

# Virtual Distillation for Quantum Error Mitigation

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

<sup>1</sup>Google Quantum AI, Venice, CA 90291, United States

<sup>2</sup>Department of Chemistry, University of California, Berkeley, CA 94607, United States

<sup>3</sup>Department of Materials, University of Oxford, Parks Road, Oxford OX1 3PH, United Kingdom

<sup>4</sup>Instituut-Lorentz, Universiteit Leiden, 2300 RA Leiden, The Netherlands

<sup>5</sup>Department of Chemistry, Columbia University, New York, NY, USA

(Dated: November 16, 2020)

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  $\rho$ . 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  $\rho$ , exponentially quickly. We analyze the effectiveness of virtual distillation and find that it is governed in many regimes by the behaviour of this pure state (corresponding to the dominant eigenvector of  $\rho$ ). 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.

## 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]. We find it interesting to ask if these additional qubits can be used fruitfully without employing the full machinery of fault-tolerance. This question becomes especially relevant as larger noisy intermediate-scale quantum (NISQ) devices are constructed along the way to a full fault-tolerant device [3]. 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 against errors on NISQ devices, i.e., to efficiently approximate the output that would be produced in the absence of noise using data collected from a noisy device. One broad class of approaches attempts to collect data at a variety of error rates, determine the function which relates the measured value of an observable to the error rate, and extrapolate to the zero noise limit [4–6]. An alternative strategy, introduced in Ref. 4, corrects the estimated expectation values by assuming a particular model for the noise

and expressing the inverse channel as a quasi-probability distribution over modified copies of the original circuit. More recently proposed are techniques that compare classical simulations of Clifford circuits with evaluations of the same circuits on a noisy device. These methods aim to learn enough about the impact of the noise to accurately predict the noise-free expectation values for structurally similar circuits with arbitrary (non-Clifford) gates [7, 8]. Ref. 9 applies a related technique experimentally, using a different family of structurally similar circuits which are easy to simulate by mapping to free fermions. In Ref. 10, O’Brien et al. put forward a version of 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. The machinery of phase estimation can be applied in conjunction with this approach to obtain error-mitigated expectation values for arbitrary observables.

Besides these methods, tools have also been developed for mitigating errors in more specific situations. The quantum subspace expansion has been shown to help mitigate against both coherent and incoherent errors when the state of interest is a ground state (or low-lying excited state) [11–13]. It relies on diagonalizing the Hamiltonian in a basis spanned by some target state and the states obtained by acting on the target state with a collection of perturbing operators. Related to this idea are error-mitigation techniques based on known symmetries of the noise-free state [14–18]. These techniques involve either explicit post-selection based on the symmetries or classical post-processing procedures that achieve the same effect at the cost of extra measurements. There are also a number of works that address the problem of mitigating

\* corresponding author: whuggins@google.com

† corresponding author: jmclean@google.com

measurement errors [19, 20].

Before the modern field of quantum error-correction was developed, an alternative proposal was put forward for stabilizing quantum computations [21–23]. The essence of this approach is to execute  $M$  redundant copies of a computation in parallel and use measurement to project into the symmetric subspace between these copies. This symmetrization can be executed repeatedly, but the simplest application of the idea is to perform it once at the end of the computation. Let  $\rho$  denote the density matrix produced by a noisy quantum computation on  $N$  qubits. As an example, we consider projecting two independently prepared copies of  $\rho$  into the symmetric subspace. We make the important assumption that the noise is stochastic and equal for both copies. This can be accomplished by performing a measurement of the swap operator between the two copies of the system and postselecting on observing the +1 eigenvalue. The postselection succeeds with probability

$$p_{+1} = \frac{1 + \text{Tr}(\rho^2)}{2}, \quad (1)$$

and the reduced state of an individual system register becomes

$$\frac{\rho + \rho^2}{\text{Tr}(\rho + \rho^2)}. \quad (2)$$

This state is purified compared to the original state  $\rho$ : under some mild assumptions about the nature of the noise, it is closer than  $\rho$  to the pure state that would have been produced in the absence of noise.

In separate works, measurements of the swap operator, denoted  $S^{(2)}$ , and its generalizations have been applied to measure Renyi entanglement entropies and other polynomial functions of the density matrix [24–26]. The identity

$$\langle S^{(2)} \rangle = \text{Tr}(\rho^2) \quad (3)$$

has been used in quantum Monte Carlo calculations [27], and in experiments with ultracold atoms to estimate the second Renyi entropy [28]. Related proposals exist for quantum algorithms [29–31] and other many-body experimental platforms [32]. In Ref. 33, Cotler et al. used a similar equality to implement an idea they call “virtual cooling.” By performing a joint measurement on  $M$  copies of a thermal state at inverse temperature  $\beta$  ( $\rho \approx e^{-\beta H}$ ), they were able to estimate expectation values with respect to the thermal state at inverse temperature  $M\beta$  ( $\rho^M \approx e^{-M\beta H}$ ). In this paper, we employ similar tools to Cotler et al., but for the purpose of suppressing errors in a noisy quantum computation.

Earlier work on using symmetrization for error mitigation focused on protocols that prepared an approximately purified state [21–23]. We shall 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

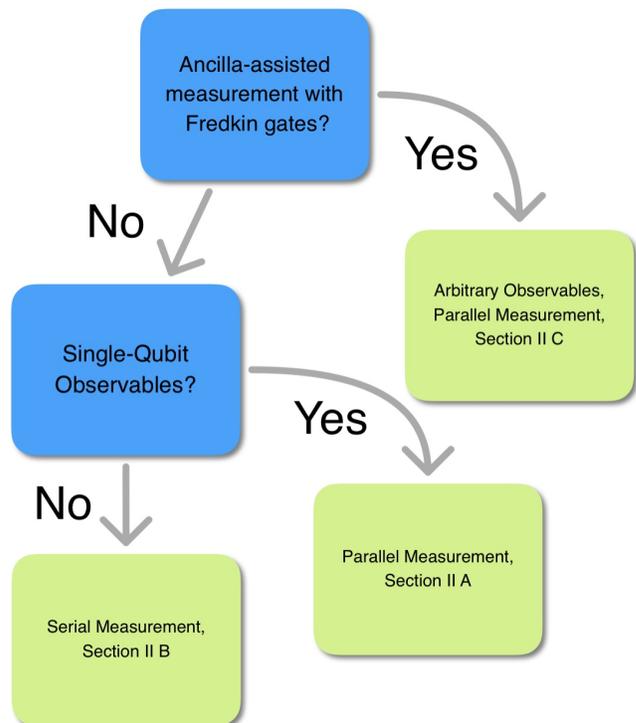


Figure 1. A flowchart that describes the choices involved in selecting between the different variants of virtual distillation. Blue boxes denote questions for the experimentalist to answer about the available quantum resources and problem to be studied, green boxes link to the relevant sections in the text.

distillation. We use the word “virtual” to distinguish our approach from typical state distillation schemes that aim to consume multiple noisy copies of a state to actually prepared a single purified copy [34–36]. To be specific, we shall use collective measurements of  $M$  copies of  $\rho$  to measure expectation values with respect to the state

$$\frac{\rho^M}{\text{Tr}(\rho^M)} = \frac{\sum_i p_i^M |i\rangle\langle i|}{\sum_i p_i^M}, \quad (4)$$

where  $\rho = \sum_i p_i |i\rangle\langle i|$  is a spectral decomposition of  $\rho$ . 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 [21–23, 37].

