

Virtual Distillation for Quantum Error Mitigation

William J. Huggins,^{1,2,*} Sam McArdle,^{1,3} Thomas E. O’Brien,^{1,4} Joonho Lee,⁵ Nicholas C. Rubin,¹ Sergio Boixo,¹ K. Birgitta Whaley,² Ryan Babbush,¹ and Jarrod R. McClean^{1,†}

¹Google Quantum AI, Venice, California 90291, USA

²Berkeley Quantum Information and Computation Center, Challenge Institute for Quantum Computation, and Department of Chemistry, University of California, Berkeley, California 94720, USA

³Department of Materials, University of Oxford, Parks Road, Oxford OX1 3PH, United Kingdom

⁴Instituut-Lorentz, Universiteit Leiden, 2300 RA Leiden, Netherlands

⁵Department of Chemistry, Columbia University, New York, New York 10027, USA



(Received 15 January 2021; revised 2 August 2021; accepted 21 September 2021; published 19 November 2021)

Contemporary quantum computers have relatively high levels of noise, making it difficult to use them to perform useful calculations, even with a large number of qubits. Quantum error correction is expected to eventually enable fault-tolerant quantum computation at large scales, but until then, it will be necessary to use alternative strategies to mitigate the impact of errors. We propose a near-term friendly strategy to mitigate errors by entangling and measuring M copies of a noisy state ρ . This enables us to estimate expectation values with respect to a state with dramatically reduced error $\rho^M/\text{Tr}(\rho^M)$ without explicitly preparing it, hence the name “virtual distillation.” As M increases, this state approaches the closest pure state to ρ exponentially quickly. We analyze the effectiveness of virtual distillation and find that it is governed in many regimes by the behavior of this pure state (corresponding to the dominant eigenvector of ρ). We numerically demonstrate that virtual distillation is capable of suppressing errors by multiple orders of magnitude and explain how this effect is enhanced as the system size grows. Finally, we show that this technique can improve the convergence of randomized quantum algorithms, even in the absence of device noise.

DOI: [10.1103/PhysRevX.11.041036](https://doi.org/10.1103/PhysRevX.11.041036)

Subject Areas: Quantum Physics, Quantum Information

I. INTRODUCTION

Performing meaningful calculations using near-term quantum computers is challenging because of the relatively high error rates of these devices. While quantum error correction promises to enable quantum computation with arbitrarily small levels of noise, the overhead required is too large to be currently practical [1,2]. The most plausible paths between today’s quantum computers and a fault-tolerant device assume a modest decrease in error rates together with a large increase in the number of qubits [2,3]. We find it interesting to ask if these additional qubits can be used fruitfully without employing the full machinery of fault tolerance. In this work, we explore an alternative to traditional quantum error correction that uses multiple

independently performed copies of a computation for error mitigation.

A variety of strategies exist to mitigate errors on noisy intermediate-scale quantum (NISQ) devices, i.e., to efficiently approximate the output that would be produced in the absence of noise. One class of approaches uses data collected at a variety of error rates to characterize the function relating the measured value of an observable to the error rate and extrapolate to the zero noise limit [4–6]. An alternative strategy proceeds by assuming a particular noise channel and expressing its inverse as a quasiprobability distribution over modified copies of the original circuit [4]. Other techniques work by comparing classically tractable simulations (tractable because they utilize a restricted set of gates) to evaluations of the same circuits on a noisy device [7–9]. These methods aim to learn enough about the impact of the noise to predict the noise-free expectation values for structurally similar circuits. In Ref. [10], O’Brien *et al.* put forward a version of a quantum phase estimation algorithm that achieves protection against errors by inverting the state preparation procedure and verifying that the system has returned to a reference state at the end of the computation. Besides these methods, more specific tools have been developed for ground-state calculations [11–13] for situations when the desired state

*Corresponding author.
whuggins@google.com

†Corresponding author.
jrmcclean@google.com

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International license](https://creativecommons.org/licenses/by/4.0/). Further distribution of this work must maintain attribution to the author(s) and the published article’s title, journal citation, and DOI.

possesses certain symmetries [14–18] and for treating errors during measurement [19–21].

Before the modern field of quantum error correction was developed, an alternative proposal was put forward for stabilizing quantum computations [22–24]. The essence of this approach is to execute M redundant copies of a computation in parallel and use a measurement to project into the symmetric subspace between these copies. Similar measurement primitives (measurements of the swap operator and its generalizations) have been applied to measure Renyi entanglement entropies and other polynomial functions of the density matrix [25–33]. One well-studied way to perform such a measurement is to use a Clebsch-Gordan or Schur transform to rotate to a basis which diagonalizes the swap operator [34]. In Ref. [35], Cotler *et al.* built on these approaches to implement an idea they call “virtual cooling.” By performing a joint measurement on M copies of a thermal state at inverse temperature β ($\rho \propto e^{-\beta H}$), they were able to estimate expectation values with respect to the thermal state at inverse temperature $M\beta$ ($\rho^M \propto e^{-M\beta H}$). In this paper, we apply the same kind of measurement techniques to the problem of mitigating errors in a noisy quantum computation.

Earlier work on using symmetrization to stabilize a noisy quantum computation focused on protocols that prepared an approximately purified state [22–24]. We abandon this goal and instead aim to reconstruct expectation values with respect to an approximately purified state without explicitly preparing it. We refer to this approach as virtual distillation, using the word “virtual” to emphasize that we do not actually prepare a purified version of the state like a typical distillation scheme would [36–38]. To be specific, we use collective measurements of M copies of ρ to measure expectation values with respect to the state

$$\frac{\rho^M}{\text{Tr}(\rho^M)} = \frac{\sum_i p_i^M |i\rangle\langle i|}{\sum_i p_i^M}, \quad (1)$$

where $\rho = \sum_i p_i |i\rangle\langle i|$ is a spectral decomposition of ρ . Under this approach, the relative weights of the non-dominant eigenvectors are suppressed exponentially in M . This represents an improvement over approaches which demand that the approximately purified state is prepared explicitly, which achieve a suppression that is merely linear in M in the general case [22–24,39].

Our proposed error mitigation technique offers the opportunity to make use of additional qubits to enhance the quality of a noisy computation without the large overhead of traditional quantum error correction. Furthermore, the technique is simple to use and analyze. If we neglect the errors that occur during measurement, it is straightforward to obtain analytic expressions for the states whose expectation values we effectively measure and for the variance of the resulting estimator. In the limit where the level of noise is small, the number of additional

measurements required by our approach goes to zero. Our error mitigation strategy, as we show, is capable of reducing the impact of stochastic errors arising from noise on a near-term device as well as stochastic errors inherent to randomized quantum algorithms implemented on an error-free device.

We begin in Sec. II by introducing the theoretical formalism of virtual distillation and presenting its simplest implementation. We continue in Sec. II B with an analysis of the sample complexity of the simple version of this technique along with a proof that there exist more efficient generalizations under certain circumstances. In Sec. III, we study the error mitigation performance of virtual distillation analytically by splitting the effect of errors into two components. We treat the shift of the leading eigenvector of the density matrix away from the target (error-free) state perturbatively (Sec. III B) and the shift of the noisy density matrix away from its dominant eigenvector using a phenomenological model of errors (Sec. III A). Although the second effect may be exponentially suppressed by increasing the number of states (M), the same is not true for the first effect, which in the worst case limits the performance of virtual distillation to providing only a constant-factor improvement in the error rate (as a function of the underlying physical noise rate). For purely coherent errors, this first effect is the only consideration, and virtual distillation offers no protection. To complement this analysis, in Sec. IV we present numerical simulations of virtual distillation applied to various noisy quantum circuits. We observe here that for some range of noise levels, virtual distillation achieves a rate of error suppression exceeding the bounds suggested in Sec. III B. Finally, in Sec. V, we consider the performance of our technique when applied to the stochastic errors that arise during randomized algorithms for real-time evolution.

II. THEORY

Virtual distillation is a protocol for using collective measurements of M copies of a state ρ to suppress incoherent errors by measuring expectation values with respect to the state $\rho^M/\text{Tr}(\rho^M)$. Virtual distillation approximates the error-free expectation value of O as

$$\langle O \rangle_{\text{corrected}} := \frac{\text{Tr}(O\rho^M)}{\text{Tr}(\rho^M)}. \quad (2)$$

The resulting estimator converges exponentially quickly toward the closest pure state to ρ as M is increased. In this section, we lay out the basic theory behind virtual distillation. We present the simplest implementation in Sec. II A and an analysis of the measurement overhead in Sec. II B. In Algorithm 1 below, we present a pseudocode for the implementation discussed in more detail in Sec. II A.

We begin by establishing some assumptions and notation. Throughout this paper, we deal with operations that

Algorithm 1. Virtual distillation, basic implementation (see Sec. II A).

Input: A number of measurement repetitions K , $2K$ copies of the N qubit state ρ (provided two at a time).
Output: An error-mitigated estimate of $\langle Z_i \rangle$ for each qubit in ρ ; $\langle Z_i \rangle_{\text{corrected}} \approx \{[\text{Tr}(Z_i \rho^2)]/[\text{Tr}(\rho^2)]\}$.
 Set $E_i = 0$ for each qubit $i \in 1, \dots, N$.
 Set $D = 0$.
for $k \in 1, \dots, K$ **do**
 Perform any SWAP operations necessary to make it possible to couple each qubit in the first copy of ρ with the corresponding qubit in the second copy.
 Apply the two-qubit gate $B_i^{(2)}$ [defined below in Eq. (10) of Sec. II A] between each qubit i in the first copy and the corresponding qubit in the second copy.
 Measure both states in the computational basis.
 Let z_i^1 and z_i^2 denote the measurement outcomes for the i th qubits in the first and second copies of ρ , respectively.
 for $i \in 1, \dots, N$ **do**
 $E_{i+} = [1/(2^N)](z_i^1 + z_i^2) \prod_{j \neq i} 1 + z_j^1 - z_j^2 + z_j^1 z_j^2$
 end for
 $D + = [1/(2^N)] \prod_{j=1}^N 1 + z_j^1 - z_j^2 + z_j^1 z_j^2$
end for
return $\{\langle Z_i \rangle_{\text{corrected}} := (E_i/D)\}$

act on multiple copies of the same state. We make the assumption that the noise experienced by the separate copies has the same form and strength. If we relax this assumption, then we still measure an effective state that corresponds to the product of the density matrices of the individual copies so long as the copies are not entangled prior to virtual distillation. We briefly explore this more general situation in the Appendix H.

We use the letter N to indicate the number of qubits in an individual system and the letter M to indicate the number of copies (which we sometimes refer to as subsystems). Superscripts with parentheses indicate an operator that acts on multiple systems. For example, we denote the cyclic shift operator between M copies by $S^{(M)}$. We use bolded superscripts without parentheses to denote which copy an operator acts on; e.g., $O^{\mathbf{1}}$ indicates the operator O acting on subsystem 1. We use superscripts without a boldfaced font or parentheses to indicate exponentiation as usual. Subscripts are used in two different ways. Subscripts on an operator generally indicate which qubit within a system the operator acts on. The exception is when the subscript is being used more generically as an index in a summation, which should always be clear from the context and the presence of the \sum symbol.

In order to evaluate the numerator and denominator of Eq. (1), we can make use of the following equality [26,27,35]:

$$\text{Tr}(O\rho^M) = \text{Tr}(O^{\mathbf{i}}S^{(M)}\rho^{\otimes M}). \quad (3)$$

Here, $O^{\mathbf{i}}$ indicates the observable O acting on (an arbitrary) subsystem i , and $S^{(M)}$ indicates the cyclic shift operator on M systems, i.e.,

$$O^{\mathbf{i}} := \mathbb{I} \otimes \mathbb{I} \cdots O \cdots \mathbb{I},$$

$$S^{(M)}|\psi_1\rangle \otimes |\psi_2\rangle \cdots |\psi_M\rangle := |\psi_2\rangle \otimes |\psi_3\rangle \cdots |\psi_1\rangle. \quad (4)$$

This identity can be proven by expanding the right-hand side, carefully keeping track of the indices. Without loss of generality, we choose $i = 1$ yielding

$$\begin{aligned} \text{Tr}(O^{\mathbf{1}}S^{(M)}\rho^{\otimes M}) &= \sum_{i_1, i_2, \dots, i_M, j_1, j_2, \dots, j_M, k} O_{k, j_1} \delta_{j_2, i_1} \cdots \delta_{j_1, i_M} \rho_{i_1, k} \cdots \rho_{i_M, j_M} \\ &= \sum_{i_1, i_2, \dots, i_M, k} \rho_{i_1, k} O_{k, i_M} \rho_{i_M, i_{M-1}} \cdots \rho_{i_2, i_1} \\ &= \text{Tr}(O\rho^M). \end{aligned} \quad (5)$$

In Fig. 1, we present a diagrammatic representation of Eq. (3) for the case where $M = 3$ (note that we commute $\rho^{\otimes 3}$ with $S^{(3)}$ in the diagram).




FIG. 1. A diagrammatic representation of Eq. (3) with $M = 3$ and $i = 1$ using tensor network notation [40–42]. The blue square represents the operator $O^{\mathbf{1}}$, each red circle represents a copy of the state ρ , and the connections between the shapes indicate indices which are summed over. The cyclic shift operator $S^{(3)}$ is naturally represented as a product of two swap operators, which are themselves indicated by the crossed wires. Note that the top diagram actually corresponds to the expression $\text{Tr}(O^{\mathbf{1}}\rho^{\otimes 3}S^{(3)})$; we commute $\rho^{\otimes 3}$ with $S^{(3)}$ before producing the figure. Rearranging the wires to yield the bottom diagram is equivalent to the simplification of the summation in Eq. (5).




FIG. 2. A flowchart that describes the choices involved in selecting between the different variants of virtual distillation presented in this work. Blue boxes denote questions for the experimentalist to answer about the available quantum resources and problem to be studied; green boxes link to the relevant sections in the text and briefly summarize the main features of each variant. The flowchart provides direction to the most flexible variant given the answers provided in the blue boxes, but the actual experimental performance depends on many factors.

The quantities in the numerator and denominator of Eq. (2) can be evaluated in a number of different ways. For simplicity, we focus our presentation on one such approach in Sec. II A. In that section, we roughly follow the work of Ref. [35], except that we use the language of qubits rather than bosonic systems. We discuss a variety of alternative protocols in Appendixes A–C. Figure 2 summarizes the differences between these variants. The practical utility of these techniques as error mitigation tools will be partly determined by the number of samples necessary to evaluate the corrected expectation values to within some target precision ϵ . We address this issue in Sec. II B and also show that there exist generalizations of our approach that can further reduce the number of circuit repetitions for a desired precision.

