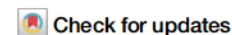




Group-theoretic error mitigation enabled by classical shadows and symmetries



Andrew Zhao^{1,2}✉ & Akimasa Miyake¹

Estimating expectation values is a key subroutine in quantum algorithms. Near-term implementations face two major challenges: a limited number of samples required to learn a large collection of observables, and the accumulation of errors in devices without quantum error correction. To address these challenges simultaneously, we develop a quantum error-mitigation strategy called *symmetry-adjusted classical shadows*, by adjusting classical-shadow tomography according to how symmetries are corrupted by device errors. As a concrete example, we highlight global U(1) symmetry, which manifests in fermions as particle number and in spins as total magnetization, and illustrate their group-theoretic unification with respective classical-shadow protocols. We establish rigorous sampling bounds under readout errors obeying minimal assumptions, and perform numerical experiments with a more comprehensive model of gate-level errors derived from existing quantum processors. Our results reveal symmetry-adjusted classical shadows as a low-cost strategy to mitigate errors from noisy quantum experiments in the ubiquitous presence of symmetry.

Quantum computers are highly susceptible to errors at the hardware level, posing a considerable challenge to realize meaningful applications in the so-called noisy intermediate-scale quantum (NISQ) era^{1,2}. One particularly promising and natural candidate for NISQ applications is the simulation of quantum many-body physics and chemistry^{3–6}. In order to minimize the accumulation of errors, such algorithms prioritize low-depth circuits, for instance, variational quantum circuits^{7–10}. However, in order to exhibit quantum advantage, these circuits must also be beyond the capabilities of classical simulation^{11–14}, resulting in noise levels that nonetheless corrupt the calculations.

While quantum error correction is the long-term solution, current state-of-the-art hardware is still a few orders of magnitude from achieving scalable, fault-tolerant quantum computation^{15–23}. In the meantime, there have been considerable theoretical and experimental efforts probing the beyond-classical potential of NISQ computers^{24–46}. Should such an application be demonstrated, quantum error mitigation (QEM) is expected to play a crucial role. Broadly speaking, QEM aims to approximately recover the output of an ideal quantum computation, given only access to noisy quantum devices and offline classical resources. We refer the reader to refs. 47,48 for a review of prominent concepts and strategies in QEM.

A related but separate challenge for NISQ algorithms is the need to learn many observables in a rudimentary fashion, i.e., by repeatedly running and sampling from quantum circuits. The number of repetitions required can be immense, both to suppress shot noise and to handle the measurement of noncommuting observables^{49,50}. One particularly promising approach is

that of classical shadows^{51,52}. In contrast to prior measurement strategies^{10,53–55}, classical shadows are remarkably simple to implement and have been shown to exhibit optimal sample complexity in certain important scenarios^{51,56}.

Classical shadows were developed primarily from the union of two themes in quantum learning theory: linear-inversion estimators for state tomography^{57,58} (closed-form solutions that admit fast postprocessing and rigorous guarantees) and the framework of shadow tomography^{59,60} (predict only a subset of observables, not the entire density matrix). The result is a simple but powerful protocol that accurately estimates a large collection of observables from relatively few samples. In terms of quantum resources, classical shadows only require the ability to measure in randomly selected bases, making the protocol particularly amenable to NISQ constraints. These desirable features have inspired a wide range of extensions and applications, for example: entanglement detection⁶¹, quantum Fisher information bounds^{62,63}, learning quantum processes^{64,65}, navigating variational landscapes^{66,67}, energy-gap estimation⁶⁸, and applications to fermions^{69–74} and bosons^{75,76}. For an overview of classical shadows and randomized measurement strategies, see ref. 77.

Due to their experimental friendliness and versatile prediction power, classical shadows naturally have been considered for QEM as well. For example, refs. 78,79 used classical shadows to approximately project a noisy quantum state toward a target subspace via classical postprocessing, the subspaces being either the logical subspace of an error-correcting code⁸⁰ and/or the dominant eigenvector (purification) of the noisy mixed state^{81,82}. These

¹Center for Quantum Information and Control, Department of Physics and Astronomy, University of New Mexico, Albuquerque, NM 87106, USA. ²Present address: Sandia National Laboratories, Livermore, CA 94550, USA. ✉e-mail: azhao@sandia.gov

shadow-based ideas circumvent some of the difficulties of performing subspace projection, at the cost of an exponential sample complexity. Meanwhile, ref. 83 intertwined classical shadows with other popular QEM strategies, with a particular focus on probabilistic error cancellation⁸⁴. They establish rigorous estimators and performance guarantees, assuming an accurate characterization of the noisy quantum device. Finally, refs. 85,86 described modifications to the classical linear-inversion step in order to mitigate errors in the randomized measurements. In particular, robust shadow estimation⁸⁵ assumes no prior knowledge of the noise, instead implementing a separate calibration experiment that learns the necessary noise features.

In this paper, we take this latter perspective^{85,86}, with an eye on a more comprehensive mitigation of errors beyond readout errors. We introduce a QEM protocol, which we refer to as *symmetry-adjusted classical shadows*, that takes advantage of known symmetries in the quantum system of interest. For example, in simulations of chemistry, the number of electrons is typically fixed. The corruption of such symmetries by noise informs us how to undo the effects of that noise. Crucially, because randomized measurements scramble the information, the other properties of the quantum system are corrupted (and therefore can be mitigated) in the same manner. Using these insights, symmetry-adjusted classical shadows appropriately modifies the linear-inversion based on the symmetry information alone.

A notable advantage of our protocol is that we do not run any extraneous calibration experiments. This has the added benefit of inherently accounting for errors that occur throughout the full quantum circuit, rather than the randomized measurements in isolation^{85–89}. Also, the simplicity of the protocol allows for additional QEM techniques to be straightforwardly applied in tandem. Finally, in contrast to other symmetry-based ideas^{83,90–92}, our approach goes beyond the concept of symmetry projection, instead utilizing a unified group-theoretic understanding of classical shadows in conjunction with symmetries. We expound on this distinction in the Supplementary Discussion, wherein we review these prior symmetry-based QEM techniques.

This paper is structured as follows. In the Results section, we begin by establishing preliminaries and background material. We then introduce our main contribution, symmetry-adjusted classical shadows, and describe its key application for mitigating local fermionic and qubit observables. We follow by highlighting additional technical results: a modification to random Pauli measurements required to tailor its irreps for use with common symmetries, called subsystem-symmetrized Pauli shadows; an improved design for compiling fermionic Gaussian unitaries with lower circuit depth and fewer gates than prior art; and a symmetry adaptation to fermionic classical shadows which reduces the quantum resources required, applicable to fermionic systems with spin symmetry. Finally, we close the Results section with a series of numerical experiments, demonstrating the effectiveness of our error-mitigation protocol under realistic scenarios. This includes simulations of a noise model based on existing superconducting-qubit platforms⁹³. In the Discussion section, we summarize our findings and discuss future prospects. In the Methods section, we illustrate the general theory of symmetry-adjusted classical shadows, and we provide further technical details regarding the applications to fermion and qubit systems with global U(1) symmetries. Details regarding the mathematical proofs and numerical simulations are provided in the Supplementary Information, and code for the latter is available at our open-source repository (<https://github.com/zhao-andrew/symmetry-adjusted-classical-shadows>).

Results

Background

First, we provide a review of classical shadows^{51,52} and robust shadow estimation⁸⁵ necessary to understand our technical results. Readers familiar with this background material can skip to the subsection “Symmetry-adjusted classical shadows,” after familiarizing themselves with the notation that we establish below.

Notation and preliminaries. For any integer $N > 1$, we define $[N] := \{0, \dots, N - 1\}$ (note that we index starting from 0). We use $i \equiv \sqrt{-1}$ for the imaginary unit.

Throughout this paper, we consider an n -qubit system with Hilbert space $\mathcal{H} := (\mathbb{C}^2)^{\otimes n}$. Its dimension is denoted by $d \equiv 2^n$ unless otherwise specified. We often work with the space of linear operators $\mathcal{L}(\mathcal{H}) \cong \mathbb{C}^{d \times d}$ as a vector space, so it will be convenient to employ the Liouville representation: for any operator $A \in \mathcal{L}(\mathcal{H})$, its vectorization $|A\rangle \in \mathbb{C}^{d^2}$ in some orthonormal operator basis $\{B_1, \dots, B_{d^2} : \text{tr}(B_i^\dagger B_j) = \delta_{ij}\}$ is defined by the components $\langle\langle B_i | A \rangle\rangle := \text{tr}(B_i^\dagger A)$. Under this representation, superoperators are mapped to $d^2 \times d^2$ matrices: any $\mathcal{E} \in \mathcal{L}(\mathcal{L}(\mathcal{H}))$ can be specified by its matrix elements $\mathcal{E}_{ij} := \langle\langle B_i | \mathcal{E} | B_j \rangle\rangle = \text{tr}(B_i^\dagger \mathcal{E}(B_j))$. We let \mathcal{E} denote both the superoperator and its matrix representation, and in a similar fashion we sometimes write $|A\rangle$ = A .

For systems of qubits, the normalized Pauli operators $\mathcal{P}(n)/\sqrt{d}$ are a convenient basis for $\mathcal{L}(\mathcal{H})$, where

$$\mathcal{P}(n) := \{\mathbb{I}, X, Y, Z\}^{\otimes n}. \tag{1}$$

This choice is called the Pauli transfer matrix (PTM) representation. The weight, or locality, of a Pauli operator $P \in \mathcal{P}(n)$ is the number of its non-trivial tensor factors, denoted by $|P|$. For each $i \in [n]$, we define $W_i \in \mathcal{P}(n)$ which acts as $W \in \{X, Y, Z\}$ on the i th qubit and trivially on the rest of the system.

For fermions in second quantization, a natural choice of basis is the set of Majorana operators, defined as $\{\Gamma_\mu/\sqrt{d} : \mu \subseteq [2n]\}$ where

$$\Gamma_\mu := (-i)^{\binom{|\mu|}{2}} \prod_{\mu \in \mu} \gamma_\mu. \tag{2}$$

The Hermitian generators $\{\gamma_\mu : \mu \subseteq [2n]\} \subset \mathcal{L}(\mathcal{H})$ obey the anticommutation relation $\gamma_\mu \gamma_\nu + \gamma_\nu \gamma_\mu = 2\delta_{\mu\nu} \mathbb{I}$ (we will use \mathbb{I} to denote any identity operator whose dimension is clear from context). They are related to the fermionic creation and annihilation operators a_p^\dagger, a_p via

$$\gamma_{2p} = a_p + a_p^\dagger, \quad \gamma_{2p+1} = -i(a_p - a_p^\dagger). \tag{3}$$

By convention, the elements of μ and the product in Eq. (2) are in strictly ascending order. We call $|\mu|$ the degree of Γ_μ , or equivalently refer to them as $(|\mu|/2)$ -body operators whenever the degree is even. It is straightforward to check that Majorana operators are isomorphic to Pauli operators, in particular satisfying the orthogonality relation $\langle\langle \Gamma_\mu | \Gamma_\nu \rangle\rangle = d\delta_{\mu\nu}$.

For any unitary $U \in U(d)$, its corresponding channel is denoted by $\mathcal{U}(\cdot) := U(\cdot)U^\dagger$. For any $|\varphi\rangle \in \mathcal{H}$, $|\varphi\rangle\rangle$ is the vectorization of $|\varphi\rangle\langle\varphi|$. We use tildes to indicate objects affected by quantum noise, e.g., $\tilde{\mathcal{U}}$ denotes a noisy implementation of the \mathcal{U} . Hats indicate statistical estimators, e.g., \hat{o} denotes an estimate for $o = \text{tr}(O\rho)$. Asymptotic upper and lower bounds are denoted by $\mathcal{O}(\cdot)$ and $\Omega(\cdot)$ respectively, and $f(x) = \Theta(g(x))$ means that $f(x)$ is both $\mathcal{O}(g(x))$ and $\Omega(g(x))$.

Classical shadows. We summarize the method of classical shadows as formalized by Huang et al.⁵¹, borrowing the PTM language of Chen et al.⁸⁵ which will make the robust extension clear later. Our task is to estimate the expectation values $\text{tr}(O_j\rho) = \langle\langle O_j | \rho \rangle\rangle$ of a collection of L observables $O_1, \dots, O_L \in \mathcal{L}(\mathcal{H})$, ideally using as few copies of ρ as possible. Classical shadows is based on a simple measurement primitive: for each copy of ρ , apply a unitary U randomly drawn from a distribution of unitaries and measure in the computational basis. This produces a sample $b \in \{0, 1\}^n$ with probability $\langle\langle b | \mathcal{U} | \rho \rangle\rangle$. One then inverts the unitary on the outcome $|b\rangle$ in postprocessing, which amounts to storing a classical representation of $U^\dagger|b\rangle$.

The unitary distribution determines the efficiency of this protocol with respect to the properties of interest. Throughout this paper, we assume that the distribution is a finite group equipped with the uniform probability distribution (it is straightforward to generalize to compact groups, using their Haar measures). Specifically, let $U : G \rightarrow U(\mathcal{H})$ be a unitary

representation of a group G . The measurement primitives averaged over all random unitaries and measurement outcomes implement the quantum channel

$$\mathcal{M} := \mathbb{E}_{g \sim G} \mathcal{U}_g^\dagger \mathcal{M}_Z \mathcal{U}_g \equiv \frac{1}{|G|} \sum_{g \in G} \mathcal{U}_g^\dagger \mathcal{M}_Z \mathcal{U}_g, \quad (4)$$

where

$$\mathcal{M}_Z = \sum_{b \in \{0,1\}^n} |b\rangle\langle b| \quad (5)$$

describes the effective process of computational-basis measurements. The channel \mathcal{U}_g is the random unitary acting on the target state ρ , while \mathcal{U}_g^\dagger is its classically computed inversion on the measurement outcomes $|b\rangle$. Thus in expectation we produce the state

$$\mathcal{M}|\rho\rangle\rangle = \mathbb{E}_{g \sim G, b \sim \mathcal{U}_g|\rho\rangle} \mathcal{U}_g^\dagger |b\rangle. \quad (6)$$

If \mathcal{M} is invertible (corresponding to informational completeness of the measurement primitive), then applying \mathcal{M}^{-1} to Eq. (6) recovers the state:

$$|\rho\rangle\rangle = \mathcal{M}^{-1} \mathcal{M}|\rho\rangle\rangle = \mathbb{E}_{g \sim G, b \sim \mathcal{U}_g|\rho\rangle} \mathcal{M}^{-1} \mathcal{U}_g^\dagger |b\rangle. \quad (7)$$

The objects $|\hat{\rho}_{g,b}\rangle\rangle := \mathcal{M}^{-1} \mathcal{U}_g^\dagger |b\rangle$ are called the classical shadows of $|\rho\rangle\rangle$, for which they serve as unbiased estimators. Hence by construction they can predict expectation values,

$$\mathbb{E}_{g \sim G, b \sim \mathcal{U}_g|\rho\rangle} \langle\langle O_j | \hat{\rho}_{g,b} \rangle\rangle = \langle\langle O_j | \rho \rangle\rangle, \quad (8)$$

as well as nonlinear functions of ρ^{51} . While \mathcal{M}^{-1} is not a physical map (it is not completely positive), it only appears as classical postprocessing. Such a computation can be accomplished, for instance, by first deriving a closed-form expression for \mathcal{M} .

One systematic approach to deriving such an expression is through the representation theory of G . First, note that the d -dimensional unitary U is promoted to a d^2 -dimensional representation \mathcal{U} . Equation (4) reveals that \mathcal{M} is a twirl of \mathcal{M}_Z by the group G under the action of \mathcal{U} . Such objects are well studied: assuming that the irreducible components of \mathcal{U} have no multiplicities, an application of Schur’s lemma implies that⁹⁴

$$\mathcal{M} = \sum_{\lambda \in R_G} f_\lambda \Pi_\lambda. \quad (9)$$

Note that the general expression with multiplicities can be found in [ref. 85, Eq. (A6)]. Here, R_G is the set of labels λ for the irreducible representations (irreps) of G . The superoperators Π_λ are orthogonal projectors onto the irreducible subspaces $V_\lambda \subseteq \mathcal{L}(\mathcal{H})$. Choosing an orthonormal basis $\{|B_\lambda^j\rangle\rangle : j = 1, \dots, \dim V_\lambda\}$ for each subspace, we can write the projectors as

$$\Pi_\lambda = \sum_{j=1}^{\dim V_\lambda} |B_\lambda^j\rangle\rangle\langle\langle B_\lambda^j|. \quad (10)$$

The eigenvalues f_λ of \mathcal{M} can be computed using the orthogonality of projectors:

$$f_\lambda = \frac{\text{tr}(\mathcal{M}_Z \Pi_\lambda)}{\text{tr}(\Pi_\lambda)}. \quad (11)$$

Note that $\text{tr}(\Pi_\lambda) = \dim V_\lambda$. From this diagonalization, we immediately acquire an expression for the desired inverse:

$$\mathcal{M}^{-1} = \sum_{\lambda \in R_G} f_\lambda^{-1} \Pi_\lambda. \quad (12)$$

If some $f_\lambda = 0$, then we may instead define \mathcal{M}^{-1} as the pseudoinverse on the subspaces where f_λ is nonvanishing. This implies that the measurement primitive is informationally complete only within those subspaces.

To analyze the sample efficiency of this protocol, suppose we have performed T experiments, yielding a collection of independent classical shadows $\hat{\rho}_1, \dots, \hat{\rho}_T$ where each $|\hat{\rho}_\ell\rangle\rangle = \mathcal{M}^{-1} \mathcal{U}_{g_\ell}^\dagger |b_\ell\rangle\rangle$. From this data we can construct estimates

$$\hat{o}_j(T) = \frac{1}{T} \sum_{\ell=1}^T \langle\langle O_j | \hat{\rho}_\ell \rangle\rangle, \quad (13)$$

which by linearity converge to $\text{tr}(O_j \rho)$. The single-shot variance of \hat{o}_j can be bounded in terms of the so-called shadow norm:

$$\begin{aligned} \text{Var}[\hat{o}_j] &\leq \max_{\text{states } \sigma} \mathbb{E}_{g \sim G, b \sim \mathcal{U}_g|\sigma\rangle} \langle\langle O_j | \mathcal{M}^{-1} \mathcal{U}_g^\dagger |b\rangle\rangle^2 \\ &=: \|\mathcal{O}_j\|_{\text{shadow}}^2. \end{aligned} \quad (14)$$

This variance controls the prediction error, rigorously established via probability tail bounds. In particular, taking a number of samples

$$T = \mathcal{O}\left(\frac{\log(L/\delta)}{\epsilon^2} \max_{1 \leq j \leq L} \|\mathcal{O}_j\|_{\text{shadow}}^2\right) \quad (15)$$

ensures that, with probability at least $1 - \delta$, each estimate exhibits at most ϵ additive error:

$$|\hat{o}_j(T) - \langle\langle O_j | \rho \rangle\rangle| \leq \epsilon. \quad (16)$$

Note that for simplicity we employ the mean estimator throughout this paper, which suffices whenever the ensemble is either local Cliffords or matchgates and the observables are Pauli or Majorana operators [ref. 69, Supplemental Material, Theorem 12]. In general, a median-of-means estimator can guarantee the advertised sample complexity regardless of ensemble.

Finally, we comment on the classical computation of \hat{o}_j . In order to evaluate Eq. (13), one may use Eqs. (10) and (12) to express the ℓ th-sample estimate as

$$\langle\langle O_j | \hat{\rho}_\ell \rangle\rangle = \sum_{\lambda \in R_G} f_\lambda^{-1} \sum_{k=1}^{\dim V_\lambda} \langle\langle O_j | B_\lambda^k \rangle\rangle \langle\langle B_\lambda^k | \mathcal{U}_{g_\ell}^\dagger |b_\ell\rangle\rangle. \quad (17)$$

Thus it suffices to be able to efficiently compute the expansion coefficients $\langle\langle O_j | B_\lambda^k \rangle\rangle = \text{tr}(O_j B_\lambda^k)$ of the observable O_j in a basis of V_λ , as well as the matrix elements $\langle\langle B_\lambda^k | \mathcal{U}_g^\dagger |b\rangle\rangle = \langle b | U_g (B_\lambda^k)^\dagger U_g^\dagger |b\rangle$. Note that this does not require explicitly representing the classical shadow $\mathcal{M}^{-1} \mathcal{U}_g^\dagger |b\rangle\rangle$; we only need to determine the diagonal entry of the rotated operator $U_g (B_\lambda^k)^\dagger U_g^\dagger$ for a given basis state $|b\rangle$.

Robust shadow estimation. We now summarize the robust shadow estimation protocol by Chen et al.⁸⁵; we note that refs. 87–89 describe analogous ideas in the case of random single-qubit measurements. The basic premise is the fact that Schur’s lemma applies to the twirl of any channel, not just \mathcal{M}_Z . Suppose that instead of \mathcal{U}_g , the quantum computer implements a noisy channel $\tilde{\mathcal{U}}_g$ which obeys the following assumptions:

Assumptions 1. ([ref.⁸⁵, Simplifying noise assumption A1]). The noise in \tilde{U}_g is gate independent, time stationary, and Markovian. Hence there exists the decomposition $\tilde{U}_g = \mathcal{E}U_g$, where \mathcal{E} is a completely positive, trace-preserving map, independent of both the ideal unitary and the experimental time.

They also assume the ability to prepare the state $|0^n\rangle$ with sufficiently high fidelity. Given these conditions, the noisy version of the shadow channel implemented in experiment becomes

$$\tilde{\mathcal{M}} := \mathbb{E}_{g \sim G} U_g^\dagger \mathcal{M}_Z \tilde{U}_g = \frac{1}{|G|} \sum_{g \in G} U_g^\dagger \mathcal{M}_Z \mathcal{E} U_g, \tag{18}$$

which is now a twirl over the composite channel $\mathcal{M}_Z \mathcal{E}$. Although \mathcal{E} is unknown, Schur’s lemma implies that the eigenbasis is preserved, as we now have

$$\tilde{\mathcal{M}} = \sum_{\lambda \in R_G} \tilde{f}_\lambda \Pi_\lambda, \tag{19}$$

where the eigenvalues depend on \mathcal{E} ,

$$\tilde{f}_\lambda = \frac{\text{tr}(\mathcal{M}_Z \mathcal{E} \Pi_\lambda)}{\text{tr}(\Pi_\lambda)}. \tag{20}$$

Therefore if one knows \tilde{f}_λ , then one can perform the correct linear inversion in the presence of noise, i.e., by replacing f_λ^{-1} with \tilde{f}_λ^{-1} in Eq. (17).

Because \mathcal{E} depends on the details of the quantum hardware, it is not possible to determine \tilde{f}_λ without an a priori accurate characterization of the noise. Absent such information, a calibration protocol is proposed to experimentally estimate the value of \tilde{f}_λ . This proceeds by performing the classical shadows protocol on a fiducial state $|0^n\rangle$, rather than the unknown target state ρ . This enables the study of errors in the random circuits U_g . Because $|0^n\rangle$ is known exactly, one can compare its noiseless properties against the noisy experimental data to determine a calibration factor.

Specifically, Chen et al.⁸⁵ construct an estimator $\text{NoiseEst}_G(\lambda, g, b)$ for each sample (U_g, b) of the calibration experiment, which converges to \tilde{f}_λ in expectation over g and b . Although they do not prescribe a generic expression for NoiseEst_G (instead considering particular choices of G), it is straightforward to derive one following their ideas. Let $D_\lambda \in V_\lambda$ be an observable supported exclusively by a single irrep such that $\langle 0^n | D_\lambda | 0^n \rangle \neq 0$. Then we have

$$\langle D_\lambda | \tilde{\mathcal{M}} | 0^n \rangle = \tilde{f}_\lambda \langle 0^n | D_\lambda | 0^n \rangle. \tag{21}$$

On the other hand, using the fact that

$$\begin{aligned} \langle D_\lambda | \tilde{\mathcal{M}} | 0^n \rangle &= \langle D_\lambda | \mathbb{E}_{g \sim G, b \sim \mathcal{U}_g | 0^n} U_g^\dagger | b \rangle \\ &= \mathbb{E}_{g \sim G, b \sim \mathcal{U}_g | 0^n} \langle b | U_g D_\lambda U_g^\dagger | b \rangle, \end{aligned} \tag{22}$$

it follows that the random variable

$$\text{NoiseEst}_G(\lambda, g, b) = \frac{\langle b | U_g D_\lambda U_g^\dagger | b \rangle}{\langle 0^n | D_\lambda | 0^n \rangle} \tag{23}$$

obeys $\mathbb{E}_{g, b} [\text{NoiseEst}_G(\lambda, g, b)] = \tilde{f}_\lambda$.

One can recover the definitions for NoiseEst_G introduced by Chen et al.⁸⁵ as follows. The global Clifford group $\text{Cl}(n)$ has two irreps: the span of the identity operator, $V_0 = \text{span}\{\mathbb{I}\}$ (which is trivial), and its orthogonal complement $V_1 = V_0^\perp$ (the set of all traceless operators). Choosing $D_1 =$

$d|0^n\rangle\langle 0^n| - \mathbb{I}$ gives

$$\text{NoiseEst}_{\text{Cl}(n)}(1, U, b) = \frac{d|\langle b | U | 0^n \rangle|^2 - 1}{d - 1}, \tag{24}$$

where $U \in \text{Cl}(n)$.

On the other hand, the local Clifford group $\text{Cl}(1)^{\otimes n}$ has 2^n irreps, labeled by all subsets $I \subseteq [n]$. Each I indexes a subsystem of qubits, and each subspace V_I is the span of all n -qubit Pauli operators which act nontrivially on exactly that subsystem. Defining

$$D_I := \prod_{i \in I} Z_i, \tag{25}$$

one obtains

$$\begin{aligned} \text{NoiseEst}_{\text{Cl}(1)^{\otimes n}}(I, U, b) &= \frac{\langle b | U D_I U^\dagger | b \rangle}{\langle 0^n | D_I | 0^n \rangle} \\ &= \prod_{i \in I} \langle b_i | C_i Z C_i^\dagger | b_i \rangle \end{aligned} \tag{26}$$

where now $U = \bigotimes_{i \in [n]} C_i \in \text{Cl}(1)^{\otimes n}$.

Any QEM strategy necessarily incurs a sampling overhead dependent on the amount of noise^{95–98}. For global Clifford shadows, Chen et al.⁸⁵ show that the sample complexity is augmented by a factor of $\mathcal{O}(F_Z(\mathcal{E})^{-2})$ for estimating observables with constant Hilbert–Schmidt norm, where $F_Z(\mathcal{E}) = 2^{-n} \sum_{b \in \{0,1\}^n} \langle b | \mathcal{E} | b \rangle$ is the average Z -basis fidelity of \mathcal{E} . Meanwhile for local Clifford shadows, they prove that product noise of the form $\mathcal{E} = \bigotimes_{i \in [n]} \mathcal{E}_i$, satisfying $\min_{i \in [n]} F_Z(\mathcal{E}_i) \geq 1 - \xi$, exhibits an overhead factor of $e^{\mathcal{O}(k\xi)}$ for estimating k -local qubit observables.

Symmetry-adjusted classical shadows

The primary contribution of this paper, symmetry-adjusted classical shadows, is visualized in Fig. 1. We describe it in detail now. Consider a classical shadows protocol over G with target observables O_1, \dots, O_L . Without loss of generality, let each $O_j \in V_\lambda$ for some subset of irreps $\lambda \in R' \subseteq R_G$. Suppose the experiment experiences an unknown noise channel \mathcal{E} obeying Assumptions 1.

We show that, if ρ obeys symmetries which are “compatible” with the irreps in R' , then it is possible to construct an estimator which accurately predicts the ideal, noiseless observables. By compatible, we mean that there exist symmetry operators $S_\lambda \in V_\lambda$ for each $\lambda \in R'$ for which their ideal expectation values

$$s_\lambda := \text{tr}(S_\lambda \rho) \tag{27}$$

are known a priori. In general, there is no reason to expect that a physical system has symmetries which exactly fit into the irreps of a classical-shadow measurement scheme. However, given a symmetry operator S , it is always possible to project it to V_λ using the superoperator projector Π_λ , i.e., $S_\lambda = \Pi_\lambda(S)$.

Then, using noisy classical shadows $\hat{\rho}(T)$ of size T , we construct error-mitigated estimates as

$$\hat{o}_j^{\text{EM}}(T) := \frac{\text{tr}(O_j \hat{\rho}(T))}{\text{tr}(S_\lambda \hat{\rho}(T)) / s_\lambda}. \tag{28}$$

We find that the relevant noise characterization in this scenario is

$$F_{Z, R'}(\mathcal{E}) := \min_{\lambda \in R'} \frac{\text{tr}(\mathcal{E} \mathcal{M}_Z \Pi_\lambda)}{\text{tr}(\mathcal{M}_Z \Pi_\lambda)}, \tag{29}$$

which can be seen as a generalization of the noise fidelity $F_Z(\mathcal{E})$ described in the “Background” subsection. Here, $F_{Z, R'}(\mathcal{E})$ only considers how the noise channel acts within the irreducible subspaces of interest.

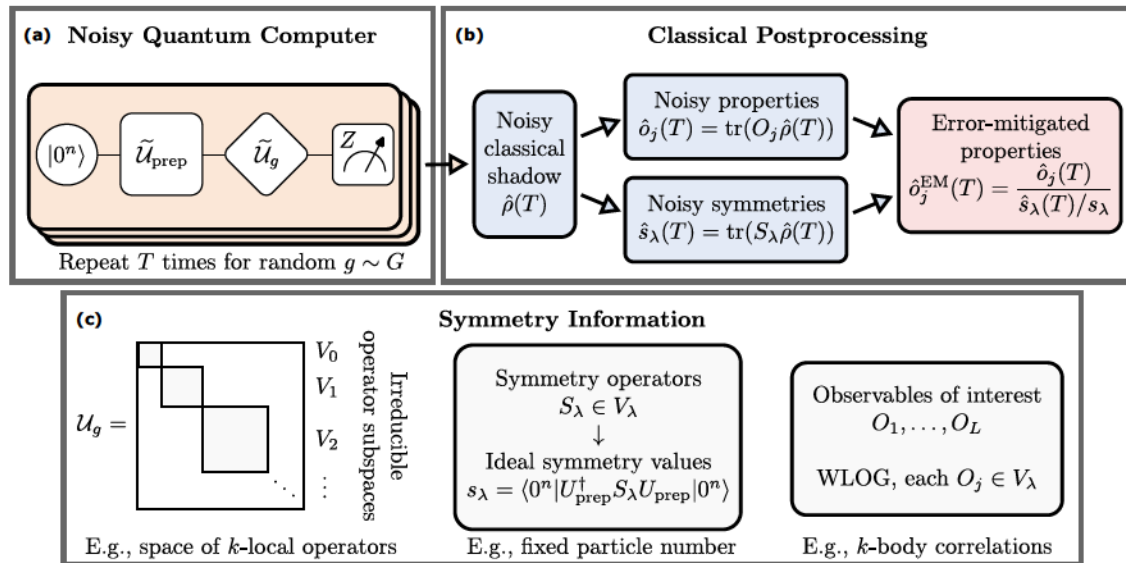


Fig. 1 | Schematic of the symmetry-adjusted classical shadows protocol. **a** Given an ideal unitary $\mathcal{U}(\cdot) = U(\cdot)U^\dagger$, its noisy implementation is denoted by $\tilde{\mathcal{U}}$. Assuming the target state $\rho = U_{\text{prep}}(|0^n\rangle\langle 0^n|)$ obeys certain symmetries S_λ , **b** we can construct error-mitigated estimates using classical shadows produced by the noisy quantum computer. (While we depict the preparation of a pure state here, our formalism is

equally valid if the target state is mixed.) In contrast to prior approaches that only address the noise in $\tilde{\mathcal{U}}$, our protocol additionally incorporates the errors within \tilde{U}_{prep} . **c** Our method is applicable whenever the symmetry is compatible with the irreps V_λ of the group G describing the classical shadows protocol.

As two key applications, we study how symmetry-adjusted classical shadows perform in simulations of fermionic and qubit systems. For fermions, we consider G corresponding to fermionic Gaussian unitaries⁶⁹ (also known as matchgate shadows⁷⁰). We establish the following performance bound for fermionic systems with particle-number symmetry, $N = \sum_{p \in [n]} a_p^\dagger a_p$.

Theorem 1. (Fermions with particle-number symmetry, informal). Let ρ be an n -mode state with $\text{tr}(N\rho) = \eta$ fermions. Under the noise model \mathcal{E} satisfying Assumptions 1 and assuming $\eta = \mathcal{O}(n)$, matchgate shadows of size

$$T = \mathcal{O}(n^2 \log(n) \epsilon^{-2} F_{Z, \{2,4\}}(\mathcal{E})^{-2}) \quad (30)$$

suffice to achieve prediction error

$$|\hat{o}_j(T) - \text{tr}(O_j \rho)| \leq \epsilon + \mathcal{O}(\epsilon^2) \quad (31)$$

with high probability, where the observables O_j can be taken as all one- and two-body Majorana operators.

The dependence on system size n and prediction error ϵ matches noiseless estimation with matchgate shadows^{69,70}. Meanwhile, the overhead of error mitigation is $\mathcal{O}(F_{Z,R'}(\mathcal{E})^{-2})$, analogous to prior related results^{85,86}. The irreps $R' = \{2, 4\}$ correspond to the Majorana degree of the k -body observables.

For qubit systems, we consider G essentially corresponding to the local Clifford group (i.e., random Pauli measurements)^{51,52}. In order to make the irreducible structure compatible with commonly encountered symmetries, we introduce a technical modification that we call subsystem-symmetrized Pauli shadows (see the subsection “Subsystem-symmetrized Pauli shadows” for a summary). The symmetry we consider here is generated by the total longitudinal magnetization, $M = \sum_{i \in [n]} Z_i$. For error-mitigated prediction of local qubit observables, we have the following result.

Theorem 2. (Qubits with total magnetization symmetry, informal). Let ρ be an n -qubit state with a fixed magnetization, $\text{tr}(M\rho) = m$. Under the noise model \mathcal{E} satisfying Assumptions 1 and assuming $m = \Theta(1)$, subsystem-

symmetrized Pauli shadows of size

$$T = \mathcal{O}(n \log(n) \epsilon^{-2} F_{Z, \{1,2\}}(\mathcal{E})^{-2}) \quad (32)$$

suffices to achieve prediction error

$$|\hat{o}_j(T) - \text{tr}(O_j \rho)| \leq \epsilon + \mathcal{O}(\epsilon^2) \quad (33)$$

with high probability, where the observables O_j can be taken as all one- and two-local Pauli operators.

