

Ansatz-Free Hamiltonian Learning with Heisenberg-Limited Scaling

Hong-Ye Hu^{1,*†}, Muzhou Ma^{2,*‡}, Weiyuan Gong³, Qi Ye^{4,3}, Yu Tong^{5,6,7},
Steven T. Flammia^{8,9} and Susanne F. Yelin^{1,§}

¹Department of Physics, [Harvard University](#), Cambridge, Massachusetts 02138, USA

²Institute for Quantum Information and Matter, [California Institute of Technology](#), Pasadena, California 91125, USA

³School of Engineering and Applied Sciences, [Harvard University](#), Allston, Massachusetts 02134, USA

⁴Center for Quantum Information, IIIS, [Tsinghua University](#), Beijing, China

⁵Department of Mathematics, [Duke University](#), Durham, North Carolina 27708, USA

⁶Department of Electrical and Computer Engineering, [Duke University](#), Durham, North Carolina 27708, USA

⁷Duke Quantum Center, [Duke University](#), Durham, North Carolina 27701, USA

⁸Department of Computer Science, [Virginia Tech](#), Alexandria, Virginia 22314, USA

⁹Phasecraft Inc., Washington, DC, USA



(Received 31 March 2025; accepted 1 August 2025; published 22 October 2025)

Learning the unknown interactions that govern a quantum system is crucial for quantum information processing, device benchmarking, and quantum sensing. The problem, known as Hamiltonian learning, is well understood under the assumption that interactions are local, but this assumption may not hold for arbitrary Hamiltonians. Previous methods all require high-order inverse polynomial dependency with precision, unable to surpass the standard quantum limit and reach the gold-standard Heisenberg-limited scaling. Whether Heisenberg-limited Hamiltonian learning is possible without prior assumptions about the interaction structures, a challenge we term *ansatz-free Hamiltonian learning*, remains an open question. In this work, we present a quantum algorithm to learn arbitrary sparse Hamiltonians without any structure constraints using only black-box queries of the system's real-time evolution and minimal digital controls to attain Heisenberg-limited scaling in estimation error. Our method is also resilient to state-preparation-and-measurement errors, enhancing its practical feasibility. We numerically demonstrate our ansatz-free protocol for learning physical Hamiltonians and validating analog quantum simulations, benchmarking our performance against the state-of-the-art Heisenberg-limited learning approach. Moreover, we establish a fundamental trade-off between total evolution time and quantum control on learning arbitrary interactions, revealing the intrinsic interplay between controllability and total evolution-time complexity for any learning algorithm. These results pave the way for further exploration into Heisenberg-limited Hamiltonian learning in complex quantum systems under minimal assumptions, potentially enabling new benchmarking and verification protocols.

DOI: [10.1103/j7b8-pb77](https://doi.org/10.1103/j7b8-pb77)

I. INTRODUCTION

Understanding the interactions that govern nature is a central goal in physics. In quantum systems, these

interactions are described by the Hamiltonian, which dictates both the static and dynamic properties of the system. Consequently, given access to a quantum system with an unknown Hamiltonian, a fundamental question arises: what is the most efficient method to learn the interactions of such a system? While this question underpins much of quantum many-body physics, it has become increasingly relevant in practice due to the remarkable progress in quantum science and technology, notably the emergence of programmable analog quantum simulators [1–3] and early fault-tolerant quantum computers [4–8]. These platforms promise to simulate complex quantum phenomena that remain intractable with classical computation. Nevertheless, they also introduce a pressing challenge: how

*Alphabetical ordering and equal contributions.

†Contact author: hongyehu@g.harvard.edu

‡Contact author: muzhouma2002@gmail.com

§Contact author: syelin@g.harvard.edu

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.

to rigorously validate and benchmark such engineered quantum devices [9]. Therefore, Hamiltonian learning is a critical tool to not only probe the unknown interactions but also characterize and control these engineered quantum systems [10–24]. Beyond programmable quantum simulators, Hamiltonian learning also arises in quantum metrology and sensing, where one aims to determine an unknown field (i.e., the Hamiltonian) to precision ϵ at the so-called $\mathcal{O}(1/\epsilon)$ Heisenberg limit. Refining these learning strategies will not only enable the certification of next-generation quantum hardware but also open new avenues in precision sensing and the broader landscape of quantum technologies.

Traditional methods for Hamiltonian learning often rely on preparing either an eigenstate or the thermal (Gibbs) state of the underlying Hamiltonian [25–30]. The coefficients of the unknown Hamiltonian are determined by solving a system of polynomial equations involving the expectation values of numerous Pauli observables. However, the state-preparation step is nontrivial, and these methods are constrained by the so-called standard quantum limit, where achieving a precision ϵ in the learned coefficients requires a total experimental time scaling as $\mathcal{O}(1/\epsilon^2)$.

Recently, inspired by quantum metrology, a new class of Hamiltonian learning algorithms has been proposed that achieves Heisenberg-limited scaling [11–16]. These methods require only simple initial state preparation and black-box queries of the Hamiltonian dynamics. Despite their efficiency, these approaches require a crucial assumption that the interactions are either geometrically local or k local. However, in many scenarios, the exact interaction structure is not known in advance, allowing for potentially arbitrary interactions. Consequently, the search space for the unknown Hamiltonian structure becomes exponentially large, making it challenging to identify interaction terms. Moreover, the possibility of noncommuting terms further complicates the accurate estimation of each coefficient. Efforts to extend these existing methods to arbitrary Hamiltonians have encountered significant obstacles: some approaches demand highly complex quantum controls, such as block encoding and the time-reversal evolutions [31], while others fail to reach the optimal Heisenberg-limited scaling [32]. Therefore, it remains a fundamental open question whether one can achieve Heisenberg-limited Hamiltonian learning with only simple black-box queries to the unitary dynamics and no prior assumptions of the interaction structure—a task we refer to as *ansatz-free Hamiltonian learning*.

In this work, we propose a novel Hamiltonian learning algorithm that overcomes these limitations. Our method achieves Heisenberg-limited scaling for arbitrary Hamiltonians in the following qubit form, including nonlocal ones, with the total experimental time scaling polynomially with the number of Pauli terms in the Hamiltonian. To be more

explicit, any n -qubit Hamiltonian H can be expressed as

$$H = \sum_{s \in \mathcal{S}} \mu_s P_s, \quad (1)$$

where P_s are the n -qubit traceless Pauli operators, $\mathcal{S}$ is the set of Pauli operators that constitute H and μ_s are the unknown coefficients. Since the coefficients μ_s can always be rescaled by the evolution time, we normalize them as $|\mu_s| \leq 1$, following standard practice in the analysis of Hamiltonian learning complexity [11,25]. We use M to denote the number of Pauli terms with nonzero coefficients in H . Importantly, we impose no structural constraints on H : the Pauli terms P_s can be nonlocal and can even have support on up to $\mathcal{O}(n)$ qubits. Moreover, any fermionic Hamiltonian can be mapped into the form of Eq. (1), allowing our results to apply directly to fermionic systems as well. This ansatz-free characterization sets our approach apart from previous works in a fundamental way.

Without prior knowledge of which terms are in the Hamiltonian and what the coefficient values are, there are two specific difficulties for learning such a Hamiltonian: (1) what is the *structure* $\mathcal{S}$, the M Pauli terms, of this Hamiltonian; (2) what are the *coefficients* μ_s with respect to each term in this Hamiltonian. We tackle these difficulties by cycling through an alternating hierarchy of two steps: structure learning and coefficient learning. In the structure learning phase, we identify the Pauli terms with large coefficients by directly sampling a simple quantum circuit, where terms with large coefficients will dominate the outcome. In the coefficient learning phase, we isolate the identified terms by applying ensembles of single-qubit Pauli gates, similar to techniques such as dynamical decoupling [33–36] or Hamiltonian reshaping [11,12]. We then estimate the coefficients through robust frequency estimation [37]. Our algorithm not only achieves Heisenberg-limited scaling in terms of total experimental time but is also resilient to state-preparation-and-measurement (SPAM) errors.

To the best of our knowledge, this is the *first* quantum algorithm capable of learning arbitrary Hamiltonians with Heisenberg-limited scaling using only product state inputs, single-qubit measurements, and black-box access to the Hamiltonian dynamics. This work not only resolves a long-standing theoretical question about Hamiltonian learning but also introduces a practical algorithm with minimal experimental requirements.

II. THE LEARNING PROTOCOL

We consider the task of learning an unknown many-body Hamiltonian H , as given in Eq. (1), through black-box access to its time evolution e^{-iHt} for arbitrary t , and a programmable quantum computer. A natural figure of merit for evaluating the efficiency of any learning protocol

is the total experimental time T . Suppose a protocol uses J experiments to estimate the Hamiltonian coefficients μ to precision ϵ , where the j th experiment involves an evolution of duration t_j and is repeated L_j times. Then, the total experimental time is

$$T = \sum_{j=1}^J L_j t_j. \quad (2)$$

A protocol achieves Heisenberg-limited scaling if $T = \mathcal{O}(1/\epsilon)$, and the standard quantum limit if $T = \mathcal{O}(1/\epsilon^2)$. In what follows, we describe in greater detail how our protocol alternates between structure learning and coefficient estimation in a hierarchical manner, with a total evolution time reaching the Heisenberg-limited scaling.

In the structure-learning step, we identify the dominant interaction terms by determining the support of the coefficient vector $\mu = (\mu_1, \mu_2, \dots, \mu_{4^n-1})^T$. To achieve this, we introduce two approaches: one (denoted as $\mathcal{A}^I$) employs n pairs of two-qubit Bell states shared between the original system and an ancillary system of the same size, while the other (denoted as $\mathcal{A}^{II}$) uses only product state inputs and single-qubit measurements, eliminating the need for ancillary systems at the cost of a moderate increase in M dependence. In the coefficient learning (denoted as $\mathcal{A}^{II}$) step, we estimate the coefficients of Pauli operators identified in the preceding structure-learning step with robust frequency estimation. Specifically, we first determine all the Pauli operators with coefficients $1/2 < |\mu_i| \leq 1$, and learn their coefficients μ_i with $\mathcal{A}^I$ and $\mathcal{A}^{II}$. We then repeat those two steps for smaller coefficient ranges $1/4 < |\mu_i| \leq 1/2$ and so on. In the k th iteration, we learn coefficients that are $(1/2)^k < |\mu_s| \leq (1/2)^{k-1}$, continuing until $k = \lceil \log_2(1/\epsilon) \rceil$, where ϵ is the desired learning precision. This hierarchical learning strategy achieves the gold-standard Heisenberg-limited scaling, requiring a total experimental time having $1/\epsilon$ dependence up to a polylogarithmic factor to reach ϵ -learning accuracy.

In the follows, we use $\tilde{\mathcal{O}}(f)$ to omit $\text{polylog}(f)$ scaling factors. The main results of the hierarchical learning algorithm are summarized as follows.

Result 1 (Informal version of Theorem B1). There exists a quantum algorithm for learning the unknown Hamiltonian H as in Eq. (1) taking n pairs of two-qubit Bell state as input for each experiment instance, querying to real-time evolution of H , and performing Bell-basis measurements that outputs estimation $\hat{\mu}$ such that

$$\|\hat{\mu} - \mu\|_\infty \leq \epsilon \quad (3)$$

with high probability. The total experimental time is

$$T = \tilde{\mathcal{O}}(M^2/\epsilon). \quad (4)$$

This algorithm has trivial classical postprocessing and is robust against SPAM errors.

Result 2 (Informal version of Theorem B2). There exists an ancilla-free quantum algorithm for learning the unknown Hamiltonian H as in Eq. (1) with product-state input, queries to real-time evolution of the Hamiltonian, and single-qubit measurements that outputs $\hat{\mu}$ achieving Eq. (3) with high probability. The total experimental time T is

$$T' = \tilde{\mathcal{O}}(M^3 \log(n)/\epsilon). \quad (5)$$

This algorithm needs classical postprocessing with time

$$T^C = \tilde{\mathcal{O}}(M^5 n \log(n)) \quad (6)$$

and is robust against a restricted class of SPAM errors [38].

In the following, we will outline the proof for both results by introducing the algorithms on structure learning ($\mathcal{A}^I$ and $\mathcal{A}^{II}$) and coefficients learning ($\mathcal{A}^{II}$), and the corresponding proof ideas.

A. Structure-learning algorithm

We first consider the structure-learning algorithm $\mathcal{A}^I$, which takes as input n pairs of Bell states shared between the original system and an n -qubit ancillary register. As a warm-up, consider the case where all unknown Hamiltonian coefficients are bounded away from zero, i.e., $\mu_m := \min \mu_s = \mathcal{O}(1)$. A key challenge in identifying the support $\mathcal{S}$ of the Hamiltonian arises from the fact that the constituent Pauli operators generally do not commute. To address this, we combine the Bell sampling circuit with the coherent evolution driven by $H = \sum_s \mu_s P_s$ for time τ . We prepare n pairs of two-qubit Bell states $|\Phi^+\rangle^{\otimes n} = ((|00\rangle + |11\rangle)/\sqrt{2})^{\otimes n}$ with transversal Hadamard and CNOT gates. The first qubit of each Bell pair undergoes evolution under H for time τ , while the second remains idle. A final Bell-basis measurement is performed by reapplying CNOT and Hadamard gates. The full circuit is illustrated in Fig. 1(a), where the blue box labeled H_* can be ignored for now. This setup enables direct sampling from the distribution $p(s|\mu)$, which is intrinsically linked to the support of μ .

To illustrate the mechanism, consider a single-qubit Hamiltonian $H = \mu_x X + \mu_y Y + \mu_z Z$. Applying H for a short time τ to the first qubit of a Bell pair yields

$$\begin{aligned} & |\Phi^+\rangle - i\tau(\mu_x XI + \mu_y YI + \mu_z ZI) |\Phi^+\rangle + \mathcal{O}(\tau^2) \\ &= |\Phi^+\rangle - i\tau\mu_x |\Psi^+\rangle - i\tau\mu_y |\Psi^-\rangle - i\tau\mu_z |\Phi^-\rangle \\ &+ \mathcal{O}(\tau^2), \end{aligned} \quad (7)$$




FIG. 1. (a) Quantum circuit for the structure-learning subroutine $\mathcal{A}^I$. It prepares n pairs of two-qubit Bell states between the original and ancillary systems via transversal gates. The original system then evolves coherently under the unknown Hamiltonian H and the engineered Hamiltonian H_* , where H_* consists of the large terms in H learned in previous steps with an opposite sign. The combined system is then measured on the n -pair Bell basis. Nontrivial outcomes have probabilities proportional to μ_s^2 , enabling inference of the interaction structure. (b) Quantum circuit for the coefficient-learning subroutine $\mathcal{A}^{II}$. By inserting random Pauli gates from a designed set into the unknown Hamiltonian's evolution, the time evolution of a specific term is approximated, allowing the interaction strength μ_s to be extracted via robust frequency estimation. (c) Combining these subroutines enables hierarchical coefficient estimation, achieving Heisenberg-limited scaling.

where $|\Phi^\pm\rangle$ and $|\Psi^\pm\rangle$ are the four Bell basis states. A Bell-basis measurement then yields outcomes with probabilities proportional to μ_x^2 , μ_y^2 , and μ_z^2 . For example, outcome $(1, 0)$ corresponds to $|\Psi^+\rangle$ and occurs with probability proportional to μ_x^2 , and so on. Thus, the presence of any nontrivial Bell-basis measurement outcome directly implies that the corresponding coefficient is nonzero. This quantum circuit effectively performs importance sampling, using quantum coherence to extract structural information about the Hamiltonian. While it is, in principle, possible to estimate the actual values of μ_x , μ_y , and μ_z by learning the corresponding outcome probabilities, this would require many samples to build up accurate statistics and would not achieve Heisenberg-limited scaling. Instead, the goal of the structure-learning phase is simply to sample each nontrivial outcome *at least once*, thereby identifying the nonzero components of the Hamiltonian. The precise values of the coefficients are then learned in the subsequent coefficient-learning step.

More specifically, by setting $\tau = \mathcal{O}((M\mu_m)^{-1})$, we show the outcome distribution can be lower bounded as $p(s|\mu) > \Omega(\mu_m^4/M^2)$ (see Appendix C for details). Consequently, by applying a union bound, one can ensure that all terms in the support set $\mathcal{S} = \{s : |\mu_s| \geq \mu_m\}$ are sampled at least once with high probability by repeating this quantum circuit $\mathcal{O}(\log(M)/\mu_m^2\tau^2)$ times. The total evolution time under H then scales as $\mathcal{O}(\log(M)/\mu_m^2\tau) = \mathcal{O}(M \log(M)/\mu_m^3)$.

Due to the finite-time evolution, different terms in $\mathcal{S}$ get multiplied together in second and higher orders as a result of Taylor expansion, leading to the possibility of spurious contributions. In rare cases, this can result in a false-positive detection of a Pauli string s with vanishing or

negligible coefficient μ_s . For instance, in the single-qubit example, even if $\mu_x = 0$, a Bell-basis measurement might still yield the readout of state $|\Psi^+\rangle$, arising from higher-order terms like $(\mu_y\tau Y)(\mu_z\tau Z)$ in $\mathcal{O}(\tau^2)$ in Eq. (7), even though they are orders of magnitude smaller. However, such terms are easily ruled out in the subsequent coefficient learning step, which accurately estimates their amplitudes and filters out any below the desired threshold μ_m . These false-positive detections are also illustrated in the Application section, for example in Fig. 2(a). Importantly, the worst-case number of such false-positive events is also upper bounded by the number of samples $\mathcal{O}(M^2 \log(M))$.

To aid readers, we also provide a high-level pseudocode in Algorithm 1. The inputs to this subroutine are (1) black-box access to the time evolution generated by the unknown Hamiltonian H ; (2) an approximately learned partial Hamiltonian $\hat{H}$; (3) the coefficient threshold μ_m ; (4) the number of terms in the Hamiltonian M ; and (5) the number of measurements $\mathfrak{M} = \mathcal{O}(1)$. Notably, the second input $\hat{H}$ plays a central role in the hierarchical learning framework: previously learned components of the Hamiltonian are used to cancel their contribution in future queries to H , as will be detailed in later sections.

It is natural to ask whether the entanglement in the structure learning step is necessary. Surprisingly, we provide a negative answer to this question by showing an alternative algorithm with product state input and single-qubit measurement at the cost of a moderate increase in M dependence. The key observation is that applying random Pauli gates before and after the evolution $e^{-iH\tau}$ will transform it into an effective Pauli channel, a process called Pauli twirling [39–43]. In the single-qubit instance, the




FIG. 2. Ansatz-free Hamiltonian learning with nonlocal and many-body interactions. (a) Learning a 1D disordered XY model with nearest-neighbor interactions ($X_i X_{i+1} + Y_i Y_{i+1}$), supplemented by additional nonlocal crosstalk and an all-to-all coupling. The ansatz-free method (this work) is compared with the vanilla approach based on a nearest-neighbor ansatz [11]. Each run uses 2000 measurements for structure learning and 1000 for coefficient estimation. The inset shows total measurement time versus target precision; a fit yields $T_{\text{total}} \propto \text{precision}^{-0.96}$, consistent with Heisenberg-limited scaling. The red dashed box highlights failure of the vanilla method to recover nonlocal and many-body terms. False positives in structure learning (Pauli strings) are shown; coefficient learning confirms them as negligible. (b) Learning the Hamiltonian of a 1D neutral atom chain with 10 μm spacing, uniform Rabi drive ($\Omega = 1.5/2\pi$ MHz), and global detuning ($\Delta = -4/2\pi$ MHz). Our method identifies the correct power-law decay from dipolar interactions after five rounds of structure and coefficient learning, using the same measurement counts as in (a). In contrast, the vanilla method captures only nearest-neighbor terms. All error bars indicate one standard deviation.

effective channel after applying random Pauli gates is

$$\Lambda_P(\rho) = \mathbb{E}_{\sigma_T \sim \mathbb{P}} \sigma_T^\dagger e^{-iH\tau} \sigma_T \rho \sigma_T^\dagger e^{iH\tau} \sigma_T \\ \approx \rho + \mu_x^2 \tau^2 X \rho X + \mu_y^2 \tau^2 Y \rho Y + \mu_z^2 \tau^2 Z \rho Z, \quad (8)$$

which is a Pauli channel with Pauli error rates $\mu_i^2 \tau^2$. It can be generalized to multiqubit systems by applying random Pauli gates on each qubit. Employing the Pauli

ALGORITHM 1. Structure learning subroutine $\mathcal{A}^I$ (two copy).