A. Measurement by diagonalization

In this section, we present a straightforward strategy applicable when the operator O is the Pauli Z operator




FIG. 3. A circuit diagram of our approach applied to a six-qubit circuit. We use twice the number of qubits to independently perform two copies of the original circuit. We then apply a single layer of the two-qubit gates specified in Eq. (10) before measuring each qubit in the computational basis. This allows us to estimate the error-mitigated expectation values for all single-site Z operators.

acting on a single qubit and $M = 2$. Other single-qubit observables can be accessed by applying the appropriate single-qubit rotations before the virtual distillation procedure. This realization of our error mitigation technique requires only a single additional layer of two-qubit gates followed by measurement in the computational basis. We present a schematic of this approach in Fig. 3.

Rather than using the relation in Eq. (3) directly, we instead define a symmetrized version of our observable,

$$O^{(M)} = \frac{1}{M} \sum_{i=1}^M O^i. \quad (6)$$

For the specific case we consider here, that means we take

$$Z_k^{(2)} = \frac{1}{2} (Z_k^1 + Z_k^2). \quad (7)$$

It is straightforward to use Eq. (3) to show that

$$\frac{\text{Tr}(O\rho^M)}{\text{Tr}(\rho^M)} = \frac{\text{Tr}(O^{(M)}S^{(M)}\rho^{\otimes M})}{\text{Tr}(S^{(M)}\rho^{\otimes M})}. \quad (8)$$

Using the symmetrized observable is advantageous because

$$[O^{(M)}, S^{(M)}] = 0, \quad (9)$$

or, in our case, $[Z_k^{(2)}, S^{(2)}] = 0$.

Both $S^{(2)}$ and $Z_k^{(2)}$ factorize into tensor products of operators that act separately on each pair of qubits, where

the i th pair consists of the i th qubit from each system. Therefore, we may simultaneously diagonalize $S^{(2)}$ and $Z_k^{(2)}S^{(2)}$ using an operator that factorizes with the same structure. We denote the two-qubit unitary that performs this diagonalization on the i th pair $B_i^{(2)}$. We give a matrix representation for this gate below, noting that there is some freedom in the choice of phases for the matrix elements,

$$B_i^{(2)} := \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \frac{\sqrt{2}}{2} & -\frac{\sqrt{2}}{2} & 0 \\ 0 & \frac{\sqrt{2}}{2} & \frac{\sqrt{2}}{2} & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}. \quad (10)$$

We then define

$$B^{(2)} := \bigotimes_{i=1}^M B_i^{(2)}. \quad (11)$$

As desired, this unitary diagonalized the individual factors that make up the observables,

$$B^{(2)}S_i^{(2)}B^{(2)\dagger} \rightarrow \frac{1}{2}(1 + Z_i^1 - Z_i^2 + Z_i^1Z_i^2), \quad (12)$$

$$B^{(2)}Z_k^{(2)}S_k^{(2)}B^{(2)\dagger} \rightarrow \frac{1}{2}(Z_k^1 + Z_k^2). \quad (13)$$

This diagonalization is particularly easy to implement when each qubit from the first copy of ρ is adjacent to the corresponding qubit from the second copy. The procedure for measuring the observables required to estimate the numerator and denominator of Eq. (8) then reduces to applying a single layer of N two-qubit gates in parallel and measuring in the computational basis. In fact, because $B^{(2)}$ diagonalizes $Z_k^{(2)}S^{(2)}$ for all N values of k , we naturally collect the data required to estimate the error-mitigated expectation values for all N of the operators Z_k simultaneously. By applying the appropriate single-qubit rotations before performing virtual distillation, we can instead access an arbitrary single-qubit observable on each qubit. We capture this process diagrammatically in Fig. 3.

In order to develop some intuition, it is helpful to express $\rho^{\otimes 2}$ using a spectral decomposition of ρ and consider two separate components of the resulting sum,

$$\begin{aligned} \rho^{\otimes 2} &= \sum_{ij} p_i p_j |i\rangle\langle i| \otimes |j\rangle\langle j| \\ &= \sum_i p_i^2 |i\rangle\langle i| \otimes |i\rangle\langle i| + \sum_{i \neq j} p_i p_j |i\rangle\langle i| \otimes |j\rangle\langle j|. \end{aligned} \quad (14)$$

The calculation of measurement probabilities and expectation values is a linear operation on the density matrix; we can therefore consider these two components separately.

The component of the state with $i = j$ is in the $+1$ eigenspace of $S^{(2)}$ and leads to measurements of $S^{(2)}$ which yield the $+1$ eigenvalue with probability $p = \sum_i p_i^2 = \text{Tr}(\rho^2)$. In the case where $i \neq j$, $|i\rangle\langle i| \otimes |j\rangle\langle j|$ is an even superposition of symmetric and antisymmetric states,

$$|i\rangle|j\rangle = \frac{1}{2}(|i\rangle|j\rangle + |j\rangle|i\rangle) + \frac{1}{2}(|i\rangle|j\rangle - |j\rangle|i\rangle). \quad (15)$$

For this component of the state, measurements of $S^{(2)}$ yield $+1$ and -1 with equal probability and $\langle S^{(2)} \rangle = 0$. Combining these two cases, we have the expected equality, $\text{Tr}(S^{(2)}\rho^{\otimes 2}) = \text{Tr}(\rho^2)$. Measurements of $S^{(2)}O^{(2)}$ follow a similar pattern.

We find it interesting to contrast this behavior with the stabilizer theory of quantum error correction. In the stabilizer formalism, errors are detected by projecting through measurement into the -1 eigenspace of one or more symmetries. In our approach, we instead rely on errors being equally supported on the eigenspaces of the symmetry we measure.

B. Sample efficiency

The number of circuit repetitions required to determine the error-mitigated expectation values within a precision ϵ depends on the variance of our estimator. In this section, we present expressions for this variance. We focus on the $M = 2$ case and the methods discussed in Sec. II A. The calculations are also applicable to the variant protocol we present in Appendix C. We also show how there exists an extension to our protocol that makes more efficient use of multiple copies when the noise level is sufficiently high.

We like to determine the variance of our estimator for the error-mitigated expectation value

$$\langle O \rangle_{\text{corrected}} = \frac{\text{Tr}(O^{(2)}S^{(2)}\rho^{\otimes 2})}{\text{Tr}(S^{(2)}\rho^{\otimes 2})}. \quad (16)$$

We leave the derivation to Appendix D and simply give an (approximate) expression for the variance,

$$\begin{aligned} \text{Var}(\langle O \rangle_{\text{corrected}}) &\approx \frac{1}{R} \left[\frac{1}{\text{Tr}(\rho^2)^2} \left(\frac{1}{2} \text{Tr}(\rho O^2) + \frac{1}{2} \text{Tr}(\rho O)^2 - \text{Tr}(\rho^2 O^2) \right) \right. \\ &\quad - 2 \frac{\text{Tr}(\rho^2 O)}{\text{Tr}(\rho^2)^3} [\text{Tr}(\rho O) - \text{Tr}(\rho^2 O)\text{Tr}(\rho^2)] \\ &\quad \left. + \frac{\text{Tr}(\rho^2 O)^2}{\text{Tr}(\rho^2)^4} [1 - \text{Tr}(\rho^2)^2] \right], \end{aligned} \quad (17)$$

where R refers to the number of measurement repetitions. It is useful to consider what happens in the limit where ρ is a pure state. In that case, the second and third lines are zero, and the variance reduces to

$$\text{Var}(\langle O \rangle_{\text{corrected}}) = \frac{1}{2R} [\text{Tr}(\rho O^2) - \text{Tr}(\rho O)^2], \quad (18)$$

exactly what one would expect when averaging $2R$ -independent measurements of O . As the purity of ρ decreases, the variance and the number of circuit repetitions increases.

The rest of this section focuses on laying the groundwork to improve the sample efficiency of these techniques. This is an important goal because the number of samples required can grow large given sufficiently noisy circuits. At high enough error rates, we are highly likely to find ourselves in a situation where

$$\text{Tr}(\rho^3) \ll \text{Tr}(\rho^2) \ll 1. \quad (19)$$

We now make the assumption that the level of error mitigation offered by measuring ρ^M is sufficient, but we have $2K \gg M$ copies of ρ available. For simplicity, we focus on the case where $M = 2$, and O is a Pauli operator acting on one or more qubits. We present a generalization of our approach involving a collective measurement of all $2K$ copies of ρ that performs better than a naive parallelization.

The naive approach we hope to beat consists of taking K pairs and running the protocol described above in parallel, averaging the results. For simplicity, we focus on the variance of our estimator for the quantity that appears in the numerator of Eq. (8) rather than the ratio itself. In Appendix D, we show that the variance of our estimator for $S^{(2)}O^{(2)}$ is $\frac{1}{2}\text{Tr}(\rho O^2) + \frac{1}{2}\text{Tr}(\rho O)^2 - \text{Tr}(\rho^2 O)^2$. Therefore, the variance obtained when using $2K$ copies in parallel is exactly

$$\begin{aligned} \text{Var}(\langle S^{(2)}O^{(2)} \rangle) \\ = \frac{1}{2K} [\text{Tr}(\rho O^2) + \text{Tr}(\rho O)^2 - 2\text{Tr}(\rho^2 O)^2]. \end{aligned} \quad (20)$$

We prove below that it is possible in some situations to obtain a more sample-efficient estimator for the corrected expectation value by performing a joint measurement on all $2K$ copies. We do so by providing an operator \tilde{O} with the desired expectation value and calculating its variance.

First, we define the operator

$$\tilde{O} = \frac{1}{\binom{2K}{2}} \sum_{i=1}^{2K} \sum_{j>i} \frac{1}{2} (O^i + O^j) S^{(i,j)}, \quad (21)$$

where we use $S^{(i,j)}$ to denote the swap operator specifically between subsystems i and j . It is simple to show that

$$\text{Tr}(\tilde{O} \rho^{\otimes 2K}) = \text{Tr}(O \rho^2). \quad (22)$$

We compute the variance of \tilde{O} with respect to the state $\rho^{\otimes 2K}$ in Appendix E, finding that

$$\text{Var}(\langle \tilde{O} \rangle) \leq \frac{1 + 7(K-1)\text{Tr}(\rho^3)}{K(2K-1)}. \quad (23)$$

When $\text{Tr}(\rho^3)$ is small, the second term in Eq. (23) is suppressed, and there is a regime where the variance of this operator shrinks quadratically with K . The naive approach, where we perform K -independent calculations on separate pairs results in an estimator whose variance is suppressed only linearly in K . We do not suggest a particular strategy, let alone one that is NISQ friendly, for implementing the measurement of \tilde{O} . We hope that future work can address this issue. Furthermore, while we establish that generalizations of the simplest virtual distillation procedure can outperform a naive parallel strategy, we do not establish a comprehensive theory on the limitations of virtual distillation. It would be useful to quantify the minimum number of samples required to resolve $\text{Tr}(O \rho^M)$ given access to a large number of copies of ρ under various assumptions about the spectrum of the density matrix.

III. PERFORMANCE UNDER DIFFERENT NOISE MODELS

In the numerical studies, we present evidence that the performance of virtual distillation can be essentially predicted by the combination of two contributions. Here we find it instructive to consider them separately using simple analytical models. To understand the potential benefit of our approach using the minimal setup, we consider the fidelity of

$$\rho_{\text{corrected}} := \frac{\rho^2}{\text{Tr}(\rho^2)} \quad (24)$$

with the ideal state generated by noiseless evolution (neglecting error introduced by the measurement procedure). We first consider the performance under noise that maps the ideal state to states orthogonal to it, leaving the dominant eigenvector of the density matrix as the ideal state. We then turn toward the effect of errors that lead to states nonorthogonal to the ideal state, causing a drift in the dominant eigenvector of the density matrix. The essential behavior of virtual distillation is to remove errors of the first kind rapidly, while converging to a floor determined by the drift in the dominant eigenvector that enables a large constant-factor improvement over the erred state.

A. Orthogonal errors

We first consider idealized errors that leave the dominant eigenvector as the ideal state. We consider a phenomenological error model motivated by the assumption that we can think of errors as discrete events that occur locally in space and time with some probability. For simplicity, we model every gate as a stochastic quantum map where with probability p an error occurs, and we assume that every

new error sends the quantum evolution to a new orthogonal state. The resulting density matrix for a circuit with G gates is

$$\begin{aligned} \rho &= (1-p)^G \rho_0 + (1-p)^{G-1} p \sum_{j=1}^G \rho_{j_1} \\ &+ (1-p)^{G-2} p^2 \sum_{j_1 \neq j_2} \rho_{j_1, j_2} \\ &+ (1-p)^{G-3} p^3 \sum_{j_1 \neq j_2 \neq j_3} \rho_{j_1, j_2, j_3} + \dots \end{aligned} \quad (25)$$

The operator for ρ^2 is similar to all the coefficients squared, as all the states are assumed to be orthogonal. Therefore,

$$\text{Tr}(\rho^2) = [(1-p)^2 + p^2]^G. \quad (26)$$

The fidelity with the ideal state ρ_0 is

$$\frac{\text{Tr}(\rho_0 \rho^2)}{\text{Tr}(\rho^2)} = \frac{(1-p)^{2G}}{[(1-p)^2 + p^2]^G} \quad (27)$$

$$\simeq 1 - Gp^2 + O(Gp^3). \quad (28)$$

Therefore, we expect a quadratic suppression of errors in the most favorable case.