Note that the irreps of subsystem-symmetrized Pauli shadows are labeled by Pauli weight. The variance bound we advertise here is linear in n , resulting from the extensive nature of the symmetry M . Specifically, we show that when $m = \Theta(1)$, $\|M\|_{\text{shadow}}^2 = \mathcal{O}(n)$ dominates the asymptotic complexity over the k -local Pauli observables (for which our protocol exhibits the usual $\|O_j\|_{\text{shadow}}^2 = 3^k$). This is consistent with standard Pauli shadows, wherein the shadow norm of arbitrary k -local observables scales at most linearly with spectral norm and exponentially in k ^{51,52}.

Besides these two examples, we describe symmetry-adjusted classical shadows for a more general class of groups G , and we establish accompanying bounds in Theorem 4 in the Methods section (proven in Supplementary Note 1). This allows for applications to other systems and unitary distributions. See “Theory of symmetry-adjusted classical shadows” in the Methods section for the general theory, and subsections “Application to fermionic (matchgate) shadows” and “Application to qubit (Pauli) shadows” for the details regarding Theorems 1 and 2, respectively.

Because our protocol always runs the full noisy quantum circuit, it has the potential to mitigate a wider range of errors than those covered by Assumptions 1, albeit without the rigorous theoretical guarantees. This is a significant feature of the method, as the preparation of ρ often dominates the total circuit complexity (i.e., U_{prep} in Fig. 1). We explore this broader mitigation potential with a series of numerical experiments below, wherein we simulate noisy Trotter circuits for systems of interacting fermions and spin-1/2 particles, respectively.

Subsystem-symmetrized Pauli shadows

While random Pauli measurements are efficient for predicting local qubit observables, the irreducible structure of the local Clifford group $\text{Cl}(1)^{\otimes n}$ is

difficult to reconcile with common symmetries under symmetry adjustment, such as the $U(1)$ symmetry generated by $M = \sum_{i \in [n]} Z_i$. To remedy this issue, we modify the protocol by what we call *subsystem symmetrization*: define the group

$$\text{Cl}(1)_{\text{Sym}}^{\otimes n} := \text{Sym}(n) \times \text{Cl}(1)^{\otimes n}, \quad (34)$$

which has the unitary representation $U_{(\pi, C)} = S_\pi C$ where S_π permutes the qubits according to $\pi \in \text{Sym}(n)$ and $C \in \text{Cl}(1)^{\otimes n}$. The circuit for S_π can be obtained as a sequence of $\mathcal{O}(n^2)$ nearest-neighbor SWAP gates in $\mathcal{O}(n)$ depth via an odd–even decomposition of π^{op} . The following theorem summarizes its group-theoretic properties relevant to classical shadows.

Theorem 3. (Irreducible representations of the subsystem-symmetrized local Clifford group). The representation $\mathcal{U} : \text{Cl}(1)_{\text{Sym}}^{\otimes n} \rightarrow \text{U}(\mathcal{L}(\mathcal{H}))$, defined by $\mathcal{U}_{(\pi, C)}(\rho) = S_\pi \rho C^\dagger S_\pi^\dagger$, decomposes into the irreps

$$V_k = \text{span}\{P \in \mathcal{P}(n) : |P| = k\}, \quad 0 \leq k \leq n. \quad (35)$$

Under this group, the (noiseless) expressions for \mathcal{M} and $\text{Var}[\hat{\rho}]$ coincide with those of standard Pauli shadows.

This modification therefore reduces the number of irreps from 2^n to $n + 1$, achieved by symmetrizing, for each k , over all k -qubit subsystems. Meanwhile, the desirable estimation properties from standard Pauli shadows are retained: for instance, the shadow norm obeys $\|P\|_{\text{shadow}}^2 = 3^k$ for k -local Pauli operators P .

The upshot is that the symmetry M is now compatible with this group, thereby enabling results such as Theorem 2. We describe this construction in ‘‘Application to qubit (Pauli) shadows’’ in the Methods section, with technical proofs in Supplementary Note 2.

Spin-adapted matchgate shadows

Systems of spinful fermions often obey a spin symmetry, which allows for compressed block-diagonal representations according to the spin sectors. Such techniques are referred to as symmetry adaptation. We introduce such an adaptation of the matchgate shadows protocol wherein the random distribution is restricted to block-diagonal orthogonal transformations,

$$Q = \begin{pmatrix} Q_\uparrow & 0 \\ 0 & Q_\downarrow \end{pmatrix} \in \text{O}(n) \oplus \text{O}(n). \quad (36)$$

We call this protocol *spin-adapted* matchgate shadows. This restricted group remains informationally complete over operators which respect the spin sectors, thus sufficing for learning properties in systems with this symmetry. In fact, we show that the shadow norms for k -fermion operators under the spin-adapted protocol scale identically as in the unadapted setting. The main advantage of spin adaptation is that the block-diagonal transformation $Q = Q_\uparrow \oplus Q_\downarrow$ can be implemented as $U_{Q_\uparrow} \otimes P_\downarrow^\dagger U_{Q_\downarrow}$, where $P_\downarrow = Z^{\otimes n/2}$ is the parity operator on the spin-down sector and $s = \delta_{-1, \det Q_\downarrow}$. This tensor-product unitary requires roughly half the number of gates and circuit depth compared to implementing a dense element of $\text{O}(2n)$. We prove the necessary details in Supplementary Note 3 and implement this modified protocol in our numerical experiments wherever applicable.

Improved circuit design for fermionic Gaussian unitaries

Fermionic Gaussian unitaries are a broad class of free-fermion rotations, and they are ubiquitous primitives in algorithms for simulating (interacting) fermions. In the context of classical shadows, they form the basis for randomized measurements in matchgate shadows^{69–71}. Such unitaries can be described by an orthogonal transformation $Q \in \text{O}(2n)$ of the Majorana

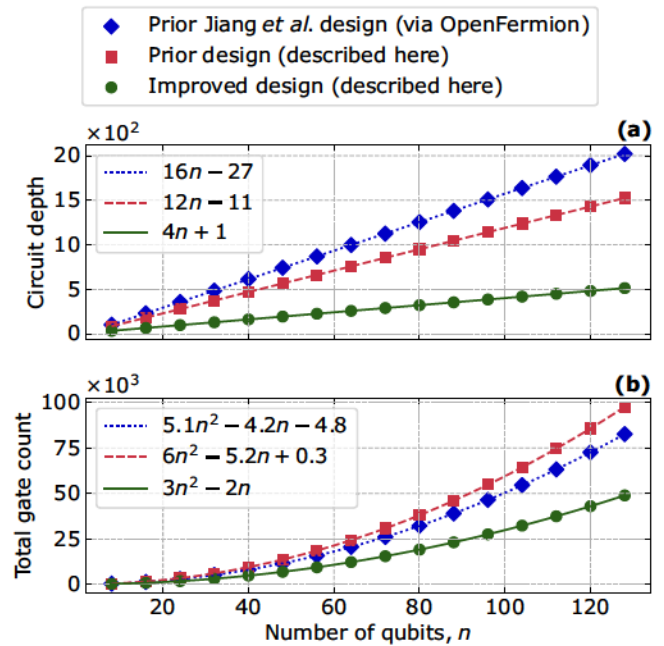


Fig. 2 | Resource comparison of our improved fermionic Gaussian circuit design versus prior designs. Our circuit costs (green) are compared against those of Jiang et al.¹⁰⁰ (blue), implemented in OpenFermion¹³², and a naive scheme described in Supplementary Note 4 (red). All circuits were optimized to a superconducting gate set consisting of \sqrt{i} SWAP gates and single-qubit rotations. We compare both a circuit depth and b gate count on random inputs $Q \in \text{O}(2n)$. Linear and quadratic fits are made, demonstrating a roughly 3 \times and 2 \times savings, respectively.

operators,

$$U_Q(\gamma_\mu) = U_Q \gamma_\mu U_Q^\dagger = \sum_{\nu \in [2n]} Q_{\nu\mu} \gamma_\nu \quad (37)$$

for each $\mu \in [2n]$. The quantum circuits implementing these transformations take $\mathcal{O}(n^2)$ gates in $\mathcal{O}(n)$ depth^{100,101}. While this scaling is necessary in general by parameter counting, constant-factor savings can substantially improve performance in practice, especially on noisy quantum computers.

To this end, we introduce a more efficient compilation scheme for fermionic Gaussian unitaries, given an arbitrary $Q \in \text{O}(2n)$. Our circuit design improves the parallelization of gates compared to prior art^{100,101}. The key idea is to observe that two Majorana modes essentially correspond to one qubit under the Jordan–Wigner transformation¹⁰². Thus, the optimal approach to compiling U_Q into single- and two-qubit gates involves decomposing the matrix Q into elementary blocks of 4×4 transformations, rather than the 2×2 Givens rotations utilized in prior designs.

The details of this scheme are described in Supplementary Note 4 and implemented in code at our open-source repository (<https://github.com/zhao-andrew/symmetry-adjusted-classical-shadows>). We make use of this improved design in our numerical simulations. We demonstrate the improvements in circuit size in Fig. 2, with respect to a gate set native to superconducting platforms. From these results we numerically infer roughly 1/3 reduction in depth and 1/2 reduction in gate count over prior designs.

Numerical experiments

We now demonstrate the error-mitigation capabilities of symmetry-adjusted classical shadows through numerical simulations. We focus on the task of estimating one- and two-body observables in both fermion and qubit systems which obey the global $U(1)$ symmetries described in the Methods section.

For each type of system, we first present results when the noise models obey Assumptions 1 (readout errors). We demonstrate the successful

mitigation at varying sample sizes, noise rates, and system sizes, confirming the correctness of our theory.

Next, we investigate how symmetry adjustment performs under a more comprehensive noise model based on superconducting-qubit platforms. These simulations were performed using the Quantum Virtual Machine (QVM) within the Cirq open-source software package^{93,103}. It uses existing hardware data on a native gate set (single-qubit rotations and two-qubit \sqrt{i} SWAP gates on a square lattice) to mimic the realistic performance of a noisy quantum computer. We use the calibration data provided of Google’s 23-qubit Rainbow processor based on the Sycamore architecture, which was used in quantum experiments simulating quantum chemistry and strongly correlated materials^{33,35}. The noise model consists of depolarizing channels, two-qubit coherent errors, single-qubit idling noise, and readout errors. Error rates vary across the chip; on the 2×4 grid that we simulated, the average single- and two-qubit Pauli error rates are $\sim 0.15\%$ and $\sim 1.5\%$, respectively. A precise description of the noise model can be found in Supplementary Note 6.

Throughout, we use the following conventions for figures. Noiseless data (blue squares) correspond to simulations of an ideal quantum computer, which experiences no noise channel and only exhibits the fundamental sampling error. Unmitigated data (black X’s) are simulations of classical shadows on a noisy quantum computer, using standard post-processing routines. The mitigated estimates (red diamonds) are instead postprocessed as symmetry-adjusted classical shadows, as described in the Methods section. In some experiments, we also compare against robust shadow estimation⁸⁵ (RShadow, green crosses), which involves simulating the calibration protocol on $|0^n\rangle$ under the same noise model. Finally, the true values (teal curves) are the ground truth, against which we determine the prediction error.

Uncertainty bars represent one standard deviation of the combined sampling and postprocessing, computed by empirical bootstrapping¹⁰⁴. To ease the computational load, we slightly modify the procedure by batching samples; see Supplementary Note 6 for details.

Fermionic systems. Our first set of numerical experiments consider the application to matchgate shadows to learn and mitigate noise in one- and two-body fermionic observables. The symmetry we consider is fixed particle number, $\text{tr}(N\rho) = \eta$. As we show in the Methods section, this symmetry projects into the relevant irreps $R' = \{2, 4\}$ of the matchgate shadows as

$$S_2 = \Pi_2(N) = -\frac{1}{2} \sum_{p \in [n]} Z_p, \tag{38}$$

$$S_4 = \Pi_4(N^2) = \frac{1}{2} \sum_{p < q} Z_p Z_q, \tag{39}$$

represented under the Jordan–Wigner transformation¹⁰² for simplicity. Their ideal values are

$$s_2 = \text{tr}(S_2\rho) = \eta - \frac{n}{2}, \tag{40}$$

$$s_4 = \text{tr}(S_4\rho) = \frac{1}{2} \binom{n}{2} - \eta(n - \eta). \tag{41}$$

Readout noise model (fermions). First, we consider the reconstruction of the fermionic two-body reduced density matrix (2-RDM) from matchgate shadows. The 2-RDM elements of a state ρ are given by

$${}^2D_{rs}^{pq} = \text{tr}(a_p^\dagger a_q^\dagger a_s a_r \rho), \quad p, q, r, s \in [n]. \tag{42}$$

In general, knowledge of the k -RDM allows one to calculate any k -body observable of the system. By anticommutation relations, there are only $\binom{n}{2}^2$ unique matrix elements, corresponding to the indices $p < q$ and $r < s$.

We therefore represent 2D as an $\binom{n}{2} \times \binom{n}{2}$ Hermitian matrix, flattening along those index pairs. Estimates ${}^2\hat{D}_{rs}^{pq}(T) = \text{tr}(a_p^\dagger a_q^\dagger a_s a_r \hat{\rho}(T))$ are computed from T matchgate-shadow samples. Here, our figure of merit for the prediction error is the spectral-norm difference between the reconstructed and the numerically exact 2-RDMs, $\epsilon = \|{}^2\hat{D} - {}^2D\|_\infty$.

We demonstrate 2-RDM reconstruction on an ensemble of 20 random Slater determinants (noninteracting-fermion states with fixed particle number). An η -fermion Slater determinant is specified by the first η columns of an $n \times n$ unitary matrix, so we generate the random states by uniformly drawing elements of $U(n)$. This $n \times n$ representation is then lifted to the $2n \times 2n$ fermionic Gaussian representation, which allows us to apply the random matchgate transformations $Q \in B(2n)$ efficiently. This simulates the action of $\rho \mapsto U_Q \rho U_Q^\dagger$. The measurement is then simulated using the algorithm of [ref. 105, Sec. 5.1]. Finally, to simulate the readout noise we implement the effective noise channel on the sampled bit strings offline.

While the 2-RDM of free-fermion states can be computed from the 1-RDM using Wick’s theorem, we do not employ any such tricks here; we use Slater determinants simply to facilitate fast classical simulation. We also do not use any additional error-mitigation strategies, such as RDM positivity constraints¹⁰⁶, that could in principle be applied in tandem.

The results are presented in Fig. 3. We consider a small system size, $n = 8$ and $\eta = 2$, and simulate three types of single-qubit noise channels before readout: depolarizing, amplitude damping, and bit flip. The noise rate p represents the probability of such an error occurring, independently on each qubit (defined in Supplementary Note 6). In the top row, we show how the prediction error varies with the total number of samples T . As expected, the noiseless estimates (corresponding to $p = 0$) converge as $\sim T^{-1/2}$, which is the standard shot-limited behavior. Then, setting $p = 0.2$, we see how the unmitigated data experiences an error floor beyond which taking additional samples does not improve the accuracy. On the other hand, the mitigated results clearly bypass this error floor and recover the shot-noise scaling with T , thus validating the theory of symmetry-adjusted classical shadows. Compared to the noiseless simulations, our mitigated data exhibit a constant factor increase in the sampling cost, corresponding to the $\mathcal{O}(F_{Z,R}^2)$ overhead of error mitigation, as it appears in Theorem 4.

For these experiments, we also compare to the performance of robust shadow estimation (RShadow) by Chen et al.⁸⁵, which requires simulating the calibration procedure on $|0^n\rangle$. For a fair comparison, we allocate $T/2$ samples to the calibration step and $T/2$ samples to the estimation step, so that the total number of samples is the same. While Chen et al.⁸⁵ did not originally consider matchgate shadows, from our generalization in Eq. (23) we can construct $\text{NoiseEst}_{B(2n)}$ by taking $D_\lambda = S_{2k}$, which obeys $\langle 0^n | S_{2k} | 0^n \rangle = -n/2$ and $\langle 0^n | S_{4k} | 0^n \rangle = n(n-1)/4$. The single-shot estimator is then

$$\text{NoiseEst}_{B(2n)}(2k, Q, b) = \frac{\langle b | U_Q S_{2k} U_Q^\dagger | b \rangle}{\langle 0^n | S_{2k} | 0^n \rangle}. \tag{43}$$

As expected, RShadow behaves similarly to symmetry-adjusted classical shadows in this scenario (wherein the noise obeys Assumptions 1). However, even here we observe that our approach exhibits a constant-factor advantage in the sampling cost. We attribute the performance of RShadow to its calibration procedure, which our method avoids.

In the bottom row of Fig. 3, we simulate the same collection of random Slater determinants, but now varying the noise rate p at a fixed shadow size $T = 10^6$. While the unmitigated errors quickly grow with increasing noise rate as expected, the mitigated estimates remain under control. Note that the

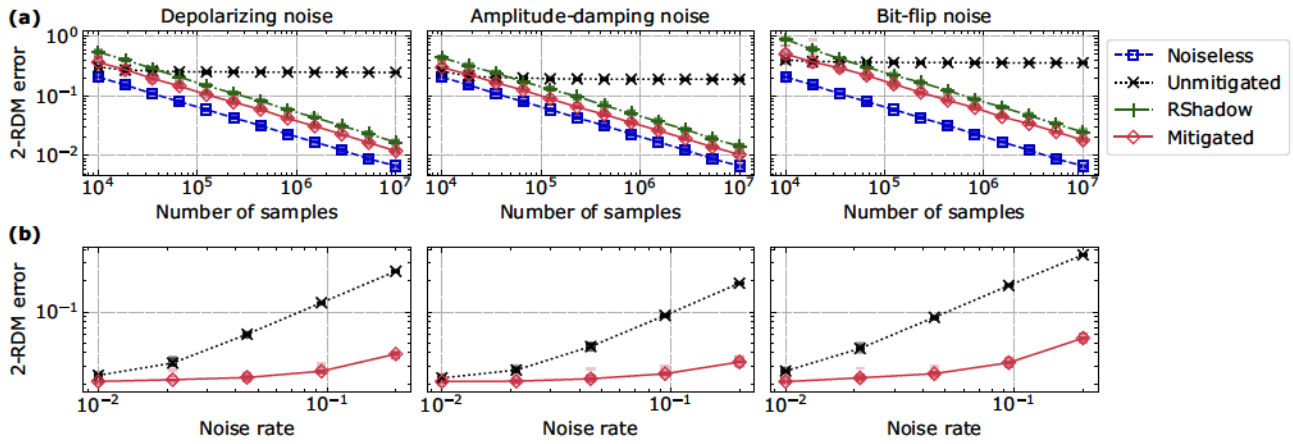


Fig. 3 | Estimation error of the fermionic 2-RDM reconstructed from matchgate shadows. Error is quantified by the spectral-norm difference between the estimated and true 2-RDM. We simulate the protocol on a collection of 20 random Slater determinants on $n = 8$ modes with $\eta = 2$ fermions; faint dashes are results for

individual states, markers indicate the median error across the 20 states. **a** Scaling of estimation error with the total number of samples T , fixing the noise rate to $p = 0.2$ for all noise models. **b** Scaling of the estimation error with the noise rate p , fixing the total number of samples to $T = 10^6$.

mitigated errors still grow modestly because we have fixed the number of samples; in order to achieve a constant prediction error, one would need to grow T proportional to $F_{Z,R}^{-2}$ (which is p -dependent). Our key takeaway is that the combination of both rows of plots indicates the ability to handle a range of common noise channels at fairly high error rates. Indeed, the growing errors seen in the bottom row can be suppressed by simply taking more samples, which is precisely what the top row demonstrates.

Next, we consider the simulation of a 1D spinful Fermi–Hubbard chain of $L = n/2$ sites (for a total of n fermionic modes/qubits). Under open boundary conditions, the Hamiltonian for this model is

$$H = J + V, \tag{44}$$

where

$$J = -t \sum_{i \in [L-1]} \sum_{\sigma \in \{\uparrow, \downarrow\}} a_{i,\sigma}^\dagger a_{i+1,\sigma} + \text{h.c.}, \tag{45}$$

$$V = U \sum_{i \in [L]} N_{i,\uparrow} N_{i,\downarrow}, \tag{46}$$

are the hopping and interaction terms, respectively. The creation operators $a_{i,\sigma}^\dagger$ produce an electron at site i with spin σ , and $N_{i,\sigma} = a_{i,\sigma}^\dagger a_{i,\sigma}$ is the associated occupation-number operator. We set units such that the hopping strength is $t = 1$.

For the target state, we use the ground state of the noninteracting term J , which is also a Slater determinant. This allows us to use the same simulation techniques as before to efficiently simulate up to 20 sites. The number of electrons in each spin sector is $\eta_\sigma = L/2$, for a total of $\eta = \eta_\uparrow + \eta_\downarrow = n/2$ electrons. Thus the system is at half filling, which requires the use of ancilla qubits to avoid division by zero (see the Methods section). In fact, we simulate $n' = n + 2$ qubits because we append an ancilla qubit to each spin sector. This is because we furthermore employ spin-adapted matchgate shadows, as described previously in the Results section. This modification essentially treats each spin sector independently when performing the randomized measurements, so each sector is at half filling.

The Fermi–Hubbard results are shown in Fig. 4. We consider the estimation of energy per electron, $\langle H \rangle / \eta$. We set the interaction strength to $U/t = 4$ and the noise model to single-qubit bit-flip errors, with probabilities $p \in \{0.01, 0.03, 0.05\}$. The energy per electron (top) and absolute estimation error (bottom) are plotted as the system size grows, keeping the number of samples fixed to $T = 2 \times 10^6$. Again, these results serve to validate our theory, this time highlighting the performance as the system size grows. This also

demonstrates the use of spin-adapted matchgate shadows and the successful use of ancillas to avoid division by zero in \hat{o}_j^{EM} .

QVM noise model (fermions). Now we turn to the gate-level noise model simulated through the QVM^{93,103}. This model strongly violates Assumptions 1, reflecting the fact that the state-preparation circuit U_{prep} is typically the dominant source of errors.

As our testbed fermionic system, we again consider the 1D spinful Fermi–Hubbard chain with open boundary conditions and interaction strength $U/t = 4$. Rather than the static problem, here we simulate Trotterized time evolution of the Hamiltonian. The number of Trotter steps provides a systematic way to increase the circuit depth (and hence the cumulative amount of noise) within the same model. Note that because we are studying the behavior of error mitigation, the ground truth of these simulations corresponds to the noiseless Trotter circuit with a finite step size (i.e., we are not comparing to the exact, non-Trotterized dynamics).

We closely follow the setup of the experiment performed in ref.³⁵ (which was in fact performed on the Sycamore processor that our noise model is based on), using code made available by the authors at ref.¹⁰⁷. Because simulating the full noisy circuit is exponentially expensive, we restrict to a four-site instance ($n = 2L = 8$). The initial state is the ground state in the $\eta_\uparrow, \eta_\downarrow = 1$ sector of the noninteracting Hamiltonian

$$H_0 = J + \sum_{i \in [L]} \sum_{\sigma \in \{\uparrow, \downarrow\}} \varepsilon_{i,\sigma} N_{i,\sigma}, \tag{47}$$

where J is the hopping term defined in Eq. (45) and we set the on-site potentials to have a Gaussian form, $\varepsilon_{i,\sigma} = -\lambda_\sigma e^{-\frac{1}{2}(i+1-c)^2/s^2}$. This generates a Slater determinant whose charge density

$$q_i = \langle N_{i,\uparrow} + N_{i,\downarrow} \rangle. \tag{48}$$

has a Gaussian profile, centered around c with width s and magnitude λ_σ . We set the parameters to $c = L/2 + 1/2 = 2.5$, $s = 7/3$, and $\lambda_\sigma = 4\delta_{\sigma,\uparrow}$. This initial state is prepared by the appropriate single-particle basis rotations^{100,108,109} on the state $|1000\rangle$ within each spin sector. Denote this unitary by $U(H_0)$. The system is then evolved by Trotterized dynamics according to H , with $R \in \{0, 1, \dots, 5\}$ steps of size $\delta t = 0.2$. Let J_{even} (resp., J_{odd}) be the terms in J with i even (resp., odd), and similarly for V_{even} , V_{odd} . One Trotter step is ordered as

$$U_{\text{Trot}} = e^{-iJ_{\text{odd}}\delta t} e^{-iV_{\text{odd}}\delta t} e^{-iV_{\text{even}}\delta t} e^{-iJ_{\text{even}}\delta t}, \tag{49}$$

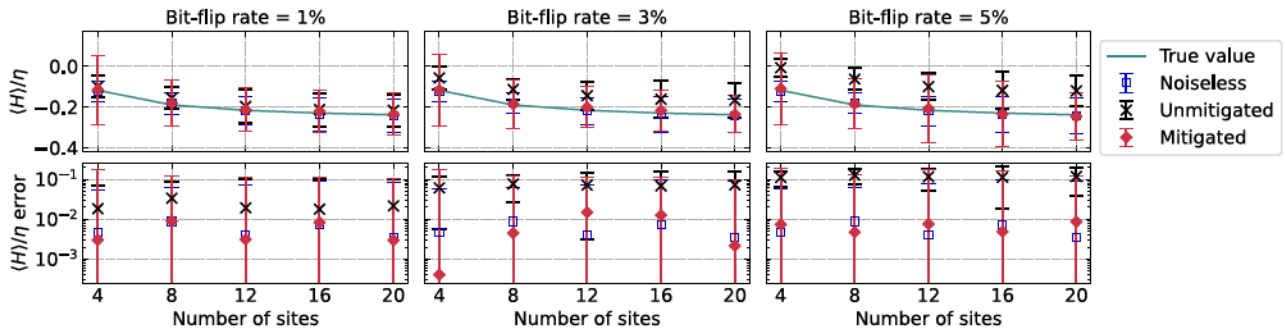


Fig. 4 | Estimation of the energy per electron of a 1D spinful Fermi-Hubbard model. We consider an interaction strength $U/t = 4$ on $4 \leq L \leq 20$ sites, and we measure $\langle H \rangle / \eta$ using spin-adapted matchgate shadows with $T = 2 \times 10^6$ samples. We simulate the protocol on the ground state of the non-interacting component J of the Hamiltonian, Eq. (45), at half filling in each spin

sector, which is a Slater determinant with $\eta_\sigma = L/2$ electrons per sector. These simulations involve $n' = 2L + 2$ noisy qubits experiencing readout bit-flip errors with probabilities 1%, 3%, and 5%. The additional qubit per spin sector is prepared in $|0\rangle$ to avoid division by zero (see the Methods section). Error bars denote one standard error of the mean.

which is then compiled into the native gate set. The full state-preparation circuit is then

$$U_{\text{prep}}(R) = U_{\text{Trot}}^R U(H_0) X_{0,\downarrow} X_{0,\uparrow}, \quad (50)$$

where $X_{0,\sigma}$ places a spin- σ electron on the first site from the vacuum (i.e., prepares $|1000\rangle$) in each spin sector). Note that $R = 0$ corresponds to only preparing the initial Slater determinant, which still has nontrivial circuit depth. Further details on the construction of these circuits can be found in refs. 35,107.

One final detail of ref. 35 that we follow is their method of qubit assignment averaging (QAA). This technique is employed as a means of ameliorating inhomogeneities in error rates across the quantum device. QAA works by identifying a collection of different assignments for the physical qubit labels and uniformly averaging over them (keeping the Jordan-Wigner convention fixed). For example, one may vary qubit assignments by selecting a different portion of the chip, or rotating/flipping the layout. Here, we fix a 2×4 grid of qubits and perform QAA over four different orderings of those eight qubits; see Supplementary Note 6 for the specific assignments chosen.

For each target state $U_{\text{prep}}(R)|0^n\rangle$, we collect $T = 9.6 \times 10^5$ spin-adapted matchgate shadow samples. In Fig. 5, we plot the Trotterized time evolution of charge density throughout the chain, as well as the charge spread

$$\kappa = \sum_{i \in [L]} |i - (L - 1)/2| \rho_i, \quad (51)$$

which quantifies how the density spreads away from the center of the chain. These quantities are only one-body observables, so as an exemplary two-body observable we also estimate the energy per electron, $\langle H \rangle / \eta$.

Because Assumptions 1 no longer hold, we no longer have the guarantees of Theorem 4 and we do not observe an arbitrary amount of error mitigation. We see that as the circuit size grows, so too do the prediction error and uncertainty. This behavior is a reflection of the noise assumptions being increasingly violated. Nonetheless, our results still show a substantial amount of noise reduction, and overall we maintain the qualitative features of the dynamics compared to the unmitigated protocol. In Supplementary Note 6, we provide a quantitative estimate of how much the QVM noise model violates Assumptions 1. There we observe a fundamental error floor, which is roughly an order of magnitude below the mitigated errors actually achieved here, indicating the potential for further mitigation beyond what we have presently demonstrated.

Qubit systems. Next, we study the application of symmetry-adjusted classical shadows to subsystem-symmetrized Pauli shadows, to predict one- and two-body qubit observables in the presence of noise. We consider a fixed magnetization symmetry $\text{tr}(M\rho) = m$, which, as we show in

the Methods section, projects into the relevant irreps $R' = \{1, 2\}$ as

$$S_1 = \Pi_1(M) = \sum_{i \in [n]} Z_i, \quad (52)$$

$$S_2 = \Pi_2(M^2) = 2 \sum_{i < j} Z_i Z_j. \quad (53)$$

The ideal symmetry values in this case are

$$s_1 = m, \quad (54)$$

$$s_2 = m^2 - n. \quad (55)$$

Readout noise model (qubits). For our first demonstration, we simulate random matrix product states (MPS) with maximum bond dimension $\chi \leq n$, lying in the $m = 0$ symmetry sector of M . We use the definition of a random MPS from refs. 110,111. Numerically, we implement all MPS calculations using the open-source software ITensor¹¹², which can guarantee the correct symmetry sector using efficient tensor-network representations. Within such representations, it is straightforward to apply random local Clifford gates and SWAP gates, and to sample measurements in the computational basis.

Unlike fermions, qubits are not symmetrized, so their 2-RDMs

$${}^2D_{ij} = \text{tr}_{[n] \setminus \{i,j\}} \rho \quad (56)$$

are in general distinct between different two-qubit subsystems. Our accuracy metric here is therefore the mean 2-RDM error over all pairs of qubits:

$$e = \frac{1}{\binom{n}{2}} \sum_{i < j} \|\hat{D}_{ij}(T) - {}^2D_{ij}\|_\infty. \quad (57)$$

From subsystem-symmetrized Pauli shadows $\hat{\rho}(T)$ of size T , we reconstruct the qubit 2-RDMs by estimating all one- and two-local Pauli expectation values and forming the 4×4 matrices

$${}^2\hat{D}_{ij}(T) = \frac{1}{4} \sum_{W, W' \in \mathcal{P}(1)} \text{tr}(W_i W'_j \hat{\rho}(T)) W \otimes W'. \quad (58)$$

The results are shown in Fig. 6. Similar to the conclusions drawn from Fig. 3 for the fermionic case, we observe that our theory is validated in two important parameters (number of samples and error rate). We note here that this simple demonstration also validates our subsystem-symmetrized

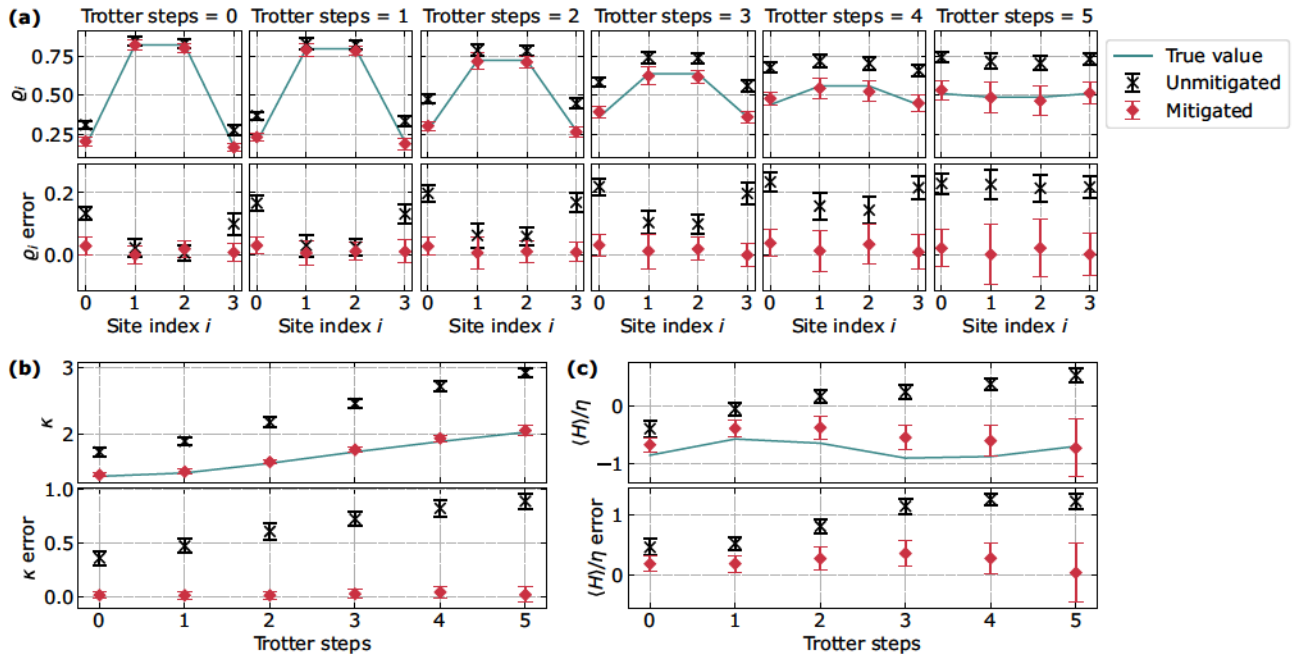


Fig. 5 | Prediction of local properties in the four-site 1D Fermi-Hubbard model undergoing Trotterized time evolution. The interaction strength is set to $U/t = 4$ and the Trotter step size is $\delta t = 0.2$. The noise model is that of the Google Sycamore Rainbow processor^{33,35}, implemented within the Cirq QVM^{93,103}. The size of the noisy

circuit grows systematically with the number of Trotter steps. We use spin-adapted matchgate shadows, taking $T = 9.6 \times 10^5$ samples. **a** Charge density Q_i at each site $i \in [L]$. **b** Charge spread κ . **c** Energy per electron $\langle H \rangle / \eta$. Error bars denote one standard error of the mean.