Input:

1. Unknown Hamiltonian H with black-box access to its evolution $U_H(t) = e^{-iHt}$
2. Approximately learned part of the Hamiltonian $\hat{H}$
3. Coefficient threshold μ_m
4. An estimated number of terms in the Hamiltonian M
5. Number of measurements $\mathfrak{M}$

Output: Estimated support set $\mathcal{S} = \{P_s : |\mu_s| \geq \mu_m\}$

Pseudo-code:

1. Initialize the set of measurements: $\mathcal{B} \leftarrow \emptyset$
2. Set evolution time $\tau \leftarrow \Theta((M\mu_m)^{-1})$
3. Set Trotterization steps $r_1 \leftarrow \Theta(M^2/\mu_m^2)$
4. **For** $i = 1$ **to** $\mathfrak{M}$ **do**:
 - (a) Prepare N Bell pairs: $|\Phi^+\rangle^{\otimes N}$
 - (b) Simulate evolution: $|\psi\rangle \leftarrow e^{-i(H-\hat{H}) \otimes I^{\otimes N} \cdot \tau} |\Phi^+\rangle^{\otimes N}$ by interleaving black-box queries of $U_H(\tau/r_1)$ with application of $e^{-i\hat{H}\tau/r_1}$ via Trotter-Suzuki decomposition
 - (c) Perform Bell-basis measurement on $|\psi\rangle$
 $\mathcal{B} \leftarrow \mathcal{B} \cup \{b_i\}$ (store measurement outcome)
5. Post-process: Return support set $\mathcal{S}$ corresponding to frequently observed Pauli terms in $\mathcal{B}$

error-rate estimation protocol using product-state inputs and single-qubit measurements [38], we obtain an alternative approach $\mathcal{A}^I$ for structure learning. In Appendix F, we provide the details of this approach.

B. Coefficient-learning algorithm

The coefficient-learning algorithm $\mathcal{A}^I$ takes the structure $\mathcal{S}$ as input and outputs the coefficients $\{\mu_s : s \in \mathcal{S}\}$. The key idea is to isolate each term $\mu_s P_s$ during the evolution and learn each individual coefficient μ_s . The first step can be achieved by *Hamiltonian reshaping* [11,12], which inserts random single-qubit Pauli gates between evolutions of H . Each μ_s can be learned with Heisenberg-limited scaling using *robust frequency estimation* [12,37].

1. Hamiltonian reshaping

The goal of Hamiltonian reshaping is to approximate the time evolution of a specific term or a subset of commuting

Pauli terms in H . We focus on the case where the target Hamiltonian after reshaping is a single traceless Pauli operator P_s for illustration. Define $\mathcal{K}_{P_s}$ as the set containing all Pauli operators that commute with P_s . For any traceless Pauli operator $P_l \neq P_s$, the ensemble average of the transformed operator $P_l^\dagger P_l P$ over set $\mathcal{K}_{P_s}$ will result in zero, since half of the Pauli operators in $\mathcal{K}_{P_s}$ commute with P_l and the other half anticommute with it. On the other hand, for P_s , since every Pauli operator in $\mathcal{K}_{P_s}$ commutes with it, this transformation will leave P_s unchanged. We implement this reshaping strategy by interleaving short-time evolutions under H , denoted $e^{-iH\tau}$, with random conjugations by Pauli operators $Q_k \in \mathcal{K}_{P_s}$. The resulting circuit takes the form:

$$Q_{r_2} e^{-iH\tau} Q_{r_2} \cdots Q_2 e^{-iH\tau} Q_2 Q_1 e^{-iH\tau} Q_1. \quad (9)$$

Averaging over all possible choices of Q_k , this sequence implements a quantum channel that approximates evolution under an effective Hamiltonian, which is P_s . This can be seen by analyzing a single time-step quantum channel:

$$\begin{aligned} \rho &\mapsto \frac{1}{2^{2n-1}} \sum_{Q \in \mathcal{K}_{P_s}} Q e^{-iH\tau} Q \rho Q e^{iH\tau} Q \\ &= \rho - i\tau [H_{\text{eff}}, \rho] + \mathcal{O}(\tau^2), \end{aligned} \quad (10)$$

where the effective Hamiltonian is

$$H_{\text{eff}} = \frac{1}{2^{2n-1}} \sum_{Q \in \mathcal{K}_{P_s}} Q H Q = P_s. \quad (11)$$

In the context of Hamiltonian learning, this procedure allows us to isolate the contribution of a single Pauli term P_s , as identified during the structure-learning phase. By applying this randomized reshaping channel over r_2 short-time steps, we approximate the desired evolution $e^{-i\mu_s P_s t}$, enabling accurate estimation of the coefficient μ_s . This approach is closely related to randomized Hamiltonian simulation techniques such as qDRIFT [44,45], and achieves diamond-norm error scaling of $\mathcal{O}(M^2 t^2 / r_2)$, where M is the total number of Pauli terms in H (see Appendix D 1 for details).

2. Robust frequency estimation

With a good approximation of the time evolution channel under a single-term Hamiltonian $H_{\text{eff}} = \mu_s P_s$ by Hamiltonian reshaping, we can robustly estimate μ_s to an accuracy ϵ with Heisenberg-limited scaling using the robust frequency estimation protocol. This protocol follows the same idea as the robust phase estimation protocol [37], but is modified to use only product-state input and to a good confidence interval rather than minimizing the mean-squared error.

ALGORITHM 2. Coefficient learning subroutine $\mathcal{A}''$.

Input:

1. Unknown Hamiltonian H with black-box access to its evolution $U_H(t) = e^{-iHt}$
2. Desired Pauli operator P_s
3. Desired learning accuracy ϵ for coefficient μ_s
4. An estimated number of terms in the Hamiltonian M

Output: Estimated coefficient $\hat{\mu}_s$ which satisfies $|\hat{\mu}_s - \mu_s| \leq \epsilon$

Pseudo-code:

Set $a = -\pi$ and $b = \pi$

For $l = 1$ **to** $\log_{3/2}(2\pi/\epsilon)$ **do**:

1. Prepare $|\psi_+\rangle$ and $|\psi_-\rangle$ state in parallel. (see Appendix D 3 for definitions of $|\psi_+\rangle, |\psi_-\rangle, O_+, O_-$).
2. Set evolution time $\tau \leftarrow \frac{\pi}{2(b-a)}$
3. Set Trotterization steps $r_2 \leftarrow \Theta(M^2 \tau^2)$
4. Apply evolution with Hamiltonian reshaping:
 $|\psi'_+\rangle \leftarrow \prod_{i=1}^{r_2} (Q_i e^{-iH\tau/r_2} Q_i) |\psi_+\rangle$
 $|\psi'_-\rangle \leftarrow \prod_{i=1}^{r_2} (Q_i e^{-iH\tau/r_2} Q_i) |\psi_-\rangle$
 where $Q_i \in \mathcal{K}_{P_s}$ are sampled uniformly.
5. Measure O_+ on $|\psi'_+\rangle$ and O_- on $|\psi'_-\rangle$ each $\mathfrak{M}$ times and forms an estimate $\hat{S} = \langle O_+ \rangle_{\psi'_+} + i \langle O_- \rangle_{\psi'_-}$ of $\exp\left(-i \frac{2\mu_s \pi}{b-a}\right)$.
6. If $\text{Im}\left(e^{-i \frac{(a+b)\pi}{2(b-a)}} \hat{S}\right) \leq 0$, then $b \leftarrow (a+2b)/3$; else $a \leftarrow (2a+b)/3$

Return $\hat{\mu}_s = (b-a)/4$

Consider a simple example with $P_s = IZZX$ and unknown μ_s . We design the following two experiments: (1) the (+) experiment with input state $|\psi_+\rangle = |0\rangle \otimes |+\rangle \otimes |0\rangle \otimes |+\rangle$ and measuring observable $O_+ = IXZX$ after evolution time t ; (2) the (−) experiment with input state $|\psi_-\rangle = |0\rangle \otimes |-\rangle \otimes |0\rangle \otimes |+\rangle$ and measuring observable $O_- = IYZX$ after evolution time t . It can be easily verified that the expectation values are $\langle O_+ \rangle = \cos(2\mu_s t)$ and $\langle O_- \rangle = \sin(2\mu_s t)$, which together form a simple oscillation $e^{i2\mu_s t}$. The idea of robust frequency estimation is to narrow down the frequency θ of an oscillation signal $e^{i\theta t}$ in a hierarchical manner. Given $\theta \in [a, b]$, one can correctly distinguish whether θ is in $\theta \in [a, (a+2b)/3]$ or in $\theta \in [(2a+b)/3, b]$ by finite sampling in $t = \pi/(b-a)$. This process is then repeated $\mathcal{O}(\log_{3/2}((b-a)/\epsilon))$ times with the range divided by 2/3 each time. The choice of the overlapped region makes

this method robust against errors. In Appendix D 2, we describe the setup of the robust frequency estimation for arbitrary Pauli operator and show it is robust against finite-sample error and channel approximation error from the Hamiltonian reshaping.

We provide a high-level pseudocode in Algorithm 2. The inputs to this subroutine are (1) black-box access to the time evolution generated by the unknown Hamiltonian H ; (2) the desired Pauli term P_s for coefficient estimation; (3) the desired learning accuracy ϵ for coefficient μ_s ; and (4) the number of terms in the Hamiltonian M .

C. Hierarchical learning

As we have described above, $\mathcal{A}^I$ (or $\mathcal{A}^{II}$) can learn the structure of the Hamiltonian and $\mathcal{A}^{II}$ can estimate the coefficients. However, directly concatenating $\mathcal{A}^I$ (or $\mathcal{A}^{II}$) with $\mathcal{A}^{II}$ will give an algorithm with a complexity of $\tilde{O}(1/\epsilon^3)$ dependence on the target accuracy ϵ , not enough for achieving Heisenberg-limited scaling. We address this issue by introducing a hierarchical learning protocol, where we divide the terms in H into $J = \lceil \log_2(1/\epsilon) \rceil$ levels. For the j th level, we learn all terms with coefficient $2^{-(j+1)} < |\mu_s| \leq 2^{-j}$ to ϵ accuracy by applying $\mathcal{A}^I$ (or $\mathcal{A}^{II}$) and $\mathcal{A}^{II}$. Crucially, since we already have an estimation of coefficients for terms with $2^{-j} \leq \mu_s$, we can approximately cancel out those terms by interfacing the black-box evolution with the quantum computer in Hamiltonian simulation.

More concretely, let $\hat{H}_j$ denote the partial Hamiltonian composed of terms with $|\mu_s| > 2^{-j}$ that have been learned in previous rounds. Using this classical description, we can approximate the evolution under the residual Hamiltonian $H - \hat{H}_j$ by interleaving the black-box evolution $e^{-iH\tau}$ with the engineered evolution $e^{i\hat{H}_j\tau}$ implemented on a quantum computer. Specifically, we construct the sequence $(e^{-iHt/r} e^{i\hat{H}_j t/r})^r$, which approximates $e^{-i(H-\hat{H}_j)t}$ via the Lie product formula. This procedure effectively isolates the remaining unlearned terms with $|\mu_s| \leq 2^{-j}$, without requiring access to the time-reversed evolution e^{+iHt} .

With the approximated time-evolution channel with all learned large terms cancelled out, we then rescale the Hamiltonian by extending the evolution time from t to $t/2^{-j}$, thereby boosting the initially small coefficients in j th level to a constant scale. The steps of this procedure are depicted in Fig. 1(c). Both the term cancellation and rescaling can be achieved using Trotterization (see Appendix C for detailed error analysis). We also provide a high-level pseudocode in Algorithm 3.

D. Total time-complexity analysis

We now analyze the scaling of the total experimental time. As defined in Eq. (2), this quantity is given by the

ALGORITHM 3. Ansatz-free Hamiltonian learning algorithm.

Input:

1. Unknown Hamiltonian H with black-box access to its evolution $U_H(t) = e^{-iHt}$
2. Desired learning accuracy ϵ for coefficients
3. An estimated number of terms in the Hamiltonian M
4. Constant number of measurements $\mathfrak{M}$

Output: Estimated coefficient $\hat{\mu}_s$ which satisfies $|\hat{\mu}_s - \mu_s| \leq \epsilon$

Pseudo-code:

Set $\hat{H} = 0$,

For $j = 0$ **to** $\lceil \log_2(1/\epsilon) \rceil - 1$ **do**:

1. $\mu_m = 2^{-(j+1)}$
2. $\mathcal{S}_j = \mathcal{A}^I((H, U_H(t)), \hat{H}, \mu_m, M, \mathfrak{M})$
3. **For** $s \in \mathcal{S}_j$ **do**:
 $\hat{\mu}_s = \mathcal{A}^{II}((H, U_H(t)), P_s, \epsilon, M)$
4. Update $\hat{H} \leftarrow \hat{H} + \sum_{s \in \mathcal{S}_j} \hat{\mu}_s P_s$

Return $\hat{H}$

sum of all evolution times multiplied by the number of experimental repetitions. At each level j , we first perform structure learning to identify the support set

$$\mathfrak{S}_j = \{s : 2^{-(j+1)} < |\mu_s| \leq 2^{-j}\}. \quad (12)$$

In each such experiment, the evolution time scales as $t_{j,l}^{(1)} = \mathcal{O}(2^j/M)$, where the 2^j factor arises from rescaling the Hamiltonian to boost small coefficients. As discussed previously, we require $L = \tilde{O}(M^2)$ repetitions to ensure that each $s \in \mathfrak{S}_j$ is sampled *at least once* with high probability. Thus, the total experimental time spent on structure learning at level j is

$$T_1^{\text{th}} = L t_{j,l}^{(1)} = \tilde{O}(2^j M). \quad (13)$$

Next, we perform coefficient learning using robust frequency estimation to estimate each μ_s for $s \in \mathfrak{S}_j$ to within precision ϵ . Each experiment requires evolution time $t_{j,l}^{(2)} = \tilde{O}(1/\epsilon)$, and the number of repetitions per term is constant due to the robustness of the frequency estimation method, which tolerates a constant additive error. In the worst case, we need to estimate up to $L_j = \mathcal{O}(M^2)$ coefficients at level j , so the total time for coefficient learning at level j is

$$T_2^{\text{th}} = L_j t_{j,l}^{(2)} = \tilde{O}(M^2/\epsilon). \quad (14)$$

Combining the contributions from structure and coefficient learning across all levels $j \in \{0, 1, \dots, \lceil \log_2(1/\epsilon) \rceil - 1\}$,

the total experimental time is

$$\begin{aligned}
 T &= \sum_{j=0}^{\lceil \log_2(1/\epsilon) \rceil - 1} (T_1^{j\text{th}} + T_2^{j\text{th}}) \\
 &= \sum_{j=0}^{\lceil \log_2(1/\epsilon) \rceil - 1} (\tilde{\mathcal{O}}(2^j M) + \tilde{\mathcal{O}}(M^2/\epsilon)) \\
 &= \tilde{\mathcal{O}}(M/\epsilon) + \tilde{\mathcal{O}}(M^2/\epsilon), \tag{15}
 \end{aligned}$$

which confirms that the overall protocol achieves Heisenberg-limited scaling, $T = \tilde{\mathcal{O}}(1/\epsilon)$, in total experimental time. In addition, we show this protocol is robust against SPAM errors (see Appendix C) and the time complexity with the alternative approach is provided in Appendix F.

III. APPLICATIONS

To demonstrate the effectiveness of our method, we consider three representative Hamiltonian learning tasks and benchmark our performance against the state-of-the-art Heisenberg-limited learning approach introduced in Ref. [11]. Our ansatz-free method achieves the same gold-standard Heisenberg scaling—where the total experimental time T_{total} scales inversely with the target precision ϵ as $T_{\text{total}} \sim \epsilon^{-1}$ —while offering significantly broader applicability. In particular, our method accurately recovers *non-local* and *multi-qubit* interactions without relying on a predefined ansatz. This is especially important in scenarios where the structure of the effective Hamiltonian is unknown or difficult to model, making ansatz-free learning a critical tool for validating analog pulse design and analog quantum simulation protocols.

Our first example involves a disordered XY Hamiltonian with nearest-neighbor interactions, augmented by additional nonlocal couplings (mimicking crosstalk) and an all-to-all interaction term. We applied both our ansatz-free method and the state-of-the-art Heisenberg-limited learning approach [11] (referred to as the vanilla method) to learn the Hamiltonian coefficients. The results are shown in Fig. 2(a) as a bar plot. Due to the geometric locality ansatz, the vanilla method fails to capture the nonlocal and many-body terms, effectively zeroing them out during its divide-and-conquer procedure. This is highlighted by the missing yellow bars within the red dashed box. In contrast, our ansatz-free method accurately recovers all significant interactions, including the nonlocal and multiqubit terms, with just one iteration of structure and coefficient learning. The inset of Fig. 2(a) shows the total experimental time versus target precision ϵ for learning one Pauli operator. An empirical fit yields $T_{\text{total}} \sim \epsilon^{-0.96}$, in close agreement with the theoretically optimal Heisenberg-limited scaling as we proved in the previous sections.

In many natural quantum systems, interactions often exhibit long-range tails, and uncovering the full interaction profile—free from ansatz-induced bias—is essential for revealing the underlying physics. A prominent example arises in neutral atoms or polar molecules trapped in optical tweezers, where atoms in highly excited Rydberg states interact via long-range forces. If one encodes a qubit in the ground and the highly excited states also known as Rydberg states, the dominant interaction between nearby atoms resembles an Ising-type $Z_i Z_j$ coupling. However, the interaction is not strictly local—it decays algebraically with distance, and the precise power-law exponent depends on the physical origin, such as van der Waals ($1/r^6$) or dipole-dipole ($1/r^3$) interactions. Capturing these subtle but physically significant long-range features requires an ansatz-free Hamiltonian learning protocol that does not assume a structure of the interactions. Moreover, to resolve small-magnitude couplings efficiently, such a method must attain Heisenberg-limited scaling, ensuring maximal learning efficiency and minimal experimental time within the fundamental limits set by quantum mechanics.

To illustrate this, we simulate a physical Hamiltonian for a chain of five neutral atoms driven by a uniform Rabi frequency Ω and a global detuning Δ . The system is described by the Hamiltonian

$$\begin{aligned}
 H/\hbar &= \frac{\Omega}{2} \sum_l (|g_l\rangle\langle r_l| + |r_l\rangle\langle g_l|) - \Delta \sum_l |r_l\rangle\langle r_l| \\
 &\quad + \sum_{j < l} V_{jl} |r_j\rangle\langle r_j| \otimes |r_l\rangle\langle r_l|, \tag{16}
 \end{aligned}$$

where $V_{jl} = C/|\vec{x}_j - \vec{x}_l|^6$ describes the van der Waals interaction between atoms j and l . We use $C = 862\,690 \times 2\pi$ MHz μm^6 , corresponding to the interaction strength for Rydberg states $70S_{1/2}$ of ^{87}Rb atoms. The atoms are equally spaced with a separation of $10\,\mu\text{m}$, and the drive parameters are set to $\Omega = 1.5/2\pi$ MHz and $\Delta = -4/2\pi$ MHz. Due to the $1/r^6$ decay of the van der Waals interaction, the interaction strength rapidly diminishes with distance.

As shown in Fig. 2(b), the vanilla learning method successfully captures the dominant nearest-neighbor interactions (yellow bars) but fails to identify the weaker long-range components, as evident in the left panel. In contrast, our ansatz-free method accurately reconstructs all significant terms, including the long-range tails, after five iterations of structure and coefficient learning. Remarkably, the recovered interaction strengths closely follow the theoretical $1/r^6$ decay expected from van der Waals interactions. A key advantage of our Heisenberg-limited learning framework is that the number of measurements per iteration remains fixed at 2000 for structure learning and 1000 for coefficient estimation, regardless of the interaction strength. This allows the method to efficiently

resolve weak interactions that are several orders of magnitude smaller than the leading terms. In contrast, learning algorithms that do not achieve Heisenberg-limited scaling must require an increasing number of measurements to resolve smaller couplings, making them substantially less efficient for identifying subtle features in the Hamiltonian. In Appendix H, we provide additional details of the numerical simulations and include plots of all Pauli terms identified at each iteration, including false positives.

Ansatz-free Hamiltonian learning also serves as a powerful tool for validating analog quantum simulations, acting as an unbiased verifier of the effective dynamics [9,20]. A central challenge—and opportunity—in analog quantum simulation lies in designing robust pulse sequences that realize effective Hamiltonians beyond the hardware’s native interactions. This task is closely connected to multi-qubit gate synthesis and analog gadget engineering. To illustrate this application, we consider a time-dependent Hamiltonian implemented in a one-dimensional Rydberg atom array, as used in a recent experiment [46], where the control fields $\Omega(t)$ and $\Delta(t)$ are shown in Fig. 3(a). This pulse sequence is engineered to approximate the effective time evolution of the cluster-Ising model with three-body ZXZ interactions:

$$U = \mathcal{T} \exp \left(-i \int H(t) dt \right) = \exp(-i\theta H_{\text{eff}}), \quad (17)$$

where $H_{\text{eff}} = \sum_i Z_i X_{i+1} Z_{i+2}$ and $\theta = 0.1$. Due to hardware limitations, the resulting evolution achieves a unitary fidelity of approximately 91% in the three-atom system. This motivates the need to independently learn and verify the effective Hamiltonian realized by the pulse.

