Low-resource quantum energy gap estimation via randomization

Hugo Pages [1,2,*] Chusei Kiumi [3] Yuto Morohoshi [1] Bálint Koczor [4] and Kosuke Mitarai [1,3,†]

[1]*Graduate School of Engineering Science, The University of Osaka, 1-3 Machikaneyama, Toyonaka, Osaka 560-8531, Japan*

[2]*Télécom Physique Strasbourg, Strasbourg University, 300 Bd Sébastien Brant, Illkirch Graffenstaden 67412, France*

[3]*Center for Quantum Information and Quantum Biology, The University of Osaka, 1-2 Machikaneyama, Toyonaka, Osaka 560-0043, Japan*

[4]*Mathematical Institute, University of Oxford, Woodstock Road, Oxford OX2 6GG, United Kingdom*

(Received 11 February 2026; accepted 14 October 2026; published 18 November 2026)

Estimating the energy spectra of quantum many-body systems is a fundamental task in quantum physics, with applications ranging from chemistry to condensed matter. Algorithmic shadow spectroscopy is a recent method that leverages randomized measurements on time-evolved quantum states to extract spectral information. However, implementing accurate time evolution with low-depth circuits remains a key challenge for near-term quantum hardware. In this work, we propose a hybrid quantum-classical protocol that integrates time evolution via probabilistic angle interpolation (TE-PAI) into the shadow spectroscopy framework. TE-PAI enables the simulation of time evolution using shallow stochastic circuits while preserving unbiased estimates through quasiprobability sampling. We construct the combined estimator and derive its theoretical properties. Through numerical simulations, we demonstrate that our method accurately resolves energy gaps and exhibits enhanced robustness to gate noise compared to first-order Trotter-based shadow spectroscopy. We further validate the protocol experimentally on up to 20 qubits using IBM quantum hardware. This makes TE-PAI shadow spectroscopy a promising tool for spectral analysis on noisy intermediate-scale quantum devices.

DOI: 10.1103/mxgl-96mv

I. INTRODUCTION

Quantum many-body systems are governed by the principles of quantum mechanics. A fundamental aspect of understanding these systems is the characterization of their energy spectra. In particular, the differences between eigenvalues of the Hamiltonian, often referred to as energy gaps, play an important role in determining their physical properties [1]. For instance, in quantum chemistry, excitation energies are directly related to molecular optical responses and chemical reactivity [2]. However, accurately computing energy gaps remains a challenging task for classical computers, primarily because the dimension of the Hilbert space that must be taken into account grows exponentially with the system size [3]. Quantum computers provide a promising alternative, as they are capable of simulating quantum systems efficiently with polynomially scaling resources with respect to the system size [4]—and indeed the earliest examples of quantum advantage are expected to be in quantum simulation [5,6].

One of the earliest algorithms developed for estimating energy spectra on a quantum computer is quantum phase estimation (QPE) [7,8]. It is often considered the main method within the framework of fault-tolerant quantum computing, as it can extract eigenvalues of a Hamiltonian with high precision. However, QPE requires deep circuits with ancillary qubits and error correction, which are far beyond the capabilities of current noisy intermediate-scale quantum (NISQ) devices. The limited qubit counts, short coherence times, and high gate error rates in NISQ hardware make direct implementation of QPE infeasible in the near term, motivating the exploration of alternative approaches that are better suited to present-day hardware.

As such, researchers have proposed several quantum algorithms for estimating spectral properties, aiming for approaches that reduce circuit depth and resource requirements while remaining compatible with the constraints of current hardware. One promising method among them is algorithmic shadow spectroscopy [9]. In shadow spectroscopy, we apply the time-evolution operator of a target Hamiltonian to qubits for various durations and perform randomized measurements at each time point to generate classical shadows. By collecting and analyzing these time-dependent signals, one can reconstruct spectral information through classical postprocessing. A notable advantage of this method is that it requires only straightforward time evolution of the Hamiltonian without the need for ancilla qubits or controlled time evolution. Because of its simplicity, shadow spectroscopy offers better compatibility with near-term quantum devices compared to more demanding approaches such as quantum phase estimation [7,8]. Also, by leveraging the classical-shadow framework [10], it can estimate spectral properties with relatively few measurement shots, efficiently extracting relevant information

*Contact author: hugo.pages@etu.unistra.fr

†Contact author: mitarai.kosuke.es@osaka-u.ac.jp

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.

from limited data. This feature provides a potential advantage over variational approaches such as the variational quantum eigensolver, which often require extensive optimization and large numbers of measurements to achieve comparable accuracy. However, this method still relies on precise time evolution, which can make its implementation challenging depending on the capabilities of the quantum hardware. One approach to address this limitation is to apply quantum error mitigation techniques [11], such as probabilistic error cancellation (PEC), zero-noise extrapolation, or symmetry verification [12]. These approaches, however, require a detailed characterization of the noise and often lead to an exponential increase in the number of measurements.

To overcome this issue and further reduce the quantum resource requirements for spectral estimation, we propose an algorithm that integrates the recently developed technique of time evolution via probabilistic angle interpolation (TE-PAI) into the framework of algorithmic shadow spectroscopy [13]. TE-PAI is a stochastic method for implementing time evolution under a given Hamiltonian based on Trotter decomposition. Instead of applying all gates in the Trotterized sequence, TE-PAI probabilistically samples a subset of rotation gates according to carefully chosen quasiprobability distributions [13–15]. This approach allows us to reproduce the effect of a long and accurate Trotter sequence using only a small number of quantum gates, significantly reducing circuit depth while preserving fidelity. Furthermore, TE-PAI requires only Pauli rotation gates with fixed angles of Δ and π, followed by measurement. This approach, similar to shadow spectroscopy, does not rely on ancillary qubits or controlled time evolution. The simplicity and fundamental similarity of these randomized techniques make them naturally compatible and well aligned for integration. Although there is a trade-off between the measurement overhead and the number of gates, it is well understood and can be controlled explicitly [13]: Tuning the single parameter Δ allows a user to increase or decrease the expected circuit depth at the cost of increasing the measurement overhead. Thus, the significant advantage of shadow spectroscopy—efficiently extracting spectral information with fewer shots—alleviates the increased shot requirement of TE-PAI.

By combining TE-PAI with algorithmic shadow spectroscopy, we construct an effective protocol for estimating energy spectra. Both TE-PAI and classical shadows involve sampling from randomized procedures, that is, one over gate sequences and the other over measurement bases. A key theoretical contribution of this work is the formulation of the combined estimator, which involves a nested quasiprobability sampling structure. This analysis ensures that the full protocol yields unbiased estimators of spectral quantities and clarifies the statistical properties of the resulting estimates.

Since our algorithm is theoretically shown to offer a significant reduction in quantum resource requirements, we further highlight its practicality through extensive numerical simulations, which form the main contribution of this work, and demonstrative experiments on existing quantum hardware. Thanks to the simplicity of its circuit construction, our algorithm has the potential to serve as a key enabler of practical quantum advantage for NISQ and early fault-tolerant quantum computers (early-FTQC) [16].

While Ref. [13] focused on the unique advantages of TE-PAI for early fault-tolerant machines, we demonstrate a direct implementation of our algorithm on real quantum hardware through the simulation of the transverse Ising model. We benchmark its performance against conventional algorithmic shadow spectroscopy on a system of 20 qubits. Our results confirm the practical viability of our approach on contemporary noisy quantum devices, highlighting its potential for simulating many-body quantum systems within the current NISQ era.

This article is organized as follows. Section II introduces the theoretical background. In Sec. III, we present the TE-PAI shadow spectroscopy protocol and provide a formal analysis of its unbiasedness and variance properties. Section IV reports numerical simulations, and the following subsection reports experimental results on quantum hardware, benchmarking the proposed method against standard Trotter-based shadow spectroscopy. Finally, in Sec. V, we conclude with a discussion of the results, practical implications, and future research directions. Additional technical details and proofs are provided in the Appendixes.

II. PRELIMINARIES

In this section, we introduce the necessary background for understanding our approach. Throughout this paper, we write boldface for quantum channels acting on density matrices. For a unitary U, the associated channel $\boldsymbol{U}$ has the action $\boldsymbol{U}(\rho) = U\rho U^\dagger$ for any density matrix ρ. For a $d \times d$ matrix $\rho = \sum_{i,j} \rho_{ij} |i\rangle \langle j|$, we also introduce a vectorized notation as $|\rho\rangle\rangle := \sum_{i,j=1} \rho_{ij} |i\rangle \otimes |j\rangle$. With this convention, a channel corresponding to a unitary U can be written as $|\boldsymbol{U}(\rho)\rangle\rangle = (U \otimes U^*)|\rho\rangle\rangle$. We also introduce the Hilbert-Schmidt inner product between two operators A and B as $\langle\langle A|B\rangle\rangle := \mathrm{Tr}(A^\dagger B)$ in accordance with the Liouville space representation [17].

Our central goal is to enable the estimation of transition energy gaps from time-evolved classical-shadow data using resource-efficient circuits optimized for the constraints of NISQ and early-FTQC quantum hardware.

A. Energy gap estimation

Consider an N-qubit Hamiltonian $H = \sum_{j=1}^{J} h_j P_j$, where each $P_j \in \{I, X, Y, Z\}^{\otimes N}$ is a Pauli string. The time-evolution operator generated by H for a duration t is denoted by $U_{H,t} := e^{-iHt}$. Letting $\{E_\alpha\}$ and $\{|E_\alpha\rangle\}$ denote the eigenvalues and eigenstates of H, respectively, we define the transition energy gap between levels α and β as

$$\Delta E_{\alpha\beta} := |E_\alpha - E_\beta|.$$

This definition includes the spectral gap as a special case. When the eigenvalues are ordered as $E_0 \leqslant E_1 \leqslant \cdots$, the spectral gap usually refers to $E_1 - E_0$, which plays an important role in physical phenomena such as phase transitions, thermalization timescales, and adiabatic evolution rates. In this work, we consider the more general problem of estimating transition energy gaps $\Delta E_{\alpha\beta}$, which appear as frequencies in dynamical signals. Which transition gaps are visible depends on the initial state and on the measured observable.

In practice, quantum algorithms for gap estimation rely on extracting dynamical information from a time-evolved quantum state. Suppose a system with an initial state $|\psi_{\text{init}}\rangle = \sum_\alpha c_\alpha |E_\alpha\rangle$ evolves under H for time t as $|\psi(t)\rangle = e^{-iHt}|\psi_{\text{init}}\rangle = \sum_\alpha c_\alpha e^{-iE_\alpha t}|E_\alpha\rangle$. The expectation value of an observable O at time t becomes

$$\begin{aligned} S(t) &:= \langle\psi(t)|\, O\, |\psi(t)\rangle \\ &= \sum_{\alpha,\beta} c_\alpha^* c_\beta e^{i(E_\alpha - E_\beta)t} \langle E_\alpha|\, O\, |E_\beta\rangle . \end{aligned} \tag{1}$$

The oscillatory terms in $S(t)$ have frequencies equal to the energy differences $\Delta E_{\alpha\beta}$, and thus a Fourier transform of the time signal reveals peaks at these gaps. The intensity of each peak is determined by the spectral weight $c_\alpha^* c_\beta \langle E_\alpha|O|E_\beta\rangle$. Therefore, a prominent peak need not correspond to an adjacent-level spacing, but rather to a transition with large weight for the chosen initial state and observable.

Exact determination of $\Delta E_{\alpha\beta}$ requires sufficiently precise and long time evolution to achieve the desired frequency resolution. On NISQ devices, however, implementing long and accurate time evolutions is challenging due to limited coherence times and gate imperfections.

B. Time evolution

From Eq. (1), efficiently simulating the time-evolved state $|\rho(t)\rangle\rangle := |\boldsymbol{U}_{H,t}(\rho_{\text{init}})\rangle\rangle$ and estimating expectation values of many observables O denoted as

$$S(t) = \langle\langle O|\rho(t)\rangle\rangle$$

are key components of energy gap estimation algorithms in quantum computing. Since our target is NISQ and early-FTQC devices, we focus on methods that can implement time evolution with low quantum resource requirements.

One of the most straightforward approaches is Trotterization [4,18,19], which approximates $U_{H,t}$ as a product of exponentials of individual terms in the Hamiltonian. Let K denote the number of Trotter steps and define the step size $\delta t := t/K$. The K-step first-order Trotter circuit is denoted by $U_{H,t}^{(K)}$ and is given by

$$U_{H,t}^{(K)} := \left[\prod_{j=1}^{J} e^{-ih_j P_j \delta t}\right]^K = \left[\prod_{j=1}^{J} R_{P_j,\theta_j}\right]^K, \tag{2}$$

where $\theta_j = 2h_j\delta t$ and $R_{P,\theta} := \exp(-i\theta P/2)$ is a Pauli rotation gate. The expectation value obtained by evolving ρ_{init} with the K-step Trotter circuit $\langle O\rangle_t^{(K)} := \langle\langle O|\boldsymbol{U}_{H,t}^{(K)}(\rho_{\text{init}})\rangle\rangle$ approximates the true expectation value as $\epsilon_T := |S(t) - \langle O\rangle_t^{(K)}| \leqslant \mathcal{O}(t^2/K)$. Also, the sample complexity scaling as $\mathcal{O}(\epsilon_s^{-2})$ is required to estimate $\langle O\rangle_t^{(K)}$ within an additive statistical error ϵ_s.

We can further reduce the circuit depth required for time evolution by employing the recently proposed method of TE-PAI [13,14]. This method is based on quasiprobability sampling [12,20], which we summarize in Appendix B.

C. Time-evolution by probabilistic angle interpolation

This section briefly reviews the result on TE-PAI method proposed in Ref. [13], including the asymptotic analysis on expected gate count and variance. TE-PAI method applies probabilistic angle interpolation (PAI) [14] to Pauli rotation gates.

