by computing the distance between pairs of protocols belonging to different LMC runs, we can estimate the number of connected components in the optimal level set. In Fig. 10, we show a schematic representation of the idea behind the distance analysis, involving protocols within and across runs, and we discuss details in Appendix D3. In the single-qubit control landscape, our results suggest the presence of a single connected component for all $T \in [T_{\text{QSL}}, 3.5]$: As the number of protocols per run is increased, the distance steadily decreases to zero following an algebraic decay. The two-qubit problem, where the situation is more complex and multiple connected components are present, is discussed in detail in Ref. [41].

VI. CRITICAL BEHAVIOR AT THE QUANTUM SPEED LIMIT

In Sec. V, we presented numerical results showing the relation between the control landscape phase transition at $T = T_{\text{QSL}}$ and the formation of an infinite-dimensional

optimal level set lying in the zero-infidelity hyperplane. So far, this analysis was purely numerical, and little intuition about the general structure of this infinite-dimensional manifold was given.

In this section, we complement the numerical study, limited by construction to finite-size ($L < \infty$) systems, with an analytical characterization of the optimal level set and the associated T_{QSL} transition. In particular, we shed new light on the nature of the nonanalytic points in the order parameters $q_{\text{BB}}(T)$ and $q(T)$ visible in Figs. 3 and 9.

In what follows, we focus on the properties of the infidelity landscape that are independent of the controlled quantum system at the quantum speed limit transition. Similar to the Ginzburg-Landau theory of phase transitions, we set up an analytical theory and relate the scaling in $\Delta T = T - T_{\text{QSL}}$ of coefficients in the infidelity expansion (carrying information about the local geometry of the landscape) to the critical behavior of the order parameters $q(T)$ and $q_{\text{BB}}(T)$, in the limit $\Delta T \rightarrow 0^+$.

A. Critical behavior of the order parameter $q(T)$

The order parameters $q(T)$ and $q_{\text{BB}}(T)$, introduced in Eqs. (4) and (5), capture statistical properties of local minima in the quantum control landscape. For their evaluation, a method to compute the expectation values $\langle s(t) \rangle$ and $\langle s(t)^2 \rangle$ has to be introduced. So far, we used numerical algorithms (SD and LMC) to sample control protocols from the low-infidelity region, and we computed the expectation values by averaging over the samples. Now, as a first step toward an analytical approach, we specify a rule for the computation of the expectation values.

1. Introducing the partition function

Given the set of admissible control protocols $\mathcal{X}$, expectation values can be analytically evaluated once a set of weights $w(x)$, $x \in \mathcal{X}$ is introduced. Then, the expectation value of a generic observable $f: \mathcal{X} \rightarrow \mathbb{R}$ is computed as the ratio

$$\langle f \rangle_{\mathcal{X}} = \frac{\sum_{x \in \mathcal{X}} f(x) w(x)}{\sum_{x \in \mathcal{X}} w(x)}.$$

Defining the partition function $Z[f] = \sum_{x \in \mathcal{X}} f(x) w(x)$, we can compactly rewrite $\langle f \rangle_{\mathcal{X}} = Z[f]/Z[1]$. In our problem, $\mathcal{X}$ is the infinite-dimensional space of bounded protocols $\{s(t): |s(t)| \leq 1\}$, and the sum is a functional integral (path integral) formally written as

$$Z[f] = \int_{-1}^{+1} \prod_t ds_t f[s] w[s].$$

Since we are interested in the order parameters $q(T)$ and $q_{\text{BB}}(T)$ and because they are related to a structural change of the low-infidelity region, we can, for example, choose

the Boltzmann-Gibbs weights $w_\beta[s] = \exp(-\beta I[s])$, with $I[s] = I(T)[s]$ the infidelity and β a free parameter [67]. In this way, as $\beta \rightarrow \infty$, the dominant contributions to the partition function have the lowest possible infidelity, and the weight $w_\beta[s]$ concentrates on globally optimal protocols:

$$Z[f] = \int_{-1}^{+1} \prod_t ds_t f[s] \delta(I[s] - I_{\min}),$$

where $I_{\min} = I_{\min}(T)$. For $T \geq T_{\text{QSL}}$, by definition $I_{\min} = 0$ and the partition function reduces to

$$Z[f] = \int_{-1}^{+1} \prod_t ds_t f[s] \delta(I[s]), \quad T > T_{\text{QSL}}. \quad (29)$$

In the next sections, we explicitly evaluate the partition function and relate, for $T - T_{\text{QSL}} = \Delta T \rightarrow 0^+$, the non-analytic behavior of the order parameters $q(T)$ and $q_{\text{BB}}(T)$ to the structural changes occurring in the low-infidelity region of the landscape.

2. Optimal level set parametrization

Starting just before the quantum speed limit T_{QSL} , by definition, a subset of control protocols reaches the minimum value of the infidelity, as T is increased. In what follows, we assume a *finite* number of locally optimal protocols, which we denote $\{s_0^{(i)}\}_{i=1}^{N_0}$. This is the case, for example, in the single- and two-qubit problems introduced in Sec. II (cf. Secs. IV and V), and we expect the same situation in generic controlled quantum systems.

Let us focus on the structure of the control landscape around a single locally optimal protocol s_0 from the set $\{s_0^{(i)}\}$. Our goal is to use the infidelity expansion centered at the protocol s_0 to obtain a parametrization of the optimal level set around s_0 , for $T \geq T_{\text{QSL}}$.

The infidelity expanded around a given protocol s_0 reads as

$$\begin{aligned} I(T)[s_0 + \delta s] = & c + \int dt b_t \delta s_t + \frac{1}{2} \int d^2 t J_{t_1 t_2} \delta s_{t_1} \delta s_{t_2} \\ & + \frac{1}{3!} \int d^3 t d_{t_1 t_2 t_3} \delta s_{t_1} \delta s_{t_2} \delta s_{t_3} \\ & + \frac{1}{4!} \int d^4 t g_{t_1 t_2 t_3 t_4} \delta s_{t_1} \delta s_{t_2} \delta s_{t_3} \delta s_{t_4} \\ & + \dots, \end{aligned} \quad (30)$$

where δs represents the deviation from the protocol s_0 . The T and s_0 dependence in the expansion coefficients $c, b_t, J_{t_1 t_2}, \dots$ is omitted for notational simplicity. We expand δs_t in the eigenbasis $\{f_t^{(n)}\}_n$ of the Hessian operator $J_{t_1 t_2}$ as $\delta s_t = \sum_{n=1}^{\infty} \delta s_n f_t^{(n)}$, with δs_n the new degrees of

freedom. Using the orthonormality of the eigenbasis, $\int_0^T f_t^{(n)} f_t^{(m)} dt/T = \delta_{nm}$, we obtain

$$\begin{aligned} I(T)[s_0 + \delta s] = & c + T \sum_{n=1}^{\infty} b_n \delta s_n + \frac{T^2}{2} \sum_{n=1}^{\infty} \lambda_n \delta s_n^2 \\ & + \frac{T^3}{3!} \sum_{n,m,k} d_{nmk} \delta s_n \delta s_m \delta s_k \\ & + \frac{T^4}{4!} \sum_{n,m,k,l} g_{nmkl} \delta s_n \delta s_m \delta s_k \delta s_l + \dots, \end{aligned} \quad (31)$$

where $\{\lambda_n\}$ are the eigenvalues of $J_{t_1 t_2}$. Notice that we define $b_n = \int_0^T dt f_t^{(n)} b_t/T$ and analogously for higher-order terms: $J_{nm} = \delta_{nm} \lambda_n, d_{nmk}, g_{nmkl}, \dots$

Next, we restrict the analysis to a neighborhood of the quantum speed limit. We study the behavior of the coefficients of the expansion in Eq. (31), in the limit $\Delta T = T - T_{\text{QSL}} \rightarrow 0^-$. In this regime, there are important constraints on these coefficients, originating from the infidelity lower bound being attained at zero and the assumption of local optimality of the protocol s_0 .

Since $I_T[s] \in [0, 1]$, odd-order terms $b_n, d_{nmk}, \dots$ vanish at least as fast as ΔT ; otherwise, the infidelity bound $I_T[s] \geq 0$ would be severely violated at $\Delta T = 0$. The zeroth-order term $c \in [0, 1]$ has a zero at $\Delta T = 0$ by assumption, so that necessarily $c \sim \Delta T^2$. The eigenvalues of the second-order term λ_n separate into two subsets: A number n_+ of them remain positive, while the rest vanish: $\lambda_{n \leq n_+} \sim 1$ and $\lambda_{n > n_+} \sim \Delta T$. In general, n_+ depends on the controlled quantum system [26]; for example, in the single- and two-qubit cases introduced in Sec. II, we found $n_+ = 2$ and $n_+ = 4$, respectively (cf. Appendix C). Finally, consistent with the optimality for $T \leq T_{\text{QSL}}$ of the protocol s_0 , we necessarily have $b_{n > n_+} \sim \Delta T^2$; otherwise, the protocol s_0 would not be a minimum of the infidelity for $\Delta T = 0$.

In summary, the most general scaling behaviors in ΔT of the infidelity expansion coefficients are

$$\begin{aligned} c & \sim \Delta T^2, \\ b_{n \leq n_+} & \sim \Delta T, & b_{n > n_+} & \sim \Delta T^2, \\ d_{nmk} & \sim \Delta T, & g_{nmkl} & \sim 1 \end{aligned} \quad (32)$$

and

$$\lambda_{n \leq n_+} \sim 1, \quad \lambda_{n > n_+} \sim \Delta T, \quad (33)$$

with n_+ a positive integer.

We obtain an approximate parametrization of the optimal level set $\{s: I(T)[s] = 0\}$ around s_0 by truncating the infidelity expansion in Eq. (30) to a given order and consider the analytical continuation of the scaling in Eqs. (32) and (33), for $\Delta T \rightarrow 0^+$.

3. Hard-boundary constraint approximation

The hard-boundary constraint $|s(t)| \leq 1$ imposed on the protocol space amounts to a restriction of the integration domain for each coordinate s_t . Unfortunately, this constraint does not allow for an explicit evaluation of the partition function, so we resort to a suitable approximation.

As a first step, we rewrite the constraint as

$$\begin{aligned} \int_{-1}^1 ds_t &= \int_{-1}^1 d\lambda_t \int_{\mathbb{R}} ds_t \delta(s_t - \lambda_t) \\ &= \int_{-1}^1 d\lambda_t \int_{\mathbb{R}} ds_t \int_{\mathbb{R}} \frac{d\nu_t}{2\pi} e^{i\nu_t(s_t - \lambda_t)}, \end{aligned} \quad (34)$$

where we introduce the Fourier representation of the Dirac-delta distribution $\delta(s_t - \lambda_t)$. Integrating over λ_t gives

$$\int_{-1}^1 d\lambda_t e^{-i\nu_t \lambda_t} = 2 \frac{\sin(\nu_t)}{\nu_t} \approx 2e^{-\nu_t^2/6}.$$

The above approximation trades the $\sin(x)/x$ function for a Gaussian displaying the same behavior for $\nu_t \rightarrow 0$. Notice that both functions are peaked around $\nu_t = 0$ and vanish for $|\nu_t| \rightarrow \infty$ [68]. Using this approximation, we finally integrate over ν_t in Eq. (34). Eventually, we obtain

$$\int_{-1}^1 ds_t \approx \int_{\mathbb{R}} ds_t e^{-(3/2)s_t^2}, \quad (35)$$

which approximates the constraint $|s(t)| \leq 1$ (we neglect the overall prefactor, as it will not affect expectation values).

Equation (35) is the Gaussian approximation of the hard-boundary constraint, $|s(t)| \leq 1$. We observe the following important fact. If we relax the infidelity constraint, this approximation of the hard-boundary constraint does not affect the first two moments $\langle s_t \rangle$ and $\langle s_t^2 \rangle$. Namely,

$$\begin{aligned} \frac{\int_{-1}^1 ds_t s_t}{\int_{-1}^1 ds_t} &= \frac{\int_{\mathbb{R}} ds_t e^{-(3/2)s_t^2} s_t}{\int_{\mathbb{R}} ds_t e^{-(3/2)s_t^2}} = 0, \\ \frac{\int_{-1}^1 ds_t s_t^2}{\int_{-1}^1 ds_t} &= \frac{\int_{\mathbb{R}} ds_t e^{-(3/2)s_t^2} s_t^2}{\int_{\mathbb{R}} ds_t e^{-(3/2)s_t^2}} = \frac{1}{3}. \end{aligned} \quad (36)$$

However, higher moments will differ as the two distributions are distinct.

4. Evaluating $q(T)$

Over the space of piecewise continuous protocols, the order parameter $q(T)$, defined in Eq. (4), reads as

$$q(T) = \frac{1}{T} \int_0^T dt \left(\langle s_t^2 \rangle - \langle s_t \rangle^2 \right).$$

As we ultimately want to compare with the numerical results shown in Sec. V, we perform the averages appearing in $q(T)$ over each local minimum in $\{s_0^{(i)}\}_1^{N_0}$, separately.

That is,

$$q(T) = \frac{1}{N_0} \sum_{i=1}^{N_0} q_i(T), \quad (37)$$

$$q_i(T) = \frac{1}{T} \int_0^T dt \left(\langle s_t^2 \rangle_i - \langle s_t \rangle_i^2 \right), \quad (38)$$

where $\langle \cdot \rangle_i$ is the average relative to the minimum $s_0^{(i)}$. The remainder of this section is dedicated to the calculation of $q_i(T)$. Therefore, except for $q_i(T)$, we omit the dependence on the label i , for notational convenience.

As we saw in Sec. VI A 2, the parametrization of the optimal level set $\{s: I[s] = 0\}$ is conveniently written in terms of the local Hessian eigenbasis $\{f^{(n)}\}$. In this basis, the order parameter becomes

$$q_i(T) = \sum_{n=1}^{\infty} \left(\langle s_n^2 \rangle - \langle s_n \rangle^2 \right). \quad (39)$$

Thus, to estimate $q_i(T)$ we need to compute $\langle s_n \rangle = Z[s_n]/Z[1]$ and $\langle s_n^2 \rangle = Z[s_n^2]/Z[1]$.

Considering the approximations introduced in the last two sections, the exact partition function in Eq. (29) is now replaced by

$$Z[f] \approx \int_{\mathbb{R}} \prod_t^L ds_t e^{-(3/2)(L/T) \int_0^T dt s_t^2} \delta(I(T)[s]) f[s]. \quad (40)$$

We notice that the effective system size L represents the dimensional cutoff in the path integral, namely, $\prod_t^L ds_t \equiv ds_{t_1} ds_{t_2} \dots ds_{t_L}$. As we are interested in the case of piecewise continuous protocols, we ultimately consider the $L \rightarrow \infty$ limit.

We give the evaluation of the partition function in Appendix F to keep the discussion clear and present here only the main results. Eventually, from Eq. (40), we obtain the simple expression

$$q_i(T) \sim \frac{1}{3L} \sum_{n=1}^L (1 + y_* \lambda_n)^{-1}, \quad L \rightarrow \infty. \quad (41)$$

The quantity y_* depends on ΔT and is restricted to the interval $(-\lambda_+^{-1}, -\lambda_-^{-1})$, with $\lambda_{+,-}$ the positive and negative eigenvalues of largest absolute value. For $\Delta T \leq 0$, we find that the eigenvalues $\{\lambda_n\}$ are positive and $y_* = \infty$ so that $q_i(T) = 0$.

From this result, we deduce that the jump discontinuity of the order parameter $q_i(T)$ essentially depends on the existence of the negative eigenvalue(s); in this situation, the quantity y_* acquires a finite value and the set of optimal protocols $\{s: I[s] = 0\}$ becomes degenerate. To see this, we observe the two possible behaviors:

$$\begin{aligned} (1 + y_* \lambda_n)^{-1} &\rightarrow 0, & y_* \lambda_n &\rightarrow \infty, \\ (1 + y_* \lambda_n)^{-1} &\rightarrow 1, & y_* \lambda_n &\rightarrow 0. \end{aligned}$$

Since $|y_*| < |\lambda_-^{-1}|$ and we have $|\lambda_n| \rightarrow 0$ as $n, L \rightarrow \infty$ from general properties of the spectrum of the Hessian operator J_{nm} (cf. Sec. IV), we deduce that $q_i(T)$ acquires a finite value once $|\lambda_-^{-1}| < \infty$ exists (otherwise, $y_* = \infty$). In this case, we obtain $q(T) \rightarrow 1/3$ as $L \rightarrow \infty$ (see Appendix F for details).

The above result has the following physical interpretation. The requirement of vanishing infidelity imposes a number n_+ of constraints on the L degrees of freedom of the control protocol, when the quantum duration T is larger than the quantum speed limit T_{QSL} . Therefore, in the $L \rightarrow \infty$ limit, the extra $L - n_+$ degrees of freedom introduced are redundant—the control protocol is underconstrained. In the control landscape, these extra degrees of freedom are associated with *soft modes*, deformation of optimal control protocols that preserve the infidelity within second order. The order parameter $q(T)$ measures the average size of allowed deformations between different optimal protocols, so that its value grows with L , for $T > T_{\text{QSL}}$. In particular, the limit value $1/3$ arises from the freedom left by the infidelity expansion truncated at the second order on these extra degrees of freedom, as $L \rightarrow \infty$ [cf. Eqs. (36) and (39)].

Considering higher-order terms in the infidelity expansion introduces additional constraints on the L degrees of freedom. In a path-integral language, higher-order terms introduce interactions between the different modes that ultimately affect the value of the jump discontinuity of $q(T)$. Thus, we understand that higher-order terms in the infidelity expansion are necessary to accurately predict the size of the jump discontinuity of $q(T)$ at the quantum speed limit.

In conclusion, our analysis shows that the universal feature of the order parameter $q(T)$ is the presence of a jump discontinuity, whereas the size of the jump depends on the details of the controlled quantum system.

B. Critical behavior of the order parameter $q_{\text{BB}}(T)$

In this section, we show that the partition function introduced in the last section, defined over the continuous protocol space, can be used for the analytical estimation of the critical exponent of $q_{\text{BB}}(T)$, defined within the space of bang-bang control protocols.