Using our ansatz-free method, we performed Hamiltonian learning on both three-atom and five-atom chains driven by the same global time-dependent controls, assuming only that the target rotation angle is $\theta \approx 0.1$. As in previous examples, we used 2000 measurements for structure learning and 1000 for coefficient estimation. The results, shown in Figs. 3(b) and 3(c), confirm that all relevant terms in the effective Hamiltonian are accurately recovered. For the three-atom system, the ZXZ interaction emerges as the dominant term, accompanied by smaller perturbative corrections. However, when the same control sequence is applied to a five-atom system, the resulting effective Hamiltonian differs, despite still featuring global ZXZ components. Additional details regarding the numerical simulations are provided in Appendix H.

This observation highlights a crucial insight: ansatz derived from small systems (e.g., three atoms) may not generalize to larger systems. Imposing such an ansatz can introduce significant bias and lead to the omission of important interaction terms. As demonstrated in Fig. 2(a), our ansatz-free approach avoids this pitfall and




FIG. 3. Verifying effective Hamiltonians in engineered analog quantum simulations. (a) A 1D array of three Rydberg atoms spaced by $8.9 \mu\text{m}$ is driven by time-dependent Rabi frequency $\Omega(t)$ and global detuning $\Delta(t)$, engineered to approximate evolution under a symmetry-protected topological (SPT) Hamiltonian with three-body terms $Z_i X_{i+1} Z_{i+2}$, achieving a unitary fidelity of 90%. Verifying the resulting effective Hamiltonian is essential for validating analog quantum simulations and pulse engineering. Applying our ansatz-free learning method, we identify ZXZ as the leading interaction. (b) Extending the same pulse to a five-atom array results in a different effective Hamiltonian. Our method successfully identifies all significant Pauli terms, revealing that the three-atom ansatz fails to generalize. True values are obtained via exact simulation; error bars represent one standard deviation.

enables unbiased reconstruction of complex, engineered Hamiltonians.

IV. TRADE-OFF BETWEEN TOTAL EVOLUTION TIME AND QUANTUM CONTROL

Compared to previous approaches that learn Hamiltonians from eigenstates or Gibbs states, our protocol surpasses the standard quantum limit and achieves Heisenberg-limited scaling. This is accomplished by sequentially querying the real-time dynamics of the system, interleaved with discrete quantum controls implemented via Trotterized quantum evolution on a quantum computer. A natural question arises: are quantum controls or having a quantum computer essential for achieving Heisenberg-limited




FIG. 4. (a) Learning tree representation. Each node u represents an experiment. Starting from the root experiment r , the number of child nodes depends on the possible POVM measurements. The transition probabilities are determined by Born's rule [also see Eq. (G2)]. After N_{exp} experiments, one arrives at the leaves l of the learning tree. (b) In each node u , the learning model prepares an arbitrary state ρ_u , applies discrete control channels C_k^u , and queries real-time Hamiltonian evolutions with time τ_k^u multiple times. The protocols can also incorporate ancilla qubits (quantum memory).

scaling? We answer this affirmatively by proving a trade-off between the total evolution time required to learn an unknown Hamiltonian H and the number of discrete control operations. We emphasize that here, “control” refers specifically to unitary operations we apply during time evolution, not controlled-unitary gates.

Intuitively, we can generalize experiments with discrete quantum controls and sequential queries to the Hamiltonian as the following theoretical model. Given a Hamiltonian H with coefficient vector μ defined in Eq. (1), the protocol performs multiple experiments and measures at the end of each experiment. In each experiment, the protocol prepares an input state (possibly with ancilla qubits) and queries the Hamiltonian multiple times with a discrete quantum control channel between every two neighboring queries, as illustrated in Fig. 4(b).

The entire process, including all experimental branches and outcomes, can be organized into a tree structure—referred to as the learning tree representation, in which each node u [see Fig. 4(a)] corresponds to a particular experiment. The protocol can be adaptive in the sense that it can dynamically decide how to prepare the input state, query the real-time evolutions, perform quantum controls, and measure the final state based on the history of the previous experiments. Each node in the tree is associated with a probability $p^\mu(u)$, with the root node having $p^\mu(r) = 1$. The probability of reaching any other node is determined by the full measurement trajectory leading from the root, in accordance with Born's rule. A formal definition of the learning tree model is provided in Appendix G.

In a recent work [47], it is shown through quantum Fisher information [48–50] that for *unbiased* learning algorithm, Heisenberg-limited scaling is not possible in the absence of quantum control. In contrast, we establish a rigorous and fully general lower bound that applies to any adaptive protocol—whether biased or unbiased—that may employ ancillary quantum memory. Our result leverages tools from quantum Fisher information theory and the

learning tree representation to characterize the fundamental trade-offs in Hamiltonian learning.

Result 3 (Informal version of Theorem G1). For any protocol that can be possibly adaptive, biased, and ancilla assisted, with total evolution time T and at most $\mathcal{L}$ discrete quantum controls per experiment, there exists some Hamiltonian H that requires $T = \Omega(\mathcal{L}^{-1}\epsilon^{-2})$ to estimate the coefficient of the unknown Hamiltonian H within additive error ϵ .

To establish this result, we consider a distinguishing task between two Hamiltonians: the original Hamiltonian $H(\mu)$ and a perturbed version $H(\mu + \delta\mu)$, where the perturbation satisfies $|\delta\mu|_\infty \propto \epsilon$. The objective is to determine which of the two Hamiltonians governs the underlying dynamics. It is straightforward to see that any Hamiltonian learning algorithm capable of estimating μ can also be used to solve this distinguishing problem.

Within the learning tree framework, a necessary condition for successful discrimination between these two hypotheses—according to Le Cam's two-point method [51]—is that the total variation distance between the probability distributions over the leaf nodes of the two learning trees must be large. Quantitatively, this condition requires

$$\frac{1}{2} \sum_{\ell \in \text{leaf}(\mathcal{T})} |p^{\mu+\delta\mu}(\ell) - p^\mu(\ell)| = \Theta(1). \quad (18)$$

Using the martingale likelihood ratio technique developed in a series of recent works [52–54], the total variation lower bound can be translated into a quantitative statement via Lemma 7 of Ref. [55]. Specifically, if there exists a constant $\Delta > 0$ such that

$$\mathbb{E}_{s_u \sim p^\mu(s_u|u)} [(L_{\delta\mu}(u, s_u) - 1)^2] \leq \Delta, \quad (19)$$

where

$$L_{\delta\mu}(u, s_u) = \frac{p^{\mu+\delta\mu}(s_u | u)}{p^{\mu}(s_u | u)}, \quad (20)$$

then the number of experiments must satisfy $N_{\text{exp}} \geq \Omega(1/\Delta)$. The physical intuition behind this result is as follows: when the distributions $p^{\mu}(s_u | u)$ and $p^{\mu+\delta\mu}(s_u | u)$ are nearly indistinguishable at all internal nodes of the learning tree, a deep tree, corresponding to a large number of experimental queries, is required to amplify these differences to a detectable level at the leaves, as captured in Eq. (18). In Appendix G, we show that

$$\mathbb{E}_{s_u \sim p^{\mu}(s_u | u)} [(L_{\delta\mu}(u, s_u) - 1)^2] \leq \max_{\mathbf{v}: \|\mathbf{v}\|_2 = 3\sqrt{M}\epsilon} \mathbf{v}^{\top} I_{\mu,u}^{(Q)} \mathbf{v}, \quad (21)$$

where $I_{\mu,u}^{(Q)}$ is the quantum Fisher information matrix associated with node u . Combining this result with the Lemma 17 from Ref. [47], we establish the claim stated in Result 3.

This result implies that, in general, no protocol can achieve Heisenberg-limited scaling without employing at least $\mathcal{L} = \Omega(1/\epsilon)$ discrete quantum controls per experiment. This highlights a fundamental trade-off between quantum controllability and learning efficiency. However, this trade-off should not be misinterpreted as allowing Hamiltonian learning protocols to surpass the Heisenberg limit by simply increasing the number of discrete controls. In fact, as shown in Refs. [11,12,47], there exists a strict lower bound on the total evolution time, namely $T = \Omega(1/\epsilon)$, that applies regardless of the number of available controls.

V. OUTLOOK

In this work, we have introduced the *first* ansatz-free Hamiltonian learning protocol for an n -qubit system that makes no assumptions about the interaction structure. Our protocol achieves the gold standard of Heisenberg-limited scaling, with total experimental time scaling as $T = \mathcal{O}(1/\epsilon)$, and remains polynomial in the number of relevant interactions. Importantly, the protocol relies solely on black-box access to the time evolution generated by the unknown Hamiltonian, along with elementary digital controls on a quantum computer. We demonstrate numerically that our method can efficiently learn nonlocal and many-body interactions in disordered spin chains and Rydberg atom arrays. Furthermore, we establish a fundamental trade-off between total evolution time and quantum control in general Hamiltonian learning protocols. This result offers an intriguing bridge between quantum learning theory and quantum metrology and may serve as a foundation for future theoretical exploration.

Our work opens several promising directions for future research. First, ansatz-free learning provides a critical tool

for the verification of analog quantum simulators. In quantum many-body physics, objects such as the entanglement Hamiltonian and the entanglement negativity Hamiltonian play a key role in characterizing quantum correlations in both closed and open systems [20,56–59]. While these Hamiltonians are analytically known in only certain solvable models, their structure remains largely unexplored in generic interacting systems. Ansatz-free Hamiltonian learning offers a new tool to empirically probe and verify these objects on programmable quantum devices. In particular, the entanglement negativity Hamiltonian is generally non-Hermitian, motivating the development of new learning protocols tailored to non-Hermitian Hamiltonians [60].

Another important direction is to understand the impact of noise during black-box Hamiltonian queries on learning efficiency. Preliminary analysis suggests the existence of a threshold for learning precision ϵ , such that Heisenberg-limited scaling persists when $\epsilon > \mathcal{O}(\eta)$, where η characterizes the strength of the noise. Below this threshold, the learning performance degrades, and Heisenberg scaling breaks down. Investigating this transition is an exciting avenue for future work. In particular, techniques from quantum metrology may offer valuable insights in this direction [61,62]. Finally, extending ansatz-free learning to the open-system setting is another promising direction. Developing ansatz-free Lindbladian learning protocols would have broad applications in quantum device benchmarking, quantum noise characterization, and open-system dynamics more generally [63,64].

ACKNOWLEDGMENTS

We are thankful for the insightful discussions with Sitan Chen, Friedrich Liyuan Chen, Soonwon Choi, Dong-Ling Deng, Nik O. Gjonbalaj, Yingfei Gu, Hsin-Yuan Huang, Christian Kokail, Yunchao Liu, Francisco Machado, Daniel K. Mark, and Pengfei Zhang. H.Y.H. and S.F.Y. acknowledge the support from DOE through the QUACQ program (DE-SC0025572).

H.Y.H., M.M., and Q.Y. conceived the project during Ma's and Ye's research stay at Harvard. H.Y.H. and M.M. carried out the majority of the calculations. W.G. and Y.T. contributed to the theoretical proofs. S.T.F. and S.F.Y. supervised the project throughout.

DATA AVAILABILITY

No data were created or analyzed in this study.

APPENDIX A: NOTATIONS

In this work, the Pauli matrices are denoted by $\sigma^x, \sigma^y, \sigma^z$. We use the following notation to denote the

Pauli eigenstates:

$$\begin{aligned} |1, z\rangle &= |0\rangle, \quad |-1, z\rangle = |1\rangle, \quad |1, x\rangle = |+\rangle, \\ |-1, x\rangle &= |-\rangle, \\ |1, y\rangle &= \frac{1}{\sqrt{2}}(|0\rangle + i|1\rangle), \quad |-1, y\rangle = \frac{1}{\sqrt{2}}(|0\rangle - i|1\rangle). \end{aligned} \quad (\text{A1})$$

We denote the set of all N -fold tensor products of single-qubit Pauli matrices (and the identity) by $\mathbb{P}_n$:

$$\mathbb{P}_n = \left\{ \bigotimes_{i=1}^n P_i : P_i = I, \sigma^x, \sigma^y, \text{ or } \sigma^z \right\}. \quad (\text{A2})$$

We use the following notation for the four maximally entangled two-qubit states:

$$\begin{aligned} |\Phi^+\rangle &= \frac{|00\rangle + |11\rangle}{\sqrt{2}} \\ |\Phi^-\rangle &= \frac{|00\rangle - |11\rangle}{\sqrt{2}} \\ |\Psi^+\rangle &= \frac{|10\rangle + |01\rangle}{\sqrt{2}} \\ |\Psi^-\rangle &= i \frac{|10\rangle - |01\rangle}{\sqrt{2}}, \end{aligned} \quad (\text{A3})$$

which are also the four eigenstates of the Bell-basis measurement. In the following, we will also use $|\text{EPR}_n\rangle = |\Phi^+\rangle^{\otimes n}$.

We consider the N -qubit traceless Hamiltonian to be learned as

$$H = \sum_s \mu_s P_s, \quad |\mu_s| \leq 1, \quad (\text{A4})$$

where $P_s \in \mathbb{P}_N \setminus \{I^{\otimes N}\}$. Let M be the number of terms in the coefficient vector $\boldsymbol{\mu} = (\mu_1, \mu_2, \dots, \mu_{4^N-1})^T$, which are larger than the threshold constant $\epsilon \in (0, 1)$; this is the *sparsity* of H .

Let $S \subseteq \mathbb{P}_n$ be the set of the indices s of terms P_s in H with coefficients $|\mu_s| > \epsilon$. For $j \in \{0, 1, 2, \dots, \lceil \log_2(1/\epsilon) \rceil\}$, define the set S_j of all indices s of terms P_s in H such that their coefficients μ_s satisfy $|\mu_s| \leq 2^{-j}$:

$$\begin{aligned} S_j &= \{s : s \in S, |\mu_s| \leq 2^{-j}\}, \\ j &\in \{0, 1, 2, \dots, \lceil \log_2(1/\epsilon) \rceil\}, \end{aligned} \quad (\text{A5})$$

note that $S = S_0$. Let the set $\mathfrak{S}_j = S_j \setminus S_{j+1}$ be the set of all indices s of terms P_s in H such that their coefficients μ_s

satisfy $2^{-(j+1)} < |\mu_s| \leq 2^{-j}$:

$$\begin{aligned} \mathfrak{S}_j &= \{s : s \in S, 2^{-(j+1)} < |\mu_s| \leq 2^{-j}\}, \\ j &\in \{0, 1, 2, \dots, \lceil \log_2(1/\epsilon) \rceil - 1\}. \end{aligned} \quad (\text{A6})$$

Let $U_H(t) = e^{-iHt}$ represent the time-evolution operator under H . Let $|\text{EPR}\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$ be the two-qubit EPR state, and let $|\text{EPR}_n\rangle = \bigotimes_{j=1}^n |\text{EPR}\rangle$ be the $2n$ -qubit EPR state.

Consider a general quantum channel Λ . One can write it in its Kraus operator expression as

$$\Lambda(\rho) = \sum_k K_k \rho K_k^\dagger, \quad (\text{A7})$$

where $\sum_k K_k^\dagger K_k = I$.

Expanding the Kraus operator in the Pauli basis, we find

$$K_j = \sum_{\sigma_k \in \{I, \sigma_x, \sigma_y, \sigma_z\}} \alpha_{j,k} \sigma_k. \quad (\text{A8})$$

APPENDIX B: MAIN RESULTS AND PROOF IDEAS

The main result of Hamiltonian learning with ancillary systems can be summarized as follows:

Theorem B1 (Two-copy Heisenberg-limited Hamiltonian learning algorithm). For an arbitrary n -qubit unknown Hamiltonian $H = \sum_{s \in S} \mu_s P_s$, with $|\mu_s| \leq 1$ and $|S| \leq M$, there exists a hierarchical learning quantum algorithm which only queries the black-box forward evolution of H , and a fault-tolerant quantum computer with n ancillary qubits that outputs a classical description $\hat{\mu}_s$ of μ_s such that $|\hat{\mu}_s - \mu_s| < \epsilon$ with probability at least $1 - \delta$, and the total experimental time is

$$T = \mathcal{O} \left(\frac{M^2 \log(M/\delta) [\log(1/\epsilon)]^2}{\epsilon} \right).$$

This algorithm requires no nontrivial classical postprocessing and is robust against SPAM error.

In addition, we proposed a single-copy Hamiltonian learning algorithm without an ancillary system using only single-qubit operations, which can be summarized as follows.

Theorem B2 (Single-copy Heisenberg-limited Hamiltonian learning algorithm). For an arbitrary n -qubit unknown Hamiltonian $H = \sum_{s \in S} \mu_s P_s$, with $|\mu_s| \leq 1$ and $|S| \leq M$, there exists a hierarchical learning quantum algorithm, which only queries the black-box forward evolution of H , and a fault-tolerant quantum computer with no

ancillary qubits that outputs a classical description $\hat{\mu}_s$ of μ_s such that $|\hat{\mu}_s - \mu_s| < \epsilon$ with probability at least $1 - \delta$, and the total experimental time is

$$T = \mathcal{O}\left(\frac{M^3(\log(Mn)\log(1/\epsilon) + \log(M/\delta)[\log(1/\epsilon)]^2)}{\epsilon}\right).$$

This algorithm requires a total classical postprocessing time

$$T^C = \mathcal{O}(M^5 n \log(Mn/\delta) \log(1/\epsilon)).$$

The main subroutine of our learning algorithms is defined as follows:

Definition B1 (Hierarchical learning subroutine). Given (1) an unknown Hamiltonian $H = \sum_{s \in S} \mu_s P_s$, with $|\mu_s| \leq 1$ and $|S| \leq M$, and (2) the classical description of $\hat{H}_j$ (its Pauli operator terms and coefficients), an estimation of all coefficients larger than 2^{-j} up to ϵ term-wise error. There exists a hierarchical learning subroutine $\mathcal{A}_j$ to obtain $\hat{H}_{j+1}$ which is an estimation of all coefficients larger than $2^{-(j+1)}$ up to ϵ term-wise error, for $j = 0, 1, 2, \dots, \lceil \log_2(1/\epsilon) \rceil - 1$. The final estimation $\hat{H}_{\lceil \log_2(1/\epsilon) \rceil - 1}$ will be a characterization of H to ϵ term-wise error.

The main ideas: Based on the definitions provided in the previous section, we partition all unknown coefficients into the following sets: $S = \mathfrak{S}_0 \cup \mathfrak{S}_1 \cup \dots \cup \mathfrak{S}_{\lceil \log_2(1/\epsilon) \rceil - 1}$, where $\mathfrak{S}_j = \{s : s \in S, 2^{-(j+1)} < |\mu_s| \leq 2^{-j}\}$. At each step j , the hierarchical learning subroutine $\mathcal{A}_j$ first identifies the index $s \in \mathfrak{S}_j$ with high probability using the Hamiltonian structure learning algorithm $\mathcal{A}_j^I$. With this information, the corresponding coefficients can then be learned using the Hamiltonian coefficient learning algorithm $\mathcal{A}_j^{II}$. Below, we provide an overview of the key components of both algorithms, with a detailed analysis deferred to the subsequent sections.

In the Hamiltonian structure learning $\mathcal{A}_j^I$, we first create the EPR state between the unknown system with N -qubit ancillary system using transversal gates. Then we perform the time evolution of the unknown system under $\tilde{H}_j = (H - \hat{H}_j)/2^{-j}$ for time t by querying unknown system and a fault-tolerant quantum computer iteratively. Lastly, we perform the Bell-basis measurement on the $2N$ -qubit system. We prove that if $t = \mathcal{O}(1/CM)$ where C is a big constant to suppress the error, then all elements $s \in \mathfrak{S}_j$ will be sampled once with high probability. There are several approximation errors involved in this step: (1) Trotterization error ($\epsilon_{\text{Trotter}}$) for simulating time evolution under $\tilde{H}_j$, (2) truncation error (ϵ_{Taylor}) from the Taylor expansion of time $t = \mathcal{O}(1/CM)$ evolution, and SPAM errors (ϵ_{SPAM}). In [lem:lower bound for prob dist]Lemma 3, we prove that for elements in $\mathfrak{S}_j$, their probability is lower

bounded by $\tilde{\Omega}(1/4C^2M^2 - 1/C^3M^2 - 2^{2j}/C^2r - \epsilon_{\text{SPAM}})$, where r is the Trotter steps. Then one can sample all elements in $\mathfrak{S}_j$ at least once with probability $1 - \delta$ by querying this algorithm $\mathcal{O}(C^2M^2 \log(M/\delta))$ times. Then the total evolution time of all $\mathcal{A}_j^I$ is $\mathcal{O}(M \log(M/\delta)/\epsilon)$, which reaches the Heisenberg-limited scaling.