For any $\theta \in (-\pi, \pi)$ and a fixed angle $\Delta \in [|\theta|, \pi)$, set $\phi = \text{sign}(\theta)\Delta$. A superoperator representation of Pauli rotation $R_{P,\theta} = \exp(-i\theta P/2)$ admits the decomposition

$$\boldsymbol{R}_{P,\theta} = a_1(\theta)\boldsymbol{R}_{P,1} + a_2(\theta)\boldsymbol{R}_{P,2} + a_3(\theta)\boldsymbol{R}_{P,3}, \tag{3}$$

where $\boldsymbol{R}_{P,1}, \boldsymbol{R}_{P,2}, \boldsymbol{R}_{P,3}$ are superoperator representations of Pauli rotation gates $I, R_{P,\phi}, R_{P,\pi}$ and the analytic expression of the coefficients is explicitly given as follows:

$$\begin{aligned} a_1(\theta) &= \csc\left(\frac{\Delta}{2}\right)\cos\left(\frac{|\theta|}{2}\right)\sin\left(\frac{\Delta}{2} - \frac{|\theta|}{2}\right), \\ a_2(\theta) &= \csc(\Delta)\sin(|\theta|), \\ a_3(\theta) &= -\sec\left(\frac{\Delta}{2}\right)\sin\left(\frac{|\theta|}{2}\right)\sin\left(\frac{\Delta}{2} - \frac{|\theta|}{2}\right). \end{aligned}$$

Let $\gamma(\theta) := |a_1(\theta)| + |a_2(\theta)| + |a_3(\theta)|$ be the normalization factor. We define the random rotation channel outputting one of these three channels as $\boldsymbol{R}_{P,l}$, where $l \in \{1, 2, 3\}$ is sampled with probabilities $p_l(\theta) := |a_l(\theta)|/\gamma(\theta)$. Fixing Δ and applying the random-rotation construction independently to each Pauli-rotation gate R_{P_j,θ_j} in the circuit $U_{H,t}^{(K)}$ in Eq. (2) yields a random TE-PAI circuit

$$\hat{\boldsymbol{U}}_{l,t}^{(K)} := \prod_{k=1}^{K}\prod_{j=1}^{J} \boldsymbol{R}_{P_j,l(j,k)}, \tag{4}$$

along with an accumulated classical weight $\Gamma_{\boldsymbol{l}} := \Gamma \prod_{k=1}^{K}\prod_{j=1}^{J} \text{sign}(a_{l(j,k)})$, where $\Gamma := (\prod_{j=1}^{J}\gamma(\theta_j))^K$. Here, $l(j,k) \in \{1,2,3\}$ is sampled independently according to probability $p_l(\theta_j)$ and by definition, $\Gamma_{\boldsymbol{l}}^2 = \Gamma^2$.

Since this is quasiprobability decomposition by construction, the following random map is an unbiased estimator of the ideal rotation channel $\boldsymbol{R}_{P,\theta}$:

$$\mathbb{E}[\gamma(\theta)\,\text{sign}(a_l(\theta))\boldsymbol{R}_{P,l}] = \boldsymbol{R}_{P,\theta}.$$

Subsequently, the TE-PAI channel [Eq. (4)] along with classical weight form an unbiased estimator of the Trotterized time-evolution channel as

$$\mathbb{E}\left[\Gamma_{\boldsymbol{l}}\hat{\boldsymbol{U}}_{l,t}^{(K)}\right] = \boldsymbol{U}_{H,t}^{(K)}.$$

Let $\hat{o}_{\boldsymbol{l}}$ be a single-shot measurement estimator of the observable O obtained by executing $\hat{\boldsymbol{U}}_{l,t}^{(K)}$ on ρ_{init}. The multiplication of $\Gamma_{\boldsymbol{l}}$ in postprocessing yields an unbiased estimator of the original Trotter expectation value:

$$\mathbb{E}[\Gamma_{\boldsymbol{l}}\hat{o}_{\boldsymbol{l}}] = \langle O\rangle_t^{(K)} \approx S(t), \tag{5}$$

which approximates the true expectation value for sufficiently large K.

In implementation, each gate R_{P_j,θ_j} is replaced by one of $\{I, R_{P_j,\phi}, R_{P_j,\pi}\}$ and sampling I corresponds to skipping the gate.

As K increases, the rotation angle θ_j in each Pauli rotation becomes smaller, so the probability of sampling the

identity gate approaches 1. This mechanism probabilistically skips most operations, yielding a sampled circuit with depth $\nu_I \ll JK$, much shallower than the original Trotter sequence. Consequently, the expected number of nonidentity gates $\mathbb{E}[\nu_I]$ scales linearly with t and converges to a constant in the large-K limit:

$$\mathbb{E}[\nu_I] = \csc(\Delta)\,(3 - \cos\Delta)\|H\|_1 t + O(1/K),$$

where $\|H\|_1 := \sum_j |h_j|$ is the L^1 norm of the Hamiltonian coefficients. In contrast, the number of gates in the corresponding Trotter circuit scales as $O(t^2)$ and diverges as $K \to \infty$.

This depth reduction comes at the cost of a statistical sampling overhead, captured by the accumulated weight Γ, which leads to a sample complexity scaling of $\mathcal{O}(\Gamma^2 \epsilon_s^{-2})$. Although the overhead factor Γ^2 scales exponentially with the evolution time t, it converges to a constant in the large-K limit:

$$\Gamma^2 = \exp\left[4t\|H\|_1 \tan\left(\frac{\Delta}{2}\right)\right] + O(1/K).$$

The parameter Δ therefore provides a configurable trade-off between circuit depth and statistical sampling cost, allowing the overhead to remain manageable.

The advantage of TE-PAI lies in its tunable parameter Δ, which explicitly governs the trade-off between circuit depth and sampling overhead. In the large-K limit, increasing Δ reduces the expected circuit depth but increases the sampling overhead; decreasing Δ has the opposite effect, lowering the sampling cost at the expense of deeper circuits. As remarked in Ref. [13], by setting Δ to a sufficiently small value defined by $\Delta = 2\arctan(\frac{Q}{4\|H\|_1 t})$, TE-PAI achieves a constant measurement overhead $\exp(Q)$ while the gate count scales as $\mathcal{O}(\frac{\|H\|_1^2 t^2}{Q})$, which is comparable to the standard Trotter approach.

Furthermore, Ref. [13] demonstrated that TE-PAI is highly advantageous for the early-FTQC regime. Since the only non-Clifford operation required is $R_Z(\Delta)$, the protocol offers the potential to significantly reduce the number of T gates needed to execute a single circuit instance.

Higher-order Trotter-Suzuki formulas improve the product-formula error scaling. In particular, symmetric second-order Trotterization achieves this improvement without increasing the gate count relative to two first-order steps [18,19,21]. As discussed in Appendix C, TE-PAI can also be constructed from the second-order formula and inherits its improved Trotter error. Since our focus is the trade-off between gate count and sampling cost for a given product formula, we use the first-order formula as a simple common baseline for our comparisons.

D. Shadow spectroscopy

While the sampling overhead of TE-PAI is manageable, it becomes a critical bottleneck when scaling to a large number of observables. To circumvent this, we integrate the classical-shadow framework [10]; by employing randomized measurements, this approach allows us to estimate exponentially many observables simultaneously with only logarithmic sample scaling. In particular, we utilize the Pauli-basis variant, which is well suited for near-term hardware and enables the efficient reconstruction of local Pauli observables.

Given a quantum state ρ, we apply single-qubit Clifford unitaries $U_1, \ldots, U_N$ independently and uniformly at random, and then measure in the computational basis to obtain a bitstring $b = b_1 \cdots b_N$. From this outcome, we construct a snapshot

$$|\hat{\rho}\rangle\rangle = \bigotimes_{n=1}^{N} \left(3\,\boldsymbol{U}_n^{-1}|b_n\rangle\rangle - |I\rangle\rangle\right). \tag{6}$$

It follows that $\mathbb{E}[|\hat{\rho}\rangle\rangle] = |\rho\rangle\rangle$, so for any observable O the quantity $\langle\langle O|\hat{\rho}\rangle\rangle$ is an unbiased estimator of $\mathrm{Tr}(O\rho)$. By repeating the procedure many times, one obtains snapshots that allow simultaneous estimation of a large set of observables with rigorous sample complexity guarantees [10]. The key notion for analyzing the sample complexity is the shadow norm, which quantifies the variance of the estimator for a given observable. We define the squared shadow norm $\|O\|_{\mathrm{sh}}^2$ as the maximum second moment of the estimator over all possible states:

$$\|O\|_{\mathrm{sh}}^2 := \max_{\rho} \mathbb{E}[\langle\langle O|\hat{\rho}\rangle\rangle^2], \tag{7}$$

where the expectation $\mathbb{E}$ is taken over the random unitary transformations and measurement outcomes inherent to the shadow protocol.

For the specific case of Pauli-basis measurements and observables that are q-local Pauli strings, the squared shadow norm is given by 3^q [10]. The same collection of classical-shadow snapshots can be reused to estimate all observables in the target set $\{O_i\}_{i=1}^{N_o}$. Hence, the sample complexity is controlled by the shadow norm for a single observable, while requiring all N_o estimates to be accurate simultaneously introduces only a union-bound factor logarithmic in N_o [9,10]. The overall sample complexity is therefore

$$\mathcal{O}(3^q \epsilon^{-2} \ln N_o).$$

Finally, we employ the framework of algorithmic shadow spectroscopy [9]. We consider a sequence of time points $\{t_n\}_{n=1}^{N_t}$. Using the classical-shadow snapshots $\{\hat{\rho}(t_n)\}$ given by Eq. (6) from the state $\rho(t_n)$ at each instance and a set of target observables $\{O_i\}_{i=1}^{N_o}$, we construct the estimated time series $\hat{S}_i(t_n) = \langle\langle O_i|\hat{\rho}(t_n)\rangle\rangle$. This procedure yields a total of $N_o N_t$ data points. Since the above simultaneous-estimation guarantee applies at each time point, the total sample complexity scales as $\mathcal{O}(3^q N_t \epsilon^{-2} \ln N_o)$ to estimate all N_o time series. Following standardization and statistical filtering, we compute the cross-correlations of these signals and perform a Fourier transform; the resulting spectral peaks reveal the transition energy gaps of the Hamiltonian. A detailed description of the process is provided in Appendix D.

III. TE-PAI SHADOW SPECTROSCOPY

The central contribution of this work is a resource-efficient algorithm that integrates TE-PAI with algorithmic shadow spectroscopy. Explicitly designed for current and near-future quantum hardware, our protocol features a highly configurable trade-off between circuit depth and sampling overhead, minimizing quantum resource requirements while maintaining the flexibility to adapt to specific hardware constraints.

A. Our protocol

The following is the overview of the full procedure of TE-PAI shadow spectroscopy:

(1) *Step 1: Generate TE-PAI circuits.* Set the total Trotter steps as $K_{\text{step}} = K \times N_t$ with positive integer K and the number of time points N_t, and decompose the time-evolution operator e^{-iHt} into Trotter circuit $U_{H,t_s}^{(sK)}$ for each time point $t_s = s\frac{t}{N_t}$. For each gate in the sequence, probabilistically replace it with a fixed-angle rotation $\{0, \Delta, \pi\}$ to generate $M_{\text{TE-PAI}}$ TE-PAI circuits along with classical weights $(\Gamma_{\boldsymbol{l}_m,t_s}, \hat{U}_{\boldsymbol{l}_m,t_s}^{(sK)})$ according to Eq. (4).

(2) *Step 2: Obtaining shadow snapshots.* Apply random single-qubit Clifford unitaries $\{U_n\}_{n=1}^N$ to the state $\hat{U}_{\boldsymbol{l}_m,t_s}^{(sK)}(\rho_{\text{init}})$, and then measure in the computational basis to obtain a bitstring $b = b_1 \cdots b_N$. We then classically store the snapshot as $(\Gamma_{\boldsymbol{l},t_s}, \{U_n\}_{n=1}^N, \{b_n\}_{n=1}^N)$.

(3) *Step 3: Classical signal processing.* For N_o observables $\{O_i\}_{i=1}^{N_o}$, compute estimates of time series signals $S_i(s) = \langle\langle O_i|\rho(t_s)\rangle\rangle$ from estimator $\Gamma_{\boldsymbol{l}_m,t_s}\langle\langle O_i|\hat{\rho}_{\boldsymbol{l}_m}(t_s)\rangle\rangle$, which is constructed from the stored data according to Eq. (6). These time series are then processed, as described in Appendix D, to extract the energy spectrum of the Hamiltonian.

The detailed pseudocode for Steps 1 and 2 is provided in Algorithm 1. In our actual implementation, the process for generating these snapshots slightly differs from a simple repetition of the fundamental algorithm described in Algorithm 1. Specifically, our implementation generates one time-evolution circuit via the TE-PAI process and reuses it with N_s different randomized measurement bases to obtain N_s snapshots for each TE-PAI sampled circuit. Consequently, if we generate $M_{\text{TE-PAI}}$ distinct time-evolution circuits via TE-PAI sampling, the total number of unique quantum circuits ultimately executed is $M_{\text{TE-PAI}} \times N_s$. The detailed pseudocode for this implementation is provided in Appendix E as Algorithm 2.

This modification allows us to control the total number of circuit executions directly, while independently varying $M_{\text{TE-PAI}}$ and N_s. We show in the next subsection that this modified implementation preserves the unbiased nature of the estimator.

ALGORITHM 1. Obtaining a shadow snapshot of $e^{-iHt}\rho_{\text{init}}e^{iHt}$ using TE-PAI.

Input: N-qubit Hamiltonian $H = \sum_{j=1}^J h_j P_j$, initial state ρ_{init}, total time t, number of time steps K_{steps}, angle Δ
Output: Tuple $(\Gamma, \{U_n\}_{n=1}^N, \{b_n\}_{n=1}^N)$, where $\Gamma \in \mathbb{R}$ is the postprocessing weight, U_n is the local Clifford gate applied to qubit n, and $b_n \in \{0, 1\}$. The shadow is constructed as $\Gamma \bigotimes_{n=1}^N [3\boldsymbol{U}_n^{-1}|b_n\rangle\rangle - |I\rangle\rangle]$.