The average $\langle s_i \rangle$, contained in Eq. (5) defining the order parameter $q_{\text{BB}}(T)$, can be formally written as

$$\langle s_i \rangle = \sum_{\{s_i\}} w(\{s_i\}) s_i,$$

where the sum runs over the bang-bang control protocols space and $w(\{s_i\})$ are the statistical weights.

Nonetheless, in the limit of an infinite number of bang-bang steps $N \rightarrow \infty$, a coarse-grain average of bang-bang protocols would give

$$\langle s_i \rangle = \int_{-1}^{+1} \prod_t ds_t w[s] \rho[s] s_{t_i},$$

where the density of states $\rho[s]$ keeps track of how many bang-bang protocols are mapped to the same continuous protocol by the coarse-grain average procedure (cf. Fig. 11). Therefore, we now consider the partition function

$$Z_{\text{BB}} = \int_{-1}^{+1} \prod_t ds_t \delta(I[s]) \rho[s] \quad (42)$$

and estimate $\rho[s]$ in the limit $N \rightarrow \infty$ by subdividing the bang-bang protocol $\{s_i\}$ in L bins, each containing N/L consecutive bang-bang steps. Assuming the spins $s_i = \pm 1$ in each bin are uncorrelated, each s_i becomes an independent random variable, and the number of equivalent configurations within each bin depends on only the number density of up or down spins (cf. Fig. 11).

In the limit $N, L \rightarrow \infty, N/L \rightarrow \infty$ we obtain

$$\rho[s] = \prod_t \rho(s_t) = e^{\sum_t S(s_t)}, \quad (43)$$

$$S(s_t) \sim -(N/L)[s_t^+ \log s_t^+ + s_t^- \log s_t^-],$$

where t labels the bin, $s_t^\pm = (1 \pm s_t)/2$ is the density of up and down spins in the bin t , and $S(s_t)$ is the entropy associated with the variable s_t . Consequently, for a given protocol s , we obtain the density of states $\rho[s] = e^{S[s]}$ associated with the Shannon entropy




FIG. 11. Schematic representation of the coarse-graining average performed on a bang-bang protocol with N steps, into a piecewise-constant protocol with L time steps. The average value in each bin depends on the number of $+1$ and -1 that it contains; in the limit $N \rightarrow \infty$, different bang-bang protocols can approximate the same coarse-grained continuous protocol. This redundancy is responsible for the density of state factor $\rho[s]$ appearing in partition function (42) and estimated in Eq. (43).

$$S[s] = -\frac{N}{T} \int_0^T dt [s_t^+ \log s_t^+ + s_t^- \log s_t^-]. \quad (44)$$

This result is independent of the particular binning subdivision: Changing the number of bang-bang steps s_i contained in each bin s_t yields the same density of states provided that spins in each bin remain decorrelated in the limit $N \rightarrow \infty$.

1. Shannon entropy approximation

We would like to explicitly evaluate the partition function in Eq. (42). Unfortunately, the Shannon entropy makes the integral intractable by means of elementary analytical tools. Since we are interested in the simplest possible analytical estimate of the order parameter $q_{\text{BB}}(T)$, we search for an approximation of the Shannon entropy that allows the partition function to be evaluated by elementary means.

To this end, we replace the nonlinear Shannon entropy with the quadratic function

$$S[s] \mapsto S_2[s] = \frac{N\alpha}{2} \left(1 - \frac{1}{T} \int_0^T dt s_t^2 \right), \quad (45)$$

where α is a positive free parameter. $S_2[s]$ is commonly known as the Tsallis entropy with entropic index 2 [69]. Notice that α can be fixed by imposing a further constraint (for example, $\alpha = 1$ approximates the Shannon entropy density around $s_t = 0$ up to second order). Still, its precise value will not affect the critical scaling of q_{BB} , as we explain below.

Finally, following the Shannon entropy approximation, we extend the integration domain of $\{s_t\}$ from $[-1, 1]$ to $\mathbb{R}$.

2. Evaluating $q_{\text{BB}}(T)$

Evaluating the partition function in Eq. (42) within a second-order approximation of the infidelity expansion, we obtain the critical behavior (see Appendix H)

$$\Delta q_{\text{BB}}(T) = q_{\text{BB}}(T) - q_{\text{BB},*}(T) \sim c_0 + c_1 \Delta T, \quad (46)$$

with $q_{\text{BB},*}(T)$ given by the analytical behavior of $q_{\text{BB}}(T)$ for $T < T_{\text{QSL}}$ and $c_{0,1}$ constant in ΔT . This result deviates from the continuous curve observed in Fig. 12 by the presence of the constant term c_0 . In Appendix H, we show that this jump discontinuity in the analytical prediction of $\Delta q_{\text{BB}}(T)$ is an artifact of the second-order truncation of the infidelity expansion. In addition, we show that including quartic-order terms yields (cf. Appendix H)

$$\Delta q_{\text{BB}}(T) = q_{\text{BB}}(T) - q_{\text{BB},0}(T) \sim c_2 \Delta T^{\zeta_a} + c_3 \Delta T^{2\zeta_a}, \quad (47)$$

with $c_{2,3}$ constant in ΔT and an undetermined critical exponent in the allowed interval $\zeta_a \in [0, 1/2]$. We notice that this interval for ζ_a is obtained within the approximation




FIG. 12. Numerical estimate of the critical behavior of the order parameter $q_{\text{BB}}(T)$, for the single- and two-qubit system. (a) and (c) show the curves $q_{\text{BB}}(T)$ obtained from stochastic descent with a finite number of bang-bang steps, near the quantum speed limit transition in T_{QSL} . (b) and (d) show the collapsed curves $L^{\zeta/\nu} \Delta q_{\text{BB}}(L^{1/\nu} \Delta T)$, obtained by fitting T_{QSL} , ζ_n , ν , and $q_{\text{BB},0}(T)$ ($\Delta T = T - T_{\text{QSL}}$). The critical exponents ζ_n extracted in the two cases are in reasonable agreement with the analytical result in Eq. (47) (see the text). The scaling collapse locates the quantum speed limit T_{QSL} with approximately 0.1%–1% precision.

of the coarse-grain Shannon entropy, so that small deviations are possible. Our analysis shows that the precise value of ζ_a depends on the details of the landscape, around the optimal control protocol(s) at the quantum speed limit. More precisely, ζ_a depends on the fourth-order terms in the infidelity expansion.

For comparison, in Fig. 12, we perform a numerical scaling collapse around the quantum speed limit T_{QSL} of the curves $q_{\text{BB}}(T)$ obtained in the single- and two-qubit example discussed in Sec. II (see Appendix I). The results return the numerical critical scaling $\Delta q_{\text{BB}}(T) \sim \Delta T^{\zeta_n}$, with $\zeta_n \approx 0.7$ and an estimated error of the order of approximately 0.1, in both cases.

Overall, we find the numerical critical exponent ζ_n in reasonable agreement with the analytically predicted range for ζ_a . The possible deviation between the two results may have different origins. Our analytical result relies on the approximation of the entropy associated with the coarse-grain average, where we neglected correlations at different times t and adopted a quadratic approximation for the Shannon entropy. Our numerical results are affected by systematic errors, which limit the precision of the estimated critical exponent (cf. Appendix I). Finally, in practice, the value of coefficients c_2 and c_3 may be important for a precise comparison with the numerical results.

From our theoretical analysis, we deduce that the universal feature of the order parameter $q_{\text{BB}}(T)$ is the presence of a cusp, $\Delta q_{\text{BB}}(T) \sim \Delta T^\zeta$, with a fractional critical exponent ζ , whose precise value depends on the details of the controlled quantum system.

In conclusion, let us reemphasize the simplicity and generality of our analysis of the landscape at the quantum speed limit transition. Similar to the Ginzburg-Landau theory of phase transitions, our method focuses on the generic scaling behavior of the coefficients in the landscape expansion and, therefore, applies to generic controlled quantum systems. Using this approach, we connect the changes in the geometry of the landscape to the nonanalytic behavior of order parameters and explain their universal properties at the quantum speed limit transition.

VII. EXTENSION TO GENERIC MANY-BODY CONTROLLED QUANTUM SYSTEMS

Although the formal extension of our methods to many-body controlled quantum systems is straightforward, the exponentially growing Hilbert space dimension (in the number of controlled degrees of freedom) affects their practical implementation. Therefore, we now briefly discuss the broader applicability of our methods to generic quantum control problems, highlighting some potential challenges and presenting useful ideas on how to tackle them.

As discussed in Sec. III, the convergence of the landscape expansions is controlled by the norm of the Hamiltonian and the total protocol duration. Therefore, for many-body quantum systems, in particular, we anticipate a larger order of truncation in the landscape expansions to reliably capture the transitions in the control landscape. Nevertheless, there exist classes of many-body problems possessing polynomially large algebras [70]. In this case, the matrices associated with the generators of the dynamics are sparse, and higher-order terms of the expansion can be efficiently computed. In addition, when interested in a particular region of the control protocol space, we can reduce the order of truncation by changing the center of the analytical expansion using, e.g., an educated guess for a (close-to-optimal) protocol. This is particularly relevant for the study of CLPTs, as they originate from sudden changes in the subset of locally optimal protocols.

The adiabatic tracing method and the numerical exploration algorithm, used to study the structure of the optimal level set in Secs. IV and V, require that (i) the evolution of the quantum system can be computed and (ii) locally optimal protocols can be obtained through optimization procedure. This limits the applicability of these methods to quantum systems whose evolution can be efficiently performed; in addition, the optimization of control protocols has to be feasible, within numerical precision. Recently, it was demonstrated that both conditions are satisfied in many-body control problems with polynomially large algebras (in the number of controlled degrees of

freedom): This allows one to numerically compute almost-optimal quantum control protocols that prepare macroscopically entangled Greenberger-Horne-Zeilinger states and the topologically ordered ground state of the cluster Ising model without the need to store the quantum state [70]. We also notice that a similar idea to adiabatic tracing was recently applied as an efficient optimization strategy for the quantum approximate optimization algorithm—a parametrized *Ansatz* used to control the evolution in quantum many-body circuits [71,72]. Finally, we remark that our methods are formally independent of the method used to evolve the quantum system: The evolution does not necessarily have to be *exact* or evaluated on a *classical* computer. For instance, the quantum dynamics could also be evaluated experimentally on a quantum simulator, via circuit dynamics in a quantum computer, or using approximate numerical methods on a classical computer (such as time-dependent density matrix renormalization group or time-evolving block decimation for one-dimensional chains). In summary, experimentally relevant many-body control problems exist where the techniques we develop can be applied.

The analytical framework based on path integrals, introduced in Sec. VI, does not require detailed information about the quantum control problem. Its predictions apply to control landscapes, where a finite number of locally optimal protocols become global optima, as the protocol duration T crosses the quantum speed limit. Note that this requires a finite quantum speed limit [73].

In conclusion, extending the methods developed in this work to many-body controlled quantum systems presents an exciting and challenging open direction for future research.

VIII. CONCLUSION

A. Summary of the main results

In this work, we combined different analytical and numerical techniques that shed new light on quantum control phase transitions. At its core, our contribution is fourfold.

First, we derive three different analytical expansions that approximate the infidelity functional (or similar objective functions) of the single- and two-qubit control problems and capture the underlying control phase transitions. In particular, the Dyson and cumulant expansions admit an effective energy interpretation where the protocol configuration plays the role of a spin or field configuration (depending on the constraints imposed on the protocol; cf. bang-bang vs piecewise continuous protocol space). Control phase transitions are captured by the approximate landscape, provided the order of truncation is large enough. Nonetheless, we show that changing the center of the expansion allows one to reduce the order of truncation.

Second, we introduce the adiabatic tracing method and study the near-optimal region of the control landscape as the protocol duration T is varied. Starting from an

infinitesimal T and incrementing its value adiabatically (i.e., while optimizing the control protocol), we relate control phase transitions to precise changes in the structure of the landscape around the optimal protocol(s). This analysis relies on the knowledge of optimal protocol(s) of the landscape; in our case, we find them using either analytical *Ansätze* (motivated by the landscape expansions) or numerical optimization algorithms (simulated annealing with LMC dynamics). More generally, available tools in the quantum optimal control community (such as gradient-based or gradient-free optimizers or machine-learning methods) can be adapted to trace optimal protocols as the protocol duration is varied. In this context, notice that the adiabatic tracing method can be regarded as a simulated annealing procedure, where, instead of the temperature β^{-1} , the annealing parameter is the control protocol duration T . This suggests a new direction for improving currently available optimization procedures (both analytical and numerical) in quantum control problems.

Third, we use a gradient-free version of the Langevin–Monte Carlo algorithm (LMC) to explore the optimal level set, defined in the infinite-dimensional protocol space by the equation $I(T)[s] = 0$. Using LMC, we characterize its topological and geometrical properties (connectedness and bulk sizes), and we obtain a numerical estimate of the order parameter $q(T)$. Notice that LMC explores the optimal level set by considering stochastic fluctuations at finite inverse temperature β , in the space of piecewise-constant functions with L steps. To overcome these numerical limitations, we explicitly check for finite-size effects by studying the results of the simulations while varying L and β .

Fourth, inspired by statistical field theory methods, we rephrase the optimal control problem in path-integral language and use this new formulation to estimate the critical behavior of order parameters $q(T)$ and $q_{\text{BB}}(T)$ for $T \rightarrow T_{\text{QSL}}$. This framework does not require detailed information about the specific quantum control problem. Instead, similar to the Ginzburg-Landau treatment of phase transitions, the method takes advantage of the analyticity of the cost function and considers the most general scaling behavior for the coefficients in the landscape expansion. This method allows us to connect the critical behavior of the statistical order parameters with both the changes occurring in the geometry of the control landscape and the constraints imposed on the protocol space. In addition, it captures the jump discontinuity of $q(T)$ and the fractional critical exponent of $q_{\text{BB}}(T)$, raising the question about universality in the behavior of the order parameters at the quantum speed limit transition, irrespective of the controlled quantum system.

B. Physical significance of control landscape phase transitions

Let us now turn to the physical implications of CLPTs for both the control of quantum systems and the

optimization of control parameters. In what follows, we discuss the three transitions at T_c , T_{QSL} , and T_{sb} separately.

The T_c transition signals the lack of local controllability of the quantum system. At T_c , the shape of the optimal control protocol changes: from a bang-bang to a singular type (see Fig. 5 for an example). The singular arcs occurring in the control protocol reflect the experimental limitations in selecting the control fields; in particular, they appear when the applied fields cannot fully steer the system's evolution [1,26,74]. Depending on the accessible family of control protocols, the T_c transition may have a sizable effect on the optimization landscape and introduce a large degeneracy of locally optimal protocols. For the class of bang-bang control protocols we investigated, the T_c transition sets the separation between a convex landscape and a landscape exhibiting multiple local minima.

The quantum speed limit transition (T_{QSL}) sets the boundary between an unsolvable (overconstrained) and a solvable (underconstrained) optimization problem; its existence determines whether the target state can be reached in a finite time, given the accessible controls. Physically, the existence of a finite quantum speed limit depends on the physical constraints of the applied control fields, such as limited amplitude (as in our case), bandwidth, or total power output. For a generic controlled quantum system, no efficient way to determine the quantum speed limit is known, which makes structural changes in the control landscape a viable alternative. In our work, we show that order parameters capturing correlations between different optimal protocols exhibit universal features (independent of the controlled quantum system) at T_{QSL} . Therefore, even though the quantum speed limit is typically characterized by the minimal cost achieved by optimal control protocols, here we precisely quantify how it affects correlations between optimal control protocols. In addition, we show that a finite number of local minima, present before the transition, generate a multitude of optimal control protocols. This observation allows us to approximately reconstruct the set of optimal protocols for $T > T_{\text{QSL}}$, using information from optimal protocols at $T \leq T_{\text{QSL}}$; the procedure can have useful practical applications, as some minima within the optimal level set may have preferable properties (e.g., robustness to experimental noise) and can find application in experiments.

Finally, the T_{sb} transition marks the onset of a control-symmetry-broken phase. Here, optimal protocols violate the control symmetry dictated by the combination of the Hamiltonian and the initial and target states; consequently, for $T > T_{\text{sb}}$, the search for optima in the landscape cannot be restricted to the space of symmetric control protocols. We emphasize that, unlike for T_c and T_{QSL} , the existence of control symmetry breaking at $T = T_{\text{sb}}$ cannot be deduced *a priori* from any property of the quantum control problem alone; instead, it is observed via direct exploration of the control landscape. This is a prime example for the

manifestation of the physics underlying CLPTs. Moreover, our analysis shows that the T_{sb} transition is associated with a characteristic instability in the infidelity landscape, which occurs and can be observed right before the transition. This instability is a direct manifestation of the symmetry-breaking transition and provides information about the shape of optimal control protocols beyond the transition, which, in turn, are important for experimental optimal control. Namely, even though there exist two degenerate optimal control protocols, the two trajectories explore different parts of Hilbert space, and they are inequivalent in terms of physical observables [41]. In this case, blindly optimizing the infidelity may produce a protocol of suboptimal performance in experiments.

In summary, the CLPT at T_c and T_{sb} signals changes in the shape of optimal protocols and can help in making an informed choice when selecting the control protocol family in experiments. The CLPT at T_{QSL} signals a significant reorganization of the structure of local minima in the optimal level set and exhibits features which are independent of the controlled quantum system. These properties help locate the quantum speed limit and approximately identify the set of optimal control protocols beyond the transition. Overall, we have shown that CLPTs carry a direct significance for both controlled quantum systems and control variables used to manipulate them and can be used to reveal and understand their physical properties.

C. Outlook

By proposing a general framework for an analytical theory of quantum control landscapes, our work sheds new light on the theory of optimal control, which allows one to go beyond purely numerical studies. The perturbative expansions we develop for the effective landscape models reveal the control propagator as a central object in control landscape theory: Indeed, the landscape of any specific control problem can be obtained by projecting this propagator onto the initial and target states, emphasizing its fundamental importance. Crucially, these expansions provide means to analytically derive controlled approximations for the quantum control landscape.