Then we use the Hamiltonian coefficient learning $\mathcal{A}_j^{II}$ to learn the coefficients μ_s for s in the set $\mathfrak{S}_j$, which we identified in the previous step. There are at most $|\mathfrak{S}_j| \equiv L_j = \mathcal{O}(M^2 \log(M/\delta))$ terms sampled in the previous step. We first use the Hamiltonian reshaping technique as introduced in Appendix D1 to approximate the time evolution under a single term P_s for $s \in \mathfrak{S}_j$. We show the approximation error is bounded by $\mathcal{O}(M^2 t^2 / r_2)$ for total evolution time t with r_2 Trotter steps. Then we adapt the robust frequency estimation (RFE) as introduced in Appendix D2 to estimate μ_s within ϵ accuracy. In Theorem D2, we show the total evolution time in each RFE procedure is $\mathcal{O}((\log(1/\epsilon) + \log \log(2^{-j}/\epsilon))/\epsilon)$. As there are at most L_j terms, the total evolution time for $\mathcal{A}_j^{II}$ is $\mathcal{O}((M^2 \log(M/\delta)(\log(1/\epsilon) + \log \log(2^{-j}/\epsilon)))/\epsilon)$, which also reaches the Heisenberg-limited scaling. Moreover, since we estimate each coefficient at a time, the effective Hamiltonian H_{eff} after reshaping is a single term Hamiltonian, we only need product state input in RFE instead of highly entangled Bell state input as in Ref. [12].

In Appendix F, we investigate whether quantum entanglement of the two copy in structure learning is necessary. Surprisingly, we found a novel structure learning algorithm that only uses product-state input without any ancillary qubits that also achieves the Heisenberg limit when used in the hierarchical learning subroutine.

In Appendix G, we provide rigorous proofs for the lower bound of the general Hamiltonian learning algorithm in Theorem G1, where we utilized the learning tree representation and applied Le Cam's two-point method for Hamiltonian distinguishing problem. The lower bound implies that for a Hamiltonian learning algorithm, there is a trade-off between the total evolution-time complexity and the number of quantum controls.

APPENDIX C: HAMILTONIAN STRUCTURE LEARNING $\mathcal{A}_j^I$

We here provide details for Hamiltonian structure learning $\mathcal{A}_j^I$. In $\mathcal{A}_j^I$, we first use Trotterization to approximate the time-evolution channel under $\tilde{H}_j = (H - \hat{H}_j)/2^{-j}$. Then, we evolve an N -qubit EPR state into this channel and perform a Bell-basis measurement at the end. By analyzing the Taylor expansion of the time-evolution operator, and analyzing how considering the probability distribution of Bell-basis measurement will deviate if only considering the first-order truncation. Taking the Trotterization error $\epsilon_{\text{Trotter}}$, first-order truncation error ϵ_{Taylor} , error caused by

the imperfections of $\hat{H}_j$, and the SPAM error $\varepsilon_{\text{SPAM}}$ into consideration, we show that $\mathcal{A}_j^I$ can identify all the terms in H that is larger than $2^{-(j+1)}$ with high probability.

1. Hamiltonian simulation with Trotterization

By the Lie product formula, the time-evolution operator $e^{-i\tilde{H}_j t}$ can be approximated by interleaving the time-evolution operator $e^{-iHt/(2^{-j}r)}$ and $e^{i\tilde{H}_j t/(2^{-j}r)}$ in the asymptotic limit

$$e^{-i\tilde{H}_j t} = e^{-i(H-\hat{H}_j)t/2^{-j}} = \lim_{r \rightarrow \infty} \left(e^{-iHt/(2^{-j}r)} e^{i\tilde{H}_j t/(2^{-j}r)} \right)^r, \quad (\text{C1})$$

where $e^{i\tilde{H}_j t/(2^{-j}r)}$ is realizable through Hamiltonian simulation since the classical characterization of $\hat{H}_j$ is known.

We consider r is a finite integer, then by the Trotter-Suzuki formula [65–68], we have

$$e^{-i\tilde{H}_j t} = e^{-i(H-\hat{H}_j)t/2^{-j}} = \left(e^{-iHt/(2^{-j}r)} e^{i\tilde{H}_j t/(2^{-j}r)} \right)^r + \mathcal{O} \left([H, \hat{H}_j] \frac{(t/2^{-j})^2}{r} \right). \quad (\text{C2})$$

To characterize the accuracy of approximating the time-evolution channel, denote the time evolution channel under the Hamiltonian $\tilde{H}_j$ as

$$\mathcal{U}_{\tilde{H}_j, t}(\rho) = e^{-i\tilde{H}_j t} \rho e^{i\tilde{H}_j t}, \quad (\text{C3})$$

and the first-order Trotter-Suzuki time-evolution unitary a Hamiltonian $H - \hat{H}_j$ as

$$U_{\text{TS}, \tilde{H}_j}^{(1)}(t) = U_{\text{TS}, H-\hat{H}_j}^{(1)}(t/2^{-j}) = e^{-iHt/2^{-j}} e^{i\tilde{H}_j t/2^{-j}} \quad (\text{C4})$$

with the first-order Trotter-Suzuki time-evolution channel be

$$\mathcal{U}_{\text{TS}, H-\hat{H}_j, t/2^{-j}}^{(1)}(\rho) = U_{\text{TS}, H-\hat{H}_j}^{(1)}(t/2^{-j}) \rho U_{\text{TS}, H-\hat{H}_j}^{(1)\dagger}(t/2^{-j}). \quad (\text{C5})$$

By Eq. (C2), we can bound the first-order approximation error $\varepsilon_{\text{Trotter}}^{(1)}$,

$$\begin{aligned} \varepsilon_{\text{Trotter}}^{(1)} &= \|\mathcal{U}_{\tilde{H}_j, t} - (\mathcal{U}_{\text{TS}, H-\hat{H}_j, t/(2^{-j}r)}^{(1)})^r\|_{\diamond} \\ &= \mathcal{O} \left([\tilde{H}_j, H - \hat{H}_j] \frac{(t/2^{-j})^2}{r} \right), \end{aligned} \quad (\text{C6})$$

where $\|\cdot\|_{\diamond}$ is the diamond norm. Note that if both $H - \hat{H}_j$ and $\tilde{H}_j$ has at most M Pauli terms with coefficients bounded by 1, then $\varepsilon_{\text{Trotter}}$ can be bounded by

$$\varepsilon_{\text{Trotter}} = \mathcal{O} \left(\frac{M^2 (t/2^{-j})^2}{r} \right). \quad (\text{C7})$$

2. Taylor expansion of the time-evolution operator

Consider the Taylor expansion of the time-evolution operator under $\tilde{H}_j$:

$$\begin{aligned} e^{-i\tilde{H}_j t} &= e^{-i(H-\hat{H}_j)t/2^{-j}} = I - it \frac{(H - \hat{H}_j)}{2^{-j}} \\ &\quad - \frac{t^2}{2} \left(\frac{(H - \hat{H}_j)}{2^{-j}} \right)^2 + \mathcal{O}(t^3). \end{aligned} \quad (\text{C8})$$

Recall that $H = \sum_s \mu_s P_s$ and we only care about the coefficients such that $|\mu_s| \geq \epsilon$, thus we can express $H - \hat{H}_j$ as a rescaled Hamiltonian

$$\tilde{H}_j = \frac{H - \hat{H}_j}{2^{-j}} = \sum_{s \in S_j} \frac{\mu_s}{2^{-j}} P_s + \sum_{s \in S_0 \setminus S_j} \frac{\mu_s - \hat{\mu}_s}{2^{-j}} P_s, \quad (\text{C9})$$

where S_j is defined in Eq. (A5), and $\hat{\mu}_s$ are the estimated coefficients in $\hat{H}_j$. Note that the total number of Pauli terms in $\tilde{H}_j$ would still be bounded by M and by definition $|\mu_s - \hat{\mu}_s| \leq \epsilon$, indicating that all rescaled coefficients $(\mu_s - \hat{\mu}_s)/2^{-j}$ in the second terms on the right-hand side of Eq. (C9) with $s \in S_0 \setminus S_j$ will be upper bounded by $1/2$ since $\hat{\mu}_s$ is at most ϵ away from the terms in H_j with $s \in S_0 \setminus S_j$ in the ℓ^∞ -norm, and $2^j \geq 1/2 \cdot \epsilon$. This means that the error caused by inaccurate estimation of the terms in $\hat{H}_j$ will still be bounded even after rescaling. With this, we can rewrite Eq. (C8) as

$$\begin{aligned} e^{-i\tilde{H}_j t} &= I - it \left(\sum_{s \in S_j} \frac{\mu_s}{2^{-j}} P_s + \sum_{s \in S_0 \setminus S_j} \frac{\mu_s - \hat{\mu}_s}{2^{-j}} P_s \right) - \frac{t^2}{2} \left(\sum_{s, s' \in S_j} \frac{\mu_s \mu_{s'}}{2^{-2j}} P_s P_{s'} + \sum_{\substack{s \in S_j \\ s' \in S_0 \setminus S_j}} \frac{\mu_s (\mu_{s'} - \hat{\mu}_{s'})}{2^{-2j}} P_s P_{s'} \right. \\ &\quad \left. + \sum_{\substack{s' \in S_j \\ s \in S_0 \setminus S_j}} \frac{\mu_{s'} (\mu_s - \hat{\mu}_s)}{2^{-2j}} P_s P_{s'} + \sum_{s, s' \in S_0 \setminus S_j} \frac{(\mu_s - \hat{\mu}_s)(\mu_{s'} - \hat{\mu}_{s'})}{2^{-2j}} P_s P_{s'} \right) + o(t^3). \end{aligned} \quad (\text{C10})$$

3. Bell sampling

Consider performing Bell sampling to the time evolution under $e^{-i\tilde{H}_j t}$, with a n -qubit ancillary system. Input the $2n$ -qubit EPR state into the n -qubit time-evolution channel tensor a n -qubit identity channel, by Eq. (C10), the output state is

$$\begin{aligned}
 |\psi_{\text{out}}\rangle &= \left(e^{-i\tilde{H}_j t} \otimes I_n \right) |\text{EPR}_n\rangle \\
 &= \left(\left(I_n - it\tilde{H}_j - \frac{t^2}{2}\tilde{H}_j^2 + o(t^3) \right) \otimes I_n \right) |\text{EPR}_n\rangle \\
 &= \left(\left(I_n - it \left(\sum_{s \in \mathcal{S}_j} \frac{\mu_s}{2^{-j}} P_s + \sum_{s \in \mathcal{S}_0 \setminus \mathcal{S}_j} \frac{\mu_s - \hat{\mu}_s}{2^{-j}} P_s \right) - \frac{t^2}{2}\tilde{H}_j^{(2)} + o(t^3) \right) \otimes I_n \right) |\text{EPR}_n\rangle \\
 &= |\text{EPR}_n\rangle - it \left(\sum_{s \in \mathcal{S}_j} \frac{\mu_s}{2^{-j}} \bigotimes_{j=1}^n ((\sigma_{s,j} \otimes I) |\text{EPR}\rangle) + \sum_{s \in \mathcal{S}_0 \setminus \mathcal{S}_j} \frac{\mu_s - \hat{\mu}_s}{2^{-j}} \bigotimes_{j=1}^n ((\sigma_{s,j} \otimes I) |\text{EPR}\rangle) \right) \\
 &\quad + \left(\left(-\frac{t^2}{2}\tilde{H}_j^2 + o(t^3) \right) \otimes I_n \right) |\text{EPR}_n\rangle
 \end{aligned} \tag{C11}$$

where

$$\tilde{H}_j = \left(\sum_{s \in \mathcal{S}_j} \frac{\mu_s}{2^{-j}} P_s + \sum_{s \in \mathcal{S}_0 \setminus \mathcal{S}_j} \frac{\mu_s - \hat{\mu}_s}{2^{-j}} P_s \right) \tag{C12}$$

and

$$\begin{aligned}
 \tilde{H}_j^2 &= \left(\sum_{s, s' \in \mathcal{S}_j} \frac{\mu_s \mu_{s'}}{2^{-2j}} P_s P_{s'} + \sum_{\substack{s \in \mathcal{S}_j \\ s' \in \mathcal{S}_0 \setminus \mathcal{S}_j}} \frac{\mu_s (\mu_{s'} - \hat{\mu}_{s'})}{2^{-2j}} P_s P_{s'} \right. \\
 &\quad \left. + \sum_{\substack{s' \in \mathcal{S}_j \\ s \in \mathcal{S}_0 \setminus \mathcal{S}_j}} \frac{\mu_{s'} (\mu_s - \hat{\mu}_s)}{2^{-2j}} P_s P_{s'} + \sum_{s, s' \in \mathcal{S}_0 \setminus \mathcal{S}_j} \frac{(\mu_s - \hat{\mu}_s)(\mu_{s'} - \hat{\mu}_{s'})}{2^{-2j}} P_s P_{s'} \right)
 \end{aligned} \tag{C13}$$

are the second-order Pauli terms. Consider measuring $|\psi_{\text{out}}\rangle$ in the Bell basis states, and denote the probability distributions of the measurement outcomes as $\mathcal{D}_{\text{Bell}}$.

Notice that $\mathcal{D}_{\text{Bell}}$ is a discrete probability distribution supporting on the 4^n indices set of $\mathbb{P}_n$, each corresponding to a n -qubit Pauli operator. The goal is to sample all indices in $\mathcal{S}_j$ from $\mathcal{D}_{\text{Bell}}$ efficiently. As we will show in later sections, the probability of elements in $\mathcal{S}_j$ will be lower bounded by an inverse polynomial of M .

Intuitively, if one only considers the first order terms in $|\psi_{\text{out}}\rangle$, the probability of those elements corresponding to a term P_s in H is approximately $t^2 \mu_s^2 / 2^{-2j}$, and the probability of those elements correspond to a traceless Pauli operator P not in H is zero. However, one needs to take

the higher-order terms into consideration, and the probability distribution $\mathcal{D}_{\text{Bell}}$ is not entirely supported on the M elements correspond to the M Pauli terms in H .

Fortunately, our protocol does not require estimating this probability distribution to high accuracy, it works well as long as the probability of those terms with indices in $\mathcal{S}_j$ are at least inverse polynomials of M . Consider the second-order terms, the absolute values of all coefficients in $\tilde{H}_j^2$ are upper bounded by 1. For each term P_s with $s \in \mathcal{S}_j$, there are at most $M/2$ pairs of $\{s_1, s_2\}$ such that $P_{s_1} P_{s_2} = e^{i\theta} P_s$ in the second-order terms, where θ is a phase. Similarly, for the l th order terms there will be at most $\mathcal{O}(M^{l-1})$ l -element sets $\{s_1, s_2, \dots, s_l\}$ such that $\prod_{k=1}^l P_{s_k} = e^{i\theta} P_s$, where θ is a phase. Thus, for an index $s \in \mathcal{S}_j$, the probability of

sampling this element is lower bounded by

$$\Omega \left(\left(\frac{t}{2} - \sum_{l=2}^{\infty} (M^{l-1} t^l) \right)^2 \right), \quad (\text{C14})$$

where $\frac{1}{2}$ is the lower bound of $\mu_s/2^{-j}$ for $s \in \mathfrak{S}_j$, and for higher-ordered terms we use the upper bound 1 for all coefficients. Note that for the terms $s \in S_0 \setminus S_j$, their coefficients are the remanent $|\mu_s - \hat{\mu}_s|$ caused by the inaccurate estimations in previous steps will still be upper bounded by $\epsilon/2^{-j} \leq 1/2$ and do not contribute more than the other terms. If

$$t < \frac{1}{CM}, \quad (\text{C15})$$

where C is a large constant, the probability will be lower bounded by

$$\begin{aligned} & \Omega \left(\left(\frac{1}{2CM} - \sum_{l=2}^{\infty} \frac{1}{C^l M} \right)^2 \right) \\ &= \Omega \left(\left(\frac{1}{2CM} - \frac{1}{C^2 M} \right)^2 \right) \\ &= \Omega \left(\frac{1}{4C^2 M^2} - \frac{1}{C^3 M^2} + \frac{1}{C^4 M^2} \right) \\ &= \Omega \left(\frac{1}{4C^2 M^2} - \frac{1}{C^3 M^2} \right). \end{aligned} \quad (\text{C16})$$

Thus, of the measurement outcome will deviate from $1/4C^2 M^2$ by at most

$$\varepsilon_{\text{Taylor}} = \mathcal{O} \left(\frac{1}{C^3 M^2} \right). \quad (\text{C17})$$

4. SPAM error

We define the SPAM error in the following way.

Definition C1. Take the noise in the preparation of the initial state as an error channel $\mathcal{E}_{\text{prep}}$ applied after the ideal state-preparation channel, and the noise of measurement as an error channel $\mathcal{E}_{\text{meas}}$ applied before the ideal measurement channel. We assume that

$$\|\mathcal{E}_{\text{prep}} - \mathcal{I}\|_{\diamond} + \|\mathcal{E}_{\text{meas}} - \mathcal{I}\|_{\diamond} \leq \varepsilon_{\text{SPAM}},$$

where $\varepsilon_{\text{SPAM}}$ is the bound of the SPAM error.

Therefore, consider a ideal quantum channel $\mathcal{E}$, the channel with SPAM error can be written as

$$\tilde{\mathcal{E}} = \mathcal{E}_{\text{meas}} \circ \mathcal{E} \circ \mathcal{E}_{\text{prep}}, \quad (\text{C18})$$

and

$$\begin{aligned} \|\mathcal{E} - \tilde{\mathcal{E}}\|_{\diamond} &= \|\mathcal{I} \circ \mathcal{E} \circ \mathcal{I} - \mathcal{E}_{\text{meas}} \circ \mathcal{E} \circ \mathcal{E}_{\text{prep}}\|_{\diamond} \\ &= \|\mathcal{I} \circ \mathcal{E} \circ \mathcal{I} - \mathcal{E}_{\text{meas}} \circ \mathcal{E} \circ \mathcal{I} + \mathcal{E}_{\text{meas}} \circ \mathcal{E} \circ \mathcal{I} - \mathcal{E}_{\text{meas}} \circ \mathcal{E} \circ \mathcal{E}_{\text{prep}}\|_{\diamond} \\ &= \|(\mathcal{I} - \mathcal{E}_{\text{meas}}) \circ \mathcal{E} \circ \mathcal{I} + \mathcal{E}_{\text{meas}} \circ \mathcal{E} \circ (\mathcal{I} - \mathcal{E}_{\text{prep}})\|_{\diamond} \\ &\leq \|(\mathcal{I} - \mathcal{E}_{\text{meas}}) \circ \mathcal{E} \circ \mathcal{I}\|_{\diamond} + \|\mathcal{E}_{\text{meas}} \circ \mathcal{E} \circ (\mathcal{I} - \mathcal{E}_{\text{prep}})\|_{\diamond} \\ &= \|(\mathcal{I} - \mathcal{E}_{\text{meas}}) \circ \mathcal{E} \circ \mathcal{I}\|_{\diamond} + \|(\mathcal{E}_{\text{meas}} - \mathcal{I}) \circ \mathcal{E} \circ (\mathcal{I} - \mathcal{E}_{\text{prep}}) + \mathcal{I} \circ \mathcal{E} \circ (\mathcal{I} - \mathcal{E}_{\text{prep}})\|_{\diamond} \\ &\leq \|(\mathcal{I} - \mathcal{E}_{\text{meas}}) \circ \mathcal{E} \circ \mathcal{I}\|_{\diamond} + \|\mathcal{I} \circ \mathcal{E} \circ (\mathcal{I} - \mathcal{E}_{\text{prep}})\|_{\diamond} + \|(\mathcal{E}_{\text{meas}} - \mathcal{I}) \circ \mathcal{E} \circ (\mathcal{I} - \mathcal{E}_{\text{prep}})\|_{\diamond} \\ &\leq \|\mathcal{I} - \mathcal{E}_{\text{meas}}\|_{\diamond} \cdot \|\mathcal{E}\|_{\diamond} + \|\mathcal{E}\|_{\diamond} \cdot \|\mathcal{I} - \mathcal{E}_{\text{prep}}\|_{\diamond} + \|\mathcal{E}_{\text{meas}} - \mathcal{I}\|_{\diamond} \cdot \|\mathcal{E}\|_{\diamond} \cdot \|\mathcal{I} - \mathcal{E}_{\text{prep}}\|_{\diamond} \\ &= (\|\mathcal{I} - \mathcal{E}_{\text{meas}}\|_{\diamond} + \|\mathcal{I} - \mathcal{E}_{\text{prep}}\|_{\diamond} + \|\mathcal{I} - \mathcal{E}_{\text{meas}}\|_{\diamond} \cdot \|\mathcal{I} - \mathcal{E}_{\text{prep}}\|_{\diamond}) \cdot \|\mathcal{E}\|_{\diamond} \\ &\leq \left(\varepsilon_{\text{SPAM}} + \frac{\varepsilon_{\text{SPAM}}^2}{4} \right) \|\mathcal{E}\|_{\diamond} \\ &\leq \varepsilon_{\text{SPAM}} + \frac{\varepsilon_{\text{SPAM}}^2}{4}. \end{aligned} \quad (\text{C19})$$

5. Lower bound for the probability of elements in $\mathfrak{S}_j$

In our protocol, we do not have direct access to the time-evolution channel under $\tilde{H}_j$. Instead, we use the Trotterization method in Appendix C 1. Moreover, SPAM errors in the experiments need to be considered. In analogy with the probability distribution $\tilde{\mathcal{D}}_{\text{Bell}}$ as the probability distribution of Bell sampling in Appendix C 3, define the probability distribution of doing Bell-basis measurement on the output state of input $|\text{EPR}_n\rangle$ to the Trotterization channel for simulation time evolution $e^{-i\tilde{H}_j t}$ as described in Appendix C 1 and $t < 1/CM$ as in Eq. (C15). Consider the probability of the elements in $\mathfrak{S}_j$, we provide the following lemma for a lower bound.