Set $\delta = t/K_{\text{steps}}$
Initialize $\Gamma \leftarrow 1$
Initialize an empty quantum circuit $\mathcal{C}$
for $k = 1$ to K_{steps}
  for $j = 1$ to J
    Compute $\theta = 2h_j\delta$
    Compute a_1, a_2, a_3 according to Eq. (3)
    Set $\gamma = |a_1| + |a_2| + |a_3|$, sample $d \in \{1, 2, 3\}$ with probability $|a_d|/\gamma$
    Update $\Gamma \leftarrow \Gamma \times \gamma \times \text{sign}(a_d)$
    Append gate $R_{P_j,\phi}$ [$\phi = \text{sign}(\theta)\Delta$ if $d = 2$] or $R_{P_j,\pi}$ ($d = 3$) to $\mathcal{C}$; if $d = 1$, append nothing
  end for
end for
for $n = 1$ to N
  Sample a Clifford gate U_n uniformly at random and add to qubit n
end for
Apply $\mathcal{C}$ to ρ_{init}, measure in the computational basis to obtain $b = (b_1, \ldots, b_N)$
return Γ, $\{U_n\}$, $\{b_n\}$

B. Formal analysis

In this subsection, we analyze the statistical properties of the estimator obtained from TE-PAI shadow spectroscopy. For generality, we state the result in a form that applies not only to TE-PAI but also to any quasiprobability sampling method.

We use the notation from Appendix B, where we consider a sequence of channels $\boldsymbol{C} = \boldsymbol{C}_K \cdots \boldsymbol{C}_1$, with each $\boldsymbol{C}_k$ decomposed as $\boldsymbol{C}_k = \sum_{d=1}^D a_{k,d}\boldsymbol{C}_{k,d}$. Let $l_k \in \{1, \ldots, D\}$ be a random variable drawn according to $p_d = |a_{k,d}|/\gamma(\boldsymbol{C}_k)$, where $\gamma(\boldsymbol{C}_k) = \sum_d |a_{k,d}|$. Define the random channel $\hat{\boldsymbol{C}}_{\boldsymbol{l}} := \boldsymbol{C}_{K,l_K} \cdots \boldsymbol{C}_{2,l_2}\boldsymbol{C}_{1,l_1}$ along with an accumulated classical weight $\Gamma_{\boldsymbol{l}} := \Gamma \prod_{k=1}^K \text{sign}(a_{k,l_k})$, where $\Gamma := \prod_{k=1}^K \gamma(\boldsymbol{C}_k)$. We generate M independent samples of classical weight and circuit denoted as $\Gamma_{\boldsymbol{l}_m}$, $\hat{\boldsymbol{C}}_{\boldsymbol{l}_m}$, with $m \in \{1, \ldots, M\}$. For each sampled channel applied on the initial state ρ, we obtain N_s independent classical-shadow snapshots $\hat{\rho}_{m,s}$ ($s \in \{1, \ldots, N_s\}$) according to Eq. (6). We define the estimator $o_{m,s} := \Gamma_{\boldsymbol{l}_m}\langle\langle O|\hat{\rho}_{m,s}\rangle\rangle$ and the overall estimator as

$$\langle\hat{O}\rangle := \frac{1}{MN_s}\sum_{m=1}^{M}\sum_{s=1}^{N_s} o_{m,s}. \tag{8}$$

Then, the following theorem shows that the estimator $\langle\hat{O}\rangle$ is an unbiased estimator of the true expectation value $\langle\langle O|\boldsymbol{C}(\rho)\rangle\rangle$. For the specific case of TE-PAI, $\boldsymbol{C}$ corresponds to the Trotterized time-evolution channel $\boldsymbol{U}_{H,t}^{(K)}$ and the true expectation value is $\langle O\rangle_t^{(K)}$.

Theorem 1 (Unbiasedness). The estimator defined in Eq. (8) is an unbiased estimator of the true expectation value:

$$\mathbb{E}\big[\langle\hat{O}\rangle\big] = \langle\langle O|\boldsymbol{C}(\rho)\rangle\rangle,$$

where the expectation $\mathbb{E}[\cdot]$ is taken over both the quasiprobability sampling and the randomized measurements for classical shadow.

The detailed proof is provided in Appendix E 1. This demonstrates that the estimator remains unbiased regardless of M and N_s. The convergence to this true value is guaranteed by the law of large numbers.

The variance of the estimator is dominated by the statistical fluctuations inherent in the quasiprobability sampling. Increasing N_s reduces the variance from the shadow measurements for a fixed evolution circuit, but it does not reduce the variance arising from the choice of different evolution circuits. The latter is suppressed only by increasing the number of quasiprobability samples, M. Therefore, by increasing M, the estimator is guaranteed to converge to the true expectation value.

Theorem 2 (Variance of the overall estimator). The variance of this estimator is upper bounded as

$$\mathrm{Var}[\langle\hat{O}\rangle] \leqslant \Gamma^2\left(\frac{\|O\|_{\mathrm{sh}}^2 - \|O\|^2}{MN_s} + \frac{\|O\|^2}{M}\right).$$

For the specific case of Pauli-basis measurements and normalized q-local Pauli observables, this becomes

$$\mathrm{Var}[\langle\hat{O}\rangle] \leqslant \Gamma^2\left(\frac{3^q - 1}{MN_s} + \frac{1}{M}\right).$$

The detailed proof is provided in Appendix E 2.

By fixing the total number of circuit executions $N_{\mathrm{total}} := MN_s$, we can further simplify the variance bound as

$$\mathrm{Var}[\langle\hat{O}\rangle] \leqslant \frac{\Gamma^2}{N_{\mathrm{total}}}(3^q - 1 + N_s).$$

This bound reveals a favorable trade-off for experimental implementation: As long as the number of shots per circuit satisfies $N_s \lesssim 3^q$, the statistical penalty for reusing the same circuit instance is small. Consequently, we can reuse a single compiled TE-PAI circuit instance up to $N_s \ll 3^q$ times with negligible impact on the overall estimation error. This allows for significant amortization of circuit compilation and loading costs without compromising statistical precision. We numerically verify this behavior in Sec. IV A, demonstrating that the resulting spectrum depends mainly on the total shot budget.

C. Resource comparison with prior approaches

Energy gap estimation can be performed by several established quantum algorithms, but many of them require heavier quantum resources or stronger assumptions than the setting considered here. Quantum phase estimation can estimate eigenphases with accuracy δ, but it requires coherent controlled time evolution, ancilla registers, and long circuits whose size scales at least as $O(1/\delta)$ [7,8]. We therefore focus on ancilla-free and control-free protocols. Variational and subspace-based approaches estimate gaps by preparing approximate low-energy states and measuring energies or effective Hamiltonian matrix elements, but their accuracy depends on the ansatz, state-preparation quality, and classical optimization, which may suffer from barren plateaus [22–24]. Another ancilla-free route is real-time spectroscopy, where one extracts gaps from the oscillation frequencies of time-dependent observables or correlation functions [25,26]. Related ancilla-free and control-free approaches to energy and energy gap estimation have also been developed [27,28]. Our work builds on shadow spectroscopy in this real-time framework, using randomized classical-shadow measurements to estimate many observables over many time points [9,10], and focuses on reducing the quantum cost of the required time evolution by replacing deterministic first-order Trotter circuits with TE-PAI sampled circuits.

Next, we compare our protocol with standard shadow spectroscopy combined with first-order Trotterized time evolution. Let N_t be the number of time points, N_o the number of observables, q the locality of each observable, and ϵ the target additive accuracy for estimating each time-dependent observable. Standard shadow spectroscopy requires

$$N_{\mathrm{meas}}^{\mathrm{SS}} = O(3^q N_t \epsilon^{-2} \ln N_o)$$

randomized measurements [9,10]. This scaling describes the measurement cost, but not the cost of implementing the time evolution.

For a first-order Trotter formula, reducing the Trotter error to ϵ_{Trot} requires

$$r = O\big(t^2 \epsilon_{\mathrm{Trot}}^{-1}\big)$$

Trotter steps, up to Hamiltonian-dependent constants [18,21]. If G_{step} denotes the gate count of one Trotter step, the gate count of the deterministic Trotter circuit therefore scales as

$$G_{\mathrm{Trot}}(t, \epsilon_{\mathrm{Trot}}) = O\big(G_{\mathrm{step}} t^2 \epsilon_{\mathrm{Trot}}^{-1}\big).$$

Thus, suppressing the systematic Trotter bias requires increasingly deep circuits.

We first consider TE-PAI at a fixed finite Trotter step size. In this setting, TE-PAI can be viewed as a randomized implementation of the same first-order Trotterized evolution. It therefore has the same Trotter bias as the corresponding deterministic product formula, but reduces the expected gate count of each sampled circuit by probabilistically omitting rotations. If $G_{\mathrm{samp}}(t)$ denotes the expected gate count of a sampled TE-PAI circuit, then

$$G_{\mathrm{samp}}(t) \leqslant G_{\mathrm{Trot}}(t, \epsilon_{\mathrm{Trot}}).$$

This gate-count reduction is obtained at the cost of a sampling overhead. If Γ denotes the TE-PAI normalization factor, the effective measurement cost becomes

$$N_{\mathrm{meas}}^{\text{TE-PAI+SS}} = O\big(\Gamma^2 3^q N_t \epsilon^{-2} \ln N_o\big).$$

Hence, at finite Trotter step size, TE-PAI trades expected gate count per circuit for a multiplicative sampling overhead Γ^2, while retaining the same first-order Trotter bias.

The stronger advantage appears when the zero-Trotter-error limit is considered. In deterministic first-order Trotterization, taking $\epsilon_{\mathrm{Trot}} \to 0$ makes the circuit cost diverge as

$$G_{\mathrm{Trot}}(t, \epsilon_{\mathrm{Trot}}) = O\big(G_{\mathrm{step}} t^2 \epsilon_{\mathrm{Trot}}^{-1}\big).$$

In contrast, the infinite-step TE-PAI limit gives an unbiased estimator of the exact time evolution. In this limit, the expected number of sampled elementary rotations scales only linearly with the simulated time,

$$G_{\text{TE-PAI}}(t) = O(t),$$

up to Hamiltonian-dependent constants and the cost per sampled term. Therefore,

$$G_{\text{TE-PAI}}(t) \ll G_{\mathrm{Trot}}(t, \epsilon_{\mathrm{Trot}}) \quad \text{as} \quad \epsilon_{\mathrm{Trot}} \to 0.$$

Thus, TE-PAI has a two-level advantage: At finite Trotter step size, it reduces the expected gate count while preserving the same discretization bias, and in the infinite-step limit it removes the Trotter bias entirely, replacing the accuracy-dependent $O(t^2\epsilon_{\mathrm{Trot}}^{-1})$ first-order Trotter gate scaling by an unbiased sampled evolution with linear-in-time expected gate scaling at fixed Δ and sampling overhead Γ^2.

TE-PAI is also particularly relevant for early fault-tolerant quantum computing. In this regime, a major cost comes from

implementing non-Clifford rotations, usually through magic-state resources or T-gate synthesis. TE-PAI directly reduces this cost by decreasing the number of rotations that must be implemented in each sampled circuit [13]. This feature is especially compatible with recent proposals for efficient analog rotation gates, such as STAR-magic mutation [29]. These approaches are most effective when many logical rotations share the same angle, because the cost of preparing and reusing the corresponding rotation resource can be amortized. TE-PAI naturally produces this favorable structure: Instead of requiring a large set of unrelated rotation angles, it repeatedly samples rotations from a small set of fixed angles. Thus, in early-FTQC architectures, TE-PAI can reduce both the number of implemented rotations and the effective T gate or magic-state compilation overhead.

IV. EXPERIMENTS

In this section, we validate the performance of our proposed TE-PAI shadow spectroscopy protocol through a series of experiments conducted numerically and on quantum devices. We benchmark our method against the standard algorithmic shadow spectroscopy [9], using first-order Trotterized time evolution. Hereafter, we refer to the spectra obtained by the standard and proposed methods as the "Trotter-based spectrum" and the "TE-PAI-based spectrum," respectively.

A. Classical simulations

In this subsection, we demonstrate the efficacy and robustness of TE-PAI shadow spectroscopy through classical simulations. We first conduct a noise-free simulation to validate that our protocol can accurately recover the energy gaps of a target Hamiltonian under ideal conditions. Subsequently, we introduce a depolarizing noise model to assess the protocol's resilience to gate errors, a critical factor for practical applications on near-term devices. Through these simulations, performed on a Heisenberg model using the Qiskit Aer simulator, we compare the performance of our method against the conventional Trotter-based approach, highlighting the advantages of TE-PAI in noisy environments.

For the noise-free simulations, we consider a ten-qubit one-dimensional Heisenberg model. The Hamiltonian used in our experiments is

$$H = \sum_{i=1}^{N-1} (J_x X_i X_{i+1} + J_y Y_i Y_{i+1} + J_z Z_i Z_{i+1}), \tag{9}$$

where X_i, Y_i, Z_i are the Pauli operators acting on the ith qubit. J_x, J_y, and J_z are the coupling strength. They are set to 1 in our simulation.

To isolate a specific transition energy gap, the initial state is prepared as a superposition of the ground state and the tenth excited state. With this choice of initial state, the dominant oscillation frequency in the dynamical signal is expected to correspond to

$$\Delta E_{0,10} = E_{10} - E_0.$$

For the finite system studied here, exact diagonalization gives $\Delta E_{0,10} \approx 4.3595$, as detailed in Appendix F. This value is not the spectral gap $E_1 - E_0$, but the target transition energy gap selected by our state preparation.