The result is similar in the case of M copies:

$$\frac{\text{Tr}(\rho_0 \rho^M)}{\text{Tr}(\rho^M)} = \frac{(1-p)^{MG}}{[(1-p)^M + p^M]^G} \quad (29)$$

$$\simeq 1 - Gp^M + O(Gp^{M+1}). \quad (30)$$

The other factor affecting the performance of virtual distillation besides the fidelity is the sample complexity. We analyze the general case in more detail in Appendix D, but it is instructive to briefly consider the performance under this simplified model of errors. For simplicity, we assume that we aim to measure the error-mitigated expectation value of a Pauli operator O at the $M = 2$ level using R -independent experiments to estimate the numerator and denominator of Eq. (16) (for a total of $2R$ experiments). Then the variance of our estimators for the numerator and denominator are upper bounded by 1, and we have

$$\text{Var}(\langle O \rangle_{\text{corrected}}) \lesssim \frac{1}{R} \left(\frac{1}{\text{Tr}(\rho^2)^2} + \frac{\text{Tr}(O\rho^2)^2}{\text{Tr}(\rho^2)^4} \right). \quad (31)$$

Because of our assumption that O is a Pauli operator, and therefore $\|O\| = 1$, we have $\text{Tr}(O\rho^2) \leq \text{Tr}(\rho^2)$, implying $\text{Tr}(O\rho^2)^2 \leq \text{Tr}(\rho^2)^2$. Therefore,

$$\text{Var}(\langle O \rangle_{\text{corrected}}) \lesssim \frac{2}{R\text{Tr}(\rho^2)^2} = \frac{2}{R((1-p)^2 + p^2)^{2G}}. \quad (32)$$

When p is small, we can neglect the p^2 term in the denominator. Therefore, taking

$$R \lesssim (1-p)^{-4G} \quad (33)$$

is sufficient to estimate $\langle O_{\text{corrected}} \rangle$ to within a fixed additive error.

B. Nonorthogonal error floor

The analysis of the previous section made the simplifying assumption that the dominant eigenvector of the density matrix $\rho_0 = |0\rangle\langle 0|$ corresponds exactly to the ideal state generated by noiseless evolution. In practice, errors will lead to population in states that may not be orthogonal to the target state, leading to a drift in the dominant eigenvector of the density matrix. We see in our numerical studies that this drift limits the maximum potential upside of virtual distillation. In this second, we develop an understanding of this drift by using perturbation theory to consider the first-order change in the dominant eigenvector of the density matrix.

Let us consider a state ρ in the middle of a noisy preparation circuit, allowing for ρ to already be somewhat distorted by noise. Writing ρ in its eigenbasis, we have

$$\rho = \sum_i \lambda_i |i\rangle\langle i|, \quad (34)$$

where we order the eigenvalues in descending order. Note that we use the symbol λ_i for the i th eigenvector of the density matrix rather than p_i throughout this section, reserving the symbol p for the coefficients associated with a Kraus operator decomposition of our noise channel. We wish to consider the impact of a subsequent noise channel defined in terms of a set of Kraus operators,

$$\rho \rightarrow p_0 \rho + \sum_{j \neq 0} p_j K_j \rho K_j^\dagger. \quad (35)$$

Note that we demand a representation of the channel where K_0 is the identity matrix in order to simplify our analysis. Now let ΔV denote the change in the density matrix induced by this channel ($\rho \rightarrow \rho + \Delta V$),

$$\Delta V := (p_0 - 1)\rho + \sum_{j \neq 0} p_j K_j \rho K_j^\dagger, \quad (36)$$

where we define the scale Δ by taking $\|\Delta V\|$ to be $O(1)$.

Now we make the assumption that we are in the low-error regime. Specifically, we assume that $\lambda_0 \gg \lambda_1$ and that $\Delta \ll |\lambda_0 - \lambda_1|$. Under this assumption, we satisfy the necessary conditions for applying matrix perturbation theory to the dominant eigenvector [43]. We can therefore proceed by expressing the dominant eigenvector of $\rho + \Delta V$ as a convergent power series in Δ . This yields

$$|0\rangle = |0^{(0)}\rangle + \Delta|0^{(1)}\rangle + \Delta^2|0^{(2)}\rangle + \mathcal{O}(\Delta^3), \quad (37)$$

where $|0\rangle$ denotes the dominant eigenvector of $\rho + \Delta V$, $|0^{(0)}\rangle$ denotes the dominant eigenvector of the unperturbed ρ , and $|0^{(i)}\rangle$ denotes the correction at i th order. Likewise, we can also express the eigenvalue corresponding to the dominant eigenvector as a power series in Δ ,

$$\lambda_0 = \lambda_0^{(0)} + \Delta\lambda_0^{(1)} + \Delta^2\lambda_0^{(2)} + \mathcal{O}(\Delta^3). \quad (38)$$

We can then proceed in the usual way, expanding the eigenvalue equation,

$$(\rho + \Delta V)|0\rangle = \lambda_0|0\rangle, \quad (39)$$

and equating terms order by order. This leads to a familiar expression for the first-order correction to the dominant eigenvector in terms of the zeroth-order eigenvalues and eigenvectors,

$$|0^{(1)}\rangle = \sum_{i \neq 0} \frac{\langle i^{(0)}|V|0^{(0)}\rangle}{\lambda_0^{(0)} - \lambda_i^{(0)}} |i^{(0)}\rangle. \quad (40)$$

At this point, it is useful to carefully consider the normalization of $|0\rangle$. Let $|D\rangle$ denote the normalized form of $|0\rangle$,

$$\begin{aligned} |D\rangle &:= \frac{|0^{(0)}\rangle + \Delta|0^{(1)}\rangle + \Delta^2|0^{(2)}\rangle + \mathcal{O}(\Delta^3)}{\sqrt{1 + \Delta^2\langle 0^{(1)}|0^{(1)}\rangle + \mathcal{O}(\Delta^3)}} \\ &= |0^{(0)}\rangle + \Delta|0^{(1)}\rangle \\ &\quad + \Delta^2|0^{(2)}\rangle - \frac{\Delta^2}{2}\langle 0^{(1)}|0^{(1)}\rangle|0^{(0)}\rangle + \mathcal{O}(\Delta^3), \end{aligned} \quad (41)$$

where we make use of the fact that the first- and second-order corrections are both orthogonal to the unperturbed eigenvector.

We can now compute the trace distance between $|D\rangle$ and the dominant eigenvector of the unperturbed state,

$$\begin{aligned} T(|D\rangle, |0^{(0)}\rangle) &= \frac{1}{2} \text{Tr} \left(\sqrt{[\Delta|0^{(0)}\rangle\langle 0^{(1)}| + \Delta|0^{(1)}\rangle\langle 0^{(0)}| + \Delta^2|0^{(0)}\rangle\langle 0^{(2)}| + \Delta^2|0^{(2)}\rangle\langle 0^{(0)}| - \Delta^2\langle 0^{(1)}|0^{(1)}\rangle|0^{(0)}\rangle\langle 0^{(0)}| + \mathcal{O}(\Delta^3)]^2} \right) \\ &= \frac{1}{2} \text{Tr} \left(\sqrt{\Delta^2\langle 0^{(1)}|0^{(1)}\rangle|0^{(0)}\rangle\langle 0^{(0)}| + \Delta^2|0^{(1)}\rangle\langle 0^{(1)}| + \mathcal{O}(\Delta^3)} \right) \\ &= \frac{1}{2} \text{Tr} \left(\sqrt{\Delta^2\langle 0^{(1)}|0^{(1)}\rangle|0^{(0)}\rangle\langle 0^{(0)}| + \Delta^2\langle 0^{(1)}|0^{(1)}\rangle \frac{|0^{(1)}\rangle\langle 0^{(1)}|}{\langle 0^{(1)}|0^{(1)}\rangle} + \mathcal{O}(\Delta^3)} \right) \\ &= \Delta \sqrt{\frac{\langle 0^{(1)}|0^{(1)}\rangle}{2}} + \mathcal{O}(\Delta^2). \end{aligned} \quad (42)$$

Now let us expand $\langle 0^{(1)}|0^{(1)}\rangle$ in terms of the Kraus operators of our noise model,

$$\begin{aligned} \langle 0^{(1)}|0^{(1)}\rangle &= \sum_{i \neq 0} \frac{1}{(\lambda_0 - \lambda_i)^2} \langle 0|V^\dagger|i\rangle\langle i|V|0\rangle \\ &= \sum_{i \neq 0} \frac{1}{(\lambda_0 - \lambda_i)^2} \\ &\quad \times \left| \langle i| \left((p_0 - 1)\mathbb{I} + \sum_{j \neq 0} p_j K_j |0\rangle\langle 0|K_j^\dagger \right) |0\rangle \right|^2 \\ &= \sum_{i \neq 0} \frac{1}{(\lambda_0 - \lambda_i)^2} \left| \sum_{j \neq 0} p_j \langle i|K_j|0\rangle\langle 0|K_j^\dagger|0\rangle \right|^2, \end{aligned} \quad (43)$$

where we omit the (0) superscripts of the eigenvalues and eigenvectors on the right-hand side for readability.

We can see that, in the general case, we expect a nonzero contribution to the trace distance at first order in Δ . Because $\rho^2/\text{Tr}(\rho^2) \approx |D\rangle\langle D|$ in the low-noise regime, this will effectively set a floor for how well our method can correct errors. Therefore, without further constraints on the state,

the noise model, or the observables being measured, our method will not achieve a quadratic suppression in errors in the low noise limit but rather a constant-factor improvement whose magnitude depends on the typical size of a quantity we denote by the symbol γ ,

$$\gamma := \left| \sum_{j \neq 0} p_j \langle i|K_j|0\rangle\langle 0|K_j^\dagger|0\rangle \right|. \quad (44)$$

Interestingly, when we examine the data from our numerical simulations, we do obtain an improvement consistent with a quadratic suppression of errors at intermediate error rates. Additionally, γ has no lower bound; it can in some cases be zero, in which case, we expect to recover the quadratic suppression of error predicted from Sec. III A. As the trace distance is an upper bound for the error in any observable, particular observables of particular states may recover this performance even when $\gamma \neq 0$.

In order to shed some light on the error floor set by the drift in the dominant eigenvector of the density matrix, it can be helpful to ask when we might expect γ to be near

zero. It is clear that this quantity must be zero if one of two conditions hold:

$$K_j|0\rangle \propto |0\rangle, \quad (45)$$

$$\langle 0|K_j^\dagger|0\rangle = 0. \quad (46)$$

One way that this can occur is if the state and the circuit have a natural set of symmetries. The first condition holds if the error is drawn from such a symmetry group, while the second is satisfied if it violates it strictly. For an example of the second case, consider a bit-flip or amplitude-damping error channel acting on a state with a definite number of excitations. There are other situations where the second equality is approximately satisfied. For example, in circuits exhibiting the limits of quantum chaos, apart from a small light cone at the end of the circuit, any local errors lead to a state nearly indistinguishable from a Haar random state. Therefore, the matrix elements in Eq. (46) are exponentially small in the number of qubits. This sensitivity to local perturbations in random circuits is used in the cross-entropy benchmarking technique [44] and explains the improved behavior of our technique in numerical tests on random circuits.

IV. NUMERICAL EXPERIMENTS

In this section, we present numerical simulations of virtual distillation applied to three model systems. We first consider two classes of random circuits chosen because they are simple limits where the behavior of virtual distillation is easy to analyze. We then turn toward the application of virtual distillation to the simulation of the dynamics of a one-dimensional spin chain following a quantum quench. This example allows us to study the behavior of virtual distillation in the context of quantum simulation, an application which is a promising candidate for the eventual demonstration of practical quantum advantage in the NISQ era. We choose to focus on time evolution rather than the ground-state problem mainly because ground states have additional structure which enables specialized error mitigation techniques, and we are interested in how virtual distillation behaves in the absence of this structure.

We find it illuminating to characterize the effectiveness of our approach as a function of the expected number of errors in a particular circuit. This tends to allow more universal prediction of performance when trading between the error rate per gate and number of gates. We consider a noise model that focuses on stochastic errors in two-qubit gates. Specifically, after each two-qubit gate, we apply a single-qubit depolarizing channel to both qubits acted on by the gate. The expected number of errors (E) can be expressed simply as a function of the number of two-qubit gates in the circuit (G) and the single-qubit depolarizing probability [p defined in the usual way in Eq. (F3)],

$$E = 2pG. \quad (47)$$

To quantify the error, we focus mainly on the trace distance between the ideal state that would be obtained with noise-free evolution and the effective state accessed by virtual distillation. The trace distance leads to a natural bound in the error for the expectation value of an arbitrary observable,

$$|\text{Tr}(\rho O) - \langle \psi_{\text{ideal}} | O | \psi_{\text{ideal}} \rangle| \leq 2 \|O\| T(\rho, |\psi_{\text{ideal}}\rangle), \quad (48)$$

where O is an observable with operator norm $\|O\|$, and $T(-, -)$ denotes the trace distance.

A. Scrambling circuits

Both classes of random circuits that we simulate are related to the scrambling circuits used to demonstrate beyond classical computation in Ref. [44]. The first class is essentially a one-dimensional version of the circuit family considered in that work. The second class of circuits is the same as the first class, except that we remove the two-qubit gates. We provide some additional details in Appendix F. For these nonentangling random circuits, we still perform the noisy simulations of these circuits by applying single-qubit depolarizing channels in the same locations where the two-qubit gates would have been.

Because the behavior of the nonentangling random circuits is particularly simple to understand, we consider this class of circuits first. In the absence of entangling gates, we can commute the applications of the single-qubit depolarizing channel to the end of the circuit. We can then combine them into a single application per qubit with a larger effective error rate. We carry out this procedure analytically in Appendix F, showing that the dominant eigenvector of the density matrix corresponds exactly to the ideal state. This leads us to expect behavior that is similar to that of the phenomenological noise model we consider in Sec. III A.

In Fig. 4, we plot the trace distances between the ideal states generated by noiseless evolution and the states obtained by noisy evolution of these nonentangling random circuits (blue curve). We consider a variety of different circuit depths and error rates for both six-qubit systems (thin curves) and ten-qubit systems (thick curves). For each of these simulations, we also calculate the trace distance between the ideal state and the states we are effectively accessing by using virtual distillation with $M = 2$ [orange dotted curve, $\rho^2/\text{Tr}(\rho^2)$] or $M = 3$ [green dotted curve, $\rho^3/\text{Tr}(\rho^3)$] copies. For each particular setting of circuit depth and error rate, we consider a single randomly chosen member from the ensemble of nonentangling scrambling circuits described above.

We see that the data from this variety of simulations collapse together when we plot the error (in terms of trace distance) as a function of the expected number of gate errors. When the expected number of errors is not too large,




FIG. 4. The error in the unmitigated noisy states ($M = 1$) and the states accessed by virtual distillation ($M = 2, 3$) for a variety of nonentangling random circuits at two different system sizes (differentiated by thickness of markers). We plot the error quantified by the trace distance to the state obtained from noiseless evolution as a function of the expected number of single-qubit depolarizing errors resulting from varying both the error rate and number of gates. Unlike other cases, for these nonentangling circuits, the eigenvalue floor vanishes, and we see exponential suppression in the number of copies.

the curves for $M = 1$, $M = 2$, and $M = 3$ are nearly linear with slopes 1, 2, and 3, respectively. Although the noise model in this case does not exactly match the phenomenological model of Sec. III A, the results are broadly consistent. For these nonentangling random circuits, we observe a level of error suppression that is an exponential in M .