Lemma C1. For the probability distribution $\mathcal{D}_{\text{Bell}}$ defined on the 4^n indices set of $\mathbb{P}_n$, the probability of elements in $\mathfrak{S}_j$ is lower bounded by

$$\gamma_j = \tilde{\Omega} \left(\frac{1}{4C^2M^2} - \frac{1}{C^3M^2} - \frac{2^{2j}}{C^2r} - \varepsilon_{\text{SPAM}} \right), \quad (\text{C20})$$

where C is a large constant as in Eq. (C15), M is the number of terms in H , r is the number of steps in Trotterization, $\varepsilon_{\text{SPAM}}$ is the SPAM error, and $\tilde{\Omega}(f) = \Omega(f) - \mathcal{O}(\varepsilon_{\text{SPAM}}^2)$ only keep the leading order for the simplicity of the expression.

Proof. We first set a target probability of any term with index $s \in \mathfrak{S}_j$ as

$$\tilde{\gamma}_j = \frac{(\min\{\mu_s : s \in \mathfrak{S}_j\})^2 \cdot (1/CM)^2}{2^{-2j}}. \quad (\text{C21})$$

By definition, $\mu_s/2^{-j} \geq 1/2$ for all $s \in \mathfrak{S}_j$,

$$\tilde{\gamma}_j = \Omega \left(\frac{1}{4C^2M^2} \right). \quad (\text{C22})$$

Now we analyze how γ_j deviate from $\tilde{\gamma}_j$. There are three sources that can contribute to the deviation, which are the Trotterization error as analyzed in Appendix C 1, error originated from Bell sampling as in Appendix C 3, and the SPAM error as in Appendix C 4. Notice that the Trotterization error $\varepsilon_{\text{Trotter}}$ in Eq. (C7), and the SPAM error $\varepsilon_{\text{SPAM}}$ are defined for the diamond distance between the ideal and actual quantum channels. By definition, the diamond distance directly implies an upper bound as the total variational distance between the probability distributions of the measurement outcome of the output state between the ideal and actual quantum channels with any input state, which is again by definition an upper bound of the difference between the probability on any element. Thus, the deviation caused by Trotterization and SPAM error is upper

bounded by

$$\mathcal{O} \left(\varepsilon_{\text{Trotter}} + \varepsilon_{\text{SPAM}} + \frac{\varepsilon_{\text{SPAM}}^2}{4} \right).$$

Taking the error in Bell sampling $\varepsilon_{\text{Taylor}}$ as in Eq. (C17) into consideration, we can upper bound the difference

$$|\tilde{\gamma}_j - \gamma_j| = \mathcal{O} \left(\frac{1}{C^3M^2} + \frac{2^{2j}}{C^2r} + \varepsilon_{\text{SPAM}} + \frac{\varepsilon_{\text{SPAM}}^2}{4} \right). \quad (\text{C23})$$

Thus, we have

$$\begin{aligned} \gamma_j &\geq \tilde{\gamma}_j - \mathcal{O} \left(\frac{1}{C^3M^2} + \frac{2^{2j}}{C^2r} + \varepsilon_{\text{SPAM}} + \frac{\varepsilon_{\text{SPAM}}^2}{4} \right) \\ &= \tilde{\Omega} \left(\frac{1}{4C^2M^2} - \frac{1}{C^3M^2} - \frac{2^{2j}}{C^2r} - \varepsilon_{\text{SPAM}} \right). \end{aligned} \quad (\text{C24})$$

■

Note that in the worst case $j = \lceil \log_2(1/\epsilon) \rceil$, and $2^{2j} \approx 1/\epsilon^2$, one need to choose $r = \mathcal{O}(1/\epsilon^2)$ to make the lower bound γ_j independent of ϵ .

6. Determine the high-probability elements of a sparse probability distribution

Lemma C2. Consider a probability distribution $\mathcal{D}$ on a discrete space $\mathfrak{S}$ of size N , let the support of $\text{supp}(\mathcal{D})$ be the set of elements in $\mathfrak{S}$ on which $\mathcal{D}$ has nonzero probability, and let M be the *sparsity* of $\mathcal{D}$, i.e. $M = |\text{supp}(\mathcal{D})| \leq N$. Further define the set of high-probability elements $\mathfrak{S}_\gamma \subseteq \mathfrak{S}$ as

$$\text{supp}_\gamma(\mathcal{D}) = \{e \in \text{supp}(\mathcal{D}) : \Pr(e) \geq \gamma\}. \quad (\text{C25})$$

Note that $|\text{supp}_\gamma(\mathcal{D})| \leq |\text{supp}(\mathcal{D})| = M$. Then with $L = \mathcal{O}(\log(M/\delta_\gamma)/\gamma)$ samples, except for $1 - \delta_\gamma$ probability, each element in $\text{supp}_\gamma(\mathcal{D})$ is sampled at least once.

Proof. If we independently sample from $\mathcal{D}$ for L times, the probability that an element $j \in \text{supp}_\gamma(\mathcal{D})$ is not sampled for at least one time is upper bounded by $(1 - \gamma)^L$. If we want the probability that all M elements in the support to be sampled for at least one time to exceed $1 - \delta_\gamma$, using a union bound, we can set L to satisfy

$$|\text{supp}_\gamma(\mathcal{D})|(1 - \gamma)^L \leq M(1 - \gamma)^L \leq \delta_\gamma. \quad (\text{C26})$$

Thus, the required sample complexity to sample each element at least once for $1 - \delta_\gamma$ probability is

$$L = \mathcal{O} \left(\frac{\log(M/\delta_\gamma)}{\gamma} \right). \quad (\text{C27})$$

■

7. The complexities of $\mathcal{A}_j^I$

Combining the lower bound γ_j in Theorem C1 and the sample complexity in Theorem C2, we can directly give the sample complexity for sampling all indices $s \in \mathfrak{S}_j$ with probability at least δ_j as

$$L_j = \mathcal{O}\left(\frac{\log(M/\delta_j)}{\gamma_j}\right) = \mathcal{O}(M^2 \log(M/\delta_j)). \quad (\text{C28})$$

For each sample, the total evolution time under H is

$$t_{j,l} = \mathcal{O}\left(\frac{1}{CM} \cdot \frac{1}{2^{-j}}\right), \quad (\text{C29})$$

and the total evolution time for $\mathcal{A}_j^I$ is

$$T_{1,j} = \sum_{l=1}^{L_j} t_{j,l} = \mathcal{O}(2^j M \log(M/\delta_j)),$$

$$j = \{0, 1, 2, \dots, \lceil \log_2(1/\epsilon) \rceil\}. \quad (\text{C30})$$

Note that in when $j = \lceil \log_2(1/\epsilon) \rceil$,

$$T_{1, \lceil \log_2(1/\epsilon) \rceil} = \mathcal{O}\left(\frac{M \log(M/\delta_j)}{\epsilon}\right),$$

which reaches the Heisenberg-limited scaling.

APPENDIX D: HAMILTONIAN COEFFICIENT LEARNING $\mathcal{A}_j^H$

In this section we provide proofs for learning Hamiltonian coefficients in $\mathcal{A}_j^H$, which constitutes of two parts Hamiltonian reshaping and robust frequency estimation as we will further illustrate in the followings. The basic idea is that once we identified all possible terms in $\mathfrak{S}_j$ from $\mathcal{A}_j^I$, there are at most L_j possible terms sampled. Let the set of all sampled indices from $\mathcal{A}_j^I$ as $\mathfrak{S}_{L_j}$, with $|\mathfrak{S}_{L_j}| \leq L_j = \mathcal{O}(M^2 \log(M/\delta_j))$. We first use the Hamiltonian reshaping technique as introduced in Appendix D1 to approximate the time evolution under a single term P_s for $s \in \mathfrak{S}_{L_j}$. Later, with the time evolution under the single term P_s , we adapt the robust frequency estimation protocol as introduced in Appendix D2 to estimate μ_s to ϵ accuracy.

1. Hamiltonian reshaping

In this work we will need to use the Hamiltonian reshaping technique to obtain an effective Hamiltonian that consists of only a single Pauli operator. More precisely, given a Pauli operator $P_a \in \{I, X, Y, Z\}^{\otimes n} \setminus \{I^{\otimes n}\}$, we want the effective Hamiltonian to contain only this term, with its coefficient value the same as in the original

Hamiltonian H . To achieve this, we apply, with an interval of τ , Pauli operators randomly drawn from the set

$$\mathcal{K}_{P_a} = \{P \in \mathbb{P}_n : [P_a, P_b] = 0\}, \quad (\text{D1})$$

where $\mathbb{P}_n$ is the set of all n -qubit Pauli operators. Here we can see that $|\mathcal{K}_{P_a}| = 2^{2n-1}$ and $\mathcal{K}_{P_a}$ are the set of all n -qubit Pauli operators that commute with P_a . Also, τ needs to be sufficiently small as will be analyzed in Theorem D1. More precisely, the evolution of the quantum system is described by

$$Q_{r_2} e^{-iH\tau} Q_{r_2} \cdots Q_2 e^{-iH\tau} Q_2 Q_1 e^{-iH\tau} Q_1, \quad (\text{D2})$$

where Q_k , $k = 1, 2, \dots, r_2$, is uniformly randomly drawn from the set $\mathcal{K}_{P_a}$. In one time step of length τ , the quantum state evolves under the quantum channel as

$$\begin{aligned} \rho &\mapsto \frac{1}{2^{2n-1}} \sum_{Q \in \mathcal{K}_{P_a}} Q e^{-iH\tau} Q \rho Q e^{iH\tau} Q \\ &= \rho - \frac{1}{2^{2n-1}} \sum_{Q \in \mathcal{K}_{P_a}} i\tau [QH Q, \rho] + \mathcal{O}(\tau^2) \\ &= \rho - i\tau [H_{\text{eff}}, \rho] + \mathcal{O}(\tau^2) = e^{-iH_{\text{eff}}\tau} \rho e^{iH_{\text{eff}}\tau} + \mathcal{O}(\tau^2), \end{aligned} \quad (\text{D3})$$

where H_{eff} , the effective Hamiltonian, is

$$H_{\text{eff}} = \frac{1}{2^{2n-1}} \sum_{Q \in \mathcal{K}_{P_a}} Q H Q. \quad (\text{D4})$$

The above can be interpreted as a linear transformation applied to the Hamiltonian H . Consequently, it is natural to examine the impact of this transformation on each Pauli term within the Hamiltonian. For a term P , we note that there are two possible outcomes:

Lemma D1. Let P be a Pauli operator and let $\mathcal{K}_{P_a}$ be as defined in Eq. (D1). Then

$$\frac{1}{2^{2n-1}} \sum_{Q \in \mathcal{K}_{P_a}} Q P Q = \begin{cases} P_a, & \text{if } P = P_a, \\ 0, & \text{if } P \neq P_a. \end{cases} \quad (\text{D5})$$

Proof. If $P = P_a$, then $Q P Q = P_a$ for all $Q \in \mathcal{K}_{P_a}$, averaging over all Q therefore yields P_a . If $P \neq P_a$, then there exists $Q_0 \in \mathcal{K}_{P_a}$ such that $P Q_0 = -Q_0 P$ (because $\mathcal{K}_{P_a}$ contains all Pauli operators that commute with P_a and two Pauli matrices either commute or anticommute). Consider the mapping $\phi : \mathcal{K}_{P_a} \rightarrow \mathcal{K}_{P_a}$ defined by $\phi(Q) = Q_0 Q$. This is a bijection from $\mathcal{K}_{P_a}$ to itself, and it can be readily checked that if Q commutes with P , then $\phi(Q)$ anticommutes with P ; if Q anticommutes with P , then $\phi(Q)$ commutes with P . Consequently $Q P Q = P$ for half of all

$Q \in \mathcal{K}_{P_a}$ and $QPQ = -P$ for the other half. Thus taking the average yields 0. ■

This lemma allows us to on average single out the Pauli term P_a we want to preserve in the Hamiltonian and discard all other terms. We have

$$H_{\text{eff}} = \mu_a P_a. \quad (\text{D6})$$

Note that the coefficients μ_a of P_a are preserved in this effective Hamiltonian.

While Lemma D1 concerns the uniform average over all 2^{2n-1} elements of a set of Pauli operators, in our learning protocol we will randomly sample from this set. To assess the protocol's accuracy, we will use the following theorem.

Theorem D1 (Theorem 3 of Ref. [12]). Let P_a traceless n -qubit Pauli operator, and let $\mathcal{K}_{P_a}$ be as defined in Eq. (D1). Let U be the random unitary defined in Eq. (D2). Let $V = e^{-iH_{\text{eff}}}$ for H_{eff} given in Eq. (D6), and $t = r_2 \tau$. We define the quantum channels $\mathcal{U}$ and $\mathcal{V}$ be

$$\mathcal{U}(\rho) = \mathbb{E}[U\rho U^\dagger], \quad \mathcal{V}(\rho) = V\rho V^\dagger.$$

Then

$$\|\mathcal{U} - \mathcal{V}\|_\diamond \leq \frac{4M^2 t^2}{r_2}.$$

Here the norm is the diamond norm (or completely bounded norm), given in terms of the Schatten L_1 norm by $\|\mathcal{U}\|_\diamond := \max_{X: \|X\|_1 \leq 1} \|(\mathcal{U} \otimes I_n)(X)\|_1$.

Proof. (Appendix A of Ref. [12])

Recall from Eq. (D2) that

$$U = Q_r e^{-iH\tau} Q_r \cdots Q_1 e^{-iH\tau} Q_1.$$

We define

$$U_j = Q_j e^{-iH\tau} Q_j,$$

and the quantum channel $\mathcal{U}_j$

$$\mathcal{U}_j(\rho) = \mathbb{E}_{Q_j} U_j \rho U_j^\dagger.$$

Because each Q_j is chosen independently, we have

$$\mathcal{U} = \mathcal{U}_{r_2} \mathcal{U}_{r_2-1} \cdots \mathcal{U}_1.$$

We then define $\mathcal{V}_\tau(\rho) = e^{-iH_{\text{eff}}\tau} \rho e^{iH_{\text{eff}}\tau}$, which then gives us $\mathcal{V} = \mathcal{V}_\tau^2$. We will then focus on obtaining a bound for $\|\mathcal{U}_j - \mathcal{V}_\tau\|_\diamond$.

Consider a quantum state ρ on the joint system consists of the current system and an auxiliary system denoted by

α , and an observable O on the combined quantum system such that $\|O\| \leq 1$. We have

$$\begin{aligned} & \text{Tr}[O(I_\alpha \otimes U_j)\rho(I_\alpha \otimes U_j^\dagger)] \\ &= \text{Tr}[O(I_\alpha \otimes (Q_j e^{-iH\tau} Q_j))\rho(I_\alpha \otimes (Q_j e^{iH\tau} Q_j))] \\ &= \text{Tr}[O\rho] - i\tau \text{Tr}[O[I_\alpha \otimes (Q_j H Q_j), \rho]] \\ &\quad + \tau^2 \text{Tr}[O(I_\alpha \otimes (Q_j H Q_j))\rho(I_\alpha \otimes (Q_j H Q_j))] \\ &\quad - \frac{\tau^2}{2} \text{Tr}[O(I_\alpha \otimes (Q_j H^2 e^{-iH\xi} Q_j))\rho] \\ &\quad - \frac{\tau^2}{2} \text{Tr}[O\rho(I_\alpha \otimes (Q_j H^2 e^{+iH\xi} Q_j))], \end{aligned}$$

where we are using Taylor's theorem with the remainder in the Lagrange form and $\xi \in [0, \tau]$. This tells us that

$$\begin{aligned} & \left| \text{Tr}[O(I_\alpha \otimes U_j)\rho(I_\alpha \otimes U_j^\dagger)] \right. \\ & \quad \left. - (\text{Tr}[O\rho] - i\tau \text{Tr}[O[I_\alpha \otimes (Q_j H Q_j), \rho]]) \right| \leq 2\|H\|^2 \tau^2. \end{aligned}$$

Because the two terms on the left-hand side, when we take the expectation value with respect to Q_j , become

$$\begin{aligned} & \mathbb{E}_{Q_j} \text{Tr}[O(I_\alpha \otimes U_j)\rho(I_\alpha \otimes U_j^\dagger)] = \text{Tr}[O(\mathcal{I}_\alpha \otimes \mathcal{U}_j)(\rho)] \\ & \mathbb{E}_{Q_j} [\text{Tr}[O\rho] - i\tau \text{Tr}[O[I_\alpha \otimes (Q_j H Q_j), \rho]]] \\ &= \text{Tr}[O\rho] - i\tau \text{Tr}[O[I_\alpha \otimes H_{\text{eff}}, \rho]], \end{aligned}$$

we have

$$\begin{aligned} & \left| \text{Tr}[O(\mathcal{I}_\alpha \otimes \mathcal{U}_j)(\rho)] - (\text{Tr}[O\rho] - i\tau \text{Tr}[O[I_\alpha \otimes H_{\text{eff}}, \rho]]) \right| \\ & \leq 2\|H\|^2 \tau^2. \end{aligned}$$

Because this is true for all O such that $\|O\| \leq 1$,

$$\|(\mathcal{I}_\alpha \otimes \mathcal{U}_j)(\rho) - (\rho - i\tau[I_\alpha \otimes H_{\text{eff}}, \rho])\|_1 \leq 2\|H\|^2 \tau^2. \quad (\text{D7})$$

Similarly we can also show that

$$\|(\mathcal{I}_\alpha \otimes \mathcal{V}_\tau)(\rho) - (\rho - i\tau[I_\alpha \otimes H_{\text{eff}}, \rho])\|_1 \leq 2\|H_{\text{eff}}\|^2 \tau^2. \quad (\text{D8})$$

Combining Eqs. (D7) and (D8), and by the triangle inequality, we have

$$\begin{aligned} & \|(\mathcal{I}_\alpha \otimes \mathcal{U}_j)(\rho) - (\mathcal{I}_\alpha \otimes \mathcal{V}_\tau)(\rho)\|_1 \\ & \leq 2\|H\|^2 \tau^2 + 2\|H_{\text{eff}}\|^2 \tau^2. \end{aligned}$$

Because the above is true for all auxiliary system α and for all quantum states on the combined quantum system, we

have

$$\|\mathcal{U}_j - \mathcal{V}_\tau\|_\diamond \leq 2\|H\|^2\tau^2 + 2\|H_{\text{eff}}\|^2\tau^2. \quad (\text{D9})$$

Next, we observe that

$$\mathcal{U} - \mathcal{V} = \sum_{k=1}^{r_2} \mathcal{U}_r \cdots \mathcal{U}_{k+1} (\mathcal{U}_k - \mathcal{V}_\tau) \mathcal{V}_\tau^{k-1}.$$

Therefore,

$$\begin{aligned} \|\mathcal{U} - \mathcal{V}\|_\diamond &\leq \sum_{k=1}^{r_2} \|\mathcal{U}_k - \mathcal{V}_\tau\|_\diamond \\ &\leq r_2(2\|H\|^2\tau^2 + 2\|H_{\text{eff}}\|^2\tau^2). \end{aligned}$$

The bound in the theorem follows by observing that $\|H\|, \|H_{\text{eff}}\| \leq M$, and $\tau = t/r_2$. ■

2. Robust frequency estimation

In this section we introduce the robust frequency-estimation protocol in Ref. [13], which largely follows the idea of the robust phase-estimation protocol in Ref. [37] and is also used in Ref. [12]. However, unlike in the previous works, which require highly entangled Bell-state input, as described in Sec. II B, we only require product state input in this work, since the effective Hamiltonian as in Eq. (D6) has only one Pauli term, which means that we only need to set the input state to have maximum sensitivity on one qubit. Here we restate the proofs in Ref. [12].