A key insight brought in by our theory is the reconciliation of seemingly disparate control problems under a unified analytical framework: For example, the landscapes of fidelity maximization, energy minimization, and observable extremization all admit a natural analytical description within the same formalism. Another interesting feature is the natural separation of quantum (i.e., physical) and control degrees of freedom, which makes the classical nature of control phase transitions evident. Whether these transitions are of equilibrium (i.e., thermodynamic) or nonequilibrium (i.e., spin-glass-like) character remains an open question for many-body control problems; nevertheless, in the single- or two-qubit problem, we show that (i) CLPTs are associated with critical changes in the optimal level set and (ii) local

traps are absent or can be easily avoided; therefore, in these two cases, the landscape does not possess glassy complexity. In addition, our theory allows us to handle on equal footing continuous and discrete (e.g., bang-bang) control protocols: While the latter result in theoretically appealing $\mathbb{Z}_2$ Ising models, the former are often more relevant for experiments. Importantly, the coupling constants in these classical landscape models do not depend on the family of control protocols used.

To analyze the critical behavior of the control landscape in the vicinity of the quantum speed limit (T_{QSL}), we develop new techniques based on statistical field theory. For bang-bang protocols of duration less than T_{QSL} , we show that multiple local minima can appear in the landscape as a result of the entropy associated with the number of different bang-bang approximations to a single continuous optimal protocol; hence, the bang-bang control landscape may be sensitive to the details of the time discretization and exhibits features not necessarily representative of the behavior of the controlled physical system. Our theory circumvents this issue by providing a uniform description for bang-bang and continuous drives. This is expected to shed new light on the performance of state-of-the-art algorithms, such as the quantum approximate optimization algorithm [33,71,72,75].

Beyond T_{QSL} , we overcome the difficulty of taking into account boundary constraints (e.g., boundedness, bang-bang, etc.) analytically by developing a novel path-integral approach, inspired by techniques used to study spin-glass physics. This allows us to obtain an analytical estimate of the critical scaling behavior of the protocol order parameter in the vicinity of T_{QSL} . Our results pave the way for future applications of statistical field theory methods in the characterization of quantum control landscapes and their phase transitions.

While we focused on two toy models where we computed the coefficients of the control landscape exactly, we expect our methodology to be broadly applicable. The formalism suggests straightforward extensions with little-to-no modifications to optimally control open few-qubit systems subject to dissipation or decoherence (Lindblad control) or exactly solvable many-body systems, such as the transverse-field Ising model, which admit a mapping to a collection of independent two-level systems. Moreover, by considering quantum models in which $H_{0,1}$ are random matrices drawn from a Gaussian unitary ensemble, the ensemble average of the effective landscape model may give rise to a classical spin-glass system already for a single-qubit control problem (cf. Sec. III). When it comes to investigating the properties of optimal control landscapes in nonintegrable quantum many-body systems, we believe one can develop suitable low-rank approximations to the control propagator [76] and then apply our techniques. Hence, our work lays the foundations for a comprehensive theory of control phase transitions.

ACKNOWLEDGMENTS

We thank F. Balducci, S. Cacciatori, M. Ciarchi, L. Piroli, M. Radice, M. Serbyn, and X. Turkeshi for insightful discussions. Funded by the European Union (ERC, QuSimCtrl, 101113633). M.B. was supported by the Marie Skłodowska-Curie Grant Agreement No. 890711. Numerical simulations were performed on the MIPKs HPC cluster.

Views and opinions expressed are, however, those of the authors only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency.

DATA AVAILABILITY

The data and code that support the findings of this article are openly available [77].

APPENDIX A: INFIDELITY EXPANSIONS

In this appendix, we discuss in more detail the analytical expansions introduced in Sec. III. In the first two subsections, we introduce the infidelity Taylor expansion and show how it connects with the Dyson expansion. In the last subsection, we comment on the radius of convergence of the Dyson and Magnus expansions.

We assume a generic Hamiltonian with the form in Eq. (1). The procedure can be generalized to the case with multiple control parameters.

1. Taylor expansions centered around the arbitrary protocol $s_0(t)$

In what follows, we adopt the round and square brackets convention to distinguish between, respectively, the parametric dependence on the quantum duration T and the functional dependence on the control protocol s . Consider the Taylor expansion for the infidelity functional $I[s](T)$:

$$\begin{aligned} I[s](T) &= I[s_0](T) + \int_0^T dt_1 \left. \frac{\delta I}{\delta s(t_1)} \right|_{s_0} \delta s(t_1) \\ &\quad + \frac{1}{2!} \int_0^T dt_1 \int_0^T dt_2 \left. \frac{\delta^2 I}{\delta s(t_1) \delta s(t_2)} \right|_{s_0} \delta s(t_1) \delta s(t_2) + \dots \\ &= \sum_{n=0}^{\infty} \frac{1}{n!} \int_0^T dt_1 \dots dt_n \left. \frac{\delta^n I}{\delta s(t_1) \dots \delta s(t_n)} \right|_{s_0} \delta s(t_1) \dots \delta s(t_n), \end{aligned} \quad (\text{A1})$$

where $\delta s(t) = s(t) - s_0(t)$ denotes deviation from protocol $s_0(t)$ at time t . The generic variational derivative

$$\left. \frac{\delta^n I}{\delta s(t_1) \cdots \delta s(t_n)} \right|_{s_0}$$

can be obtained explicitly from Eq. (2). The evolution operator $\hat{U}[s](T, 0) = \mathcal{T} e^{-i \int_0^T \hat{H}(t) dt}$, acting on the left, evolves quantum states from time $t = 0$ to time $t = T$. Notice that $\mathcal{T}$ represents the time-ordering operator. From the definition, we obtain

$$\begin{aligned} \left. \frac{\delta^n \langle \psi_* | \hat{U}[s](T, 0) | \psi_0 \rangle}{\delta s(t_1) \cdots \delta s(t_n)} \right|_{s_0} \\ = (-i)^n \langle \psi_* | \mathcal{T} \hat{U}[s_0](T, t_1) \partial_s \hat{H}(t_1) \hat{U}[s_0](t_1, t_2) \cdots \\ \times \hat{U}[s_0](t_{n-1}, t_n) \partial_s \hat{H}(t_n) \hat{U}[s_0](t_n, 0) | \psi_0 \rangle, \end{aligned} \quad (\text{A2})$$

where $\partial_s \hat{H}$ denotes the partial derivative of $\hat{H}$ with respect to the control parameter s . The time-ordering operator $\mathcal{T}$ contained in the evolution operator ensures the correct time ordering of the time labels $\{t_n\}$. As a consequence, the functional derivative of the unitary operator is symmetric under permutations of the time labels. For example, for $n = 1$, we have

$$\begin{aligned} \left. \frac{\delta \langle \psi_* | \hat{U}[s](T, 0) | \psi_0 \rangle}{\delta s(t_1)} \right|_{s_0} \\ = (-i) \langle \psi_* | \hat{U}[s_0](T, t_1) \partial_s \hat{H}(t_1) \hat{U}[s_0](t_1, 0) | \psi_0 \rangle. \end{aligned} \quad (\text{A3})$$

For $n = 2$,

$$\begin{aligned} \left. \frac{\delta^2 \langle \psi_* | \hat{U}[s](T, 0) | \psi_0 \rangle}{\delta s(t_1) \delta s(t_2)} \right|_{s_0} \\ = (-i)^2 \langle \psi_* | \hat{U}[s_0](T, t_1) \partial_s \hat{H}(t_1) \\ \times \hat{U}[s_0](t_1, t_2) \partial_s \hat{H}(t_2) \hat{U}[s_0](t_2, 0) | \psi_0 \rangle, \end{aligned} \quad (\text{A4})$$

where we assume $t_1 > t_2$ (if $t_2 > t_1$ one exchanges $t_1 \leftrightarrow t_2$ in the right-hand side to preserve the correct time ordering). The infidelity Taylor expansion can be easily written in terms of the variational derivative of the unitary operator $\hat{U}[s](T)$ given above. Introducing the variable $z = \langle \psi_* | \hat{U}[s](T, 0) | \psi_0 \rangle$ and its complex conjugated $\bar{z}$ as short-hand notation, we have $I[s](T) = 1 - \bar{z}z$ and

$$\left. \frac{\delta I}{\delta s(t_1)} \right|_{s_0} = -2\text{Re} \left[\bar{z} \frac{\delta z}{\delta s(t_1)} \right]_{s_0}, \quad (\text{A5})$$

$$\begin{aligned} \left. \frac{\delta^2 I}{\delta s(t_2) \delta s(t_1)} \right|_{s_0} &= -2\text{Re} \left[\bar{z} \frac{\delta^2 z}{\delta s(t_2) \delta s(t_1)} \right]_{s_0} \\ &+ \frac{\delta \bar{z}}{\delta s(t_2)} \bigg|_{s_0} \frac{\delta z}{\delta s(t_1)} \bigg|_{s_0}, \end{aligned} \quad (\text{A6})$$

where Re indicates the real part. Notice that in $t_1 = t_2$ (or in

$t_i = t_j$ when $n > 2$) the variational derivative has a removable discontinuity. Our analysis is not affected by these discontinuities: In the functional expansions, the derivatives appear under integration, and the discontinuities involve a zero-measure set.

In Sec. III, we conveniently rewrite the infidelity in terms of the orthogonal evolution operator [cf. Eq. (10)]. Let us now consider the Taylor expansion in the orthogonal operator space. Repeating the same steps as above, we obtain

$$\begin{aligned} \left. \frac{\delta^n I[s](T)}{\delta s(t_1) \cdots \delta s(t_n)} \right|_{s_0} \\ = -(1/2) \vec{n}_* \cdot \mathcal{T} M[s_0](T, t_1) \partial_s m M[s_0](t_1, t_2) \cdots \\ \times M[s_0](t_{n-1}, t_n) \partial_s m M[s_0](t_n, 0) \vec{n}_0, \end{aligned} \quad (\text{A7})$$

where $m(t)$ is the generator of the orthogonal evolution operator associated with the Hamiltonian $\hat{H}(t)$. By comparison with the same expansion in the unitary operator space, we see that the orthogonal operator space allows a considerable simplification of each term in the series: The n th-order coefficients of the infidelity expansion simply correspond to the n th functional derivative of the infidelity. In particular, in the orthogonal operator space, each term is simply given by an alternating product of evolution operators $M[s_0](t_i, t_{i+1})$ and generators $\partial_s m(t_i)$. For example, for $n = 1$,

$$\begin{aligned} \left. \frac{\delta I[s](T)}{\delta s(t_1)} \right|_{s_0} &= -(1/2) \vec{n}_* \cdot M[s_0](T, t_1) \\ &\times \partial_s m(t_1) M[s_0](t_1, 0) \vec{n}_0, \end{aligned} \quad (\text{A8})$$

and for $n = 2$,

$$\begin{aligned} \left. \frac{\delta^2 I[s](T)}{\delta s(t_1) \delta s(t_2)} \right|_{s_0} \\ = -(1/2) \vec{n}_* \cdot M[s_0](T, t_1) \partial_s m(t_1) \\ \times M[s_0](t_1, t_2) \partial_s m(t_2) M[s_0](t_2, 0) \vec{n}_0 \end{aligned} \quad (\text{A9})$$

when $t_1 > t_2$.

2. Connection between Taylor and Dyson expansions

As in Sec. III, let us now perform the reference-frame transformation

$$\begin{aligned} \hat{H}(t) &\mapsto \hat{H}'(t) = \hat{U}_0^\dagger(t, 0) [\hat{H}(t) - i\partial_t] \hat{U}_0(t, 0), \\ \hat{U}[s](T, 0) &\mapsto \hat{U}_0(T, 0) \hat{U}'[s](T, 0), \end{aligned}$$

where

$$\begin{aligned}\hat{U}_0(T, 0) &= \exp(-iT\hat{H}_0), \\ \hat{U}'[s](T, 0) &= \mathcal{T} e^{-i \int_0^T dt \hat{H}'(t)}.\end{aligned}$$

We obtain $\hat{H}'(t) = U_0^\dagger(t, 0)\hat{H}_1(t)\hat{U}_0(t, 0)$ so that the new Hamiltonian $\hat{H}'(t)$ is proportional to the protocol $s(t)$. Consequently, for $s = 0$, the new evolution operator $\hat{U}'[s](T, 0)$ reduces to the identity operator: $\hat{U}'[0](T, 0) = \mathbb{1}$.

Similarly, in the dual space, we perform the reference-frame transformation

$$m(t) \mapsto m'(t) = M_0(T, 0)^\dagger [m(t) - \partial_t] M_0(T, 0),$$

$$M[s](T, 0) \mapsto M_0(T, 0) M'[s](T, 0),$$

where

$$\begin{aligned}M_0(T, 0) &= \exp(Tm_0), \\ M'[s](T, 0) &= \mathcal{T} e^{\int_0^T dt m'(t)}.\end{aligned}$$

Again, the new generator $m'(t)$ in the orthogonal operator space is proportional to the protocol $s(t)$: $m'(t) = M_0(T, 0)^\dagger m_1(t) M_0(T, 0)$. Hence, for $s = 0$, the evolution operator reduces to the identity operator: $M'[0](T, 0) = \mathbb{1}$.

The properties $\hat{U}'[0](T, 0) = \mathbb{1}$ and $M'[0](T, 0) = \mathbb{1}$ considerably simplify the coefficients appearing in the infidelity expansion. From Eqs. (A2) and (A7), we obtain

$$\begin{aligned}\frac{\delta^n \langle \psi_* | \hat{U}'[s](T, 0) | \psi_0 \rangle}{\delta s(t_1) \cdots \delta s(t_n)} \Big|_{s_0} \\ = (-i)^n \langle \psi_* | \hat{U}_0(T, 0) \mathcal{T} \partial_s \hat{H}'(t_1) \cdots \partial_s \hat{H}'(t_n) | \psi_0 \rangle\end{aligned}\quad (\text{A10})$$

and

$$\begin{aligned}\frac{\delta^n I[s](T)}{\delta s(t_1) \cdots \delta s(t_n)} \\ = -(1/2) \vec{n}_* \cdot M_0(T, 0) \mathcal{T} \partial_s m'(t_1) \cdots \partial_s m'(t_n) \vec{n}_0.\end{aligned}$$

Notice that for $n = 1, 2$ this last result corresponds to the “interaction terms” in Eq. (15); this expression generalizes the Dyson expansion terms presented in Sec. III to an arbitrary order $n \in \mathbb{N}$.

3. Convergence of the landscape expansions

The convergence properties of the landscape expansion introduced in Sec. III sensibly depend on the properties of the controlled quantum system and the space of allowed control protocols. The quantum control problems introduced in Sec. II are characterized by the following two properties: They (i) have a finite-dimensional Hilbert space and (ii) are defined for bounded control protocols, $|s(t)| \leq 1$, $\forall t \in [0, T]$. As a consequence, in our case, the norm of the Hamiltonian is bounded at each time, $\max_t ||H(s(t))|| < \infty$.

In this subsection, for the sake of generality, we can relax hypothesis (ii) and allow the more general class of control protocols satisfying

$$C_H \equiv \int_0^T dt ||H[s(t)]|| < \infty. \quad (\text{A11})$$

This second class contains the class of bounded control protocols, as in the latter case we have $C_H = T \max_t ||H[s(t)]|| < \infty$.

a. Dyson and Taylor expansions

Under the above assumptions, the Dyson (and Taylor) expansions are absolutely convergent for any fixed $T < \infty$. To see this, consider the generic evolution operator

$$\begin{aligned}\hat{U}(T, 0) &= \mathcal{T}_t \exp\left(-i \int_0^T dt \hat{H}(t)\right) \quad \text{or} \\ M(T, 0) &= \mathcal{T}_t \exp\left(\int_0^T dt m(t)\right)\end{aligned}$$

and observe that each term in the associated Dyson series is bounded. For example,

$$\begin{aligned}\left\| \mathcal{T}_t \exp\left(\int_0^T dt m(t)\right) \right\| &= \left\| \sum_{n=0}^{\infty} \frac{1}{n!} \int_0^T dt^n \prod_{i=0}^n m(t_i) \right\| \\ &\leq \sum_{n=0}^{\infty} \left\| \frac{1}{n!} \prod_{i=1}^n \int_0^T dt_i m(t_i) \right\| \\ &\leq \sum_{n=0}^{\infty} \frac{1}{n!} \left(\int_0^T dt ||m(t_i)|| \right)^n,\end{aligned}$$

where we use the triangular inequality and the submultiplicative property of the operator (matrix) norm. Thus, for any finite T , the Dyson and Taylor expansions are bounded in norm by the absolutely convergent numerical series

$$\sum_{n=0}^{\infty} a_n = \sum_{n=0}^{\infty} \frac{(C_m)^n}{n!} = e^{C_m}.$$

From this result, we extract (i) a *conservative* estimate of the order of the expansion, P , after which higher-order terms decrease in norm, and (ii) a bound on the error associated with a finite order of truncation, ΔI_{err} .

Let P be the integer for which the numerical series $\{a_n\}$ achieves its maximum value, $P = \arg\max_{n \in \mathbb{N}} |a_n|$. In particular, we estimate P by imposing the condition $a_{P-1} = a_P$, that yields

$$P = C_m. \quad (\text{A12})$$

Thus, P grows linearly with C_m (cf. $C_m = T \max_t ||m[s(t)]||$, for bounded protocols). For example, in the single-qubit problem in Eq. (18), we have $\max_s ||m(s)|| \simeq h_x = \sqrt{5}$ and

$T_{\text{QSL}} \simeq 2.5$ so that $P(T_{\text{QSL}}) \simeq 6$. In the two-qubit case in Eq. (19), $\max_s \|m(s)\| \simeq \sqrt{2}h_x = \sqrt{10}$ and $T_{\text{QSL}} \simeq 3.0$ yield $P(T_{\text{QSL}}) \simeq 10$. In both cases, the analytical estimate (A12) returns $P(T_{\text{QSL}}) > 3$, consistent with the numerical results in Fig. 3.