Fig. 6 | Average estimation error of the 2-RDM over all two-qubit subsystems, reconstructed from (subsystem-symmetrized) Pauli shadows. Error is quantified by the spectral-norm difference between the estimated and true 2-RDM for each subsystem, taking the mean over all $\binom{n}{2}$ subsystems. We simulate the protocol on a collection of 20 random eight-qubit matrix product states, with maximum bond

dimension $\chi \leq n$ and fixed Z magnetization, $m = 0$. Because $m = 0$ implies $s_1 = 0$, we simulate a system of nine qubits with the ancilla prepared in $|0\rangle$. **a** Scaling of estimation error with the total number of samples T , fixing the noise rate to $p = 0.2$ for all noise models. **b** Scaling of the estimation error with the noise rate p , fixing the total number of samples to $T = 10^6$.

Pauli shadows protocol and the use of ancillas in this scenario as well (recall that the random states we study here have vanishing symmetry value, $m = s_1 = 0$).

Our next set of numerical experiments are performed on the ground state of an antiferromagnetic XXZ Heisenberg chain with open boundary conditions:

$$H = J \sum_{i \in [n-1]} (X_i X_{i+1} + Y_i Y_{i+1} + \Delta Z_i Z_{i+1}). \quad (59)$$

Throughout, we set units such that $J = 1$ and consider an anisotropy of $\Delta = 1.5$. This Hamiltonian commutes with the symmetry operator M , and in

particular the ground state obeys $m = 0$ (assuming the number of spins n is even). We find the ground state via the density-matrix renormalization group (DMRG) algorithm¹¹³, represented as an MPS; therefore we can employ the same classical simulation algorithms as before. Although $m = 0$ implies a vanishing conserved quantity for the one-body subspace, $s_1 = 0$, we do not employ the ancilla technique for these simulations because we will only be interested in predicting strictly two-body observables (for which $s_2 = m^2 - n \neq 0$).

In Fig. 7 we show the mitigation of energy per spin $\langle H \rangle / n$ at different system sizes and bit-flip rates on each qubit. For these experiments, the number of samples taken is $T = 10^6$. Again, the results validate our theory for Pauli-shadow symmetry adjustment over a range of noise rates and system

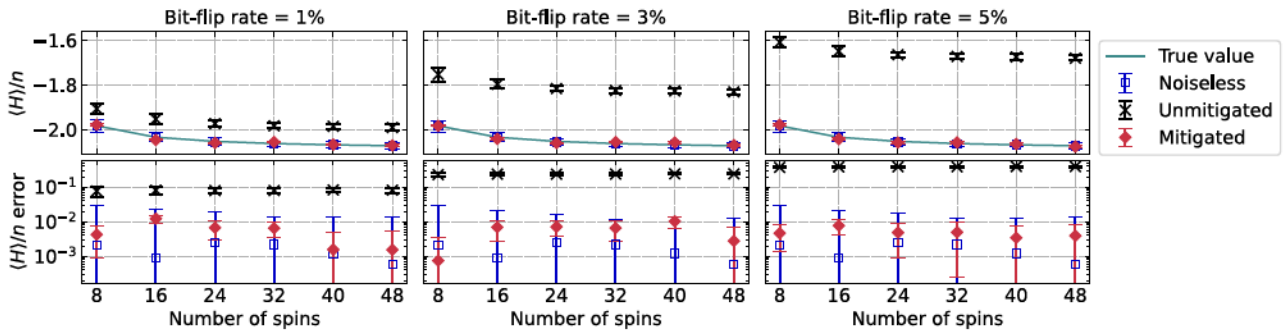


Fig. 7 | Estimation of energy per particle in the 1D XXZ ground state. We measure $\langle H \rangle/n$ from $T = 10^6$ subsystem-symmetrized Pauli shadows. The anisotropy of the XXZ model is $\Delta = 1.5$. Note that, as an antiferromagnetic model, the ground state

exhibits an $m = 0$ net magnetization symmetry. The noise model is single-qubit bit flip with probabilities $p = 1\%$, 3% , and 5% . Error bars denote one standard error of the mean.

sizes. In particular, although we require estimating the symmetry operator M which has variance $\mathcal{O}(n)$ (as opposed to H/n , which has constant variance), we see that in practice it suffices to take a number of samples constant in system size. This may indicate that our analysis of the worst-case sampling bounds for symmetry-adjusted classical shadows may be overly pessimistic in typical settings.

QVM noise model (qubits). We now turn to simulations using the QVM noise model, taking the same XXZ Heisenberg spin chain ($\Delta = 1.5$ and $n = 8$) as our testbed system. Similar to our numerical experiments with the Fermi–Hubbard model, we simulate Trotter circuits of the XXZ model starting from a product state within the symmetry sector of $m = 0$. Again, we will only be interested in strictly two-local observables so we do not employ the ancilla trick here.

Our initial state is a Néel-ordered product state, $|01010101\rangle = \prod_{j \text{ odd}} X_j |0^n\rangle$. Defining H_{even} and H_{odd} as the terms in H with i even and odd, respectively, a single Trotter step is given by

$$U_{\text{Trot}} = e^{-iH_{\text{odd}}\delta t} e^{-iH_{\text{even}}\delta t}, \tag{60}$$

where we take the step size to be $\delta t = 0.2$. Hence, the full state-preparation circuit for R steps is

$$U_{\text{prep}}(R) = U_{\text{Trot}}^R \prod_{j \text{ odd}} X_j, \tag{61}$$

which is then compiled into the native gate set. For each R , we collect $T = 4.8 \times 10^5$ samples using subsystem-symmetrized Pauli shadows. Because the initial state is a simple basis state, we only display results for $R \in \{1, \dots, 5\}$ for these studies. In line with our Fermi–Hubbard simulations on the QVM, we perform QAA here as well, averaging over twelve different assignments of the same 2×4 qubits (see Supplementary Note 6 for details).

First, we consider the spin–spin correlations $\langle S_i \cdot S_j \rangle$, where

$$S_i = \frac{1}{2} \begin{pmatrix} X_i \\ Y_i \\ Z_i \end{pmatrix}, \tag{62}$$

for all qubit pairs (i, j) throughout the chain. We plot the prediction errors of these correlation functions in Fig. 8, with the unmitigated data in the first row and mitigated data in the second row. We observe that, while the shallower Trotter circuits are well handled by symmetry-adjusted classical shadows, the mitigation power diminishes as the circuit grows deeper. To examine this effect closer, we plot in the bottom two rows of Fig. 8 the correlation functions between the first spin and the rest of the chain. We see that the $\langle S_0 \cdot S_1 \rangle$ errors are particularly dominant due to the magnitude of its true value. Although the absolute error is only marginally improved by

symmetry adjustment for some of these pairs, the qualitative behavior is more faithfully recovered than in the unmitigated data (wherein the increasing circuit noise washes out the antiferromagnetic correlations).

Next, we consider macroscopic observables in Fig. 9: the Néel order parameter

$$\langle S_{\text{AF}}^2 \rangle = \frac{1}{n^2} \sum_{i,j \in [n]} (-1)^{i+j} \langle S_i \cdot S_j \rangle, \tag{63}$$

and the energy per spin $\langle H \rangle/n$. Again we see general trends similar to the other QVM simulations: the mitigated results are in closer qualitative agreement with the true values than the unmitigated data, at the cost of larger uncertainty bars, and without arbitrary amounts of error suppression. Symmetry adjustment consistently reduces the absolute error compared to the unmitigated data, although we note that some of the energy estimates are still a few standard deviations away from the true value. This is attributed to the violation of Assumptions 1, and we leave it an open problem of how to further ameliorate this property.

Discussion

In this paper, we have introduced symmetry-adjusted classical shadows, a QEM protocol applicable to quantum systems with known symmetries. Our approach builds on the highly successful classical-shadow tomography^{51,52}, modifying the classically computed linear-inversion step according to symmetry information in the presence of noise. Because our strategy is performed in postprocessing on the noisy measurement data, it allows for straightforward combinations with other QEM strategies. As opposed to prior related works^{85–89}, the main advantage of our approach is the use of the entire noisy circuit, thereby bypassing the need for calibration experiments and accounting for errors in state preparation. Meanwhile, in contrast with other symmetry-based strategies^{83,90–92}, we require no additional quantum resources, utilize finer-grained symmetry information, and can easily take advantage of a wider range of symmetries (e.g., particle number as opposed to only parity conservation). We note that, while this work has focused on local observables (linear functions of ρ), classical shadows can seamlessly be used for nonlinear observable estimation as well⁵¹. Because symmetry adjustment works at the level of the shadow channel inversion, our QEM strategy applies within that context just as well.

Overall, our findings reveal that as a low-cost scheme, symmetry-adjusted classical shadows by itself is already potent for practical error mitigation. Our analytical results guarantee the accuracy of prediction under readout noise assumptions. Even when these assumptions are violated in practice, we expect these results to still provide intuition regarding the mitigation behavior. Indeed, this expectation is validated by our numerical experiments with superconducting-qubit noise models on the Cirq QVM^{93,103}. From these simulations, we have observed substantial quantitative improvement when the cumulative circuit noise is sufficiently weak, and qualitative improvements across all experiments performed.

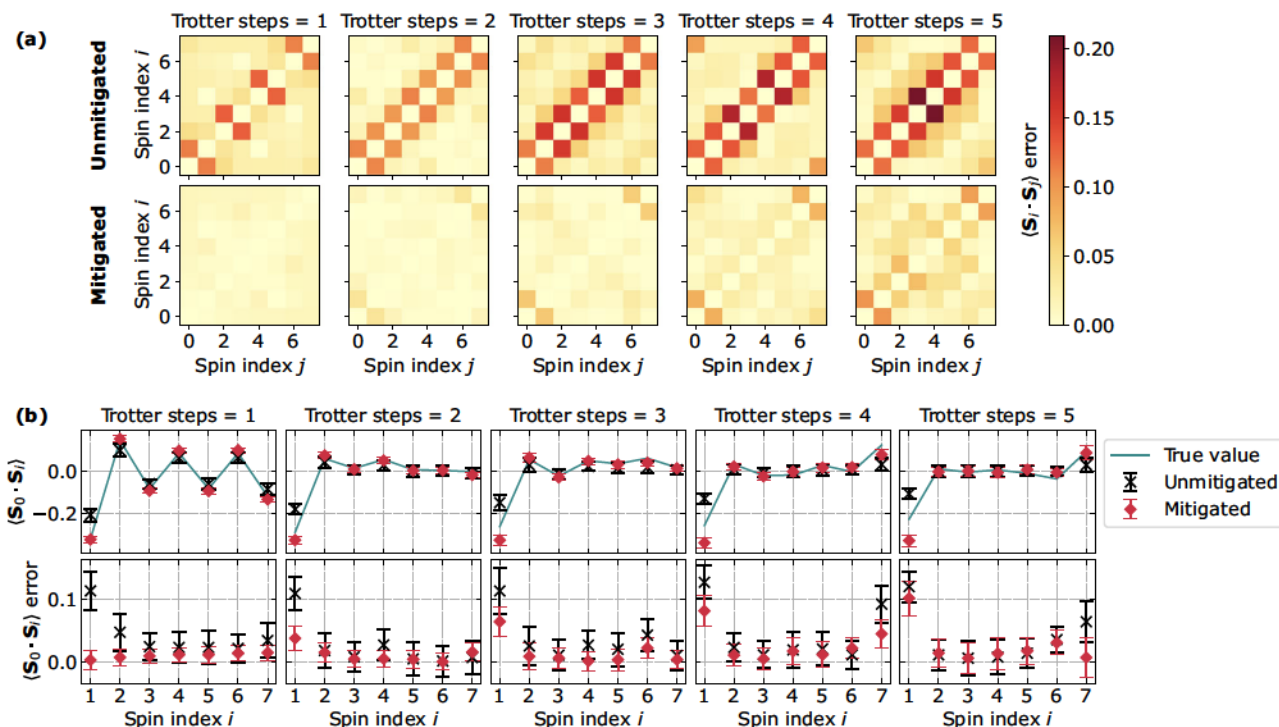


Fig. 8 | Prediction of spin–spin correlations between qubits in the XXZ chain undergoing Trotterized time evolution. The anisotropy in the model is $\Delta = 1.5$, and we take a Trotter step size of $\delta t = 0.2$. The noise model is the superconducting-hardware model implemented within the QVM. We take $T = 4.8 \times 10^5$ subsystem-symmetrized Pauli shadows to estimate the observables. **a** Heatmaps of estimation errors for $\langle S_i \cdot S_j \rangle$, under unmitigated versus our symmetry-enabled mitigated postprocessing. **b** Plots of the $\langle S_0 \cdot S_i \rangle$ correlation functions, to show further details of that particular row of the heatmaps. Error bars denote one standard error of the mean.

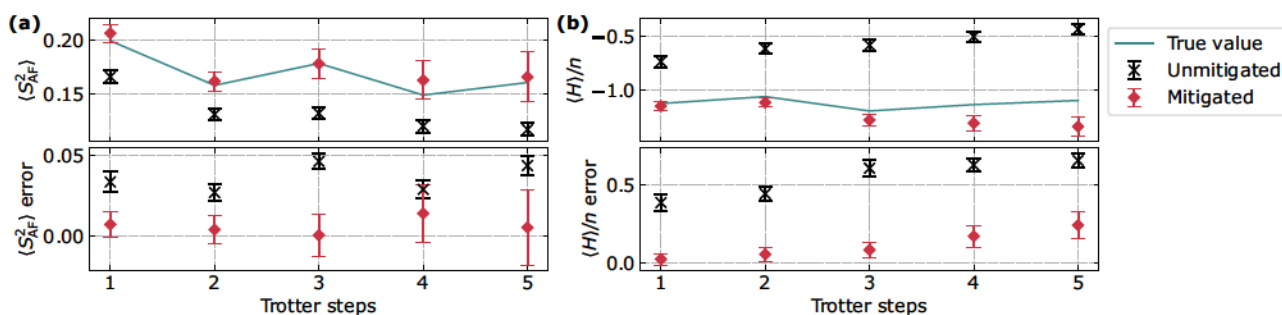


Fig. 9 | Prediction of macroscopic observables in the Trotterized XXZ model under the QVM noise model. Observables are estimated with $T = 4.8 \times 10^5$ subsystem-symmetrized Pauli shadows. **a** Néel order parameter $\langle S_{AF}^2 \rangle$. **b** Energy per spin $\langle H \rangle/n$. Error bars denote one standard error of the mean.

Along the way, we have developed a number of ancillary results that may also be of independent interest. Of note are (1) the subsystem-symmetrized Pauli shadows, which uniformly symmetrizes the irreps of the local Clifford group among subsystems; (2) an improved circuit compilation scheme for fermionic Gaussian unitaries, which treats Majorana modes on a more natural footing to improve two-qubit gate parallelization; and (3) symmetry-adapted matchgate shadows, which uses block-diagonal transformations within spin sectors to reduce the size of the random matchgate circuits. We expect that these techniques will find broader applicability in quantum simulation beyond the scope of this paper.

A number of pertinent open questions and future directions remain. For simplicity of the protocol, and because of the examples that we focused on, we restricted attention to multiplicity-free groups. However, tools to generalize to nonmultiplicity-free groups already exist, and in the context of character randomized benchmarking¹¹⁴ such an extension has been developed successfully¹¹⁵. It would therefore be useful to extend our ideas

similarly, and investigate what effect (if any) multiplicities have on symmetry-adjusted classical shadows.

Regarding the protocols considered, we have focused on local observable estimation in systems with global $U(1)$ symmetry. However, it is worth noting that the n -qubit Clifford group possesses only one nontrivial irrep, making it essentially compatible with any symmetry. Because its shadow norm is exponentially large for local observables, it is an unfavorable choice for typical quantum simulation applications. One may wonder whether the desirable universality of this irrep can nonetheless be harnessed, analogous to how we constructed the subsystem-symmetrized Pauli shadows. A particularly interesting candidate for future studies would be the global $SU(2)/Cl(1)$ control, introduced in ref. 116, which features both a group with high amounts of symmetry and low sample complexity for local observable estimation. Similarly, the single-fermion $U(n)$ basis rotations used in ref. 72 would also be a promising option to study. Alternatively, one may consider different classes of symmetries, such as local (rather than global) symmetries.

One key advantage of symmetry adjustment is its flexibility, allowing for easy integration with other error-mitigation strategies. Investigating this interplay is a clear target for future work. Particularly valuable would be other techniques to massage the circuit noise into approximately satisfying Assumptions 1, for instance by randomized compiling¹⁷. From our usage of QAA³⁵ in the numerical experiments, we have already shown heuristically that the mere choice of qubit assignments appears to have such an effect.

Indeed, the reliance on such assumptions for rigorous guarantees may be viewed as a limitation of this work. While our numerical results are encouraging, it behooves one to seek a more comprehensive error analysis applicable to a wider range of noise models. For example, while gate-dependent errors are particularly detrimental to our method, they have been closely studied in the context of randomized benchmarking^{118–121}. The tools developed therein may be valuable to this setting as well. Establishing a better understanding here may also inspire extensions to surpass the limitations of the current theory. We leave such goals to future work.

Note added.—Shortly after our manuscript appeared on the arXiv preprint server, two related works^{122,123} subsequently appeared. The former develops a calibration estimator equivalent to our Eq. (43), while the latter analytically studies the effects of gate-dependent noise on Clifford shadow protocols. The formulation and analyses of symmetry-adjusted classical shadows remain original to our manuscript.

Methods

Theory of symmetry-adjusted classical shadows

Here we describe the theory behind the symmetry-adjusted classical shadows estimator. This approach uses known symmetry information about the ideal, noiseless state ρ that we wish to prepare (but are only able to produce a noisy version of). In this section, we describe the idea for an arbitrary multiplicity-free group G ; in the subsequent subsections, we will provide concrete applications to the efficient estimation of local fermionic and qubit observables, respectively.

Suppose ρ is a quantum state obeying a known symmetry, corresponding to a collection of operators $S_\lambda \in V_\lambda$ for which the values $s_\lambda := \langle\langle S_\lambda | \rho \rangle\rangle$ are known a priori. For example, if the system has a symmetry operator S which spans multiple irreps, then we can construct S_λ using the projectors Π_λ :

$$|S_\lambda\rangle\rangle = \Pi_\lambda |S\rangle\rangle. \tag{64}$$

By construction, S_λ is an eigenoperator of both \mathcal{M} and $\widetilde{\mathcal{M}}$:

$$\begin{aligned} \mathcal{M}|S_\lambda\rangle\rangle &= f_\lambda |S_\lambda\rangle\rangle, \\ \widetilde{\mathcal{M}}|S_\lambda\rangle\rangle &= \widetilde{f}_\lambda |S_\lambda\rangle\rangle. \end{aligned} \tag{65}$$

If one is interested in only a subset $R' \subseteq R_G$ of the irreps, then it suffices to only know those symmetries S_λ for which $\lambda \in R'$.

Because the ideal values of s_λ and f_λ are already known, we can use the estimated noisy expectation value of S_λ to build an estimate for \widetilde{f}_λ . We start with the standard postprocessing of classical shadows: applying \mathcal{M}^{-1} to the measurement outcomes of the noisy quantum experiments produces, in expectation, the effective state

$$|\widetilde{\rho}\rangle\rangle := \mathcal{M}^{-1} \widetilde{\mathcal{M}} |\rho\rangle\rangle = \mathbb{E}_{g \sim G, b \sim \widetilde{\mathcal{U}}_g |\rho\rangle\rangle} \mathcal{M}^{-1} \mathcal{U}_g^\dagger |b\rangle\rangle, \tag{66}$$

which clearly differs from $|\rho\rangle\rangle$ when $\widetilde{\mathcal{M}} \neq \mathcal{M}$. Nonetheless, we can use this noisy data to estimate the value of $\langle\langle S_\lambda | \widetilde{\rho} \rangle\rangle$, which is equal to

$$\langle\langle S_\lambda | \widetilde{\rho} \rangle\rangle = \frac{\widetilde{f}_\lambda}{f_\lambda} s_\lambda \tag{67}$$

by Eq. (65). In fact, this relation applies to any $O \in V_\lambda$:

$$\langle\langle O | \widetilde{\rho} \rangle\rangle = \frac{\widetilde{f}_\lambda}{f_\lambda} \langle\langle O | \rho \rangle\rangle. \tag{68}$$

Hence while we use Eq. (67) to learn \widetilde{f}_λ from the symmetry S_λ , this is in turn applicable to all other operators within the same irrep. This leads to the recovery of the ideal expectation values as

$$\langle\langle O | \rho \rangle\rangle = \frac{\langle\langle O | \widetilde{\rho} \rangle\rangle}{\langle\langle S_\lambda | \widetilde{\rho} \rangle\rangle / s_\lambda}. \tag{69}$$

Having established the theory in expectation, we now analyze the implementation in practice. Let T be the number of classical-shadow snapshots, $|\widehat{\rho}_\ell\rangle\rangle = \mathcal{M}^{-1} \mathcal{U}_{g_\ell}^\dagger |b_\ell\rangle\rangle$ for $\ell = 1, \dots, T$, obtained by sampling the noisy quantum computer. Recall that these snapshots converge to $\widetilde{\rho}$ rather than ρ . From their empirical average, $\widehat{\rho}(T) = (1/T) \sum_{\ell=1}^T |\widehat{\rho}_\ell\rangle\rangle$, we can estimate the lefthand side of Eq. (67) as

$$\widehat{s}_\lambda(T) := \langle\langle S_\lambda | \widehat{\rho}(T) \rangle\rangle. \tag{70}$$

This in turn provides an estimate for \widetilde{f}_λ ,

$$\widehat{f}_\lambda(T) := f_\lambda \frac{\widehat{s}_\lambda(T)}{s_\lambda}. \tag{71}$$

This can be understood as a generalization of NoiseEst_G(λ, g, b) from Eq. (23), making the replacements $D_\lambda \rightarrow S_\lambda$ and $|0^n\rangle\langle 0^n| \rightarrow \rho$. Indeed, one can view the calibration state $|0^n\rangle$ as obeying the symmetries given by its stabilizer group.

Consider the estimation of observables O_1, \dots, O_L with symmetry-adjusted classical shadows. If any observable is supported over multiple irreps, then we can always decompose it as a linear combination of basis elements across those irreps. Thus without loss of generality we suppose that each $O_j \in V_\lambda$ for some $\lambda \in R'$. From the same noisy classical shadow $\widehat{\rho}(T)$, we also have estimates for their noisy expectation values: $\mathbb{E} \langle\langle O_j | \widehat{\rho}(T) \rangle\rangle = \langle\langle O_j | \widetilde{\rho} \rangle\rangle$. Then, following Eq. (69) we can directly construct error-mitigated estimators as

$$\widehat{o}_j^{\text{EM}}(T) := \frac{\widehat{o}_j(T)}{\widehat{s}_\lambda(T) / s_\lambda}, \tag{72}$$

which converges to $\langle\langle O_j | \rho \rangle\rangle$ in the $T \rightarrow \infty$ limit (if Assumptions 1 hold). Because $\mathbb{E}[X/Y] \neq \mathbb{E}[X]/\mathbb{E}[Y]$ (for nontrivial random variables X and Y), Eq. (72) describes a biased estimator. In the following theorem, we quantify this bias by bounding the total prediction error of $\widehat{o}_j^{\text{EM}}(T)$. This in turn bounds the number of symmetry-adjusted classical-shadow samples T required.

Theorem 4. Fix accuracy and confidence parameters $\epsilon, \delta \in (0, 1)$. Let O_1, \dots, O_L be a collection of observables, each supported on an irrep of $\mathcal{U} : G \rightarrow \text{U}(\mathcal{L}(\mathcal{H}))$ as $O_j \in V_\lambda$ for $\lambda \in R' \subseteq R_G$. Let $S_\lambda \in V_\lambda$ be a symmetry operator for each $\lambda \in R'$, for which the ideal values $s_\lambda = \text{tr}(S_\lambda \rho)$ of the target state ρ are known a priori. Suppose that each noisy unitary satisfies Assumptions 1, $\widetilde{\mathcal{U}}_g = \mathcal{E} \mathcal{U}_g$, and define the quantities

$$F_{Z, R'}(\mathcal{E}) := \min_{\lambda \in R'} \frac{\text{tr}(\mathcal{E} \mathcal{M}_Z \Pi_\lambda)}{\text{tr}(\mathcal{M}_Z \Pi_\lambda)}, \tag{73}$$

$$\sigma^2 := \max_{1 \leq j \leq L, \lambda \in R'} \left\{ \text{Var}[\widehat{o}_j], \text{Var} \left[\frac{\widehat{s}_\lambda}{s_\lambda} \right] \right\}. \tag{74}$$

Then, a (noisy) classical shadow $\hat{\rho}(T)$ of size

$$T = \mathcal{O}\left(\frac{\log((L + |R'|)/\delta)}{F_{Z,R'}(\mathcal{E})^2 \epsilon^2}\right) \quad (75)$$

can be used to construct error-mitigated estimates

$$\hat{\rho}_j^{\text{EM}}(T) := \frac{\text{tr}(O_j \hat{\rho}(T))}{\text{tr}(S_\lambda \hat{\rho}(T))/s_\lambda} \quad (76)$$

which obey

$$|\hat{\rho}_j^{\text{EM}}(T) - \text{tr}(O_j \rho)| \leq (\|O_j\|_\infty + 1)\epsilon + \mathcal{O}(\|O_j\|_\infty \epsilon^2) \quad (77)$$

for all $1 \leq j \leq L$, with success probability at least $1 - \delta$.

The proof of this statement is provided in Supplementary Note 1. Note that $\|\cdot\|_\infty$ denotes the spectral (operator) norm. We phrase this result in terms of variances, rather than the state-independent shadow norm, because knowledge about ρ (namely, its symmetries) can potentially provide tighter bounds. Note that the variance is with respect to the effective noisy state $\tilde{\rho}$, which was defined in Eq. (66).

Let us make a few remarks on this result. First, although the symmetry operators appear in the denominator of Eq. (72), they affect the sample complexity as usual for classical shadows, albeit normalized by the value s_λ of the symmetry sector. Thus the division by $\hat{s}_\lambda(T)$ does not hinder our control over the variance, except when $s_\lambda = 0$ (we furthermore show how to handle such pathological cases in the application examples below). For typical applications, the variance $\text{Var}[\hat{s}_\lambda/s_\lambda] \leq \|S_\lambda/s_\lambda\|_{\text{shadow}}^2$ will be comparable to the baseline variance of estimation, $\text{Var}[\hat{\rho}_j] \leq \max_j \|O_j\|_{\text{shadow}}^2$. Additionally, the number of irreps considered is typically $|R'| \ll L$ (for instance, in the concrete examples considered in this work, $|R'|$ is a constant). Thus, we expect that the inclusion of symmetry operators incurs negligible overheads for most applications.

Instead, the primary overhead arises from the fact that error-mitigated estimation necessarily comes at the cost of larger overall variances^{95–98}. The quantity

$$F_{Z,R'}(\mathcal{E}) = \min_{\lambda \in R'} \frac{\text{tr}(\mathcal{E} \mathcal{M}_Z \Pi_\lambda)}{\text{tr}(\mathcal{M}_Z \Pi_\lambda)} \quad (78)$$

characterizes an effective noise strength, and it can be seen as a generalization of the average Z -basis fidelity of \mathcal{E} ,

$$F_Z(\mathcal{E}) = \frac{\text{tr}(\mathcal{E} \mathcal{M}_Z)}{\text{tr}(\mathcal{M}_Z)} = \frac{1}{2^n} \sum_{b \in \{0,1\}^n} \langle\langle b|\mathcal{E}|b\rangle\rangle, \quad (79)$$

which appears in prior works on noise-robust classical shadows^{85,86}. In contrast to $F_Z(\mathcal{E})$, the quantity $F_{Z,R'}(\mathcal{E})$ is a more fine-grained characterization of the noise channel, averaged within the relevant subspaces V_λ . Similar to prior results^{85–89}, the sampling overhead of our error-mitigated estimates also depends inverse quadratically on this noise fidelity.

Finally, the error bound we obtain is $\mathcal{O}(\|O_j\|_\infty \epsilon)$ when $\epsilon < 1$. Note that $\|O_j\|_\infty = 1$ for Pauli and Majorana operators. Our result also features error terms of order $\mathcal{O}(\|O_j\|_\infty \epsilon^2)$, which reflect the biased nature of $\hat{\rho}_j^{\text{EM}}(T)$ as a ratio of two random variables. Nonetheless, our theorem establishes that this bias vanishes as $\epsilon^2 \sim 1/T$, so that for sufficiently large T the prediction error is dominated by the standard shot-noise scaling of $\epsilon \sim 1/\sqrt{T}$.

Application to fermionic (matchgate) shadows

The first application of symmetry-adjusted classical shadows that we consider is the estimation of local fermionic observables. This is achieved efficiently by fermionic classical shadows⁶⁹, wherein the group G corresponds

fermionic Gaussian unitaries (also referred to as matchgate shadows⁷⁰). We will consider a commonly encountered symmetry in fermionic systems: fixed particle number. However, it will be clear how the general idea can apply to other symmetries, such as spin.

Background on matchgate shadows. We begin with a review of matchgate shadows. Let a_p, a_p be creation and annihilation operators for a system of n fermionic modes, $p \in [n]$. The associated Majorana operators are

$$\gamma_{2p} = a_p + a_p^\dagger, \quad \gamma_{2p+1} = -i(a_p - a_p^\dagger). \quad (80)$$

Under the Jordan–Wigner transformation¹⁰², these are mapped to Pauli operators as

$$\gamma_{2p} = \left(\prod_{q < p} Z_q\right) X_p, \quad \gamma_{2p+1} = \left(\prod_{q < p} Z_q\right) Y_p. \quad (81)$$

Recall from Eq. (2) that all d^2 basis operators are generated by taking arbitrary products:

$$\Gamma_\mu = (-i)^{\binom{m}{2}} \gamma_{\mu_1} \cdots \gamma_{\mu_m}, \quad (82)$$

where $\mu = (\mu_1, \dots, \mu_m) \subseteq [2n]$. By convention, we order $\mu_1 < \dots < \mu_m$. We can group all the m -degree Majorana indices by defining the set

$$\mathcal{C}_{2n,m} := \{\mu \subseteq [2n] : |\mu| = m\}. \quad (83)$$

Physical fermionic observables have even degree $m = 2k$. An important subset of such operators comprises those which are diagonal in the standard basis, corresponding to the index set

$$\mathcal{D}_{2n,2k} := \{(2p_1, 2p_1 + 1, \dots, 2p_k, 2p_k + 1) : \mathbf{p} \in \mathcal{C}_{n,k}\}. \quad (84)$$

Using Eq. (81), each $\tau \in \mathcal{D}_{2n,2k}$ corresponds to the Pauli- Z operator $\Gamma_\tau = Z_{p_1} \cdots Z_{p_k}$ under the Jordan–Wigner mapping.

The group of fermionic Gaussian unitaries is the image of the homomorphism $U : \text{O}(2n) \rightarrow \text{U}(d)$ whose adjoint action obeys

$$U_Q \gamma_\mu U_Q^\dagger = \sum_{\nu \in [2n]} Q_{\nu\mu} \gamma_\nu, \quad Q \in \text{O}(2n). \quad (85)$$

These unitaries are equivalent to (generalized) matchgate circuits¹²⁴ and constitute a class of classically simulatable circuits^{125–130}. Fermionic (matchgate) shadows then randomize over certain subgroups $G \subseteq \text{O}(2n)$ of these Gaussian unitaries. The measurement channel takes the form

$$\mathcal{M} = \sum_{k=0}^n f_{2k} \Pi_{2k}, \quad (86)$$

where the eigenvalues are

$$f_{2k} = \binom{n}{k} / \binom{2n}{2k} \quad (87)$$

and each irrep is the image of

$$\Pi_{2k} = \frac{1}{d} \sum_{\mu \in \mathcal{C}_{2n,2k}} |\Gamma_\mu\rangle\langle\Gamma_\mu|. \quad (88)$$

While \mathcal{U} carries $2n + 1$ unique irreps (each labeled by a Majorana degree m)^{115,124}, only the $n + 1$ irreps $\lambda = 2k$ have nonvanishing f_λ ^{69,70}. Therefore \mathcal{M}^{-1} is formally the pseudoinverse restricted to those

subspaces. Finally, the shadow norm of k -body Majorana operators is⁶⁹

$$\|\Gamma_\mu\|_{\text{shadow}}^2 = f_{2k}^{-1} = \mathcal{O}(n^k). \tag{89}$$

Variance expressions for arbitrary observables can be found in refs.^{70,71}. For the postprocessing of T shadows into estimates of all k -body Majorana observables, we describe an algorithm in Supplementary Note 5 which runs in time $\mathcal{O}(n^k T)$.

We now comment on the choice of $G \subseteq O(2n)$. Fermionic classical shadows were introduced in ref. 69, which initially considered the intersection of proper matchgate circuits [the special orthogonal group $SO(2n)$] with n -qubit Clifford unitaries $Cl(n)$. The result is the group of all $2n \times 2n$ signed permutation matrices with determinant 1, denoted by $B^+(2n) \subset SO(2n)$. They also showed that its unsigned subgroup, $Alt(2n) \subset B^+(2n)$, possesses the same irrep structure [ref. 69, Supplemental Material, Theorem 11]. While the full, continuous group $SO(2n)$ has not yet been analyzed for classical shadows, it was studied for character randomized benchmarking¹⁴ in [ref. 115, Sec. VI], wherein they demonstrated the presence of multiplicities. These multiplicities can be avoided by enlarging to the generalized matchgate group, i.e., all of $O(2n)$ [ref. 124, Lemma 3]. Ref. 70 applied these generalized matchgates to fermionic classical shadows, and in particular they prove that the Clifford intersection in this setting (now yielding the subgroup $B(2n) \subset O(2n)$ of signed permutation matrices with either determinant ± 1) is a 3-design for $O(2n)$. This implies that $B(2n)$ is also multiplicity-free.

Due to the variety of options, for the rest of this paper we assume matchgate shadows under any G with the desired irreps. We note that ref. 71 introduced a smaller subset of $B(2n)$ based on perfect matchings, which has the same channel \mathcal{M} and variances; however, its connection to representation theory was not explored.