Theorem D2 (Robust frequency estimation [12]). Let $\theta \in [-A, A]$. Let $X(t)$ and $Y(t)$ be independent random variables satisfying

$$\begin{aligned} |X(t) - \cos(\theta t)| &< 1/\sqrt{2}, \text{ with probability at least } 2/3, \\ |Y(t) - \sin(\theta t)| &< 1/\sqrt{2}, \text{ with probability at least } 2/3. \end{aligned} \quad (\text{D10})$$

Then with K independent nonadaptive [69] samples $X(t_1), X(t_2), \dots, X(t_K)$ and $Y(t_1), Y(t_2), \dots, Y(t_K)$, $t_j \geq 0$, for

$$K = \mathcal{O}(\log(A/\epsilon)(\log(1/q) + \log \log(A/\epsilon))), \quad (\text{D11})$$

$$T = \sum_{j=1}^K t_j = \mathcal{O}((1/\epsilon)(\log(1/q) + \log \log(A/\epsilon))),$$

$$\max_j t_j = \mathcal{O}(1/\epsilon), \quad (\text{D12})$$

we can obtain a random variable $\hat{\theta}$ such that

$$\Pr[|\hat{\theta} - \theta| > \epsilon] \leq q. \quad (\text{D13})$$

We modify the protocol in Ref. [13] because our goal is to obtain an accurate estimate with large probability rather than having an optimal mean-squared error scaling. The key tool is the following lemma that allows us to incrementally refine the frequency estimate:

Lemma D2. Let $\theta \in [a, b]$. Let $Z(t)$ be a random variable such that

$$|Z(t) - e^{i\theta t}| < 1/2. \quad (\text{D14})$$

Then we can correctly distinguish between two overlapping cases $\theta \in [a, (a+2b)/3]$ and $\theta \in [(2a+b)/3, b]$ with one sample of $Z(\pi/(b-a))$.

Proof. These two situations can be distinguished by looking at the value of

$$f(\theta) = \sin\left(\frac{\pi}{b-a}\left(\theta - \frac{a+b}{2}\right)\right).$$

We know from Eq. (D14) that

$$\left| \text{Im}\left(e^{-i\frac{(a+b)\pi}{2(b-a)}} Z(\pi/(b-a))\right) - f(\theta) \right| < 1/2,$$

where t in Eq. (D14) is substituted by $\pi/(b-a)$ and a phase factor is added.

If

$$\text{Im}\left(e^{-i\frac{(a+b)\pi}{2(b-a)}} Z(\pi/(b-a))\right) \leq 0,$$

then $f(\theta) < 1/2$, which implies $\theta \in [a, (a+2b)/3]$. If

$$\text{Im}\left(e^{-i\frac{(a+b)\pi}{2(b-a)}} Z(\pi/(b-a))\right) > 0,$$

then $f(\theta) > -1/2$, which implies $\theta \in [(2a+b)/3, b]$. ■

Using this lemma we will prove Theorem D2, which we restate below:

Theorem. Let $\theta \in [-A, A]$. Let $X(t)$ and $Y(t)$ be independent random variables satisfying

$$\begin{aligned} |X(t) - \cos(\theta t)| &< 1/\sqrt{2}, \text{ with probability at least } 2/3, \\ |Y(t) - \sin(\theta t)| &< 1/\sqrt{2}, \text{ with probability at least } 2/3. \end{aligned} \quad (\text{D15})$$

Then with independent nonadaptive samples $X(t_1), X(t_2), \dots, X(t_K)$ and $Y(t_1), Y(t_2), \dots, Y(t_K)$, $t_j \geq 0$, for

$$K = \mathcal{O}(\log(A/\epsilon)(\log(1/q) + \log \log(A/\epsilon))), \quad (\text{D16})$$

$$T = \sum_{j=1}^K t_j = \mathcal{O}((1/\epsilon)(\log(1/q) + \log \log(A/\epsilon))),$$

$$\max_j t_j = \mathcal{O}(1/\epsilon), \quad (\text{D17})$$

we can obtain a random variable $\hat{\theta}$ such that

$$\Pr[|\hat{\theta} - \theta| > \epsilon] \leq q. \quad (\text{D18})$$

Proof. We let $\lambda = \epsilon/3$. We build a random variable $S(t)$ satisfying Eq. (D14), with which we will iteratively narrow down the interval $[a, b]$ containing θ until $|b - a| \leq 2\lambda$, at which point we choose $\hat{\theta} = (a + b)/2$. If $\theta \in [a, b]$ we will then ensure $|\hat{\theta} - \theta| \leq \lambda$. However, each iteration will involve some failure probability, which we will analyze later.

To build the random variable $S(t)$, we first use m -independent samples of $X(t)$ and then take median $X_{\text{median}}(t)$, which satisfies

$$|X_{\text{median}}(t) - \cos(\theta t)| \leq 1/\sqrt{2},$$

with probability at least $1 - \delta/2$, where by Eq. (D15) and the Chernoff bound

$$\delta = c_1 e^{-c_2 m},$$

for some universal constant c_1, c_2 . Similarly we can obtain $Y_{\text{median}}(t)$ such that

$$|Y_{\text{median}}(t) - \cos(\theta t)| \leq 1/\sqrt{2}$$

with probability at least $1 - \delta/2$. With these medians we then define

$$S(t) = X_{\text{median}}(t) + iY_{\text{median}}(t).$$

This random variable satisfies

$$|S(t) - e^{i\theta t}| \leq 1/2$$

with probability at least $1 - \delta$ using the union bound. It therefore allows us to solve the discrimination task in Lemma D2 with probability at least $1 - \delta$.

Whether each iteration proceeds correctly or not, the algorithm terminates after

$$L = \lceil \log_{3/2}(A/\lambda) \rceil$$

iterations. In the l th iteration we sample $X(s_l)$ and $Y(s_l)$ where $s_l = (\pi/2A)(3/2)^{l-1}$. We use m samples of $X(s_l)$ and $Y(s_l)$ for computing the median in each iteration, and therefore the failure probability is at most $\delta = c_1 e^{-c_2 m}$. With probability at most $L\delta$ using the union bound, one of

the iteration fails. In order to ensure the protocol succeeds with probability at least $1 - q$, it suffices to let $\delta \leq q/L$, and it therefore suffices to choose

$$m = \lceil c_3 \log(L/q) \rceil = \mathcal{O}(\log(1/q) + \log \log(A/\epsilon)).$$

The total number of samples (for either $X(t)$ or $Y(t)$) is therefore as described in Eq. (D16). All the t in each sample added together is

$$T = m \sum_{l=1}^L s_l = \frac{m\pi}{2A} \sum_{l=1}^L \left(\frac{3}{2}\right)^{l-1} \leq \frac{3\pi m}{2\epsilon} = \mathcal{O}((1/\epsilon)(\log(1/q) + \log \log(A/\epsilon))), \quad (\text{D19})$$

thus giving us Eq. (D17). ■

3. The experimental setup of $\mathcal{A}_j''$

With the aforementioned lemmas, we now describe the experimental setup for robustly estimating the coefficient of each term with ϵ accuracy. In an experiment for learning the coefficient μ_s of P_s , let $\beta_s \in \{I, x, y, z\}^{\otimes n}$ satisfies $P_s = \bigotimes_{j=1}^n \sigma^{\beta_{s,j}}$, where $\beta_{s,j}$ is the j th entries of β_s . Denote the set of qubit indices where P_s nontrivially act on as S_{P_s} . As defined in Eq. (A1), we set

$$\begin{aligned} |1, I\rangle &= |0\rangle, & |1, x\rangle &= |+\rangle, & |1, y\rangle &= \frac{1}{\sqrt{2}}(|0\rangle + i|1\rangle), \\ |1, z\rangle &= |0\rangle, & |-1, I\rangle &= |0\rangle, & |-1, x\rangle &= |-\rangle, \\ |-1, y\rangle &= \frac{1}{\sqrt{2}}(|0\rangle - i|1\rangle), & |-1, z\rangle &= |1\rangle, \end{aligned} \quad (\text{D20})$$

and we define the following ± 1 eigenstates of P_s as

$$|1, \beta_s\rangle = \bigotimes_{j=1}^n |1, \beta_{s,j}\rangle, \quad |-1, \beta_s\rangle = \bigotimes_{j=1}^n |-1, \beta_{s,j}\rangle. \quad (\text{D21})$$

We set the input state as

$$\begin{aligned} |\phi_0^s\rangle &= \frac{1}{\sqrt{2}}(|1, \beta_s\rangle + |-1, \beta_s\rangle) = \left(\bigotimes_{j=1}^{s^*-1} |1, \beta_{s,j}\rangle \right) \\ &\otimes \frac{1}{\sqrt{2}}(|1, \beta_{s,s^*}\rangle + |-1, \beta_{s,s^*}\rangle) \\ &\otimes \left(\bigotimes_{j=s^*+1}^n |1, \beta_{s,j}\rangle \right), \end{aligned} \quad (\text{D22})$$

which is a product state, and is the equal-weight superposition of two eigenstates of $H_{\text{eff}} = \mu_s P_s$ and thus has maximal sensitivity to the time evolution under H_{eff} . Evolve

$|\phi_0^s\rangle$ under the effective Hamiltonian H_{eff} as described in Appendix D 1, at time t we will approximately obtain the state

$$|\phi_t^s\rangle = \frac{1}{\sqrt{2}}(e^{-i\mu_s t} |1, \beta_s\rangle + e^{i\mu_s t} |-1, \beta_s\rangle). \quad (\text{D23})$$

In the end, we then measure $|\phi_t^s\rangle$ with observables

$$\begin{aligned} O_s^+ &= \left(\bigotimes_{j=1}^{s^*-1} |1, \beta_{s,j}\rangle \right) \otimes Q_{s,j^*}^+ \otimes \left(\bigotimes_{j=s^*+1}^n |1, \beta_{s,j}\rangle \right), \\ O_s^- &= \left(\bigotimes_{j=1}^{s^*-1} |1, \beta_{s,j}\rangle \right) \otimes Q_{s,j^*}^- \otimes \left(\bigotimes_{j=s^*+1}^n |1, \beta_{s,j}\rangle \right), \end{aligned} \quad (\text{D24})$$

where Q_{s,j^*}^+ and Q_{s,j^*}^- are chosen to be single-qubit Pauli operators such that

$$\begin{aligned} Q_{s,j^*}^+ |1, \beta_{s,j^*}\rangle &= |-1, \beta_{s,j^*}\rangle, \quad Q_{s,j^*}^+ |-1, \beta_{s,j^*}\rangle = |1, \beta_{s,j^*}\rangle; \\ Q_{s,j^*}^- |1, \beta_{s,j^*}\rangle &= i |-1, \beta_{s,j^*}\rangle, \\ Q_{s,j^*}^- |-1, \beta_{s,j^*}\rangle &= -i |1, \beta_{s,j^*}\rangle. \end{aligned} \quad (\text{D25})$$

Therefore, the expectation value is

$$\langle \phi_t^s | O_s^+ | \phi_t^s \rangle = \cos(2\mu_s t), \quad \langle \phi_t^s | O_s^- | \phi_t^s \rangle = \sin(2\mu_s t). \quad (\text{D26})$$

Note that in order to measure O_s^+ and O_s^- , we only need to measure the j^* th qubit since all other qubits are in the $+1$ eigenstates and would not affect the phase.

We summarize the above experiment in the following definition:

Definition D1 (Frequency estimation experiment $\mathcal{A}^H$). We call the procedure below a (s, t, τ) -phase estimation experiment:

- (1) Prepare the initial product state $|\phi_0^s\rangle = \frac{1}{\sqrt{2}}(|1, \beta_s\rangle + |-1, \beta_s\rangle)$ as in Eq. (D22).
- (2) Let the system evolve for time t while applying random Pauli operators from $\mathcal{K}_{P_s}$ [defined in Eq. (D1)] with interval τ .
- (3) Measure the observables O_s^+ or O_s^- (by measuring the j^* th qubit in the Q_{s,j^*}^+ or Q_{s,j^*}^- basis) as defined above to obtain a ± 1 outcome.

The goal of the above experiment is to estimate μ_s , for we need the robust frequency-estimation protocol as introduced in Appendix D 2. We will discuss how this is done and analyze the effect of the Hamiltonian reshaping error

and the SPAM error as defined in Definition C1 in the following.

Through a (s, t, τ) -phase estimation experiment, if we measure O_s^+ in the end, we will obtain a random variable $v_s^+(t) \in \pm 1$. If we have the exact state $|\phi_t^s\rangle$ as defined in Eq. (D23), then we can have $\mathbb{E}[v_s^+(t)] = \cos(2\mu_s t)$. However, due to the Hamiltonian reshaping error and the SPAM error, we have

$$|\mathbb{E}[v_s^+(t)] - \cos(2\mu_s t)| \leq \frac{4M^2 t^2}{r_2} + \epsilon_{\text{SPAM}}. \quad (\text{D27})$$

Notice that the first term on the right-hand side comes from the Hamiltonian reshaping error, where we set $t = r_2 \tau$ as in Theorem D1. Moreover the variance of $v_s^+(t)$ is at most 1 because it can only take values ± 1 . We assume that $\epsilon_{\text{SPAM}} \leq 1/(3\sqrt{2})$, and choose r_2 to be $r_2 = \mathcal{O}(M^2 t^2)$ so that

$$\frac{4M^2 t^2}{r_2} < \frac{1}{3\sqrt{2}}.$$

Note that τ and r_2 are related through $\tau = t/r_2$. Then we have

$$|\mathbb{E}[v_s^+(t)] - \cos(2\mu_s t)| < \frac{2}{3\sqrt{2}}.$$

We then take 54 independent samples of $v_s^+(t)$ and average them, denoting the sample average by $X_s(t)$.

By Chebyshev's inequality, this ensures

$$\begin{aligned} \Pr[|X_s(t) - \mathbb{E}[v_s^+(t)]| \geq 1/(3\sqrt{2})] \\ &= \Pr[|X_s(t) - \mathbb{E}[X_s(t)]| \geq 1/(3\sqrt{2})] \\ &\leq \frac{1}{54 \times 1/(3\sqrt{2})^2} = \frac{1}{3}. \end{aligned}$$

Therefore, combining the above with Eq. (D27) we have

$$\Pr[|X_s(t) - \cos(2\mu_s t)| \geq 1/\sqrt{2}] \leq 1/3.$$

This guarantees estimating $\cos(2\mu_s t)$ to a constant $1/\sqrt{2}$ accuracy with at least $2/3$ probability, which gives us the $X_s(t)$ required in the robust frequency estimation protocol in Theorem D2. The $Y_s(t)$ in Theorem D2 can be similarly obtained by using O_s^- as the observable. We also note that because $|2\mu_s| \leq 2$, we can set $A = 2$ in Theorem D2. Therefore we can state the following for the phase-estimation experiment (using the notation of $\mathcal{K}_{P_s}$ as defined in Eq. (D1):

Theorem D3 (Theorem D2 for the Hamiltonian learning task). We assume that the quantum system is evolving under a Hamiltonian H with M terms with the absolute value of the coefficient of each term bounded by 1. With N_{exp} independent nonadaptive (s_l, t_l, τ_l) -phase-estimation

experiments (Definition D1), $j = 1, 2, \dots, L$ (L is the number of coefficients of different Pauli terms in H to be estimated), with the SPAM error (Definition C1) satisfying $\varepsilon_{\text{SPAM}} \leq 1/(3\sqrt{2})$, we can obtain an estimate $\hat{\theta}$ such that

$$\Pr[|\hat{\theta} - 2\mu_s| \geq \eta] \leq q,$$

where μ_s are eigenvalues of the effective Hamiltonian $H_{\text{eff}} = \mu_s P_s$ as defined in Eq. (D6). In the above N_{exp} , $\{t_l\}$, and $\{\tau_l\}$ satisfy

$$\begin{aligned} N_{\text{exp}} &= \mathcal{O}(\log(1/\eta)(\log(1/q) + \log \log(1/\eta))), \\ T &= \sum_{l=1}^{N_{\text{exp}}} t_l = \mathcal{O}\left(\frac{1}{\eta}(\log(1/q) + \log \log(1/\eta))\right), \\ \max_l t_l &= \mathcal{O}(1/\eta), \end{aligned} \quad (\text{D28})$$

and $\tau_l = \Omega(1/(M^2 t_l))$.

4. The complexity of $\mathcal{A}_j''$

In $\mathcal{A}_j''$, we learn each coefficient in $\mathfrak{S}_{L_j}$ to accuracy ϵ . There are $|\mathfrak{S}_{L_j}|$ terms to be estimated by the robust

frequency estimation as described in Appendix D2. By Eq. (D17), the total evolution time under H in each robust frequency-estimation instance is

$$\begin{aligned} t_{j,l} &= \mathcal{O}((1/\epsilon)(\log(1/\epsilon) + \log \log(2^{-j}/\epsilon))), \\ l &= \{0, 1, 2, \dots, |\mathfrak{S}_{L_j}|\}. \end{aligned} \quad (\text{D29})$$

Therefore, the total evolution time under H in $\mathcal{A}_j''$ is

$$\begin{aligned} T_{2,j} &= \sum_{l=1}^{|\mathfrak{S}_{L_j}|} t_{j,l} \\ &= |\mathfrak{S}_{L_j}| \cdot \mathcal{O}((1/\epsilon)(\log(1/\epsilon) + \log \log(2^{-j}/\epsilon))) \\ &= \mathcal{O}\left(\frac{M^2 \log(M/\delta_j)}{\epsilon}(\log(1/\epsilon) + \log \log(2^{-j}/\epsilon))\right). \end{aligned} \quad (\text{D30})$$

APPENDIX E: TOTAL EVOLUTION-TIME COMPLEXITY

Combining the results in Appendixes C7 and D4, the total evolution time in $\mathcal{A}_j$ is

$$T_j = T_{1,j} + T_{2,j} = \mathcal{O}\left(2^j M \log(M/\delta_j) + \frac{M^2 \log(M/\delta_j)}{\epsilon}(\log(1/\epsilon) + \log \log(2^{-j}/\epsilon))\right). \quad (\text{E1})$$

Thus, the total number of experiments is and the total evolution time complexity of our protocol is

$$L = \sum_{j=0}^{\lceil \log_2(1/\epsilon) \rceil - 1} L_j = \mathcal{O}(M^2 \log(M/\delta) \log(1/\epsilon)), \quad (\text{E2})$$

and the total evolution-time complexity of our protocol is

$$T = \sum_{j=0}^{\lceil \log_2(1/\epsilon) \rceil - 1} T_j = \mathcal{O}\left(\frac{M \log(M/\delta) \log(1/\epsilon)}{\epsilon} + \frac{M^2 \log(M/\delta) \log(1/\epsilon)}{\epsilon}(\log(1/\epsilon) + \log \log(1/\epsilon))\right), \quad (\text{E3})$$

where $\delta = \min\{\delta_j\}$.

APPENDIX F: ALTERNATIVE SINGLE-COPY PRODUCT-STATE INPUT APPROACH $\mathcal{A}_j'''$

We here provide the proof for the alternative approach of realizing the large-term identification in $\mathcal{A}_j'''$, which requires no ancillary qubit and only product-state input. With this approach, the entire Hamiltonian learning protocol requires no ancillary qubits, only product-state input, and still achieves Heisenberg-limited scaling. This

approach replaced the step of identifying terms in $\mathfrak{S}_j$ by using the population recovery protocol for estimating the Pauli error rate of a general quantum channel in Ref. [38]. Instead of using this protocol to directly learn the Pauli error rates to a high accuracy, we learn them to a low accuracy and only take those large terms over a threshold to be estimated to high accuracy in $\mathcal{A}_j''$.

1. Pauli error rate for time-evolution channel

Consider a general quantum channel Λ , one can write it in its Kraus operator expression as in Eq. (A7).