Let $N \in \mathbb{N}$ be the finite order of truncation of the Dyson expansion. A bound on the error ΔI_{err} between the exact infidelity and the truncated Dyson expansion is given by the norm of the terms neglected by the truncation:

$$\begin{aligned} \Delta I_{\text{err}} &= \left\| \sum_{n=N+1}^{\infty} \frac{1}{n!} \int_0^T dt \prod_{i=0}^n m(t_i) \right\| \\ &\leq \sum_{n=N+1}^{\infty} a_n = \sum_{n=N+1}^{\infty} \frac{(TC)^n}{n!}. \end{aligned}$$

b. Magnus expansion

The well-known sufficient condition for the convergence of the Magnus expansion reads [60]

$$\int_0^T dt \|m(t)\| \leq \pi. \quad (\text{A13})$$

In the case of Eq. (16), this yields a relatively short radius where convergence is guaranteed. In particular, for the single- and two-qubit problems, we have, respectively, $T \lesssim 1.4$ and $T \lesssim 1.0$ (to be compared with $T_{\text{QSL}} \simeq 2.5$ and $T_{\text{QSL}} \simeq 3.0$).

However, the following remarks are in order. First, Eq. (A13) gives a *sufficient* condition so that the Magnus expansion may still converge even when the bound is violated [58]. Second, the right-hand side of Eq. (A13) depends on the specific protocol $s(t)$ appearing in $m(t)$. Consequently, although not uniformly, the Magnus expansion may still converge in a subset of the protocol space. As a specific example, consider the single-qubit problem, where the optimal protocol $s_{\Delta_0(T)}$ for $T \leq T_{\text{QSL}}$ yields (cf. Sec. IV)

$$\int_0^T dt \|m(t)[s_{\Delta_0(T)}]\| = T_c h_x < \pi, \quad T \leq T_{\text{QSL}}. \quad (\text{A14})$$

Third, remember that the Magnus expansion can be centered around an arbitrary protocol $s_0(t)$ via reference-frame transformation. This property can be used to obtain a convergent Magnus expansion around a desired control protocol.

c. Cumulants expansion

We use the cumulants expansion to obtain, from the Magnus expansion, an infidelity expansion with scalar-valued coefficients; hence, the sufficient criterion (A13) also applies in this case.

There are no other constraints on the radius of convergence set by the cumulants expansion. To see this, observe that (i) the cumulants expansion is a Taylor expansion of the

quantity $\log\{1 - I(T)[s]\}$, $I(T)[s] \in [0, 1]$, and (ii) the radius of convergence of the Maclaurin series of the function $\log(1 + x)$, $x \in \mathbb{R}$, is unity.

APPENDIX B: STOCHASTIC DESCENT SIMULATIONS

In this section, we provide additional information on the SD simulations used in Sec. III to sample locally optimal bang-bang protocols in the quantum control landscape.

In the stochastic descent algorithm, the initial bang-bang protocol is drawn from a uniform distribution site per site: $s_i = \pm 1$. At each iteration, a spin flip $s_j \mapsto -s_j$ at random site j is performed—the modified protocol is accepted whenever it leads to decrement in the infidelity; otherwise, the spin flip is undone. The algorithm terminates when a single spin flip in all sites does not decrease infidelity. For each T , we collect 250 locally optimal protocols from 250 independent SD runs. From the set of locally optimal protocols, $\mathcal{S}$, we estimate the average protocol $\langle s_i \rangle_{\mathcal{S}}$ and compute $q_{\text{BB}}(T)$ [cf. Eq. (5)]. Notice that we do not observe any sensible change in the curve $q_{\text{BB}}(T)$ by doubling the number of locally optimal protocols sampled. In the




FIG. 13. Stochastic descent results complementing Fig. 3. In the plot, we compare the minimum infidelity $\min_s I(T)[s]$ achieved in the approximated control landscapes with respect to the exact control landscape (showed in red). Magnus is the only expansion preserving the infidelity's domain $[0, 1]$ and closely following the exact (red) curve. The Dyson and cumulants expansions break by construction the unitarity of the evolution operator, and this leads to strong deviations in the corresponding curves (for $T \gtrsim 2.5$ and $T \gtrsim 2.0$ in the single- and two-qubit problem, respectively). The letters D, M, and C in the legend stand for Dyson, Magnus, and cumulant expansion, while the adjacent integer specifies the order of truncation.

exact infidelity landscape, we let SD search in the space of bang-bang protocols with $N_{\text{ex}} = 5000$ bang-bang steps. In the case of the approximated infidelity landscape, SD searches in the space of bang-bang protocols with $N = 500$ bang-bang steps. The different numbers N_{ex} and N of bang-bang steps in the two infidelity landscape (exact and approximated, respectively) is motivated by the different scalings in the number of bang-bang steps exhibited by the numerical $q_{\text{BB}}(T)$ curve. Intuitively, the different scaling behavior is related to the different type of discretization: In the exact infidelity, the bang-bang protocol directly enters the unitary operator; instead, the infidelity expansions are derived in the more general space of piecewise continuous protocols, and the bang-bang constraint is imposed after the expansion is performed.

In Secs. III B and III C, we show the comparison between the curve $q_{\text{BB}}(T)$ obtained from SD performed over the exact infidelity landscape and the approximated infidelity landscape. It is also interesting to compare the minimum infidelity obtained. In the single-qubit case, from Fig. 13 (left column), we see that for small T all the expansions at the highest order have minimum infidelity close to the exact landscape result. For $T \gtrsim 3.0$, the Dyson expansion [Fig. 13(a)] breaks the unitarity of the evolution operator and infidelity is not bounded within $[0, 1]$. On contrary, the Magnus and cumulants expansions [Figs. 13(b) and 13(c)] approximate the minimum infidelity curve in the whole interval $T \in [0, 3.5]$. In the two-qubit case, the Dyson and

cumulant expansions [Figs. 13(d) and 13(e)] break down beyond $T \simeq 2$, while the Magnus expansion [Fig. 13(f)] reproduces the exact landscape minimum infidelity curve in the whole interval $T \in [0, 4]$.

Let us now show how the number of bang-bang steps N affects the order parameter curve $q_{\text{BB}}(T)$. In Fig. 14, the curve $q_{\text{BB}}(T)$ is plotted for the Magnus expansion truncated at third order for different numbers of bang-bang steps N . As N increases, the main change in $q_{\text{BB}}(T)$ occurs around the critical points: T_c , T_{QSL} for the single-qubit [Fig. 14(a)] and T_{sb} , T_{QSL} for the two-qubit problem [Fig. 14(b)]. Notice that, in the two-qubit case, convergence of $q_{\text{BB}}(T)$ is significantly slower around the critical points $T = T_{\text{QSL}}$. Nevertheless, in both cases, the qualitative behavior of the curve $q_{\text{BB}}(T)$ is already visible for $N \gtrsim 250$.

Finally, in Fig. 15, we directly compare the average protocols $\langle s_i \rangle_S$ extracted from the different SD simulations in the single-qubit problem. Besides the order parameter $q(T)$ and fidelity $F(T)$, which effectively reduce the complexity of the system to a single value, in this plot we observe precisely how much the approximated landscapes reproduce the correct optimal protocol. Even though




FIG. 14. Stochastic descent results complementing Fig. 3. Finite-size scaling of the curve $q_{\text{BB}}(T)$ obtained via Magnus expansion, as the number of bang-bang steps N is increased. For reference, in black we report the curve $q_{\text{BB}}(T)$ obtained in the exact control landscape (with $N_{\text{ex}} = 5000$ bang-bang steps). In the legend, “M3” stands for Magnus expansion truncated at third order.




FIG. 15. Single-qubit problem: average protocols $\langle s_i \rangle$ for different quantum durations $T = 0.80, 1.50, 2.50, 2.60$. The average is performed over 250 bang-bang protocols sampled by SD in the exact and approximated control landscapes—D3, Dyson expansion truncated at third order; M3, Magnus expansion truncated at third order; C5, cumulants expansion truncated at fifth order (cf. Sec. III). Interestingly, from (c) we observe that the shape of optimal protocols may qualitatively differ in the approximated and exact landscapes. This suggests some care when using the infidelity expansions to predict the shape of optimal protocols, even when they reproduce quantitatively well the behavior of other observables [such as $q_{\text{BB}}(T)$ in Fig. 3].

the Magnus expansion at third order better approximates the order parameter curve $q_{\text{BB}}(T)$, the average sampled protocols in Magnus do not always follow the qualitative behavior of the exact result [cf. Fig. 15(c)]. On the contrary, the average protocols sampled by the Dyson and cumulants expansions possess the qualitative structure of the exact result. From this result, we understand that the quality of the approximation associated with the analytical expansions may sensibly depend on the features one is interested in comparing.

APPENDIX C: QUADRATIC STABILITY ANALYSIS

In this section, we provide details about the quadratic stability analysis presented in Sec. IV. The analysis is based on the numerical diagonalization of the Hessian operator $J_{t_1 t_2}(T)$ appearing in Eqs. (15) and (22). Numerically, we discretize the operator in an $L \times L$ matrix and account for the scaling prefactor $1/L$.

First, we show how the Hessian eigenvalues scale as a function of the discretization points. Let us consider explicitly the single-qubit problem and the quadratic operator relative to the protocol $s_{\Delta_0(T)}$, defined in Eq. (25) and




FIG. 16. Single-qubit problem. (a), (b) Spectrum of the discretized second-order operator in Eq. (9) relative to the protocol $s_{\Delta_0(T)}$ [cf. Eq. (24)], optimal for $T \rightarrow T_{\text{QSL}} \simeq 2.51$. We observe the characteristic scaling of eigenvalues $\lambda_n \sim n^{-2}$ in the $L \rightarrow \infty$ limit. (c), (d) Eigenfunctions associated with the two positive and the first four negative eigenvalues at $T = 2.52$. Notice the close similarity to the Fourier basis.

optimal in the $[T_c, T_{\text{QSL}}]$ interval. In Figs. 16(a) and 16(b), we show how eigenvalues scale as $L = 16, 32, 64$ at $T = 2.50, 2.52$. The convergence in L is relatively fast for both positive and negative eigenvalues. Notice the change in sign happening exactly after T_{QSL} : Beyond this critical point, $L - 2$ negative eigenvalues are present. It is also interesting to visualize the eigenfunctions associated with the numerical diagonalization. In Figs. 16(c) and 16(d), we show the first six eigenfunctions of the quadratic operator for $L = 64, T = 2.52$. Except for the eigenfunction $f^{(2)}$, the eigenbasis $\{f^{(n)}\}$ appears to be a (relatively small) deformation of the standard Fourier basis on the $[0, T]$ domain. The role of the $f^{(2)}$ eigenfunctions becomes clear when considering deformations of $s_{\Delta_0(T)}$ along the directions of the eigenfunctions (see Appendix E).

Second, in Fig. 17, we show how eigenvalues of the Hessian operator change around the $T_{\text{QSL}} \simeq 2.95$ in the two-qubit problem. In particular, the infidelity expansion is centered around one of the two isolated optimal protocols found numerically by the LMC algorithm, for each fixed T (cf. Sec. V). In LMC, the time interval $[0, T]$ is divided in $L = 64$ uniform points, and protocols are then piecewise constant on L steps. To reduce finite-size effects in the computation of the spectrum, the Hessian operator is instead discretized in $L' = 256$ steps: The optimal protocol with L' steps is obtained by interpolating the optimal protocol with L steps found by LMC (using function interp from PYTHON's package NUMPY). Notice that, for $T_{\text{sb}} \leq T \leq T_{\text{QSL}}$, the two isolated optimal protocols present in the landscape are related to each other via the control problem symmetry $s(t) \leftrightarrow -s(T - t)$ [40]. As the infidelity is invariant under this symmetry transformation, the Hessian operators relative to the two optimal protocols possess the same spectrum. The Hessian spectrum shows qualitatively similar behavior to the single-qubit case [cf. Fig. 6(a)] except for the number of nonvanishing




FIG. 17. Two-qubit problem: spectrum of the Hessian operator around $T_{\text{QSL}} \simeq 2.95$ (vertical gray line). The infidelity expansion is centered around one of the two optimal protocols found by the LMC algorithm introduced in Sec. V (the spectrum is identical for the two optimal protocols as discussed in Appendix C). From the plot, we see $L - 4$ eigenvalues vanishing for $T \rightarrow T_{\text{QSL}}$, up to deviations due to numerical errors. The result is analogous to the single-qubit case, shown in Fig. 6.

eigenvalues: four for the two-qubit and two for the single-qubit. Notice that, in the two-qubit case, the $L - 4$ eigenvalues are not exactly zero for $T \geq T_{\text{QSL}}$, since the protocol chosen as the center of the infidelity expansion (at each fixed T) is optimal within numerical error; in this case, the error depends on the finite number of discretization steps, L , L' , and the finite value of the parameter β appearing in the LMC update rule (28).

APPENDIX D: GRADIENT-FREE LANGEVIN-MONTE CARLO SIMULATIONS

In this appendix, we present in more detail results from the LMC dynamics. We divide the discussion into three sections, respecting the three consecutive stages of LMC dynamics: thermal relaxation, diffusion along the optimal level set valleys, and equilibrium sampling, as each one of the three stages is affected by CLPTs.

First, let us clarify the reason behind the simple dynamics of the algorithm. We discarded the gradient term, normally present in the Langevin update rule, because we are interested in exploring the optimal ($I(T)[s] = 0$) level set, where the gradient term (as well as $L - n_+$ eigenvalues of the Hessian operator, with the positive integer n_+ dependent on the quantum control problem) are expected to vanish (cf. Sec. IV). Reinstating it results in a faster convergence to the optimal level set in the first part of the LMC simulations. However, its contribution during the level set exploration is expected to be less relevant.

Second, this algorithm uses purely stochastic moves to explore the optimal level set. Note that the update rule above is likely not the most efficient one: For example, one may exploit the Hessian analysis and restrict the stochastic move to the Hessian eigenspace associated with vanishing eigenvalues. This modification would probably result in a higher acceptance rate during the Metropolis rule (since in the update rule the infidelity would be preserved up to second-order corrections), albeit at the cost of requiring more computational time due to implementing an update in a rotated basis. Nevertheless, as a first approach, we prefer to avoid any bias in the numerical exploration of the optimal level set at the cost of a less efficient algorithm.

We tested LMC dynamics varying parameters in the following intervals: $T \in [0, 3]$, $\sigma \in [10^{-4}, 10^{-3/2}]$, and $\beta \in [10^2, 10^8]$. At the beginning of each LMC run, the initial protocol configuration $s_0 = (s_{0,1}, \dots, s_{0,L})$ is chosen randomly by drawing each component $s_{0,i}$ from the uniform distribution in $[-1, 1]$. We define a single LMC iteration as L local attempted moves.

1. Thermal relaxation stage

In Figs. 18 and 19, we show how infidelity changes as the LMC dynamics evolves the protocol for $\beta = 10^3, 10^4, \dots, 10^8$, $\sigma = 10^{-3/2}$, in the single- and two-qubit problem, respectively. In the single-qubit case, we plot the




FIG. 18. Single-qubit problem. The plots show the infidelity during thermal relaxation, averaged over ten independent LMC runs. The parameter T , different in the four panels, is changed to investigate the effects of CLPTs (cf. $T_c \simeq 0.98$, $T_{\text{QSL}} \simeq 2.51$). Each curve refers to a different Monte Carlo inverse temperature β , while $\sigma = 10^{-3/2}$ and $L = 512$ are fixed. In particular, we plot the absolute value of the difference between the infidelity and the exact optimal infidelity, analytically known in the single-qubit case (but not in the two-qubit case; cf. Fig. 19). Interestingly, in each panel, the six curves lie on top of each other until thermalization occurs. A simple effective one-dimensional toy model [black dashed line in (a) and (b); cf. Eq. (D2)] reproduces the typical scaling observed in the numerical data when T is sufficiently different from the critical value $T_{\text{QSL}} \simeq 2.51$. At $T = 2.50$ [(c); just before T_{QSL}], the $\simeq 1/2$ slope differs from the $= 1$ slope characteristic of the late-time scaling of the one-dimensional model in Eq. (D4). On the contrary, at $T = 2.52$ [(d); just beyond T_{QSL}], LMC dynamics thermalizes much faster.

infidelity difference between the LMC and the theoretical optimal value computed from the exact solution of the problem [39]. In the two-qubit case, the exact solution is not known, and we plot the discrete derivative of the infidelity as a function of the LMC iteration. In both cases, we average over a set of ten independent LMC runs. When the Monte Carlo inverse temperature β is large enough, the average motion in the thermal relaxation stage does not depend sensibly on it. In this case, the infidelity (derivative) curve displays, in general, two consecutive transient behaviors: The first one characterizes the dynamics at early iterations, whereas the second one originates with a sudden transition (occurring around LMC iteration $n = 10^2$ for $\sigma = 10^{-3/2}$) and terminates once the system reaches the thermal equilibrium in the low-infidelity region. As visible from the plots, these two substages possess distinct characteristic scaling behaviors independent of β .




FIG. 19. Two-qubit problem. The plots show the discrete derivative of infidelity, dI_n/dn during thermal relaxation, averaged over ten independent LMC runs. The parameter T , different in the four panels, is changed to investigate the effects of CLPTs (cf. $T_{\text{sb}} \simeq 1.57, T_{\text{QSL}} \simeq 2.95$). Each curve refers to a different Monte Carlo inverse temperature β , while $\sigma = 10^{-3/2}$ and $L = 512$ are fixed. As in the single-qubit system, a simple effective one-dimensional toy model [black dashed line in (a) and (d); cf. Eq. (D2)] can approximate the typical scaling observed in data away from the critical values. The change in slope (from $\simeq 1$ to $\simeq 1/2$) observed just before T_{QSL} in the single-qubit case (cf. Fig. 18) is also visible here around T_{sb} [(b) and (c)] and just before T_{QSL} [(e)]. The discrete numerical derivative is computed with the function gradient from PYTHON's package NUMPY. We decreased the opacity of each curve to 10% after thermalization occurred.