Utilizing particle-number symmetry. Suppose the ideal state we wish to prepare lies in the η -particle sector of \mathcal{H} . This is a $U(1)$ symmetry generated by the fermion-number operator, $N = \sum_{p \in [n]} a_p^\dagger a_p$. In particular, powers of N obey

$$\text{tr}(N^k \rho) = \eta^k, \tag{90}$$

which provides us a collection of conserved quantities with which to perform symmetry adjustment. Recall from Eq. (88) that Π_m projects onto the irrep

$$V_m = \text{span}\{\Gamma_\mu : \mu \in \mathcal{C}_{2n,m}\}. \tag{91}$$

Then, projecting N^k onto V_{2k} yields the symmetry operators S_{2k} , and solving the resulting linear system of equations recovers the ideal values for $s_{2k} = \text{tr}(S_{2k} \rho)$. For ease of exposition we will consider only $k = 1, 2$, but one may generalize to higher k using these ideas.

Concretely, we start with the fact that $a_p^\dagger a_p = (\mathbb{I} - \Gamma_{(2p, 2p+1)})/2$, and $\Gamma_{(2p, 2p+1)} \Gamma_{(2q, 2q+1)} = \Gamma_{(2p, 2p+1, 2q, 2q+1)}$ for $p < q$. Then, expanding N and N^2 into a linear combination of Majorana operators, one finds

$$S_2 = \Pi_2(N) = -\frac{1}{2} \sum_{\mu \in \mathcal{D}_{2n,2}} \Gamma_\mu, \tag{92}$$

$$S_4 = \Pi_4(N^2) = \frac{1}{2} \sum_{\mu \in \mathcal{D}_{2n,4}} \Gamma_\mu. \tag{93}$$

Using Eq. (90) and the relations between S_2 and S_4 to N and N^2 (for example, $N = n\mathbb{I}/2 + S_2$), we arrive at:

$$s_2 = \text{tr}(S_2 \rho) = \eta - \frac{n}{2}, \tag{94}$$

$$s_4 = \text{tr}(S_4 \rho) = \frac{1}{2} \binom{n}{2} - \eta(n - \eta). \tag{95}$$

For the sampling cost incurred by these symmetry operators, we argue that the typical shadow norms of these symmetries are $\|S_{2k}/s_{2k}\|_{\text{shadow}}^2 = \mathcal{O}(n^k)$, which is the same as the base estimation. To see this, consider a triangle inequality on the shadow norm:

$$\begin{aligned} \|S_{2k}\|_{\text{shadow}} &\leq \frac{1}{2} \sum_{\mu \in \mathcal{D}_{2n,2k}} \|\Gamma_\mu\|_{\text{shadow}} \\ &= \frac{1}{2} \binom{n}{k} \sqrt{\frac{\binom{2n}{2k}}{\binom{n}{k}}} \\ &= \mathcal{O}(n^{3k/2}). \end{aligned} \tag{96}$$

Thus $\|S_2\|_{\text{shadow}}^2 = \mathcal{O}(n^3)$ and $\|S_4\|_{\text{shadow}}^2 = \mathcal{O}(n^6)$. Next, we need to examine how s_{2k}^2 scales with system size. Assuming that $s_2, s_4 \neq 0$ and that the number of electrons is $\eta = \mathcal{O}(n)$, then from Eqs. (94) and (95) we see that $s_2^2 = \Theta(n^2)$ and $s_4^2 = \Theta(n^4)$. Thus

$$\|S_2/s_2\|_{\text{shadow}}^2 = \mathcal{O}(n), \tag{97}$$

$$\|S_4/s_4\|_{\text{shadow}}^2 = \mathcal{O}(n^2). \tag{98}$$

Avoiding division by zero. One potential obstruction to symmetry adjustment is when some $s_{2k} = 0$. This can occur whenever the particle number takes a specific value:

$$s_2 = 0 \text{ if } \eta = \frac{n}{2}, \tag{99}$$

$$s_4 = 0 \text{ if } \eta = \frac{n \pm \sqrt{n}}{2}. \tag{100}$$

Equation (99) occurs at half-filling, which is fairly common. On the other hand, Eq. (100) occurs only when the number of modes n is a perfect square and the number of particles η is one of two specific values, so it is less likely to occur. Nonetheless, there is a straightforward way to circumvent both possibilities by introducing a single ancilla qubit.

To do so, append an additional fermion mode initialized in the unoccupied state $|0\rangle$, so that the ideal state is now the $(n+1)$ -mode state $\rho' = \rho \otimes |0\rangle\langle 0|$. Given that ρ has η particles on n modes, ρ' is an η -particle state on $n+1$ modes. The new symmetry operators on the $(n+1)$ -mode Hilbert space are

$$S'_2 = -\frac{1}{2} \sum_{\mu \in \mathcal{D}_{2(n+1),2}} \Gamma_\mu, \tag{101}$$

$$S'_4 = \frac{1}{2} \sum_{\mu \in \mathcal{D}_{2(n+1),4}} \Gamma_\mu, \tag{102}$$

which have ideal values

$$s'_2 = \text{tr}(S'_2 \rho') = \eta - \frac{n+1}{2}, \tag{103}$$

$$s'_4 = \text{tr}(S'_4 \rho') = \frac{1}{2} \binom{n+1}{2} - \eta(n+1 - \eta). \tag{104}$$

It is straightforward to check that, if either condition Eq. (99) or Eq. (100) holds, then s'_2 and s'_4 are always nonzero for $n > 1$.

Under the Jordan-Wigner mapping, this modification is easily achieved by initializing a single ancilla qubit in $|0\rangle$. Recall that the terms in the symmetries S_2, S_4 are the diagonal operators $\Gamma_{(2p, 2p+1)} = Z_p$ and $\Gamma_{(2p, 2p$

$+1, 2q, 2q+1) = Z_p Z_q$. Note also that the ancilla qubit is acted on only during the random unitary U_Q (where now Q has dimension $2n + 2$) and otherwise does not interact with the n system qubits.

Application to qubit (Pauli) shadows

Now we turn to the application for local observable estimation in systems of spin-1/2 particles (qubits). Random Pauli measurements are efficient for this task; however, for compatibility with the global $U(1)$ symmetry considered in this work, we must slightly modify the protocol to accommodate its irreps. We begin with a review of the standard Pauli shadows protocol, followed by our modification.

Background on standard Pauli shadows. The local Clifford group $Cl(1)^{\otimes n}$ is implemented by uniformly drawing a single-qubit Clifford gate for each qubit independently. It has 2^n irreducible representations, corresponding to all k -qubit subsystems $I \subseteq [n]$, where $|I| = k \in \{0, 1, \dots, n\}$ ¹³¹. Twirling \mathcal{M}_Z by this group yields

$$\mathcal{M} = \sum_{I \subseteq [n]} f_I \Pi_I, \tag{105}$$

where $f_I = 3^{-|I|}$ and Π_I projects onto the subspace of operators which act nontrivially on precisely the subsystem I . The squared shadow norm for k -local Pauli operators P is⁵¹

$$\|P\|_{\text{shadow}}^2 = 3^k. \tag{106}$$

A more general variance bound was derived in ref. 52: a simple loose bound of their result can be stated as $\text{Var}[\hat{o}] \leq 3^k R \|O\|_{\infty}^2$, where O is an arbitrary k -local traceless observable and R is the number of terms in its Pauli decomposition. However, they argue that a tighter expression, essentially $3^k \|O\|_{\infty}^2$, is typically a good approximation to the variance.

Subsystem symmetrization of Pauli shadows. The irreps of $Cl(1)^{\otimes n}$ are difficult to reconcile with commonly encountered symmetries. For example, consider a conserved total magnetization $M = \sum_{i \in [n]} Z_i$. In terms of qubits, this is equivalent to the different Hamming-weight sectors. Each term Z_i lies in a different irrep $I = \{i\}$, so M spans multiple irreps rather than having a single conserved quantity per irrep.

To remedy this conflict, we introduce what we call subsystem-symmetrized Pauli shadows, which randomizes over a group whose irreps are labeled only by the qubit locality k , rather than any specific subsystem I of k qubits. (This is analogous to how the matchgate irreps depend only on fermionic locality, due to the inherent antisymmetry of fermions.) We formalize the group as follows.

Definition 5. The subsystem-symmetrized local Clifford group is defined as $Cl(1)_{\text{Sym}}^{\otimes n} := \text{Sym}(n) \times Cl(1)^{\otimes n}$, where $\text{Sym}(n)$ is the symmetric group and $Cl(1)$ is the single-qubit Clifford group. Its unitary action on \mathcal{H} is given by

$$U_{(\pi, C)} = S_{\pi} C, \tag{107}$$

where $C = \bigotimes_{i \in [n]} C_i \in Cl(1)^{\otimes n}$ and $\pi \in \text{Sym}(n)$ is represented by a permutation of the n qubits:

$$S_{\pi} |b_0\rangle \cdots |b_{n-1}\rangle = |b_{\pi^{-1}(0)}\rangle \cdots |b_{\pi^{-1}(n-1)}\rangle \tag{108}$$

for all $b_i \in \{0, 1\}$, $i \in [n]$.

The unitaries S_{π} can be implemented with $\mathcal{O}(n^2)$ gates and depth $\mathcal{O}(n)$, for example by constructing a parallelized network of nearest-neighbor SWAP gates according to an odd-even sorting algorithm⁹⁹ applied to π . Representing π as an array of the permuted elements of $[n]$, the sorting algorithm returns a sequence of adjacent transpositions $i \leftrightarrow i + 1$ which maps π to $(0, 1, \dots, n - 1)$. This sequence therefore implements π^{-1} as

desired. Each such transposition then maps to a SWAP _{$i, i+1$} gate to construct the quantum circuit. For the postprocessing of T shadows into k -local Pauli estimates, we review in Supplementary Note 5 the algorithm which runs in time $\mathcal{O}(n^k T)$.

We prove the relevant properties of subsystem-symmetrized Pauli shadows in Supplementary Note 2, namely its irreps and the shadow norm of local observables. We summarize the results here: each irrep is the space of all k -local operators,

$$V_k = \text{span}(\mathcal{B}_k) \text{ where } \mathcal{B}_k := \{P \in \mathcal{P}(n) : |P| = k\}, \tag{109}$$

for each $k \in \{0, 1, \dots, n\}$. Hence the (noisy) measurement channel is

$$\widetilde{\mathcal{M}} = \sum_{k=0}^n \widetilde{f}_k \Pi_k, \tag{110}$$

where $\widetilde{f}_k = \text{tr}(\mathcal{M}_Z \mathcal{E} \Pi_k) / \binom{n}{k}$ and

$$\Pi_k = \frac{1}{d} \sum_{P \in \mathcal{B}_k} |P\rangle\langle P|. \tag{111}$$

When \mathcal{E} is the identity channel, we recover $f_k = 3^{-k}$. Also in the absence of noise, the variance formulas are exactly the same as in standard Pauli shadows.

The canonical example we have considered in this paper is a $U(1)$ symmetry generated by a total magnetization $M = \sum_{i \in [n]} Z_i$. Suppose the ideal state has a known value of $m = \text{tr}(M\rho)$ (equivalently, ρ lives in a sector of fixed Hamming weight $(n - m)/2$). The symmetries projected into the irreps of $Cl(1)_{\text{Sym}}^{\otimes n}$ are then

$$S_1 = \Pi_1(M) = \sum_{i \in [n]} Z_i, \tag{112}$$

$$S_2 = \Pi_2(M^2) = 2 \sum_{i < j} Z_i Z_j, \tag{113}$$

whose ideal values are

$$s_1 = m, \tag{114}$$

$$s_2 = m^2 - n. \tag{115}$$

As in the fermionic setting, we encounter issues if s_1 or s_2 vanish (i.e., $m = 0$ or $m = \pm \sqrt{n}$, respectively). In this case, we can perform the same ancilla trick, appending a qubit in $|0\rangle$ and modifying the conserved quantities to

$$s'_1 = m + 1, \tag{116}$$

$$s'_2 = (m + 1)^2 - (n + 1). \tag{117}$$

The variances of the symmetry operators are

$$\text{Var}[\hat{s}_1/s_1] = \mathcal{O}(n), \tag{118}$$

$$\text{Var}[\hat{s}_2/s_2] = \mathcal{O}(1), \tag{119}$$

whenever the ideal state lives in a symmetry sector of constant $m = \Theta(1)$. We show this in Supplementary Note 2, along with general m -dependent expressions. This n -dependent variance bound reflects the fact that the symmetries are extensive properties. While local Pauli operators have variances bounded by a constant, we point out that many local observables of interest are linear combinations of an extensive number of Pauli terms. As such, their shadow norms typically grow with system size as well (which can also be seen by the fact that the shadow norm scales with operator norm).

Data availability

The data used in this work are available from the corresponding author upon reasonable request.

Code availability

The code for running the numerical experiments is available at this link (<https://github.com/zhao-andrew/symmetry-adjusted-classical-shadows>).

Received: 7 November 2023; Accepted: 23 May 2024;

Published online: 08 June 2024

References

- Preskill, J. Quantum computing in the NISQ era and beyond. *Quantum* **2**, 79 (2018).
- Bharti, K. et al. Noisy intermediate-scale quantum algorithms. *Rev. Mod. Phys.* **94**, 015004 (2022).
- Feynman, R. P. Simulating physics with computers. *Int. J. Theor. Phys.* **21**, 467–488 (1982).
- Georgescu, I. M., Ashhab, S. & Nori, F. Quantum simulation. *Rev. Mod. Phys.* **86**, 153–185 (2014).
- McArdle, S., Endo, S., Aspuru-Guzik, A., Benjamin, S. C. & Yuan, X. Quantum computational chemistry. *Rev. Mod. Phys.* **92**, 015003 (2020).
- Bauer, B., Bravyi, S., Motta, M. & Chan, G. K.-L. Quantum algorithms for quantum chemistry and quantum materials science. *Chem. Rev.* **120**, 12685–12717 (2020).
- Peruzzo, A. et al. A variational eigenvalue solver on a photonic quantum processor. *Nat. Commun.* **5**, 4213 (2014).
- McClean, J. R., Romero, J., Babbush, R. & Aspuru-Guzik, A. The theory of variational hybrid quantum-classical algorithms. *New J. Phys.* **18**, 023023 (2016).
- Yuan, X., Endo, S., Zhao, Q., Li, Y. & Benjamin, S. C. Theory of variational quantum simulation. *Quantum* **3**, 191 (2019).
- Cerezo, M. et al. Variational quantum algorithms. *Nat. Rev. Phys.* **3**, 625–644 (2021).
- Osborne, T. J. Efficient approximation of the dynamics of one-dimensional quantum spin systems. *Phys. Rev. Lett.* **97**, 157202 (2006).
- Bravyi, S., Gosset, D. & Movassagh, R. Classical algorithms for quantum mean values. *Nat. Phys.* **17**, 337–341 (2021).
- Napp, J. C., La Placa, R. L., Dalzell, A. M., Brandão, F. G. S. L. & Harrow, A. W. Efficient classical simulation of random shallow 2D quantum circuits. *Phys. Rev. X* **12**, 021021 (2022).
- Wild, D. S. & Alhambra, A. M. Classical simulation of short-time quantum dynamics. *PRX Quantum* **4**, 020340 (2023).
- Fowler, A. G., Mariantoni, M., Martinis, J. M. & Cleland, A. N. Surface codes: Towards practical large-scale quantum computation. *Phys. Rev. A* **86**, 032324 (2012).
- Kelly, J. et al. State preservation by repetitive error detection in a superconducting quantum circuit. *Nature* **519**, 66–69 (2015).
- Egan, L. et al. Fault-tolerant control of an error-corrected qubit. *Nature* **598**, 281–286 (2021).
- Postler, L. et al. Demonstration of fault-tolerant universal quantum gate operations. *Nature* **605**, 675–680 (2022).
- Zhao, Y. et al. Realization of an error-correcting surface code with superconducting qubits. *Phys. Rev. Lett.* **129**, 030501 (2022).
- Sundaresan, N. et al. Demonstrating multi-round subsystem quantum error correction using matching and maximum likelihood decoders. *Nat. Commun.* **14**, 2852 (2023).
- Google Quantum AI. Suppressing quantum errors by scaling a surface code logical qubit. *Nature* **614**, 676–681 (2023).
- Sivak, V. V. et al. Real-time quantum error correction beyond break-even. *Nature* **616**, 50–55 (2023).
- Ni, Z. et al. Beating the break-even point with a discrete-variable-encoded logical qubit. *Nature* **616**, 56–60 (2023).
- O'Malley, P. J. J. et al. Scalable quantum simulation of molecular energies. *Phys. Rev. X* **6**, 031007 (2016).
- Kandala, A. et al. Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. *Nature* **549**, 242–246 (2017).
- Colless, J. I. et al. Computation of molecular spectra on a quantum processor with an error-resilient algorithm. *Phys. Rev. X* **8**, 011021 (2018).
- Dumitrescu, E. F. et al. Cloud quantum computing of an atomic nucleus. *Phys. Rev. Lett.* **120**, 210501 (2018).
- Hempel, C. et al. Quantum chemistry calculations on a trapped-ion quantum simulator. *Phys. Rev. X* **8**, 031022 (2018).
- Kandala, A. et al. Error mitigation extends the computational reach of a noisy quantum processor. *Nature* **567**, 491–495 (2019).
- Kokail, C. et al. Self-verifying variational quantum simulation of lattice models. *Nature* **569**, 355–360 (2019).
- Nam, Y. et al. Ground-state energy estimation of the water molecule on a trapped-ion quantum computer. *npj Quantum Inf.* **6**, 33 (2020).
- Arute, F. et al. Quantum supremacy using a programmable superconducting processor. *Nature* **574**, 505–510 (2019).
- Rubin, N. C. et al. Hartree-Fock on a superconducting qubit quantum computer. *Science* **369**, 1084–1089 (2020).
- Harrigan, M. P. et al. Quantum approximate optimization of non-planar graph problems on a planar superconducting processor. *Nat. Phys.* **17**, 332–336 (2021).
- Jiang, Z. et al. Observation of separated dynamics of charge and spin in the Fermi-Hubbard model. *arXiv:2010.07965* <https://arxiv.org/abs/2010.07965> (2020).
- Zhong, H.-S. et al. Quantum computational advantage using photons. *Science* **370**, 1460–1463 (2020).
- Huggins, W. J. et al. Unbiasing fermionic quantum Monte Carlo with a quantum computer. *Nature* **603**, 416–420 (2022).
- Kim, Y. et al. Scalable error mitigation for noisy quantum circuits produces competitive expectation values. *Nat. Phys.* **19**, 752–759 (2023).
- Huang, H.-Y. et al. Quantum advantage in learning from experiments. *Science* **376**, 1182–1186 (2022).
- Stanisic, S. et al. Observing ground-state properties of the Fermi-Hubbard model using a scalable algorithm on a quantum computer. *Nat. Commun.* **13**, 5743 (2022).
- Tazhigulov, R. N. et al. Simulating models of challenging correlated molecules and materials on the Sycamore quantum processor. *PRX Quantum* **3**, 040318 (2022).
- Madsen, L. S. et al. Quantum computational advantage with a programmable photonic processor. *Nature* **606**, 75–81 (2022).
- Motta, M. et al. Quantum chemistry simulation of ground-and excited-state properties of the sulfonium cation on a superconducting quantum processor. *Chem. Sci.* **14**, 2915–2927 (2023).
- O'Brien, T. E. et al. Purification-based quantum error mitigation of pair-correlated electron simulations. *Nat. Phys.* (2023).
- Morvan, A. et al. Phase transition in random circuit sampling. *arXiv:2304.11119* <https://arxiv.org/abs/2304.11119> (2023).
- Kim, Y. et al. Evidence for the utility of quantum computing before fault tolerance. *Nature* **618**, 500–505 (2023).
- Endo, S., Cai, Z., Benjamin, S. C. & Yuan, X. Hybrid quantum-classical algorithms and quantum error mitigation. *J. Phys. Soc. Jpn.* **90**, 032001 (2021).
- Cai, Z. et al. Quantum error mitigation. *Rev. Mod. Phys.* **95**, 045005 (2023).
- Wecker, D., Hastings, M. B. & Troyer, M. Progress towards practical quantum variational algorithms. *Phys. Rev. A* **92**, 042303 (2015).
- Gonthier, J. F. et al. Measurements as a roadblock to near-term practical quantum advantage in chemistry: resource analysis. *Phys. Rev. Res.* **4**, 033154 (2022).

51. Huang, H.-Y., Kueng, R. & Preskill, J. Predicting many properties of a quantum system from very few measurements. *Nat. Phys.* **16**, 1050–1057 (2020).
52. Pains, M., Kalev, A., Padilha, D. & Ruck, B. Estimating expectation values using approximate quantum states. *Quantum* **5**, 413 (2021).
53. Cotler, J. & Wilczek, F. Quantum overlapping tomography. *Phys. Rev. Lett.* **124**, 100401 (2020).
54. Bonet-Monroig, X., Babbush, R. & O'Brien, T. E. Nearly optimal measurement scheduling for partial tomography of quantum states. *Phys. Rev. X* **10**, 031064 (2020).
55. Tilly, J. et al. The variational quantum eigensolver: a review of methods and best practices. *Physics Reports* **986**, 1–128 (2022).
56. Zhao, A. *Learning, Optimizing, and Simulating Fermions with Quantum Computers*. Ph.D. thesis, University of New Mexico (2023). <https://arxiv.org/abs/2312.10399>.
57. Sugiyama, T., Turner, P. S. & Mura, M. Precision-guaranteed quantum tomography. *Phys. Rev. Lett.* **111**, 160406 (2013).
58. Guță, M., Kahn, J., Kueng, R. & Tropp, J. A. Fast state tomography with optimal error bounds. *J. Phys. A: Math. Theor.* **53**, 204001 (2020).
59. Aaronson, S. Shadow tomography of quantum states. *SIAM J. Comput.* **49**, 368–394 (2020).
60. Aaronson, S. & Rothblum, G. N. Gentle measurement of quantum states and differential privacy. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, 322–333 (Association for Computing Machinery, New York, 2019).
61. Elben, A. et al. Mixed-state entanglement from local randomized measurements. *Phys. Rev. Lett.* **125**, 200501 (2020).
62. Rath, A., Branciard, C., Minguzzi, A. & Vermersch, B. Quantum Fisher information from randomized measurements. *Phys. Rev. Lett.* **127**, 260501 (2021).
63. Vitale, V. et al. Estimation of the quantum Fisher information on a quantum processor. *arXiv:2307.16882* <https://arxiv.org/abs/2307.16882> (2023).
64. Levy, R., Luo, D. & Clark, B. K. Classical shadows for quantum process tomography on near-term quantum computers. *arXiv:2110.02965* <https://arxiv.org/abs/2110.02965> (2021).
65. Kunjummen, J., Tran, M. C., Carney, D. & Taylor, J. M. Shadow process tomography of quantum channels. *Phys. Rev. A* **107**, 042403 (2023).
66. Sack, S. H., Medina, R. A., Michailidis, A. A., Kueng, R. & Serbyn, M. Avoiding barren plateaus using classical shadows. *PRX Quantum* **3**, 020365 (2022).
67. Boyd, G. & Koczor, B. Training variational quantum circuits with CoVaR: Covariance root finding with classical shadows. *Phys. Rev. X* **12**, 041022 (2022).
68. Chan, H. H. S., Meister, R., Goh, M. L. & Koczor, B. Algorithmic shadow spectroscopy. *arXiv:2212.11036* <https://arxiv.org/abs/2212.11036> (2022).
69. Zhao, A., Rubin, N. C. & Miyake, A. Fermionic partial tomography via classical shadows. *Phys. Rev. Lett.* **127**, 110504 (2021).
70. Wan, K., Huggins, W. J., Lee, J. & Babbush, R. Matchgate shadows for fermionic quantum simulation. *Commun. Math. Phys.* **404**, 629–700 (2023).
71. O’Gorman, B. Fermionic tomography and learning. *arXiv:2207.14787* <https://arxiv.org/abs/2207.14787> (2022).
72. Low, G. H. Classical shadows of fermions with particle number symmetry. *arXiv:2208.08964* <https://arxiv.org/abs/2208.08964> (2022).
73. Babbush, R. et al. Quantum simulation of exact electron dynamics can be more efficient than classical mean-field methods. *Nat. Commun.* **14**, 4058 (2023).
74. Denzler, J., Mele, A. A., Derbyshire, E., Guaita, T. & Eisert, J. Learning fermionic correlations by evolving with random translationally invariant Hamiltonians. *arXiv:2309.12933* <https://arxiv.org/abs/2309.12933> (2023).
75. Gu, T., Yuan, X. & Wu, B. Efficient measurement schemes for bosonic systems. *Quantum Sci. Technol.* **8**, 045008 (2023).
76. Becker, S., Datta, N., Lami, L. & Rouzé, C. Classical shadow tomography for continuous variables quantum systems. *IEEE Trans. Inf. Theory* **70**, 3427–3452 (2024).
77. Elben, A. et al. The randomized measurement toolbox. *Nat. Rev. Phys.* **5**, 9–24 (2023).
78. Seif, A., Cian, Z.-P., Zhou, S., Chen, S. & Jiang, L. Shadow distillation: Quantum error mitigation with classical shadows for near-term quantum processors. *PRX Quantum* **4**, 010303 (2023).
79. Hu, H.-Y., LaRose, R., You, Y.-Z., Rieffel, E. & Wang, Z. Logical shadow tomography: Efficient estimation of error-mitigated observables. *arXiv:2203.07263* <https://arxiv.org/abs/2203.07263> (2022).
80. McClean, J. R., Jiang, Z., Rubin, N. C., Babbush, R. & Neven, H. Decoding quantum errors with subspace expansions. *Nat. Commun.* **11**, 636 (2020).
81. Koczor, B. Exponential error suppression for near-term quantum devices. *Phys. Rev. X* **11**, 031057 (2021).
82. Huggins, W. J. et al. Virtual distillation for quantum error mitigation. *Phys. Rev. X* **11**, 041036 (2021).
83. Jnane, H., Steinberg, J., Cai, Z., Nguyen, H. C. & Koczor, B. Quantum error mitigated classical shadows. *PRX Quantum* **5**, 010324 (2024).
84. Temme, K., Bravyi, S. & Gambetta, J. M. Error mitigation for short-depth quantum circuits. *Phys. Rev. Lett.* **119**, 180509 (2017).
85. Chen, S., Yu, W., Zeng, P. & Flammia, S. T. Robust shadow estimation. *PRX Quantum* **2**, 030348 (2021).
86. Koh, D. E. & Grewal, S. Classical shadows with noise. *Quantum* **6**, 776 (2022).
87. Karalekas, P. J. et al. A quantum-classical cloud platform optimized for variational hybrid algorithms. *Quantum Sci. Technol.* **5**, 024003 (2020).
88. Van Den Berg, E., Mineev, Z. K. & Temme, K. Model-free readout-error mitigation for quantum expectation values. *Phys. Rev. A* **105**, 032620 (2022).
89. Arrasmith, A., Patterson, A., Boughton, A. & Pains, M. Development and demonstration of an efficient readout error mitigation technique for use in NISQ algorithms. *arXiv:2303.17741* <https://arxiv.org/abs/2303.17741> (2023).
90. Bonet-Monroig, X., Sagastizabal, R., Singh, M. & O’Brien, T. E. Low-cost error mitigation by symmetry verification. *Phys. Rev. A* **98**, 062339 (2018).
91. McArdle, S., Yuan, X. & Benjamin, S. Error-mitigated digital quantum simulation. *Phys. Rev. Lett.* **122**, 180501 (2019).
92. Cai, Z. Quantum error mitigation using symmetry expansion. *Quantum* **5**, 548 (2021).
93. Isakov, S. V. et al. Simulations of quantum circuits with approximate noise using qsim and Cirq. *arXiv:2111.02396* <https://arxiv.org/abs/2111.02396> (2021).
94. Fulton, W. & Harris, J. *Representation Theory: A First Course* (Springer-Verlag, New York, 2004).
95. Takagi, R., Endo, S., Minagawa, S. & Gu, M. Fundamental limits of quantum error mitigation. *npj Quantum Inf.* **8**, 114 (2022).
96. Takagi, R., Tajima, H. & Gu, M. Universal sampling lower bounds for quantum error mitigation. *Phys. Rev. Lett.* **131**, 210602 (2023).
97. Tsubouchi, K., Sagawa, T. & Yoshioka, N. Universal cost bound of quantum error mitigation based on quantum estimation theory. *Phys. Rev. Lett.* **131**, 210601 (2023).
98. Quek, Y., França, D. S., Khatri, S., Meyer, J. J. & Eisert, J. Exponentially tighter bounds on limitations of quantum error

- mitigation. *arXiv:2210.11505* <https://arxiv.org/abs/2210.11505> (2022).
99. Habermann, A. N. Parallel neighbor-sort (or the glory of the induction principle). *Carnegie Mellon University Technical Report No. AD-759-248* https://kilthub.cmu.edu/articles/journal_contribution/Parallel_neighbor-sort_or_the_glory_of_the_induction_principle_/6608258/files/12099395.pdf (1972).
 100. Jiang, Z., Sung, K. J., Kechedzhi, K., Smelyanskiy, V. N. & Boixo, S. Quantum algorithms to simulate many-body physics of correlated fermions. *Phys. Rev. Applied* **9**, 044036 (2018).
 101. Oszmaniec, M., Dangniam, N., Morales, M. E. S. & Zimborás, Z. Fermion sampling: a robust quantum computational advantage scheme using fermionic linear optics and magic input states. *PRX Quantum* **3**, 020328 (2022).
 102. Jordan, P. & Wigner, E. Über das Paulische Äquivalenzverbot. *Z. Phys.* **47**, 631–651 (1928).
 103. Cirq Developers. Cirq <https://github.com/google/cirq> (2023).
 104. Efron, B. Bootstrap methods: Another look at the jackknife. In Kotz, S. & Johnson, N. L. (eds.) *Breakthroughs in Statistics*, 569–593 (Springer, New York, 1992).
 105. Bravyi, S. & König, R. Classical simulation of dissipative fermionic linear optics. *Quantum Inf. Comput.* **12**, 925–943 (2012).
 106. Rubin, N. C., Babbush, R. & McClean, J. Application of fermionic marginal constraints to hybrid quantum algorithms. *New J. Phys.* **20**, 053020 (2018).
 107. Quantum AI team and collaborators. ReCirq <https://doi.org/10.5281/zenodo.4091471> (2020).
 108. Wecker, D. et al. Solving strongly correlated electron models on a quantum computer. *Phys. Rev. A* **92**, 062318 (2015).
 109. Kivlichan, I. D. et al. Quantum simulation of electronic structure with linear depth and connectivity. *Phys. Rev. Lett.* **120**, 110501 (2018).
 110. Garnerone, S., de Oliveira, T. R. & Zanardi, P. Typicality in random matrix product states. *Phys. Rev. A* **81**, 032336 (2010).
 111. Garnerone, S., de Oliveira, T. R., Haas, S. & Zanardi, P. Statistical properties of random matrix product states. *Phys. Rev. A* **82**, 052312 (2010).
 112. Fishman, M., White, S. R. & Stoudenmire, E. M. The ITensor software library for tensor network calculations. *Sci. Post Phys. Codebases* **4** (2022).
 113. White, S. R. Density matrix formulation for quantum renormalization groups. *Phys. Rev. Lett.* **69**, 2863–2866 (1992).
 114. Helsen, J., Xue, X., Vandersypen, L. M. K. & Wehner, S. A new class of efficient randomized benchmarking protocols. *npj Quantum Inf.* **5**, 71 (2019).
 115. Claes, J., Rieffel, E. & Wang, Z. Character randomized benchmarking for non-multiplicity-free groups with applications to subspace, leakage, and matchgate randomized benchmarking. *PRX Quantum* **2**, 010351 (2021).
 116. Van Kirk, K., Cotler, J., Huang, H.-Y. & Lukin, M. D. Hardware-efficient learning of quantum many-body states. *arXiv:2212.06084* <https://arxiv.org/abs/2212.06084> (2022).
 117. Wallman, J. J. & Emerson, J. Noise tailoring for scalable quantum computation via randomized compiling. *Phys. Rev. A* **94**, 052325 (2016).
 118. Proctor, T., Rudinger, K., Young, K., Sarovar, M. & Blume-Kohout, R. What randomized benchmarking actually measures. *Phys. Rev. Lett.* **119**, 130502 (2017).
 119. Wallman, J. J. Randomized benchmarking with gate-dependent noise. *Quantum* **2**, 47 (2018).
 120. Carignan-Dugas, A., Boone, K., Wallman, J. J. & Emerson, J. From randomized benchmarking experiments to gate-set circuit fidelity: how to interpret randomized benchmarking decay parameters. *New J. Phys.* **20**, 092001 (2018).
 121. Merkel, S. T., Pritchett, E. J. & Fong, B. H. Randomized benchmarking as convolution: Fourier analysis of gate dependent errors. *Quantum* **5**, 581 (2021).
 122. Wu, B. & Koh, D. E. Error-mitigated fermionic classical shadows on noisy quantum devices. *npj Quantum Inf.* **10**, 39 (2024).
 123. Brieger, R., Heinrich, M., Roth, I. & Kliesch, M. Stability of classical shadows under gate-dependent noise. *arXiv:2310.19947* <https://arxiv.org/abs/2310.19947> (2023).
 124. Helsen, J., Nezami, S., Reagor, M. & Walter, M. Matchgate benchmarking: Scalable benchmarking of a continuous family of many-qubit gates. *Quantum* **6**, 657 (2022).
 125. Valiant, L. G. Quantum computers that can be simulated classically in polynomial time. In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing*, 114–123 (2001).
 126. Knill, E. Fermionic linear optics and matchgates. *arXiv:quant-ph/0108033* <https://arxiv.org/abs/quant-ph/0108033> (2001).
 127. Terhal, B. M. & DiVincenzo, D. P. Classical simulation of noninteracting-fermion quantum circuits. *Phys. Rev. A* **65**, 032325 (2002).
 128. Bravyi, S. Lagrangian representation for fermionic linear optics. *Quantum Inf. Comput.* **5**, 216–238 (2005).
 129. DiVincenzo, D. P. & Terhal, B. M. Fermionic linear optics revisited. *Found. Phys.* **35**, 1967–1984 (2005).
 130. Jozsa, R. & Miyake, A. Matchgates and classical simulation of quantum circuits. *Proc. R. Soc. A* **464**, 3089–3106 (2008).
 131. Gambetta, J. M. et al. Characterization of addressability by simultaneous randomized benchmarking. *Phys. Rev. Lett.* **109**, 240504 (2012).
 132. McClean, J. R. et al. OpenFermion: the electronic structure package for quantum computers. *Quantum Sci. Technol.* **5**, 034014 (2020).