In Figs. 5 and 6, we present plots that explore the behavior of entangling random circuits on a one-dimensional line of qubits. When we consider the nonentangling random circuits, we find that errors (quantified by the trace distance to the ideal state) depend mostly on the system size and the expected number of gate errors. This is true regardless of whether the expected number of errors is varied by changing the circuit depth or by changing the error rate per gate. Here we observe slightly different behavior between these two cases and therefore consider them separately. These two figures also differ from Fig. 4 in that they include a red dashed curve corresponding to trace distance between the dominant eigenvector of the density matrix [$\lim_{M \rightarrow \infty} \rho^M / \text{Tr}(\rho^M)$] and the ideal state, a quantity which is nonzero for the richer family of circuits we now consider.

Figure 5 presents two plots that show the effects of varying the circuit depth at two different fixed error rates. We see that the error in the dominant eigenvector effectively sets a floor for the minimal error achievable by virtual distillation for any value of M . This floor grows slowly with increasing circuit depth. Furthermore, both the absolute magnitude and the rate of growth appear to be suppressed with system size. In Sec. III B, we show that the leading-order contributions to the trace distance between the




FIG. 5. The error in the unmitigated noisy states ($M = 1$), the states accessed by virtual distillation ($M = 2, 3$), and the dominant eigenvector of the density matrices ($M \rightarrow \infty$) for a variety of entangling random circuits. We plot the trace distance to the state obtained from noiseless evolution as a function of the expected number of single-qubit depolarizing errors for two different system sizes (represented by thickness of marker). Here we vary the expected number of errors by varying the number of gates, fixing the single-qubit depolarizing probabilities to 5×10^{-4} (left panel) or 5×10^{-3} (right panel). We see that the dominant eigenvector determines the noise floor beyond which we cannot improve, independent of the number of copies, and that this floor drops as the size of the system increases.




FIG. 6. The error in the unmitigated noisy states ($M = 1$), the states accessed by virtual distillation ($M = 2, 3$), and the dominant eigenvector of the density matrices ($M \rightarrow \infty$) for a variety of entangling random circuits. We plot the trace distance to the state obtained from noiseless evolution, as a function of the expected number of single-qubit depolarizing errors, for six- and ten-qubit systems (demarcated by the thickness of the symbols). We vary the expected number of errors by varying the error rate per gate, fixing the number of two-qubit gates to be 450. It is clear that there is a maximum number of expected errors for which the technique is effective, and below a certain error rate, the achievable improvement is fixed by the drift in the dominant eigenvector ($M \rightarrow \infty$).

dominant eigenvector and the ideal state can be understood in terms of the matrix elements of the Kraus operators [see Eq. (44)]. As the circuit depth of the random circuit increases, we expect a 1D random circuit to approach a Haar random circuit at a depth proportional to the number of qubits N . Once this approximation is sufficient, all but a small fraction of errors in the light cone of the observable at the end of the circuit will lead to matrix elements that contribute to the drift in the dominant eigenvector that are exponentially small in the number of qubits. This observation may explain the scaling we see in Fig. 5.

In Fig. 6, we plot the error in terms of trace distance as we vary the expected number of gate errors by varying the per-gate error rate for a fixed circuit. At low error rates, we see that the errors in the dominant eigenvectors (red dashed curves) scale linearly with the error rate but are orders of magnitude smaller than the errors in the unmitigated state. This matches the behavior we would expect from the analysis of Sec. III B. As in Fig. 5, the error in the dominant eigenvector sets a floor for the performance of our method at finite M and that this floor is suppressed as the system size increases.

B. Heisenberg quench

The properties of random circuits can be somewhat unique in their ability to scramble errors. It is thus important to consider how the approach works for other circuits of interest, like the quantum simulation of physical systems. In this section, we explore the performance of our approach applied to the simulation of time evolution following a quantum quench in a spin model. We initialize the system in an antiferromagnetic state, e.g., $|0101\rangle$, and simulate the time evolution under the Hamiltonian

$$H = \sum_{i=1}^{N-1} (J_x X_i X_{i+1} + J_y Y_i Y_{i+1} + J_z Z_i Z_{i+1}) + \sum_{i=1}^N h X_i. \quad (49)$$

Here we choose the parameters $J_x = J_y = 1.0$, $J_z = 1.5$, $h = 1.0$ in order to match a previously studied family of nonintegrable models [45], although we take open-boundary conditions rather than periodic ones. We approximate the time evolution under this Hamiltonian by Trotterization with a time step of $\Delta t = 0.2$. Specifically, we use alternating layers of single-qubit gates, two-qubit gates between odd-even pairs of qubits, and two-qubit gates between even-odd pairs of qubits. As above, we simulate the resulting circuits with single-qubit depolarizing noise applied after every two-qubit gate.

In Fig. 7, we plot the bounds on the error of an arbitrary observable (normalized so that $\|O\| = 1$) derived from the trace distance to the noiseless state. We vary the expected number of errors by varying the circuit depth of a six-qubit




FIG. 7. The average error in the single-site magnetization in the unmitigated noisy states ($M = 1$), the states accessed by our error mitigation technique ($M = 2, 3$), and the dominant eigenvectors of the density matrices ($M \rightarrow \infty$) for states generated by the Trotterized time evolution of a Heisenberg model. We plot the actual average errors we calculate using the blue dots ($M = 1$), orange crosses ($M = 2$), green squares ($M = 3$), red diamonds ($M \rightarrow \infty$), and large purple diamonds ($M = 2$, noisy distillation) alongside bounds determined by the trace distance to the ideal state using Eq. (48) (various curves). We plot these quantities as a function of the expected number of single-qubit depolarizing errors, which we vary by varying the number of gates, fixing the single-qubit depolarizing probability to 5×10^{-3} . We see that for this specific observable, the trace distance bounds are pessimistic by roughly an order of magnitude, though they generally respect the behavior of the eigenvector floor. Furthermore, we notice an almost perfect coincidence between the orange crosses and purple diamonds, indicating that performing the virtual distillation circuits of Sec. II A with noise has a negligible effect on the corrected expectation value.

system with a fixed single-qubit depolarizing probability of 5×10^{-3} . Alongside these bounds (plotted using solid and dashed curves), we also plot the actual average error in the single-site magnetization (averaged over the six sites) at various points throughout the circuit. For the two-copy ($M = 2$) version of our proposal, we plot the error calculated directly from the state $\rho^2 / \text{Tr}(\rho^2)$ using yellow crosses, and the error we would obtain by applying the destructive measurement described in Sec. II A using large purple diamonds. For this second calculation, we simulate the application of the six two-qubit gates [Eq. (10)] required to diagonalize the observables using the same noise model as the rest of the circuit. From the nearly perfect overlap of the yellow crosses with the purple diamonds, we can see that circuit noise during the diagonalization step has barely any effect on the reconstructed expectation values. It is also apparent that, although the average error in the magnetization does not saturate the bounds implied by the trace distance, our approach suppresses the errors in the actual expectation values to a similar degree that it suppresses the trace distance.

Figure 8 offers a different look at the same system. As with Fig. 6, in this figure we fix the number of two-qubit



FIG. 8. The error in the unmitigated noisy states ($M = 1$), the states accessed by virtual distillation ($M = 2, 3$), and the dominant eigenvectors of the density matrices ($M \rightarrow \infty$), for states generated by the Trotterized time evolution of a Heisenberg model. We plot the trace distance to the state obtained from noiseless evolution as a function of the expected number of single-qubit depolarizing errors. The expected number of errors is varied by changing the error rate per gate, fixing the number of two-qubit gates to be 450. We show these data for six- and ten-qubit systems (differentiated by the size of the markers). As we increase the system size from six qubits to ten, we observe that the error (quantified by the trace distance to the ideal state) decreases for the error-mitigated states ($M > 1$).

gates to be 450, and we vary the expected number of gate errors by sweeping over a range of per-gate error rates. Here we clearly see the impact of the floor set by the drift in the dominant eigenvector (red dotted curve). At low-gate-error rates, the error (in terms of trace distance to the ideal state) for the virtually distilled states ($M > 1$) is suppressed by a constant factor relative to the error in the unmitigated state ($M = 1$). The constant-factor improvement is substantial and appears to increase with system size. Both the size of the improvement and its sensitivity to the system size are smaller than we observe for the one-dimensional scrambling circuits of Fig. 6.

In Fig. 9, we consider the cost of performing the two-copy ($M = 2$) version of virtual distillation for the same systems considered in Fig. 8. We do this using the expression for the variance presented in Eq. (17) and derived in Appendix D. Using this expression, we calculate the average variance of our error-mitigated estimators for the single-site magnetization, $\{Z_i\}$. We consider the ratio of this average variance (for the error-mitigated expectation values) with the average variance of the same measurements without error mitigation. Because the number of measurements required for some fixed precision scales linearly with the variance, this ratio is also the ratio between the number of measurements required to use virtual distillation and the number of measurements required to measure the unmitigated expectation values (assuming the same target precision). This quantity therefore encapsulates the overhead incurred by our scheme.




FIG. 9. The average overhead in the number of measurement repetitions required to measure the single-site magnetization (of a Heisenberg model) with virtual distillation ($M = 2$). The overhead is calculated by taking the ratio of the time required using our technique to achieve a fixed target precision and the time required when measuring the same quantity to the same precision with respect to the unmitigated noisy state ($M = 1$). We vary the expected number of errors by varying the error rate per gate, fixing the number of two-qubit gates to be 450, and plot the overhead as a function of the expected number of single-qubit depolarizing errors. We show data for six- and ten-qubit systems, denoting the larger system using larger markers. As the error rate grows beyond $O(1)$ errors, the overhead increases dramatically, with the larger system seeing the greatest inflation.

When the expected number of errors is small, we see that virtual distillation barely increases the number of measurements required. It is only as the number of errors grows larger than one that the measurement cost rises dramatically. We note that Eq. (17) implicitly assumes that we perform a number of measurements R such that

$$R \gg \frac{1}{\text{Tr}(\rho^2)^2}. \quad (50)$$

This assumption will break down when the target precision is low and the expected number of errors in Fig. 9 is large, but the qualitative conclusion remains the same.

V. MITIGATING ALGORITHMIC ERRORS

To date, most error mitigation methods have focused on the reduction of errors caused by imperfections in a device implementation, such as decoherence or control errors. Here we explore the idea that some of these techniques can be applied to algorithmic errors incurred during otherwise noise-free implementations of randomized algorithms. Previous works have used extrapolation [46] or randomized symmetry application [47] to mitigate coherent errors in evolution; we extend this concept to incoherent errors.

Recent developments in Hamiltonian simulation have led to the development of randomized evolution methods such as qDRIFT (quantum stochastic drift protocol) [48], randomized Trotter [49], and combinations thereof [50],

which have benefits in some situations over their deterministic counterparts. As these methods are randomized, they output mixed states rather than pure states, even in the absence of noise. Moreover, they depend on an approximation parameter with a natural limit in which they converge to the pure state generated by exact evolution. In this section, we show numerically that virtual distillation applied to qDRIFT can suppress this deviation from the exact evolution. For the particular model system we consider, we find that virtual distillation can reduce the coherent space-time volume required to reach a particular accuracy threshold by a factor of 8 or more compared with the standard qDRIFT.

A. qDRIFT

We briefly introduce some background on the qDRIFT method. qDRIFT simulates time evolution under a Hamiltonian H by constructing product formulas using a randomized selection rule. Terms are chosen from H at random, with a selection probability proportional to their interaction strength in the Hamiltonian. One then evolves the system forward in time under this Hamiltonian term for a fixed time step and repeats this process a number of times, generating a product formula that provides an approximation to the time-evolution operator. When averaged over the classical randomness (in the choice of interaction terms), qDRIFT generates a quantum channel that closely approximates the exact evolution more closely than the individual product formulas. Importantly, unlike most deterministic Trotter methods, the scaling of this approach does not depend explicitly on the number of terms in the Hamiltonian, but rather the 1-norm of the coefficients.

More precisely, we consider a Hamiltonian that we may decompose as $H = \sum_i h_i H_i$, where all h_i are made real and positive by absorbing signs into H_i , and the spectral norm of H_i is bounded by 1. Defining $\lambda = \sum_i h_i$, the diamond norm distance between the qDRIFT channel and the true time evolution is bounded by

$$\epsilon = \frac{2\lambda^2 t^2}{\eta}, \quad (51)$$

where η is the number of qDRIFT selection steps performed to generate each instance of the qDRIFT channel, and hence controls the amount of coherent evolution required. As η increases, the resulting quantum channel converges to the unitary corresponding to the exact evolution. It will be our aim to understand how our virtual distillation technique can reduce the coherent space-time volume required by reducing this factor η required to achieve the same error in practice.

B. Virtual distillation applied to qDRIFT

Here we study the application of virtual distillation to qDRIFT numerically. Specifically, we investigate how virtual distillation can impact the number of coherent steps




FIG. 10. qDRIFT coherent cost reduction through virtual distillation in the Heisenberg model. Here we show the number of coherent qDRIFT steps required to reach a target trace distance, with and without virtual distillation using two copies. We see that there is a consistent reduction of at least $16\times$ in the number of required steps. When accounting for the overhead of using two copies, this amounts to an $8\times$ reduction in the coherent space-time volume used to reach the same error rate.

η required to reach a target accuracy. For this, we choose a Heisenberg Hamiltonian with up to six qubits per copy. The Hamiltonian is given by

$$H = \sum_{i=1}^N X_i X_{i+1} + Y_i Y_{i+1} + Z_i Z_{i+1} + h_i Z_i, \quad (52)$$

where $h_i \in \{-h, h\}$ are randomly chosen Z magnetic field strengths, and periodic boundary conditions are applied such that site $N+1$ is site 1. For our studies here, we choose a time-evolution length of $t = N$ and let $h = 1$. We numerically investigate the number of coherent qDRIFT steps required to achieve a trace distance of 0.01 to the ideal state for evolutions under such a Heisenberg model. The results of this analysis are shown in Fig. 10. We see for this system that virtual distillation consistently reduces the required number of coherent steps to achieve the desired trace distance by a factor of more than $16\times$. If we account for the space overhead of using two copies, this still amounts to a space-time advantage of $8\times$. These results suggest that the use of error mitigation techniques may be further developed to yield practical algorithmic improvements for real systems, especially in the NISQ regime.