Our proposed error mitigation technique complements this existing body of literature by offering the opportunity (besides the usual quantum error correction formalism) to make use of additional qubits to enhance the quality of a noisy computation. Furthermore, the technique is extremely 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 are effectively measuring 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 [Section II](#) by giving an overview of the theoretical formalism to be used in this work. We then present three different methods (Secs. [IIB](#), [IIA](#), and [IIC](#)) for virtual distillation that differ by their resource requirements and the types of observables that may be measured. These differences are summarized in [Figure 1](#) for ease of access to a reader wishing to use these techniques. We continue in [Section IID](#) with an analysis of the sample complexity of this technique in the simple case (where one consumes two copies of  $\rho$  to measure  $\text{Tr}(\rho^2 O)$ ) and in a more complex situation (where one wishes to reduce the sample complexity by consuming additional copies).

In [Section III](#), we study the error mitigation performance of virtual distillation in simple cases where analytic expressions can be found. We split the effect of errors into the shift of the closest pure state to the density matrix (the leading eigenvector) away from the target (error-free) state ([Section IIIB](#)), and the shift of the noisy density matrix away from this state ([Section IIIA](#)). Although the second effect may be exponentially suppressed by increasing the number of samples  $K$ , the same is not true for the first effect, which in the worst case limits the performance of virtual distillation to only providing a constant-factor improvement in error rate (as a function of the underlying physical noise rate). To complement this analysis, in [Section IV](#) we present numerical simulations of the virtual distillation of various noisy quantum circuits. We observe here that for some range of noise levels, virtual distillation does achieve an asymptotic error suppression, below the upper bounds suggested in [Section IIIB](#). Finally, in [Section VB](#), we consider the performance of our technique when applied to the stochastic errors that arise during randomized algorithms for real-time evolution.

## II. THEORY

We first establish some assumptions and notation. Throughout this paper we deal with operations that 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 [Appendix G](#).

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 shall 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 bold-faced 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.

The error-mitigated expectation value of an operator  $O$  is the expectation value of  $O$  with respect to the state from [Eq. 4](#),

$$\frac{\text{Tr}(O\rho^M)}{\text{Tr}(\rho^M)}. \quad (5)$$

In order to evaluate the numerator and denominator of this equation, we can make use of the following equality [[25](#), [26](#), [33](#)],

$$\text{Tr}(O\rho^M) = \text{Tr}(O^{\mathbf{i}} S^{(M)} \rho^{\otimes M}). \quad (6)$$

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.,

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

$$S^{(M)} |\psi_1\rangle \otimes |\psi_2\rangle \cdots |\psi_M\rangle = |\psi_2\rangle \otimes |\psi_3\rangle \cdots |\psi_1\rangle. \quad (7)$$

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

$$\begin{aligned} \text{Tr}(O^{\mathbf{1}} S^{(M)} \rho^{\otimes M}) &= \\ \sum_{i_1, i_2, \dots, i_M, j_1, j_2, \dots, j_M, k} O_{k, j_1} \delta_{j_2, i_1} \cdots \delta_{j_1, i_M} \rho_{i_1, k} \cdots \rho_{i_M, j_M} &= \\ \sum_{i_1, i_2, \dots, i_M, k} \rho_{i_1, k} O_{k, i_M} \rho_{i_M, i_{M-1}} \cdots \rho_{i_2, i_1} &= \\ \text{Tr}(O\rho^M). \end{aligned} \quad (8)$$

In [Figure 2](#) we present a diagrammatic representation of [Eq. 6](#) for the case where  $M = 3$  (note that we have commuted  $\rho^{\otimes 3}$  with  $S^{(3)}$  in the diagram).

In the following subsections, we shall present three strategies for measuring the numerator and denominator of [Eq. 5](#). We capture the main differences between these methods in [Figure 1](#). In [Section IIA](#) and [Section IIB](#) we outline two related strategies that don't require ancilla qubits or Fredkin gates, and are thus especially suitable for near-term implementation. The approach of [Section IIA](#) allows for the simultaneous measurement of the error-mitigated value of the  $Z$  operator on each qubit,

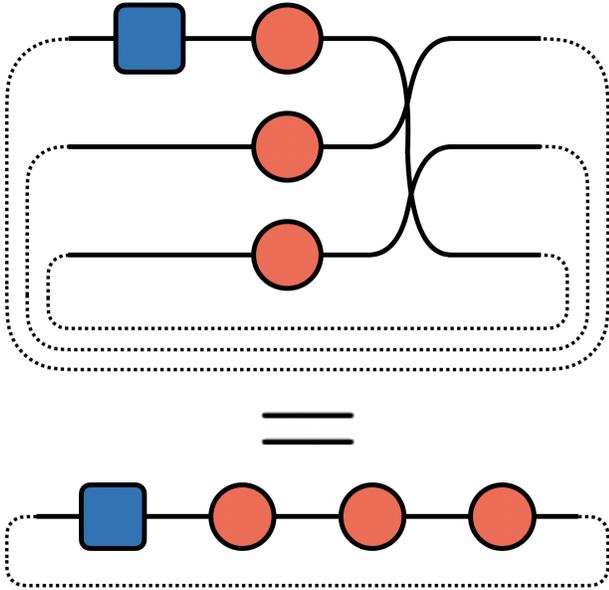


Figure 2. A diagrammatic representation of Eq. 6 with  $M = 3$  and  $i = 1$  using tensor network notation [38–40]. The blue square represents the operator  $O^1$ , each red circle represent a copy of the state  $\rho$ , and the connections between the shapes indicate indices which are summed over. 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  $\text{Tr}(O^1 \rho^{\otimes 3} S^{(3)})$ ; we commuted  $\rho^{\otimes 3}$  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. 8.

whereas Section II B presents a strategy that doesn't restrict the form of the operator but also doesn't support simultaneous measurement. The version presented in Section II C requires an ancilla-assisted measurement, but it has the advantage that it allows for any of the usual techniques for measuring a large number of commuting observables simultaneously. Finally in Section II D, we determine the variance of the estimators from Section II A and Section II C. We conclude that subsection by proving that there exists generalizations of our approach that can further reduce the variance (and therefore, the number of circuit repetitions required for a desired precision).