Acknowledgements

This work was supported by the National Science Foundation STAQ Project (PHY-1818914, PHY-2325080) and CHE-2037832. Support is also acknowledged from the U.S. Department of Energy, Office of Science National Quantum Information Science Research Center, Quantum Systems Accelerator. The authors thank the UNM Center for Advanced Research Computing, supported in part by the National Science Foundation, for providing the high-performance computing and large-scale storage resources used in this work.

Author contributions

Project design and conceptualization were envisioned by A.Z. and A.M. The analyses and numerical experiments were led by A.Z. and discussed with A.M. The manuscript was written by A.Z. and A.M.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41534-024-00854-5>.

Correspondence and requests for materials should be addressed to Andrew Zhao.

Reprints and permissions information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2024

Supplementary Information for “Group-theoretic error mitigation enabled by classical shadows and symmetries”

Andrew Zhao¹ and Akimasa Miyake¹

¹*Center for Quantum Information and Control, Department of Physics and Astronomy,
University of New Mexico, Albuquerque, New Mexico 87106, USA*

(Dated: May 14, 2024)

CONTENTS

| | |
|--------------------------|----|
| Supplementary Note 1 | 1 |
| Supplementary Note 2 | 4 |
| Supplementary Note 3 | 12 |
| Supplementary Note 4 | 15 |
| Supplementary Note 5 | 19 |
| Supplementary Note 6 | 21 |
| Supplementary Discussion | 26 |
| Supplementary References | 28 |

SUPPLEMENTARY NOTE 1

Error analysis. Here we provide the proof for Theorem 4 from the main text, restated below for convenience.

Theorem 4 (Restated from main text). Fix accuracy and confidence parameters $\epsilon, \delta \in (0, 1)$. Let O_1, \dots, O_L be a collection of observables, each supported on an irrep of $\mathcal{U} : G \rightarrow \text{U}(\mathcal{L}(\mathcal{H}))$ as $O_j \in V_\lambda$ for $\lambda \in R' \subseteq R_G$. Let $S_\lambda \in V_\lambda$ be a symmetry operator for each $\lambda \in R'$, for which the ideal values $s_\lambda = \text{tr}(S_\lambda \rho)$ of the target state ρ are known *a priori*. Suppose that each noisy unitary satisfies Assumptions 1, $\tilde{\mathcal{U}}_g = \mathcal{E}\mathcal{U}_g$, and define the quantities

$$F_{Z, R'}(\mathcal{E}) := \min_{\lambda \in R'} \frac{\text{tr}(\mathcal{E}\mathcal{M}_Z\Pi_\lambda)}{\text{tr}(\mathcal{M}_Z\Pi_\lambda)}, \quad (1)$$

$$\sigma^2 := \max_{1 \leq j \leq L, \lambda \in R'} \left\{ \text{Var}[\hat{o}_j], \text{Var} \left[\frac{\hat{s}_\lambda}{s_\lambda} \right] \right\}. \quad (2)$$

Then, a (noisy) classical shadow $\hat{\rho}(T)$ of size

$$T = \mathcal{O} \left(\frac{\log((L + |R'|)/\delta)}{F_{Z, R'}(\mathcal{E})^2 \epsilon^2} \sigma^2 \right) \quad (3)$$

can be used to construct error-mitigated estimates

$$\hat{o}_j^{\text{EM}}(T) := \frac{\text{tr}(O_j \hat{\rho}(T))}{\text{tr}(S_\lambda \hat{\rho}(T))/s_\lambda} \quad (4)$$

which obey

$$|\hat{o}_j^{\text{EM}}(T) - \text{tr}(O_j \rho)| \leq (\|O_j\|_\infty + 1)\epsilon + \mathcal{O}(\|O_j\|_\infty \epsilon^2) \quad (5)$$

for all $1 \leq j \leq L$, with success probability at least $1 - \delta$.

Proof. Let $\hat{\rho}_1, \dots, \hat{\rho}_T$ be the T noisy classical shadows. Construct the mean of these snapshots,

$$\hat{\rho}(T) = \frac{1}{T} \sum_{\ell=1}^T \hat{\rho}_\ell. \quad (6)$$

(It is straightforward to replace this by a median-of-means estimator if necessary.) In expectation we have $\mathbb{E}[\hat{\rho}(T)] = \tilde{\rho}$, where the effective noisy state can be described as

$$\tilde{\rho} = \mathcal{M}^{-1}(\tilde{\mathcal{M}}(\rho)). \quad (7)$$

Let $O \in V_\lambda$, with symmetry $s_\lambda = \text{tr}(S_\lambda \rho)$ in the same irrep. Define estimates of the noisy expectation values using $\hat{\rho}(T)$:

$$\bar{X} = \text{tr}(O \hat{\rho}(T)) \quad (8)$$

$$\bar{Y} = \text{tr}\left(\frac{S_\lambda}{s_\lambda} \hat{\rho}(T)\right). \quad (9)$$

In expectation, these random variables obey $\mathbb{E}[\bar{X}] = \text{tr}(O \tilde{\rho})$ and $\mathbb{E}[\bar{Y}] = \text{tr}(S_\lambda \tilde{\rho})/s_\lambda = \tilde{f}_\lambda/f_\lambda$. Therefore as established from the main text, we have

$$r(\mathbb{E}[\bar{X}], \mathbb{E}[\bar{Y}]) = \frac{\text{tr}(O \tilde{\rho})}{\text{tr}(S_\lambda \tilde{\rho})/s_\lambda} = \text{tr}(O \rho), \quad (10)$$

where we have defined the function $r(x, y) := x/y$. From a finite number of samples, however, we can only construct $\hat{\rho}^{\text{EM}}(T) := r(\bar{X}, \bar{Y})$, which is generally a biased estimator since $\mathbb{E}[r(\bar{X}, \bar{Y})] \neq r(\mathbb{E}[\bar{X}], \mathbb{E}[\bar{Y}])$.

To quantify the estimation error, we employ Taylor's remainder theorem: expanding $r(x, y)$ to first order about a point (x_0, y_0) , we have

$$r(x, y) = r(x_0, y_0) + \partial_x r(x_0, y_0)(x - x_0) + \partial_y r(x_0, y_0)(y - y_0) + h_1(x, y), \quad (11)$$

where the remainder term is

$$h_1(x, y) = \frac{1}{2!} [\partial_x^2 r(a, b)(x - x_0)^2 + \partial_y^2 r(a, b)(y - y_0)^2 + 2\partial_{xy} r(a, b)(x - x_0)(y - y_0)] \quad (12)$$

for some points $a \in [\min(x, x_0), \max(x, x_0)]$ and $b \in [\min(y, y_0), \max(y, y_0)]$. The relevant partial derivatives of $r(x, y)$ are enumerated below:

$$\partial_x r(x, y) = 1/y, \quad (13)$$

$$\partial_y r(x, y) = -x/y^2, \quad (14)$$

$$\partial_x^2 r(x, y) = 0, \quad (15)$$

$$\partial_{xy} r(x, y) = -1/y^2, \quad (16)$$

$$\partial_y^2 r(x, y) = 2x/y^3. \quad (17)$$

Suppose T is large enough such that (with high probability) the estimation error of all *noisy* observables are uniformly bounded by some $\tilde{\epsilon} \in (0, 1)$:

$$|\bar{X} - \mathbb{E}[\bar{X}]| \leq \tilde{\epsilon}, \quad (18)$$

$$|\bar{Y} - \mathbb{E}[\bar{Y}]| \leq \tilde{\epsilon}. \quad (19)$$

This is achieved by standard classical shadow arguments, which we will elaborate on later. For now, assuming these error bounds hold, we rearrange Eq. (11), set $(x, y) = (\bar{X}, \bar{Y})$ and $(x_0, y_0) = (\mathbb{E}[\bar{X}], \mathbb{E}[\bar{Y}])$, and apply a triangle inequality to obtain

$$|r(\bar{X}, \bar{Y}) - r(\mathbb{E}[\bar{X}], \mathbb{E}[\bar{Y}])| \leq \frac{1}{|\mathbb{E}[\bar{Y}]|} \tilde{\epsilon} + \frac{|\mathbb{E}[\bar{X}]|}{|\mathbb{E}[\bar{Y}]|^2} \tilde{\epsilon} + |h_1(\bar{X}, \bar{Y})|. \quad (20)$$

To proceed with this error bound, we make the following observations. First, note that

$$\mathbb{E}[\bar{Y}] = \frac{\tilde{f}_\lambda}{f_\lambda} = \frac{\text{tr}(\mathcal{M}_Z \mathcal{E} \Pi_\lambda)}{\text{tr}(\mathcal{M}_Z \Pi_\lambda)} \in [0, 1], \quad (21)$$

which we will denote by ξ_λ . We assume that that noise channel \mathcal{E} is such that $\xi_\lambda > 0$, as otherwise the quantity $r(\mathbb{E}[\bar{X}], \mathbb{E}[\bar{Y}])$ diverges. Next, because $\mathbb{E}[\bar{X}]/\mathbb{E}[\bar{Y}] = \text{tr}(O\rho)$, we have the bound

$$\left| \frac{\mathbb{E}[\bar{X}]}{\mathbb{E}[\bar{Y}]} \right| \leq \|O\|_\infty. \quad (22)$$

Thus Eq. (20) becomes

$$|r(\bar{X}, \bar{Y}) - r(\mathbb{E}[\bar{X}], \mathbb{E}[\bar{Y}])| \leq \frac{1}{\xi_\lambda} \tilde{\epsilon} + \frac{\|O\|_\infty}{\xi_\lambda} \tilde{\epsilon} + |h_1(\bar{X}, \bar{Y})|. \quad (23)$$

We can bound the remainder term $|h_1(\bar{X}, \bar{Y})|$ as follows. Applying a triangle inequality to Eq. (12) yields

$$|h_1(\bar{X}, \bar{Y})| \leq \frac{|a|}{|b|^3} \tilde{\epsilon}^2 + \frac{1}{b^2} \tilde{\epsilon}^2. \quad (24)$$

Taylor's remainder theorem tells us that the value of a (resp., b) lies between \bar{X} and $\mathbb{E}[\bar{X}]$ (resp., \bar{Y} and $\mathbb{E}[\bar{Y}]$), which we know are at most $\tilde{\epsilon}$ apart. We can therefore bound

$$\begin{aligned} |a| &\leq \max\{|\mathbb{E}[\bar{X}]|, |\bar{X}|\} \\ &\leq \max\{|\mathbb{E}[\bar{X}]|, |\mathbb{E}[\bar{X}] + \tilde{\epsilon}|, |\mathbb{E}[\bar{X}] - \tilde{\epsilon}|\} \\ &\leq |\mathbb{E}[\bar{X}]| + \tilde{\epsilon} \\ &\leq \xi_\lambda \|O\|_\infty + \tilde{\epsilon}. \end{aligned} \quad (25)$$

Similarly for b , using the fact that $\mathbb{E}[\bar{Y}] > 0$,

$$\begin{aligned} |b| &\geq \min\{|\mathbb{E}[\bar{Y}]|, |\bar{Y}|\} \\ &\geq \min\{|\mathbb{E}[\bar{Y}]|, |\mathbb{E}[\bar{Y}] + \tilde{\epsilon}|, |\mathbb{E}[\bar{Y}] - \tilde{\epsilon}|\} \\ &= \min\{\xi_\lambda, |\xi_\lambda - \tilde{\epsilon}|\}. \end{aligned} \quad (26)$$

If $\tilde{\epsilon} < \xi_\lambda$, then $|b| \geq \xi_\lambda - \tilde{\epsilon} > 0$ always holds. We will see later that this condition is always justified; for now, we will just suppose that this lower bound on $|b|$ holds. Then the remainder obeys

$$|h_1(\bar{X}, \bar{Y})| \leq \frac{1}{(\xi_\lambda - \tilde{\epsilon})^2} \left(\frac{\xi_\lambda \|O\|_\infty + \tilde{\epsilon}}{\xi_\lambda - \tilde{\epsilon}} + 1 \right) \tilde{\epsilon}^2. \quad (27)$$

Combining Eqs. (23) and (27), we arrive at

$$|\hat{\delta}^{\text{EM}}(T) - \text{tr}(O\rho)| \leq \frac{1}{\xi_\lambda} (\|O\|_\infty + 1) \tilde{\epsilon} + \frac{1}{(\xi_\lambda - \tilde{\epsilon})^2} \left(\frac{\xi_\lambda \|O\|_\infty + \tilde{\epsilon}}{\xi_\lambda - \tilde{\epsilon}} + 1 \right) \tilde{\epsilon}^2. \quad (28)$$

In order to bound this error by $\mathcal{O}(\|O\|_\infty \epsilon)$ for some desired $\epsilon \in (0, 1)$, we can choose $\tilde{\epsilon} = \xi_\lambda \epsilon$, yielding

$$|\hat{\delta}^{\text{EM}}(T) - \text{tr}(O\rho)| \leq (\|O\|_\infty + 1) \epsilon + \frac{1}{(1 - \epsilon)^2} \left(\frac{\|O\|_\infty + \epsilon}{1 - \epsilon} + 1 \right) \epsilon^2. \quad (29)$$

Thus, by demanding $\epsilon < 1$ we ensure that the required technical condition $\tilde{\epsilon} < \xi_\lambda$ is met. Now we need to verify that the remainder term is bounded by $\mathcal{O}(\|O\|_\infty \epsilon^2)$, so that the $\mathcal{O}(\|O\|_\infty \epsilon)$ term dominates asymptotically as $\epsilon \rightarrow 0$. Indeed, as long as ϵ is bounded away from 1 then

$$\frac{1}{(1 - \epsilon)^2} \left(\frac{\|O\|_\infty + \epsilon}{1 - \epsilon} + 1 \right) = \mathcal{O}(\|O\|_\infty). \quad (30)$$

Finally, we analyze the sample complexity required to achieve the error bound of Eq. (29). In Eqs. (18) and (19) we required that the number of samples T be such that the shot noise of \bar{X} and \bar{Y} are at most $\tilde{\epsilon}$. These random variables correspond to the observables $\{O_j\}_{j=1}^L \cup \{S_\lambda/s_\lambda\}_{\lambda \in R'}$. Standard classical-shadows theory informs us that

$$T = \mathcal{O} \left(\frac{\log((L + |R'|)/\delta)}{\tilde{\epsilon}^2} \max_{\substack{1 \leq j \leq L \\ \lambda \in R'}} \left\{ \text{Var}[\hat{o}_j], \text{Var} \left[\frac{\hat{s}_\lambda}{s_\lambda} \right] \right\} \right) \quad (31)$$

suffices to accomplish this task (with probability at least $1 - \delta$) [1]. Then, setting $\tilde{\epsilon} = \min_{\lambda \in R'} \xi_\lambda \epsilon$ ensures that $\tilde{\epsilon}$ is small enough for Eq. (29) to apply to all target observables. \square

SUPPLEMENTARY NOTE 2

Subsystem-symmetrized Pauli shadows. Here we prove the properties of the subsystem-symmetrized Pauli shadows introduced in the main text. In “Irreducible representations” we identify the irreps, and in “Variance of symmetry operators” we bound the variance of observables under this protocol, particularly the symmetry operators S_k obtained from $M = \sum_{i \in [n]} Z_i$.

Irreducible representations. Recall that the subsystem-symmetrized local Clifford group is the direct product

$$\text{Cl}(1)_{\text{Sym}}^{\otimes n} := \text{Sym}(n) \times \text{Cl}(1)^{\otimes n}, \quad (32)$$

where $\pi \in \text{Sym}(n)$ acts on $(\mathbb{C}^2)^{\otimes n}$ as

$$S_\pi |b\rangle = |\pi^{-1}(b)\rangle. \quad (33)$$

For shorthand, we write $\pi(b)$ for the n -bit string $b_{\pi(0)} \cdots b_{\pi(n-1)}$. It is clear that the adjoint representation $\mathcal{U}_{(\pi, C)}$ block diagonalizes into subspaces spanned by k -local Pauli operators:

$$V_k := \text{span}\{P \in \mathcal{P}(n) : |P| = k\}. \quad (34)$$

This can be seen from the fact that neither single-qubit nor SWAP gates can change the operator locality; however, SWAP gates *can* map between equally sized subsystems on which the operator nontrivially acts. What remains is to show that each of these subspaces is irreducible.

First, we define the twirling map.

Definition 1. Let $\phi : G \rightarrow \text{U}(V)$ be a unitary representation of a compact group G on a vector space V , and let $\Phi : G \rightarrow \text{U}(\mathcal{L}(V))$ be its adjoint action, i.e., $\Phi_g(\cdot) = \phi_g(\cdot)\phi_g^\dagger$. The t -fold twirl by Φ is defined as

$$\mathcal{T}_{t, \Phi} := \mathbb{E}_{g \sim G} \Phi_g^{\otimes t}, \quad (35)$$

which is a linear map on $\mathcal{L}(V)^{\otimes t}$.

Twirls have a number of convenient properties, mostly arising from the fact that Φ is a group homomorphism. For example, they are G -invariant from the left and right:

$$\Phi_h^{\otimes t} \circ \mathcal{T}_{t, \Phi} = \mathcal{T}_{t, \Phi} = \mathcal{T}_{t, \Phi} \circ \Phi_h^{\otimes t} \quad (36)$$

for all $h \in G$. This furthermore implies that they are in fact projectors:

$$\mathcal{T}_{t, \Phi}^2 = \mathcal{T}_{t, \Phi}. \quad (37)$$

The study of twirls also allows us to determine the irreducible representations of a group. This can be seen by the following well-known result for multiplicity-free groups, which for completeness we provide a self-contained proof of at the end of this subsection.

Proposition 2. Let G , V , ϕ , and Φ be as in Definition 1. For any $X \in \mathcal{L}(V)$, the 1-twirl of X by Φ takes the form

$$\mathcal{T}_{1, \Phi}(X) = \sum_{\lambda \in R_G} \frac{\text{tr}(X \Pi_\lambda)}{\text{tr}(\Pi_\lambda)} \Pi_\lambda \quad (38)$$

if and only if ϕ decomposes irreducibly as $V = \bigoplus_{\lambda \in R_G} V_\lambda$, where Π_λ is the orthogonal projector onto V_λ .

Our strategy for determining the irreps of $G = \text{Cl}(1)_{\text{Sym}}^{\otimes n}$ is therefore to directly compute $\mathcal{T}_{1,\Phi}$, from which we can infer the irreps from its block-diagonal structure. To use Proposition 2, we will take ϕ as the unitary channel \mathcal{U} , so that $V = \mathcal{L}(\mathcal{H})$ and $\Phi(\cdot) = \mathcal{U}(\cdot)\mathcal{U}^\dagger$ (note that this is a superchannel). For technical reasons, it will be easier to first compute $\mathcal{T}_{2,\mathcal{U}}$, from which the desired twirl $\mathcal{T}_{1,\Phi}$ can be evaluated. The relation between these two twirls is given by the following lemma.

Lemma 3. *Let $\mathcal{U} : G \rightarrow \text{U}(\mathcal{L}(\mathcal{H}))$ be a unitary representation and $\Phi : G \rightarrow \text{U}(\mathcal{L}(\mathcal{L}(\mathcal{H})))$ its adjoint representation, i.e., $\Phi_g(\mathcal{A}) = \mathcal{U}_g^\dagger \mathcal{A} \mathcal{U}_g$ for any superoperator \mathcal{A} . The 1-twirl by Φ can be computed from the 2-twirl by \mathcal{U} as*

$$\mathcal{T}_{1,\Phi}(\mathcal{A})(X) = \text{tr}_2[\mathcal{T}_{2,\mathcal{U}}(\mathcal{A})(\mathbb{I} \otimes X)], \quad (39)$$

for all $\mathcal{A} \in \mathcal{L}(\mathcal{L}(\mathcal{H}))$ and $X \in \mathcal{L}(\mathcal{H})$. Here, the domain of $\mathcal{T}_{2,\mathcal{U}}$ is understood with respect to the isomorphism $\mathcal{L}(\mathcal{L}(\mathcal{H})) \cong \mathcal{L}(\mathcal{H})^{\otimes 2}$, given by

$$|A\rangle\langle\langle B| \cong A \otimes B^\dagger. \quad (40)$$

Proof. Write $\mathcal{A} = \sum_{i,j \in [d^2]} \mathcal{A}_{ij} |B_i\rangle\langle\langle B_j|$, where $\mathcal{A}_{ij} \in \mathbb{C}$ and $\{|B_i\rangle\}_{i \in [d^2]}$ is an orthonormal operator basis. By a direct calculation:

$$\begin{aligned} \mathcal{T}_{1,\Phi}(\mathcal{A})(X) &= \mathbb{E}_{g \sim G} \Phi_g(\mathcal{A})(X) \\ &= \mathbb{E}_{g \sim G} \mathcal{U}_g^\dagger \mathcal{A} \mathcal{U}_g(X) = \mathbb{E}_{g \sim G} \mathcal{U}_g^\dagger \mathcal{A} (U_g X U_g^\dagger) \\ &= \mathbb{E}_{g \sim G} \mathcal{U}_g^\dagger \sum_{i,j \in [d^2]} \mathcal{A}_{ij} |B_i\rangle\langle\langle B_j| U_g X U_g^\dagger \\ &= \mathbb{E}_{g \sim G} \sum_{i,j \in [d^2]} \mathcal{A}_{ij} U_g^\dagger B_i U_g \text{tr}(B_j^\dagger U_g X U_g^\dagger) \\ &= \mathbb{E}_{g \sim G} \text{tr}_2 \left[\sum_{i,j \in [d^2]} \mathcal{A}_{ij} (U_g^\dagger B_i U_g) \otimes (U_g^\dagger B_j^\dagger U_g X) \right] \\ &= \text{tr}_2 \left[\sum_{i,j \in [d^2]} \mathcal{A}_{ij} \mathcal{T}_{2,\mathcal{U}}(B_i \otimes B_j^\dagger)(\mathbb{I} \otimes X) \right] \\ &= \text{tr}_2 \left[\mathcal{T}_{2,\mathcal{U}} \left(\sum_{i,j \in [d^2]} \mathcal{A}_{ij} |B_i\rangle\langle\langle B_j| \right) (\mathbb{I} \otimes X) \right] \\ &= \text{tr}_2[\mathcal{T}_{2,\mathcal{U}}(\mathcal{A})(\mathbb{I} \otimes X)]. \end{aligned} \quad (41)$$

□

Before we can compute $\mathcal{T}_{2,\mathcal{U}}$ for the subsystem-symmetrized local Clifford group, we will need a small result about the group orbit of a k -local Pauli operator P under the action of \mathcal{U} . The orbit is defined as

$$G \cdot P := \{\mathcal{U}_g(P) \mid g \in G\}. \quad (42)$$

This will help us determine how the twirl acts on Pauli operators, which as a basis is used to compute the matrix elements of $\mathcal{T}_{2,\mathcal{U}}$. To this end, we define an orthonormal basis of k -local Pauli operators,

$$\mathcal{B}_k := \{P \in \mathcal{P}(n)/\sqrt{d} : |P| = k\}, \quad (43)$$

which contains $|\mathcal{B}_k| = 3^k \binom{n}{k}$ elements.

Lemma 4. *Let $G = \text{Cl}(1)_{\text{Sym}}^{\otimes n}$. The orbit $G \cdot P$ of any $P \in \mathcal{B}_k$ is equal to $\pm \mathcal{B}_k$, i.e., the set of all signed k -local Pauli operators.*

Proof. Let the nontrivial support of P be $I \subseteq [n]$, $|I| = k$. For each (normalized) Pauli matrix acting on subsystem I , its orbit by all single-qubit Clifford gates is $\pm\{X, Y, Z\}/\sqrt{2}$. Meanwhile, the trivial factors $\mathbb{I}/\sqrt{2}$ acting on $[n] \setminus I$ are invariant to any unitary transformation. Therefore $\text{Cl}(1)^{\otimes n} \cdot P$ is the set of all normalized Pauli operators acting nontrivially only on the qubits in I (with both signs ± 1).

Then, conjugation by S_π for arbitrary $\pi \in \text{Sym}(n)$ permutes the k nontrivial factors of P among the n qubits. The orbit over all permutations yields all possible $\binom{n}{k}$ supports. Taking the direct product of both these Clifford- and symmetric-group actions therefore yields all k -local Pauli operators, with prefactors ± 1 . □

We are now ready to compute the 2-fold twirl by \mathcal{U} . We comment that the high-level proof structure of this lemma is inspired by that of Ref. [2, Section IV A 1].

Lemma 5. *Let $\mathcal{U} : G \rightarrow \text{U}(\mathcal{L}(\mathcal{H}))$ be the unitary representation of $G = \text{Cl}(1)_{\text{Sym}}^{\otimes n}$, defined by $\mathcal{U}_{(\pi,C)}(\rho) = S_\pi C \rho C^\dagger S_\pi^\dagger$. Its 2-fold twirl is the projector*

$$\mathcal{T}_{2,\mathcal{U}} = \sum_{k=0}^n |\Sigma_k^{(2)}\rangle\rangle\langle\langle \Sigma_k^{(2)}|, \quad (44)$$

where $|\Sigma_k^{(2)}\rangle\rangle \in \mathcal{L}(\mathcal{H})^{\otimes 2}$ is defined as

$$|\Sigma_k^{(2)}\rangle\rangle = \frac{1}{\sqrt{3^k \binom{n}{k}}} \sum_{P \in \mathcal{B}_k} |P\rangle\rangle|P\rangle\rangle. \quad (45)$$

Proof. First, we will establish that for any two basis Pauli operators $P \neq Q$, we have $\mathcal{T}_{2,\mathcal{U}}|P\rangle\rangle|Q\rangle\rangle = 0$. Thus we only need to consider basis elements of $\mathcal{L}(\mathcal{H})^{\otimes 2}$ of the form $|P\rangle\rangle|P\rangle\rangle$. Next, we will show that $\mathcal{T}_{2,\mathcal{U}}|P\rangle\rangle|P\rangle\rangle = \mathcal{T}_{2,\mathcal{U}}|P'\rangle\rangle|P'\rangle\rangle$ whenever $|P| = |P'|$. Finally, using these two properties we can derive Eq. (44).

Fix the basis of Pauli operators such that $P, Q \in \bigcup_{0 \leq k \leq n} \mathcal{B}_k$. If $P \neq Q$, then there exists at least one qubit $i \in [n]$ on which P and Q act as a different Pauli matrix. Hence there always exists some $W \in \mathcal{P}(1)$ which anticommutes with one and commutes with the other, e.g., $\mathcal{W}_i|P\rangle\rangle = W_i P W_i^\dagger = -P$ and $\mathcal{W}_i|Q\rangle\rangle = Q$. Note that \mathcal{W}_i is equal to $\mathcal{U}_{(e,W_i)}$ where $e \in \text{Sym}(n)$ is the identity permutation. Thus using the property that $\mathcal{T}_{2,\mathcal{U}} \circ \mathcal{W}_i^{\otimes 2} = \mathcal{T}_{2,\mathcal{U}}$, we have

$$\begin{aligned} \mathcal{T}_{2,\mathcal{U}}|P\rangle\rangle|Q\rangle\rangle &= \mathcal{T}_{2,\mathcal{U}}\mathcal{W}_i^{\otimes 2}|P\rangle\rangle|Q\rangle\rangle \\ &= -\mathcal{T}_{2,\mathcal{U}}|P\rangle\rangle|Q\rangle\rangle, \end{aligned} \quad (46)$$

implying that $\mathcal{T}_{2,\mathcal{U}}|P\rangle\rangle|Q\rangle\rangle = 0$.

Now let P, P' be k -local Pauli operators for any k . If they act nontrivially on different subsets $I, I' \subseteq [n]$ of qubits, then let $\pi \in \text{Sym}(n)$ be a permutation that maps I to I' . Given this permutation, if they act as different Pauli matrices on their new shared support I' , then furthermore let $C_i \in \text{Cl}(1)$ for $i \in I'$ be Clifford gates that map each one to the other. Writing $C = \bigotimes_{i \in I'} C_i \otimes \mathbb{I}^{\otimes (n-k)}$, this transformation acts as $\mathcal{U}_{(\pi,C)}^{\otimes 2}|P\rangle\rangle|P\rangle\rangle = |P'\rangle\rangle|P'\rangle\rangle$, which implies that

$$\begin{aligned} \mathcal{T}_{2,\mathcal{U}}|P\rangle\rangle|P\rangle\rangle &= \mathcal{T}_{2,\mathcal{U}}\mathcal{U}_{(\pi,C)}^{\otimes 2}|P\rangle\rangle|P\rangle\rangle \\ &= \mathcal{T}_{2,\mathcal{U}}|P'\rangle\rangle|P'\rangle\rangle. \end{aligned} \quad (47)$$

We are now ready to derive Eq. (44). As established by Eq. (46), we only need to expand the 2-fold twirl in the basis of $|P\rangle\rangle|P\rangle\rangle$:

$$\begin{aligned} \mathcal{T}_{2,\mathcal{U}} &= \sum_{P, P' \in \mathcal{P}(n)} \langle\langle P|\langle\langle P|\mathcal{T}_{2,\mathcal{U}}|P'\rangle\rangle|P'\rangle\rangle|P\rangle\rangle|P\rangle\rangle\langle\langle P'|\langle\langle P'| \\ &= \sum_{k=0}^n \sum_{P, P' \in \mathcal{B}_k} \langle\langle P|\langle\langle P|\mathcal{T}_{2,\mathcal{U}}|P'\rangle\rangle|P'\rangle\rangle|P\rangle\rangle|P\rangle\rangle\langle\langle P'|\langle\langle P'|, \end{aligned} \quad (48)$$

where the second simplification is due to the fact that $\mathcal{U}_{(\pi,C)}$ preserves Pauli locality, hence $\langle\langle P|\langle\langle P|\mathcal{T}_{2,\mathcal{U}}|P'\rangle\rangle|P'\rangle\rangle = 0$ whenever $|P| \neq |P'|$. Now we invoke Eq. (47), which implies that $\langle\langle P|\langle\langle P|\mathcal{T}_{2,\mathcal{U}}|P'\rangle\rangle|P'\rangle\rangle = c'_k$ for all $P, P' \in \mathcal{B}_k$ (i.e., the matrix element does not depend on the particular choice of P, P'). Hence

$$\begin{aligned} \mathcal{T}_{2,\mathcal{U}} &= \sum_{k=0}^n c'_k \sum_{P, P' \in \mathcal{B}_k} |P\rangle\rangle|P\rangle\rangle\langle\langle P'|\langle\langle P'| \\ &= \sum_{k=0}^n c_k |\Sigma_k^{(2)}\rangle\rangle\langle\langle \Sigma_k^{(2)}|, \end{aligned} \quad (49)$$

where we have rescaled $c_k = c'_k 3^k \binom{n}{k}$ to account for the normalization of $|\Sigma_k^{(2)}\rangle\rangle$.

Finally, we show that all $c_k = 1$ by proving that $\mathcal{T}_{2,\mathcal{U}}|\Sigma_k^{(2)}\rangle\rangle = |\Sigma_k^{(2)}\rangle\rangle$. Expand the expression:

$$\mathcal{T}_{2,\mathcal{U}}|\Sigma_k^{(2)}\rangle\rangle = \frac{1}{|G|} \sum_{g \in G} \frac{1}{\sqrt{|\mathcal{B}_k|}} \sum_{P \in \mathcal{B}_k} \mathcal{U}_g|P\rangle\rangle \otimes \mathcal{U}_g|P\rangle\rangle. \quad (50)$$

We first compute the average over the group for some fixed P . By Lemma 4, we know that the orbit $G \cdot P = \{\mathcal{U}_g|P\rangle\rangle \mid g \in G\} = \pm\mathcal{B}_k$. Thus

$$\sum_{g \in G} \mathcal{U}_g|P\rangle\rangle \otimes \mathcal{U}_g|P\rangle\rangle = 2 \frac{|G|}{|G \cdot P|} \sum_{Q \in \mathcal{B}_k} (\pm 1)^2 |Q\rangle\rangle |Q\rangle\rangle, \quad (51)$$

where the factor of 2 is due to the fact that for each $Q \in \mathcal{B}_k$, both $\pm Q \in G \cdot P$, and the factor of $|G|/|G \cdot P|$ takes care of double counting when summing over all elements of G . Noting that $|G \cdot P| = 2|\mathcal{B}_k|$, we can plug this result into Eq. (50) to find that

$$\begin{aligned} \mathcal{T}_{2,\mathcal{U}}|\Sigma_k^{(2)}\rangle\rangle &= \frac{1}{|\mathcal{B}_k|^{3/2}} \sum_{P \in \mathcal{B}_k} \sum_{Q \in \mathcal{B}_k} |Q\rangle\rangle |Q\rangle\rangle \\ &= \frac{1}{\sqrt{|\mathcal{B}_k|}} \sum_{Q \in \mathcal{B}_k} |Q\rangle\rangle |Q\rangle\rangle \\ &= |\Sigma_k^{(2)}\rangle\rangle, \end{aligned} \quad (52)$$

as desired. \square

We are now ready to prove the main result of this section: the irreps of $\text{Cl}(1)_{\text{Sym}}^{\otimes n}$ are labeled by the Pauli weights $k \in \{0, 1, \dots, n\}$. The proof structure is as follows: from the expression for $\mathcal{T}_{2,\mathcal{U}}$ from Lemma 5, we can compute $\mathcal{T}_{1,\Phi}$ by using Lemma 3. Then by examining $\mathcal{T}_{1,\Phi}$, we use Proposition 2 to infer the irreps.

Theorem 6. *The representation $\mathcal{U} : \text{Cl}(1)_{\text{Sym}}^{\otimes n} \rightarrow \text{U}(\mathcal{L}(\mathcal{H}))$, defined by $\mathcal{U}_{(\pi,C)}(\rho) = S_\pi C \rho C^\dagger S_\pi^\dagger$, decomposes into the irreps*

$$V_k = \text{span}(\mathcal{B}_k), \quad k \in \{0, 1, \dots, n\}. \quad (53)$$

Proof. From Lemma 5, we have

$$\mathcal{T}_{2,\mathcal{U}} = \sum_{k=0}^n |\Sigma_k^{(2)}\rangle\rangle\langle\langle \Sigma_k^{(2)}|, \quad (54)$$

where $|\Sigma_k^{(2)}\rangle\rangle$ is defined in Eq. (45). Using Lemma 3, we compute $\mathcal{T}_{1,\Phi}(\mathcal{A})$ by evaluating $\mathcal{T}_{2,\mathcal{U}}(\mathcal{A})$ for arbitrary superoperators \mathcal{A} . Let us express \mathcal{A} in the Pauli basis:

$$\mathcal{A} = \sum_{k,\ell=0}^n \sum_{P \in \mathcal{B}_k} \sum_{Q \in \mathcal{B}_\ell} \mathcal{A}_{PQ} |P\rangle\rangle\langle\langle Q|. \quad (55)$$

Recall from Eq. (40) that in order to evaluate $\mathcal{T}_{2,\mathcal{U}}(\mathcal{A})$, we need $\mathcal{T}_{2,\mathcal{U}}(Q \otimes P)$ for every P, Q . But because $\mathcal{T}_{2,\mathcal{U}}$ projects onto symmetrized basis elements $P \otimes P$, we only have to consider the case where $P = Q$:

$$\begin{aligned} \mathcal{T}_{2,\mathcal{U}}|P \otimes P\rangle\rangle &= |\Sigma_k^{(2)}\rangle\rangle\langle\langle \Sigma_k^{(2)}|P \otimes P\rangle\rangle \\ &= |\Sigma_k^{(2)}\rangle\rangle \frac{1}{\sqrt{|\mathcal{B}_k|}} \sum_{P' \in \mathcal{B}_k} \langle\langle P'|P\rangle\rangle \langle\langle P'|P\rangle\rangle \\ &= \frac{1}{\sqrt{|\mathcal{B}_k|}} |\Sigma_k^{(2)}\rangle\rangle, \end{aligned} \quad (56)$$

where $|\mathcal{B}_k| = 3^k \binom{n}{k}$.