VI. CONCLUSION

In this work, we show how techniques for using multiple copies of a state to access polynomials of the density matrix can be used to mitigate incoherent errors. We study the effectiveness of this approach for two analytically tractable noise models and characterize its limit in terms of the dominant eigenvector of the noisy density matrix. We numerically demonstrate reductions in the error (quantified by the trace distance to the noise-free state) of up to 3 orders

of magnitude for a collection of small model systems. Furthermore, we show that this error suppression is enhanced as the system size or the speed of information scrambling grows. We also consider the application of our error mitigation approach to the incoherent algorithmic error that arises when approximating time evolution using the qDRIFT algorithm, finding a substantial constant-factor improvement.

Our proposed strategy for error mitigation, which we refer to as virtual distillation, is simple to use and analyze. It provides a natural way to take advantage of the surplus of qubits that we expect to have available as the NISQ era continues to suppress the effects of incoherent errors. We expect that our technique will prove complementary to other error mitigation and calibration techniques, especially those capable of addressing coherent errors. The effective state accessed by virtual distillation approaches the dominant eigenvector of the density matrix exponentially quickly as the number of copies M increases. Therefore, the utility of our technique depends mainly on the error in the dominant eigenvector and the number of samples required. In particular, we devote a significant amount of attention to the question of sample complexity. This is due to the fact that proposed NISQ applications, especially in quantum simulation, already face a daunting cost in this regard [17]. Our analytical work and numerical simulations indicate that our strategy is most effective and affordable when the number of errors expected in the circuit is $O(1)$. The reach of our approach will therefore naturally grow throughout the NISQ era as hardware platforms continue to improve.

There are several other important considerations relevant to the performance of this technique in the NISQ era besides the drift in the dominant eigenvector and the sample complexity. First of all, virtual distillation requires collective measurements that couple the qubits of one copy with the corresponding qubits of additional copies. For hardware platforms based on a 2D grid of qubits, these measurements are easiest to perform when the connectivity of the original circuit is linear and would require a substantial number of additional gates in the general case. Fortunately, a linearly connected array of qubits is known to be sufficient to achieve the optimal gate complexity in some cases, including certain approaches for the simulation of quantum chemistry [51,52]. A second important consideration is that we are often interested in measuring observables with support on more than one qubit. We discuss some options for performing such measurements in Appendixes A–C, but these options come with their own overheads in gate complexity or the number of measurements repetitions. Third, and perhaps most importantly, the assumption that we are able to measure expectation values with respect to $\rho^M / \text{Tr}(\rho^M)$ is violated when the individual copies are not in the same state or when errors occur during execution of the measurements. We discuss some aspects of the breakdown

of this assumption in Appendix H. In general, we do not expect virtual distillation to be able to correct errors during the measurement process, but we note that we see a reasonable level of robustness to such errors in our numerical simulations (see Fig. 7).

Even as we begin to leave the NISQ era behind and approach devices that start to incorporate quantum error correction, in the early days the desire to use as many logical qubits as possible means we may perform some computations that still have an appreciable number of logical errors. Given that our technique can provide a substantial improvement in error at negligible overhead compared to traditional quantum error correction techniques, there may be some advantageous interplay between the two, where small distance codes are used in conjunction with this technique before more qubits are available. We explore this opportunity in more detail in Appendix G.

Our technique builds upon a long tradition of work that uses the symmetric group for stabilizing quantum computations and mitigating errors. We believe that this research direction continues to hold promise, and we identify a few directions that we find particularly intriguing. Virtual distillation is based on a simple collective measurement of M copies of ρ . In Sec. II B, we show that there exists more sophisticated collective measurements whose sample complexity improves quadratically upon our approach in some regimes. It would be interesting to investigate this further, both from a fundamental perspective, and with an eye toward practical implementation. Besides this potential improvement, another clear question arises from our work. The drift in the dominant eigenvector of the density matrix is a coherent error. As with the coherent errors that occur directly in the execution of circuits on a NISQ device, we are hopeful that variational optimization or other complementary error mitigation techniques may prove useful in addressing this additional source of coherent error. Studying this in the context of the real noise experienced in hardware will be especially important and illuminating.

ACKNOWLEDGMENTS

The authors are extremely grateful to Dave Bacon for key discussions around the symmetric subspace and diagrammatic derivations, Gian-Luca Anselmetti for suggestions around performing virtual distillation with two distinct states, and Nathan Wiebe for helpful advice at various stages. W. J. H. and K. B. W. acknowledge support from the NSF QLCI program through Grant No. OMA-2016245.

Note added.—Recently, a related paper, Ref. [53], appeared in the literature. Our work highlights different conclusions than theirs in error scaling due to our explicit consideration of errors that induce a drift in the dominant eigenvector of the density matrix.

APPENDIX A: MEASURING MULTIQUBIT OBSERVABLES BY DIAGONALIZATION

In Sec. II A, we present a straightforward strategy for applying the $M = 2$ version of virtual distillation to single-qubit observables. Measuring observables with support on more than one qubit is more challenging. Reference [35] solved this issue by using an approach like the one we describe below in Appendix C. Alternatively, one could apply a circuit to localize an observable of interest to a single qubit before performing virtual distillation. Here we present an alternative solution that does not require the use of ancilla-assisted measurement or circuit depth.

The challenge arises due to the use of Eq. (8), in particular, the choice to use the symmetrized version of an observable, a notion defined in Eq. (6). Using the symmetrized version of a multiqubit observable means that it is not possible to perform the required diagonalization using a tensor product of separate unitaries across each pair of qubits. As an example, we consider the operator $O = Z_i Z_j$. Our arguments hold equally well for any other operator composed of a tensor product of (more than one) single-qubit Pauli operators. Taking the product of the symmetrized observable and the swap operator yields

$$O^{(2)}S^{(2)} = \frac{1}{2}(Z_i^1 Z_j^1 + Z_i^2 Z_j^2)S^{(2)}. \quad (\text{A1})$$

This operator does not factorize into a tensor product of operators with support on the individual pairs of qubits, nor can it be diagonalized by an operator that factors this way. However, instead of using Eq. (8) to determine the corrected expectation value of O , we can instead use the nonsymmetrized form introduced in Eq. (2). Returning to our example where $O = Z_i Z_j$, we see that we need to estimate the numerator and denominator of

$$\frac{\text{Tr}(Z_i^1 Z_j^1 S^{(2)} \rho^{\otimes 2})}{\text{Tr}(S^{(2)} \rho^{\otimes 2})}. \quad (\text{A2})$$

Unlike the symmetrized observable of Eq. (A1), the operator $Z_i^1 Z_j^1 S^{(2)}$ factorizes into a tensor product over the N pairs of qubits (a pair being one qubit from the first system and the corresponding qubit from the second system). $Z_i^1 Z_j^1 S^{(2)}$ is not Hermitian, but because it is unitary, we can still estimate $\text{Tr}(Z_i^1 Z_j^1 S^{(2)} \rho^{\otimes 2})$ by applying a circuit to diagonalize it and measuring in the computational basis. As $Z_i^1 Z_j^1 S^{(2)}$ factorizes into a product of two-qubit operators, the circuit that diagonalizes it does as well.

Note that because $Z_i^1 Z_j^1$ does not commute with $S^{(2)}$, we are unable to simultaneously estimate the numerator and denominator of Eq. (A2). We are also unable to simultaneously measure the corrected expectation value corresponding to different choices of i and j . More generally, we

are able to measure any single tensor product of one-qubit operators at a time, regardless of the number of qubits it acts on. The details of the diagonalization will, of course, depend upon the operator to be measured. We do not carefully analyze the number of measurements required by this flavor of virtual distillation, but we note that the inability to parallelize the measurement of commuting multiqubit observables would make it challenging to profitably combine this approach with sophisticated NISQ measurement strategies, such as the one presented in Ref. [17]. In particular, an individual multiqubit operator must be measured separately using this approach, even if it commutes. This fact increases the overall number of circuit repetitions required for many applications.

APPENDIX B: MEASUREMENT BY DIAGONALIZATION WITH THREE OR MORE COPIES

In Sec. II A and Appendix A, we describe protocols for measuring the expectation value of O with respect to the state $\{[\rho^2]/[\text{Tr}(\rho^2)]\}$ by diagonalizing $S^{(2)}$ and either $O^1 S^{(2)}$ or $O^{(2)} S^{(2)}$. Here we describe how these approaches can be generalized to higher powers of ρ in a natural way. Like the swap operator, the cyclic shift operator between M N -qubit systems $S^{(M)}$ factorizes into a tensor product of N M -qubit gates. Specifically, it factorizes into the tensor product of N single-qubit cyclic shift operators. The symmetrized operator $Z_k^{(M)} = (1/M) \sum_{i=1}^M Z_k^i$ commutes with the operator $S^{(M)}$. Therefore, $Z_k^{(M)} S^{(M)}$ and $S^{(M)}$ are simultaneously diagonalizable even though $S^{(M)}$ is unitary but not Hermitian for $M > 2$. Because $Z_k^{(M)}$ and $S^{(M)}$ both factorize into tensor products over N M -tuples of qubits, the unitary that diagonalizes these operators then factorizes into a tensor product of M -qubit operators in the same way.

The same concerns about correcting the expectation values of multiqubit observables that we discuss in Appendix A for the two-copy ($M = 2$) case apply to this generalized proposal. The tools developed so far allow us to simultaneously estimate $\text{Tr}(Z_k \rho^M)$ for all values of m and also $\text{Tr}(\rho^M)$. If we are interested in reconstructing $\text{Tr}(P \rho^M)$ for some multiqubit Pauli operator P , we can do so using a generalization of Eq. (A2), but we would be limited to measuring the operators required for one particular P at a time.

For the specific case of $M = 3$, we numerically optimize the quantum circuit of Fig. 11 to simultaneously diagonalize $Z_k^{(3)} S_k^{(3)}$ and $S_k^{(3)}$. We obtain parameters for the four two-qubit gates that allow for an error (measured in the Frobenius norm of the difference between the exact and approximate matrices) of approximately $5E - 5$ when the following equations are used:




FIG. 11. The ansatz that we numerically optimize to approximately diagonalize $S_k^{(3)}$ and $Z_k^{(3)}$. The two-qubit gates parametrized by the $\vec{\theta}_i$'s are arbitrary two-qubit gates. We perform the numerical optimization using the JULIA language.

$$B_k^{(3)} S_k^{(3)} B_k^{(3)\dagger} \rightarrow \frac{1}{8} [2 + (-3 - \sqrt{3}i)Z_k^1 + (3 - \sqrt{3}i)Z_k^1 Z_k^2 + (3 + \sqrt{3}i)Z_k^3 + 2i\sqrt{3}Z_k^1 Z_k^3 + (3 - \sqrt{3}i)Z_k^2 Z_k^3], \quad (\text{B1})$$

$$B_k^{(3)} Z_k^{(3)} B_k^{(3)\dagger} \rightarrow \frac{1}{3} (Z_k^1 + Z_k^2 + Z_k^3). \quad (\text{B2})$$

APPENDIX C: ANCILLA-ASSISTED MEASUREMENT

In this section, we present the approach one may take if ancilla-assisted measurement is feasible in the experimental setup. This is a simplified version of the proposal for ancilla-assisted measurement protocol found in Ref. [35] for estimating the expectation value of an observable O with respect to the state $\rho^2/\text{Tr}(\rho^2)$. As with the method we discuss in Sec. II A, we do this by approximating the numerator and denominator of Eq. (8). Unlike that method, this approach uses a nondestructive measurement of the swap operator ($S^{(2)}$). The main reward for this added complexity is that this variant of virtual distillation does not restrict the form of the operators being measured, nor does it prevent simultaneous measurement of operators acting on overlapping subsets of qubits. Therefore, it is compatible with some of the recently developed techniques for efficiently measuring a large collection of commuting operators [17,54,55]. While we focus on the $M = 2$ copy version here, we also briefly discuss the generalization to $M \geq 3$ copies.

To use this method, we begin with two system registers, each in the state ρ , as well as an ancilla qubit in the $|0\rangle$ state. We then perform a nondestructive measurement of $S^{(2)}$ in the standard way, using the so-called swap or Hadamard test [25,26,56]. Specifically, we apply a Hadamard gate to the ancilla qubit, apply $S^{(2)}$ conditioned on the ancilla qubit being in the $|1\rangle$ state, and measure the ancilla qubit in the X basis. The expectation value of X on the ancilla qubit is then equal to $\langle S^{(2)} \rangle$. Because $S^{(2)}$ factorizes into a tensor product of two-qubit SWAP gates, its controlled version likewise factorizes into a series of N Fredkin (controlled-SWAP) gates. Compiling this circuit may necessitate some extra steps (such as expanding the single ancilla qubit into a

Greenberger-Horne-Zeilinger state using a series of controlled-NOT gates) in order to deal with the restricted connectivity of a near-term device.

It is not technically necessary but it simplifies the analysis and reduces the variance of the resulting estimator to focus on the symmetrized form of O , $O^{(2)} = \frac{1}{2}(O^1 + O^2)$. As in Sec. II A, this is beneficial because the symmetrized observable $O^{(2)}$ commutes with $S^{(2)}$. We can therefore measure the product $O^{(2)}S^{(2)}$ by first measuring $S^{(2)}$ using the Hadamard test described above and then measuring $O^{(2)}$ on the system registers. This protocol does not require a separate estimation of $\text{Tr}(\rho O)$ like the original proposal of Ref. [35]. Furthermore, it allows us to make use of measurements of O on both copies of ρ and also simultaneously estimate the numerator and denominator of Eq. (8), leading to a relatively sample-efficient scheme.

Now let us consider the case with three or more copies of ρ . $S^{(N)}$ is not Hermitian for $N > 2$, but the natural generalization to the above strategy still works as expected. Specifically, we can use a controlled version of the cyclic shift operator $S^{(N)}$ to sample an observable whose expectation value is equal to $\text{Re}[\text{Tr}(S^{(N)}\rho^{\otimes N})]$ [26,27,35].

Because the symmetrized observable $O^{(N)}$ commutes with $S^{(N)}$, it also commutes with the observable measured by this generalization of the swap test. Therefore, we can sample from an observable whose expectation value is $\text{Tr}(S^{(N)}O^{(N)}\rho^{\otimes N})$ by first performing the higher-order swap test and then a measurement of $O^{(N)}$.