Definition F1 (Pauli twirling of a quantum channel, Definition 28 in Ref. [38]). Let Λ denote an arbitrary n -qubit quantum channel. Its Pauli twirl Λ_P is the n -qubit quantum channel defined by

$$\Lambda_P(\rho) = \mathbb{E}_{\sigma_T \in \{I, \sigma_x, \sigma_y, \sigma_z\}} \left[\sigma_T^\dagger \left(\Lambda \left(\sigma_T \rho \sigma_T^\dagger \right) \right) \sigma_T \right]. \quad (\text{F1})$$

Fact F1 (Pauli error rates of general quantum channels). Suppose we write K_j for the Kraus operators of Λ , so $\Lambda\rho = \sum_j K_j \rho K_j^\dagger$. Further suppose that K_j is represented in the Pauli basis as in Eq. (A8). Then Λ 's Pauli error rates are given by $p(k) = \sum_j |\alpha_{j,k}|^2$.

Consider the time-evolution channel under Hamiltonian $\tilde{H}_j$ as in Appendix C3, there is only one Kraus operator in this Pauli channel $K = e^{-i\tilde{H}_j t}$. We obtain the expansion of this Kraus operator in the Pauli basis by considering its Taylor expansion as in Eq. (C10). Similar to Eq. (C15), we set $t < 1/CM$. Then, by Appendix C5 the Pauli error rate of a term P_s where $s \in \mathfrak{S}_j$ is lower bounded by $\tilde{\Omega}(1/4C^2M^2 - 1/C^3M^2 - \varepsilon_{\text{SPAM}})$.

2. Estimating Pauli error rates via population recovery

Theorem F1 (Theorem 1 in Ref. [38]). There is a learning algorithm that, given parameters $0 < \delta, \epsilon_1 < 1$, as well as access to an n -qubit channel with Pauli error rates p , has the following properties:

- (1) It prepares $m = \mathcal{O}(1/\epsilon_1^2) \cdot \log(n/\epsilon_1\delta)$ unentangled n -qubit pure states, where each of the mn 1-qubit states is chosen uniformly at random from $\{|0\rangle, |1\rangle, |+\rangle, |-\rangle, |i\rangle, |-i\rangle\}$.
- (2) It passes these m states through the Pauli channel.
- (3) It performs unentangled measurements on the resulting states, with each qubit being measured in either the $\{|0\rangle, |1\rangle\}$ basis, the $\{|+\rangle, |-\rangle\}$ basis, or $\{|i\rangle, |-i\rangle\}$ basis.
- (4) It performs an $\mathcal{O}(mn/\epsilon_1)$ -time classical postprocessing algorithm on the resulting mn measurement outcome bits.
- (5) It outputs hypothesis Pauli error rates $\hat{p}$ in the form of a list of at most $\frac{4}{\epsilon_1}$ pairs $(k, \hat{p}(k))$, with all unlisted $\hat{p}$ values treated as 0.

The algorithm's hypothesis $\hat{p}$ will satisfy $\|\hat{p} - p\|_\infty \leq \epsilon_1$ except with probability at most δ .

Fact F2 (Fact 29 in Ref. [38]). Applying the population recovery protocol in Ref. [38] to a general quantum

channel Λ can learn its Pauli error rates with the properties mentioned in Theorem F1.

Lemma F1 (Lower bound for Pauli error rates). For any term P_s where $s \in \mathfrak{S}_j$, the Pauli error rate of the time-evolution channel under $e^{-i\tilde{H}_j t}$ realized by Trotterization is lower bounded by

$$\gamma_j = \tilde{\Omega} \left(\frac{1}{4C^2M^2} - \frac{1}{C^3M^2} - \frac{2^{2j}}{C^2r} - \varepsilon_{\text{SPAM}} \right) \quad (\text{F2})$$

for $t = \Theta(1/CM)$, and $\varepsilon_{\text{SPAM}}$ is the SPAM error in experiments.

Proof. The proof is exactly the same as the proof of Theorem C1 since the analysis for Taylor expansion in Appendix F1 is the same with the analysis in Appendix C3. ■

Let the precision of learning Pauli errors be

$$\epsilon_1 = \left(\frac{1}{2} - c' \right) \frac{1}{CM}, \quad (\text{F3})$$

where c' satisfies that $(\frac{1}{2} - c') \frac{1}{CM} < \gamma_j/2$ and is only introduced for simplifying the expression, it would not contribute to the leading terms in the following complexity analysis. Then, for a term P_s where $s \in \mathfrak{S}_j$ with the absolute value its rescaled coefficient in $\tilde{H}_j$ larger than $1/2$, by applying the population recovery protocol in Theorem F1, the learned coefficient recovered from the estimated Pauli error rate will be lower bounded by $1/4$ with probability at least $1 - \delta$. By setting the threshold to be $\frac{1}{2}\gamma_j$ and pass all terms over this threshold to be estimated to high accuracy in $\mathcal{A}_j''$, we can cover all terms in $\mathfrak{S}_j$ with high probability.

3. Complexity of $\mathcal{A}_j''$

By Theorem F1, the number of input states $m = \mathcal{O}(1/\epsilon_1^2) \cdot \log(n/\epsilon_1\delta)$, and from Appendix F2, we set $\epsilon_1 = \frac{1}{2}\gamma_j$, then

$$m = \mathcal{O} \left(\frac{C^4M^4}{(\frac{1}{2} - c')^2} \log \left(\frac{CMn}{\delta} \right) \right). \quad (\text{F4})$$

For each input, the total evolution time under H is

$$t'_{j,1} = t/2^{-j} = \Theta \left(\frac{1}{CM2^{-j}} \right), \quad (\text{F5})$$

so the total evolution time of $\mathcal{A}_j''$ is

$$\begin{aligned} T'_{1,j} &= m \cdot t'_{j,1} = \mathcal{O} \left(\frac{C^3M^3}{2^{-j}(\frac{1}{2} - c')} \log \left(\frac{CMn}{\delta} \right) \right) \\ &\approx \mathcal{O} \left(\frac{C^3M^3}{2^{-j}} \log \left(\frac{CMn}{\delta} \right) \right). \end{aligned} \quad (\text{F6})$$

Meanwhile, by Theorem F1, $\mathcal{A}_j^I$ requires a T_j^C -time classical postprocessing algorithm on the resulting mn measurement outcome bits

$$T_j^C = \mathcal{O}\left(\frac{mn}{\epsilon_1}\right) = \mathcal{O}\left(\frac{C^5 M^5 n}{(\frac{1}{2} - c')^2} \log\left(\frac{CMn}{\delta}\right)\right). \quad (\text{F7})$$

$$T' = \sum_{j=0}^{\lceil \log_2(1/\epsilon) \rceil - 1} (T_{1,j} + T_{2,j}) = \mathcal{O}\left(\frac{C^3 M^3}{\epsilon} \log\left(\frac{CMn}{\delta}\right) \log(1/\epsilon) + \frac{M^2 \log(M/\delta) \log(1/\epsilon)}{\epsilon} (\log(1/\epsilon) + \log \log(1/\epsilon))\right), \quad (\text{F8})$$

and the total classical postprocessing time is

$$T^C = \sum_{j=0}^{\lceil \log_2(1/\epsilon) \rceil - 1} T_j^C = \mathcal{O}\left(C^5 M^5 n \log\left(\frac{CMn}{\delta}\right) (\lceil \log_2(1/\epsilon) \rceil - 1)\right). \quad (\text{F9})$$

APPENDIX G: PROOF FOR THE TRADE-OFF

In this section, we provide the proof for the lower bound in Result 3, which is formulated as follows.

Theorem G1. For any protocol that can be possibly adaptive, biased or unbiased, and possibly ancilla-assisted with total evolution time T and at most $\mathcal{L}$ discrete quantum controls per experiment (as shown in Fig. S1), there exists some parametrized Hamiltonian H that requires $T = \Omega(\mathcal{L}^{-1} \epsilon^{-2})$ to estimate the coefficient of the unknown Hamiltonian H within additive error ϵ .

We consider the problem of learning a Hamiltonian of M terms

$$H(\boldsymbol{\mu}) = \sum_{s=1}^M \mu_s P_s \quad (\text{G1})$$

4. Total evolution-time complexity of the single-copy product-state input Hamiltonian learning protocol

By replacing $\mathcal{A}_j^I$ by $\mathcal{A}_j^{I'}$, we obtain the single-copy product-state input Hamiltonian learning protocol. The total evolution time complexity is

with $\boldsymbol{\mu} = (\mu_1, \dots, \mu_M)$ and $|\mu_s| \leq 1$ for any $s = 1, \dots, M$.

The high-level strategy of our proof is to combine the learning tree framework [70] equipped with the martingale trick [52,53,55] and quantum Fisher information.

Here, we generalize any experimental settings with discrete quantum controls and sequential queries to the Hamiltonian as the following theoretical model. Given a Hamiltonian H with coefficient vector $\boldsymbol{\mu}$ defined in Eq. (1), the protocol performs multiple experiments as shown in Fig. 5 and measures at the end of each experiment. In each experiment (indexed by u), the protocol prepares an input state ρ_u (possibly with ancilla qubits) and queries the Hamiltonian m_u times with evolution time $\tau_1^u, \dots, \tau_{m_u}^u$. We can also apply discrete quantum control unitaries $C_1^u, \dots, C_{m_u-1}^u$ between every two neighboring queries. The protocol can be adaptive in the sense that it can dynamically decide how to prepare the input state, query the




FIG. 5. (a) Learning tree representation. Each node u represents an experiment. Starting from the root experiment r , the number of child nodes depends on the possible POVM measurements. The transition probabilities are determined by Eq. (G2). After N_{exp} experiments, one arrives at the leaves l of the learning tree. (b) In each node u , the learning model prepares an arbitrary state ρ_u , applies discrete control channels C_k^u , and queries real-time Hamiltonian evolutions with time τ_k^u multiple times. The protocols can also incorporate ancilla qubits (quantum memory).

real-time evolutions, perform quantum controls, and measure the final state based on the history of the previous experiments. We denote the summation of the evolution times for querying the Hamiltonian throughout *all* experiment to be the total evolution time T and the maximal number of discrete quantum controls in *one* experiment as $\mathcal{L}$.

Formally, we describe the protocol scheme in Fig. 5 as the learning tree model [55,70] as follows:

Definition G1 (Learning tree representation). Given a Hamiltonian $H(\mathbf{u})$ as defined in Eq. (G1) with M terms, a learning protocol using discrete control channels and queries to the Hamiltonian real-time evolutions as shown in Fig. 5 can be represented as a rooted tree $\mathcal{T}$ of depth N_{exp} (corresponding to N_{exp} measurements) with each node representing the measurement outcome history so far. In addition, the following conditions are satisfied:

- (1) We assign a probability $p^\mu(u)$ for any node u on the tree $\mathcal{T}$. The probability assigned to the root r is $p^\mu(r) = 1$.
- (2) At each nonleaf node u , we input the state ρ_u . By convexity, we can assume $\rho_u = |\psi_u\rangle\langle\psi_u|$ a pure state. We then query m_t rounds of real-time Hamiltonian evolution of time $\tau_1^u, \dots, \tau_{m_t}^u$, interleaved with m_t discrete control channels $\mathcal{C}_1^u, \dots, \mathcal{C}_{m_t}^u$ (the last one, $\mathcal{C}_{m_t}^u$, can be absorbed into the measurement as shown in the figure). We then perform a POVM $\{\mathcal{M}_{s_u}^u\}_{s_u}$, which by the fact that any POVM can be simulated by rank-1 POVM [70], can be assumed to be a rank-1 POVM as $\{w_{s_u}^u |\phi_{s_u}^u\rangle\langle\phi_{s_u}^u|\}_{s_u}$ with $\sum w_{s_u}^u = 2^n$, and get classical output s_u . The child node v corresponding to the classical outcome s_u of the node u is connected through the edge e_{u,s_u} . The probability associated with v is given by

$$p^\mu(v) \triangleq p^\mu(u) \cdot \text{tr}(w_{s_u}^u |\phi_{s_u}^u\rangle\langle\phi_{s_u}^u| \cdot \mathcal{U}_{m_u} \mathcal{C}_{m_u-1} \mathcal{U}_{m_u-1} \cdots \mathcal{U}_2 \mathcal{C}_1 \mathcal{U}_1 (|\psi_u\rangle\langle\psi_u|)), \quad (\text{G2})$$

where $\mathcal{U}_i = e^{-iH(\mu)\tau_i}$ is the unitary real-time evolution of $H(\mu)$ for time τ_i .

- (3) Each root-to-leaf path is of length N_{exp} . For a leaf node ℓ , $p^\mu(\ell)$ is the probability of reaching this leaf ℓ at the end of the learning protocol. The set of leaves is denoted as $\text{leaf}(\mathcal{T})$.

We consider a point-versus-point distinguishing task between two cases for a hyperparameter μ :

- (1) The Hamiltonian for the real-time evolution is exactly $H(\mu)$.
- (2) The Hamiltonian for the real-time evolution is actually $H(\mu + \delta\mu)$ with $\|\delta\mu\|_\infty = 3\epsilon$.

The goal is to distinguish which case is happening. It is straightforward to see that given an algorithm that can learn the Hamiltonian, we can solve this distinguishing problem.

Now, we consider how hard this distinguishing problem is. In the learning tree framework, the necessary condition for distinguishing between these two cases according to Le Cam's two-point method [51] is that the total variation distance between the probability distributions on the leaves of the learning trees for the two cases is large. Quantitatively,

$$\frac{1}{2} \sum_{\ell \in \text{leaf}(\mathcal{T})} |p^{\mu+\delta\mu}(\ell) - p^\mu(\ell)| = \Theta(1). \quad (\text{G3})$$

Using the martingale likelihood ratio argument developed by a series of works [52–54], the above lower bound can be converted to Lemma 7 of Ref. [55]:=. If there is a $\Delta > 0$ such that

$$\mathbb{E}_{s_u \sim p^\mu(s_u|u)} [(L_{\delta\mu}(u, s_u) - 1)^2] \leq \Delta, \quad (\text{G4})$$

where

$$L_{\delta\mu}(u, s_u) = \frac{p^{\mu+\delta\mu}(s_u|u)}{p^\mu(s_u|u)}, \quad (\text{G5})$$

then $N_{\text{exp}} \geq \Omega(1/\Delta)$. The physical intuition behind this result is that when the difference between the probability $p^\mu(s_u|u)$ and the perturbed probability $p^{\mu+\delta\mu}(s_u|u)$ is very small across all levels of the learning tree [Eq. (G4)], a very deep tree is required—i.e., $N_{\text{exp}} \geq \Omega(1/\Delta)$ —to ensure that the measurement distributions at the leaf nodes are distinct [Eq. (G3)].

Now, we focus on this quantity

$$\mathbb{E}_{s_u \sim p^\mu(s_u|u)} [(L_{\delta\mu}(u, s_u) - 1)^2] \quad (\text{G6})$$

for a specific u . Recall that $\|\delta\mu\|_\infty = 3\epsilon$, we thus have $\|\delta\mu\|_2 \leq 3\sqrt{M}\epsilon$. We can thus compute this quantity by extending it to the first-order term as

$$\begin{aligned} & \mathbb{E}_{s_u \sim p^\mu(s_u|u)} [(L_{\delta\mu}(u, s_u) - 1)^2] \\ &= \mathbb{E}_{s_u \sim p^\mu(s_u|u)} \left[\left(\frac{p^{\mu+\delta\mu}(s_u|u)}{p^\mu(s_u|u)} - 1 \right)^2 \right] \\ &\simeq \mathbb{E}_{s_u \sim p^\mu(s_u|u)} \left[\left(\frac{p^\mu(s_u|u) + \nabla_\mu p^\mu(s_u|u) \cdot \delta\mu}{p^\mu(s_u|u)} - 1 \right)^2 \right] \\ &= \mathbb{E}_{s_u \sim p^\mu(s_u|u)} \left[\left(\frac{\nabla_\mu p^\mu(s_u|u) \cdot \delta\mu}{p^\mu(s_u|u)} \right)^2 \right] \\ &= \delta\mu^\top I_{\mu,u}^{(C)} \delta\mu \\ &\leq \max_{\mathbf{v}: \|\mathbf{v}\|_2 \leq 3\sqrt{M}\epsilon} \mathbf{v}^\top I_{\mu,u}^{(C)} \mathbf{v}, \end{aligned} \quad (\text{G7})$$

where

$$\begin{aligned} [I_{\mu,u}^{(C)}]_{a,a'} &= \mathbb{E}_{s_u \sim p^\mu(s_u|u)} \left[\left(\frac{\partial p^\mu(s_u|u)}{\partial \lambda_a} \right) \left(\frac{\partial p^\mu(s_u|u)}{\partial \lambda_{a'}} \right) \right] \\ &= \mathbb{E}_{s_u \sim p^\mu(s_u|u)} \left[\left(\frac{\partial \ln p^\mu(s_u|u)}{\partial \lambda_a} \right) \left(\frac{\partial \ln p^\mu(s_u|u)}{\partial \lambda_{a'}} \right) \right] \end{aligned} \quad (\text{G8})$$

is the classical Fisher information matrix [48–50]. Note that quantum Fisher information $\mathbf{v}^\top I_{\mu,u}^{(Q)} \mathbf{v}$ is the supremum of classical Fisher information $\mathbf{v}^\top I_{\mu,u}^{(C)} \mathbf{v}$ over all possible measurements or observables [71]. By Lemma 17 of Ref. [47], assuming $H(\mu)$ can be diagonalized as $H(\mu) = W(\mu)^\dagger D(\mu) W(\mu)$ with unitary W and diagonal D , we have the quantum Fisher information satisfies

$$\begin{aligned} \max_{\mathbf{v}: \|\mathbf{v}\|_2 \leq \sqrt{M}\epsilon} \mathbf{v}^\top I_{\mu,u}^{(Q)} \mathbf{v} &\leq M\epsilon^2 \left(\min \{t_u \|\partial_\mu D\|_{\text{HS}} \right. \\ &\quad \left. + 2m_u \|\partial_\mu W\|_{\text{HS}}, t_u \|\partial_\mu H\|_{\text{HS}} \} \right)^2, \end{aligned} \quad (\text{G9})$$

where $t_u = \sum_{i=1}^{m_u} \tau_i^u$ and $\|\cdot\|_{\text{HS}}$ is the Hilbert Schmidt norm (spectral norm). By applying Eq. (G4), we have

$$N_{\text{exp}} = \Omega \left(\left(M\epsilon^2 \left(\min \{2m_u \|\partial_\mu W\|_{\text{HS}}, t_u \|\partial_\mu H\|_{\text{HS}}\} \right)^2 \right)^{-1} \right) \quad (\text{G10})$$

for any protocols that can solve this distinguishing problem with a high probability. The RHS is maximized when $t_u \leq 2m_u \|\partial_\mu W\|_{\text{HS}} / \|\partial_\mu H\|_{\text{HS}}$. Denote $\mathcal{L} = \max_u m_u$ as the maximal number of discrete quantum controls in one experiment and $T = \sum_{u=1}^{N_{\text{exp}}} t_u$ as the total evolution time, we have

$$2M\epsilon^2 \mathcal{L} T \|\partial_\mu W\|_{\text{HS}} \|\partial_\mu H\|_{\text{HS}} = \Theta(1), \quad (\text{G11})$$

which indicates that any protocol that can solve this distinguishing problem with a high probability requires $T = \Omega(\mathcal{L}^{-1}\epsilon^{-2})$ as claimed in Theorem G1.

APPENDIX H: DETAILS OF THE NUMERICAL SIMULATION

In this Appendix, we provide additional details on our numerical simulations. For structure learning, we perform 2000 measurement shots, while for coefficient learning, we use 1000 simulated measurements throughout. A key advantage of our Heisenberg-limited learning algorithm is that the number of required measurements does not scale with the target learning accuracy. This allows us to use a constant number of measurements even when estimating small coefficients. For short-time evolutions, we set the

evolution time to $\tau = 0.001$, and we find that the results are not particularly sensitive to the precise choice of τ .

For the first example with learning disordered XY model, we perform the hierarchical learning with $J = 1$. The total evolution time in structure learning is chosen as $T^{j=1} = 0.3$, and the learning precision in the coefficient learning is set to be $\epsilon = 0.005$.