Inserting this result into Eq. (39) yields

$$\begin{aligned}
\mathcal{T}_{1,\Phi}(\mathcal{A})(X) &= \text{tr}_1[\mathcal{T}_{2,\mathcal{U}}(\mathcal{A})(X \otimes \mathbb{I})] \\
&= \text{tr}_1 \left[\sum_{k=0}^n \sum_{P \in \mathcal{B}_k} \mathcal{A}_{PP} \frac{1}{\sqrt{|\mathcal{B}_k|}} \Sigma_k^{(2)}(X \otimes \mathbb{I}) \right] \\
&= \text{tr}_1 \left[\sum_{k=0}^n \sum_{P \in \mathcal{B}_k} \mathcal{A}_{PP} \frac{1}{\sqrt{|\mathcal{B}_k|}} \right. \\
&\quad \left. \times \frac{1}{\sqrt{|\mathcal{B}_k|}} \sum_{P' \in \mathcal{B}_k} (P' \otimes P')(X \otimes \mathbb{I}) \right] \\
&= \sum_{k=0}^n \frac{\sum_{P \in \mathcal{B}_k} \mathcal{A}_{PP}}{|\mathcal{B}_k|} \sum_{P' \in \mathcal{B}_k} \text{tr}(P'X)P'.
\end{aligned} \tag{57}$$

We make a number of observations here. First, note that $\sum_{P \in \mathcal{B}_k} \mathcal{A}_{PP} = \text{tr}(\mathcal{A}\Pi_k)$, where $\Pi_k = \sum_{P \in \mathcal{B}_k} |P\rangle\rangle\langle\langle P|$. Also, $|\mathcal{B}_k| = \text{tr}(\Pi_k)$. Finally, the sum over P' can be represented as

$$\begin{aligned}
\sum_{P' \in \mathcal{B}_k} \text{tr}(P'X)P' &= \sum_{P' \in \mathcal{B}_k} |P'\rangle\rangle\langle\langle P'|X\rangle\rangle \\
&= \Pi_k |X\rangle\rangle.
\end{aligned} \tag{58}$$

Because this holds for all $X \in \mathcal{L}(\mathcal{H})$, we can say that

$$\mathcal{T}_{1,\Phi}(\mathcal{A}) = \sum_{k=0}^n \frac{\text{tr}(\mathcal{A}\Pi_k)}{\text{tr}(\Pi_k)} \Pi_k. \tag{59}$$

By Proposition 2, we know that the twirl has this expression if and only if the irreducible subspaces of \mathcal{U} are $V_k = \text{span}(\mathcal{B}_k)$. \square

Finally, we close this subsection with the deferred proof of the well-known result Proposition 2, for completeness.

Proof (of Proposition 2). For the forward direction, suppose $\phi = \bigoplus_{\lambda \in R_G} \phi^{(\lambda)}$ where each $\phi^{(\lambda)} : G \rightarrow \text{U}(V_\lambda)$ is irreducible. (This is guaranteed by Maschke's theorem, and generalizes to the Peter–Weyl theorem for compact groups [3].) Because $\mathcal{T}_{1,\Phi}(X)$ commutes with all ϕ_g , they are simultaneously block diagonal, so $\mathcal{T}_{1,\Phi}(X) = \bigoplus_{\lambda \in R_G} \mathcal{T}_{1,\Phi^{(\lambda)}}(X)$ where $\Phi^{(\lambda)}(\cdot) = \phi^{(\lambda)}(\cdot)\phi^{(\lambda)\dagger}$. Because $\phi^{(\lambda)}$ is irreducible, by Schur's lemma $\mathcal{T}_{1,\Phi^{(\lambda)}}(X)$ must be a multiple of the identity on V_λ . Therefore

$$\mathcal{T}_{1,\Phi}(X) = \bigoplus_{\lambda \in R_G} c_\lambda(X) \mathbb{I}_{V_\lambda} = \sum_{\lambda \in R_G} c_\lambda(X) \Pi_\lambda. \tag{60}$$

From the orthogonality of projectors $\Pi_\lambda \Pi_{\lambda'} = \delta_{\lambda\lambda'} \Pi_\lambda$, the scalar $c_\lambda(X)$ is determined by

$$\begin{aligned}
c_\lambda(X) \text{tr}(\Pi_\lambda) &= \text{tr}(\Pi_\lambda \mathcal{T}_{1,\Phi}(X)) \\
&= \mathbb{E}_{g \sim G} \text{tr}(\phi_g^\dagger \Pi_\lambda \phi_g X) \\
&= \mathbb{E}_{g \sim G} \text{tr}(\Pi_\lambda X) = \text{tr}(\Pi_\lambda X).
\end{aligned} \tag{61}$$

For the reverse direction, suppose the twirl takes the form

$$\mathcal{T}_{1,\Phi}(X) = \bigoplus_{\lambda} \frac{\text{tr}(X \Pi_\lambda)}{\text{tr}(\Pi_\lambda)} \mathbb{I}_{V_\lambda}, \tag{62}$$

where we denote each block by $\mathcal{T}_{1,\Phi^{(\lambda)}}(X)$. Again because $\mathcal{T}_{1,\Phi}(X)$ and ϕ_g commute, the matrix ϕ_g is block diagonal in the subspaces V_λ for all $g \in G$. We need to show that each block $\phi^{(\lambda)}$ is irreducible.

Recall that $\phi^{(\lambda)}$ is irreducible if the only subspaces $W \subseteq V_\lambda$ for which $\phi_G^{(\lambda)}(W) \subseteq W$ are $W = \{0\}$ or $W = V_\lambda$. Indeed, let $W \subseteq V_\lambda$ be a subspace such that for any $|v\rangle \in W$ and $g \in G$, $\phi_g^{(\lambda)}|v\rangle \in W$. Suppose there exists a vector $|x\rangle \in V_\lambda$ that is orthogonal to W and set $X = |x\rangle\langle x|$. Then

$$\mathcal{T}_{1,\Phi^{(\lambda)}}(|x\rangle\langle x|)|v\rangle = \mathbb{E}_{g \sim G} \phi_g^{(\lambda)}|x\rangle\langle x|(\phi_g^{(\lambda)})^\dagger|v\rangle = 0 \quad (63)$$

because all $(\phi_g^{(\lambda)})^\dagger|v\rangle \in W$. However, from $\mathcal{T}_{1,\Phi^{(\lambda)}}(X) = c_\lambda(X)\mathbb{I}_{V_\lambda}$ we see that also

$$\mathcal{T}_{1,\Phi^{(\lambda)}}(|x\rangle\langle x|)_\lambda|v\rangle = \frac{\langle x|x\rangle}{\dim V_\lambda}|v\rangle. \quad (64)$$

Supposing $|x\rangle \neq 0$, we see that $|v\rangle = 0$ is the only possible element of W to satisfy Eqs. (63) and (64) simultaneously. Hence $W = \{0\}$. Otherwise, $|x\rangle = 0$ is the only element of V_λ orthogonal to W , implying that there is in fact no nontrivial subspace orthogonal to W . Thus $W = V_\lambda$ in this case. \square

Variance of symmetry operators. In this section, we analyze the variance associated with the symmetry operators,

$$\frac{S_1}{s_1} = \frac{1}{m} \sum_{i \in [n]} Z_i, \quad (65)$$

$$\frac{S_2}{s_2} = \frac{2}{m^2 - n} \sum_{i < j} Z_i Z_j. \quad (66)$$

Because our error analysis of symmetry-adjusted classical shadows bootstraps from the variance of unmitigated estimation, we only need to compute quantities related to the noiseless protocol. In this case, the subsystem-symmetrized local Clifford group yields the same channel and variances as the standard local Clifford group because the random permutations have no effect on the twirling on computational basis states.

Specifically, using the two- and three-fold twirls we can express

$$\mathcal{M}(\rho) = \text{tr}_1 \left[\sum_{b \in \{0,1\}^n} \mathcal{T}_{2,G}(|b\rangle\langle b|^{\otimes 2})(\rho \otimes \mathbb{I}) \right] \quad (67)$$

and

$$\begin{aligned} \text{Var}_\rho[\hat{\theta}] &= \text{tr} \left[\sum_{b \in \{0,1\}^n} \mathcal{T}_{3,G}(|b\rangle\langle b|^{\otimes 3})(\rho \otimes \mathcal{M}^{-1}(O)^{\otimes 2}) \right] \\ &\quad - \text{tr}(O\rho)^2, \end{aligned} \quad (68)$$

where the t -fold twirl by $U : G \rightarrow \text{U}(\mathcal{H})$ is defined as $\mathcal{T}_{t,G} := \mathbb{E}_{g \sim G} U_g^{\otimes t}$. These expressions are the same whether we take $G = \text{Cl}(1)_{\text{Sym}}^{\otimes n}$ or $\text{Cl}(1)^{\otimes n}$, due to the following equivalence:

$$\begin{aligned} \sum_{b \in \{0,1\}^n} \mathcal{T}_{t,\text{Cl}(1)_{\text{Sym}}^{\otimes n}}(|b\rangle\langle b|^{\otimes t}) &= \sum_{b \in \{0,1\}^n} \mathbb{E}_{(\pi,C) \sim \text{Cl}(1)_{\text{Sym}}^{\otimes n}} \left[(C^\dagger S_\pi^\dagger |b\rangle\langle b| S_\pi C)^{\otimes t} \right] \\ &= \mathbb{E}_{C \sim \text{Cl}(1)^{\otimes n}} \left[(C^\dagger)^{\otimes t} \mathbb{E}_{\pi \sim \text{Sym}(n)} \left[\sum_{b \in \{0,1\}^n} |\pi(b)\rangle\langle \pi(b)|^{\otimes t} \right] C^{\otimes t} \right] \\ &= \mathbb{E}_{C \sim \text{Cl}(1)^{\otimes n}} \left[(C^\dagger)^{\otimes t} \sum_{b \in \{0,1\}^n} |b\rangle\langle b|^{\otimes t} C^{\otimes t} \right] \\ &= \sum_{b \in \{0,1\}^n} \mathcal{T}_{t,\text{Cl}(1)^{\otimes n}}(|b\rangle\langle b|^{\otimes t}). \end{aligned} \quad (69)$$

The third equality follows due to the fact that permutations are bijections, hence each $\sum_{b \in \{0,1\}^n} |\pi(b)\rangle\langle \pi(b)|^{\otimes t}$ is just a reordering of the terms in $\sum_{b \in \{0,1\}^n} |b\rangle\langle b|^{\otimes t}$.

As an immediate consequence, we see that the variance of observables under subsystem symmetrization are exactly the same as with standard Pauli shadows. For the rest of this section, we will explicitly compute the variance of S_k using known Haar-averaging formulas over the Clifford group [1, Eqs. (S35) and (S36)]:

$$\mathbb{E}_{U \sim \text{Cl}(1)} U^\dagger |x\rangle\langle x| U \langle x| U A U^\dagger |x\rangle = \frac{A + \text{tr}(A)\mathbb{I}}{6} \quad (70)$$

$$\mathbb{E}_{U \sim \text{Cl}(1)} U^\dagger |x\rangle\langle x| U \langle x| U B_0 U^\dagger |x\rangle \langle x| U C_0 U^\dagger |x\rangle = \frac{\text{tr}(B_0 C_0)\mathbb{I} + B_0 C_0 + C_0 B_0}{24}, \quad (71)$$

for all unit vectors $|x\rangle \in \mathbb{C}^2$ and Hermitian matrices $A, B_0, C_0 \in \mathbb{C}^{2 \times 2}$, with $\text{tr} B_0 = \text{tr} C_0 = 0$. The extension to $\text{Cl}(1)^{\otimes n}$ follows by linearity and statistical independence. We first apply these formulas to S_1 to compute its variance. Writing $C = \bigotimes_{i \in [n]} C_i$ and $|b\rangle = \bigotimes_{i \in [n]} |b_i\rangle$, we have

$$\begin{aligned} \mathbb{E}[\hat{s}_1^2] &= \text{tr} \left(\rho \sum_{b \in \{0,1\}^n} \mathbb{E}_{C \sim \text{Cl}(1)^{\otimes n}} C^\dagger |b\rangle\langle b| C \langle b| C \mathcal{M}^{-1}(S_1) C^\dagger |b\rangle^2 \right) \\ &= \text{tr} \left(\rho \sum_{b \in \{0,1\}^n} \mathbb{E}_{C \sim \text{Cl}(1)^{\otimes n}} C^\dagger |b\rangle\langle b| C \times 3^2 \sum_{i,j \in [n]} \langle b| C Z_i C^\dagger |b\rangle \langle b| C Z_j C^\dagger |b\rangle \right) \\ &= 9 \text{tr} \left(\rho \sum_{i \in [n]} \sum_{b_i \in \{0,1\}} \mathbb{E}_{C_i \sim \text{Cl}(1)} C_i^\dagger |b_i\rangle\langle b_i| C_i \langle b_i| C_i Z_i C_i^\dagger |b_i\rangle^2 \right) \\ &\quad + 9 \times 2 \text{tr} \left(\rho \sum_{i < j} \sum_{b_i, b_j \in \{0,1\}} \mathbb{E}_{C_i, C_j \sim \text{Cl}(1)} C_i^\dagger |b_i\rangle\langle b_i| C_i \otimes C_j^\dagger |b_j\rangle\langle b_j| C_j \langle b_i| C_i Z_i C_i^\dagger |b_i\rangle \langle b_j| C_j Z_j C_j^\dagger |b_j\rangle \right) \\ &= 9 \text{tr} \left(\rho \sum_{i \in [n]} \frac{\mathbb{I}}{3} \right) + 18 \text{tr} \left(\rho \sum_{i < j} \frac{Z_i Z_j}{3 \cdot 3} \right) \\ &= 3n + \text{tr}(S_2 \rho). \end{aligned} \quad (72)$$

Thus the variance is

$$\begin{aligned} \text{Var}_\rho \left[\frac{\hat{s}_1}{s_1} \right] &= \frac{1}{s_1^2} (\mathbb{E}[\hat{s}_1^2] - \mathbb{E}[\hat{s}_1]^2) \\ &= \frac{1}{m^2} (3n + \text{tr}(S_2 \rho) - \text{tr}(S_1 \rho)^2). \end{aligned} \quad (73)$$

If ρ is the ideal state with symmetries $\text{tr}(S_1 \rho) = m$ and $\text{tr}(S_2 \rho) = m^2 - n$, then $\text{Var}_\rho[\hat{s}_1/s_1] = 2n/m^2$. On the other hand, if we make the noisy replacement $\rho \rightarrow \tilde{\rho} = \mathcal{M}^{-1} \tilde{\mathcal{M}}(\rho)$, then we can obtain a bound

$$\begin{aligned} \text{Var}_{\tilde{\rho}} \left[\frac{\hat{s}_1}{s_1} \right] &= \frac{1}{m^2} (3n + F_{Z,2}(m^2 - n) - (F_{Z,1}m)^2) \\ &\leq \frac{2n}{m^2} + 1. \end{aligned} \quad (74)$$

Next we compute the variance of estimating S_2 . Analogous to the calculation presented in Eq. (72), we expand $\langle b| C \mathcal{M}^{-1}(S_2) C^\dagger |b\rangle^2$ and group terms based on the overlapping of indices:

$$\begin{aligned} \langle b| C \mathcal{M}^{-1}(S_2) C^\dagger |b\rangle^2 &= (3^2 \times 2)^2 \sum_{i < j} \sum_{k < l} \langle b| C Z_i Z_j C^\dagger |b\rangle \langle b| C Z_k Z_l C^\dagger |b\rangle \\ &= (3^2 \times 2)^2 \left(\sum_{\substack{(i=k) < (j=l) \\ j \neq l}} + \sum_{\substack{(i=k) < j, l \\ j \neq l}} + \sum_{k < (i=l) < j} \right. \\ &\quad \left. + \sum_{i < (j=k) < l} + \sum_{\substack{i, k < (j=l) \\ i \neq k}} + \sum_{\substack{i < j, k < l \\ k \neq i \neq l; k \neq j \neq l}} \right) \langle b| C Z_i Z_j C^\dagger |b\rangle \langle b| C Z_k Z_l C^\dagger |b\rangle. \end{aligned} \quad (75)$$

We now go through each summation and evaluate the expectations:

$$\begin{aligned} & \sum_{(i=k)<(j=l)} \sum_{b_i, b_j \in \{0,1\}} \mathbb{E}_{C_i, C_j \sim \text{Cl}(1)} C_i^\dagger |b_i\rangle \langle b_i| C_i \otimes C_j^\dagger |b_j\rangle \langle b_j| C_j \langle b_i| C_i Z_i C_i^\dagger |b_i\rangle^2 \langle b_j| C_j Z_j C_j^\dagger |b_j\rangle^2 \\ &= \sum_{i<j} \frac{\mathbb{I}}{3^2} = \frac{1}{3^2} \binom{n}{2} \mathbb{I}, \end{aligned} \quad (76)$$

$$\begin{aligned} & \sum_{\substack{(i=k)<j,l \\ j \neq l}} \sum_{b_i, b_j, b_l \in \{0,1\}} \mathbb{E}_{C_i, C_j, C_l \sim \text{Cl}(1)} C_i^\dagger |b_i\rangle \langle b_i| C_i \otimes C_j^\dagger |b_j\rangle \langle b_j| C_j \otimes C_l^\dagger |b_l\rangle \langle b_l| C_l \\ & \quad \times \langle b_i| C_i Z_i C_i^\dagger |b_i\rangle^2 \langle b_j| C_j Z_j C_j^\dagger |b_j\rangle \langle b_l| C_l Z_l C_l^\dagger |b_l\rangle \\ &= 2 \sum_{i<j<l} \frac{\mathbb{I}}{3} \frac{Z_j}{3} \frac{Z_l}{3} = \frac{2}{3^3} \sum_{i<j<l} Z_j Z_l, \end{aligned} \quad (77)$$

$$\begin{aligned} & \sum_{k<(i=l)<j} \sum_{b_i, b_j, b_k \in \{0,1\}} \mathbb{E}_{C_i, C_j, C_k \sim \text{Cl}(1)} C_i^\dagger |b_i\rangle \langle b_i| C_i \otimes C_j^\dagger |b_j\rangle \langle b_j| C_j \otimes C_k^\dagger |b_k\rangle \langle b_k| C_k \\ & \quad \times \langle b_i| C_i Z_i C_i^\dagger |b_i\rangle^2 \langle b_j| C_j Z_j C_j^\dagger |b_j\rangle \langle b_k| C_k Z_k C_k^\dagger |b_k\rangle \\ &= \frac{1}{3^3} \sum_{k<i<j} Z_k Z_j, \end{aligned} \quad (78)$$

$$\begin{aligned} & \sum_{i<(j=k)<l} \sum_{b_i, b_j, b_l \in \{0,1\}} \mathbb{E}_{C_i, C_j, C_l \sim \text{Cl}(1)} C_i^\dagger |b_i\rangle \langle b_i| C_i \otimes C_j^\dagger |b_j\rangle \langle b_j| C_j \otimes C_l^\dagger |b_l\rangle \langle b_l| C_l \\ & \quad \times \langle b_i| C_i Z_i C_i^\dagger |b_i\rangle \langle b_j| C_j Z_j C_j^\dagger |b_j\rangle^2 \langle b_l| C_l Z_l C_l^\dagger |b_l\rangle \\ &= \frac{1}{3^3} \sum_{i<j<l} Z_i Z_l, \end{aligned} \quad (79)$$

$$\begin{aligned} & \sum_{\substack{i,k<(j=l) \\ i \neq k}} \sum_{b_i, b_j, b_k \in \{0,1\}} \mathbb{E}_{C_i, C_j, C_k \sim \text{Cl}(1)} C_i^\dagger |b_i\rangle \langle b_i| C_i \otimes C_j^\dagger |b_j\rangle \langle b_j| C_j \otimes C_k^\dagger |b_k\rangle \langle b_k| C_k \\ & \quad \times \langle b_i| C_i Z_i C_i^\dagger |b_i\rangle \langle b_j| C_j Z_j C_j^\dagger |b_j\rangle^2 \langle b_k| C_k Z_k C_k^\dagger |b_k\rangle \\ &= \frac{2}{3^3} \sum_{i<k<j} Z_i Z_k, \end{aligned} \quad (80)$$

$$\begin{aligned} & \sum_{\substack{i<j;k<l \\ k \neq i \neq l; k \neq j \neq l}} \bigotimes_{q \in \{i,j,k,l\}} \sum_{b_q \in \{0,1\}} \mathbb{E}_{C_q \sim \text{Cl}(1)} C_q^\dagger |b_q\rangle \langle b_q| C_q \langle b_q| C_q Z_q C_q^\dagger |b_q\rangle = \frac{1}{3^4} \sum_{\substack{i<j;k<l \\ k \neq i \neq l; k \neq j \neq l}} Z_i Z_j Z_k Z_l \\ & \quad = \frac{6}{3^4} \sum_{i<j<k<l} Z_i Z_j Z_k Z_l. \end{aligned} \quad (81)$$

Eqs. (77) to (80) can be combined by relabeling the indices and recognizing that the resulting three-index summation has $3 \binom{n}{3} = (n-2) \binom{n}{2}$ terms of the form $Z_p Z_q$ (symmetric across the index pairs $p, q \in \{i, j, k\}$ with $p \neq q$). Hence there are only $\binom{n}{2}$ unique terms, all of which are repeated $n-2$ times:

$$\frac{2}{3^3} \sum_{i<j<k} (Z_i Z_j + Z_i Z_k + Z_j Z_k) = \frac{2}{3^3} (n-2) \sum_{i<j} Z_i Z_j. \quad (82)$$

Combining these expressions, we obtain

$$\sum_{b \in \{0,1\}^n} \mathbb{E}_{C \sim \text{Cl}(1)^{\otimes n}} C^\dagger |b\rangle \langle b| C \langle b| C M^{-1} (S_2) C^\dagger |b\rangle^2 = 36 \binom{n}{2} \mathbb{I} + 24(n-2) \sum_{i<j} Z_i Z_j + 24 \sum_{i<j<k<l} Z_i Z_j Z_k Z_l. \quad (83)$$

Due to the presence of the four-body term, we will need the conserved quantity associated with M^4 :

$$\begin{aligned} M^4 &= (M^2)^2 = \left(n\mathbb{I} + 2 \sum_{i<j} Z_i Z_j \right)^2 \\ &= n^2 \mathbb{I} + 4n \sum_{i<j} Z_i Z_j + 4 \sum_{i<j} \sum_{k<l} Z_i Z_j Z_k Z_l. \end{aligned} \quad (84)$$

The four-index sum here can be grouped as we did in Eq. (75), and the conditions can be simplified as before. Along with the fact that $Z_i^2 = \mathbb{I}$, a straightforward calculation reveals

$$\sum_{i<j} \sum_{k<l} Z_i Z_j Z_k Z_l = \binom{n}{2} \mathbb{I} + 2(n-2) \sum_{i<j} Z_i Z_j + 6 \sum_{i<j<k<l} Z_i Z_j Z_k Z_l. \quad (85)$$

Plugging this into Eq. (84), combined with $\text{tr}(M^4 \rho) = m^4$ and $\text{tr}(S_2 \rho) = m^2 - n$, we arrive at

$$\text{tr}(S_4 \rho) = 24 \text{tr} \left(\rho \sum_{i<j<k<l} Z_i Z_j Z_k Z_l \right) = m^4 - n^2 - 4 \binom{n}{2} - (6n-8)(m^2 - n), \quad (86)$$

where $S_4 := \Pi_4(M^4) = 4! \sum_{i<j<k<l} Z_i Z_j Z_k Z_l$. Finally, applying this result to Eqs. (83), we can compute the variance with respect to an ideal state lying in the symmetry sector:

$$\begin{aligned} \text{Var}_\rho \left[\frac{\hat{s}_2}{s_2} \right] &= \frac{1}{(m^2 - n)^2} \left[36 \binom{n}{2} + 12(n-2) \text{tr}(S_2 \rho) + \text{tr}(S_4 \rho) - \text{tr}(S_2 \rho)^2 \right] \\ &= \frac{1}{(m^2 - n)^2} \left[32 \binom{n}{2} + (6n-16)(m^2 - n) + m^4 - n^2 \right] - 1. \end{aligned} \quad (87)$$

Again making the replacement $\rho \rightarrow \tilde{\rho}$, we instead have the following bound:

$$\begin{aligned} \text{Var}_{\tilde{\rho}} \left[\frac{\hat{s}_2}{s_2} \right] &= \frac{1}{(m^2 - n)^2} \left[36 \binom{n}{2} + 12(n-2) F_{Z,2} \text{tr}(S_2 \rho) + F_{Z,4} \text{tr}(S_4 \rho) \right] - F_{Z,2}^2 \\ &\leq \frac{1}{(m^2 - n)^2} [18n(n-1) + 12(n-2)m^2 + m^4 + 6n^2 + 8m^2] \\ &= \frac{1}{(m^2 - n)^2} [6n(4n-3) + 12nm^2 + m^2(m^2 - 16)]. \end{aligned} \quad (88)$$

If $m = \Theta(1)$, which is the case in our numerical experiments of antiferromagnetic spin systems, then $\text{Var}_{\tilde{\rho}}[\hat{s}_2/s_2] = \mathcal{O}(1)$. Recall from Eq. (74) that $\text{Var}_{\tilde{\rho}}[\hat{s}_1/s_1] = \mathcal{O}(n/m^2)$. Therefore the variance overhead of estimating these symmetry operators is, asymptotically,

$$\max \left\{ \text{Var}_{\tilde{\rho}} \left[\frac{\hat{s}_1}{s_1} \right], \text{Var}_{\tilde{\rho}} \left[\frac{\hat{s}_2}{s_2} \right] \right\} = \mathcal{O}(n) \quad (89)$$

when $\text{tr}(M \rho) = m = \Theta(1)$.

Systems wherein m depends on n will require a case-by-case analysis, which we leave to the reader. As a pathological example, consider two different functions which are both $\Theta(\sqrt{n})$: if $m = \alpha\sqrt{n}$ for some constant $\alpha \neq 1$, then $\text{Var}_{\tilde{\rho}}[\hat{s}_2/s_2] = \mathcal{O}(1)$. However, if instead $m = \sqrt{n+c}$ for some constant $c \neq 0$, then $\text{Var}_{\tilde{\rho}}[\hat{s}_2/s_2] = \mathcal{O}(n^2)$. In both cases, $\text{Var}_{\tilde{\rho}}[\hat{s}_1/s_1] = \mathcal{O}(1)$. Thus the specific form of $m = f(n)$ can drastically affect the asymptotic bounds here.

SUPPLEMENTARY NOTE 3

Spin-adapted fermionic shadows. It is well known that number-conserving fermion basis rotations which preserve spin symmetries can be block diagonalized according to the spin sectors, leading to savings in both classical and quantum resources. Here we show how to leverage spin symmetries for the broader class of fermionic Gaussian

transformations, and in particular we construct a spin-adapted matchgate shadows protocol. As such, this scheme will be informationally complete only over spin-conserving observables.

Let n_\uparrow and n_\downarrow be the number of spin-up and spin-down fermionic modes, respectively. The total number of modes is $n = n_\uparrow + n_\downarrow$, and we order the labels such that all spin-up modes come first. Gaussian transformations which do not mix between different spin types are block diagonal,

$$Q = \begin{pmatrix} Q_\uparrow & 0 \\ 0 & Q_\downarrow \end{pmatrix} \in O(2n), \quad (90)$$

where $Q_\sigma \in O(2n_\sigma) =: G_\sigma$ for each $\sigma \in \{\uparrow, \downarrow\}$. Let $G_{\text{spin}} := G_\uparrow \times G_\downarrow$ be the spin-adapted group, i.e., the set of all elements of the form of Eq. (90). In order to calculate properties of this ensemble for classical shadows, we shall use the fact that U_Q is nearly equivalent to the tensor product $U_{Q_\uparrow} \otimes U_{Q_\downarrow}$, up to a factor that depends on the determinant of Q_\uparrow . More precisely, we have the following.

Lemma 7. *For all block-diagonal Q of the form of Eq. (90), the unitary can be written as*

$$U_Q = U_{Q_\uparrow} \otimes (P_\downarrow^{s(Q_\uparrow)} U_{Q_\downarrow}) \quad (91)$$

where

$$s(Q_\uparrow) = \begin{cases} 1 & \text{if } \det(Q_\uparrow) = -1, \\ 0 & \text{else,} \end{cases} \quad (92)$$

and $P_\sigma = (-i)^{n_\sigma} \prod_{\mu=0}^{2n_\sigma-1} c_{\mu,\sigma}$ is the parity operator on the σ -spin sector.

For technical reasons, we have introduced two new sets of Majorana operators $\{c_{\mu,\sigma} \mid \mu \in [2n_\sigma]\}$ on each n_σ -mode Hilbert space, in order to talk about the different spin sectors in terms of the standard tensor product. Because the tensor product does not respect the antisymmetry of fermions, these new Majorana operators are related to the usual Majorana operators (acting on the full n -mode Hilbert space) via

$$\gamma_{\mu,\uparrow} = c_{\mu,\uparrow} \otimes \mathbb{I}, \quad (93)$$

$$\gamma_{\mu,\downarrow} = P_\uparrow \otimes c_{\mu,\downarrow}. \quad (94)$$

Lemma 7 therefore addresses this technicality of maintaining the anticommutation relations when expressing U_Q as a tensor product.

Proof (of Lemma 7). It is clear that conjugation by U_{Q_σ} transforms each $c_{\mu,\sigma}$ as desired. What we need to ensure is that the Majorana operators $\gamma_{\mu,\sigma}$ on the full Hilbert space transform properly. Indeed, Eq. (91) performs the desired transformation; for the spin-up sector, we simply have

$$U_Q \gamma_{\mu,\uparrow} U_Q^\dagger = U_{Q_\uparrow} c_{\mu,\uparrow} U_{Q_\uparrow}^\dagger \otimes \mathbb{I}. \quad (95)$$

For the spin-down sector, we will make use of following commutation relations:

$$P_\sigma U_{Q_\sigma} = \det(Q_\sigma) U_{Q_\sigma} P_\sigma \quad (96)$$

$$P_\sigma c_{\mu,\sigma} = -c_{\mu,\sigma} P_\sigma. \quad (97)$$

Consider the case $\det(Q_\uparrow) = 1$. Then $s(Q_\uparrow) = 0$, and so $U_Q = U_{Q_\uparrow} \otimes U_{Q_\downarrow}$:

$$\begin{aligned} U_Q \gamma_{\mu,\downarrow} U_Q^\dagger &= U_{Q_\uparrow} P_\uparrow U_{Q_\uparrow}^\dagger \otimes U_{Q_\downarrow} c_{\mu,\downarrow} U_{Q_\downarrow}^\dagger \\ &= P_\uparrow \otimes U_{Q_\downarrow} c_{\mu,\downarrow} U_{Q_\downarrow}^\dagger, \end{aligned} \quad (98)$$

as desired. If instead $\det(Q_\uparrow) = -1$, then $U_Q = U_{Q_\uparrow} \otimes (P_\downarrow U_{Q_\downarrow})$ and so

$$\begin{aligned} U_Q \gamma_{\mu,\downarrow} U_Q^\dagger &= U_{Q_\uparrow} P_\uparrow U_{Q_\uparrow}^\dagger \otimes P_\downarrow U_{Q_\downarrow} c_{\mu,\downarrow} U_{Q_\downarrow}^\dagger P_\downarrow^\dagger \\ &= (-P_\uparrow) \otimes (-\det(Q_\downarrow)^2 P_\downarrow P_\downarrow^\dagger U_{Q_\downarrow} c_{\mu,\downarrow} U_{Q_\downarrow}^\dagger) \\ &= P_\uparrow \otimes U_{Q_\downarrow} c_{\mu,\downarrow} U_{Q_\downarrow}^\dagger, \end{aligned} \quad (99)$$

where we have used the fact that $P_\downarrow P_\downarrow^\dagger = P_\downarrow^2 = \mathbb{I}$. \square

In the context of classical shadows, the appearance of P_\downarrow is inconsequential because it merely acts as a phasing operator which appears directly before measurement. To see this, first observe that $\mathcal{M}_Z = \mathcal{M}_{Z,\uparrow} \otimes \mathcal{M}_{Z,\downarrow}$, where $\mathcal{M}_{Z,\sigma} = \sum_{b \in \{0,1\}^{n_\sigma}} |b\rangle\langle b|$. Using the fact that $\mathcal{P}_\downarrow |b\rangle = P_\downarrow |b\rangle \langle b| P_\downarrow^\dagger = |b\rangle \langle b| = |b\rangle$, we have that that shadow channel of the spin-adapted ensemble is

$$\mathcal{M}_{G_{\text{spin}}} = \mathbb{E}_{Q_\uparrow, Q_\downarrow} \left[\mathcal{U}_{Q_\uparrow}^\dagger \otimes (\mathcal{U}_{Q_\downarrow}^\dagger \mathcal{P}_\downarrow^{s(Q_\uparrow)}) (\mathcal{M}_{Z,\uparrow} \otimes \mathcal{M}_{Z,\downarrow}) \mathcal{U}_{Q_\uparrow} \otimes (\mathcal{P}_\downarrow^{s(Q_\uparrow)} \mathcal{U}_{Q_\downarrow}) \right] \quad (100)$$

$$\begin{aligned} &= \mathbb{E}_{Q_\uparrow, Q_\downarrow} \left[(\mathcal{U}_{Q_\uparrow}^\dagger \mathcal{M}_{Z,\uparrow} \mathcal{U}_{Q_\uparrow}) \otimes (\mathcal{U}_{Q_\downarrow}^\dagger \mathcal{M}_{Z,\downarrow} \mathcal{U}_{Q_\downarrow}) \right] \\ &= \mathcal{M}_\uparrow \otimes \mathcal{M}_\downarrow. \end{aligned} \quad (101)$$

Therefore the spin-adapted matchgate shadows behaves as two independent instances on each spin sector. The estimators and variance bounds also follow straightforwardly; first, the shadow channel is

$$\begin{aligned} \mathcal{M}_{G_{\text{spin}}} &= \mathcal{M}_\uparrow \otimes \mathcal{M}_\downarrow \\ &= \sum_{j=0}^{n_\uparrow} \sum_{\ell=0}^{n_\downarrow} f_{n_\uparrow, j} f_{n_\downarrow, \ell} \Pi_j \otimes \Pi_\ell, \end{aligned} \quad (102)$$

where for ease of notation in this section, we define

$$f_{n_\sigma, j} := \binom{n_\sigma}{j} / \binom{2n_\sigma}{2j}. \quad (103)$$

For the variance, we use the property that the shadow norm of a tensor-product distribution is the product of shadow norms on each subsystem. This can be seen from the fact that shadow norm of an operator A is the spectral norm of a related operator $A^G := \mathbb{E}_{U \sim G} \sum_{b \in \{0,1\}^n} U^\dagger |b\rangle \langle b| U \langle b| U \mathcal{M}_G^{-1}(A) U^\dagger |b\rangle$:

$$\begin{aligned} \|A\|_G^2 &= \max_{\text{states } \rho} \text{tr}(\rho A^G) \\ &= \|A^G\|_\infty, \end{aligned} \quad (104)$$

which holds because A^G is positive semidefinite. For clarity, in this section we use the notation $\|\cdot\|_G$ for the shadow norm associated with the group G . Thus for any shadow channel formed as a tensor product $\mathcal{M}_{G_1 \oplus G_2} = \mathcal{M}_{G_1} \otimes \mathcal{M}_{G_2}$, we have

$$\begin{aligned} \|A_1 \otimes A_2\|_{G_1 \oplus G_2}^2 &= \max_{\text{states } \rho} \text{tr} \left(\rho \left(A_1^{G_1} \otimes A_2^{G_2} \right) \right) \\ &= \|A_1^{G_1} \otimes A_2^{G_2}\|_\infty \\ &= \|A_1^{G_1}\|_\infty \|A_2^{G_2}\|_\infty \\ &= \|A_1\|_{G_1}^2 \|A_2\|_{G_2}^2. \end{aligned} \quad (105)$$

(This argument generalizes to multiple tensor products.) Within the context of our spin-adapted ensemble, this implies that any spin-respecting Majorana operator

$$\begin{aligned} \Gamma_{\mathbf{p}, \mathbf{q}} &= (-i)^{j+\ell} \gamma_{p_1, \uparrow} \cdots \gamma_{p_{2j}, \uparrow} \gamma_{q_1, \downarrow} \cdots \gamma_{q_{2\ell}, \downarrow} \\ &= (-i)^j c_{p_1, \uparrow} \cdots c_{p_{2j}, \uparrow} \otimes (-i)^\ell c_{q_1, \downarrow} \cdots c_{q_{2\ell}, \downarrow}, \end{aligned} \quad (106)$$

has a squared shadow norm of

$$\|\Gamma_{\mathbf{p}, \mathbf{q}}\|_{G_{\text{spin}}}^2 = f_{n_\uparrow, j}^{-1} f_{n_\downarrow, \ell}^{-1}. \quad (107)$$

Thus, for Majorana operators of constant degree $2(j+\ell) \leq 2k$, the variance scales as $\mathcal{O}(n_\uparrow^j n_\downarrow^\ell) = \mathcal{O}(n^k)$, just as in the unadapted setting. Note that spin-respecting here means that the operator factorizes into an even-degree Majorana operator on each spin sector.