The characteristic scaling behaviors observed during the LMC thermal relaxation stage for $\beta \gtrsim 10^3$ are captured to some extent by an effective one-dimensional toy model. Consider a one-dimensional system $x_n \in \mathbb{R}$ evolving according to LMC dynamics in Eq. (27) in the landscape $I(x) = \alpha|x - x_{\min}|$ from a random initial position $x_0 \neq x_{\min}$. The transition probability for a single iteration is then

$$P(x_{n+1}|x_n) = \mathcal{N}_{x_{n+1}}(x_n, \sigma) p_{\text{acc}}(\Delta I_n), \quad (\text{D1})$$

where $\mathcal{N}_x(x_n, \sigma)$ is a Gaussian distribution centered at x_n with variance σ^2 , $p_{\text{acc}}(\Delta I_n) = \max(1, e^{-\beta \Delta I_n})$ is the

Metropolis acceptance probability, and $\Delta I_n = I(x_{n+1}) - I(x_n)$ is the change in infidelity during the iteration. On average, the final state is then

$$\mathbb{E}(x_{n+1}|x_n) = \int_{\mathbb{R}} dx_{n+1} x_{n+1} P(x_{n+1}|x_n).$$

For $\beta \rightarrow \infty$, thermal effects play no role, since $p_{\text{acc}}(\Delta I_n) \rightarrow \text{sgn}(\Delta I_n)$. In this case, the dynamics depends only on σ and the relative distance $\ell_n = |x_n - x_{\min}|$ between x_n and the minimum of the infidelity at $x_{\min}$. Through a change of coordinate, we fix $x_n = 0$ and $x_{\min} > 0$. Then, in the limit $\beta \rightarrow \infty$, we obtain

$$\begin{aligned} \mathbb{E}(x_{n+1}|x_n = 0) &= \int_{\mathbb{R}} dx_{n+1} x_{n+1} \mathcal{N}_{x_{n+1}}(0, \sigma) \text{sgn}(\Delta I_n) \\ &= \int_0^{2\ell_n} dx_{n+1} x_{n+1} \mathcal{N}_{x_{n+1}}(0, \sigma) \\ &= \sqrt{\frac{\sigma^2}{2\pi}} (1 - e^{-(1/2)(2\ell_n)^2/\sigma^2}). \end{aligned}$$

Thus, while averaging over different LMC realizations, we obtain the differential equation for ℓ_n :

$$\frac{d\ell_n}{dn} \approx -\mathbb{E}(x_{n+1}|x_n = 0), \quad (\text{D2})$$

where we promote n to a continuous variable assuming sufficiently small changes in ℓ_n over single LMC iterations. This equation admits simple closed-form solutions in the two limiting cases

$$\ell_n/\sigma \gg 1: \quad \ell_n \sim -\alpha_1 n, \quad (\text{D3})$$

$$\ell_n/\sigma \ll 1: \quad \ell_n \sim \alpha_2 n^{-1}, \quad (\text{D4})$$

where the two prefactors $\alpha_{1,2}$ depend on σ . These two analytical limiting cases are related with the two characteristic scaling behaviors visible in Figs. 18 and 19. In the single-qubit system, the prediction in Eq. (D2) fits qualitatively well the whole numerical curves away from the $T_{\text{QSL}} \simeq 2.51$ transition [black dashed curve in Figs. 18(a) and 18(b)]; in the two-qubit problem, the same is true away from the $T_{\text{sb}} \simeq 1.57$ transition and again before the $T_{\text{QSL}} \simeq 2.95$ transition [black dashed curve in Figs. 19(a) and 19(d)]. In this way, we gain intuition about the two characteristic scaling behaviors visible during LMC thermal relaxation stage: The first (second) transient behavior characterizes the motion when the protocol $\mathbf{s}_n$ is sufficiently far away from (close to) the low-infidelity region with respect to the average jump size σ of the Langevin update rule.

Interestingly, the scaling behavior changes across the CLPT at the quantum speed limit. For example, in the single-qubit system, compare $T = 2.50$ with $T = 2.52$

(cf. $T_{\text{QSL}} \simeq 2.51$) or in the two-qubit system $T = 2.94$ with $T = 2.97$ (cf. $T_{\text{QSL}} \simeq 2.95$). We observe that the early-time characteristic behavior remains qualitatively unaffected by the T_{QSL} transition: We deduce that there are no significant changes around T_{QSL} in the infidelity landscape in the high-infidelity region. On the contrary, the T_{QSL} transition heavily affects the scaling behavior in the late-time thermal relaxation stage. Observe that the late-time characteristic scaling at the two sides of the T_{QSL} transition has opposite behavior: It slows down when $T \rightarrow T_{\text{QSL}}^-$, and it becomes fast when $T \rightarrow T_{\text{QSL}}^+$. On a qualitative level, the different behavior at the two sides of the T_{QSL} transition can be understood as follows. In the limit $T \rightarrow T_{\text{QSL}}^-$, the stochastic search for the unique optimal protocols is complicated by an increasing number of suboptimal protocols having similar infidelity as the true optimal protocol. In the other case, for $T \rightarrow T_{\text{QSL}}^+$, there is an extensive number of optimal protocols, and the stochastic dynamics easily converges to one of the many possible minima. In more detail, at $T = T_{\text{QSL}}^-$, the late-time characteristic scaling is $n^{-1/2}$ (notice the scaling n^{-1} observed away from the T_{QSL} transition). This different scaling behavior is phenomenologically explained by the one-dimensional effective model by assuming an effective potential $V(|x|) \propto |x|^{1/2}$ in the $\ell_n \ll \sigma$ case. On the other side, at $T = T_{\text{QSL}}^+$, the second transient is not visible at all: All curves converge approximately simultaneously to the low-infidelity region shortly after the end of the first transient. This other behavior is due to the presence of many optimal protocols and is effectively described in the one-dimensional model by an extended flat valley $V = 0$ with a length greater than the σ .

The two-qubit problem presents additional features not observed in the single-qubit case. Around the $T_{\text{sb}} \simeq 1.57$ transition, the early-time scaling behavior is different than what is observed for other values of T . In particular, for $T = 1.56$, the scaling during the thermal relaxation phase is approximately algebraic (with a slightly different slope in the early- and late-time thermal relaxation substages). For $T = 1.60$, the same scaling persists, but an additional “bump” in the infidelity curve is now separating the two substages. Although a more detailed analysis is needed to better understand this other behavior, intuitively the bump may be related to a saddle point present in the landscape, slowing down the thermal relaxation dynamics. Finally, the characteristic late-time scaling $n^{1/2}$, observed in the single-qubit problem just before the T_{QSL} transition, appears also in the two-qubit problem not only before its T_{QSL} transition, but also around the T_{sb} transition.

Last, let us point out that, in all our simulations in the single-qubit landscape, LMC always thermalizes around a global optimal protocol without getting trapped in any local minimum. On the contrary, in the two-qubit problem for $L = 16, 32, 64$, we sometimes found the LMC dynamics

trapped in a local minimum of the landscape for $T \in [2.0, 3.3]$, approximately. In this case, a simulated annealing protocol combined with a postselection among a set of five thermal relaxation runs is enough to guarantee the convergence of LMC to the optimal level set in all of our simulations.

2. Diffusion stage

In order to develop some intuition about the diffusion stage, it is convenient to consider the theoretical limit $\sigma \rightarrow 0$ at fixed $\beta \gg \sigma^2$. In this limit, the acceptance probability P_{acc} approaches unity, and the dynamics becomes effectively a free random walk in L dimensions. In practice, we consider a small but finite σ so that $P_{\text{acc}} \simeq 1$ only during the very early diffusion stage. As $\mathbf{s}_n$ moves, P_{acc} suppresses the motion in the directions associated with increasing infidelity. As long as the dynamics involves a neighborhood of the starting point, the motion of $\mathbf{s}_n$ can be decomposed along the local flat directions (*tangent* directions) and the local infidelity-increasing directions (*orthogonal* directions): The tangential motion is approximately free, whereas the motion along orthogonal directions is suppressed by P_{acc} . Hence, LMC becomes effectively a random walk in L dimensions restricted by the infidelity landscape to move along the valley’s flat directions. Asymptotically, $\mathbf{s}_n$ diffuses through the valleys and explores a global portion of the optimal level set.

Using the random walk as a theoretical limiting behavior, we characterize the diffusive motion along the optimal level set by considering the covariance matrix associated with a protocol’s fluctuations. More precisely, we sample protocols at different iterations $\mathcal{D}_M = \{\mathbf{s}_{n_1} \dots \mathbf{s}_{n_M}\}$ and compute the covariance matrix

$$\sigma_{M,ij} = \langle s(t_i)s'(t_j) \rangle_M - \langle s(t_i) \rangle_M \langle s'(t_j) \rangle_M, \quad (\text{D5})$$

where the average is performed over the set $\mathcal{D}_M$. As we discuss above, the diffusion stage is expected to have three different timescales, each associated with some properties in the covariance matrix (D5).

- (1) *Microscopic timescale*—At early iterations, the dynamics can be approximated by a free random walk in L dimensions. In this regime, the growth of covariance matrix eigenvalues λ_i has a diffusive scaling behavior $\lambda_i \sim t$.
- (2) *Mesoscopic timescale*—When the motion involves a neighborhood of the starting point, the dynamics can be decomposed into a free random walk along the tangent directions and a bounded fluctuating motion in the orthogonal directions. Upon diagonalization of the covariance matrix, we expect the tangent (orthogonal) $[\parallel (\perp)]$ subspace decomposition $\sigma_n \cong D_{\parallel,n} \oplus D_{\perp}$, where $D_{\perp}$ is a diagonal matrix (approximately) independent from n and $D_{\parallel,n}$ is a

diagonal matrix whose eigenvalues exhibits diffusive scaling behavior.

- (3) *Macroscopic timescale*—Asymptotically, the dynamics involves the whole optimal level set. In this regime, covariance matrix eigenvectors specify the principle axis decomposition (of the Gaussian approximation) of the distribution of optimal protocols, while the eigenvalues specify the lengths of each semiaxis.

Global properties of the optimal level set can be extracted by letting LMC run on a macroscopic timescale. LMC works at finite L and β , so we eventually need to study how results scale as $L, \beta \rightarrow \infty$. To this end, we run LMC at fixed L and β and let the covariance matrix eigenvalues stabilize to constant values. In Fig. 20, the asymptotic eigenvalues of the covariance matrix are shown for $\beta = 10^4, 10^5, 10^6$. The asymptotic independence of $(L - n_+) > 0$ eigenvalues from β for $T > T_{\text{QSL}}$ gives a direct numerical confirmation of the presence of an optimal level set of continuously connected optimal protocols. In particular, we see that $n_+ = 2, 4$ in the single- and two-qubit case, respectively. We notice that the eigenvalue λ_2 (λ_4) in the single- (two-) qubit case shows a weak dependence with β [cf. Figs. 20(c) and 20(d)]: This is




FIG. 20. LMC diffusion stage, covariance-matrix asymptotic eigenvalues for a single LMC run. The eigenvalues capture fluctuations in the protocol dynamics and depend, in general, on β , the Monte Carlo effective inverse temperature in Eq. (28). For $T > T_{\text{QSL}}$ [(c) and (d)], protocol fluctuations along the optimal level set become independent from β ($T_{\text{QSL}} \simeq 2.51$ and 2.98 in single- and two-qubit problems, respectively). Notice the different number n_+ of β -dependent eigenvalues for $T > T_{\text{QSL}}$ in the two problems: the optimal level set expands in $L - 2$ ($L - 4$) directions of the L -dimensional protocol space in the single-qubit (two-qubit) problem.

compatible with the observed magnitude of the smallest positive eigenvalue of the Hessian, visible in Figs. 6(b) and 17(b). Notice that the presence of n_+ β -dependent eigenvalues means that along these directions (in the L -dimensional space) the protocol fluctuations still depend on β . Geometrically, this implies that the optimal level set lies in an n_+ -dimensional hyperplane of the L -dimensional space.

From the LMC diffusive stage analysis, we estimate the number of LMC iterations Δn_{dec} required during diffusion to explore a global portion of the optimal level set. In particular, we define Δn_{dec} as the number of iterations required for half of covariance matrix eigenvalues to stabilize to a constant value. Hence, the set of protocols $\mathcal{E} = \{\mathbf{s}_{n_1} \dots \mathbf{s}_{n_M}\}$ sampled during the LMC equilibrium sampling stage are separated by Δn_{dec} iterations (namely, $n_{i+1} - n_i = \Delta n_{\text{dec}}$). Based on our observations, we set $\Delta n_{\text{dec}} = 2^{12}$ in the single-qubit problem and $\Delta n_{\text{dec}} = 2^{14}$ in the two-qubit problem.

3. Equilibrium sampling stage

Eventually, in the LMC equilibrium sampling stage, we collect r sets of protocols $\mathcal{E}_a$, $a = 1, 2, \dots, r$ from r independent LMC runs.

LMC allows one to estimate the number of connected components in the optimal level set. To this end, we consider the Euclidean distance between two protocols defined by

$$d(\mathbf{s}_1, \mathbf{s}_2) = \sqrt{\frac{1}{T} \int_0^T dt [\mathbf{s}_1(t) - \mathbf{s}_2(t)]^2}.$$

The distance $d(\mathcal{E}_a, \mathcal{E}_b)$ between sets $\mathcal{E}_a$ and $\mathcal{E}_b$ is defined as the minimum distance found between any two of their protocols:

$$d(\mathcal{E}_a, \mathcal{E}_b) = \min_{\mathbf{s} \in \mathcal{E}_a, \mathbf{s}' \in \mathcal{E}_b} d(\mathbf{s}, \mathbf{s}'). \quad (\text{D6})$$

The presence of multiple disconnected components can be detected from the (asymptotic) distribution of distances between pairs of LMC runs. If only one connected component is present, we expect to observe a single-peaked distribution of distances with a peak location that shifts to zero as a larger number of protocols are considered within each set. Otherwise, when multiple disconnected components are present, the distribution of distances is expected to possess multiple peaks, while its average value converges to a nonzero value as more protocols are considered. Here, we assume a single LMC run to sample from a single connected component via (approximately) homotopic transformations, for sufficiently large β [cf. Eq. (27)]. For the single-qubit system, we show how the mean value $\langle d(\mathcal{E}_a, \mathcal{E}_b) \rangle_{a,b}$ scales as the number of protocol samples is increased in Fig. 21. Across the T_{QSL} transition, the minimum distance curves show a similar behavior. For




FIG. 21. Single-qubit problem, LMC equilibrium sampling stage. The average distance between pairs of LMC runs (out of a total of 20) is shown as a function of $|\mathcal{E}_a|^{-1} = |\mathcal{E}_b|^{-1}$, namely, the inverse number of protocols contained in each run. The same slope at the two sides of the $T_{\text{QSL}} \simeq 2.51$ transition [compare (a) with (c) and (b) with (d)] suggests the presence of a single connected component in the optimal level set for $T \simeq T_{\text{QSL}}$.

$T > T_{\text{QSL}}$, distances are independent from β and decrease with the same slope as in $T < T_{\text{QSL}}$, thus suggesting the presence of a single connected component in the optimal level set. The situation in the two-qubit problem is richer, and it is discussed in Ref. [41]. In general, we expect different quantum control problems to have different numbers of connected components.

On a different note, in Fig. 22, we report the average protocols obtained from LMC equilibrium sampling at different Monte Carlo inverse temperature β in the single-qubit problem. The average shapes of the optimal protocols collected from LMC match the average optimal protocols obtained with SD (cf. Fig. 15) in the exact infidelity landscape. Notice the protocol fluctuations, represented by the colored area in the plot, affected by β only for $T = 2.50 < T_{\text{QSL}}$.

APPENDIX E: OPTIMAL LEVEL SET PARAMETRIZATION

In this section, we test the analytical parametrization proposed in Sec. VI for the optimal level set after the quantum speed limit T_{QSL} . In particular, we consider the single-qubit problem defined in Sec. II, which possesses a single optimal control protocol $s_{\Delta_0(T)}$ up to the quantum speed limit [cf. Eq. (24)].




FIG. 22. Single-qubit problem. Average protocol resulting from a single LMC run during the equilibrium sampling stage. The colored area shows fluctuations generated by the LMC dynamics in the optimal level set. Notice that the size of fluctuations depends on β only for $T < T_{\text{QSL}} \simeq 2.51$ [(a)–(c)]. It is interesting to compare this result with the average protocols in SD in Fig. 15.

We consider the parametrization of the optimal level set obtained from the infidelity expansion centered at $s_{\Delta_0(T)}$ and truncated at second order. We explicitly evaluate the infidelity for a continuous family of protocol deformations of the form

$$s_n(x_n) = s_{\Delta_0(T)} + x_n f^{(n)}, \quad (\text{E1})$$

where $n \in \mathbb{N}$ labels orthogonal eigenfunctions and $x_n \in \mathbb{R}$ controls the magnitude of the deformation; $f^{(n)}(t)$ are the eigenfunctions of the Hessian operator in the infidelity expansion centered in $s_{\Delta_0(T)}$ (cf. Fig. 16).

In the analytical parametrization of the optimal level set truncated at second order, along each direction n associated with a negative Hessian eigenvalue $\lambda_n < 0$, the infidelity eventually crosses the zero-infidelity hyperplane. Since fourth-order terms are neglected, we use these intersection points to estimate the locations of global minima in the exact infidelity landscape.