#### A. Measurement by Diagonalization with Symmetrized Observables

In this subsection, we present a straightforward strategy for evaluating the quantities in the numerator and denominator of Eq. 5 applicable when the operator  $O$  is the Pauli  $Z$  operator acting on a single qubit. Other single-qubit observables can be accessed by applying the appropriate single-qubit rotations before the virtual distillation procedure. Here we follow the approach pre-

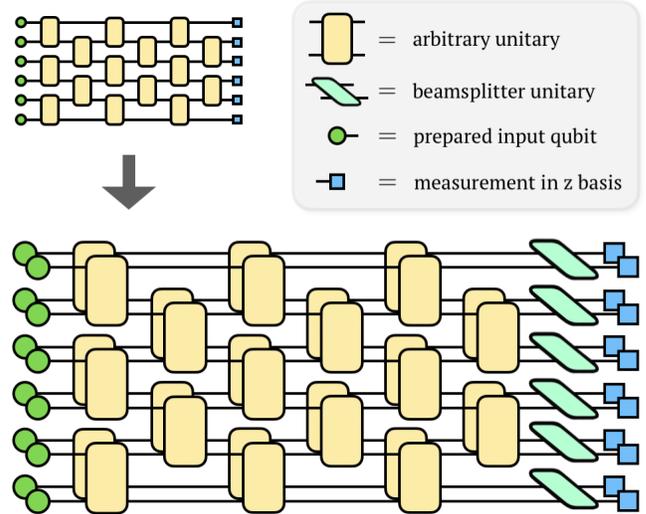


Figure 3. A circuit diagram of our approach applied to a six-qubit circuit. We use twice the number of qubits to independently perform two copies of the original circuit. We then apply a single layer of the two-qubit gates specified in Eq. 14 before measuring each qubit in the computational basis. This allows us to estimate the error-mitigated expectation values for all single-site  $Z$  operators.

sented in Ref. 33, except that we translate their work into the language of qubits rather than bosonic systems. We focus in this section on the version that uses  $M = 2$  copies of  $\rho$ . In Appendix A we discuss the generalization to higher values of  $M$ . We show a schematic of this approach in Figure 3. Essentially, all that we have to do is apply a single layer of two-qubit gates and measure each of the  $2N$  qubits in the computational basis.

Rather than using the relation in Eq. 6 directly, we instead define a symmetrized version of our observable,

$$O^{(M)} = \frac{1}{M} \sum_{i=1}^M O^i. \quad (9)$$

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

$$Z_k^{(2)} = \frac{1}{2} (Z_k^1 + Z_k^2). \quad (10)$$

It is straightforward to use Eq. 6 to show that

$$\frac{\text{Tr}(O \rho^M)}{\text{Tr}(\rho^M)} = \frac{\text{Tr}(O^{(M)} S^{(M)} \rho^{\otimes M})}{\text{Tr}(S^{(M)} \rho^{\otimes M})}. \quad (11)$$

Using the symmetrized observable is advantageous because

$$[O^{(M)}, S^{(M)}] = 0, \quad (12)$$

or, specifying to the case we treat in this subsection,  $[Z_k^{(2)}, S^{(2)}] = 0$ .

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 then define

$$B^{(2)} = \bigotimes_{i=1}^M B_i^{(2)}. \quad (13)$$

We give a matrix representation for this gate below, noting that there is some freedom in the choice of phases for the matrix elements,

$$B_i^{(2)} := \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \frac{\sqrt{2}}{2} & -\frac{\sqrt{2}}{2} & 0 \\ 0 & \frac{\sqrt{2}}{2} & \frac{\sqrt{2}}{2} & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}. \quad (14)$$

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

$$B^{(2)}S_i^{(2)}B^{(2)\dagger} \rightarrow \frac{1}{2}(1 + Z_i^1 - Z_i^2 + Z_i^1Z_i^2), \quad (15)$$

$$B^{(2)}Z_k^{(2)}S_k^{(2)}B^{(2)\dagger} \rightarrow \frac{1}{2}(Z_k^1 + Z_k^2). \quad (16)$$

This diagonalization is particularly easy to implement when each qubit from the first copy of  $\rho$  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. 11 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 could instead access an arbitrary single-qubit observable on each qubit. We capture this process diagrammatically in Figure 3.

## B. Measurement by Diagonalization with a Non-Symmetrized Observable

Measuring observables with support on more than one qubit is less straightforward. Ref. 33 solved this issue by using an approach like the one we describe below in Section II C. Here we present an alternative solution that doesn't require the use of ancilla-assisted measurement. The challenge arises due to the use of Eq. 11, in particular, the choice to use the symmetrized version of  $O$  defined in Eq. 9. Using the symmetrized version of a multi-qubit 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  $O = Z_iZ_j$ . Our arguments hold equally well for any other operator composed of a tensor product of single-qubit Paulis. Taking the product of the symmetrized observable and the swap operator yields

$$O^{(2)}S^{(2)} = \frac{1}{2}(Z_i^1Z_j^1 + Z_i^2Z_j^2)S^{(2)}. \quad (17)$$

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. 11 to determine the corrected expectation value of  $O$ , we can instead use the non-symmetrized form introduced in Eq. 5. Returning to our example where  $O = Z_iZ_j$ , we see that we need to estimate the numerator and denominator of

$$\frac{\text{Tr}(Z_i^1Z_j^1S^{(2)}\rho^{\otimes 2})}{\text{Tr}(S^{(2)}\rho^{\otimes 2})}. \quad (18)$$

Unlike the symmetrized observable of Eq. 17, the operator  $Z_i^1Z_j^1S^{(2)}$  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).  $Z_i^1Z_j^1S^{(2)}$  is not Hermitian, but because it is unitary, we can still estimate  $\text{Tr}(Z_i^1Z_j^1S^{(2)}\rho^{\otimes 2})$  by applying a circuit to diagonalize it and measuring in the computational basis. As  $Z_i^1Z_j^1S^{(2)}$  factorizes into a product of two-qubit operators, the circuit that diagonalizes it does as well.

Note that because  $Z_i^1Z_j^1$  does not commute with  $S^{(2)}$ , we will be unable to simultaneously estimate the numerator and denominator of Eq. 18. We will also be 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 our error mitigation proposal, but the inability to parallelize the measurement of commuting multi-qubit 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, individual operators would have to be measured separately even if they commute, increasing the overall number of circuit repetitions required for many applications. In Appendix A, we briefly explain how this variant of our error mitigation strategy generalizes to  $M > 2$  copies.