APPENDIX D: VARIANCE OF THE CORRECTED EXPECTATION VALUE ESTIMATOR

In this section, we calculate the variance of the estimator for the corrected expectation value obtained by applying Eq. (8) with $M = 2$. Specifically, we consider the estimation of the expectation value of an observable O with respect to the state $\{\rho^2/\text{Tr}(\rho^2)\}$ constructed by repeatedly measuring the operators $S^{(2)}O^{(2)}$ and $S^{(2)}$,

$$\frac{\text{Tr}(O\rho^2)}{\text{Tr}(\rho^2)} = \frac{\text{Tr}(O^{(2)}S^{(2)}\rho^{\otimes 2})}{\text{Tr}(S^{(2)}\rho^{\otimes 2})}. \quad (\text{D1})$$

We assume that both operators are simultaneously measured by averaging over R repetitions of state preparation and measurement. This assumption applies to the measurement by the diagonalization method presented in Sec. II A of the main text and also to the ancilla-assisted measurement approach of Appendix C.

The outcomes obtained from measurements of these operators are classical random variables, and we can proceed by determining the variance of these two random variables and their covariance. We begin by calculating the variance of the numerator (with respect to the state $\rho^{\otimes 2}$)

$$\text{Var}(S^{(2)}O^{(2)}) = \langle (S^{(2)}O^{(2)})^2 \rangle - \langle S^{(2)}O^{(2)} \rangle^2 \quad (\text{D2})$$

$$= \langle (O^{(2)})^2 \rangle - \text{Tr}(\rho^2 O)^2 \quad (\text{D3})$$

$$= \frac{1}{4} \text{Tr}[(\rho \otimes \rho)(O^2 \otimes \mathbb{I} + 2O \otimes O + \mathbb{I} \otimes O^2)] - \text{Tr}(\rho^2 O)^2 \quad (\text{D4})$$

$$= \frac{1}{2} \text{Tr}(\rho O^2) + \frac{1}{2} \text{Tr}(\rho O)^2 - \text{Tr}(\rho^2 O)^2. \quad (\text{D5})$$

The variance of the random variable in the denominator follows by taking $O = \mathbb{I}$,

$$\text{Var}(S^{(2)}) = \langle (S^{(2)})^2 \rangle - \langle S^{(2)} \rangle^2 \quad (\text{D6})$$

$$= 1 - \text{Tr}(\rho^2)^2. \quad (\text{D7})$$

We also need the covariance between the random variables representing measurements of the operators which estimate the numerator and the denominator,

$$\text{Cov}(S^{(2)}O^{(2)}, S^{(2)}) = \langle S^{(2)}O^{(2)}S^{(2)} \rangle - \langle S^{(2)}O^{(2)} \rangle \langle S^{(2)} \rangle \quad (\text{D8})$$

$$= \langle O^{(2)} \rangle - \text{Tr}(\rho^2 O) \text{Tr}(\rho^2) \quad (\text{D9})$$

$$= \text{Tr}(\rho O) - \text{Tr}(\rho^2 O) \text{Tr}(\rho^2). \quad (\text{D10})$$

There is not a closed-form expression for the variance of the ratio of two random variables [57], but we can take the standard approximation based on a Taylor series expansion,

$$\text{Var}\left(\frac{A}{B}\right) \approx \frac{1}{\langle B \rangle^2} \text{Var}(A) - 2 \frac{\langle A \rangle}{\langle B \rangle^3} \text{Cov}(A, B) + \frac{\langle A \rangle^2}{\langle B \rangle^4} \text{Var}(B). \quad (\text{D11})$$

We estimate the expectation values for the numerator and denominator of Eq. (D1) by averaging over a series of R experiments. This scales the variances calculated above by a factor of $(1/R)$. If R is sufficiently large, then the approximation presented in Eq. (D11) will be a good one. Applying this expression to determine the variance of the estimator from Eq. (D1) yields

$$\begin{aligned} \text{Var}(\text{estimator}) \approx & \frac{1}{R} \left[\frac{1}{\text{Tr}(\rho^2)^2} \left(\frac{1}{2} \text{Tr}(\rho O^2) + \frac{1}{2} \text{Tr}(\rho O)^2 - \text{Tr}(\rho^2 O)^2 \right) \right. \\ & \left. - 2 \frac{\text{Tr}(\rho^2 O)}{\text{Tr}(\rho^2)^3} [\text{Tr}(\rho O) - \text{Tr}(\rho^2 O) \text{Tr}(\rho^2)] + \frac{\text{Tr}(\rho^2 O)^2}{\text{Tr}(\rho^2)^4} [1 - \text{Tr}(\rho^2)^2] \right]. \end{aligned} \quad (\text{D12})$$

APPENDIX E: VARIANCE OF THE PROPOSED COLLECTIVE MEASUREMENT

In Sec. II B, we claim that there exists a joint measurement on $2K$ copies of ρ that allows us to estimate $\text{Tr}(O\rho^2)$ with a lower variance than performing K copies of the basic virtual distillation procedure in parallel. Specifically, we claim that the operator \tilde{O} (whose definition we reproduce below) exhibits a lower variance than the simple alternative under certain conditions. In this Appendix, we prove this claim. First, we recall the definition

$$\tilde{O} = \frac{1}{\binom{2K}{2}} \sum_{i=1}^{2K} \sum_{j>i} \frac{1}{2} (O^i + O^j) S^{(i,j)}, \quad (\text{E1})$$

where $S^{(i,j)}$ denotes the swap operator between subsystems i and j , and O is an arbitrary Pauli operator. Linearity ensures that a calculation of the expectation value of this operator can be reduced to the virtual distillation procedure applied to two copies yielding

$$\text{Tr}(\tilde{O}\rho^{\otimes 2K}) = \text{Tr}(O\rho^2). \quad (\text{E2})$$

We now bound the variance of measurements of this operator with respect to the state $\rho^{\otimes 2K}$,

$$\text{Var}(\tilde{O}) = \langle O^2 \rangle - \langle O \rangle^2 \quad (\text{E3})$$

$$= \frac{1}{4 \binom{K}{2}^2} \sum_{i=1}^{2K} \sum_{j>i} \sum_{a=1}^{2K} \sum_{b>a} \langle (O^i + O^j) S^{(i,j)} (O^a + O^b) S^{(a,b)} \rangle - \text{Tr}(O\rho^2)^2. \quad (\text{E4})$$

Note that we make use of Eq. (E2) to replace $\langle O \rangle^2$ by $\text{Tr}(O\rho^2)^2$. We proceed by breaking the summation into three cases. In the first case, $i = a$ and $j = b$. In the second case, there are only three distinct values among the indices i, j, a, b . In the fourth case, all four indices take distinct values.

Consider the first case where $i = a$ and $j = b$. Then we can simplify and bound the sum

$$\frac{1}{4\binom{2K}{2}^2} \sum_{i=1}^{2K} \sum_{j>i} \langle (O^i + O^j)S^{(i,j)}(O^i + O^j)S^{(i,j)} \rangle = \frac{1}{4\binom{2K}{2}^2} \sum_{i=1}^{2K} \sum_{j>i} \langle 2 + O^i O^j + O^j O^i \rangle \quad (\text{E5})$$

$$\leq \frac{1}{\binom{K}{2}}. \quad (\text{E6})$$

Here we use the properties that $S^{(i,j)}$, O^i , and O^j are self-inverse, and that $S^{(i,j)}$ commutes with $(O^i + O^j)$.

Next, let us consider the third case, where all four indices take distinct values. Then the operators $(O^i + O^j)S^{(i,j)}$ and $(O^a + O^b)S^{(a,b)}$ act on distinct pairs of systems. Therefore, their expectation values with respect to the tensor product $\rho^{\otimes 2K}$ can be evaluated separately and multiplied together. We can use this fact to simplify and bound this component of the sum

$$\frac{1}{4\binom{2K}{2}^2} \sum_{i=1}^{2K} \sum_{j>i} \sum_{a=1, a \neq i, a \neq j}^{2K} \sum_{b>a, b \neq i, b \neq j} \langle (O^i + O^j)S^{(i,j)}(O^a + O^b)S^{(a,b)} \rangle \quad (\text{E7})$$

$$= \frac{1}{4\binom{2K}{2}^2} \sum_{i=1}^{2K} \sum_{j>i} \sum_{a=1, a \neq i, a \neq j}^{2K} \sum_{b>a, b \neq i, b \neq j} \langle (O^i + O^j)S^{(i,j)} \rangle \langle (O^a + O^b)S^{(a,b)} \rangle \quad (\text{E8})$$

$$= \frac{1}{2\binom{2K}{2}} \sum_{i=1}^{2K} \sum_{j>i} \sum_{a=1, a \neq i, a \neq j}^{2K} \sum_{b>a, b \neq i, b \neq j} \text{Tr}(O\rho^2)^2 \quad (\text{E9})$$

$$= \frac{1}{\binom{2K}{2}} \binom{2K-2}{2} \text{Tr}(O\rho^2)^2 < \text{Tr}(O\rho^2)^2. \quad (\text{E10})$$

Now we treat the case where the indices take three distinct values. Actually, there are four subcases here. We could have any one of the four possibilities $i = a$, $i = b$, $j = a$, or $j = b$. We work out the details for the $i = a$ case below, noting that the others behave symmetrically,

$$\langle (O^i + O^j)S^{(i,j)}(O^i + O^b)S^{(i,b)} \rangle \quad (\text{E11})$$

$$= \langle (O^i + O^j)(O^j + O^b)S^{(i,j)}S^{(i,b)} \rangle \quad (\text{E12})$$

$$= \langle (O^i O^j + O^i O^b + O^j O^b + 1)S^{(i,j)}S^{(i,b)} \rangle. \quad (\text{E13})$$

Here we use the property that O^i is self-inverse. Now we note that the product $S^{(i,j)}S^{(i,b)}$ is a cyclic shift between the subsystems i, j, b and that this product commutes with the operator $(O^i O^j + O^i O^b + O^j O^b + 1)$. Computation using a tensor network diagram (see Fig. 12) establishes that

$$\langle (O^i O^j + O^i O^b + O^j O^b + 1)S^{(i,j)}S^{(i,b)} \rangle \quad (\text{E14})$$

$$= 3\text{Tr}(O\rho O\rho^2) + \text{Tr}(\rho^3). \quad (\text{E15})$$

In order to bound this quantity, let us denote the projector onto the $+1$ eigenspace of O by P_+ and the projector onto the -1 eigenspace by P_- . Then we can expand Eq. (E15) in terms of these projectors yielding




FIG. 12. A diagrammatic proof of Eq. (E15).

$$3\text{Tr}(O\rho O\rho^2) + \text{Tr}(\rho^3) \quad (\text{E16})$$

$$= 3\text{Tr}(P_+\rho P_+\rho^2) - 3\text{Tr}(P_-\rho P_+\rho^2) - 3\text{Tr}(P_+\rho P_-\rho^2) + 3\text{Tr}(P_-\rho P_-\rho^2) + \text{Tr}(\rho^3) \quad (\text{E17})$$

$$\leq 7\text{Tr}(\rho^3). \quad (\text{E18})$$

For simplicity, let us define the indicator function

$$W(i, j, a, b) = \begin{cases} 1 & \text{if } i, j, a, b \text{ take exactly three distinct values,} \\ 0 & \text{otherwise.} \end{cases} \quad (\text{E19})$$

Now we can bound the component of the sum where the indices take three distinct values. For each of the $\binom{2K}{2}$ values of i and j , there are exactly $4K - 4$ values of a and b such that $I(i, j, a, b) = 1$. Therefore, we have

$$\frac{1}{4\binom{2K}{2}^2} \sum_{i=1}^{2K} \sum_{j>i}^{2K} \sum_{a=1}^{2K} \sum_{b>a}^{2K} \langle (O^i + O^j)S^{(i,j)}(O^a + O^b)S^{(a,b)} \rangle I(i, j, a, b) \quad (\text{E20})$$

$$\leq \frac{7}{4\binom{2K}{2}^2} \sum_{i=1}^{2K} \sum_{j>i}^{2K} \sum_{a=1}^{2K} \sum_{b>a}^{2K} I(i, j, a, b) \text{Tr}(\rho^3) \quad (\text{E21})$$

$$\leq \frac{7(4K - 4)}{4\binom{2K}{2}} \text{Tr}(\rho^3) \quad (\text{E22})$$

$$\leq \frac{7(2K - 2)}{2K(2K - 1)} \text{Tr}(\rho^3). \quad (\text{E23})$$

Now we can combine the bounds from the three different cases and simplify the expression for the variance to yield

$$\text{Var}(\tilde{O}) \leq \frac{1 + 7(K - 1)\text{Tr}(\rho^3)}{K(2K - 1)}. \quad (\text{E24})$$

Note that we simplify by subtracting the $\text{Tr}(O\rho^2)^2$ term that arose from evaluating $\langle O \rangle^2$.

APPENDIX F: DETAILS REGARDING THE NUMERICAL EXPERIMENTS WITH SCRAMBLING CIRCUITS

In Sec. IV A, we briefly describe the random circuits that we simulate to produce Figs. 4–6. Here we expand upon that description.

The first class of circuits is essentially the one-dimensional analogs of the random circuits of Ref. [44]. They are constructed by alternating between layers of two-qubit gates and single-qubit gates. The two-qubit-gate layers consist of ‘‘Sycamore gates,’’ two-qubit gates that enact the unitary

$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & -i & 0 \\ 0 & -i & 0 & 0 \\ 0 & 0 & 0 & e^{-\frac{i\pi}{6}} \end{bmatrix}. \quad (\text{F1})$$

The two-qubit-gate layers themselves alternate between layers that have Sycamore gates on every even-odd pair and every odd-even pair. The ensemble of random circuits is defined by adding a layer of randomly chosen single-qubit gates between every layer of two-qubit gates in this fixed structure. These single-qubit gates are drawn from the set

$$\{X, Y, Z, \sqrt{X}, \sqrt{Y}, \sqrt{Z}\}, \quad (\text{F2})$$

with the square root of a gate being defined by taking the principal square root in the eigenbasis of the gate.

The second class of random circuits is the same as the first class, except that we effectively remove the two-qubit gates by replacing the Sycamore gates with the identity. When we perform noisy simulations of these circuits, we apply the single-qubit depolarizing channels in the same locations despite the lack of two-qubit gates. Note that because this second class of random circuits contains only single-qubit gates, the applications of the single-qubit

depolarizing noise channels can be commuted to the end of the circuit and combined.