We test the extent to which this approximation holds by computing how the exact infidelity behaves under the deformations in Eq. (E1). At fixed x_n , we numerically

solve the Schrödinger equation associated with the quantum control problem. In Fig. 23, we show the infidelity as a function of x_n for small and large values of the integer n , for fixed $T = 2.52$. At fixed n , the infidelity curve exhibits, in general, two minima $I(x_n^\pm)$, located in the two half-planes $x_n^- < 0, x_n^+ > 0$. We notice that not every minimum is numerically indistinguishable from zero: For example, for $n = 3, 5$ the infidelity at the minima is $\lesssim 10^{-8}$, whereas for $n = 7, 9$ it reaches a value of approximately 10^{-5} . In fact, for odd n , infidelity evaluated at the minima locations increases as n grows. That is, for odd n deformations, s_n are not able to reach zero infidelity, except for $n = 3, 5$. On the other hand, for even n , infidelity evaluated at the minima locations remains small ($\lesssim 10^{-8}$) as n grows. Thus, we observe an important limitation of the infidelity second-order approximation: Some predicted optimal protocols in the approximated landscape (at second order) have non-vanishing infidelity in the exact landscape.

As we saw in Appendix C, parity of n is related to the parity of the $f^{(n)}$ eigenfunction: Odd (even) n corresponds




FIG. 23. Single-qubit problem. Infidelity evaluated for different families of protocols $s_n(x_n)$, $n \in \mathbb{N}$ in Eq. (E1), parametrized by the real variable x_n , at fixed $T = 2.52 > T_{\text{QSL}}$. According to the infidelity expansion truncated at second order, along each direction $n > 2$ the infidelity should intersect the $I(T)[s] = 0$ hyperplane at two values $x_n^+ > 0$ and $x_n^- < 0$. From the plot, we observe that this is not always the case in the exact infidelity landscape: for example, for $n = 6, 7, 9$, $I(x_n^\pm) \gtrsim I(0)$ [(b)] and for odd $n \gg 1$, $x_n^\pm = 0$ [(d)]. We conclude that, in some cases, the infidelity expansion truncated at second order may fail to correctly approximate (even at a qualitative level) the exact infidelity landscape; if so, higher-order terms are important for quantitative comparison.

to an even-parity (odd-parity) eigenfunction with respect to the transformation $f(t) \mapsto -f(T - t)$. Hence, for small n , protocols obtained by deforming $s_{\Delta_0}(t)$ along even-parity $f^{(n)}$ are (within numerical precision) optimal and violate the symmetry of the quantum control problem (cf. Sec. II). Therefore, for $T > T_{\text{QSL}}$, the quantum control symmetry (respected for $T < T_{\text{QSL}}$) is broken in the optimal level set. Interestingly, deformations along eigenfunctions with even parity (i.e., violating the quantum control problem's symmetry) are exactly the ones strongly affected by higher-order corrections in the infidelity expansion.

Let us also comment on the infidelity curve along the $n = 2$ direction, which is highly nonsymmetric with respect to axis $x_2 = 0$ (see Fig. 23). To understand this, remember




FIG. 24. Single-qubit problem. Optimal level set parametrization shortly after the quantum speed limit $T_{\text{QSL}} \simeq 2.51$. The plot shows the separations $|x_n^+ - x_n^-|$ between the two minima x_n^+ and x_n^- in the infidelity curves $I[x_n]$ in Fig. 23, for different directions $n > 2$. Because of the Hessian eigenvalues, negative for $n > 2$, minima are, in general, present pairwise. Here, we compare the separation estimated by the infidelity expansion (truncated at second order; orange points) with the separation resulting from the exact landscape (obtained from numerical integration of the Schrödinger equation; blue points). At $T = 2.52$, we observe good quantitative agreement between the exact and approximated infidelity for $n \lesssim 12$; for $n \gtrsim 12$, the agreement is good at a qualitative level for even n only. Comparing $T = 2.52, 2.60$ [(a) and (b), respectively], we observe that for large n there is a better agreement between the exact infidelity and the second-order expansion results. We relate this behavior to the quantum control symmetry $s(t) \leftrightarrow -s(T - t)$, constraining optimal protocols for $T \leq T_{\text{QSL}}$.

that $n = 2$ is associated with one of the two *positive* eigenvalues visible in the Hessian spectrum (the other being $n = 1$, not shown in the four plots). Hence, the minimum visible in $x_2 \simeq -0.05$ is due to the linear-order term in the infidelity expansion.

As a second comparison between the second-order approximation and the exact infidelity, in Fig. 24, we compare the *location* of minima $x_n^\pm$ in the exact infidelity landscape with the prediction obtained from the second-order expansion of the infidelity, for $T = 2.52$ and $T = 2.60$. For $n \lesssim 12$, there is a good quantitative agreement between the exact (blue points) and approximated locations (orange points). As n increases, higher-order corrections in the infidelity expansion become important, and the two sets of points deviate from each other. In the exact landscape, the location of *odd* n minima eventually moves to zero, for $n \gtrsim 15$: Infidelity does not decrease along these directions, and higher-order terms in the infidelity expansion are important for a qualitative comparison with the exact infidelity landscape. Conversely, the second-order expansion of the infidelity underestimates the location of *even* n minima; in this case, higher-order corrections mildly change the locations of the minima: The second-order expansion and the exact landscape produce the same qualitative result. As a final remark, we notice that, as $T \rightarrow T_{\text{QSL}}$, $|x_n^+|, |x_n^-| \rightarrow 0$ for all *odd* n (corresponding to even-parity eigenfunctions). This is consistent with the quantum control symmetry $f(t) \mapsto -f(T - t)$ respected by the isolated optimal protocol for $T \leq T_{\text{QSL}}$. We deduce that the relatively large contribution of higher-order corrections to the optimal level set parametrization is a direct consequence of the symmetry of the quantum control problem we consider.

APPENDIX F: CALCULATION OF THE ORDER PARAMETER $q(T)$

In this appendix, we provide details of the calculation of the order parameter $q(T)$ discussed in Sec. VI A.

The starting point is the partition function in Eq. (40). Here and in Appendixes G and H, for notational convenience, we redefine coefficients $b_n, \lambda_n, \dots$ appearing in the infidelity expansion in Eq. (31) and absorb the prefactors $T, T^2, \dots$. This operation does not affect the critical scaling of the coefficient as $T \rightarrow T_{\text{QSL}}$ for $\Delta T \rightarrow 0$.

1. The generating functional $G[k]$

Our goal is to evaluate the moments $\langle s_n \rangle$ and $\langle s_n^2 \rangle$, appearing in the order parameter $q(T)$ in Eq. (39).

More in general, the moments $\langle s_n^\alpha \rangle$, $\alpha \in \mathbb{N}$, are conveniently obtained from the generating functional

$$G[k] = Z[e^{\int_0^T dt k_t s_t / T}]$$

via derivatives, $\langle s_n^\alpha \rangle = \partial_{k_n}^\alpha G[k]|_{k=0}/G[0]$.

Therefore, let us evaluate the following quantity:

$$G[k] = \int_{\mathbb{R}} \prod_n ds_n \int_{\mathbb{R}} dz e^{\kappa L F[s] + s \cdot k},$$

$$F[s] = -s^2/2 + izI[s]. \quad (\text{F1})$$

This expression follows from Eq. (40) with the change of variables $s_t = \sum_{n=1}^L s_n f_t^{(n)}$ and the Fourier representation of the Dirac delta distribution:

$$\delta(I) = \frac{\kappa L}{2\pi} \int_{\mathbb{R}} dz \exp[\kappa L (iIz)].$$

In addition, for later convenience, we use the parameter $\kappa \in (0, \infty)$ in place of the numerical prefactor 3 (see Appendix H), and we introduce the shorthand notation

$$s_1 \cdot s_2 = \frac{1}{T} \int_0^T dt s_1(t) s_2(t) = \sum_n s_{1n} s_{2n},$$

$$s^2 = s \cdot s.$$

2. Integration over control protocol variables s_n

As a first step for the explicit evaluation of $G[k]$, we perform the shift of control protocol variables, $s_n \mapsto s_{0n} + \delta s_n$, and let δs_n be fixed by the stationarity condition

$$\left. \frac{\delta F[s_0 + \delta s]}{\delta(\delta s_n)} \right|_{\delta s = \delta \bar{s}} = 0, \quad (\text{F2})$$

to be solved up to a given order in δs . Notice that the solution of Eq. (F2) depends on the Lagrange multiplier z , $\delta \bar{s} = \delta \bar{s}(z)$. This operation centers the functional $F[s]$ at one of its stationary points and leaves us with the functional $\bar{F}[\Delta] = F[s_0 + \delta \bar{s} + \Delta]$, where the field Δ describes fluctuations around the stationary point $s_0 + \delta \bar{s} = \bar{s}_0$.

We notice that, when $F[s]$ is truncated to second order, the stationarity condition has a unique solution (geometrically, the vertex of a paraboloid). Truncating at a higher order may introduce multiple solutions that can be distinguished by the signs of eigenvalues of the associated Hessian operator [78,79]. For the moment, we keep the discussion general and postpone the choice of the order of truncation. Eventually, we see that a second-order truncation is sufficient to understand the critical behavior of the order parameter $q(T)$. In contrast, in Sec. VIB, we show that quartic-order terms are necessary to explain the critical behavior of the order parameter $q_{\text{BB}}(T)$.

After the shift $s_0 \mapsto s_0 + \delta \bar{s}$, we obtain (ignoring the overall prefactor)

$$G[k] \approx \int_{\mathbb{R}} dz e^{-\kappa L \bar{s}_0^2/2 + k \cdot \bar{s}_0} \int_{\mathbb{R}} \prod_n d\Delta_n e^{\kappa L \bar{F}[\Delta]},$$

$$\bar{F}[\Delta] = iz\bar{c} + k \cdot \Delta/(\kappa L) + \Delta \cdot \bar{\Pi}(z) \cdot \Delta/2 + \dots,$$

where the second-order term

$$\Delta \cdot \bar{\Pi}(z) \cdot \Delta = \sum_{n,m} \Delta_n (\delta_{nm} - iz \bar{J}_{nm}) \Delta_m \quad (\text{F3})$$

contains the new Hessian operator

$$\bar{J}_{nm} = \delta_{nm} \lambda_n + \sum_k d_{nmk} \delta \bar{s}_k + \frac{1}{2} \sum_{kl} g_{nmkl} \delta \bar{s}_k \delta \bar{s}_l + \dots, \quad (\text{F4})$$

that includes corrections from higher-order terms $d_{nmk}, g_{nmkl}, \dots$ caused by the shift $\delta \bar{s}$ (analogously for the new constant term $\bar{c}$).

Next, we go to the diagonal basis of the Hessian operator, $\{\tilde{f}_i^{(p)}\}$, with eigenvalues $\{\tilde{\lambda}_p\}$. Neglecting $\mathcal{O}(\Delta^3)$ terms, we perform the Gaussian integrals over the variables $\{\Delta_p\}$ and obtain

$$G[k] \approx \int_{\mathbb{R}} dz e^{\kappa L \Omega_z[k]}, \quad (\text{F5})$$

$$\begin{aligned} \Omega(z)[k] = & -\bar{s}_0^2/2 + iz\bar{c} + \bar{s}_0 \cdot k/(\kappa L) \\ & + (1/2)k \cdot \bar{\Pi}^{-1} \cdot k/(\kappa L)^2 \\ & - (1/2)\text{Tr} \log \bar{\Pi}/(\kappa L). \end{aligned} \quad (\text{F6})$$

In the following, we ignore logarithmic corrections in $G[k]$ as they do not affect $q(T)$.

3. Integration over Lagrange multiplier z

In the $L \rightarrow \infty$ limit, the integral in Eq. (F5) is dominated by the saddle point z_* of the exponent $\Omega(z)[k]$, satisfying

$$\left. \frac{\partial \Omega(z)[k]}{\partial z} \right|_{z=z_*} = 0. \quad (\text{F7})$$

For $z \in \mathbb{R}$, the exponent is a complex number so that valid saddle points must satisfy additional conditions (cf. Appendix G). Although there is no valid saddle point for $z \in \mathbb{R}$, following the steepest-descent method, we deform the integration contour in the complex plane and intersect a complex-valued saddle point, $z_* \in \mathbb{C}$. We notice that each positive (negative) eigenvalue $\tilde{\lambda}_p$ corresponds to a branch point located on the negative (positive) imaginary semiaxis, $z = i\tilde{\lambda}_p^{-1}$. The deformation of the integration contour is valid provided we do not cross any branch point during the process [80].

The saddle-point conditions in Eqs. (F2) and (F7) have to be solved self-consistently to obtain a valid pair of solutions $\delta \bar{s}, z_*$. In general, due to the nonlinearity of the infidelity

$I[s]$, this is a difficult task. In Appendix G, we show that, within a second-order approximation in the infidelity expansion $I[\delta \bar{s}]$, a unique purely imaginary saddle point $z_* = iy_*$, $y_* \in \mathbb{R}$, always exists in the interval

$$\text{Im}(z_*) = y_* \in (-\lambda_+^{-1}, -\lambda_-^{-1}) \quad (\text{F8})$$

for $T \rightarrow T_{\text{QSL}}^+$. Here, we indicate with $\lambda_{+,-}$ the positive and negative eigenvalues among $\{\lambda_p\}$ with largest absolute value (cf. Fig. 25).

In conclusion, following the steepest-descent method, the cumulant-generating functional reads

$$\log G[k] \approx \kappa L \Omega(z_*)[k] \quad (\text{F9})$$

with $\Omega(z)[k]$ defined in Eq. (F6).

From $\log G[k]$, we obtain $q_i(T)$ by deriving two times with respect to the source terms k . This operation is cumbersome, as the quantities $\bar{s}_0$, $\bar{c}$, and $\bar{\Pi}_p$ depend on $\delta \bar{s}_0(y_*)$ and y_* depends on k through the saddle-point condition (F7), $y_* = y_*[k]$. Nevertheless, as shown in the next section, we obtain an informative result by truncating the infidelity expansion in Eq. (31) to second order.




FIG. 25. Schematic representation in the complex plane $z \in \mathbb{C}$ of the simple poles $\{-i\lambda_n^{-1}\}$ and the saddle point z_* of the complex function $\Omega(z)$ defined in Eq. (F6). $\Omega(z)$ appears in the calculation of the generating functional $G[k]$ [cf. Eq. (F5)]. The spectrum $\{\lambda_n\}$ of the second-order term in the infidelity expansion determines the poles' location. The number of positive eigenvalues n_+ depends on the particular controlled quantum system. The green path represents the initial integration path along the real axis appearing in $G[k]$. We perform a steepest-descent approximation by deforming the green into the blue path passing through the saddle point z_* of $\Omega(z)$ in the direction of steepest descent. The red segments represent branch cuts created by the square root appearing in the integrand (F6).

4. The order parameter $q(T)$

In this section, we truncate the infidelity expansion in Eq. (31) to second order and show that the only non-vanishing contribution to $q_i(T)$ in the limit $L \rightarrow \infty$ comes from the second-to-last term in Eq. (F6). In other words, we demonstrate that the dependence $y_* = y_*[k]$ does not affect the behavior of $q(T)$, in the $L \rightarrow \infty$ limit. Consequently, we obtain the result discussed in Sec. VI A:

$$q_i(T) \sim \frac{\text{Tr}(\Pi^{-1})}{\kappa L} = \frac{1}{\kappa L} \sum_n \frac{1}{1 + y_* \lambda_n}, \quad L \rightarrow \infty.$$

The starting point is the cumulant generating function $\log G[k]$ in Eq. (F9). For convenience, we introduce the compact partial derivative notation

$$\partial_n f[k]|_0 \equiv \partial_{k_n} f[k] \Big|_{k=0}.$$

We remember that $\delta\bar{s}$ and y_* are determined by the stationarity conditions in Eqs. (F2) and (F7).

Evaluating $\partial_n^2 \Omega(y_*)[k]|_0$, we obtain the terms

$$\begin{aligned} \partial_n^2 \bar{s}_0(y_*)^2|_0 &= 2[\bar{s}_0 \cdot \partial_n \delta\bar{s} + (\partial_n \delta\bar{s})^2]|_0, \\ \partial_n^2 k \cdot \bar{s}_0|_0 &= 2\partial_n \delta\bar{s}_n|_0, \\ \partial_n^2 y_* \bar{I}|_0 &= [(\partial_n^2 y_*) \bar{I} + 2\partial_n y_* \partial_n \bar{I} + y_* \partial_n^2 \bar{I}]|_0, \end{aligned} \quad (\text{F10})$$

where we use $\bar{s}_0(y_*) = s_0 + \delta\bar{s}(y_*)$ and omit the dependence of $\bar{s}_0$, $\delta\bar{s}$, and $\bar{I}$ on y_* for notational convenience. Let us explicitly write the terms appearing in Eq. (F10). From the infidelity expansion (31), we have

$$\begin{aligned} \partial_n I|_0 &= [b \cdot \partial_n \delta\bar{s} + \delta\bar{s} \cdot \Lambda \cdot \partial_n \delta\bar{s}]|_0, \\ \partial_n^2 I|_0 &= [b \cdot \partial_n^2 \delta\bar{s} + \partial_n \delta\bar{s} \cdot \Lambda \cdot \partial_n \delta\bar{s} + \delta\bar{s} \cdot \Lambda \cdot \partial_n^2 \delta\bar{s}]|_0 \end{aligned}$$

and, from the stationarity condition (F2),

$$\begin{aligned} \partial_n \delta\bar{s}_m|_0 &= (\partial_{y_*} \delta\bar{s}_m)(\partial_n y_*)|_0, \\ \partial_n^2 \delta\bar{s}_m|_0 &= (\partial_{y_*}^2 \delta\bar{s}_m)(\partial_n y_*)|_0^2 + (\partial_{y_*} \delta\bar{s}_m)(\partial_n^2 y_*)|_0 \end{aligned}$$

with

$$\begin{aligned} \partial_{y_*} \delta\bar{s}_m &= (s_{0m} - b_m/\lambda_m) \lambda_m \Pi_m^{-2}, \\ \partial_{y_*}^2 \delta\bar{s}_m &= -2(s_{0m} - b_m/\lambda_m) \lambda_m^2 \Pi_m^{-3}, \\ \Pi_m &= 1 + y_* \lambda_m. \end{aligned}$$

Finally, we obtain $\partial_n y_*|_0$ and $\partial_n^2 y_*|_0$ from the stationarity condition (F7), using the implicit function theorem:

$$\begin{aligned} \partial_n y_*|_0 &= -\frac{F_n}{G} L^{-1}, \\ \partial_n^2 y_*|_0 &= -\frac{A(\partial_n y_*|_0)^2 + B_n(\partial_n y_*|_0) + C_n}{G} L^{-2} \end{aligned}$$

with

$$\begin{aligned} F_n &= \kappa^{-1} \partial_{y_*} \delta\bar{s}_n|_0, & A &= \partial_{y_*}^3 [-\bar{s}_0^2 - y_* \bar{I}]|_0, \\ G &= \partial_{y_*}^2 [-\bar{s}_0^2 - y_* \bar{I}]|_0, & B_n &= \kappa^{-1} \partial_{y_*}^2 \delta\bar{s}_n|_0, \\ & & C_n &= \kappa^{-2} \Pi_n^{-1}. \end{aligned}$$

From this expression, we deduce that terms in Eqs. (F10) are of the order of $\mathcal{O}L^{-2}$. Therefore, their contribution to $q(T)$ is of the form $L^{-1} \sum_{n=1}^{\infty} a_n$, where the numerical sequence $\{a_n\}$ depends on the details of the landscape expansion. A sufficiently fast convergence to zero of the gradient components $\{b_n\}$ and the Hessian eigenvalues $\{\lambda_n\}$ guarantees that the sum $\sum_{n=1}^L a_n$ converges to a finite value, as $L \rightarrow \infty$, and, therefore, does not contribute to $q(T)$.