### 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. 33 for estimating the expectation value of an observable  $O$  with respect to the state  $\rho^2/\text{Tr}(\rho^2)$ . Like the method we discussed in Section II A, we shall do this by measuring the numerator and denominator of Eq. 11. Unlike that method, this approach uses a non-destructive measurement of the swap operator ( $S^{(2)}$ ). The main reward for this added complexity is that this variant of virtual distillation doesn't 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, 41, 42]. While we focus on the  $M = 2$  copy version here, in Appendix B we discuss the generalization to  $M \geq 3$  copies where there may be some additional advantages compared to the generalizations of the destructive measurement techniques presented above.

To use this method, we begin with two system registers, each in the state  $\rho$ , as well as an ancilla qubit in the  $|0\rangle$  state. We then perform a non-destructive measurement of  $S^{(2)}$  in the standard way, using the so-called swap or Hadamard test [24, 25, 43]. 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 GHZ state using a series of CNOT gates) in order to deal with the restricted connectivity of a near-term device.

It isn't 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 Section 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. 33. Furthermore, it allows us to make use of measurements of  $O$  on both copies of  $\rho$  and also simultaneously estimate the numerator and denominator of Eq. 11, leading to a relatively sample-efficient scheme.

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

sum,

$$\begin{aligned} \rho^{\otimes 2} &= \sum_{ij} p_i p_j |i\rangle\langle i| \otimes |j\rangle\langle j| \\ &= \sum_i p_i^2 |i\rangle\langle i| \otimes |i\rangle\langle i| + \sum_{i \neq j} p_i p_j |i\rangle\langle i| \otimes |j\rangle\langle j|. \end{aligned} \quad (19)$$

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 anti-symmetric states,

$$|i\rangle|j\rangle = \frac{1}{2}(|i\rangle|j\rangle + |j\rangle|i\rangle) + \frac{1}{2}(|i\rangle|j\rangle - |j\rangle|i\rangle). \quad (20)$$

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.

### D. Sample Efficiency

The practical utility of these techniques as error-mitigation tools will be partly determined by the number of samples necessary to determine the corrected expectation values to within some target precision  $\epsilon$ . In this section, we present some calculations of the variance of our estimator for the corrected expectation values of an observable  $O$ . We focus on the  $M = 2$  case and the methods discussed in Section II A and Section II C specifically. We leave the derivation to Appendix C and simply give the expression for the approximate variance,

$$\begin{aligned} &\frac{1}{R} \left( \frac{1}{\text{Tr}(\rho^2)^2} \left( \frac{1}{2} \text{Tr}(\rho O^2) + \frac{1}{2} \text{Tr}(\rho O)^2 - \text{Tr}(\rho^2 O^2) \right) \right. \\ &\quad - 2 \frac{\text{Tr}(\rho^2 O)}{\text{Tr}(\rho^2)^3} (\text{Tr}(\rho O) - \text{Tr}(\rho^2 O) \text{Tr}(\rho^2)) \\ &\quad \left. + \frac{\text{Tr}(\rho^2 O)^2}{\text{Tr}(\rho^2)^4} (1 - \text{Tr}(\rho^2)^2) \right). \end{aligned} \quad (21)$$

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

$$\frac{1}{2R} (\text{Tr}(\rho O^2) - \text{Tr}(\rho O)^2), \quad (22)$$

exactly what one would expect when averaging  $2R$  independent measurements of  $O$ . As the purity of  $\rho$  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. As we shall see, 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

$$\text{Tr}(\rho^3) \ll \text{Tr}(\rho^2) \ll 1. \quad (23)$$

We now make the assumption that the level of error mitigation offered by measuring  $\rho^M$  is sufficient but we have  $K \gg M$  copies of  $\rho$  available. For simplicity, we'll focus on the case where  $M = 2$ . We argue that there exist generalizations of our approach involving a collective measurement of all  $K$  copies of  $\rho$  that perform better than naively parallelizing our strategy. We hope that future work can develop a practical version of these generalizations in order to make good use of the larger devices we expect to see later in the NISQ era.

The naive approach we hope to beat consists of taking  $\lfloor K/2 \rfloor$  pairs and running the protocol described above in parallel, averaging the results. This would result in an estimator whose variance is scaled by a factor of  $1/\lfloor K/2 \rfloor \approx 2/K$ . For simplicity, we focus on the variance of our estimator for the quantity that appears in the numerator of Eq. 11 rather than the ratio itself. In Appendix C 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  $K$  copies in parallel is exactly

$$\frac{1}{K}(\text{Tr}(\rho O^2) + \text{Tr}(\rho O)^2 - 2\text{Tr}(\rho^2 O)^2), \quad (24)$$

for even values of  $K$ . We shall prove 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  $K$  copies. We do so by providing an operator  $\tilde{O}$  with the desired expectation value and calculating its variance.

First, we define the operator

$$\tilde{O} = \frac{1}{\binom{K}{2}} \sum_{i=1}^K \sum_{j>i} \frac{1}{2}(O^i + O^j)S^{(i,j)}, \quad (25)$$

where we use  $S^{(i,j)}$  to denote the swap operator specifically between subsystems  $i$  and  $j$ . The equation

$$\text{Tr}(\tilde{O}\rho^{\otimes K}) = \text{Tr}(O\rho^2) \quad (26)$$

follows from our work above. We shall assume here that  $O$  is a Pauli operator acting on one or more qubits. Therefore it is Hermitian, self-inverse, and has eigenvalues  $+1$  and  $-1$ .

$\tilde{O}$  is a Hermitian operator and we can compute its variance with respect to the state  $\rho^{\otimes K}$ ,

$$\text{Var}(\tilde{O}) = \langle O^2 \rangle - \langle O \rangle^2. \quad (27)$$

We carry out this calculation in Appendix D, ultimately finding that

$$\text{Var}(\tilde{O}) \leq \frac{2}{K(K-1)} + \frac{7(K-2)\text{Tr}(\rho^3)}{K(K-1)}. \quad (28)$$

When  $\text{Tr}(\rho^3)$  is small, the second term is suppressed and there is a regime where the variance of this operator shrinks quadratically with  $K$ . The naive approach, where we perform  $\lfloor \frac{K}{2} \rfloor$  independent calculations on separate pairs results in an estimator whose variance is suppressed only linearly in  $K$ .

### III. PERFORMANCE UNDER DIFFERENT NOISE MODELS

Here we develop some simple models for understanding the different factors affecting the performance of virtual distillation. Specifically, to understand the potential benefit of our approach using the minimal setup, we shall consider the fidelity of

$$\frac{\rho^2}{\text{Tr}[\rho^2]} \quad (29)$$

with the ideal state generated by noiseless evolution. Note that this ignores errors introduced by the overhead of implementing the Hadamard test. It will be shown in the numerical studies that the performance of the approach can be essentially predicted by the combination of two contributions, and it is instructive to consider them separately. The first is the performance under noise that maps the ideal states to states orthogonal to it, leaving the dominant eigenvector as the ideal state. The second is the effect of errors that lead to states non-orthogonal to the ideal state, and cause drift in the dominant eigenvector of the density matrix. The essential behavior of the approach 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. While this does not ultimately determine our noise floor, it is instructive and predictive of behaviors in some regimes.