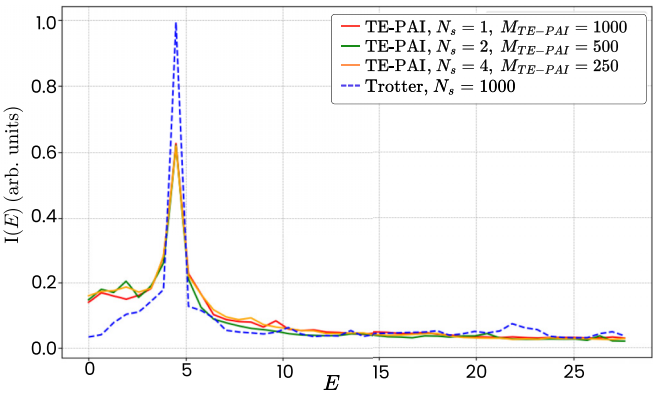


FIG. 1. Energy spectra of the ten-qubit Heisenberg model from our TE-PAI-based method (solid lines) and the Trotter-based method (dotted line), with the total number of circuit executions fixed at 1000. The legend details the combinations of TE-PAI samples ($M_{\text{TE-PAI}}$) and shadow snapshots (N_s). For the same execution cost, the Trotter-based method shows a higher peak intensity, while the TE-PAI results are consistent across different configurations.

The time evolution is simulated for $N_t = 90$ time points with a step of $dt = 0.11$. For both the TE-PAI and Trotter-based methods, the underlying Trotter decomposition uses $K_{\text{steps}} = 650$ steps, and for TE-PAI, we set $\Delta = \pi/2^7$. Although the simulation of the circuits is noise-free, i.e., free from gate or decoherence errors, we incorporate the statistical noise inherent in quantum measurements by taking only a single shot from each executed circuit. To ensure a fair comparison, we fix the total number of circuit executions to 1000 at each time point. For the standard Trotter-based method, this corresponds to simply taking $N_s = 1000$ shadow snapshots. For the TE-PAI method, we set $M_{\text{TE-PAI}} \times N_s = 1000$ and explore configurations such as $(M_{\text{TE-PAI}}, N_s) = (1000, 1), (500, 2)$, and $(250, 4)$. At each time point, we collect classical-shadow snapshots and feed them into the algorithmic shadow spectroscopy protocol. Within this post-processing, we estimate the expectation values of all 3-local Pauli operators from the snapshots and retain only the top 10% of the most significant signals, which are identified using a Ljung-Box test (see also Appendix D for detailed procedure). The Ljung-Box test and the retention of the top 10% most significant signals constitute heuristic components of the spectral postprocessing, adopted in line with prior work. A detailed assessment of the sensitivity to these choices is beyond the scope of this work; however, within the parameter ranges explored here, we do not observe qualitative changes in the extracted spectral features when these thresholds are varied moderately.

Figure 1 shows the energy spectra obtained from the noise-free simulation. As can be seen, both the Trotter-based method (blue dotted line) and our TE-PAI-based approach (solid lines) successfully identify the target transition energy gap at $\Delta E_{0,10} \approx 4.36$, which validates the fundamental principle and correctness of our protocol.

A clear difference, however, appears in the signal intensity; for a fixed total number of 1000 circuit executions, the

Trotter-based spectrum exhibits a significantly higher peak. This reduced intensity for the TE-PAI method is an expected consequence of the additional statistical variance introduced by the quasiprobability sampling of gates. As discussed in Ref. [13], this variance necessitates a larger number of total samples to achieve the same signal-to-noise ratio as the deterministic Trotter evolution. However, this higher sampling cost is the trade-off for the primary advantage of TE-PAI: a significant reduction in circuit depth, which provides crucial robustness against gate errors, as we will demonstrate later.

A second, notable finding from Fig. 1 is the internal consistency of the TE-PAI results. The peak intensity remains nearly identical regardless of how the 1000 total executions are distributed between the number of TE-PAI samples, $M_{\text{TE-PAI}}$, and the shadow snapshots per sample, N_s. This is a nontrivial observation, as it empirically validates that the statistical precision of our hybrid protocol is determined solely by the total number of samples, not the specific strategy used to obtain them. This finding has significant practical implications, as it offers the flexibility to choose the most convenient balance between generating many distinct evolution circuits (large $M_{\text{TE-PAI}}$) and performing more measurements on fewer circuits (large N_s) without affecting the final signal-to-noise ratio for a fixed total cost (see Appendix G for simulation with Large N_s).

Next, to assess the protocol's performance under more realistic conditions, we perform simulations on a six-qubit Heisenberg model with a depolarizing noise model. First, we construct the quantum circuits with a gate set consisting of $R_X, R_Z, R_{XX}, R_{YY}, R_{ZZ}$ Pauli rotations. We then introduce depolarizing noise after each gate, with an error probability of $p_1 = 10^{-4}$ for single-qubit gates and $p_2 = 10^{-3}$ for two-qubit gates. For this experiment, we set the number of Trotter steps to $K_{\text{steps}} = 300$ and the total number of circuit executions to 3000. The PAI with $\Delta = \pi/2^5$ is used to reduce circuit depth. In this setup, the Trotter-based circuits for $K_{\text{steps}} = 300$ have an average depth of about 1810, while the corresponding TE-PAI circuits have an average depth of about 751. The depth is calculated as the minimum number of sequential layers of gates needed to execute the circuit, i.e., the length of the longest path through gates with qubits not sharing a layer. It was obtained in the simulation using Qiskit QuantumCircuit.depth method [30], version = "1.4.2".

Figure 2 shows the results of the noisy simulation, comparing them against the noise-free case. The outcome reveals a stark difference in noise robustness. The TE-PAI spectrum is only slightly degraded by the noise and the spectral peak remains clearly resolved. In contrast, the Trotter-based spectrum collapses, as the noise almost entirely suppresses the peak. This result confirms the superior noise resilience of the TE-PAI method, a direct consequence of its significantly shallower circuits.

To furthere probe the interplay between sampling and gate noise, in Fig. 3, we apply TE-PAI and standard Trotter shadow spectroscopy to a six-qubit Heisenberg chain with $J_x = J_y = J_z = 1, J = 15$ Pauli terms, and $\|H\|_1 = 15$. The initial state is prepared as $|\psi_{\text{init}}\rangle = \frac{1}{\sqrt{2}}(|E_0\rangle + |E_{10}\rangle)$. For this state, the ideal time signal contains a single dominant transition frequency at $\Delta E_{0,10} = 5.917$. The time-evolved state is reconstructed from $q = 3$-local Pauli shadows, and gate noise is modeled as

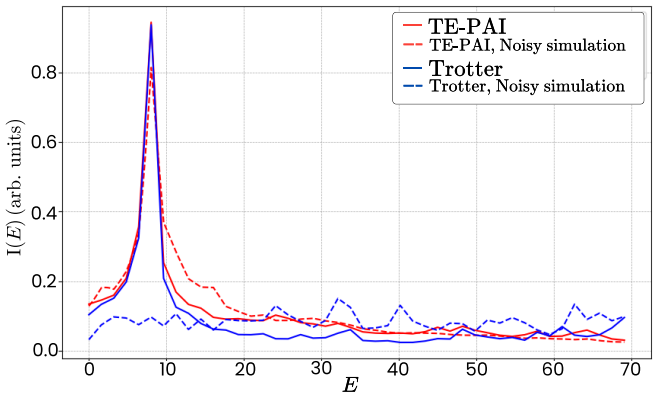


FIG. 2. Effect of depolarizing noise on the energy spectra obtained from TE-PAI (red) and Trotter-based (blue) methods. Dotted lines represent the noise-free benchmarks, while solid lines show the results under noise. The TE-PAI spectrum demonstrates strong robustness, with its peak remaining clearly visible. In contrast, the Trotter-based spectrum collapses, losing nearly all spectral information. This highlights the practical advantage of TE-PAI's shallower circuits in noisy environments.

amplitude damping, corresponding to T_1 relaxation, with rates $p_1 = 2.5 \times 10^{-5}$ and $p_2 = 2.5 \times 10^{-4}$ on one- and two-qubit gates, respectively. Both methods use the same first-order step size $\delta t = 0.008$, i.e., the same $K_{\text{step}} = 500$ steps up to $t_{\text{max}} = 4$ on a grid of $N_t = 80$ time points, with the TE-PAI angle $\Delta = \pi/2^6$. The deterministic Trotter circuit applies all J Pauli rotations at every step, giving 7.5×10^3 gates. In contrast, TE-PAI keeps only $\mathbb{E}[\nu_l] \approx 2.4 \times 10^3$ gates on average, about one third as many, at the cost of quasiprobability overhead $\Gamma \approx 7.3$.

For the Trotter circuit, the noisy limiting spectrum is already dominated by a spurious low-frequency decoherence mode, and the reconstruction fails for both shot budgets. Increasing M therefore cannot restore the true transition gap; it only reduces statistical fluctuations around an already distorted noisy signal.

The TE-PAI reconstruction behaves differently. The correct gap peak remains identifiable for both $M = 1000$ and $M = 10\,000$, confirming the benefit of the shallower sampled circuits. However, increasing the shot budget does not monotonically improve the spectrum. A larger M reduces statistical fluctuations and drives the estimator closer to the spectrum of the implemented noisy circuit ensemble; under gate noise, this limiting spectrum is not the ideal Trotter spectrum but the decohered spectrum generated by the noisy circuit. Consequently, the noise-induced low-frequency feature can also become sharper, so the reconstruction may look worse even though the statistical variance is reduced. This suggests that the quasiprobability sampling noise, usually regarded as the main drawback of TE-PAI, is not necessarily uniformly detrimental in this spectroscopy task: At lower shot budgets, finite-sampling fluctuations may partly obscure gate-noise-induced artifacts, while the true gap oscillation remains visible due to the shallower circuits. This may also indicate that, in this instance, the genuine gap peak is relatively robust against finite-sampling fluctuations.

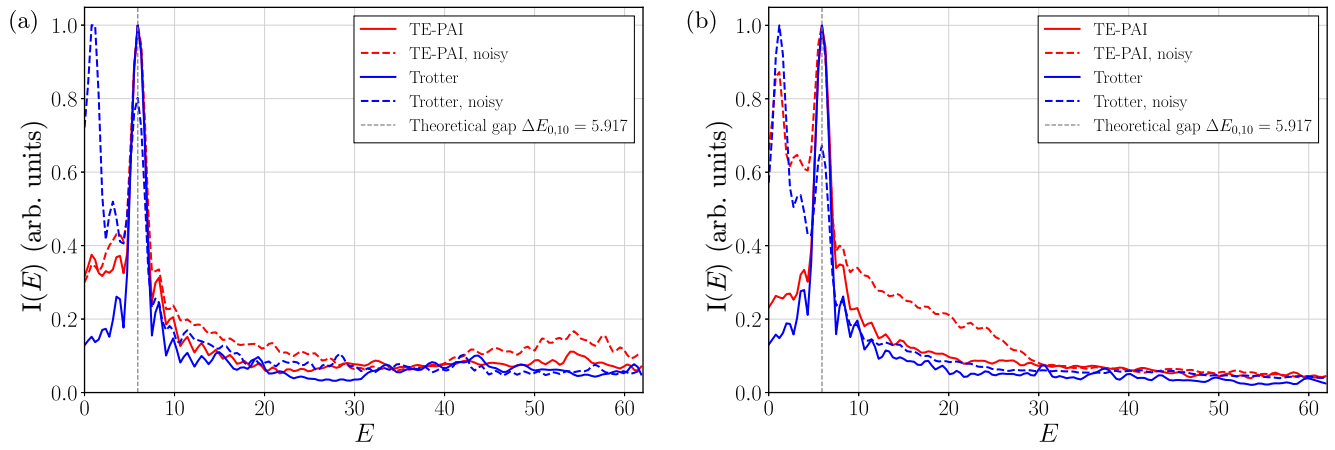


FIG. 3. Under *amplitude-damping* (T_1) gate noise, a six-qubit Heisenberg chain ($J_x = J_y = J_z = 1$, $\|H\|_1 = 15$) is prepared in $|\psi_0\rangle = \frac{1}{\sqrt{2}}(|E_0\rangle + |E_{10}\rangle)$ with target gap $\Delta E_{0,10} = 5.917$ (gray dashed line). We compare TE-PAI (red) and standard Trotterization (blue), both noise-free (solid) and subject to amplitude damping (dashed; $p_1 = 2.5 \times 10^{-5}$, $p_2 = 2.5 \times 10^{-4}$). Both methods use the same constant step width $\delta t = 0.008$ and $\Delta = \pi/2^6$, corresponding to the same 500 time steps up to $t_{\max} = 4$. They differ only in circuit depth: Standard Trotterization applies 15 gates per step, giving $15 \times 500 = 7500$ gates, whereas TE-PAI keeps only 2400 gates on average—about 33% as many—at the price of a quasiprobability sampling overhead $\Gamma \approx 7.3$, which is absent for Trotterization. Increasing the number of shots from (a) $M = 1000$ to (b) $M = 10\,000$ suppresses shot noise but does not necessarily improve the noisy spectrum. The deep Trotter circuit fails in both panels, being dominated by a spurious low-frequency decoherence peak. The shallower TE-PAI circuit keeps the correct gap peak identifiable in both panels, but the larger shot budget also sharpens a noise-induced low-frequency feature. Thus, additional sampling can make the wrong spectral feature more visible rather than removing gate-noise-induced errors.

Having established the protocol on a single energy gap, we next verify that its noise robustness persists when several transition gaps must be resolved simultaneously. Multigap simulations are more stringent because the reconstructed spectrum must separate several weak transition peaks from cross-gap contributions and noise-induced artifacts. Following the standard practice for shadow spectroscopy [9], we prepare a dominant-ground superposition $|\psi_{\text{init}}\rangle = \frac{1}{\mathcal{N}}(|E_0\rangle + \epsilon \sum_k |E_k\rangle)$ over a set of low-lying eigenstates of a five-qubit Heisenberg chain ($J_x = J_y = J_z = 1$, $\|H\|_1 = 12$), with $\epsilon = 0.25$ and normalization factor $\mathcal{N}$. For such a state, the spectral weights $c_0^* c_k \langle E_0| O |E_k\rangle$ of the ground-to-excited transitions dominate, so the time signal $S(t)$ in Eq. (1) carries several gap frequencies $\Delta E_{k,0}$ at once, while the cross-gaps ΔE_{kl} with $k, l \neq 0$ are suppressed as $\mathcal{O}(\epsilon^2)$.