APPENDIX G: STEEPEST-DESCENT METHOD AND SADDLE POINT OF $\Omega(z)$

In this appendix, we discuss the existence of a saddle point z_* of the complex function $\Omega(z) = \Omega(z)[k]|_{k=0}$ in Eq. (F6), valid for the steepest-descent method applied to Eq. (F5).

We show that, within a second-order truncation of the infidelity expansion, the integration domain $z \in \mathbb{R}$ can always be deformed to pass through a saddle point z_* lying on the imaginary axis, $\text{Re}(z) = 0$.

First, we separate the real and imaginary parts $z = x + iy$, $x, y \in \mathbb{R}$. Notice that $\text{Im} \Omega(z)$ vanishes identically on the imaginary axis, $x = 0$; the function $\Omega(z)$ restricted to the imaginary axis is the real function

$$\Omega(y) = -\frac{1}{2} \sum_n \bar{s}_{0n}^2 - y \bar{c}(y). \quad (\text{G1})$$

Truncating the infidelity expansion at second order, we obtain from Eqs. (31) and (F2)

$$\begin{aligned} \bar{c}(y) &= c + \sum_n b_n \delta\bar{s}_n + \frac{1}{2} \sum_n \lambda_n \delta\bar{s}_n^2, \\ \bar{s}_{0n} &= s_{0n} + \delta\bar{s}_n, \\ \delta\bar{s}_n &= -\frac{s_{0n} + y b_n}{1 + y \lambda_n}. \end{aligned}$$

Substituting in Eq. (G1), we obtain eventually

$$\Omega(y) = -\frac{1}{2} \sum_n s_{0n}^2 - y c + \frac{1}{2} \sum_n \frac{(s_{0n} + y b_n)^2}{1 + y \lambda_n}.$$

From this expression, we deduce that $\Omega(y)$ possesses a minimum in

$$y \in (-\lambda_+^{-1}, -\lambda_-^{-1}) \equiv A,$$

with $\lambda_{+,-}$ the positive and negative eigenvalues of the Hessian operator J_{nm} with largest absolute value. To prove this, it is sufficient to observe that $\Omega(y)$ is continuous for $y \in A$ and

$$\Omega(y) \rightarrow +\infty, \quad y\lambda_{\pm} \rightarrow (-1)^+.$$

Thus, there exists $y_* \in A$ such that y_* is a stationary point (minimum) of the real function $\Omega(y)$. Finally, as $\Omega(z)$ is holomorphic in a neighborhood of $z_* = iy_*$, z_* is also stationary point of the complex function $\Omega(z)$.

To prove that y_* is the unique saddle point of $\Omega(z)$ restricted to $x = 0$, $y \in A$, we observe that the real function $\Omega(y)$ is convex for $y \in A$; i.e., we observe $d_y^2 \Omega(y) > 0$ for $y \in A$.

As a last step, we prove that z_* is a *valid* stationary point for the steepest-descent method [81]. In particular, we prove that the integration path in Eq. (F5) may be deformed to cross the saddle point $z = iy_*$ in the tangent direction $z(x)$, for which (i) the imaginary part is constant and (ii) the second derivative of the real part is negative. The above two properties follow from the holomorphicity of $\Omega(z)$ in z_* . In particular, holomorphicity implies that $\text{Im } \Omega(z)$ is constant around the point y_* along the tangent direction $z(x) = x + iy_*$ and that $d_x^2 \Omega(z_*) < 0$ [81]. This completes the proof that $z_* = iy_*$ is a valid point for the steepest-descent method applied to Eq. (F5).

APPENDIX H: COMPUTING $q_{\text{BB}}(T)$

In this appendix, we compute the critical scaling of $q_{\text{BB}}(T)$, as $\Delta T = T - T_{\text{QSL}} \rightarrow 0^+$, predicted by the partition function in Eq. (42).

Following the same steps as in Appendix F, we arrive at the generating functional

$$\begin{aligned} G_{\text{BB}}[k] &= \sum_i G_{\text{BB}}[s_0^{(i)}; k], \\ G_{\text{BB}}[s_0; k] &\approx \int_{\mathbb{R}} \prod_{n=1}^L ds_n \int_{\mathbb{R}} dz e^{\alpha N F_{\text{BB}}[s_0; s]}, \\ F_{\text{BB}}[s_0; \delta s] &= -s^2/2 + izI[s_{0n} + \delta s] + k \cdot s/(\alpha N), \end{aligned} \quad (\text{H1})$$

where N is the number of bang-bang steps of the “microscopic” protocol and L is the number of steps of the “mesoscopic” protocols, obtained through the coarse-graining average.

Contrary to Sec. VI A, here we do not separate contributions from different locally optimal control protocols $\{s_0^{(i)}\}$. This choice is motivated by our interest in

comparing with numerical results in Sec. II, where the computation of the average $\langle \cdot \rangle$ does not separate contributions from disconnected cluster of solutions.

We notice that $G_{\text{BB}}[s_0; k]$ has the same structure as $G[k]$ in Eq. (F1) with the identification $\kappa = \alpha(N/L)$. Hence, within a second-order approximation, the Shannon entropy originating from the coarse-graining average in Eq. (44) has the same effect as the hard-boundary constraint, $|s_i| \leq 1$.

From Eqs. (5) and (F5), we have

$$q_{\text{BB}}(T) = 1 - \sum_n (\partial_n \log G[0])^2 \quad (\text{H2})$$

and

$$G[k] \approx \sum_i \exp(\kappa L \Omega(y_*)[s_0^{(i)}, k]).$$

Here, we highlighted the dependence $\Omega(y_*) = \Omega(y_*)[s_0^{(i)}, k]$ on the center of expansion $s_0^{(i)}$.

We observe that, in the $\kappa L \rightarrow \infty$ limit, the dominant contributions in

$$\partial_n \log G[0] = \frac{\sum_i \kappa L \partial_n \Omega[s_0^{(i)}, 0] \exp(\kappa L \Omega[s_0^{(i)}, 0])}{\sum_i \exp(\kappa L \Omega[s_0^{(i)}, 0])}$$

arise from protocols in $\{s_0^{(i)}\}$ minimizing the quantity $\Omega[s_0^{(i)}, 0]$. Therefore, denoting this subset with $\{s_*^{(j)}\}$, we are left with the evaluation of the quantity

$$\partial_n \log G[0] = \frac{1}{N_*} \sum_{j=1}^{N_*} \partial_n \Omega[s_*^{(j)}, 0].$$

For the evaluation of $\partial_n \Omega[s_*^{(j)}, 0]$, it is necessary to consider the dependence of the quantities $\bar{s}_*^{(j)}$, $\bar{I}$, and y_* on the source terms k [through the stationarity conditions in Eqs. (F2) and (F7)]. However, using the results from Appendix F 4, we verify that this dependence does not affect the critical scaling of $q_{\text{BB}}(T)$, within a quadratic order truncation of the infidelity. Therefore, in the following discussion, we neglect this dependence.

In this case, the order parameter $q_{\text{BB}}(T)$ solely depends on the terms $k \cdot \bar{s}_*^{(j)}/\kappa L$ with $\bar{s}_*^{(j)} = s_*^{(j)} + \delta \bar{s}_*^{(j)}$. Therefore, we obtain

$$\partial_n \log G[0] \approx \frac{1}{N_*} \sum_j \bar{s}_{*p}^{(j)} = \frac{1}{N_*} \sum_j (s_{*p}^{(j)} + \delta \bar{s}_{*p}^{(j)}),$$

with $\delta \bar{s}_*$ determined by Eq. (F2), and

$$\begin{aligned}
\Delta q_{\text{BB}}(T) &= q_{\text{BB}}(T) - q_{\text{BB},*}(T) \\
&\approx -\frac{1}{N_*^2} \sum_{j,j'} \left(s_*^{(j)} \cdot \delta \bar{s}_*^{(j')} + \delta \bar{s}_*^{(j)} \cdot \delta \bar{s}_*^{(j')} \right), \\
q_{\text{BB},*}(T) &= 1 - \left(\frac{1}{N_*} \sum_j s_*^{(j)} \right)^2.
\end{aligned} \tag{H3}$$

Here, $q_{\text{BB},*}(T)$ is the contribution to $q_{\text{BB}}(T)$ given by the locally optimal protocols $\{s_*^{(j)}\}$, found for $T \leq T_{\text{QSL}}$.

The nonanalytic behavior of $q_{\text{BB}}(T)$ at the quantum speed limit T_{QSL} is contained in the deviation $\Delta q_{\text{BB}}(T) = q_{\text{BB}}(T) - q_{\text{BB},*}(T)$. In particular, we extract its critical behavior for $\Delta T \rightarrow 0$ from the protocol deviation $\delta \bar{s}_*$.

Truncating the infidelity expansion at second order, we obtain

$$\begin{aligned}
\delta \bar{s}_{*n} &= -\frac{s_{*n} + y_* b_n}{1 + y_* \lambda_n} \\
&\sim \begin{cases} b_n & n \leq n_+ \\ s_{*n} & n > n_+ \end{cases} \sim \begin{cases} \Delta T & n \leq n_+ \\ 1 & n > n_+, \end{cases}
\end{aligned}$$

where we used Eqs. (32) and (33). Therefore, within a second-order approximation of the infidelity, we obtain the result in Eq. (46), for which $\Delta q_{\text{BB}}(T)$ is discontinuous at the quantum speed limit.

We observe that the jump discontinuity of $\Delta q_{\text{BB}}(T)$ is an artifact of the second-order truncation of the infidelity expansion. Intuitively, for $\Delta T \rightarrow 0^+$ the “entropic potential” $\int_0^T dt s(t)^2$ dominates over the infidelity constraint, since the second-order term of the infidelity expansion, $\int_0^T dt^2 J_{t_1 t_2} s_{t_1} s_{t_2}$, possesses a number n_+ of vanishing eigenvalues, $\{\lambda_n\}_{n > n_+}$. Therefore, considering terms up to fourth order in the infidelity expansion is crucial for determining the critical behavior of $q_{\text{BB}}(T)$.

Truncating the infidelity expansion at fourth order, we obtain

$$\begin{aligned}
0 &= -s_{*n} - \delta \bar{s}_{*n} \\
&+ y_* \left[b_n + \lambda_n \delta \bar{s}_{*n} + \frac{1}{2} \sum_{m,k} d_{nmk} \delta \bar{s}_{*m} \delta \bar{s}_{*k} \right. \\
&\left. + \frac{1}{3!} \sum_{m,k,l} g_{nmkl} \delta \bar{s}_{*m} \delta \bar{s}_{*k} \delta \bar{s}_{*l} \right].
\end{aligned} \tag{H4}$$

From this equation, we extract the critical scaling in ΔT of $\delta \bar{s}$ by dimensional analysis. In particular, we use the scaling behavior in Eqs. (32) and (33), assume $y_* \sim \Delta T^{-\gamma}$ with $\gamma > 0$ unknown, and let

$$\delta \bar{s}_{*n} \sim \begin{cases} \Delta T^{\zeta_+} & n \leq n_+, \\ \Delta T^{\zeta_-} & n > n_+ \end{cases}$$

with $\zeta_{+,-}$ to be determined consistently with Eq. (H4).

As a working hypothesis, let us first assume $\zeta_- \leq \zeta_+$. In this case, we start the analysis of Eq. (H4) for $n > n_+$, where we obtain the dimensional equality

$$\begin{aligned}
0 &\sim \Delta T^\gamma + \Delta T^{\gamma+\zeta_-} \\
&+ \Delta T^2 + \Delta T^{1+\zeta_-} + \Delta T^{1+2\zeta_-} + \Delta T^{3\zeta_-}.
\end{aligned} \tag{H5}$$

This dimensional equality fixes ζ_- as a function of γ . Namely, for each γ , ζ_- is fixed such that (i) there are at least two terms in the right-hand side with the same dimension (i.e., have the same exponent); (ii) the other terms are subdominant (i.e., have a larger exponent). These conditions are necessary for the leading terms in Eq. (H4) to cancel. Imposing these conditions yields (cf. Fig. 26)

$$\zeta_-(\gamma) = \begin{cases} \gamma/3 & 0 \leq \gamma \leq 3/2, \\ 1/2, \gamma - 1 & 3/2 \leq \gamma \leq 2, \\ 1/2, 1 & 2 \leq \gamma. \end{cases} \tag{H6}$$

Now, considering Eq. (H4) for $n < n_+$ yields

$$\begin{aligned}
0 &\sim \Delta T^\gamma + \Delta T^{\gamma+\zeta_+} \\
&+ \Delta T^1 + \Delta T^{\zeta_+} + \Delta T^{1+2\zeta_+} + \Delta T^{3\zeta_+}.
\end{aligned} \tag{H7}$$

We substitute $\zeta_-(\gamma)$ from Eq. (H6) and obtain

$$\zeta_+(\gamma) = \begin{cases} \gamma & 0 \leq \gamma \leq 1, \\ 1 & 1 \leq \gamma. \end{cases} \tag{H8}$$

These results for $\zeta_-(\gamma)$ and $\zeta_+(\gamma)$ are consistent with the initial assumption $\zeta_- \leq \zeta_+$. On the contrary, we notice that




FIG. 26. Geometrical representation of the dimensional analysis of Eq. (H4), used to extract the critical behavior of the order parameter $q_{\text{BB}}(T)$, for $T \rightarrow T_{\text{QSL}}^+$. (a) Plot of the dimensions of relevant terms in Eq. (H5), as a function of ζ_- and γ . Depending on the value of the parameter γ , the intersections between the different curves change. For each γ , ζ_- is fixed such that (i) at least two curves intersect at ζ_- and (ii) there are no other curves below the intersection point at ζ_- . (b) Plot of $\zeta_- = \zeta_-(\gamma)$ determined from (a) by the conditions (i) and (ii).

repeating the same procedure with the alternative hypothesis $\zeta_- > \zeta_+$ does not yield consistent results.

The allowed analytical range for ζ_- can be further narrowed, as we now show. In particular, we observe that $\zeta_- > 1/2$ is inconsistent with $\gamma > 1$. To see this, we need to remember the constraint $y_* < -\bar{\lambda}_-^{-1}$, with $\{\bar{\lambda}_n\}$ eigenvalues of the Hessian operator $\bar{J}_{nm}$ defined in Eq. (F4) and $\bar{\lambda}_-$ the negative eigenvalue with largest absolute value. Hence, substituting $\delta\bar{s}_n \sim \Delta T^{\zeta_-}$ in Eq. (F4), we can observe that corrections from third- and fourth-order terms,

$$d_{nmk}\delta\bar{s}_k \sim \Delta T^{1+\zeta_-}, \quad g_{nmkl}\delta\bar{s}_k\delta\bar{s}_l \sim \Delta T^{2\zeta_-},$$

are subdominant, when $\zeta_- > 1/2$, with respect to the lowest-order contribution λ_n [cf. Eq. (33)]. As a consequence, when $\zeta_- > 1/2$, eigenvalues $\{\bar{\lambda}_n\}$ of the Hessian operator $\bar{J}_{nm}$ necessarily possess the same scaling as $\{\lambda_n\}$ (i.e., $\lambda_n \sim \Delta T$) and we have $\gamma \leq 1$. In the other case, $\zeta_- \leq 1/2$, fourth-order corrections can change the qualitative behavior of negative eigenvalues, thus allowing for $\gamma \geq 1$.

Finally, substituting the scaling behavior of $\delta\bar{s}_{*n}$ in Eq. (H3), we obtain the result in Eq. (47).

In conclusion, we note that the key difference between the second- and fourth-order results is that there is now a continuous set of potential solutions, $\zeta_- \in [0, 1/2]$. Here, the precise value of ζ_- (and γ) is fixed by solving the stationary conditions in Eqs. (F2) and (F7). On the contrary, within a second-order truncation, we obtain the unique solution $\zeta_- = 0$ (and $\gamma = 1$).

APPENDIX I: SCALING COLLAPSE OF $q_{BB}(T)$ AT THE QUANTUM SPEED LIMIT

In this appendix, we discuss details regarding the scaling collapse shown in Fig. 12. The scaling collapse is performed on the curves $q_{BB}(T)$ shown in Fig. 2, while varying the number of bang-bang steps N in the stochastic descent simulations. We refer to Refs. [82,83] for details regarding the curve-fitting procedure.