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

$$\begin{aligned} \rho &= (1-p)^G \rho_0 + (1-p)^{G-1} p \sum_{j=1}^G \rho_{j_1} \\ &+ (1-p)^{G-2} p^2 \sum_{j_1 \neq j_2} \rho_{j_1, j_2} \\ &+ (1-p)^{G-3} p^3 \sum_{j_1 \neq j_2 \neq j_3} \rho_{j_1, j_2, j_3} + \dots \end{aligned} \quad (30)$$

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

$$\text{Tr}[\rho^2] = ((1-p)^2 + p^2)^G. \quad (31)$$

The fidelity with the ideal state  $\rho_0$  is

$$\frac{\text{Tr}[\rho_0 \rho^2]}{\text{Tr}[\rho^2]} = \frac{(1-p)^{2G}}{((1-p)^2 + p^2)^G} \quad (32)$$

$$\simeq 1 - Gp^2 + O[G^2 p^4]. \quad (33)$$

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

The result is similar in the case of  $M$  copies:

$$\frac{\text{Tr}[\rho_0 \rho^M]}{\text{Tr}[\rho^M]} = \frac{(1-p)^{MG}}{((1-p)^M + p^M)^G} \quad (34)$$

$$\simeq 1 - Gp^M + O[(Gp^M)^2]. \quad (35)$$

## B. Non-Orthogonal 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 reality, this assumption will not hold, and we will see numerically that it is this drift that limits the maximum upside of the approach. 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. The effect can both limit the ultimate effectiveness of virtual distillation, and damage the performance at smaller levels of protection. We shall explore this effect numerically in [Section IV](#), but first we consider it analytically.

Let us consider a state  $\rho$  in the middle of a noisy preparation circuit, and allow for  $\rho$  to already be somewhat distorted by noise. Writing  $\rho$  in its own eigenbasis, we have

$$\rho = \sum_i \epsilon_i |i\rangle\langle i|. \quad (36)$$

We wish to consider the impact of a subsequent noise channel defined in terms of a set of Kraus operators,

$$\rho \rightarrow p_0 \rho + \sum_{j \neq 0} p_j K_j \rho K_j^\dagger. \quad (37)$$

Note that we have demanded a representation of the channel where  $K_0$  is the identity matrix in order to simplify our analysis. Now let  $V$  denote the change in the density matrix induced by this channel,

$$V = (1-p_0)\rho - \sum_{j \neq 0} p_j K_j \rho K_j^\dagger. \quad (38)$$

Now we make the assumption that we are in the low-error regime. This implies that the eigenvalue of the dominant eigenvector is much larger than the other eigenvalues. Furthermore, assuming that the probability of errors at any particular step is low entails that  $V$  is small. We can therefore analyze the drift in the dominant eigenvector using non-degenerate perturbation theory. As noise induces excitations to other density matrix states, the drift in the dominant eigenvector is predicted to change at first order by

$$|D\rangle \approx |0\rangle + |R\rangle, \quad (39)$$

$$|R\rangle = \sum_{i \neq 0} \frac{\langle i|V|0\rangle}{\epsilon_0 - \epsilon_i} |i\rangle. \quad (40)$$

Here we use  $|D\rangle$  to denote the dominant eigenvector corrected to first order and  $|R\rangle$  to denote the correction.

Normalizing  $D$  to second order, we have

$$|D\rangle = (1 - \frac{1}{2} \langle R|R\rangle)(|0\rangle + |R\rangle), \quad (41)$$

and we can now compute the trace distance between  $|D\rangle$  and  $|0\rangle$ . Repeatedly dropping the terms inside the square root that are higher than second-order, we have

$$\begin{aligned} T(|D\rangle, |0\rangle) &\approx \frac{1}{2} \text{Tr}(\sqrt{(|0\rangle\langle R| + |R\rangle\langle 0| - \langle R|R\rangle |0\rangle\langle 0|)^2}) \\ &\approx \frac{1}{2} \text{Tr}(\sqrt{\langle R|R\rangle |0\rangle\langle 0| + |R\rangle\langle R|}) \\ &= \frac{1}{2} \text{Tr}(\sqrt{\langle R|R\rangle |0\rangle\langle 0| + \langle R|R\rangle \frac{|R\rangle\langle R|}{\langle R|R\rangle}}) \\ &= \sqrt{\frac{\langle R|R\rangle}{2}}. \end{aligned} \quad (42)$$

Now let us expand  $\langle R|R\rangle$  in terms of the Kraus operators

of our noise model.

$$\begin{aligned}
\langle R|R \rangle &= \sum_{i \neq 0} \frac{1}{(\epsilon_0 - \epsilon_i)^2} \langle 0|V^\dagger|i\rangle\langle i|V|0\rangle \\
&= \sum_{i \neq 0} \frac{1}{(\epsilon_0 - \epsilon_i)^2} \\
&\quad \left| \langle i| \left( (1-p_0)|0\rangle\langle 0| + \sum_{j \neq 0} p_j K_j |0\rangle\langle 0| K_j^\dagger \right) |0\rangle \right|^2 \\
&= \sum_{i \neq 0} \frac{1}{(\epsilon_0 - \epsilon_i)^2} \left| \sum_{j \neq 0} p_j \langle i|K_j|0\rangle\langle 0|K_j^\dagger|0\rangle \right|^2.
\end{aligned} \tag{43}$$

We can see that, in the general case, we expect a non-zero contribution to the trace distance at first order. Because  $|D\rangle\langle D| \approx \rho^2 / \text{Tr}(\rho^2)$  in the low-noise regime, this will effectively set a floor for how well our method can correct errors. We therefore expect that, in the general case, 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

$$\left| \sum_{j \neq 0} p_j \langle i|K_j|0\rangle\langle 0|K_j^\dagger|0\rangle \right|^2. \tag{44}$$

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, the quantity in Eq. 44 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 Section 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 Eq. 44 is nonzero.

In order to shed some light on when this might be true, it can be helpful to ask when we might expect the quantity in Eq. 44 to be near zero. It is clear that this quantity must be zero if one of two conditions hold:

$$K_j|0\rangle \propto |0\rangle \tag{45}$$

$$\langle 0|K_j^\dagger|0\rangle = 0. \tag{46}$$

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. Note that this reduces the error sensitivity to second-order only; pairs of nearby errors in time may lead to a combined error channel which does not satisfy Eq. 46. 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 order to examine the effectiveness of our approach on circuits of interest, we present numerical simulations of our approach applied to three different model systems. We first consider two classes of random circuits, chosen to help develop intuition about the performance of our approach in different limits. We then turn towards a condensed matter application, the simulation of the dynamics of a one dimensional spin chain following a quantum quench to an anisotropic Heisenberg Hamiltonian.