For the second example, learning long-range interaction decay in a static Rydberg Hamiltonian, we perform hierarchical learning up to level $J = 5$. The physical Hamiltonian takes the form

$$\begin{aligned} H/\hbar &= \frac{\Omega}{2} \sum_l (|g_l\rangle\langle r_l| + |r_l\rangle\langle g_l|) - \Delta \sum_l |r_l\rangle\langle r_l| \\ &\quad + \sum_{j<l} V_{jl} |r_j\rangle\langle r_j| \otimes |r_l\rangle\langle r_l|, \end{aligned} \quad (\text{H1})$$

where $V_{jl} = C/|\vec{x}_j - \vec{x}_l|^6$ describes the van der Waals interaction between atoms j and l . We use $C = 862\,690 \times 2\pi$ MHz μm^6 , corresponding to the interaction strength for Rydberg states $70S_{1/2}$ of ^{87}Rb atoms. The atoms are equally spaced with a separation of $10\,\mu\text{m}$, and the drive parameters are set to $\Omega = 1.5/2\pi$ MHz and $\Delta = -4/2\pi$ MHz. Due to the $1/r^6$ decay of the van der Waals interaction, the interaction strength rapidly diminishes with distance. We placed all atoms in a chain separated equally with $d = 10\,\mu\text{m}$. For each level j in the hierarchical learning protocol, we set the evolution time T^j and learning precision ϵ^j as follows:

- (1) Level 1: $T^{j=1} = 0.08\,\mu\text{s}$ and $\epsilon^{j=1} = 0.005\,\text{rad}/\mu\text{s}$
- (2) Level 2: $T^{j=2} = 0.2\,\mu\text{s}$ and $\epsilon^{j=2} = 0.005\,\text{rad}/\mu\text{s}$
- (3) Level 3: $T^{j=3} = 6.0\,\mu\text{s}$ and $\epsilon^{j=3} = 0.001\,\text{rad}/\mu\text{s}$
- (4) Level 4: $T^{j=4} = 60.0\,\mu\text{s}$ and $\epsilon^{j=4} = 0.0001\,\text{rad}/\mu\text{s}$
- (5) Level 5: $T^{j=5} = 1200.0\,\mu\text{s}$ and $\epsilon^{j=5} = 0.00001\,\text{rad}/\mu\text{s}$

The Pauli terms learned at each level are summarized in Table I.

For the final example: learning the effective Hamiltonian generated by a time-dependent driven pulse, we consider a Rydberg atom array with the same interaction strength as before, but with an interatomic spacing of $d = 8.9\,\mu\text{m}$. The system is driven by a time-dependent laser pulse, as shown in Fig. 3(a). We do the Hamiltonian learning with the information that this pulse engineered a unitary that is close to

$$U \approx e^{-i\theta H_{\text{ZXZ}}}, \quad (\text{H2})$$

with $\theta = 0.1$ and $H_{\text{ZXZ}} = \sum_i Z_i X_{i+1} Z_{i+2}$. Usually, this information is available because engineered pulses are optimized to reach a certain target with the target being known. Thus, the unitary implemented by the pulse can be

TABLE I. Hierarchical learning of static Rydberg atom Hamiltonian.

j th	Pauli String	Value
1	IIIZZ	1.343 ± 0.036
1	IIZZII	1.331 ± 0.047
1	IZZIII	1.359 ± 0.044
1	ZZIIII	1.355 ± 0.025
1	IIIIIZ	0.618 ± 0.029
1	IIIIZI	-0.733 ± 0.043
1	IIIIII	0.750 ± 0.019
1	IZIIII	-0.728 ± 0.043
1	IIIXII	0.741 ± 0.026
1	IXIIII	0.746 ± 0.029
1	ZIIIII	0.624 ± 0.036
1	IIIXII	0.748 ± 0.033
1	IIIIIX	0.751 ± 0.013
1	IZZXXI	0.013 ± 0.059
2	IZIIII	-0.753 ± 0.001
3	IIZZII	0.0245 ± 0.0006
3	IZIZII	0.0211 ± 0.0003
3	ZIZIII	0.0212 ± 0.00007
3	IIZIIZ	0.0214 ± 0.0008
3	IZZXXI	-0.0130 ± 0.0006
3	IIIZZ	0.0123 ± 0.0008
3	IIIXII	0.0088 ± 0.0005
3	IZZIII	-0.0040 ± 0.0005
3	ZIIZII	0.0019 ± 0.0007
4	IZIIII	-0.00525 ± 0.00007
4	IXIIII	0.00390 ± 0.00008
4	IIIIIZ	0.00374 ± 0.00004
4	ZIIIII	-0.00223 ± 0.00001
4	IZIIIZ	0.00186 ± 0.00013
4	IIIIYI	0.00199 ± 0.000005
4	IIIIIX	-0.00149 ± 0.00006
4	IXIIZ	-0.000012 ± 0.00012
4	ZZIIII	0.00022 ± 0.00009
4	IYIIII	0.00006 ± 0.00011
4	IXIIX	-0.000013 ± 0.00011
4	ZXIIII	0.000007 ± 0.00011
4	IYIIZ	-0.000001 ± 0.00012
4	IYYII	-0.000012 ± 0.00014
5	ZIIIZ	0.000331 ± 0.000001

viewed as a short-time evolution under an unknown effective Hamiltonian that is close to H_{ZZ} . Since we lack prior knowledge about the exact form of this effective Hamiltonian, the ansatz-free learning protocol is particularly well suited for this task.

In the structure learning stage, we prepare Bell pairs, apply the time-dependent Hamiltonian to one half of each pair, and then perform Bell-basis measurements. We use 1000 measurement shots for structure learning. Once the structure is identified, we treat U as a short-time evolution with $\theta = 0.1$ under an unknown effective Hamiltonian, and proceed with coefficient learning using Hamiltonian

reshaping. For this step, we also use 1000 measurement shots, with the target learning accuracy set to $\epsilon = 0.01$.

- [1] A. J. Daley, Twenty-five years of analogue quantum simulation, *Nat. Rev. Phys.* **5**, 702 (2023).
- [2] H. Bernien, S. Schwartz, A. Keesling, H. Levine, A. Omran, H. Pichler, S. Choi, A. S. Zibrov, M. Endres, M. Greiner, V. Vuletić, and M. D. Lukin, Probing many-body dynamics on a 51-atom quantum simulator, *Nature* **551**, 579 (2017).
- [3] G. Semeghini, H. Levine, A. Keesling, S. Ebadi, T. T. Wang, D. Bluvstein, R. Verresen, H. Pichler, M. Kalinowski, R. Samajdar, A. Omran, S. Sachdev, A. Vishwanath, M. Greiner, V. Vuletić, and M. D. Lukin, Probing topological spin liquids on a programmable quantum simulator, *Science* **374**, 1242 (2021).
- [4] C. Ryan-Anderson, J. G. Bohnet, K. Lee, D. Gresh, A. Hankin, J. P. Gaebler, D. Francois, A. Chernoguzov, D. Lucchetti, N. C. Brown, T. M. Gatterman, S. K. Halit, K. Gilmore, J. A. Gerber, B. Neyenhuis, D. Hayes, and R. P. Stutz, Realization of real-time fault-tolerant quantum error correction, *Phys. Rev. X* **11**, 041058 (2021).
- [5] L. Postler, S. Heußen, I. Pogorelov, M. Rispler, T. Feldker, M. Meth, C. D. Marciniak, R. Stricker, M. Ringbauer, R. Blatt, P. Schindler, M. Müller, and T. Monz, Demonstration of fault-tolerant universal quantum gate operations, *Nature* **605**, 675 (2022).
- [6] D. Bluvstein *et al.*, Logical quantum processor based on reconfigurable atom arrays, *Nature* **626**, 58 (2024).
- [7] Google Quantum AI and Collaborators, Quantum error correction below the surface code threshold, *Nature* **638**, 920 (2024).
- [8] H. Zhou, C. Zhao, M. Cain, D. Bluvstein, C. Ducker, H.-Y. Hu, S.-T. Wang, A. Kubica, and M. D. Lukin, Algorithmic fault tolerance for fast quantum computing, *arXiv:2406.17653* [quant-ph].
- [9] J. Carrasco, A. Elben, C. Kokail, B. Kraus, and P. Zoller, Theoretical and experimental perspectives of quantum verification, *PRX Quantum* **2**, 010102 (2021).
- [10] D. Hangleiter, I. Roth, J. Fuksa, J. Eisert, and P. Roushan, Robustly learning the Hamiltonian dynamics of a superconducting quantum processor, *Nat. Commun.* **15**, 9595 (2024).
- [11] H.-Y. Huang, Y. Tong, D. Fang, and Y. Su, Learning many-body Hamiltonians with Heisenberg-limited scaling, *Phys. Rev. Lett.* **130**, 200403 (2023).
- [12] M. Ma, S. T. Flammia, J. Preskill, and Y. Tong, Learning k -body Hamiltonians via compressed sensing, *arXiv:2410.18928* [quant-ph].
- [13] H. Li, Y. Tong, T. Gefen, H. Ni, and L. Ying, Heisenberg-limited Hamiltonian learning for interacting bosons, *npj Quantum Inf.* **10**, 83 (2024).
- [14] H. Ni, H. Li, and L. Ying, Quantum Hamiltonian learning for the Fermi-Hubbard model, *Acta Appl. Math.* **191**, 2 (2024).
- [15] A. Bakshi, A. Liu, A. Moitra, and E. Tang, in *2024 IEEE 65th Annual Symposium on Foundations of Computer*

- Science (FOCS)* (IEEE, Chicago, Illinois, USA, 2024), pp. 1037–1050.
- [16] A. Mirani and P. Hayden, Learning interacting fermionic Hamiltonians at the Heisenberg limit, *Phys. Rev. A* **110**, 062421 (2024).
 - [17] D. Stilck França, L. A. Markovich, V. V. Dobrovitski, A. H. Werner, and J. Borregaard, Efficient and robust estimation of many-qubit Hamiltonians, *Nat. Commun.* **15**, 311 (2024).
 - [18] W. Yu, J. Sun, Z. Han, and X. Yuan, Robust and efficient Hamiltonian learning, *Quantum* **7**, 1045 (2023).
 - [19] A. Anshu, S. Arunachalam, T. Kuwahara, and M. Soleimanifar, Sample-efficient learning of interacting quantum systems, *Nat. Phys.* **17**, 931 (2021).
 - [20] C. Kokail, R. van Bijnen, A. Elben, B. Vermersch, and P. Zoller, Entanglement Hamiltonian tomography in quantum simulation, *Nat. Phys.* **17**, 936 (2021).
 - [21] T. Olsacher, T. Kraft, C. Kokail, B. Kraus, and P. Zoller, Hamiltonian and Liouvillian learning in weakly-dissipative quantum many-body systems, *Quantum Sci. Technol.* **10**, 015065 (2025).
 - [22] R. Ott, T. V. Zache, M. Prüfer, S. Erne, M. Tajik, H. Pichler, J. Schmiedmayer, and P. Zoller, Hamiltonian learning in quantum field theories, [arXiv:2401.01308](https://arxiv.org/abs/2401.01308) [cond-mat.quant-gas].
 - [23] M. C. Caro, Learning quantum processes and Hamiltonians via the Pauli transfer matrix, *ACM Trans. Quantum Comput.* **5**, 1 (2024).
 - [24] T. Möbus, A. Bluhm, M. C. Caro, A. H. Werner, and C. Rouzé, Dissipation-enabled Bosonic Hamiltonian learning via new information-propagation bounds, [arXiv:2307.15026](https://arxiv.org/abs/2307.15026).
 - [25] J. Haah, R. Kothari, and E. Tang, in *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)* (IEEE Computer Society Technical Committee on the Mathematical Foundations of Computing (TCMF), Denver, Colorado, USA, 2022), pp. 135–146.
 - [26] A. Bakshi, A. Liu, A. Moitra, and E. Tang, Learning quantum Hamiltonians at any temperature in polynomial time, [arXiv:2310.02243](https://arxiv.org/abs/2310.02243) [quant-ph].
 - [27] A. Gu, L. Cincio, and P. J. Coles, Practical Hamiltonian learning with unitary dynamics and Gibbs states, *Nat. Commun.* **15**, 312 (2024).
 - [28] X.-L. Qi and D. Ranard, Determining a local Hamiltonian from a single eigenstate, *Quantum* **3**, 159 (2019).
 - [29] Z. Li, L. Zou, and T. H. Hsieh, Hamiltonian tomography via quantum quench, *Phys. Rev. Lett.* **124**, 160502 (2020).
 - [30] T. J. Evans, R. Harper, and S. T. Flammia, Scalable Bayesian Hamiltonian learning, [arXiv:1912.07636](https://arxiv.org/abs/1912.07636) [quant-ph].
 - [31] A. Zhao, Learning the structure of any Hamiltonian from minimal assumptions, [arXiv:2410.21635](https://arxiv.org/abs/2410.21635) [quant-ph].
 - [32] S. Arunachalam, A. Dutt, and F. Escudero Gutiérrez, Testing and learning structured quantum Hamiltonians, [arXiv:2411.00082](https://arxiv.org/abs/2411.00082) [quant-ph].
 - [33] D. A. Lidar, Review of decoherence free subspaces, noiseless subsystems, and dynamical decoupling, [arXiv:1208.5791](https://arxiv.org/abs/1208.5791) [quant-ph].
 - [34] J. Choi, H. Zhou, H. S. Knowles, R. Landig, S. Choi, and M. D. Lukin, Robust dynamic Hamiltonian engineering of many-body spin systems, *Phys. Rev. X* **10**, 031002 (2020).
 - [35] S. Choi, N. Y. Yao, and M. D. Lukin, Dynamical engineering of interactions in qudit ensembles, *Phys. Rev. Lett.* **119**, 183603 (2017).
 - [36] B. Evert, Z. Gonzalez Izquierdo, J. Sud, H.-Y. Hu, S. Grabbe, E. G. Rieffel, M. J. Reagor, and Z. Wang, Syn-copated dynamical decoupling for suppressing crosstalk in quantum circuits, [arXiv:2403.07836](https://arxiv.org/abs/2403.07836) [quant-ph].
 - [37] S. Kimmel, G. H. Low, and T. J. Yoder, Robust calibration of a universal single-qubit gate set via robust phase estimation, *Phys. Rev. A* **92**, 062315 (2015).
 - [38] S. T. Flammia and R. O'Donnell, Pauli error estimation via population recovery, *Quantum* **5**, 549 (2021).
 - [39] J. J. Wallman and J. Emerson, Noise tailoring for scalable quantum computation via randomized compiling, *Phys. Rev. A* **94**, 052325 (2016).
 - [40] J. Emerson, M. Silva, O. Moussa, C. Ryan, M. Laforest, J. Baugh, D. G. Cory, and R. Laflamme, Symmetrized characterization of noisy quantum processes, *Science* **317**, 1893 (2007).
 - [41] S. Chen, Y. Liu, M. Otten, A. Seif, B. Fefferman, and L. Jiang, The learnability of Pauli noise, *Nat. Commun.* **14**, 52 (2023).
 - [42] E. van den Berg, Z. K. Mineev, A. Kandala, and K. Temme, Probabilistic error cancellation with sparse Pauli–Lindblad models on noisy quantum processors, *Nat. Phys.* **19**, 1116 (2023).
 - [43] H.-Y. Hu, A. Gu, S. Majumder, H. Ren, Y. Zhang, D. S. Wang, Y.-Z. You, Z. Mineev, S. F. Yelin, and A. Seif, Demonstration of robust and efficient quantum property learning with shallow shadows, [arXiv:2402.17911](https://arxiv.org/abs/2402.17911) [quant-ph].
 - [44] E. Campbell, Random compiler for fast Hamiltonian simulation, *Phys. Rev. Lett.* **123**, 070503 (2019).
 - [45] C.-F. Chen, H.-Y. Huang, R. Kueng, and J. A. Tropp, Concentration for random product formulas, *PRX Quantum* **2**, 040305 (2021).
 - [46] H.-Y. Hu, A. M. Gomez, L. Chen, A. Trowbridge, A. J. Goldschmidt, F. T. Chong, A. Jaffe, and S. F. Yelin, Expressivity of analog quantum simulators with global controls (2025), in preparation.
 - [47] A. Dutkiewicz, T. E. O'Brien, and T. Schuster, The advantage of quantum control in many-body Hamiltonian learning, *Quantum* **8**, 1537 (2024).
 - [48] W. K. Wootters, Statistical distance and Hilbert space, *Phys. Rev. D* **23**, 357 (1981).
 - [49] S. L. Braunstein and C. M. Caves, Statistical distance and the geometry of quantum states, *Phys. Rev. Lett.* **72**, 3439 (1994).
 - [50] S. L. Braunstein, C. M. Caves, and G. Milburn, Generalized uncertainty relations: Theory, examples, and Lorentz invariance, *Ann. Phys.* **247**, 135 (1996).
 - [51] B. Yu, Assouad, Fano, and Le Cam, in *Research Papers in Probability and Statistics* (Springer, 1997), <https://link.springer.com/content/pdf/10.1007/978-1-4612-1880-7.pdf#page=423.00>.

- [52] S. Chen, J. Cotler, H.-Y. Huang, and J. Li, The complexity of NISQ, *Nat. Commun.* **14**, 6001 (2023).
- [53] S. Chen and W. Gong, Efficient Pauli channel estimation with logarithmic quantum memory, [arXiv:2309.14326](#) [quant-ph].
- [54] C.-F. Chen and F. G. S. L. Brandão, Fast thermalization from the eigenstate thermalization hypothesis, [arXiv:2112.07646](#) [quant-ph].
- [55] S. Chen, W. Gong, and Q. Ye, in *2024 IEEE 65th Annual Symposium on Foundations of Computer Science (FOCS)* (IEEE Computer Society Technical Committee on the Mathematical Foundations of Computing (TCMF), Chicago, Illinois, USA, 2024), pp. 1086–1105.
- [56] J. Cardy and E. Tonni, Entanglement Hamiltonians in two-dimensional conformal field theory, *J. Stat. Mech.: Theory Exp.* **2016**, 123103 (2016).
- [57] M. Dalmonte, V. Eisler, M. Falconi, and B. Vermersch, Entanglement Hamiltonians: From field theory to lattice models and experiments, *Ann. Phys.* **534**, 2200064 (2022).
- [58] S. Murciano, V. Vitale, M. Dalmonte, and P. Calabrese, Negativity Hamiltonian: An operator characterization of mixed-state entanglement, *Phys. Rev. Lett.* **128**, 140502 (2022).
- [59] R. Travaglino, C. Rylands, and P. Calabrese, Quench dynamics of negativity Hamiltonians, [arXiv:2506.09561](#) [quant-ph].
- [60] X.-D. Xie, Z.-Y. Xue, and D.-B. Zhang, Determining the non-Hermitian parent Hamiltonian from a single eigenstate, *Phys. Rev. B* **110**, 235113 (2024).
- [61] S. Zhou, M. Zhang, J. Preskill, and L. Jiang, Achieving the Heisenberg limit in quantum metrology using quantum error correction, *Nat. Commun.* **9**, 78 (2018).
- [62] R. R. Allen, F. Machado, I. L. Chuang, H.-Y. Huang, and S. Choi, Quantum computing enhanced sensing, [arXiv:2501.07625](#) [quant-ph].
- [63] A. Seif, S. Chen, S. Majumder, H. Liao, D. S. Wang, M. Malekakhlagh, A. Javadi-Abhari, L. Jiang, and Z. K. Mineev, Entanglement-enhanced learning of quantum processes at scale, [arXiv:2408.03376](#) [quant-ph].
- [64] S. Chen, S. Zhou, A. Seif, and L. Jiang, Quantum advantages for Pauli channel estimation, *Phys. Rev. A* **105**, 032435 (2022).
- [65] M. Suzuki, Generalized Trotter’s formula and systematic approximants of exponential operators and inner derivations with applications to many-body problems, *Commun. Math. Phys.* **51**, 183 (1976).
- [66] M. Suzuki, On the convergence of exponential operators—the Zassenhaus formula, BCH formula and systematic approximants, *Commun. Math. Phys.* **57**, 193 (1977).
- [67] K. Nakaji, M. Bagherimehrab, and A. Aspuru-Guzik, High-order randomized compiler for Hamiltonian simulation, *PRX Quantum* **5**, 020330 (2024).
- [68] A. M. Childs, Y. Su, M. C. Tran, N. Wiebe, and S. Zhu, Theory of Trotter error with commutator scaling, *Phys. Rev. X* **11**, 011020 (2021).
- [69] By “nonadaptive” we mean that the choice of each t_j does not depend on the value of $X(t_{j'})$ or $Y(t_{j'})$ for any j' .
- [70] S. Chen, J. Cotler, H.-Y. Huang, and J. Li, in *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)* (IEEE Computer Society Technical Committee on the Mathematical Foundations of Computing (TCMF), Denver, Colorado, USA, 2022), pp. 574–585.
- [71] M. G. A. Paris, Quantum estimation for quantum technology, *Int. J. Quantum Inf.* **07**, 125 (2009).