We can therefore write an analytical expression for the density matrix at the end of the noisy computation in terms of the noiseless single-qubit states $\{|\phi_i\rangle\}$, the single-qubit depolarizing probability p , and the depth of the circuit D . We define the single-qubit depolarizing channel in the usual way,

$$\Delta(\rho) = (1-p)\rho + \frac{p}{3} + \frac{p}{3}(X\rho X + Y\rho Y + Z\rho Z). \quad (\text{F3})$$

An equivalent formulation which is useful for our purposes is

$$\Delta(\rho) = \left(1 - \frac{4}{3}p\right)\rho + \frac{4}{3}p\frac{\mathbb{I}}{2}. \quad (\text{F4})$$

For a pure state $\rho = |\phi\rangle\langle\phi|$, we have

$$\Delta(|\phi\rangle\langle\phi|) = \left(1 - \frac{2}{3}p\right)|\phi\rangle\langle\phi| + \frac{2}{3}p|\phi^\perp\rangle\langle\phi^\perp|, \quad (\text{F5})$$

where $|\phi^\perp\rangle$ is orthogonal to $|\phi\rangle$. Equation (F4) allows us to easily analyze D -repeated applications of the channel,

$$\Delta(\rho)^D = \left(1 - \frac{4}{3}p\right)^D \rho + \left[1 - \left(1 - \frac{4}{3}p\right)^D\right] \frac{\mathbb{I}}{2}, \quad (\text{F6})$$

which tells us that D applications of the channel with an error rate p are equivalent to a single application with error rate

$$\tilde{p} = \frac{3}{4} - \frac{3}{4} \left(1 - \frac{4}{3}p\right)^D. \quad (\text{F7})$$

We can now write an expression for the density matrix at the end of the computation,

$$\bigotimes_{i=1}^N \left(1 - \frac{2}{3}\tilde{p}_{(i)}\right) |\phi_i\rangle\langle\phi_i| + \frac{2}{3}\tilde{p}_{(i)} |\phi_i^\perp\rangle\langle\phi_i^\perp|. \quad (\text{F8})$$

Here the effective single-qubit depolarizing probability $\tilde{p}_{(i)}$ depends on i because the qubits at the end of the circuit are subject only to $D/2$ applications of the single-qubit depolarizing channel instead of the D that are applied to qubits in the bulk. Therefore, employing Eq. (F7), we have

$$\tilde{p}_{(i)} = \begin{cases} \frac{3}{4} - \frac{3}{4} \left(1 - \frac{4}{3}p\right)^{\frac{D}{2}} & \text{for } i = 1 \text{ or } i = N, \\ \frac{3}{4} - \frac{3}{4} \left(1 - \frac{4}{3}p\right)^D & \text{for } 2 \leq i \leq N-1. \end{cases} \quad (\text{F9})$$

We can observe a few things from the combination of Eqs. (F8) and (F9). First of all, for any value of p smaller than the maximal $p = \frac{3}{4}$, the dominant eigenvector of the density matrix is exactly the ideal state $(|\phi_1\rangle \otimes |\phi_2\rangle \otimes \cdots \otimes |\phi_N\rangle)$. Second, when p is small,

the next largest eigenvectors of the density matrix will correspond to states with an error on a single qubit. Neglecting the subtlety caused by the two different values of \tilde{p} , we can see that there will be N such eigenvectors with eigenvalues approximately equal to $\frac{2}{3}\tilde{p}$. After these states, there will be $\binom{N}{2}$ eigenvectors corresponding to two states with two errors. This distribution does not exactly match the phenomenological noise model we assume in Sec. III A, but it is qualitatively similar.

APPENDIX G: INTERPLAY WITH THE SURFACE CODE

While the majority of this work focuses on the NISQ regime, one interesting question to ask is what role does this approach play after some degree of quantum error correction has been deployed? To explore this connection concretely, we imagine that a fault-tolerant surface-code quantum computer is in operation with typical gate-error rates on the order of 10^{-3} . For such systems, it has been determined numerically [58] that in conjunction with a minimum-weight perfect-matching decoder, the error rate of a surface-code cycle is roughly

$$\epsilon_c = 10^{-(d+3)/2}, \quad (\text{G1})$$

where d is the distance of the code protecting a given logical qubit. Including data and measurement qubits, the translation to physical qubits for a given distance is $n = 2d^2$. In order to guarantee protection against measurement (or timelike) errors up to the same distance without using an excessive number of qubits, one must repeat measurements a number of cycles proportional to d . For operations like gates, additional cycles are required to perform the operation as well. For example, many simple Clifford operations may be done in a number of cycles like $2d$ using lattice surgery techniques [59,60]. However, more complicated arbitrary rotations like the ones used in many NISQ algorithms must first be broken down into a combination of discrete gates like T and Clifford gates through gate synthesis, then those T gates consume on the order of $20d$ cycles for successful distillation. Using a coarse synthesis heuristic of roughly ten T gates and ten Clifford gates per arbitrary rotation, this gives approximately $200d$ cycles of the surface code per arbitrary rotation. If we average this coarsely, assuming an even distribution of Clifford and arbitrary rotations, as is common in NISQ approaches, then we can model on average that we require $100d$ rounds of the surface code per gate we wish to perform. While these numbers are subject to refinement and improvements, we believe these can approximately serve to understand where an advantageous combination of methods might occur. As is common with early circuit implementations, we may assume that the gates are densely packed so that additional idling error is minimal.

With these assumptions, using n physical qubits to represent a single logical qubit, we have an overall fidelity of

$$f_1 = (1 - 10^{-(\sqrt{n/2+3})/2})^{100\sqrt{n/2}G}, \quad (\text{G2})$$

where G is the number of gates performed. The virtual distillation technique uses twice the qubits to effect a large constant-factor improvement over the bare circuit. Hence, the apt comparison here is to consider the use of twice the qubits within the virtual distillation technique, or to use twice the qubits to improve the distance of the surface-code logical qubit. An asymptotic analysis would argue that the exponential returns of the error-correcting code would be the best option; however, the overhead can mean that a large constant factor could make virtual distillation advantageous in some cases. To examine this possibility, consider the error rate of G gates in the surface code using twice the qubits

$$f_2 = (1 - 10^{-(\sqrt{n}+3)/2})^{100\sqrt{n}G}. \quad (\text{G3})$$

A strict analysis would consider that we need to round these to integer distances, but for this approximate analysis, this should suffice. If we consider the ratio between the implied error rates $c_s = (1 - f_1)/(1 - f_2)$, we can find the required constant factor for a given number of qubits per logical qubit and the number of gates to make using virtual distillation advantageous. Past a certain number of qubits, this constant factor is enormous, but we find that up to distance 10–15 the empirical improvements measured in the text are sufficient to justify the use of virtual distillation in place of additional qubit protection. In particular, at distance 10 with $n = 200$ physical qubits per logical qubit, performing $G = 1000$ gates on the logical qubit, the respective error rates are about 10^{-1} and 10^{-5} , and hence, a constant improvement of about 10^4 is sufficient to justify the use of virtual distillation, which is on par with some improvements seen in the main text. To be fair, one might argue that an overall error rate of 10^{-5} would already suffice, and by a distance of 15, the required improvement is on the order of 10^7 , which is at the upper limit of what we imagine can be achieved with this technique. At smaller distances and numbers of gates, the required constant factors decrease as well. If we assume that we consistently push the limits of the number of logical qubits we use, reducing the number of physical qubits per logical qubit available, this may imply a regime in early fault tolerance where this technique is applicable. Further studies will be required to identify precisely under what conditions this may be the case.

APPENDIX H: VIRTUAL DISTILLATION APPLIED TO DISTINCT STATES

In the main body of this paper, we apply virtual distillation to a variety of systems under the assumption that we have access to multiple copies of the same noisy state. In reality,

even if we attempt to perform the same computation multiple times in parallel, the noise experienced by each copy will not be identical. We now consider the case where we apply virtual distillation to two distinct states ρ_A and ρ_B . It is straightforward to show that we then effectively measure expectation values with respect to the state

$$\rho_{\text{eff}} = \frac{\rho_A \rho_B}{\text{Tr}(\rho_A \rho_B)}. \quad (\text{H1})$$

This equality can be demonstrated in various ways, but the most straightforward is to use a diagrammatic proof of the kind we illustrated in Fig. 1. Note that we still rely on the important assumption that the two copies are unentangled prior to virtual distillation. Furthermore, some additional care must be taken when virtual distillation is applied to two distinct states. In particular, ρ_{eff} is not guaranteed to be a positive semidefinite matrix and does not, in general, correspond to a valid quantum state. This is especially important to note in the context of variational algorithms, where care would have to be taken to ensure that a nonvariational answer is not achieved due to the nonphysicality of ρ_{eff} . It will be an interesting direction for future work to address this potential challenge.

In order to explore the impact of virtually distilling two different states together, we present an additional simulation of the Heisenberg evolution that we consider in Sec. IV B. In Fig. 13, rather than employing the single-qubit depolarizing




FIG. 13. The error in the unmitigated noisy states, and the states accessed by virtual distillation, for the same ten-qubit Heisenberg evolution under two different noise models. We consider a bit-flip error model (ρ_{bit} , teal curve) and phase-flip error model (ρ_{phase} , orange dashed curve). We show the error for virtual distillation applied in the usual way to two identical copies of each noisy state (ρ_{bit}^2 and ρ_{phase}^2 , purple dotted curve and pink dashed curve). We also consider the error when virtual distillation is applied to one copy of each state ($\rho_{\text{bit}}\rho_{\text{phase}}$, green dotted curve). We quantify the error in terms of the trace distance to the state obtained from noiseless evolution as a function of the expected number of single-qubit gate errors. The expected number of errors is varied by changing the error rate per gate, fixing the number of two-qubit gates to be 450. Ultimately, we find that the performance of virtual distillation is barely affected when the two input states are generated with different noise processes.

noise model we used previously, we simulate $\rho_A = \rho_{\text{bit}}$ using an analogous application of a bit-flip error channel, while using a phase-flip channel for $\rho_B = \rho_{\text{phase}}$. We find that the error in the effective state accessed by performing virtual distillation to these two different states closely tracks the error we obtain by using two copies of either state individually.

APPENDIX I: PERFORMANCE UNDER AMPLITUDE DAMPING AND DEPHASING NOISE

In Sec. IV, we use numerical simulations of three different systems to explore the performance of virtual distillation under a noise model where we apply a single-qubit depolarizing noise channel after each two-qubit gate. It is natural to ask how virtual distillation performs under a more realistic approximation to the actual noise experienced on NISQ hardware. In this Appendix, we shed light on this question by considering a more physically motivated model of stochastic errors in two-qubit gates based on single-qubit amplitude-damping and dephasing channels. As with the error model considered in the main text, in the error model we construct here each two-qubit gate is followed by the application of a single-qubit error channel to each qubit. Here that single-qubit error channel is the concatenation of the dephasing and amplitude-damping channels described below.

Our amplitude-damping channel is parametrized by γ_1 , which represents the probability that a qubit in the $|1\rangle$ state will spontaneously decay into the $|0\rangle$ state. We can express this channel in a standard (but nonunique) way using the Kraus operators,

$$M_0 = \begin{bmatrix} 1 & 0 \\ 0 & \sqrt{1-\gamma_1} \end{bmatrix}, \quad M_1 = \begin{bmatrix} 0 & \sqrt{\gamma_1} \\ 0 & 0 \end{bmatrix}. \quad (\text{I1})$$

Our dephasing channel is parametrized by γ_2 , which represents the probability that an unintended interaction between a qubit and its environment entangles the two, effectively performing a measurement in the computational basis. One standard way of expressing this channel in terms of Kraus operators is

$$M_0 = \sqrt{1-\gamma_2} \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad M_1 = \begin{bmatrix} \sqrt{\gamma_2} & 0 \\ 0 & 1 \end{bmatrix},$$

$$M_2 = \begin{bmatrix} 0 & 0 \\ 0 & \sqrt{\gamma_2} \end{bmatrix}. \quad (\text{I2})$$

It can also be convenient to use the equivalent representation

$$M_0 = \begin{bmatrix} 1 & 0 \\ 0 & \sqrt{1-\tilde{\gamma}} \end{bmatrix}, \quad M_1 = \begin{bmatrix} 0 & 0 \\ 0 & \sqrt{\tilde{\gamma}} \end{bmatrix},$$

$$\tilde{\gamma} = 2\gamma_2 - \gamma_2^2. \quad (\text{I3})$$

The parameters γ_1 and γ_2 can be related to the T_1 and T_2 times frequently used to characterize relaxation in two-level systems [61,62]. In our simulations, we consider $\lambda_1 = \lambda_2$, a choice with a physical model where the T_1 and T_2 times are comparable. This allows us to plot the error (quantified by the trace distance to the state that would be obtained in the absence of noise) as a function of the expected number of error events. The expected number of errors (E) can be expressed in terms of the number of two-qubit gates in the circuit (G) and the two error probabilities

$$E = 2G(\lambda_1 + \lambda_2). \quad (\text{I4})$$

In Figs. 14 and 15, we present the results of two sets of simulations under this error model. Figure 14 closely follows Fig. 6 from the main text, examining the performance of virtual distillation applied to a random circuit on a one-dimensional array of qubits (see Sec. IV A). Likewise, Fig. 15 considers the same Heisenberg evolution treated in Fig. 8 and described in Sec. IV B. In both cases, we show the error in the unmitigated noisy state, the states accessed by virtual distillation with $M = 2$ and $M = 3$ copies, and the dominant eigenvector of the noisy density matrix. We quantify the error in terms of the trace distance to the ideal state that would be obtained in the absence of noise, plotting this trace distance as a function of the expected number of errors.




FIG. 14. The error in the unmitigated noisy states ($M = 1$), the states accessed by virtual distillation ($M = 2, 3$), and the dominant eigenvector of the density matrices ($M \rightarrow \infty$) for a variety of entangling random circuits (described in Sec. IV A). We plot the trace distance to the state obtained from noiseless evolution as a function of the expected number of single-qubit amplitude-damping and dephasing errors, for six- and ten-qubit systems (demarcated by the thickness of the symbols). We vary the expected number of errors by varying the error rate per gate, fixing the number of two-qubit gates to be 450. This figure is constructed to parallel Fig. 6, except that we consider an error model based on single-qubit amplitude damping and dephasing rather than a depolarizing channel. We see that the system size dependence vanishes at low error rates in this case, in contrast with the data from Fig. 6.