We have found it illuminating to focus on characterizing 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 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. We can calculate the expected number of errors ( $E$ ) in terms 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. E3),

$$E = 2pG. \tag{47}$$

Therefore, in the various figures we present, we can vary the expected number of errors in two distinct ways. In some cases, we fix a particular circuit and vary the error rate per gate. In other cases, we fix a particular error rate and vary the number of two-qubit gates by increasing the circuit depth.

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 that our error mitigation approach accesses. The trace distance leads to a natural bound in the error for the expectation value of an arbitrary observable,

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

where  $O$  is an observable with operator norm  $\|O\|$   $T(-, -)$  denotes the trace distance. We also present some data in Figure 7 that shows the actual error incurred in estimating the expectation value of particular observables using our technique. Finally, we plot the expected overhead (in terms of the number of additional measurements required for a fixed precision) for one of the concrete examples we consider.



Figure 4. The error in the unmitigated noisy states ( $M = 1$ ) and the states accessed by our error mitigation technique ( $M = 2, 3$ ) for a variety of non-entangling 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 non-entangling circuits, the eigenvalue floor vanishes and we see exponential suppression in the number of copies. In this case, more qubits offer slightly more protection as predicted.

### 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 exactly the same as the first class, except that we effectively remove the two-qubit gates by replacing them with identity operations. For these non-entangling 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 behaviour of the non-entangling 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 together into a single application per qubit with a larger effective error rate. We carry this procedure out analytically in Appendix E 1 and just summarize the result here. The most important consequence is that the dominant eigenvector of the density matrix exactly corresponds to the noiseless evolution of the state. Furthermore, the distribution of eigenvalues of the density matrix decays exponentially (albeit with some degeneracies arising from combinatorial factors). This leads us to expect behaviour similar to that of the phenomenological noise model we considered in Section III A.

In Figure 4, we plot the trace distances between

the ideal states generated by noiseless evolution and the states obtained by noisy evolution of these non-entangling random circuits (blue curve). We consider a variety of different circuit depths, 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. The results plotted here include data from simulations where we fix the size of the circuit and vary the error rate as well as simulations where we fix the error rate and vary the circuit depth. In each case, we consider a particular randomly chosen member from the ensemble of non-entangling scrambling circuits we described above.

One thing that is immediately apparent in Figure 4 is the way in which data from this variety of simulations collapses together when we plot the error in terms of trace distance on the y-axis and the expected number of gate errors on the x-axis. Notice that when the expected number of errors is not too large, 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 Section III A, the results are broadly consistent. Here we observe the same quadratic suppression for the  $M = 2$  case when compared with the behaviour of the unmitigated state  $M = 1$ . Furthermore, we observe the cubic suppression of the impact of errors that we might expect when applying our technique with  $M = 3$  copies of the state.

In Figure 5 and Figure 6, we present plots that explore the behaviour of the first family of random circuits, the entangling random circuits on a one-dimensional line of qubits. When we considered the non-entangling random circuits we found that error (quantified by the trace distance to the ideal state) depended mostly on the system size and the expected number of gate errors. This was true regardless of whether or the expected number of errors was varied by changing the circuit depth or by changing the error rate per-gate. Here we observe slightly different behaviour between these two cases. Therefore, in Figure 5 we present two plots that show the effect of varying the expected number of errors by varying the circuit depth at fixed error rates. In Figure 6 we make the opposite choice, fixing the number of gates in the circuits and varying the expected number of errors by adjusting the per-gate error rate.

Turning first to Figure 5, we see that besides this change, we have also plotted a curve corresponding to trace distance between the dominant eigenvector of the density matrix (red dashed curve,  $\lim_{M \rightarrow \infty} \rho^M/\text{Tr}(\rho^M)$ ) and the ideal state. This curve was absent from Figure 4 because the dominant eigenvector of the noisy state from those simulations doesn't drift away from the ideal state generated by noiseless evolution. Here we see that the error in the dominant eigenvector effectively sets a floor

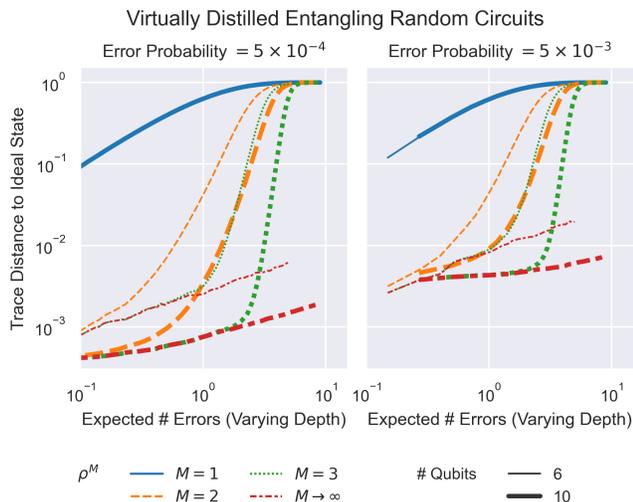


Figure 5. The error in the unmitigated noisy states ( $M = 1$ ) and the states accessed by our error mitigation technique ( $M = 2, 3$ ) for a variety of entangling random circuits. We plot the trace distance to the state obtained from noiseless evolution as a function of the expected number of single-qubit 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  $5 \times 10^{-4}$  (left panel) or  $5 \times 10^{-3}$  (right panel). Here we also plot the error for the state corresponding to the dominant eigenvector of the density matrix ( $M \rightarrow \infty$ ), and we see that even for this ideal case of random circuits, the dominant eigenvector determines the noise floor beyond which we cannot improve, independent of the number of copies. This floor drops as the size of the system increases.

for the minimal error achievable by our method for any value of  $M$ . Interestingly, we see that this distance grows slowly with increasing circuit depth. Furthermore, both the absolute magnitude and the rate of growth appear to be suppressed with system size. In Section III B we showed that this trace distance can be understood up to second order in the error rate 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 lightcone 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 Figure 5.

In Figure 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) are orders of magnitude smaller than the errors in the unmitigated state. The scaling of the



Figure 6. The error in the unmitigated noisy states ( $M = 1$ ) and the states accessed by our error mitigation technique ( $M = 2, 3$ ) 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 6 and 10 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's 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 ( $M \rightarrow \infty$ ).

trace distance, however, is linear in the error rate. This matches the behaviour we would expect from the analysis of Section III B. The leading order term in the trace distance scales linearly with the probability of error. As in Figure 5, we see that the error in the dominant eigenvector sets a floor for the performance of our method at finite  $M$ . We also see that this floor is suppressed by increasing the system size, consistent with the fact that a deep enough scrambling circuit starts to approximate a Haar random circuit to some order, where all but a few circuit errors will lead to exponentially small coupling elements.