The advantage of this ensemble is that the required circuit depth and gate count are roughly halved, since we only need to implement two independent matchgate circuits on n_σ qubits each. Furthermore, one can also check that the shadow norm constant factors in the spin-adapted setting are also slightly smaller (for example, for $k=2$ and $n_\sigma = n/2$, the ratio of spin-adapted to unadapted shadow norms is asymptotically $\lim_{n \rightarrow \infty} f_{n/2, 1}^{-2} / f_{n, 2}^{-1} = 3/4$).

SUPPLEMENTARY NOTE 4

Improved compilation of fermionic Gaussian unitaries. In this section we describe a new scheme for compiling the matchgate circuit U_Q for arbitrary $Q \in O(2n)$, under the Jordan–Wigner mapping. This approach improves upon the circuit depth of prior art [4] by optimizing the parallelization of nearest-neighbor single- and two-qubit gates. We accomplish this by modifying previously established ideas to better respect the mapping of Majorana modes to qubits. Our improved design is implemented in code at our open-source repository (<https://github.com/zhao-andrew/symmetry-adjusted-classical-shadows>).

A previous circuit design. First we will review a prior circuit design to encode the action of U_Q into a sequence of single- and two-qubit gates, from which it will become clear where there is room for improved parallelization. While the precise scheme that we describe here has not previously appeared in the literature, the high-level ideas follow from a combination of already developed results [4–9].

Recall that our convention for the Jordan–Wigner mapping is

$$\gamma_{2p} = Z_0 \cdots Z_{p-1} X_p, \quad (108)$$

$$\gamma_{2p+1} = Z_0 \cdots Z_{p-1} Y_p \quad (109)$$

for $p \in [n]$, and our convention for the Gaussian transformation is

$$U_Q \gamma_\mu U_Q^\dagger = \sum_{\nu \in [2n]} Q_{\nu\mu} \gamma_\nu. \quad (110)$$

It is straightforward to check that $\mathcal{U} : O(2n) \rightarrow U(\mathcal{L}(\mathcal{H}))$ is a group homomorphism: $\mathcal{U}_Q \mathcal{U}_{Q'} = \mathcal{U}_{QQ'}$ for any $Q, Q' \in O(2n)$. From this property, a circuit for arbitrary U_Q can be constructed by a QR decomposition of Q . Such a decomposition yields a sequence of nearest-neighbor Givens rotations, which we then map to single- and adjacent two-qubit gates.

One possible QR decomposition is

$$Q = G_1 \cdots G_L D, \quad (111)$$

where each G_j is a Givens rotation among adjacent rows and columns, and D is the upper-right triangular matrix from the QR decomposition. Because Q is an orthogonal matrix, D is guaranteed to be a diagonal matrix with ± 1 entries along the diagonal. This is equivalent to the Reck *et al.* [5] design, and the number of Givens rotations is $L = \mathcal{O}(n^2)$ in depth $\mathcal{O}(n)$. By the homomorphism property of \mathcal{U} , this matrix decomposition yields a sequence of circuit elements that implements the desired unitary:

$$\mathcal{U}_Q = \mathcal{U}_{G_1 \cdots G_L D} = \mathcal{U}_{G_1} \cdots \mathcal{U}_{G_L} \mathcal{U}_D. \quad (112)$$

Alternatively, the Clements *et al.* [8] design computes a decomposition of the form¹

$$Q = G_{R+1} \cdots G_{R+L} D G_R \cdots G_1. \quad (113)$$

The total number of Givens rotations here is the same, $R + L = \mathcal{O}(n^2)$. However, by utilizing rotations that act from both left and right, it optimizes parallelization to reduce the depth by a constant factor (roughly 1/2).

Refs. [4, 6, 7] showed how to convert these Givens rotations into number-preserving quantum gates; here we seek to generalize to fermionic Gaussian unitaries which do not necessarily conserve particle number. While Ref. [4] also considered this scenario, they maintained the representation of Givens rotations as number-preserving gates. Their circuit design breaks particle-number symmetry by interspersing particle–hole transformations throughout the decomposition.

Instead, we will use a representation that inherently features non-number-preserving rotations. Suppose that the Givens rotation G_j acts nontrivially on the axes $(\mu, \mu + 1)$ as

$$G_j = \begin{pmatrix} 1 & \cdots & 0 & 0 & \cdots & 0 \\ \vdots & \ddots & \vdots & \vdots & & \vdots \\ 0 & \cdots & \cos \theta_j & -\sin \theta_j & \cdots & 0 \\ 0 & \cdots & \sin \theta_j & \cos \theta_j & \cdots & 0 \\ \vdots & & \vdots & \vdots & \ddots & \vdots \\ 0 & \cdots & 0 & 0 & \cdots & 1 \end{pmatrix} \in \text{SO}(2n). \quad (114)$$

¹ The use of the Clements *et al.* [8] design was first pointed out in Ref. [10] by Dominic Berry, in the context of number-preserving matchgate circuits.

The quantum gate which achieves this transformation is a single- or two-qubit Pauli rotation, given by

$$U_{G_j} = e^{-(\theta_j/2)\gamma_\mu\gamma_{\mu+1}} = \begin{cases} e^{-i(\theta_j/2)Z_p} & \text{if } \mu = 2p \text{ is even,} \\ e^{-i(\theta_j/2)X_pX_{p+1}} & \text{if } \mu = 2p + 1 \text{ is odd.} \end{cases} \quad (115)$$

Indeed, one may check that $\mathcal{U}_{G_j}(\gamma_p) = \sum_{q \in [2n]} [G_j]_{qp} \gamma_q$, as desired.

To implement the diagonal matrix D of signs, we require a different scheme. In particular, we can construct U_D as a single layer of Pauli gates. Consider the 2×2 block along the diagonal

$$D^{(p)} = \begin{pmatrix} D_{2p} & 0 \\ 0 & D_{2p+1} \end{pmatrix}, \quad (116)$$

which describes the transformation

$$\gamma_{2p} \mapsto D_{2p}\gamma_{2p}, \quad (117)$$

$$\gamma_{2p+1} \mapsto D_{2p+1}\gamma_{2p+1}. \quad (118)$$

If $D_{2p} = D_{2p+1} = 1$, then clearly no operations are required. If instead $D_{2p} = D_{2p+1} = -1$, then conjugation by Z_p applies the desired signs on γ_{2p} and γ_{2p+1} while leaving all other Majorana operators invariant.

The remaining cases, $D_{2p} = -D_{2p+1}$, can be handled as follows. First, suppose $D_{2p} = 1$. We wish to find the gates which perform the transformation

$$\gamma_{2p} \mapsto \gamma_{2p}, \quad (119)$$

$$\gamma_{2p+1} \mapsto -\gamma_{2p+1}, \quad (120)$$

while leaving all other Majorana operators invariant. We can almost accomplish this with X_p , since it will map X_p to itself and Y_p to $-Y_p$. It also commutes with all Majorana operators $\gamma_{2q}, \gamma_{2q+1}$ for $q < p$. However, for $q > p$ this will accrue unwanted signs:

$$\begin{aligned} X_p\gamma_{2q}X_p &= X_p(Z_0 \cdots Z_p \cdots Z_{q-1}X_q)X_p \\ &= -Z_0 \cdots Z_p \cdots Z_{q-1}X_q \\ &= -\gamma_{2q}, \end{aligned} \quad (121)$$

$$\begin{aligned} X_p\gamma_{2q+1}X_p &= X_p(Z_0 \cdots Z_p \cdots Z_{q-1}Y_q)X_p \\ &= -Z_0 \cdots Z_p \cdots Z_{q-1}Y_q \\ &= -\gamma_{2q+1}. \end{aligned} \quad (122)$$

To correct these signs, we introduce a Pauli- Z string running in the opposite direction of the Jordan–Wigner convention. That is, define

$$P_p^{(X)} := X_p Z_{p+1} \cdots Z_{n-1}. \quad (123)$$

This unitary has the correct action on $\gamma_{2p}, \gamma_{2p+1}$ and continues to commute with the Majorana operators $\gamma_{2q}, \gamma_{2q+1}$ with $q < p$. For $q > p$, however, we now have

$$\begin{aligned} P_p^{(X)}\gamma_{2q}P_p^{(X)} &= (X_p Z_{p+1} \cdots Z_{n-1})(Z_0 \cdots Z_p \cdots Z_{q-1}X_q)(X_p Z_{p+1} \cdots Z_{n-1}) \\ &= Z_0 \cdots (-Z_p) \cdots Z_{q-1}(-X_q) \\ &= \gamma_{2q}, \end{aligned} \quad (124)$$

$$\begin{aligned} P_p^{(X)}\gamma_{2q+1}P_p^{(X)} &= (X_p Z_{p+1} \cdots Z_{n-1})(Z_0 \cdots Z_p \cdots Z_{q-1}Y_q)(X_p Z_{p+1} \cdots Z_{n-1}) \\ &= Z_0 \cdots (-Z_p) \cdots Z_{q-1}(-Y_q) \\ &= \gamma_{2q+1}. \end{aligned} \quad (125)$$

Thus $P_p^{(X)}$ implements the desired transformation by $D^{(p)} = \text{diag}(1, -1)$. For $D^{(p)} = \text{diag}(-1, 1)$, we simply replace $P_p^{(X)}$ by an analogously defined $P_p^{(Y)}$. This causes the sign of γ_{2p} , rather than γ_{2p+1} , to flip, while retaining all other properties.

Altogether, we determine these transformations for all 2×2 diagonal blocks of D , resulting in n Pauli strings of the form

$$W_p = \begin{cases} \mathbb{I} & \text{if } D^{(p)} = \text{diag}(1, 1), \\ Z_p & \text{if } D^{(p)} = \text{diag}(-1, -1), \\ P_p^{(X)} & \text{if } D^{(p)} = \text{diag}(1, -1), \\ P_p^{(Y)} & \text{if } D^{(p)} = \text{diag}(-1, 1). \end{cases} \quad (126)$$

The overall transformation is then simply the product of these Pauli strings, which can be concatenated into a single layer of Pauli gates:

$$U_D = \prod_{p \in [n]} W_p. \quad (127)$$

Note that the order of this product does not matter, since Pauli gates commute up to an unobservable global phase.

Discussion on suboptimality. Now we observe that, depending on the parity of $\mu \in \{0, \dots, 2n - 2\}$, U_{G_j} is either a single- or two-qubit gate. However, the decomposition of Q described above is implicitly optimized under the assumption that only two-qubit gates are present: each Givens rotation acts on two axes at a time, and it is assumed that this corresponds to physically acting on two wires at a time. This results in underutilized space in the quantum circuit whenever a single-qubit Z rotation occurs, as it leaves a qubit wire needlessly idle. This is true for both the Reck *et al.* [5] and Clements *et al.* [8] designs. Ultimately, this suboptimality is due to the fact that Q is a $2n \times 2n$ matrix, so there is a two-to-one correspondence between axes and qubits: the rows/columns labeled by $(2p, 2p + 1)$ correspond to two Majorana operators, both of which are in turn associated with a single qubit p . Note that this discrepancy is not present in circuit designs for the class of number-conserving rotations [4, 6, 7], which are instead more compactly represented by an $n \times n$ unitary matrix already.

Circuit design with improved parallelization. Now we introduce a circuit design which explicitly accounts for this two-to-one correspondence. The basic idea is to generalize the notion of Givens rotations, which act on a two-dimensional subspace to zero out a single matrix element, to a four-dimensional orthogonal transformation which zeroes out blocks of 2×2 at a time. Each 4×4 orthogonal transformation acts on the axes $(2p, 2p + 1, 2p + 2, 2p + 3)$, which corresponds to qubits p and $p + 1$. By performing this process according to the scheme of Clements *et al.* [8] (but now treating each 2×2 block of Q as a “single” element), we obtain a decomposition wherein the optimal parallelization of the scheme is fully preserved in terms of interactions between nearest-neighbor qubits. Finally, each 4×4 orthogonal transformation is ultimately decomposed into six rotations of the form of Eq. (115) and a layer of Pauli gates, achieved by the standard decomposition that we described in “A previous circuit design” (i.e., by bootstrapping off the prior scheme within blocks of $2n = 4$). Note that in principle one may instead implement the 4×4 orthogonal transformations using any gate set of one’s choice, rather than XX and Z rotations.

We now describe the algorithm in detail. First we compute a decomposition analogous to the Clements *et al.* [8] design,

$$Q = B_{R+1} \cdots B_{R+L} D G B_R \cdots B_1, \quad (128)$$

but instead of Givens rotations, each B_k acts nontrivially on a 4×4 block. (Note that there is a single 2×2 Givens rotation G as well, which serves to zero out a final matrix element that we will elaborate on later.) We accomplish this by treating Q as an $n \times n$ matrix of 2×2 blocks,

$$Q_{p,q} = \begin{pmatrix} Q_{2p,2q} & Q_{2p,2q+1} \\ Q_{2p+1,2q} & Q_{2p+1,2q+1} \end{pmatrix}, \quad (129)$$

for each $p, q \in [n]$. Just as Givens rotations are chosen to zero a specific matrix element, each B_k acts to zero out a particular 2×2 block $Q_{p,q}$.

Suppose we want to find a B_{R+i} which acts from the left ($i = 1, \dots, L$) to zero out the block $Q_{p,q}$. Then we perform a QR decomposition on the 4×2 submatrix which includes the target block and the block directly above it:

$$\begin{pmatrix} Q_{p-1,q} \\ Q_{p,q} \end{pmatrix} = B'_{R+i} \begin{pmatrix} * & * \\ 0 & * \\ 0 & 0 \\ 0 & 0 \end{pmatrix}, \quad (130)$$

hence zeroing out the lower block as desired. Here, $B'_{R+i} \in O(4)$ is computed from the QR decomposition, and so the orthogonal matrix $B_{R+i} \in O(2n)$ appearing in Eq. (128) is defined as B'_{R+i} along the axes $(2p-2, 2p-1, 2p, 2p+1)$ and the identity elsewhere.

Similarly, if we want a B_j which acts from the right ($j = 1, \dots, R$), then we consider instead a 2×4 submatrix with the target block on the left:

$$(Q_{p,q} | Q_{p,q+1}). \quad (131)$$

This can be zeroed out by performing an LQ decomposition (which is essentially just the transpose of the QR decomposition). For notation in this section, let tildes denote the flipping of rows in a matrix, for example

$$M = \begin{pmatrix} M_{11} & M_{12} \\ M_{21} & M_{22} \\ M_{31} & M_{32} \\ M_{41} & M_{42} \end{pmatrix} \mapsto \tilde{M} = \begin{pmatrix} M_{41} & M_{42} \\ M_{31} & M_{32} \\ M_{21} & M_{22} \\ M_{11} & M_{12} \end{pmatrix}. \quad (132)$$

Then performing an LQ decomposition on the row-flipped version of Eq. (131), we have

$$\begin{aligned} (\tilde{Q}_{p,q} | \tilde{Q}_{p,q+1}) &= \begin{pmatrix} * & 0 & 0 & 0 \\ * & * & 0 & 0 \end{pmatrix} B'_j \\ &= \begin{pmatrix} 0 & 0 & 0 & * \\ 0 & 0 & * & * \end{pmatrix} \tilde{B}'_j. \end{aligned} \quad (133)$$

Flipping the rows back to normal on the lefthand side, we get

$$(Q_{p,q} | Q_{p,q+1}) = \begin{pmatrix} 0 & 0 & * & * \\ 0 & 0 & 0 & * \end{pmatrix} \tilde{B}'_j \quad (134)$$

as desired. Then we define $B_j \in O(2n)$ acting as \tilde{B}'_j on the axes $(2q, 2q+1, 2q+2, 2q+3)$ and trivially elsewhere.

Now we address the need for the sole Givens rotation G appearing in Eq. (128). As the zeroing-out procedure described above progresses, the nonzero blocks get “pushed” towards the diagonal until the final matrix is (2×2) -block diagonal. These nonzero blocks must be triangular because they are produced by QR/LQ decompositions; but since Q is orthogonal, this implies that the final triangular blocks along the diagonal must be diagonal themselves. The exception to this is either the leftmost or rightmost block, depending on whether n is even or odd. This is because the decomposition procedure inevitably leaves one of those blocks untouched, so it was never made triangular/diagonal.

This can be visualized as follows: if n is odd, then we have

$$Q \rightarrow \left(\begin{array}{c|c|c} * & * & * \\ * & * & * \\ * & * & * \\ * & * & * \\ \hline 0 & 0 & * \\ 0 & 0 & * \end{array} \right) \rightarrow \left(\begin{array}{c|c|c} * & * & * \\ 0 & * & * \\ 0 & 0 & * \\ 0 & 0 & * \\ \hline 0 & 0 & * \\ 0 & 0 & * \end{array} \right) \rightarrow \left(\begin{array}{c|c|c} * & * & * \\ 0 & * & * \\ 0 & 0 & * \\ 0 & 0 & 0 \\ \hline 0 & 0 & 0 \\ 0 & 0 & 0 \end{array} \right) = \left(\begin{array}{c|c|c} \pm 1 & 0 & 0 \\ 0 & \pm 1 & 0 \\ 0 & 0 & \pm 1 \\ 0 & 0 & 0 \\ \hline 0 & 0 & 0 \\ 0 & 0 & 0 \end{array} \right). \quad (135)$$

We use boldface to clarify which matrix elements are newly zeroed at each step. The condition that Q is an orthogonal matrix implies the final equality. It also enforces the remaining 2×2 block to be orthogonal, so that we can diagonalize it by computing the appropriate Givens rotation acting on axes $(2n-2, 2n-1)$. This elucidates the appearance of G

Hence for operators of degree $2j$,

$$\begin{aligned} U_Q \Gamma_{\boldsymbol{\mu}} U_Q^\dagger &= (-i)^j U_Q \gamma_{\mu_1} \cdots \gamma_{\mu_{2j}} U_Q^\dagger \\ &= (-i)^j s_{\pi^{-1}(\mu_1)} \cdots s_{\pi^{-1}(\mu_{2j})} \gamma_{\pi^{-1}(\mu_1)} \cdots \gamma_{\pi^{-1}(\mu_{2j})}. \end{aligned} \quad (138)$$

We would like to retain the ordering of indices when working with the multidegree Majorana operators; therefore we introduce a further permutation as $\tilde{\pi}^{-1}(\mu_i)$, which is defined to satisfy $\tilde{\pi}^{-1}(\mu_1) < \cdots < \tilde{\pi}^{-1}(\mu_{2j})$. This incurs another sign factor $(-1)^p$ where $p \in \{0, 1\}$ is the parity of the permutation which sends $\pi^{-1}(\boldsymbol{\mu}) \mapsto \tilde{\pi}^{-1}(\boldsymbol{\mu})$. Collecting all signs as $\text{sign}_Q(\boldsymbol{\mu}) = (-1)^p s_{\pi^{-1}(\mu_1)} \cdots s_{\pi^{-1}(\mu_{2j})}$, we arrive at

$$\text{tr}(\Gamma_{\boldsymbol{\mu}} \hat{\rho}_{Q,b}) = f_{2j}^{-1} \text{sign}_Q(\boldsymbol{\mu}) \langle b | \Gamma_{\tilde{\pi}^{-1}(\boldsymbol{\mu})} | b \rangle. \quad (139)$$

The matrix element $\langle b | \Gamma_{\tilde{\pi}^{-1}(\boldsymbol{\mu})} | b \rangle$ is nonzero if and only if $\tilde{\pi}^{-1}(\boldsymbol{\mu}) \in \mathcal{D}_{2n,2j}$, from which its value of ± 1 is straightforward to determine (e.g., by mapping to Pauli- Z operators). In total, evaluating Eq. (139) takes time $\mathcal{O}(n^2 + j^2 + j)$, corresponding respectively to the inversion of $\pi \in \text{Sym}(2n)$, the calculation of $\tilde{\pi}^{-1}(\boldsymbol{\mu})$ and its parity on $2j$ indices, and evaluating the product of $2j + 1$ signs and $\langle b | \Gamma_{\tilde{\pi}^{-1}(\boldsymbol{\mu})} | b \rangle$, the latter requiring only checking $2j$ indices and j bits of b . Assuming $j \leq k = \mathcal{O}(1)$, this implies a computational complexity of $\mathcal{O}(n^2)$ per operator per sample.

To compute this estimator for all $\boldsymbol{\mu} \in \mathcal{C}_{2n,2j}$, a naive approach iterates through each $\boldsymbol{\mu}$, of which there are $\binom{2n}{2j} = \mathcal{O}(n^{2j})$ many. Repeating this for each of the T samples would therefore cost $\mathcal{O}(T(n^{2j} + n^2)) = \mathcal{O}(Tn^{2j})$ time. Noting that $T = \tilde{\mathcal{O}}(n^j \epsilon^{-2})$ suffices for ϵ -accurate estimation,² the total complexity of $\tilde{\mathcal{O}}(n^{3j} \epsilon^{-2}) \leq \tilde{\mathcal{O}}(n^{3k} \epsilon^{-2})$ would be unacceptably large.

We can speed up the computation over all operators per sample to $\mathcal{O}(n^j)$ by using the fact that many $\langle b | \Gamma_{\tilde{\pi}^{-1}(\boldsymbol{\mu})} | b \rangle$ vanish. That is, rather than compute $\tilde{\pi}^{-1}(\boldsymbol{\mu})$ for all $\boldsymbol{\mu} \in \mathcal{C}_{2n,2j}$ and checking whether each is an element of $\mathcal{D}_{2n,2j}$, we work backwards by looping over all target elements $\boldsymbol{\tau} \in \mathcal{D}_{2n,2j}$ and computing $\pi(\boldsymbol{\tau})$ to find its preimage. As before, let $\tilde{\pi}(\boldsymbol{\tau}) \in \mathcal{C}_{2n,2j}$ be the reordering of $\pi(\boldsymbol{\tau})$ with associated sign $(-1)^p$. Then for each $\boldsymbol{\tau} \in \mathcal{D}_{2n,2j}$, we compute the estimator for $\Gamma_{\tilde{\pi}(\boldsymbol{\tau})}$,

$$\text{tr}(\Gamma_{\tilde{\pi}(\boldsymbol{\tau})} \hat{\rho}_{Q,b}) = f_{2j}^{-1} \text{sign}_Q(\pi(\boldsymbol{\tau})) \langle b | \Gamma_{\boldsymbol{\tau}} | b \rangle, \quad (140)$$

where the cumulative sign is $\text{sign}_Q(\pi(\boldsymbol{\tau})) = (-1)^p s_{\tau_1} \cdots s_{\tau_{2j}}$. All other Majorana operators not in the preimage are implicitly assigned an estimate of 0. Hence we only iterate over the $\binom{n}{j} = \mathcal{O}(n^j)$ elements of $\mathcal{D}_{2n,2j}$, with each evaluation of Eq. (140) taking $\mathcal{O}(j^2) = \mathcal{O}(1)$ time. Note that this approach also avoids the need to find the inverse permutation π^{-1} .

Performing this procedure over all T samples results in a time complexity of $\mathcal{O}(n^j T) \leq \mathcal{O}(n^k T)$, running over all $j \in \{1, \dots, k\}$. We can also include an additive $\mathcal{O}(n^{2k})$ cost to preallocate storage for $\bigcup_{j \leq k} \mathcal{C}_{2n,2j}$. While not strictly necessary, this is convenient in practice, and besides when $T = \tilde{\mathcal{O}}(n^k \epsilon^{-2})$ the total complexity is $\tilde{\mathcal{O}}(n^{2k} \epsilon^{-2})$ whether or not we preallocate memory.

For the spin-adapted shadows, because the protocol factorizes across the spin sectors, we perform this algorithm on each sector independently. The estimator for operators of the form $\Gamma_{\boldsymbol{\mu}} \otimes \Gamma_{\boldsymbol{\nu}}$ is then the product of the independent estimates. Note that if either $|\boldsymbol{\mu}|$ or $|\boldsymbol{\nu}|$ are odd, then the estimator always vanishes; this reflects the fact that the spin-adapted ensemble is not informationally complete over such operators.

Pauli shadows. Because single-qubit measurements factorize, we consider each qubit $i \in [n]$ independently. Given the random Clifford $C_i \in \text{Cl}(1)$ and measurement outcome $b_i \in \{0, 1\}$, the estimator for $\sigma^{(i)} \in \{\mathbb{I}, X, Y, Z\}$ is [1]

$$\text{tr}(\sigma^{(i)} \hat{\rho}_{C_i, b_i}) = 3 \langle b_i | C_i \sigma^{(i)} C_i^\dagger | b_i \rangle - \text{tr}(\sigma^{(i)}). \quad (141)$$

Each Pauli-shadow sample is stored as (W_i, b_i) , where $W_i = C_i^\dagger Z C_i \in \pm\{X, Y, Z\}$. Evaluating Eq. (141) reduces to checking the conditions

$$\text{tr}(\sigma^{(i)} \hat{\rho}_{C_i, b_i}) = \begin{cases} \pm 3 \langle b_i | Z | b_i \rangle & \text{if } \sigma^{(i)} = \pm W_i, \\ 1 & \text{if } \sigma^{(i)} = \mathbb{I}, \\ 0 & \text{else.} \end{cases} \quad (142)$$

² The notation $\tilde{\mathcal{O}}(\cdot)$ suppresses polylogarithmic factors in the complexity.

The product over $i \in [n]$ then estimates $P = \bigotimes_{i \in [n]} \sigma^{(i)}$ for the full n -qubit system. This suffices to estimate any k -local observable, which can be decomposed into a linear combination of polynomially many ($\leq k$)-local Pauli operators.

The total time complexity of estimating all k -local Pauli operators with T snapshots is $\mathcal{O}(n^k T)$. The algorithm is as follows. For each $W = (W_0, \dots, W_{n-1})$, we take, for each $j \leq k$, all $\binom{n}{j}$ combinations W_{i_1}, \dots, W_{i_j} and compute Eq. (142) for each $\sigma^{(i)} = \pm W_i$. We assign the result as an estimate for the j -local operator $P = W_{i_1} \otimes \dots \otimes W_{i_j} \otimes \mathbb{I}^{\otimes(n-j)}$, and implicitly assign 0 to all other Pauli operators. Note that there are a total of $\sum_{j \leq k} 3^j \binom{n}{j} = \mathcal{O}(n^k)$ local Pauli operators, so preallocating storage here is asymptotically negligible.

For the subsystem-symmetrized protocol, the n -qubit estimator now takes the form

$$\text{tr}(P \hat{\rho}_{(\pi, C), b}) = 3^{|P|} \langle b | S_\pi C P C^\dagger S_\pi^\dagger | b \rangle. \quad (143)$$

Using the fact that $S_\pi^\dagger | b \rangle = \bigotimes_{i \in [n]} |\pi(b_i)\rangle$, we can simply apply the standard scheme described above, but with the replacement $b_i \rightarrow \pi(b_i)$. For each sample this is only an additive $\mathcal{O}(n)$ cost.

SUPPLEMENTARY NOTE 6

Additional details on numerical experiments. Here we supply further information regarding the numerical simulations, to both provide additional insight into our results and facilitate easier replication of our results by the motivated reader.

Readout errors. In the main text, we demonstrated our mitigation strategy under single-qubit readout errors. The noise channels occur immediately before measurement and are implemented probabilistically: independently and identically (i.i.d.) on each qubit per circuit repetition. We consider depolarizing, amplitude-damping, and bit-flip errors occurring with probability p , which are respectively

$$\mathcal{E}_{\text{dep}}(\rho) = (1-p)\rho + p\frac{\mathbb{I}}{2}, \quad (144)$$

$$\mathcal{E}_{\text{AD}}(\rho) = E_0 \rho E_0^\dagger + E_1 \rho E_1^\dagger, \quad (145)$$

$$E_0 = \begin{pmatrix} 1 & 0 \\ 0 & \sqrt{1-p} \end{pmatrix}, E_1 = \begin{pmatrix} 0 & \sqrt{p} \\ 0 & 0 \end{pmatrix}$$

$$\mathcal{E}_{\text{BF}}(\rho) = (1-p)\rho + pX\rho X. \quad (146)$$

These models obey Assumptions 1, although we comment that more complicated noise channels can also satisfy the assumptions, such as non-i.i.d. errors, correlated multiqubit errors, and even coherent gate errors [13].

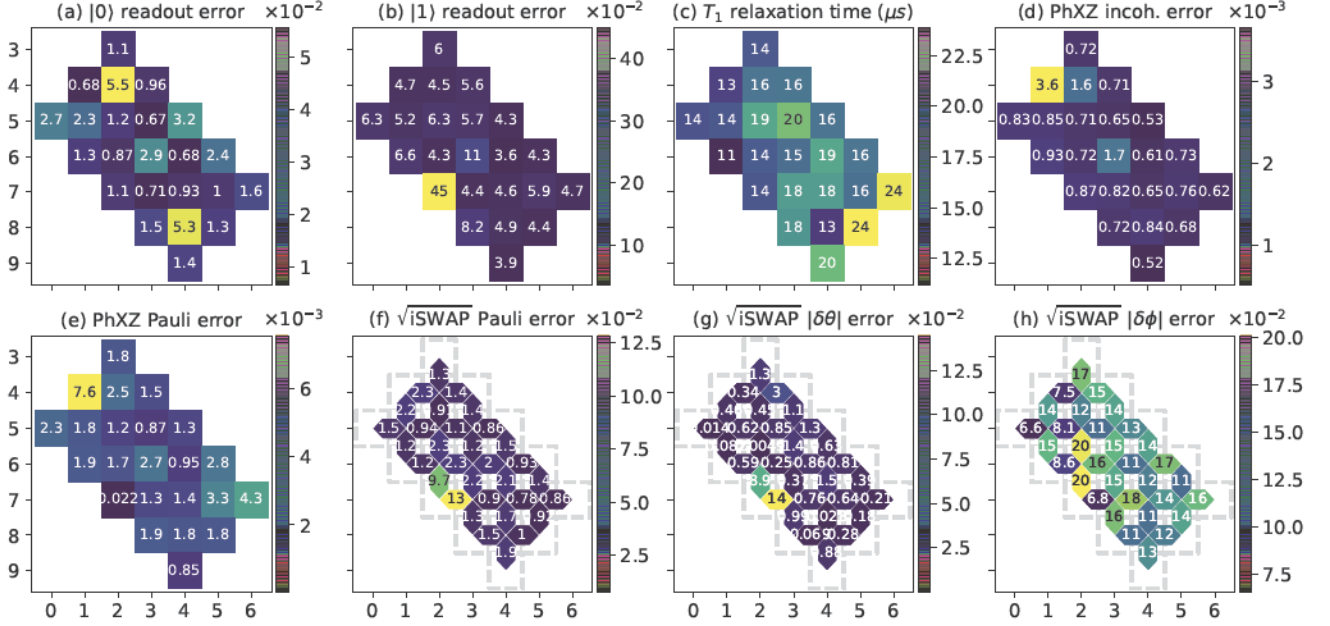
QVM gate set and noise model. The noise model we implement on the Cirq Quantum Virtual Machine is based on the Google Sycamore processor ‘‘Rainbow,’’ a 2D grid of 23 superconducting qubits. We use the calibration data obtained from November 16, 2021, which can be found in the Cirq open-source repository [14]. The native gate set that we compile our circuits to include single-qubit rotations in the form of phased XZ gates,

$$\begin{aligned} & \text{PhXZ}(x, z, a) \\ &= \begin{pmatrix} e^{i\frac{\pi x}{2}} \cos\left(\frac{\pi x}{2}\right) & -ie^{i\pi\left(\frac{x}{2}-a\right)} \sin\left(\frac{\pi x}{2}\right) \\ -ie^{i\pi\left(\frac{x}{2}+a+z\right)} \sin\left(\frac{\pi x}{2}\right) & e^{i\pi\left(\frac{x}{2}+z\right)} \cos\left(\frac{\pi x}{2}\right) \end{pmatrix} \\ &= Z^z Z^a X^x Z^{-a}. \end{aligned} \quad (147)$$

This describes a rotation by πx about an axis determined by the parameter a within the xy plane, followed by a phasing of πz . The native two-qubit gates that we use are

$$\begin{aligned} \sqrt{i\text{SWAP}} &= \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \frac{1}{\sqrt{2}} & \frac{i}{\sqrt{2}} & 0 \\ 0 & \frac{i}{\sqrt{2}} & \frac{1}{\sqrt{2}} & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \\ &= e^{i\frac{\pi}{4}(X \otimes X + Y \otimes Y)/2}, \end{aligned} \quad (148)$$

constrained to the nearest-neighbor connectivity of the chip.