FIG. 15. The error in the unmitigated noisy states ($M = 1$), the states accessed by virtual distillation ($M = 2, 3$), and the dominant eigenvectors of the density matrices ($M \rightarrow \infty$), for states generated by the Trotterized time evolution of a Heisenberg model (described in Sec. IV B). We plot the trace distance to the state obtained from noiseless evolution as a function of the expected number of single-qubit amplitude-damping and dephasing errors. The expected number of errors is varied by changing the error rate per gate, fixing the number of two-qubit gates to be 450. We show these data for six- and ten-qubit systems (differentiated by the size of the markers). This figure mirrors Fig. 8, except that we consider an error model based on single-qubit amplitude damping and dephasing rather than depolarizing noise. The two figures display similar behavior, indicating that our conclusions about virtual distillation have some robustness to changes in the noise model.

Comparing Figs. 14 and 15 with Figs. 6 and 8 from the main text, we note a few things. First of all, the qualitative behavior of virtual distillation is largely unaffected by the change in the noise model. When the expected number of errors is small enough, we still see that virtual distillation decreases the error by orders of magnitude. The performance is still limited by the drift in the dominant eigenvector, and we find that in many regimes $M = 2$ copies are sufficient to maximize the potential benefit. The largest differences are observed when one compares the two figures that analyze the performance of virtual distillation on random circuits (Figs. 14 and 6). In this case, the potential benefit of virtual distillation is smaller (roughly 2 orders of magnitude instead of 3 for the ten-qubit system) under the noise model based on amplitude damping and dephasing. Additionally, and perhaps more interestingly, we see that the dependence on system size mostly vanishes for the random circuits when we switch to the noise model considered in this Appendix, while it persists when we consider the Heisenberg evolution.

[1] D. Aharonov and M. Ben-Or, in *Proceedings of the Twenty-Ninth Annual ACM Symposium on Theory of Computing, STOC '97* (Association for Computing Machinery, New York, 1997), pp. 176–188.

- [2] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, *Surface Codes: Towards Practical Large-Scale Quantum Computation*, *Phys. Rev. A* **86**, 032324 (2012).
- [3] J. Preskill, *Quantum Computing in the NISQ Era and Beyond*, *Quantum* **2**, 79 (2018).
- [4] K. Temme, S. Bravyi, and J. M. Gambetta, *Error Mitigation for Short-Depth Quantum Circuits*, *Phys. Rev. Lett.* **119**, 180509 (2017).
- [5] S. Endo, S. C. Benjamin, and Y. Li, *Practical Quantum Error Mitigation for Near-Future Applications*, *Phys. Rev. X* **8**, 031027 (2018).
- [6] A. Kandala, K. Temme, A. D. Córcoles, A. Mezzacapo, J. M. Chow, and J. M. Gambetta, *Error Mitigation Extends the Computational Reach of a Noisy Quantum Processor*, *Nature (London)* **567**, 491 (2019).
- [7] A. Strikis, D. Qin, Y. Chen, S. C. Benjamin, and Y. Li, *Learning-Based Quantum Error Mitigation*, [arXiv:2005.07601](https://arxiv.org/abs/2005.07601).
- [8] P. Czarnik, A. Arrasmith, P. J. Coles, and L. Cincio, *Error Mitigation with Clifford Quantum-Circuit Data*, [arXiv:2005.10189](https://arxiv.org/abs/2005.10189).
- [9] F. Arute *et al.*, *Observation of Separated Dynamics of Charge and Spin in the Fermi-Hubbard Model*, [arXiv:2010.07965](https://arxiv.org/abs/2010.07965).
- [10] T. E. O’Brien, S. Polla, N. C. Rubin, W. J. Huggins, S. McArdle, S. Boixo, J. R. McClean, and R. Babbush, *Error Mitigation via Verified Phase Estimation*, *PRX Quantum* **2**, 020317 (2021).
- [11] 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).
- [12] J. I. Colless, V. V. Ramasesh, D. Dahlen, M. S. Blok, M. E. Kimchi-Schwartz, J. R. McClean, J. Carter, W. A. de Jong, and I. Siddiqi, *Computation of Molecular Spectra on a Quantum Processor with an Error-Resilient Algorithm*, *Phys. Rev. X* **8**, 011021 (2018).
- [13] J. R. McClean, Z. Jiang, N. C. Rubin, R. Babbush, and H. Neven, *Decoding Quantum Errors with Subspace Expansions*, *Nat. Commun.* **11**, 636 (2020).
- [14] 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).
- [15] S. McArdle, X. Yuan, and S. Benjamin, *Error-Mitigated Digital Quantum Simulation*, *Phys. Rev. Lett.* **122**, 180501 (2019).
- [16] R. Sagastizabal, X. Bonet-Monroig, M. Singh, M. A. Rol, C. C. Bultink, X. Fu, C. H. Price, V. P. Ostroukh, N. Muthusubramanian, A. Bruno, M. Beekman, N. Haider, T. E. O’Brien, and L. DiCarlo, *Experimental Error Mitigation via Symmetry Verification in a Variational Quantum Eigensolver*, *Phys. Rev. A* **100**, 010302(R) (2019).
- [17] W. J. Huggins, J. McClean, N. Rubin, Z. Jiang, N. Wiebe, K. B. Whaley, and R. Babbush, *Efficient and Noise Resilient Measurements for Quantum Chemistry on Near-Term Quantum Computers*, *npj Quantum Inf.* **7**, 23 (2021).
- [18] F. Arute *et al.* (Google AI Quantum and Collaborators), *Hartree-Fock on a Superconducting Qubit Quantum Computer*, *Science* **369**, 1084 (2020).

- [19] Y. Chen, M. Farahzad, S. Yoo, and T.-C. Wei, *Detector Tomography on IBM Quantum Computers and Mitigation of an Imperfect Measurement*, *Phys. Rev. A* **100**, 052315 (2019).
- [20] F. B. Maciejewski, Z. Zimborás, and M. Oszmaniec, *Mitigation of Readout Noise in Near-Term Quantum Devices by Classical Post-Processing Based on Detector Tomography*, *Quantum* **4**, 257 (2020).
- [21] S. Bravyi, S. Sheldon, A. Kandala, D. C. McKay, and J. M. Gambetta, *Mitigating Measurement Errors in Multi-Qubit Experiments*, *Phys. Rev. A* **103**, 042605 (2021).
- [22] A. Berthiaume, D. Deutsch, and R. Jozsa, in *Proceedings of the Workshop on Physics and Computation, PhysComp '94* (IEEE, New York, 1994), pp. 60–62.
- [23] A. Barenco, A. Berthiaume, D. Deutsch, A. Ekert, R. Jozsa, and C. Macchiavello, *Stabilization of Quantum Computations by Symmetrization*, *SIAM J. Comput.* **26**, 1541 (1997).
- [24] A. Peres, *Error Symmetrization in Quantum Computers*, *Int. J. Theor. Phys.* **38**, 799 (1999).
- [25] P. Horodecki and A. Ekert, *Method for Direct Detection of Quantum Entanglement*, *Phys. Rev. Lett.* **89**, 127902 (2002).
- [26] A. K. Ekert, C. M. Alves, D. K. L. Oi, M. Horodecki, P. Horodecki, and L. C. Kwek, *Direct Estimations of Linear and Nonlinear Functionals of a Quantum State*, *Phys. Rev. Lett.* **88**, 217901 (2002).
- [27] T. A. Brun, *Measuring Polynomial Functions of States*, *Quantum Inf. Comput.* **4**, 401 (2004).
- [28] M. B. Hastings, I. González, A. B. Kallin, and R. G. Melko, *Measuring Renyi Entanglement Entropy in Quantum Monte Carlo Simulations*, *Phys. Rev. Lett.* **104**, 157201 (2010).
- [29] R. Islam, R. Ma, P. M. Preiss, M. E. Tai, A. Lukin, M. Rispoli, and M. Greiner, *Measuring Entanglement Entropy in a Quantum Many-Body System*, *Nature (London)* **528**, 77 (2015).
- [30] J. C. Garcia-Escartin and P. Chamorro-Posada, *SWAP Test and Hong-Ou-Mandel Effect Are Equivalent*, *Phys. Rev. A* **87**, 052330 (2013).
- [31] S. Johri, D. S. Steiger, and M. Troyer, *Entanglement Spectroscopy on a Quantum Computer*, *Phys. Rev. B* **96**, 195136 (2017).
- [32] Y. Subaşı, L. Cincio, and P. J. Coles, *Entanglement Spectroscopy with a Depth-Two Quantum Circuit*, *J. Phys. A* **52**, 044001 (2019).
- [33] L. Banchi, A. Bayat, and S. Bose, *Entanglement Entropy Scaling in Solid-State Spin Arrays via Capacitance Measurements*, *Phys. Rev. B* **94**, 241117(R) (2016).
- [34] D. Bacon, I. L. Chuang, and A. W. Harrow, *Efficient Quantum Circuits for Schur and Clebsch-Gordan Transforms*, *Phys. Rev. Lett.* **97**, 170502 (2006).
- [35] J. Cotler, S. Choi, A. Lukin, H. Gharibyan, T. Grover, M. E. Tai, M. Rispoli, R. Schittko, P. M. Preiss, A. M. Kaufman, M. Greiner, H. Pichler, and P. Hayden, *Quantum Virtual Cooling*, *Phys. Rev. X* **9**, 031013 (2019).
- [36] S. Bravyi and A. Kitaev, *Universal Quantum Computation with Ideal Clifford Gates and Noisy Ancillas*, *Phys. Rev. A* **71**, 022316 (2005).
- [37] E. Knill, *Quantum Computing with Realistically Noisy Devices*, *Nature (London)* **434**, 39 (2005).
- [38] J. Haah and M. B. Hastings, *Codes and Protocols for Distilling T, Controlled-S, and Toffoli Gates*, *Quantum* **2**, 71 (2018).
- [39] J. I. Cirac, A. K. Ekert, and C. Macchiavello, *Optimal Purification of Single Qubits*, *Phys. Rev. Lett.* **82**, 4344 (1999).
- [40] J. Biamonte and V. Bergholm, *Tensor Networks in a Nutshell*, [arXiv:1708.00006](https://arxiv.org/abs/1708.00006).
- [41] J. C. Bridgeman and C. T. Chubb, *Hand-Waving and Interpretive Dance: An Introductory Course on Tensor Networks*, *J. Phys. A* **50**, 223001 (2017).
- [42] R. Orus, *A Practical Introduction to Tensor Networks: Matrix Product States and Projected Entangled Pair States*, *Ann. Phys. (Amsterdam)* **349**, 117 (2014).
- [43] T. Kato, *Perturbation Theory for Linear Operators* (Springer Science & Business Media, New York, 2013).
- [44] F. Arute *et al.*, *Quantum Supremacy Using a Programmable Superconducting Processor*, *Nature (London)* **574**, 505 (2019).
- [45] D. V. Dmitriev and V. Y. Krivnov, *Quasi-One-Dimensional Anisotropic Heisenberg Model in a Transverse Magnetic Field*, *J. Exp. Theor. Phys. Lett.* **80**, 303 (2004).
- [46] S. Endo, Q. Zhao, Y. Li, S. Benjamin, and X. Yuan, *Mitigating Algorithmic Errors in a Hamiltonian Simulation*, *Phys. Rev. A* **99**, 012334 (2019).
- [47] M. C. Tran, Y. Su, D. Carney, and J. M. Taylor, *Faster Digital Quantum Simulation by Symmetry Protection*, *PRX Quantum* **2**, 010323 (2021).
- [48] E. Campbell, *Random Compiler for Fast Hamiltonian Simulation*, *Phys. Rev. Lett.* **123**, 070503 (2019).
- [49] A. M. Childs, A. Ostrander, and Y. Su, *Faster Quantum Simulation by Randomization*, *Quantum* **3**, 182 (2019).
- [50] Y. Ouyang, D. R. White, and E. T. Campbell, *Compilation by Stochastic Hamiltonian Sparsification*, *Quantum* **4**, 235 (2020).
- [51] I. D. Kivlichan, J. McClean, N. Wiebe, C. Gidney, A. Aspuru-Guzik, G. K.-L. Chan, and R. Babbush, *Quantum Simulation of Electronic Structure with Linear Depth and Connectivity*, *Phys. Rev. Lett.* **120**, 110501 (2018).
- [52] B. O’Gorman, W. J. Huggins, E. G. Rieffel, and K. B. Whaley, *Generalized Swap Networks for Near-Term Quantum Computing*, [arXiv:1905.05118](https://arxiv.org/abs/1905.05118).
- [53] B. Koczor, *Exponential Error Suppression for Near-Term Quantum Devices*, *Phys. Rev. X* **11**, 031057 (2021).
- [54] J. Cotler and F. Wilczek, *Quantum Overlapping Tomography*, *Phys. Rev. Lett.* **124**, 100401 (2020).
- [55] X. Bonet-Monroig, R. Babbush, and T. E. O’Brien, *Nearly Optimal Measurement Scheduling for Partial Tomography of Quantum States*, *Phys. Rev. X* **10**, 031064 (2020).
- [56] A. Yu. Kitaev, *Quantum Measurements and the Abelian Stabilizer Problem*, [arXiv:quant-ph/9511026](https://arxiv.org/abs/quant-ph/9511026).
- [57] M. G. Kendall *et al.*, *The Advanced Theory of Statistics* (Charles Griffin and Co., Ltd., London, 1946).
- [58] D. S. Wang, A. G. Fowler, A. M. Stephens, and L. C. L. Hollenberg, *Threshold Error Rates for the Toric and Planar Codes*, *Quantum Inf. Comput.* **10**, 456 (2010).
- [59] C. Horsman, A. G. Fowler, S. Devitt, and R. Van Meter, *Surface Code Quantum Computing by Lattice Surgery*, *New J. Phys.* **14**, 123011 (2012).

- [60] D. Litinski, *A Game of Surface Codes: Large-Scale Quantum Computing with Lattice Surgery*, *Quantum* **3**, 128 (2019).
- [61] J. Preskill, *Lecture Notes for Physics 229: Quantum Information and Computation* (1998), <https://books.google.com/books?id=Miv8rQEACAAJ>.
- [62] P. Krantz, M. Kjaergaard, F. Yan, T.P. Orlando, S. Gustavsson, and W.D. Oliver, *A Quantum Engineer's Guide to Superconducting Qubits*, *Appl. Phys. Rev.* **6**, 021318 (2019).