## 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.,  $|0101\rangle$ ,

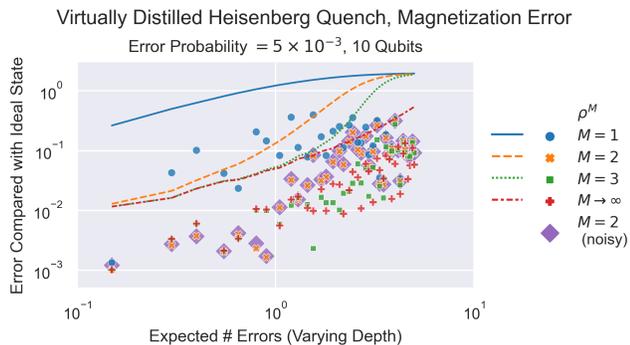


Figure 7. The average error in the single-site magnetization in the unmitigated noisy states ( $M = 1$ ) and the states accessed by our error mitigation technique ( $M = 2, 3$ ) for states generated by the Trotterized time evolution of a Heisenberg model. We plot the actual average errors we calculate 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. 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  $5 \times 10^{-3}$ . We see that for this specific observable, the trace distance bounds are pessimistic by roughly an order of magnitude, though 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 Section II A with noise has a negligible effect on the corrected expectation value.

and simulate the time evolution under the Hamiltonian

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

Here we have chosen the parameters,  $J_x = J_y = 1.0$ ,  $J_z = 1.5$ ,  $h = 1.0$ , in order to match a previously studied family of non-integrable models [45], although we take open boundary conditions rather than periodic ones. We approximate the time evolution under this Hamiltonian by Trotterization with a timestep 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 Figure 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 system with a fixed single-qubit depolarizing probability of  $5 \times 10^{-3}$ . Alongside these bounds

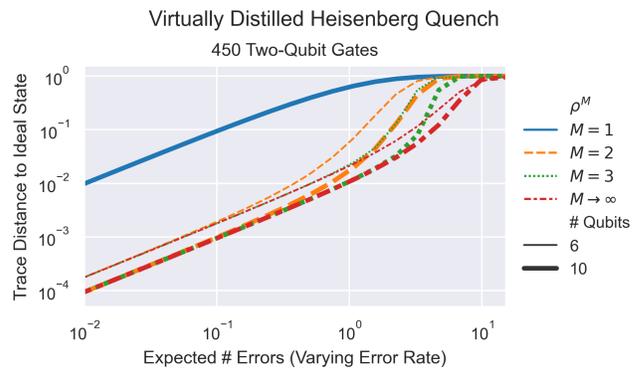


Figure 8. The error in the unmitigated noisy states ( $M = 1$ ) and the states accessed by our error mitigation technique ( $M = 2, 3$ ) for states generated by the Trotterized time evolution of a Heisenberg model. Also shown is the error for the state corresponding to the dominant eigenvector of the density matrix ( $M \rightarrow \infty$ ). We plot the trace distance to the state obtained from noiseless evolution as a function of the expected number of single-qubit depolarizing errors. We show this data for 6 and 10 qubit systems (differentiated by the size of the markers). The expected number of errors is varied by changing the error rate per-gate, fixing the number of two-qubit gates to be 450. As we increase the system size from 6 qubits to 10, we observe that the error (quantified by the trace distance to the ideal state) decreases for the error-mitigated states ( $M > 1$ ).

(plotted using solid and dashed curves) we also plot the actual average error in the single-site magnetization (averaged over the 6 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 Section II A using large purple diamonds. For this second calculation, we simulate the application of the six two-qubit gates (Eq. 14) 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 Figure 6, in this figure we fix the number of two-qubit 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 again 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 error mitigated states ( $M > 1$ ) is suppressed by a constant factor relative

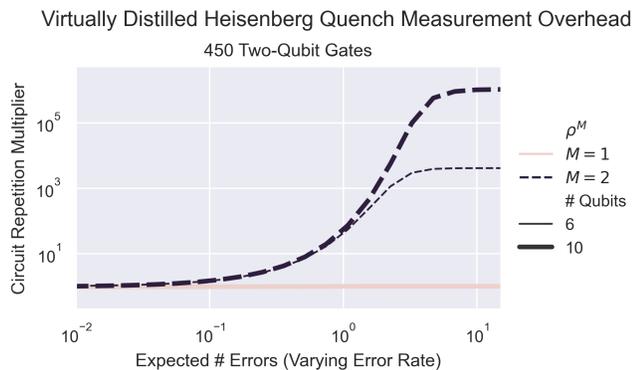


Figure 9. The average overhead in the number of measurement repetitions required to measure the single-site magnetization (of a Heisenberg model) with virtual distillation ( $M = 2$ ). The overhead is calculated by taking the ratio of the time required using our technique to achieve a fixed target precision 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 6 and 10 qubit systems, denoting the larger system using larger markers. In calculating the overhead, we make use of Eq. 21. As the error rate grows beyond  $O(1)$  errors, the overhead increases dramatically, with the larger system seeing the greatest inflation.

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 rate of increase with the system size are smaller than we observed for the one-dimensional scrambling circuits plotted in Figure 6.

In Figure 9, we consider the cost of performing the two-copy ( $M = 2$ ) version of our proposal for the same systems considered in Figure 8. We do this using the expression for the variance of our estimator presented in Eq. 21 and derived in Appendix C. Using this expression, we calculate the variance of our error-mitigated estimators for the magnetization at each site  $\{Z_i\}$  and then average the resulting variances. 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 by our scheme and the number of measurements required to measure the unmitigated expectation values (assuming the same target precision). This quantity therefore neatly encapsulates the overhead incurred by our scheme.

When the expected number of errors is small, we see that our scheme barely increases the number of measurements required. It is only as the number of errors grows larger than one that the measurement cost rises dramati-

cally. We note that Eq. 21 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 Figure 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 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 [48], randomized Trotter [49], and combinations thereof [50]. These methods have benefits in some situations over their deterministic counterparts, and are of particular interest when one is focused on minimizing the coherent depth of a circuit required for a given approximation error.