Figure 4 compares TE-PAI (red) and first-order Trotter (blue) shadow spectroscopy for this state, without gate noise on the left and with depolarizing gate noise on the right ($p_1 = 4 \times 10^{-5}$, $p_2 = 4 \times 10^{-4}$), using $q = 3$-local Pauli shadows and $M = 4000$ shots per time point. The ground-dominant initial state is designed to make the ground-to-excited transition gaps visible while keeping excited-to-excited cross-gaps relatively suppressed.

Both methods use the same step width δt ($\Delta = \pi/2^6$, $K_{\text{steps}} = 700$ steps up to $t = 6$) and therefore incur the same Trotter error; the spectra are evaluated on $N_t = 120$ time points. The two differ in circuit depth: TE-PAI realizes this accuracy with a $\sim 3\times$ shallower circuit ($\sim 2.9 \times 10^3$ versus $\sim 8.4 \times 10^3$ gates), at the cost of a sampling overhead $\Gamma \approx 10$. In the noiseless case, both resolve the low-lying gaps, but with a complementary bias: Standard Trotter captures the three lowest transitions ($\Delta E_{k,0} = 2.88,\ 4.48,\ 6.48$) most strongly while leaving the two highest (10.07, 11.71) at the floor, whereas TE-PAI under-weights the lowest transition yet retains visible peaks at the two highest gaps, where it outperforms Trotter.

Under gate noise, this difference is decisive: The deeper Trotter circuit accumulates more decoherence and its spectrum is strongly suppressed at the resolved gaps, drifting toward a spurious low-frequency mode. In contrast, the shallower TE-PAI circuit preserves visible signatures of the gap structure, with higher contrast at several resolved gaps, especially in the higher-frequency region. The advantage of TE-PAI shadow spectroscopy is therefore not specific to the single-gap setting but carries over to the simultaneous estimation of multiple transition gaps.

B. Experiments on quantum hardware

To demonstrate the practical viability of our protocol, we present experiments for a 20-qubit transverse-field Ising model on IBM quantum hardware. The Hamiltonian for this model is

$$H = -J \sum_{k=1}^{N-1} Z_k Z_{k+1} - d \sum_{k=1}^{N} X_k. \tag{10}$$

The theoretical spectral gap between the ground and first excited states is $\Delta E = 2\sqrt{J^2 + d^2 - 2Jd\cos(\pi/(N+1))}$ [31]. As a practical compromise between preparation efficiency and sufficient overlap with these states, we use the initial state $|+\cdots+0\rangle$ and set the parameters to $J = 0.1$ and $d = 2$. With this choice of initial state and the subsequent shadow-spectroscopy filtering, the dominant observed spectral feature corresponds to the ground-first-excited-state transition, while other transitions have much weaker spectral weight in

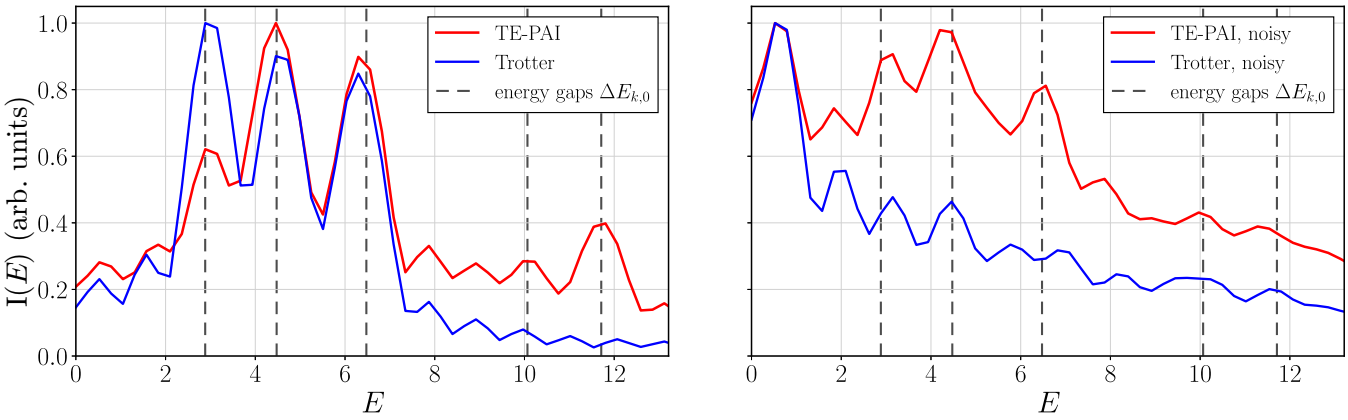


FIG. 4. Robustness of TE-PAI shadow spectroscopy to gate noise for a *multigap* initial state, with $M = 4000$ shots per time point. A five-qubit Heisenberg chain ($J_x = J_y = J_z = 1$, $\|H\|_1 = 12$) is prepared in a dominant-ground superposition $|\psi_{\text{init}}\rangle = \frac{1}{\mathcal{N}}(|E_0\rangle + \epsilon \sum_k |E_k\rangle)$ over the low-lying levels $\{1, 2, 4, 7, 9\}$, with $\epsilon = 0.25$ and normalization factor $\mathcal{N}$. The resulting dynamics contains several transition gaps $\Delta E_{k,0}$ (gray dashed lines), rather than a single frequency. The spectral intensity $I(E)$ is reconstructed using TE-PAI (red) and Trotterization (blue). The left panel shows the noiseless spectra, while the right panel includes depolarizing gate noise with $p_1 = 4 \times 10^{-5}$ and $p_2 = 4 \times 10^{-4}$. Both methods use the same step width, with $\Delta = \pi/2^6$ and $K_{\text{step}} = 700$ steps up to $t = 6$. TE-PAI realizes this evolution with a $\sim 3\times$ shallower circuit, using $\sim 2.9 \times 10^3$ gates on average instead of $\sim 8.4 \times 10^3$ Trotter gates, at a quasiprobability overhead $\Gamma \approx 10$. In the absence of gate noise, both methods resolve the low-lying transition gaps. We consistently observe that TE-PAI gives higher-contrast peaks at several higher-frequency gaps. Under gate noise, the shallower TE-PAI circuit preserves visible signatures of the gap structure, especially in the higher-frequency region, whereas the deeper Trotter circuit is strongly suppressed at the resolved gaps and drifts toward a spurious low-frequency decoherence mode.

the processed signal. We set the number of Trotter steps to $K_{\text{steps}} = 115$ and use a TE-PAI angle of $\Delta = \pi/2^5$, which limits the average circuit depth to approximately 600. For the time evolution, we use $N_t = 80$ steps with an interval of $dt = 0.037$. At each time point, we perform 3000 total circuit executions, configured as $(M_{\text{TE-PAI}}, N_s) = (3000, 1)$. From the resulting shadow snapshots, the algorithmic shadow spectroscopy process is performed in exactly the same manner as the classical simulation. We run these experiments on the *ibm_kobe* and *ibm_kingston* quantum computers, whose detailed properties are provided in Appendix H.

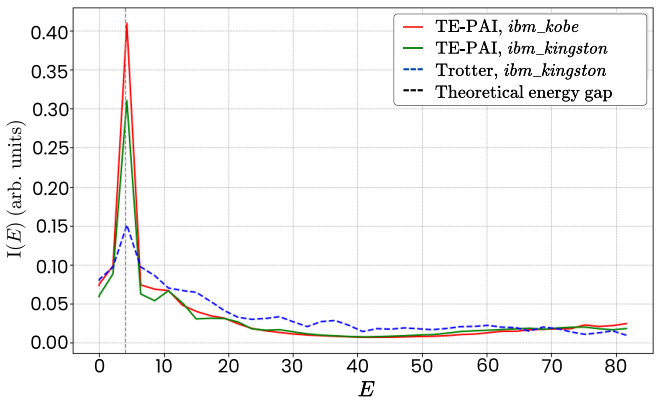


FIG. 5. Spectra of a 20-qubit Transverse Ising Hamiltonian measured on quantum hardware. The red and green lines show a TE-PAI-based spectrum performed, respectively, on *ibm_kobe* and *ibm_kingston*. The blue dotted line shows a Trotter-based spectrum performed on *ibm_kingston*. The gray dotted vertical line indicates the exact energy gap.

Figure 5 shows the spectra from our experiments on quantum hardware. The data from both the *ibm_kobe* (red line) and *ibm_kingston* (green line) devices correctly resolve the theoretical energy gap, confirming our protocol's viability under realistic noise conditions. We observe a higher peak from *ibm_kobe*, which is consistent with its lower average gate error rate as detailed in Appendix H. Crucially, the TE-PAI spectrum is more pronounced than the one from the conventional Trotter-based method (blue dotted line) run on the same device. This result provides direct experimental evidence that the shallower circuits produced by TE-PAI lead to superior robustness against hardware noise. It thereby marks a significant scale-up from the six-qubit hardware demonstration in the original work [9].

V. CONCLUSION

In this work, we introduced TE-PAI shadow spectroscopy, a hybrid quantum-classical protocol that integrates the randomized time evolution of TE-PAI with the measurement framework of algorithmic shadow spectroscopy. We have theoretically established that this combination yields an unbiased estimator for spectral properties and experimentally validated its performance. Our noise-free classical simulations confirmed that the protocol correctly identifies energy gaps, while also highlighting an inherent trade-off: The quasiprobability sampling in TE-PAI introduces statistical variance that requires a larger number of total circuit executions to achieve the same signal intensity as the first-order Trotter-based method. However, this cost is made up for by the protocol's primary advantage, which we demonstrated in noisy simulations: A significant reduction in circuit depth leads to superior robustness against gate errors. We also found that, under gate noise, a moderate shot budget can preserve the correct peak while

suppressing spurious noise-induced features. In contrast, increasing the shot budget can sharpen these spurious features and thereby worsen the spectral estimate. This suggests that the sampling overhead of TE-PAI is not merely a drawback in this task: By avoiding over-resolution of gate-noise-induced artifacts, it can sometimes help robust energy gap estimation. Also, our additional multigap simulations further showed that TE-PAI can resolve several transition gaps and tends to identify higher-frequency gaps more clearly than the Trotter-based approach. This key benefit was further confirmed through experiments on 20-qubit quantum hardware, where our method successfully resolved the target energy gap and outperformed the conventional approach, demonstrating its practical viability for spectroscopy in a significantly larger-scale experiment than previously performed.

Let us briefly address the practical time complexity of our protocol compared to standard approaches. While TE-PAI utilizes $M_{\text{TE-PAI}}$ unique circuit instances rather than a single reused circuit, the efficiency of modern control stacks prevents this from causing a significant increase in wall-clock execution time. Current quantum runtime environments (such as Qiskit Runtime) support circuit batching, allowing the full ensemble of $M_{\text{TE-PAI}}$ circuits to be submitted and executed within a single job, thereby avoiding the latency overhead associated with multiple job submissions. Furthermore, our variance analysis (Theorem 2) and numerical results (Fig. 1) demonstrate that the estimator remains unbiased even when reusing a single compiled circuit instance for N_s shots. This allows for the compilation and loading costs to be amortized over N_s measurements. By optimizing the ratio of unique circuits (M) to shots-per-circuit (N_s) for a fixed total budget, one can balance the statistical requirements against the hardware's specific compilation and loading latencies, rendering the total wall-clock time comparable to standard Trotterized evolution.

Looking forward, TE-PAI shadow spectroscopy opens several promising research avenues. The natural trade-off between circuit depth and sampling overhead constitutes a tunable design parameter that can be optimized in accordance with the specific performance characteristics of the underlying quantum hardware. In particular, quantum devices with lower error rates can reliably support deeper circuits, thereby reducing the required sampling overhead. Conversely, on hardware with higher error rates, circuit depth can be deliberately curtailed at the expense of increased sampling overhead, enabling a more favorable balance between noise resilience and statistical accuracy. An important direction for future work is to quantify how much different product formulas, including symmetric second-order and randomized Trotter formulas [32], can improve the performance of TE-PAI shadow spectroscopy on quantum hardware. A natural next step would be to combine our protocol with other error mitigation techniques, such as readout error correction or zero-noise extrapolation, to further suppress the impact of device imperfections. Furthermore, the demonstrated scalability and noise resilience make our method a strong candidate for tackling challenging problems in quantum chemistry and condensed matter physics, where deep circuits have been a major bottleneck. The flexibility in distributing the sampling budget between the number of TE-PAI circuits and shadow snapshots also suggests possibilities for hardware-aware codesign of experiments, potentially leading to even more efficient spectral analysis on near-term quantum computers.

ACKNOWLEDGMENTS

Y.M. was supported by JST SPRING (Grant No. JPMJSP2138). This work was partially supported by JST ASPIRE (Grants No. JPMJAP2319 and No. JPMJAP26A4), the New Energy and Industrial Technology Development Organization (NEDO) (Project No. JPNP20017), JST COI-NEXT (Grant No. JPMJPF2014), JST Moonshot Re&D (Grant No. JPMJMS256J), MEXT Q-LEAP (Grant No. JPMXS0120319794), and JST NEXUS (Grant No. JPMJNX26C6). K.M. was supported by JST FOREST (Grant No. JPMJFR232Z), JSPS KAKENHI (Grants No. 23H03819 and No. 24K16980), and JST CREST (Grant No. JPMJCR24I4). C.K. was supported by JST PRESTO, Japan (Grant No. JPMJPR25F1), and JSPS KAKENHI (Grant No. JP26K17052). This work, as part of the ITI 20212028 program of the University of Strasbourg, CNRS and Inserm, was supported by IdEx Unistra (Grant No. ANR-10-IDEX-0002), and by SFRI STRAT'US project (ANR-20-SFRI-0012) and EUR QMAT (ANR-17-EURE-0024) under the framework of the French Investments for the Future Program. B.K. thanks UKRI for the Future Leaders Fellowship Theory to Enable Practical Quantum Advantage (MR/Y015843/1) and EPSRC for funding through the project Software Enabling Early Quantum Advantage (SEEQA, EP/Y004655/1). This research was funded in part by UKRI (Grant No. MR/Y015843/1).