Supplementary Figure 1. Chip layout, connectivity, and error rates of the Google Sycamore Rainbow processor, as simulated by the QVM. Our eight-qubit numerical experiments use the 2×4 grid spanning from qubit (5, 1) to qubit (6, 4). (a, b) Readout error probabilities, measured in parallel. (c) Characteristic T_1 qubit decay times, measured in isolation. (d) Incoherent component of single-qubit gate errors, measured in isolation from RB and purity benchmarking data. Used in conjunction with T_1 to infer the T_2 dephasing time. (e) Total Pauli error of single-qubit gates, measured in isolation via RB. (f) Total Pauli error of two-qubit \sqrt{i} SWAP gates, measured in parallel via XEB. (g, h) Coherent errors in two-qubit gates \sqrt{i} SWAP = fSim($\theta = -\frac{\pi}{4}, \phi = 0$), measured in parallel from XEB data. We display the magnitudes $|\delta\theta|, |\delta\phi|$ for visualization purposes; full calibration data (including signs) is available at the Cirq open-source repository [14].

The QVM noise model that we simulate is not fully comprehensive of all types of errors occurring in an actual device, however it captures the most dominant error sources in the superconducting platform [15]. It consists of four categories:

1. Readout errors are modeled as asymmetric bit-flip channels on each qubit. The asymmetry reflects the fact that the probability of a $|1\rangle$ outcome being erroneously measured as $|0\rangle$ is generally higher than misreading a $|0\rangle$ outcome. Although the errors are modeled as single-qubit channels, the calibration data is taken from parallel experiments, to potentially account for effects such as readout crosstalk and other unintended interactions between qubits.
2. Decay (T_1) and dephasing (T_2) errors occur whenever a qubit idles during a moment (layer) of a circuit. Both T_1 and T_2 relaxations are incorporated into a single channel,

$$\mathcal{E}_{\text{idle}}(\rho) = \begin{pmatrix} 1 - \rho_{11}e^{-t/T_1} & \rho_{01}e^{-t/T_2} \\ \rho_{10}e^{-t/T_2} & \rho_{11}e^{-t/T_1} \end{pmatrix}. \quad (149)$$

The decay time T_1 is characterized by a simple experiment that prepares $|1\rangle$ and measures the survival probability as a function of t . This experiment is performed in isolation, i.e., one qubit at a time while all other qubits on the chip idle.

The T_2 time is determined from the equation

$$\frac{1}{T_2} = \frac{1}{2T_1} + \frac{1}{T_\phi}, \quad (150)$$

where $1/T_\phi$ is the pure dephasing rate that can in principle be measured by Ramsey interferometry. For simplicity, however, this noise model instead approximates T_ϕ from the total single-qubit incoherent error ϵ_{inc} , which is determined by purity benchmarking [16, 17] performed in isolation. To leading order, T_ϕ is approximated

using the relation

$$\epsilon_{\text{inc}} = \frac{t}{3T_1} + \frac{t}{3T_\phi} + \mathcal{O}(t^2). \quad (151)$$

The time t which appears in the model channel $\mathcal{E}_{\text{idle}}$ is the longest gate duration occurring within that moment: PhXZ gates have a duration of 25 ns, while $\sqrt{\text{iSWAP}}$ gates take 32 ns.

3. Single-qubit gate errors are modeled as depolarizing channels occurring after each gate. The depolarizing rate is set to match the total single-qubit Pauli error, which is measured from the device via randomized benchmarking (RB) [18, 19] in isolation.
4. Two-qubit gate errors are modeled with both coherent and incoherent components. The coherent contribution uses the fact that $\sqrt{\text{iSWAP}}$ is an instance of the general fermionic simulation (fSim) gate,

$$\text{fSim}(\theta, \phi) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos \theta & -i \sin \theta & 0 \\ 0 & -i \sin \theta & \cos \theta & 0 \\ 0 & 0 & 0 & e^{-i\phi} \end{pmatrix}, \quad (152)$$

which is a native, tunable interaction on the superconducting platform. The $\sqrt{\text{iSWAP}}$ gate is the instance $(\theta, \phi) = (-\frac{\pi}{4}, 0)$. Coherent errors are thus modeled as an overrotation by $(\delta\theta, \delta\phi)$, which are determined for each pair of connected qubits by fitting to cross-entropy benchmarking (XEB) data using random cycles of gates across the chip [20–22].

After the coherent overrotation, an incoherent error follows, modeled as a two-qubit depolarizing channel. The depolarizing rate $r_{\text{dep}}^{(i,j)}$ for each pair (i, j) of connected qubits is inferred as follows: from the total XEB Pauli error $r_{\text{XEB}}^{(i,j)}$, we subtract off the single-qubit incoherent error rates $r_{\text{inc}}^{(i)}, r_{\text{inc}}^{(j)}$ (determined from RB), as well as the average entangling error rate $r_{\text{ent}}^{(i,j)}$, which are calculated using the coherent errors $\delta\theta, \delta\phi$. The model’s two-qubit depolarizing rate is then set to account for the remaining amount of error:

$$r_{\text{dep}}^{(i,j)} = r_{\text{XEB}}^{(i,j)} - r_{\text{inc}}^{(i)} - r_{\text{inc}}^{(j)} - r_{\text{ent}}^{(i,j)}. \quad (153)$$

Due to the nature of XEB, both two-qubit error sources are characterized by parallel experimental data.

Further details of the noise model, its numerical implementation, and the calibration-data acquisition are described in Ref. [15], as well as in the Cirq repository [14]. For completeness, in Supplementary Figure 1 we display a series of plots which show the chip connectivity and numerical values of the calibration data used for the various errors described above.

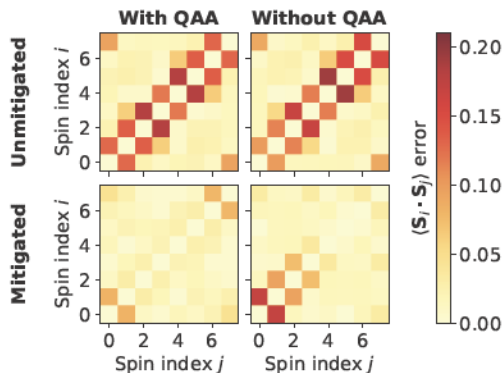
Compiling circuits to the native gate set. Single-qubit rotations are compiled into PhXZ gates according to an Euler-angle decomposition. Two-qubit unitaries are compiled into at most three $\sqrt{\text{iSWAP}}$ gates (interleaved with single-qubit rotations) by a KAK decomposition, although most two-qubit unitaries (79% with respect to the Haar measure) can be implemented with just two $\sqrt{\text{iSWAP}}$ gates [23]. After compiling the entire circuit into this gate set, single-qubit rotations are concatenated into a single PhXZ gate whenever possible. All operations besides readout are pushed as early into the circuit as possible.

One exception we make is in the random permutation circuits S_π appearing in the group $\text{Cl}(1)_{\text{Sym}}^{\otimes n}$ (for subsystem-symmetrized Pauli shadows). First, we decompose π into an parallelized network of adjacent transpositions using an odd–even sorting algorithm [24]. Each transposition $i \leftrightarrow j$ corresponds to a SWAP gate between qubits i and j . However, rather than compile SWAP to the gate set directly (which would require three $\sqrt{\text{iSWAP}}$ gates and four layers of $\text{PhXZ}^{\otimes 2}$ gates), we instead implement the unitary

$$\text{iSWAP} = \sqrt{\text{iSWAP}} \times \sqrt{\text{iSWAP}}, \quad (154)$$

which uses only two $\sqrt{\text{iSWAP}}$ gates and no single-qubit gates. The iSWAP gate differs from SWAP only by a phasing of i on the basis states $|01\rangle$ and $|10\rangle$. Such a replacement is valid because S_π occurs only at the end of the circuit, immediately before readout. Thus while this phasing is technically unwanted, it has no observable effect on the measurement outcomes.

Finally, we note that the Trotter circuits for our Fermi–Hubbard simulations are optimized for the Sycamore architecture according to Ref. [25], which we follow closely. In particular, open-source code for their implementation



Supplementary Figure 2. Behavior of qubit assignment averaging (QAA), demonstrated with the XXZ spin–spin correlations ($R = 4$ Trotter steps) and $T = 4.8 \times 10^5$ subsystem-symmetrized Pauli shadows. Experiments with QAA average over the twelve configurations described in “Qubit assignment averaging,” while experiments without QAA fix the qubit ordering $0 \mapsto (5, 1), \dots, 7 \mapsto (6, 1)$.

can be found in Ref. [26].

Qubit assignment averaging. Our eight-qubit numerical experiments on the QVM utilize the 2×4 grid spanning from qubits $(5, 1)$ to $(6, 4)$ (see Supplementary Figure 1). To map these qubits to the simulated degrees of freedom (fermion modes or spin-1/2 particles), we employ qubit assignment averaging (QAA), which was introduced in Ref. [25] in order to handle the issue of inhomogeneous error rates across a noisy quantum device. QAA works by identifying N different assignments of the n qubits and allocating T/N of the experimental repetitions to each realization. Properties are estimated by averaging over all T samples as usual. In principle, one can use a combination of shifting, rotating, and flipping the qubits throughout the chip; for our simulations, we vary qubit assignments within the same fixed 2×4 grid.

For the Fermi–Hubbard model, we assign a spin sector to each of the parallel 1×4 qubit chains. We average over $N = 4$ different qubit assignments, defined by setting either the top or bottom chain as the spin-up chain, and ordering the four site labels starting either from the left or the right.

For the XXZ Heisenberg model, the eight-spin chain is embedded into the 2×4 grid of qubits. Each qubit assignment ($N = 12$) is defined by setting one of six qubits $\in \{(5, 1), (5, 2), (5, 3), (5, 4), (6, 4), (6, 1)\}$ as either the left end (ordered clockwise) or right end (ordered counterclockwise) of the spin chain.

While QAA aims to reduce device inhomogeneities, it cannot lower the total amount of circuit noise. Thus QAA does not necessarily improve prediction accuracy with the unmitigated (standard shadow) estimators. Instead, homogenizing the noise appears to massage it into an effective form which approximately satisfies Assumptions 1 better than a single fixed configuration. We substantiate this claim with Supplementary Figure 2, using spin–spin correlations of the XXZ model ($R = 4$ Trotter steps) as a demonstrative example. We see that the unmitigated errors are virtually identical whether or not we perform QAA. On the other hand, the symmetry-adjusted estimates with QAA exhibits a more uniform error profile and overall improved noise suppression. Further investigation into this behavior is left as an open problem.

Bootstrapping uncertainty bars. To estimate uncertainty bars, we employ empirical bootstrapping [27], modified by batching together samples. First we summarize the original method: given T classical-shadow snapshots, one resamples that data T times with replacement. Then, averages $\hat{\delta}_j(T)$ (being either the unmitigated or mitigated estimators) are computed from that resampled data, yielding one bootstrap sample. Repeating this B times and computing the standard deviation among those B bootstrap samples yields the uncertainty bar.

Due to the size $T \sim 10^6$ – 10^7 from our simulations and limitations on classical compute resources, we perform bootstrapping on batches of snapshots. Split the T samples into K batches (each containing T/K samples) and compute $\hat{\delta}_j^{(k)}(T/K)$ for each batch $k = 1, \dots, K$. Because these estimates obey $\hat{\delta}_j(T) = (1/K) \sum_{k=1}^K \hat{\delta}_j^{(k)}(T/K)$, we resample the K batches (rather than all T shots) to bootstrap uncertainty bars for $\hat{\delta}_j(T)$. Depending on T , we set $K \sim 10^2$ – 10^3 , and for all cases we take $B = 200$.

Estimating the gate dependence of the QVM noise model. Here we provide an estimate of how much the QVM noise model violates Assumptions 1. We quantify this by computing a lower bound on the minimal observable error achievable by symmetry-adjusted classical shadows.

Let U_{prep} be the state-preparation circuit and U_g a random measurement circuit. For Schur's lemma to hold (Assumptions 1), we require that the entire noisy circuit take the form $\mathcal{E}U_gU_{\text{prep}}$, where \mathcal{E} is the both time- and g -independent. While the noise model that we simulate is indeed time stationary and Markovian, the effective error channel $\mathcal{E} = \mathcal{E}_g$ depends on g . (This can be seen, for example, by commuting all the individual gate-level errors throughout \tilde{U}_g and \tilde{U}_{prep} to the end of the circuit.)

In order to study this dependence on g , consider the decomposition

$$\mathcal{E}_g = \mathcal{E}_0 + \Delta_g, \quad (155)$$

where \mathcal{E}_0 is defined to be independent of $g \in G$. Although somewhat of an artificial decomposition, this is always mathematically possible with both \mathcal{E}_0 and Δ_g completely positive; indeed, a trivial choice is $\mathcal{E}_0 = 0$. Our goal is to find the ‘‘largest’’ (in some sense) valid solution for \mathcal{E}_0 . The remaining contribution Δ_g will then represent the minimal amount of assumption-violating noise in the model that our rigorous theory currently has no guarantees for.

From the decomposition above, the noisy measurement channel can be written as

$$\begin{aligned} \tilde{\mathcal{M}} &= \mathbb{E}_{g \sim G} U_g^\dagger \mathcal{M}_Z \mathcal{E}_g U_g \\ &= \tilde{\mathcal{M}}_0 + \bar{\Delta}, \end{aligned} \quad (156)$$

where

$$\begin{aligned} \tilde{\mathcal{M}}_0 &= \mathbb{E}_{g \sim G} U_g^\dagger \mathcal{M}_Z \mathcal{E}_0 U_g \\ &= \sum_{\lambda \in R_G} \tilde{f}_\lambda(\mathcal{E}_0) \Pi_\lambda \end{aligned} \quad (157)$$

is diagonal in the irreps of G , while the form of $\bar{\Delta} := \mathbb{E}_{g \sim G} U_g^\dagger \mathcal{M}_Z \Delta_g U_g$ is unknown.

Applying \mathcal{M}^{-1} and taking expectation values for the observables $\{O_j\}_{j=1}^L$ yields (assuming each $O_j \in V_\lambda$)

$$\begin{aligned} \langle\langle O_j | \mathcal{M}^{-1} \tilde{\mathcal{M}} | \rho \rangle\rangle &= \langle\langle O_j | \mathcal{M}^{-1} \tilde{\mathcal{M}}_0 | \rho \rangle\rangle + \langle\langle O_j | \mathcal{M}^{-1} \bar{\Delta} | \rho \rangle\rangle \\ &= \frac{\tilde{f}_\lambda(\mathcal{E}_0)}{f_\lambda} \langle\langle O_j | \rho \rangle\rangle + \delta_j. \end{aligned} \quad (158)$$

The terms $\delta_j := \langle\langle O_j | \mathcal{M}^{-1} \bar{\Delta} | \rho \rangle\rangle$ describe the deviation of observable estimates due to violations of the noise assumptions, which is precisely what we wish to quantify. For notation, denote the noisy expectations by $y_j := \langle\langle O_j | \mathcal{M}^{-1} \tilde{\mathcal{M}} | \rho \rangle\rangle$ and noiseless expectations by $x_j := \langle\langle O_j | \rho \rangle\rangle$. We collect these quantities into vectors of length L and define the diagonal matrix $A \in \mathbb{R}^{L \times L}$ with eigenvalues $\tilde{f}_\lambda(\mathcal{E}_0)/f_\lambda$ (in the appropriate positions corresponding to the irreps). This yields in the linear relationship

$$\boldsymbol{\delta} = \mathbf{y} - A\mathbf{x}. \quad (159)$$

This equation is underconstrained, so we opt for an estimate of $\boldsymbol{\delta}$ by bounding its norm from below. Namely, let \hat{A} be a diagonal matrix of free parameters $0 \leq \xi_\lambda \leq 1$, which we optimize by nonnegative least-squares (NNLS) minimization:

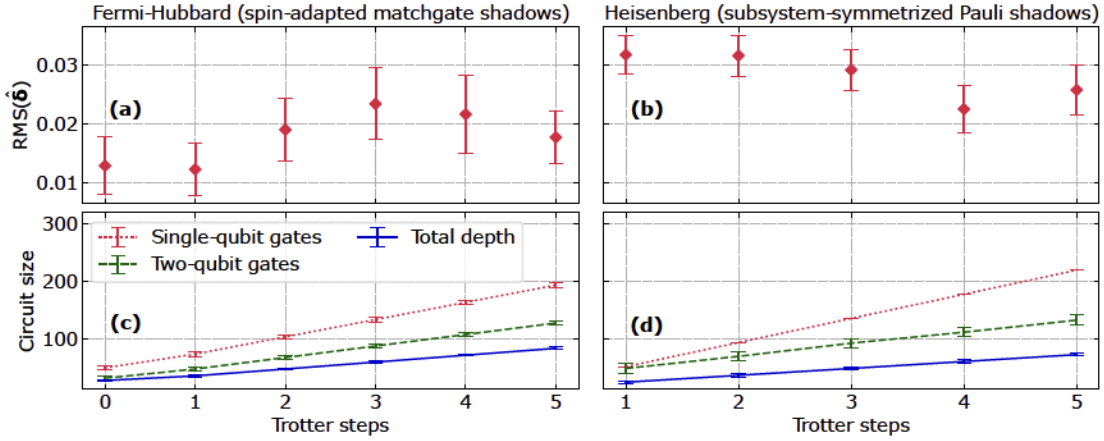
$$\|\boldsymbol{\delta}\|_2^2 \geq \min_{0 \leq \{\xi_\lambda\}_{\lambda \in R'} \leq 1} \|\mathbf{y} - \hat{A}\mathbf{x}\|_2^2. \quad (160)$$

Define $\hat{\boldsymbol{\delta}} := \mathbf{y} - \hat{A}\mathbf{x}$ as the solution to this problem. In this sense, $\hat{\boldsymbol{\delta}}$ represents an error floor beyond which our theory for symmetry adjustment cannot mitigate due to inherent violations of Assumptions 1.

In Supplementary Figure 3 we plot the root mean square of $\hat{\boldsymbol{\delta}}$,

$$\text{RMS}(\hat{\boldsymbol{\delta}}) := \frac{1}{\sqrt{L}} \|\hat{\boldsymbol{\delta}}\|_2, \quad (161)$$

which quantifies the average additive error of the estimates. The observables we choose constitute local operators depending on the type of system simulated. For fermions, we consider one- and two-body Majorana operators that respect the spin adaptation. For qubits, we take strictly two-body Pauli operators. Uncertainty bars are bootstrapped as described in ‘‘Bootstrapping uncertainty bars,’’ where each bootstrap sample is obtained from the NNLS solution



Supplementary Figure 3. Estimated deviation of the QVM noise from a gate-independent model. (a, b) Lower bound on the root-mean-square deviation of one- and two-body Majorana/Pauli expectation values as the number of Trotter steps (noisy circuit depth) grows. The estimated vector of deviations $\hat{\delta}$ is described in Eq. (160). Uncertainty bars are calculated by empirical bootstrapping. (c, d) Metrics for the size of the full circuit (state preparation via Trotterization and the random measurement unitary). We report the number of single-qubit (PhXZ) and two-qubit (\sqrt{i} SWAP) gates used, as well as the overall compiled circuit depth. Error bars denote one standard error of the mean.

of the resampled data. We also show data for the Trotter circuit size: the number of single- and two-qubit gates after compiling to the native gate set, as well as the circuit depth. Uncertainty bars here are given by one standard deviation in the size fluctuations due to the random unitaries U_g .

Overall, we assess that there is an error floor on the order of 10^{-2} per observable (recall that the observables have unit spectral norm). Interestingly, this lower bound appears roughly independent of circuit size (within uncertainty bars), perhaps indicating a saturation of the g -dependent contributions after a certain circuit size. In practice however, we have observed that symmetry-adjusted classical shadows only achieve mitigated errors on the order of 10^{-1} at the deepest circuits. We leave a closer analysis of this behavior, and whether this lower bound can actually be achieved, to future work.

SUPPLEMENTARY DISCUSSION

Review of prior symmetry-based QEM techniques. In this section we review prior work on techniques broadly known as symmetry verification (SV), introduced by Bonet-Monroig *et al.* [28] and McArdle *et al.* [29], and generalized by Cai [30]. Like symmetry-adjusted classical shadows, these approaches take advantage of inherent symmetries of the quantum system, and certain formulations also accomplish this in an offline manner. However, symmetry adjustment is a fundamentally different idea, and we will discuss this distinction here. Note that, in principle many QEM strategies (and indeed quantum error correction itself) can be understood as symmetry-based techniques wherein one artificially builds large amounts of symmetry into the system; we focus on techniques based on inherently possessed, physical symmetries here.

Symmetry verification. First we describe postprocessing symmetry verification (ppSV) [28]. Let M_S be the projector onto a symmetry subspace and ρ an ideal state such that $M_S \rho M_S = \rho$. Given the preparation of a noisy state $\tilde{\rho}$, ppSV aims to calculate the properties of the state

$$\rho^{\text{SV}} := \frac{M_S \tilde{\rho} M_S}{\text{tr}(M_S \tilde{\rho})} \quad (162)$$

in an offline manner. That is, consider an observable O which commutes with the symmetry. Then the error-mitigated estimate for $\langle O \rangle$ is

$$\text{tr}(O \rho^{\text{SV}}) = \frac{\sum_{i=1}^m \text{tr}(O S_i \tilde{\rho})}{\sum_{i=1}^m \text{tr}(S_i \tilde{\rho})}, \quad (163)$$

where we have written $M_S = \frac{1}{m} \sum_{i=1}^m S_i$ in terms of the symmetry stabilizers S_i . This quantity requires measuring the (noisy) expectation values of all S_i and $O S_i$ for $i = 1, \dots, m$. Furthermore, the sampling cost to maintain the

estimation accuracy ϵ is amplified by a factor of $\mathcal{O}(\text{tr}(M_S \tilde{\rho})^{-2})$. Bonet-Monroig *et al.* [28] show that this approach is equivalent to an instance of quantum subspace expansion [31], but with the ability to mitigate both coherent and incoherent errors outside the symmetry sector. Unfortunately, errors that commute with the symmetry cannot be projected out by this method. In general, SV projects $\tilde{\rho}$ to the closest symmetry-respecting state, with no guarantee of closeness to the ideal state ρ .

This idea can be generalized to symmetry expansion (SE) [30], which essentially replaces M_S with the symmetry expansion operator

$$E_S := \sum_{i=1}^m w_i S_i, \quad (164)$$

where w_i are nonnegative weights that sum to unity. The symmetry-expanded pseudostate $\rho^{\text{SE}} := E_S \tilde{\rho} / \text{tr}(E_S \tilde{\rho})$ is no longer guaranteed to lie in the desired symmetry sector, nor even to be a positive semidefinite operator. Hence SE exhibits estimation bias. However, by searching for weights that minimize this bias, one can heuristically achieve mitigated estimates with biased errors below that of the unmitigated noise level, while being simpler to implement than ppSV (for instance, by enforcing some of the $w_i = 0$).

Finally, direct symmetry verification (dSV) [28, 29] introduces additional quantum circuitry in order to check the symmetry value of the state throughout the course of the quantum computation. For example, in order to check the parity operator $S = Z_1 \cdots Z_n$, one can perform n CNOT gates, controlled on each system qubit and targeting an ancilla, followed by reading off the ancilla. This encodes the parity information into the ancilla, and if at any point the ancilla returns the incorrect parity value then that circuit run is therefore discarded. Other symmetries may require more complicated circuitry; for example, verifying the particle number η requires $\mathcal{O}(n \log \eta)$ CPhase gates in total to read off the binary representation of η from the ancilla. Although dSV requires significant additional coherent quantum control, the associated sampling overhead is $\mathcal{O}(\text{tr}(M_S \tilde{\rho})^{-1})$, a quadratic improvement over ppSV.

Distinction from symmetry-adjusted classical shadows. We now discuss the distinction between our method and these SV techniques. We will focus on ppSV, as it is the most comparable to symmetry-adjusted classical shadows. That is, rather than dSV which requires coherent detection of symmetry violations, both ppSV and symmetry-adjusted classical shadows use classical postprocessing as their primary mechanism for error mitigation.

The most significant conceptual difference is that symmetry adjustment does not necessarily project the noisy state into the desired symmetry subspace. Instead, the symmetry is used as a reference point to calibrate the effects of the noise on the shadow measurement channel \mathcal{M} . Recall that we aim to estimate the eigenvalues of the channel $\tilde{\mathcal{M}} = \sum_{\lambda} \tilde{f}_{\lambda} \Pi_{\lambda}$ via the relation

$$\tilde{f}_{\lambda} = f_{\lambda} \frac{\text{tr}(S_{\lambda} \tilde{\rho})}{s_{\lambda}}. \quad (165)$$

Thus, symmetry adjustment ultimately yields an estimate for the noisy channel, whose inverse is formally applied to produce the mitigated shadow

$$\hat{\rho}^{\text{EM}} = \sum_{\lambda} \frac{s_{\lambda} \Pi_{\lambda}(\hat{\rho})}{\text{tr}(S_{\lambda} \rho)}. \quad (166)$$

Here, $\hat{\rho}$ is the noisy shadow which converges to the noisy state $\tilde{\rho} = \mathbb{E}[\hat{\rho}]$, but in fact any state could in principle be used as long as the noise channel is the same. For instance, suppose we prepare a different state σ that experiences the same error channels that ρ does. (This may hold under models where both states are prepared by circuits of the same structure but different gate angles, such as in variational quantum algorithms.) Then Eq. (165) tells us that the \tilde{f}_{λ} learned by the measurements of ρ is equally applicable in the mitigation of properties of σ . Furthermore, the twirling nature of the randomized measurements allows the technique to account for errors that commute with the symmetry.

Beyond this conceptual difference, symmetry-adjusted classical shadows can easily incorporate symmetries such as particle number, whereas ppSV cannot. This is because the measurements of $\langle OS_i \rangle$ required of ppSV restrict the practical implementation to simple symmetries such as \mathbb{Z}_2 Pauli symmetries. More generally, it requires the decomposition of M_S into a tractable number of measurable terms. However, the projector into the η -particle sector is

$$M_S = \sum_{|x|=\eta} |x\rangle\langle x|, \quad (167)$$

where each $|x\rangle\langle x| = \prod_{j \in [n]} \frac{1}{2}(\mathbb{I} + (-1)^{x_j} Z_j)$ expands into exponentially many Pauli- Z terms. While all such terms mutually commute and can therefore be measured in the same basis, the upfront classical postprocessing cost of exactly computing this ppSV estimator is nonetheless exponential.

Additionally, as a classical-shadows protocol, our method is tailored to tackle the problem of multiple observable estimation. In contrast, all forms of SV can complicate the task, namely due to the inclusion of the additional OS_i operators. Recent work [32] has attempted to address this by straightforwardly combining ppSV with classical shadows. However, they showed limited success with this unification. Indeed, suppose we wish to project into a parity sector. The parity operator is n -body, while the target observables O_j are k -local for some small constant k . This means that each $O_j S$ has locality $n - k = \mathcal{O}(n)$, which leads to an exponentially large shadow norm (variance).

SUPPLEMENTARY REFERENCES

- [1] H.-Y. Huang, R. Kueng, and J. Preskill, Predicting many properties of a quantum system from very few measurements, *Nat. Phys.* **16**, 1050 (2020).
- [2] K. Wan, W. J. Huggins, J. Lee, and R. Babbush, Matchgate shadows for fermionic quantum simulation, *Commun. Math. Phys.* **404**, 629 (2023).
- [3] W. Fulton and J. Harris, *Representation Theory: A First Course* (Springer-Verlag, New York, 2004).
- [4] Z. Jiang, K. J. Sung, K. Kechedzhi, V. N. Smelyanskiy, and S. Boixo, Quantum algorithms to simulate many-body physics of correlated fermions, *Phys. Rev. Applied* **9**, 044036 (2018).
- [5] M. Reck, A. Zeilinger, H. J. Bernstein, and P. Bertani, Experimental realization of any discrete unitary operator, *Phys. Rev. Lett.* **73**, 58 (1994).
- [6] D. Wecker, M. B. Hastings, N. Wiebe, B. K. Clark, C. Nayak, and M. Troyer, Solving strongly correlated electron models on a quantum computer, *Phys. Rev. A* **92**, 062318 (2015).
- [7] I. D. Kivlichan *et al.*, Quantum simulation of electronic structure with linear depth and connectivity, *Phys. Rev. Lett.* **120**, 110501 (2018).
- [8] W. R. Clements, P. C. Humphreys, B. J. Metcalf, W. S. Kolthammer, and I. A. Walmsley, Optimal design for universal multiport interferometers, *Optica* **3**, 1460 (2016).
- [9] M. Oszmaniec, N. Dangniam, M. E. S. Morales, and Z. Zimborás, Fermion sampling: a robust quantum computational advantage scheme using fermionic linear optics and magic input states, *PRX Quantum* **3**, 020328 (2022).
- [10] W. J. Huggins *et al.*, Efficient and noise resilient measurements for quantum chemistry on near-term quantum computers, *npj Quantum Inf.* **7**, 23 (2021).
- [11] A. Zhao, N. C. Rubin, and A. Miyake, Fermionic partial tomography via classical shadows, *Phys. Rev. Lett.* **127**, 110504 (2021).
- [12] A. Chapman and A. Miyake, Classical simulation of quantum circuits by dynamical localization: Analytic results for Pauli-observable scrambling in time-dependent disorder, *Phys. Rev. A* **98**, 012309 (2018).
- [13] S. Chen, W. Yu, P. Zeng, and S. T. Flammia, Robust shadow estimation, *PRX Quantum* **2**, 030348 (2021).
- [14] Cirq Developers, *Cirq* (2023).
- [15] S. V. Isakov *et al.*, Simulations of quantum circuits with approximate noise using qsim and Cirq, [arXiv:2111.02396](https://arxiv.org/abs/2111.02396) (2021).
- [16] J. Wallman, C. Granade, R. Harper, and S. T. Flammia, Estimating the coherence of noise, *New J. Phys.* **17**, 113020 (2015).
- [17] G. Feng *et al.*, Estimating the coherence of noise in quantum control of a solid-state qubit, *Phys. Rev. Lett.* **117**, 260501 (2016).
- [18] E. Magesan, J. M. Gambetta, and J. Emerson, Scalable and robust randomized benchmarking of quantum processes, *Phys. Rev. Lett.* **106**, 180504 (2011).
- [19] E. Magesan, J. M. Gambetta, and J. Emerson, Characterizing quantum gates via randomized benchmarking, *Phys. Rev. A* **85**, 042311 (2012).
- [20] S. Boixo *et al.*, Characterizing quantum supremacy in near-term devices, *Nat. Phys.* **14**, 595 (2018).
- [21] C. Neill *et al.*, A blueprint for demonstrating quantum supremacy with superconducting qubits, *Science* **360**, 195 (2018).
- [22] F. Arute *et al.* (Google AI Quantum and collaborators), Quantum supremacy using a programmable superconducting processor, *Nature* **574**, 505 (2019).
- [23] C. Huang *et al.*, Quantum instruction set design for performance, *Phys. Rev. Lett.* **130**, 070601 (2023).
- [24] A. N. Habermann, Parallel neighbor-sort (or the glory of the induction principle), *Carnegie Mellon University Technical Report No. AD-759-248* (1972).
- [25] Z. Jiang *et al.* (Google AI Quantum and collaborators), Observation of separated dynamics of charge and spin in the Fermi-Hubbard model, [arXiv:2010.07965](https://arxiv.org/abs/2010.07965) (2020).
- [26] Quantum AI team and collaborators, *ReCirq* (2020).
- [27] B. Efron, Bootstrap methods: Another look at the jackknife, in *Breakthroughs in Statistics*, edited by S. Kotz and N. L. Johnson (Springer, New York, 1992) pp. 569–593.
- [28] X. Bonet-Monroig, R. Sagastizabal, M. Singh, and T. E. O’Brien, Low-cost error mitigation by symmetry verification, *Phys. Rev. A* **98**, 062339 (2018).
- [29] S. McArdle, X. Yuan, and S. Benjamin, Error-mitigated digital quantum simulation, *Phys. Rev. Lett.* **122**, 180501 (2019).

- [30] Z. Cai, Quantum error mitigation using symmetry expansion, [Quantum 5, 548 \(2021\)](#).
- [31] J. R. McClean, M. E. Kimchi-Schwartz, J. Carter, and W. A. de Jong, Hybrid quantum-classical hierarchy for mitigation of decoherence and determination of excited states, [Phys. Rev. A 95, 042308 \(2017\)](#).
- [32] H. Jnane, J. Steinberg, Z. Cai, H. C. Nguyen, and B. Koczor, Quantum error mitigated classical shadows, [PRX Quantum 5, 010324 \(2024\)](#).