An important preliminary step, necessary for the scaling collapse, is the estimation of the numerical behavior of $q_{BB,N}(T)$ in the $N \rightarrow \infty$ limit and for $T \leq T_{QSL}$. The procedure we adopt to estimate the limiting curve $\lim_{N \rightarrow \infty} q_{BB,N}(T \leq T_{QSL})$ consists of two steps. First, for a fixed $T_0 < T_{QSL}$, we interpolate $q_{BB,N}(T_0)$ with the curve $f(N) = AN^{-B} + C$ (A , B , and C : fitting parameters). In the limit $N \rightarrow \infty$, we estimate the limiting value $q_{BB,\infty}(T_0)$ from C . Second, we repeat the first step for different values of $T < T_{QSL}$ and obtain the points $q_{BB,\infty}(T)$. Then, we interpolate the points $q_{BB,\infty}(T)$ with a linear function in T , $q_{BB,0}(T) = q + mT$ (q and m : fitting parameters). Eventually, we perform the scaling collapse on the shifted data $\Delta q_{BB,N}(T) = q_{BB,N}(T) - q_{BB,0}(T)$ and estimate the critical exponent $\Delta q_{BB}(\Delta T) \sim \Delta T^{\zeta_n}$.

The estimate presented in Sec. VIB, $\zeta_n \approx 0.7$, has an uncertainty of the order of approximately 0.1. Here, the main source of error originates in the estimation of $q_{BB,0}(T)$ and, more specifically, in the particular choice of the interval of T used to estimate $q_{BB,0}(T)$. To produce Fig. 12, we extrapolate $q_{BB,0}(T)$ using $T \in [2.3, 2.45]$, $[2.8, 2.95]$ in the single- and two-qubit problem, respectively. Different choice of these intervals leads to slightly different values of the critical exponent within the range $[0.6, 0.8]$, in both problems.

- [1] D. D'Alessandro (2nd ed.) *Introduction to Quantum Control and Dynamics* (Chapman and Hall/CRC, London, 2021).
- [2] H. A. Rabitz, M. M. Hsieh, and C. M. Rosenthal, *Quantum optimally controlled transition landscapes*, *Science* **303**, 1998 (2004).
- [3] U. Boscain, M. Sigalotti, and D. Sugny, *Introduction to the Pontryagin maximum principle for quantum optimal control*, *PRX Quantum* **2**, 030203 (2021).
- [4] Q. Ansel, E. Dionis, F. Arrouas, B. Peaudecerf, S. Guérin, D. Guéry-Odelin, and D. Sugny, *Introduction to theoretical and experimental aspects of quantum optimal control*, *J. Phys. B* **57**, 133001 (2024).
- [5] C. P. Koch *et al.*, *Quantum optimal control in quantum technologies. strategic report on current status, visions and goals for research in Europe*, *Eur. Phys. J. Quantum Technol.* **9**, 19 (2022).
- [6] S. J. Glaser *et al.*, *Training Schrödinger's cat: Quantum optimal control: Strategic report on current status, visions and goals for research in Europe*, *Eur. Phys. J. D* **69**, 1 (2015).
- [7] A. Acín *et al.*, *The quantum technologies roadmap: A European Community view*, *New J. Phys.* **20**, 080201 (2018).
- [8] A. Ungar, P. Cappellaro, A. Cooper, and Won Kyu Calvin Sun, *Control of an environmental spin defect beyond the coherence limit of a central spin*, *PRX Quantum* **5**, 010321 (2024).
- [9] Z. Li, X. Ye, X. Kong, T. Xie, Z. Yang, P. Zhao, Y. Wang, F. Shi, and J. Du, *Zero-field quantum sensing via precise geometric controls for a spin-1 system*, *Phys. Rev. Appl.* **21**, 054011 (2024).
- [10] G. Wang, Y.-X. Liu, J. M. Schloss, S. T. Alsld, D. A. Braje, and P. Cappellaro, *Sensing of arbitrary-frequency fields using a quantum mixer*, *Phys. Rev. X* **12**, 021061 (2022).
- [11] P. Rembold, N. Oshnik, M. M. Müller, S. Montangero, T. Calarco, and E. Neu, *Introduction to quantum optimal control for quantum sensing with nitrogen-vacancy centers in diamond*, *AVS Quantum Sci.* **2**, 024701 (2020).
- [12] C. L. Degen, F. Reinhard, and P. Cappellaro, *Quantum sensing*, *Rev. Mod. Phys.* **89**, 035002 (2017).
- [13] J. E. Gough and V. P. Belavkin, *Quantum control and information processing*, *Quantum Inf. Process.* **12**, 1397 (2013).
- [14] J. M. Gertler, B. Baker, J. Li, S. Shirol, J. Koch, and C. Wang, *Protecting a bosonic qubit with autonomous quantum error correction*, *Nature (London)* **590**, 243 (2021).

- [15] S. Zhou and L. Jiang, *Optimal approximate quantum error correction for quantum metrology*, *Phys. Rev. Res.* **2**, 013235 (2020).
- [16] K. Khodjasteh and L. Viola, *Dynamical quantum error correction of unitary operations with bounded controls*, *Phys. Rev. A* **80**, 032314 (2009).
- [17] C. Ahn, A. C. Doherty, and A. J. Landahl, *Continuous quantum error correction via quantum feedback control*, *Phys. Rev. A* **65**, 042301 (2002).
- [18] C. Altafini and F. Ticozzi, *Modeling and control of quantum systems: An introduction*, *IEEE Trans. Autom. Control* **57**, 1898 (2012).
- [19] C. Brif, R. Chakrabarti, and H. Rabitz, *Control of quantum phenomena: Past, present and future*, *New J. Phys.* **12**, 075008 (2010).
- [20] N. Khaneja, T. Reiss, C. Kehlet, T. Schulte-Herbrüggen, and S. J. Glaser, *Optimal control of coupled spin dynamics: design of NMR pulse sequences by gradient ascent algorithms*, *J. Magn. Reson.* **172**, 296 (2005).
- [21] D. M. Reich, M. Ndong, and C. P. Koch, *Monotonically convergent optimization in quantum control using Krotov's method*, *J. Chem. Phys.* **136**, 104103 (2012).
- [22] T. Caneva, T. Calarco, and S. Montangero, *Chopped random-basis quantum optimization*, *Phys. Rev. A* **84**, 022326 (2011).
- [23] P. Doria, T. Calarco, and S. Montangero, *Optimal control technique for many-body quantum dynamics*, *Phys. Rev. Lett.* **106**, 190501 (2011).
- [24] T. Mahesh, P. Batra, and M. H. Ram, *Quantum optimal control: Practical aspects and diverse methods*, *J. Indian Inst. Sci.* **103**, 591 (2023).
- [25] S. Machnes, U. Sander, S. J. Glaser, P. de Fouquières, A. Gruslys, S. Schirmer, and T. Schulte-Herbrüggen, *Comparing, optimizing, and benchmarking quantum-control algorithms in a unifying programming framework*, *Phys. Rev. A* **84**, 022305 (2011).
- [26] R. Chakrabarti and H. Rabitz, *Quantum control landscapes*, *Int. Rev. Phys. Chem.* **26**, 671 (2007).
- [27] X. Ge, R.-B. Wu, and H. Rabitz, *The optimization landscape of hybrid quantum-classical algorithms: From quantum control to NISQ applications*, *Annu. Rev. Control* **54**, 314 (2022).
- [28] X. Ge, R. Wu, and H. Rabitz, *Optimization landscape of quantum control systems*, *Complex Syst. Model. Simul.* **1**, 77 (2021).
- [29] A. G. R. Day, M. Bukov, P. Weinberg, P. Mehta, and D. Sels, *Glassy phase of optimal quantum control*, *Phys. Rev. Lett.* **122**, 020601 (2019).
- [30] S. van Frank, M. Bonneau, J. Schmiedmayer, S. Hild, C. Gross, M. Cheneau, I. Bloch, T. Pichler, A. Negretti, T. Calarco *et al.*, *Optimal control of complex atomic quantum systems*, *Sci. Rep.* **6**, 34187 (2016).
- [31] M. G. Bason, M. Viteau, N. Malossi, P. Huillery, E. Arimondo, D. Ciampini, R. Fazio, V. Giovannetti, R. Mannella, and O. Morsch, *High-fidelity quantum driving*, *Nat. Phys.* **8**, 147 (2012).
- [32] M. Dalgaard, F. Motzoi, and J. Sherson, *Predicting quantum dynamical cost landscapes with deep learning*, *Phys. Rev. A* **105**, 012402 (2022).
- [33] G. Matos, S. Johri, and Z. Papić, *Quantifying the efficiency of state preparation via quantum variational eigensolvers*, *PRX Quantum* **2**, 010309 (2021).
- [34] T. Caneva, A. Silva, R. Fazio, S. Lloyd, T. Calarco, and S. Montangero, *Complexity of controlling quantum many-body dynamics*, *Phys. Rev. A* **89**, 042322 (2014).
- [35] S. Hernández-Gómez, F. Balducci, G. Fasiolo, P. Cappellaro, N. Fabbri, and A. Scardicchio, *Optimal control of a quantum sensor: A fast algorithm based on an analytic solution*, *SciPost Phys.* **17**, 004 (2024).
- [36] A. P. Solon and J. M. Horowitz, *Phase transition in protocols minimizing work fluctuations*, *Phys. Rev. Lett.* **120**, 180605 (2018).
- [37] G. Folena, A. Manacorda, and F. Zamponi, *Introduction to the dynamics of disordered systems: Equilibrium and gradient descent*, *Physica (Amsterdam)* **631A**, 128152 (2023).
- [38] H. Nishimori, *Statistical Physics of Spin Glasses and Information Processing: An Introduction* (Oxford University Press, New York, 2001).
- [39] M. Bukov, A. G. R. Day, D. Sels, P. Weinberg, A. Polkovnikov, and P. Mehta, *Reinforcement learning in different phases of quantum control*, *Phys. Rev. X* **8**, 031086 (2018).
- [40] M. Bukov, A. G. R. Day, P. Weinberg, A. Polkovnikov, P. Mehta, and D. Sels, *Broken symmetry in a two-qubit quantum control landscape*, *Phys. Rev. A* **97**, 052114 (2018).
- [41] N. Beato, P. Patil, and M. Bukov, *Topological phase transitions in a constrained two-qubit quantum control landscape*, *Phys. Rev. Lett.* **135**, 110803 (2025).
- [42] V. Ros and Y. V. Fyodorov, *The high-d landscapes paradigm: Spin-glasses, and beyond*, *arXiv:2209.07975*.
- [43] Y. V. Fyodorov and R. Tublin, *Optimization landscape in the simplest constrained random least-square problem*, *J. Phys. A* **55**, 244008 (2022).
- [44] D. Gamarnik, C. Moore, and L. Zdeborová, *Disordered systems insights on computational hardness*, *J. Stat. Mech.* (2022) 114015.
- [45] P. Urbani, *A continuous constraint satisfaction problem for the rigidity transition in confluent tissues*, *J. Phys. A* **56**, 115003 (2023).
- [46] L. Zdeborová and F. Krzakala, *Statistical physics of inference: Thresholds and algorithms*, *Adv. Phys.* **65**, 453 (2016).
- [47] D. Hu, P. Ronhovde, and Z. Nussinov, *Stability-to-instability transition in the structure of large-scale networks*, *Phys. Rev. E* **86**, 066106 (2012).
- [48] S. C. Morampudi, B. Hsu, S. L. Sondhi, R. Moessner, and C. R. Laumann, *Clustering in Hilbert space of a quantum optimization problem*, *Phys. Rev. A* **96**, 042303 (2017).
- [49] F. Krzakala and J. Kurchan, *Landscape analysis of constraint satisfaction problems*, *Phys. Rev. E* **76**, 021122 (2007).
- [50] M. Mézard, M. Palassini, and O. Rivoire, *Landscape of solutions in constraint satisfaction problems*, *Phys. Rev. Lett.* **95**, 200202 (2005).
- [51] R. Monasson and R. Zecchina, *Statistical mechanics of the random k -satisfiability model*, *Phys. Rev. E* **56**, 1357 (1997).

- [52] M. Mézard, G. Parisi, and R. Zecchina, *Analytic and algorithmic solution of random satisfiability problems*, *Science* **297**, 812 (2002).
- [53] M. Mezard, G. Parisi, and M. Virasoro, *Spin Glass Theory and Beyond* (World Scientific, Singapore, 1986).
- [54] M. Bukov, D. Sels, and A. Polkovnikov, *Geometric speed limit of accessible many-body state preparation*, *Phys. Rev. X* **9**, 011034 (2019).
- [55] H. Schättler and U. Ledzewicz, *Geometric Optimal Control* (Springer, New York, 2012).
- [56] L. T. Brady, C. L. Baldwin, A. Bapat, Y. Kharkov, and A. V. Gorshkov, *Optimal protocols in quantum annealing and quantum approximate optimization algorithm problems*, *Phys. Rev. Lett.* **126**, 070505 (2021).
- [57] G. C. Hegerfeldt, *Driving at the quantum speed limit: Optimal control of a two-level system*, *Phys. Rev. Lett.* **111**, 260501 (2013).
- [58] M. Bukov, L. D'Alessio, and A. Polkovnikov, *Universal high-frequency behavior of periodically driven systems: From dynamical stabilization to Floquet engineering*, *Adv. Phys.* **64**, 139 (2015).
- [59] W. Magnus, *On the exponential solution of differential equations for a linear operator*, *Commun. Pure Appl. Math.* **7**, 649 (1954).
- [60] S. Blanes, F. Casas, J. Oteo, and J. Ros, *The Magnus expansion and some of its applications*, *Phys. Rep.* **470**, 151 (2009).
- [61] S. Blanes, F. Casas, J. A. Oteo, and J. Ros, *A pedagogical approach to the Magnus expansion*, *Eur. J. Phys.* **31**, 907 (2010).
- [62] Here, we use the shorthand notation $s_t = s(t)$ and $m'_t = m'_s(t)$.
- [63] Observe that, by construction, $J_{tt'}$ is the real and symmetric kernel of a Hilbert-Schmidt integral operator; i.e., it satisfies the following properties: (i) $J_{tt'} \in \mathbb{R}$, (ii) $J_{tt'} = J_{t't}$, and (iii) $\int_0^T dt \int_0^T dt' |J_{tt'}|^2 < \infty$. Then, if we search for eigenfunctions $\{f_t^{(n)}\}$ in the functional space $L^2([0, T])$, the associated eigenvalues $\{\lambda_n\}$ exist well defined, are real, and have a finite range and an accumulation point at zero. In this case, we may estimate the eigenvalues $\{\lambda_n\}$ by first discretizing the time domain $[0, T]$ in L steps and, subsequently, consider the limit $L \rightarrow \infty$.
- [64] The choice of $s(t) \equiv 0$ as a starting protocol is arbitrary but convenient. For other choices $s \neq 0$, finding the optimal protocol requires more attention.
- [65] This argument should not be considered as rigorous. In our case, it is sufficient to guess the structure of the optimal protocol for $T > T_c$. In general, a more quantitative analysis is required to obtain the optimal protocol structure past a given CLPT (see Ref. [35] for an example).
- [66] G. O. Roberts and J. S. Rosenthal, *Optimal scaling of discrete approximations to Langevin diffusions*, *J. R. Stat. Soc. Ser. B* **60**, 255 (2002).
- [67] Note that here β is different from the corresponding hyperparameter of the LMC algorithm in Sec. V.
- [68] More formally, we rewrite $\exp(i\nu_t s_t) \sin(\nu_t)/\nu_t = \exp[i\nu_t s_t + \log(\sin \nu_t/\nu_t)]$ and expand the exponent in the right-hand side at second order in ν_t .
- [69] C. Tsallis, *Possible generalization of Boltzmann-Gibbs statistics*, *J. Stat. Phys.* **52**, 479 (1988).
- [70] M. Orozco-Ruiz, N. H. Le, and F. Mintert, *Quantum control without quantum states*, *PRX Quantum* **5**, 040346 (2024).
- [71] S. H. Sack, R. A. Medina, R. Kueng, and M. Serbyn, *Recursive greedy initialization of the quantum approximate optimization algorithm with guaranteed improvement*, *Phys. Rev. A* **107**, 062404 (2023).
- [72] R. A. Medina and M. Serbyn, *A recursive lower bound on the energy improvement of the quantum approximate optimization algorithm*, [arXiv:2405.10125](https://arxiv.org/abs/2405.10125).
- [73] The quantum speed limit may be infinite, e.g., when a large number of local minima preclude any practical optimization procedure from finding optima with an arbitrarily small infidelity.
- [74] S. Kirchhoff, T. Keßler, P. J. Liebermann, E. Assémat, S. Machnes, F. Motzoi, and F. K. Wilhelm, *Optimized cross-resonance gate for coupled transmon systems*, *Phys. Rev. A* **97**, 042348 (2018).
- [75] E. Farhi, J. Goldstone, and S. Gutmann, *A quantum approximate optimization algorithm*, [arXiv:1411.4028](https://arxiv.org/abs/1411.4028).
- [76] A. Leroise, M. Sonner, and D. A. Abanin, *Influence matrix approach to many-body Floquet dynamics*, *Phys. Rev. X* **11**, 021040 (2021).
- [77] N. Beato, P. Pranay, and M. Bukov, *Toward a theory of phase transitions in quantum control landscapes*, Zenodo (2025), [10.5281/zenodo.17102346](https://zenodo.org/record/17102346).
- [78] J. M. Dominy, T.-S. Ho, and H. A. Rabitz, *Characterization of the critical sets of quantum unitary control landscapes*, *IEEE Trans. Autom. Control* **59**, 2083 (2014).
- [79] L. I. Nicolaescu *et al.*, *An Invitation to Morse theory* (Springer, New York, 2007).
- [80] S. Donaldson, *Riemann Surfaces* (Oxford University Press, New York, 2011).
- [81] J. D. Murray, *Asymptotic Analysis* (Springer Science & Business Media, New York, 2012), Vol. 48.
- [82] O. Melchert, *autoscale.py—A program for automatic finite-size scaling analyses: A user's guide*, [arXiv:0910.5403](https://arxiv.org/abs/0910.5403).
- [83] A. Sorge, *pyfssa 0.7.6* (2015), [10.5281/zenodo.35293](https://zenodo.org/record/35293).