DATA AVAILABILITY

The simulation code used in this work is available online [33].

APPENDIX A: PRODUCT FORMULAS

Throughout this work, we employ the first-order Trotter product formula. The associated Trotter error, arising from the discretization of time evolution, is an intensively studied topic [21]; here, we briefly summarize the relevant results. Let $U_{H,t}^{(K)}$ denote the Trotter circuit defined in Eq. (2) with K Trotter steps. The spectral norm error between the exact time-evolution operator $U_{H,t}$ and the approximation $U_{H,t}^{(K)}$ is bounded by

$$\left\|U_{H,t} - U_{H,t}^{(K)}\right\| \leqslant \frac{t^2\|c\|_T}{2K},$$

where $\|c\|_T := \sum_{\gamma<\beta} \|[h_\gamma P_\gamma, h_\beta P_\beta]\|$ is the commutator norm of H. The error in the estimated expectation value induced by the Trotter approximation is bounded by the operator error and the spectral norm of the observable:

$$\begin{aligned}\epsilon_T &:= \left|S(t) - \langle O\rangle_t^{(K)}\right| \\ &\leqslant 2\|O\| \left\|U_{H,t} - U_{H,t}^{(K)}\right\| \\ &\leqslant \frac{t^2\|O\|\|c\|_T}{K},\end{aligned}$$

where $\langle O\rangle_t^{(K)} := \langle\langle O|\boldsymbol{U}_{H,t}^{(K)}(\rho)\rangle\rangle$ denotes the expectation value obtained by evolving ρ with the K-step Trotter circuit. As K

increases, the Trotter approximation converges to the exact evolution, so the Trotter error vanishes in the limit $K \to \infty$ [the bound above scales as $O(t^2/K)$]; moreover, this is a worst-case upper bound, and the empirical error for specific instances is often much smaller [21].

Other than the Trotter error, finite sampling introduces statistical error. To estimate $\langle O \rangle_t^{(K)}$ on hardware, we sample N independent measurement outcomes $o_1, \ldots, o_N$ from repeated executions of the Trotter circuit $U_{H,t}^{(K)}$ and use the sample mean $\hat{o} := \frac{1}{N}\sum_{n=1}^{N} o_n$. Then, $\mathbb{E}[\hat{o}] = \langle O \rangle_t^{(K)}$ and $\mathrm{Var}(\hat{o}) \leqslant \|O\|^2/N$, leading to a sample complexity scaling as

$$N = \mathcal{O}\left(\frac{\|O\|^2 \ln(1/\delta)}{\epsilon^2}\right),$$

to achieve additive error ϵ with failure probability δ.

APPENDIX B: QUASIPROBABILITY SAMPLING

The theoretical results presented in Sec. III B are not limited to TE-PAI but apply generally to any quasiprobability sampling method. Here, we briefly review this underlying framework.

In this approach, quantum channels that are difficult to implement exactly are approximated by randomly sampling from a set of simpler operations. This strategy is mathematically grounded in quasiprobability decompositions, which serve as the foundation for many randomized compiling techniques. Prominent examples include PEC for stochastic noise mitigation [11,12,20], PAI for coherent or synthesis-error mitigation [13,14], and circuit cutting [34–36].

Suppose a target channel $\boldsymbol{C}$ can be decomposed as a linear combination

$$\boldsymbol{C} = \sum_{d=1}^{D} a_d \boldsymbol{C}_d,$$

where $a_d \in \mathbb{R}$ are real coefficients and $\{\boldsymbol{C}_d\}$ is a set of efficiently implementable basis channels. We define the normalization factor $\gamma(\boldsymbol{C}) = \sum_{d=1}^{D} |a_d|$ and assign sampling probabilities $p_d = |a_d|/\gamma(\boldsymbol{C})$.

To simulate $\boldsymbol{C}$, we sample an index $l \in \{1, \ldots, D\}$ from the distribution $\{p_l\}$, apply the channel $\boldsymbol{C}_l$, and weight the measurement outcome by $\gamma(\boldsymbol{C})\mathrm{sign}(a_l)$. This weighting is performed via classical postprocessing; specifically, for any observable O and state ρ, the expectation value satisfies

$$\langle\langle O|\boldsymbol{C}(\rho)\rangle\rangle = \gamma(\boldsymbol{C})\mathbb{E}[\mathrm{sign}(a_l)\langle\langle O|\boldsymbol{C}_l(\rho)\rangle\rangle],$$

where the expectation is taken over the random choice of l. Consequently, the resulting random estimator is unbiased, satisfying

$$\mathbb{E}[\gamma(\boldsymbol{C})\mathrm{sign}(a_l)\boldsymbol{C}_l] = \boldsymbol{C}.$$

Next, we consider a sequential channel $\boldsymbol{C} = \boldsymbol{C}_K \cdots \boldsymbol{C}_1$ with each $\boldsymbol{C}_k$ decomposed as

$$\boldsymbol{C}_k = \sum_{d=1}^{D} a_{k,d} \boldsymbol{C}_{k,d},$$

where $\boldsymbol{C}_{k,d}$ are efficiently implementable channels. Let $l_k \in \{1, \ldots, D\}$ be a random variable drawn according to $p_d = |a_{k,d}|/\gamma(\boldsymbol{C}_k)$, where $\gamma(\boldsymbol{C}_k) = \sum_d |a_{k,d}|$. Define the random channel

$$\hat{\boldsymbol{C}}_{\boldsymbol{l}} := \boldsymbol{C}_{K,l_K} \cdots \boldsymbol{C}_{2,l_2}\boldsymbol{C}_{1,l_1},$$

with the classical postprocessing weight $\Gamma_{\boldsymbol{l}} = \prod_{k=1}^{K} \gamma(\boldsymbol{C}_k)\mathrm{sign}(a_{k,l_k})$. For any initial state ρ, applying the sampled channel sequence produces an unbiased estimator of $\boldsymbol{C}(\rho)$ as

$$\mathbb{E}[\Gamma_{\boldsymbol{l}}\hat{\boldsymbol{C}}_{\boldsymbol{l}}(\rho)] = \boldsymbol{C}(\rho).$$

While this reweighting guarantees unbiasedness, it scales the estimator's variance by approximately $\Gamma^2 := \prod_{k=1}^{K} \gamma(\boldsymbol{C}_k)^2$, thereby requiring the number of samples to increase by a factor of Γ^2 to maintain a fixed estimation error.

APPENDIX C: HIGHER-ORDER TE-PAI: GATE COUNT AND SAMPLING OVERHEAD

We compare K first-order steps with $K/2$ symmetric second-order steps over the same total time t, where K is even and $\delta t = t/K$:

$$U_{H,t}^{(1,K)} = \left[\prod_{j=1}^{J} e^{-ih_j P_j \delta t}\right]^K,$$

$$U_{H,t}^{(2,K/2)} = \left[\prod_{j=1}^{J} e^{-ih_j P_j \delta t} \prod_{j=J}^{1} e^{-ih_j P_j \delta t}\right]^{K/2}.$$

The second-order step has duration $2\delta t$, with each sweep covering its half step. Both circuits contain K copies of each rotation R_{P_j,θ_j}, where $\theta_j = 2h_j\delta t$, and hence KJ rotations in total. Their global product-formula errors scale as $O(t\delta t)$ and $O(t\delta t^2)$, respectively, up to Hamiltonian-dependent constants [18,19,21].

For a TE-PAI decomposition of a rotation angle θ, the sampling overhead is given by

$$\gamma(\theta) := |a_1(\theta)| + |a_2(\theta)| + |a_3(\theta)|,$$

where a_1, a_2, a_3 are defined in Eq. (3). For $|\theta| \leqslant \Delta$,

$$\gamma(\theta) = \cos|\theta| + \tan\left(\frac{\Delta}{2}\right)\sin|\theta|,$$

and hence

$$\ln\gamma(\theta) = |\theta|\tan\left(\frac{\Delta}{2}\right) + O(\theta^2). \tag{C1}$$

The probability that a rotation of angle θ is sampled as a nonidentity gate is

$$p_{\neq I}(\theta) = \frac{|a_2(\theta)| + |a_3(\theta)|}{\gamma(\theta)} = \frac{3 - \cos\Delta}{2\sin\Delta}|\theta| + O(\theta^2).$$

Let $\nu_{\boldsymbol{l}}$ denote the number of nonidentity gates in a sampled TE-PAI circuit.

Applying TE-PAI independently to these rotations with the same $\Delta \geqslant \max_j |\theta_j|$ gives exactly the same expected

number of sampled nonidentity rotations and quasiprobability overhead:

$$\mathbb{E}\left[\nu^{(1,K)}\right] = \mathbb{E}[\nu^{(2,K/2)}] = K\sum_{j=1}^{J} p_{\neq I}(\theta_j),$$

$$\Gamma^2_{1,K} = \Gamma^2_{2,K/2} = \left[\prod_{j=1}^{J}\gamma(\theta_j)\right]^{2K}.$$

Thus, second order improves the Trotter-error scaling without increasing either of these TE-PAI costs. These identities do not imply equal actual estimator variances.

The comparison above is made before merging adjacent rotations. For example, the two central rotations in a symmetric second-order step satisfy $R_{P_J,\theta_J}R_{P_J,\theta_J} = R_{P_J,2\theta_J}$, reducing two deterministic rotations to one. If they are merged before applying TE-PAI, their contribution to the expected rotation count changes from $2p_{\neq I}(\theta_J)$ to $p_{\neq I}(2\theta_J)$, provided $2|\theta_J| \leqslant \Delta$. Thus, merging reduces the deterministic rotation count, but its effect on the TE-PAI cost depends on the merged angle and the choice of Δ.

More generally, if a product formula is written as

$$U^{(p,K)}_{H,t} = \prod_{r=1}^{L_{p,K}} R_{P_r,\theta_r},$$

then the leading TE-PAI costs are governed by the total absolute rotation angle:

$$\ln\Gamma^2_p = 2\tan\left(\frac{\Delta}{2}\right)\sum_{r=1}^{L_{p,K}}|\theta_r| + O\left(\sum_{r=1}^{L_{p,K}}\theta_r^2\right)$$

and

$$\mathbb{E}\left[\nu_I^{(p)}\right] = \frac{3-\cos\Delta}{2\sin\Delta}\sum_{r=1}^{L_{p,K}}|\theta_r| + O\left(\sum_{r=1}^{L_{p,K}}\theta_r^2\right).$$

In the comparison above, both formulas have total absolute rotation angle $K\sum_j|\theta_j| = 2t\|H\|_1$. In contrast, higher-order Suzuki formulas of order four and above can contain negative or enlarged fractional time steps, which increase the total absolute rotation angle and can therefore increase the TE-PAI gate count and sampling overhead. We note that a continuous-time version of TE-PAI has recently been studied in Ref. [37], where the limiting ensemble can be interpreted as a Dyson-series-based sampling scheme with a natural leading l_1 rate. From this viewpoint, fourth- and higher-order Suzuki formulas are not expected to converge to this optimal leading cost, because their negative or enlarged fractional time steps increase the total absolute rotation angle.

APPENDIX D: ALGORITHMIC SHADOW SPECTROSCOPY

Let H be an N-qubit Hamiltonian with eigenstates $|E_\alpha\rangle$ for $\alpha = 1,\ldots,2^N$ and corresponding eigenvalues E_α. To obtain energy differences, $E_\alpha - E_\beta$, algorithmic shadow spectroscopy uses expectation values of a set of observables $\{O_i\}_{i=1}^{N_o}$ consisting of all q-local Pauli operators, evaluated with respect to time-evolved states $|\rho(t)\rangle\rangle = \boldsymbol{U}_{H,t}|\rho_{\text{init}}\rangle\rangle$ for some initial state $|\rho_{\text{init}}\rangle\rangle$ [9].

The algorithm is as follows. Let $\{t_n\}_{n=1}^{N_t}$ be a set of time points and define $S_i(n) := \langle\langle O_i|\rho(t_n)\rangle\rangle$.

(1) Gather shadow snapshots of $\rho(t_n)$ for each t_n using N_s copies of $\rho(t_n)$. This is the only step performed on quantum hardware; the following steps are classical.

(2) Estimate $\hat{S}_i(n)$ from the snapshots.

(3) Construct a data matrix $D \in \mathbb{R}^{N_o\times N_t}$ with entries $D_{i,n} = \frac{\hat{S}_i(n)-\mu_i}{\sigma_i}$ and μ_i, σ_i are the mean and standard deviation of $\{\hat{S}_i(n)\}_{n=1}^{N_t}$, respectively.

(4) Apply the Ljung-Box test to each row of D to remove signals indistinguishable from noise [38,39].

(5) Compute the correlation matrix

$$C = DD^\top \in \mathbb{R}^{N_o\times N_o}, \tag{D1}$$

and extract the top c dominant eigenvectors $\{v_1, v_2, \ldots, v_c\}$ corresponding to the largest eigenvalues [40].

(6) Compute the spectral cross-correlation matrix

$$x_{kl}(m) = \sum_{n=1}^{N_t-m-1} v_k(n+m)\,v_l(n), \tag{D2}$$

and perform a discrete Fourier transform of $x_{kl}(m)$ with respect to the time lag index m. Denote the resulting quantity by $X_{kl}(m)$, where m now represents the frequency index.

(7) Compute the largest singular value $\lambda(m)$ of the matrix $X(m)$. Peaks in $\lambda(m)$ correspond to the transition energy gaps of the Hamiltonian.

APPENDIX E: FORMAL ANALYSIS OF THE TE-PAI SHADOW ESTIMATOR

Algorithm 2 shows the detailed pseudocode of the practical implementation of our protocol used in the experiments. As discussed in the main text, this implementation is designed to generate one time-evolution circuit via the TE-PAI process and reuse it with N_s different randomized measurement bases to obtain N_s snapshots for each TE-PAI sampled circuit. Theorems 1 and 2 establish that this modified implementation yields an unbiased estimator and provides an explicit upper bound on the variance.

Before proving Theorems 1 and 2, we first present a lemma that upper bounds the variance of an expectation value for a single quasiprobability random circuit.

Lemma D1. The expectation value of any observable O under a quasiprobability random circuit yields an unbiased estimator of the true expectation value:

$$\mathbb{E}[\Gamma_I\langle\langle O|\hat{\boldsymbol{C}}_I(\rho)\rangle\rangle] = \langle\langle O|\boldsymbol{C}(\rho)\rangle\rangle.$$

Additionally, the variance is upper bounded as follows:

$$\mathrm{Var}\left[\Gamma_I\langle\langle O|\hat{\boldsymbol{C}}_I(\rho)\rangle\rangle\right] \leqslant \Gamma^2\|O\|^2.$$

Proof. The unbiasedness directly follows from the unbiasedness of the TE-PAI decomposition:

$$\begin{aligned}\mathbb{E}[\Gamma_I\langle\langle O|\hat{\boldsymbol{C}}_I(\rho)\rangle\rangle] &= \langle\langle O|\mathbb{E}[\Gamma_I\hat{\boldsymbol{C}}_I](\rho)\rangle\rangle\\ &= \langle\langle O|\boldsymbol{C}(\rho)\rangle\rangle.\end{aligned}$$

Write $\hat{o} := \Gamma_{\boldsymbol{l}}\langle\langle O|\hat{\boldsymbol{C}}_{\boldsymbol{l}}(\rho)\rangle\rangle$ for brevity. The variance is given by $\mathrm{Var}[\hat{o}] = \mathbb{E}[\hat{o}^2] - \mathbb{E}[\hat{o}]^2 = \mathbb{E}[\hat{o}^2] - \langle\langle O|\boldsymbol{C}(\rho)\rangle\rangle^2$, and we now upper bound the first term as

$$\begin{aligned}\mathbb{E}\left[\hat{o}^2\right] &= \sum_{\boldsymbol{l}} p_{\boldsymbol{l}} \times [\Gamma_{\boldsymbol{l}}\langle\langle O|\boldsymbol{C}_{\boldsymbol{l}}(\rho)\rangle\rangle]^2 \\ &= \Gamma^2 \sum_{\boldsymbol{l}} p_{\boldsymbol{l}} \times \langle\langle O|\boldsymbol{C}_{\boldsymbol{l}}(\rho)\rangle\rangle^2 \\ &\leqslant \Gamma^2 \|O\|^2.\end{aligned}$$

Above, we have used that $|\langle\langle O|\rho\rangle\rangle| \leqslant \|O\|$ for any density matrix ρ and $\sum_{\boldsymbol{l}} p_{\boldsymbol{l}} = 1$. ■

1. Proof of Theorem 1

This modified implementation preserves the unbiased nature of the estimator.

Proof. The goal is to estimate the true expectation value of an observable O, defined with respect to the original channel $\langle\langle O|\boldsymbol{C}(\rho)\rangle\rangle$. Our final estimator, $\langle\hat{O}\rangle$, is the average over all measurement outcomes. More specifically, let $o_{m,s}$ be the outcome from the mth quasiprobability random circuit and the sth shadow snapshot, that is,

$$o_{m,s} = \Gamma_{\boldsymbol{l}_m}\langle\langle O|\hat{\rho}_{m,s}\rangle\rangle,$$

where $\Gamma_{\boldsymbol{l}_m}$ is the weight accumulated from the mth quasiprobability sample and $\hat{\rho}_{m,s}$ is the corresponding classical-shadow snapshot defined in Eq. (6). The estimator is constructed as

$$\langle\hat{O}\rangle = \frac{1}{MN_s}\sum_{m=1}^{M}\sum_{s=1}^{N_s} o_{m,s}.$$

The expectation value of this estimator, $\mathbb{E}[\langle\hat{O}\rangle]$, is taken over both the probabilistic choice of $\boldsymbol{l}_m$ and the shadow's random choice of measurement bases. We can analyze the expectation in steps, from the inside out using the *law of total expectation*. First, the inner expectation is a conditional expectation over the shadow measurement outcomes, given a fixed quasiprobability random circuit $\boldsymbol{C}_{\boldsymbol{l}_m}$. We denote this expectation by $\mathbb{E}_{\mathrm{sh}}$ and we have

$$\begin{aligned}\mathbb{E}_{\mathrm{sh}}\left[\frac{1}{N_s}\sum_{s=1}^{N_s} o_{m,s}\right] &= \frac{1}{N_s}\sum_{s=1}^{N_s}\mathbb{E}_{\mathrm{sh}}[o_{m,s}] \\ &= \langle\langle O|\Gamma_{\boldsymbol{l}_m}\boldsymbol{C}_{\boldsymbol{l}_m}(\rho)\rangle\rangle. \qquad \text{(E1)}\end{aligned}$$

Next, we consider the outer average over the M samples. The quasiprobability sampling protocol is constructed to be an unbiased estimator of the exact time-evolution channel, meaning $\mathbb{E}_{\boldsymbol{l}}[\Gamma_{\boldsymbol{l}_m}\boldsymbol{C}_{\boldsymbol{l}_m}] = \boldsymbol{C}$. By the linearity of expectation and trace, the full expectation value is

$$\begin{aligned}\mathbb{E}[\langle\hat{O}\rangle] &= \mathbb{E}_{\boldsymbol{l}}\left[\frac{1}{M}\sum_{m=1}^{M}\mathbb{E}_{\mathrm{sh}}\left[\frac{1}{N_s}\sum_{s=1}^{N_s} o_{m,s}\right]\right] \\ &= \frac{1}{M}\sum_{m=1}^{M}\mathbb{E}_{\boldsymbol{l}}[\langle\langle O|\Gamma_{\boldsymbol{l}_m}\boldsymbol{C}_{\boldsymbol{l}_m}(\rho)\rangle\rangle] \\ &= \langle\langle O|\mathbb{E}_{\boldsymbol{l}}[\Gamma_{\boldsymbol{l}_m}\boldsymbol{C}_{\boldsymbol{l}_m}](\rho)\rangle\rangle \\ &= \langle\langle O|\boldsymbol{C}(\rho)\rangle\rangle.\end{aligned}$$

■

ALGORITHM 2. Experimental implementation of shadow snapshots.

Input: Hamiltonian H, initial state ρ_{init}, total time t, Trotter steps K_{steps}, angle Δ, number of TE-PAI samples $M_{\text{TE-PAI}}$, number of shadow snapshots per sample N_s
Output: A collection of $N = M_{\text{TE-PAI}} \times N_s$ shadow snapshot tuples $(\Gamma, \{U_n\}_{n=1}^{N}, \{b_n\}_{n=1}^{N})$

Initialize an empty list of snapshots $\mathcal{S}$
for $m = 1$ to $M_{\text{TE-PAI}}$
  Initialize TE-PAI weight $\Gamma_m \leftarrow 1$
  Initialize an empty quantum circuit for evolution $\mathcal{C}_{\mathrm{evol}}$
  Set $\delta = t/K_{\mathrm{steps}}$
  for $k = 1$ to K_{steps}
    for $j = 1$ to J
      Compute $\theta_j = 2h_j\delta$ and coefficients a_1, a_2, a_3
      Set $\gamma = |a_1| + |a_2| + |a_3|$.
      Sample d with probability $|a_d|/\gamma$
      Update $\Gamma_m \leftarrow \Gamma_m \times \gamma \times \mathrm{sign}(a_d)$
      Append the corresponding gate to $\mathcal{C}_{\mathrm{evol}}$
    end for
  end for
  for $s = 1$ to N_s
    Create a copy: $\mathcal{C}_{\mathrm{total}} \leftarrow \mathcal{C}_{\mathrm{evol}}$
    for $n = 1$ to N
      Sample a Clifford gate U_n uniformly at random
      Append U_n to qubit n in $\mathcal{C}_{\mathrm{total}}$
    end for
    Apply $\mathcal{C}_{\mathrm{total}}$ to ρ_{init} and measure to get bitstring b
    Append the tuple $(\Gamma_m, \{U_n\}, b)$ to $\mathcal{S}$
  end for
end for
return $\mathcal{S}$

2. Proof of Theorem 2

Proof. Let Y_m denote the estimator obtained from the mth sampled circuit, averaged over its N_s shadow snapshots:

$$Y_m = \frac{1}{N_s}\sum_{s=1}^{N_s} o_{m,s}.$$

Since the M circuit samples are independent and identically distributed, the total variance is $\mathrm{Var}[\langle\hat{O}\rangle] = \frac{1}{M}\mathrm{Var}[Y_m]$. We apply the law of total variance to Y_m, conditioning on the choice of the random circuit $\boldsymbol{C}_{\boldsymbol{l}_m}$. Let $\mathbb{E}_{\boldsymbol{l}}$ and $\mathrm{Var}_{\boldsymbol{l}}$ denote the expectation and variance over the circuit distribution given by the quasiprobability sampling, and $\mathbb{E}_{\mathrm{sh}}$ and $\mathrm{Var}_{\mathrm{sh}}$ denote the conditional expectation and variance over the shadow measurement randomness, given a fixed random circuit $\boldsymbol{C}_{\boldsymbol{l}_m}$. The total variance decomposes as

$$\mathrm{Var}[Y_m] = \mathrm{Var}_{\boldsymbol{l}}(\mathbb{E}_{\mathrm{sh}}[Y_m]) + \mathbb{E}_{\boldsymbol{l}}[\mathrm{Var}_{\mathrm{sh}}(Y_m)].$$

Conditioned on a specific circuit $\boldsymbol{C}_m$, the shadow estimator is unbiased for that circuit's observable value as in Eq. (E1). From Lemma D1, the first term is upper bounded as

$$\mathrm{Var}_{\boldsymbol{l}}(\mathbb{E}_{\mathrm{sh}}[Y_m]) \leqslant \Gamma^2\|O\|^2 - \langle\langle O|\boldsymbol{C}(\rho)\rangle\rangle^2. \qquad \text{(E2)}$$

Conditioned on a fixed random circuit $\mathcal{C}_{\boldsymbol{l}_m}$, Y_m is the average of N_s independent snapshots. The variance of a single weighted

snapshot $\Gamma_m \mathrm{Tr}(O\hat{\rho})$ is determined by the second moment of the shadow estimator as $\mathbb{E}_{\mathrm{sh}}[\langle\langle O|\hat{\rho}\rangle\rangle^2] \leqslant \|O\|_{\mathrm{sh}}^2$. Thus, the variance of a single snapshot given $\mathcal{C}_m$ is

$$\mathrm{Var}_{\mathrm{sh}}(\Gamma_{I_m}\langle\langle O|\hat{\rho}_{m,s}\rangle\rangle) \leqslant \Gamma^2(\|O\|_{\mathrm{sh}}^2 - \langle\langle O|\boldsymbol{C}_{I_m}(\rho)\rangle\rangle^2).$$

Since Y_m averages N_s such snapshots, its conditional variance is scaled by $1/N_s$:

$$\mathrm{Var}_{\mathrm{sh}}(Y_m) = \frac{\Gamma^2}{N_s}(\|O\|_{\mathrm{sh}}^2 - \langle\langle O|\boldsymbol{C}_{I_m}(\rho)\rangle\rangle^2).$$

Taking the expectation over circuits yields the second term of the total variance:

$$\begin{aligned}\mathbb{E}_I[\mathrm{Var}_{\mathrm{sh}}(Y_m)] &= \frac{\Gamma^2}{N_s}\left(\|O\|_{\mathrm{sh}}^2 - \mathbb{E}_I[\langle\langle O|\boldsymbol{C}_{I_m}(\rho)\rangle\rangle^2]\right) \\ &\leqslant \frac{\Gamma^2}{N_s}\left(\|O\|_{\mathrm{sh}}^2 - \|O\|^2\right). \qquad \text{(E3)}\end{aligned}$$

Summing Eqs. (E2) and (E3) gives $\mathrm{Var}[Y_m]$:

$$\mathrm{Var}[Y_m] \leqslant \frac{\Gamma^2}{N_s}\|O\|_{\mathrm{sh}}^2 + \left(1 - \frac{1}{N_s}\right)\Gamma^2\|O\|^2.$$

Finally, dividing by M yields the variance of the total estimator $\langle\hat{O}\rangle$:

$$\mathrm{Var}[\langle\hat{O}\rangle] \leqslant \Gamma^2\left(\frac{\|O\|_{\mathrm{sh}}^2 - \|O\|^2}{MN_s} + \frac{\|O\|^2}{M}\right).$$

■

APPENDIX F: ENERGY SPECTRUM OF THE HEISENBERG HAMILTONIAN

In condensed matter physics, the isotropic Heisenberg model is typically defined using spin-1/2 operators $\mathbf{S}_i = \frac{1}{2}\boldsymbol{\sigma}_i$, where $\boldsymbol{\sigma} = (X, Y, Z)$ are the Pauli matrices [41,42]. The standard Hamiltonian is

$$\mathcal{H} = J\sum_{i=1}^{N-1}(\mathbf{S}i \cdot \mathbf{S}i + 1). \qquad \text{(F1)}$$

Our simulation employs Pauli operators directly, implying a scaling of $\boldsymbol{\sigma}_i \cdot \boldsymbol{\sigma}_{i+1} = 4(\mathbf{S}_i \cdot \mathbf{S}_{i+1})$. Consequently, all eigenvalues are scaled by a factor of 4 relative to the classic Bethe ansatz results [43]. For an infinite chain, the ground state energy E_0 follows the Hulthén integral [44], while low-lying excitations are described by des Cloizeaux-Pearson magnon triplets [45]. In our finite $N = 10$ system with open boundary conditions and $J = 1$, exact diagonalization reveals a discrete spectrum governed by SU(2) symmetry. The ground state and the tenth energy level, which resides in the third excited manifold, yield the following numerical results: $E_{\text{Ground state}} = -17.0321$ and $E_{10^{th}} = -12.6726$, resulting in the target transition energy gap for our state preparation, $\Delta E_{0,10} = E_{10} - E_0 \approx 4.3595$. This gap represents the energy required to sustain multimagnon excitations within the quantized momentum space of a ten-qubit chain. We emphasize that this is not an adjacent-level spacing. It is the transition energy gap selected by preparing a superposition of the ground state and the tenth energy level, and therefore it appears as the dominant frequency in the time-dependent signal.

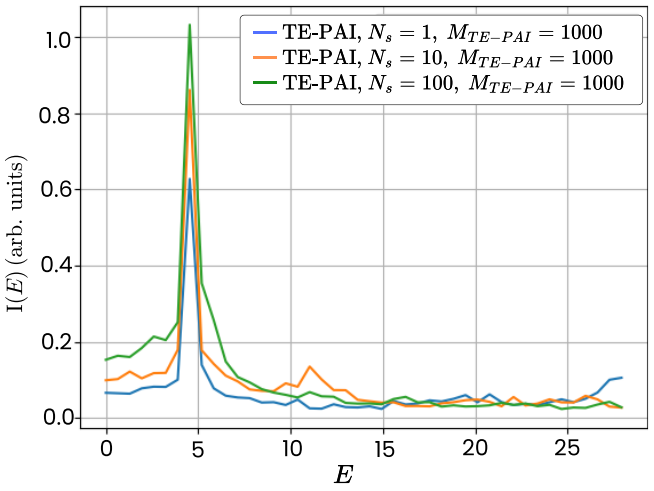


FIG. 6. Comparison of TE-PAI spectra computed with different shadow sizes. The spectra were generated using a fixed number of TE-PAI samples while varying the shadow size $N_s \in \{1, 10, 100\}$, corresponding to the blue, orange, and green spectra, respectively. While all cases resolve the target transition energy gap, increasing N_s significantly enhances the peak intensity. Simulation parameters match those detailed in Sec. IV A.

APPENDIX G: INFLUENCE OF THE SHADOW SIZE ON THE TE-PAI SPECTRUM

We now examine how shadow size influences the resulting spectrum. Figure 6 presents three TE-PAI spectra obtained using a fixed number of TE-PAI samples across varying shadow sizes, $N_s \in \{1, 10, 100\}$. All simulations utilized the parameters defined in Sec. IV A. While all three spectra successfully resolve the target transition energy gap, the peak intensity increases significantly with larger shadow sizes. However, a larger shadow size imposes a substantial computational overhead, as the total number of circuit executions scales with $M_{\text{TE-PAI}} \times N_s$. In this study, we fix the shadow size to $N_s = 1$ and only change the number of TE-PAI samples.

APPENDIX H: HARDWARE PROPERTIES

Table I summarizes the hardware characteristics of the two IBM Quantum backends used in this study: *ibm_kobe* and *ibm_kingston*. Both devices are based on the Heron r2 architecture and provide identical native basis gate sets, consisting of `cz`, `id`, `rx`, `rz`, `rzz`, `sx`, and `x`. This extensive gate set allows for greater flexibility in circuit design, particularly by enabling simulation of more complex Hamiltonians

TABLE I. Comparison of processor performance metrics between *ibm_kobe* and *ibm_kingston*.

Processor	*ibm_kobe*	*ibm_kingston*
2Q error (best)	1.03×10^{-3}	1.06×10^{-3}
2Q error (layered)	2.68×10^{-3}	3.64×10^{-3}
CLOPS	2.5×10^{5}	2.5×10^{5}
Median CZ error	1.858×10^{-3}	2.088×10^{-3}
Median SX error	2.383×10^{-4}	2.379×10^{-4}

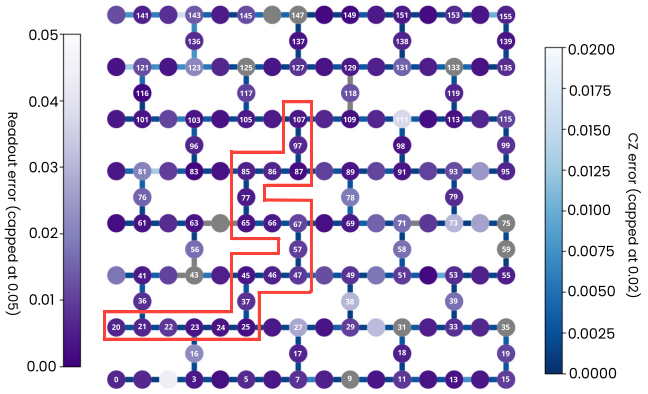


FIG. 7. Architecture and error landscape of the *ibm_kobe* quantum processor. Each node represents a qubit, with its color indicating the average readout assignment error. Each edge represents a physical connection between two qubits, and its color encodes the average controlled-Z (CZ) gate error. Gray nodes or edges indicate that the corresponding error exceeds a predefined threshold, set to 0.05 for the readout error and 0.02 for the CZ error. The qubits and connections highlighted by the red outline denote the specific subset of the device that is utilized in our experiments.

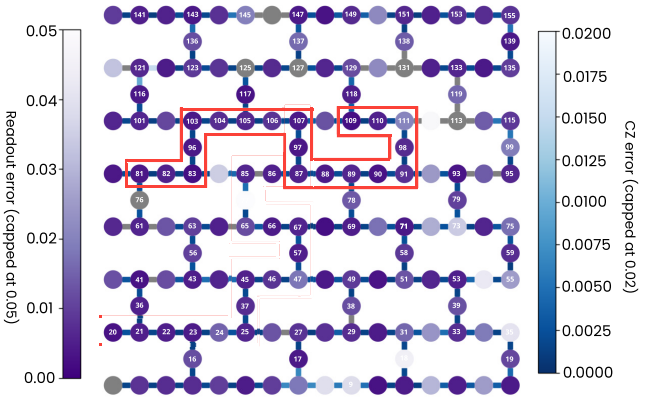


FIG. 8. Same as Fig. 7, but for the *ibm_kingston* quantum processor.

natively, without requiring excessive transpilation. Despite architectural similarities, the backends differ slightly in their two-qubit gate fidelity. *ibm_kobe* exhibits a lower two-qubit gate error, both in the best-case scenario (1.03×10^{-3}) and in layered execution (2.68×10^{-3}), compared to *ibm_kingston*, which shows 1.06×10^{-3} and 3.64×10^{-3}, respectively. This improved fidelity makes *ibm_kobe* a better candidate for experiments requiring two-qubit gates. Although the differences in fidelity are marginal, *ibm_kobe*'s slightly better two-qubit performance and its alignment with low transpilation depth make it our preferred choice for hardware-based simulations involving entanglement-heavy operations.

The logical-to-physical mapping is performed based on an optimization strategy prioritizing long T_1/T_2 and low error rates.

Figures 7 and 8 show diagram of the physical mapping of qubits.

[1] S. Sachdev, *Quantum Phase Transitions*, 2nd ed. (Cambridge University Press, Cambridge, UK, 2011).

[2] I. N. Levine, *Quantum Chemistry*, 6th ed. (Pearson Prentice Hall, London, 2009).

[3] P. A. M. Dirac, Quantum mechanics of many-electron systems, Proc. R. Soc. London Ser. A **123**, 714 (1929).

[4] S. Lloyd, Universal quantum simulators, Science **273**, 1073 (1996).

[5] Z. Zimborás, B. Koczor, Z. Holmes, E.-M. Borrelli, A. Gilyén, H.-Y. Huang, Z. Cai, A. Acín, L. Aolita, L. Banchi, *et al.*, Myths around quantum computation before full fault tolerance: What no-go theorems rule out and what they don't, arXiv:2501.05694.

[6] R. Babbush, R. King, S. Boixo, W. Huggins, T. Khattar, G. H. Low, J. R. McClean, T. O'Brien, and N. C. Rubin, The grand challenge of quantum applications, PRX Quantum **7**, 020101 (2026).

[7] A. Y. Kitaev, Quantum measurements and the Abelian stabilizer problem, arXiv:quant-ph/9511026.

[8] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*, 10th anniversary ed. (Cambridge University Press, Cambridge, UK, 2010), p. 702.

[9] H. H. S. Chan, R. Meister, M. L. Goh, and B. Koczor, Algorithmic shadow spectroscopy, PRX Quantum **6**, 010352 (2025).

[10] H.-Y. Huang, R. Kueng, and J. Preskill, Predicting many properties of a quantum system from very few measurements, Nat. Phys. **16**, 1050 (2020).

[11] H. Jnane, J. Steinberg, Z. Cai, H. C. Nguyen, and B. Koczor, Quantum error mitigated classical shadows, PRX Quantum **5**, 010324 (2024).

[12] S. Endo, S. C. Benjamin, and Y. Li, Practical quantum error mitigation for near-future applications, Phys. Rev. X **8**, 031027 (2018).

[13] C. Kiumi and B. Koczor, TE-PAI: Exact time evolution by sampling random circuits, Quantum Sci. Technol. **10**, 045071 (2025).

[14] B. Koczor, J. J. L. Morton, and S. C. Benjamin, Probabilistic interpolation of quantum rotation angles, Phys. Rev. Lett. **132**, 130602 (2024).

[15] B. Koczor, Sparse probabilistic synthesis of quantum operations, PRX Quantum **5**, 040352 (2024).

[16] Y. Kim, A. Eddins, S. Anand, K. X. Wei, E. van den Berg, S. Rosenblatt, H. Nayfeh, Y. Wu, M. Zaletel, K. Temme, and A. Kandala, Evidence for the utility of quantum computing before fault tolerance, Nature (London) **618**, 500 (2023).

[17] A. Gilchrist, N. K. Langford, and M. A. Nielsen, Distance measures to compare real and ideal quantum processes, Phys. Rev. A **71**, 062310 (2005).

[18] M. Suzuki, Fractal decomposition of exponential operators with applications to many-body theories and Monte Carlo simulations, Phys. Lett. A **146**, 319 (1990).

[19] M. Suzuki, General theory of fractal path integrals with applications to many-body theories and statistical physics, J. Math. Phys. **32**, 400 (1991).

[20] K. Temme, S. Bravyi, and J. M. Gambetta, Error mitigation for short-depth quantum circuits, Phys. Rev. Lett. **119**, 180509 (2017).
[21] 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).
[22] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O'Brien, A variational eigenvalue solver on a photonic quantum processor, Nat. Commun. **5**, 4213 (2014).
[23] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, and P. J. Coles, Variational quantum algorithms, Nat. Rev. Phys. **3**, 625 (2021).
[24] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, Barren plateaus in quantum neural network training landscapes, Nat. Commun. **9**, 4812 (2018).
[25] Y. Matsuzaki, H. Hakoshima, K. Sugisaki, Y. Seki, and S. Kawabata, Direct estimation of the energy gap between the ground state and excited state with quantum annealing, Jpn. J. Appl. Phys. **60**, SBBI02 (2021).
[26] K. Gnatenko, H. Laba, and V. Tkachuk, Energy levels estimation on a quantum computer by evolution of a physical quantity, Phys. Lett. A **424**, 127843 (2022).
[27] A. E. Russo, K. M. Rudinger, B. C. A. Morrison, and A. D. Baczewski, Evaluating energy differences on a quantum computer with robust phase estimation, Phys. Rev. Lett. **126**, 210501 (2021).
[28] H. Kuji, Y. Shingu, T. Nikuni, T. Imoto, K. Sugisaki, and Y. Matsuzaki, Robust phase estimation of the ground-state energy without controlled time evolution on a quantum device, Phys. Rev. A **111**, 042618 (2025).
[29] R. Toshio, S. Kanasugi, J. Fujisaki, H. Oshima, S. Sato, and K. Fujii, STAR-Magic Mutation: Even more efficient analog rotation gates for early fault-tolerant quantum computer, arXiv:2603.22891.
[30] Qiskit contributors, Qiskit SDK API reference: QuantumCircuit.depth, 2025, accessed 15 January 2026, https://docs.quantum.ibm.com/api/qiskit/qiskit.circuit.QuantumCircuit#depth.
[31] P. Pfeuty, The one-dimensional Ising model with a transverse field, Ann. Phys. **57**, 79 (1970).
[32] A. M. Childs, A. Ostrander, and Y. Su, Faster quantum simulation by randomization, Quantum **3**, 182 (2019).
[33] H. Pages *et al.*, TE-PAI_shadow_spectroscopy, GitHub repository, 2026, https://github.com/hugopgs/TE-PAI_shadow_spectroscopy.
[34] T. Peng, A. W. Harrow, M. Ozols, and X. Wu, Simulating large quantum circuits on a small quantum computer, Phys. Rev. Lett. **125**, 150504 (2020).
[35] K. Mitarai and K. Fujii, Constructing a virtual two-qubit gate by sampling single-qubit operations, New J. Phys. **23**, 023021 (2021).
[36] A. W. Harrow and A. Lowe, Optimal quantum circuit cuts with application to clustered Hamiltonian simulation, PRX Quantum **6**, 010316 (2025).
[37] J. W. Dai, F. Hasselgren, and C. Kiumi, Structure-aware variance reduction for unbiased randomized Hamiltonian simulation, arXiv:2606.23544.
[38] G. E. P. Box and D. A. Pierce, Distribution of residual autocorrelations in autoregressive-integrated moving average time series models, J. Am. Stat. Assoc. **65**, 1509 (1970).
[39] G. M. Ljung and G. E. P. Box, On a measure of lack of fit in time series models, Biometrika **65**, 297 (1978).
[40] R. Bronson and G. B. Costa, Eigenvalues and eigenvectors, in *Matrix Methods*, 4th ed., edited by R. Bronson and G. B. Costa (Academic Press, London, UK, 2021), Chap. 6, pp. 185–223.
[41] M. Takahashi, *Thermodynamics of One-Dimensional Solvable Models* (Cambridge University Press, Cambridge, UK, 1999).
[42] D. C. Mattis, *The Theory of Magnetism I: Statics and Dynamics* (Springer, Berlin, Germany, 1981).
[43] H. Bethe, Zur Theorie der Metalle, Z. Phys. **71**, 205 (1931).
[44] L. Hulthén, Über das Austauschproblem eines Kristalls, Arkiv für Mathematik, Astronomi och Fysik **26A**, 1 (1938).
[45] J. des Cloizeaux and J. J. Pearson, Spin-wave spectrum of the antiferromagnetic linear chain, Phys. Rev. **128**, 2131 (1962).