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. Hence, if we imagine deviations from this pure state as incoherent errors, it is natural to wonder if incoherent mitigation techniques such as virtual distillation can be effective in accelerating convergence to this ideal pure state. In this section, we establish an advantage concretely through numerical demonstrations of virtual distillation applied to qDRIFT. 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 formulae 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 forwards in time under this Hamiltonian term, for a fixed timestep. Repeating this process

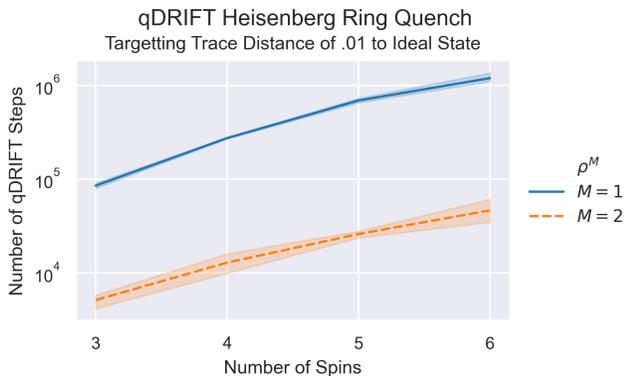


Figure 10. qDRIFT coherent cost reduction through virtual distillation in the Heisenberg model. Here we show the number of coherent qDRIFT steps required to reach a target trace distance, with and without virtual distillation using 2 copies. We see that there is a consistent reduction of at least 16x in the number of required steps. When accounting for the overhead of using two copies, this amounts to a 8x reduction in the coherent space-time volume used to reach the same error rate.

a number of times generates a product formula that provides an approximation to the time evolution operator. While it is possible to consider a single qDRIFT circuit instance in isolation, as in Ref. 51, the method is geared towards repeating the calculation a number of times. Using qDRIFT this way, with a number of independently generated circuits, results in a quantum channel that approximates the exact evolution and a final state described density matrix. 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 than 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  $\eta$  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  $\eta$  increases, the resulting density matrix converges to a pure state. It will be our aim to understand how our virtual distillation technique can reduce the coherent space-time volume required, by reducing this factor  $\eta$  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,  $\eta$ , required to reach a target accuracy. For this, we choose a Heisenberg Hamiltonian with up to 6 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  $h = 1$ . As virtual distillation focuses on mitigation of observables, a good proxy for improvement in error is the trace distance, which bounds the maximum deviation on observables in a system. 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 16x. If we account for the space overhead of using two copies, this still amounts to a space-time advantage of 8x through the use of virtual distillation. 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 showed how techniques for using multiple copies of a state to access polynomials of the density matrix can be used for error mitigation. We studied the effectiveness of this approach for two analytically tractable noise models and characterized its limit in terms of the dominant eigenvector of the noisy density matrix. We numerically demonstrated reductions in the error (quantified by the trace distance to the noise-free state) of up to three orders of magnitude for a collection of relatively small model systems. Furthermore, we showed that this error suppression is enhanced as the system size or the speed of information scrambling grows. We also considered 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. As the number of copies,  $M$ , used for virtual distillation is increased, the effective state being measured tends to the dominant eigenvector of the noisy density matrix exponentially quickly. Therefore, the practical utility of our technique depends mainly on the error in the dominant eigenvector and the number of samples required to apply the technique. Our 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. Even as we 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 F.

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  $\rho$ . In Section IID, 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 towards 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. Is it possible to remove this source of error by using variational optimization or other, complementary, error mitigation techniques? Studying this in the context of the real noise experienced in hardware will be especially important and likely lead to new insights.

In the process of preparing this work, a related paper, Ref. 52 appeared in the literature. Our work obtains different conclusions to theirs in error scaling due to our explicit consideration of errors that prepare states non-orthogonal to the original state.

#### Acknowledgements

The authors are extremely grateful to Nathan Wiebe for helpful discussions at various stages.

- 
- [1] D Aharonov and M Ben-Or, “Fault-tolerant quantum computation with constant error,” in *Proceedings of the twenty-ninth annual ACM symposium on Theory of computing*, STOC ’97 (Association for Computing Machinery, New York, NY, USA, 1997) pp. 176–188.
  - [2] Austin G Fowler, Matteo Mariantoni, John M Martinis, and Andrew N Cleland, “Surface codes: Towards practical large-scale quantum computation,” *Phys. Rev. A* **86**, 032324 (2012).
  - [3] John Preskill, “Quantum computing in the NISQ era and beyond,” *Quantum* **2**, 79 (2018).
  - [4] Kristan Temme, Sergey Bravyi, and Jay M Gambetta, “Error mitigation for Short-Depth quantum circuits,” *Phys. Rev. Lett.* **119**, 180509 (2017).
  - [5] Suguru Endo, Simon C Benjamin, and Ying Li, “Practical quantum error mitigation for Near-Future applications,” *Phys. Rev. X* **8**, 031027 (2018).
  - [6] Abhinav Kandala, Kristan Temme, Antonio D Córcoles, Antonio Mezzacapo, Jerry M Chow, and Jay M Gambetta, “Error mitigation extends the computational reach of a noisy quantum processor,” *Nature* **567**, 491–495 (2019).
  - [7] Armands Strikis, Dayue Qin, Yanzhu Chen, Simon C Benjamin, and Ying Li, “Learning-based quantum error mitigation,” [arXiv:2005.07601](https://arxiv.org/abs/2005.07601) (2020).
  - [8] Piotr Czarnik, Andrew Arrasmith, Patrick J Coles, and Lukasz Cincio, “Error mitigation with clifford quantum-circuit data,” [arXiv:2005.10189](https://arxiv.org/abs/2005.10189) (2020).
  - [9] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C Bardin, Rami Barends, Andreas Bengtsson, Sergio Boixo, Michael Broughton, Bob B Buckley, David A Buell, Brian Burkett, Nicholas Bushnell, Yu Chen, Zijun Chen, Yu-An Chen, Ben Chiaro, Roberto Collins, Stephen J Cotton, William Courtney, Sean Demura, Alan Derk, Andrew Dunsworth, Daniel Eppens, Thomas Eckl, Catherine Erickson, Edward Farhi, Austin Fowler, Brooks Foxen, Craig Gidney, Marissa Giustina, Rob Graff, Jonathan A Gross, Steve Habegger, Matthew P Harrigan, Alan Ho, Sabrina Hong, Trent Huang, William Huggins, Lev B Ioffe, Sergei V Isakov, Evan Jeffrey, Zhang Jiang, Cody Jones, Dvir Kafri, Kostyantyn Kechedzhi, Julian Kelly, Seon Kim, Paul V Klimov, Alexander N Korotkov, Fedor Kostritsa, David Landhuis, Pavel Laptev, Mike Lindmark, Erik Lucero, Michael Marthaler, Orion Martin, John M Martinis, Anika Maruszyk, Sam McArdle, Jarrod R McClean, Trevor McCourt, Matt McEwen, Anthony Megrant, Carlos Mejuto-Zaera, Xiao Mi, Masoud Mohseni, Wojciech Mroczkiewicz, Josh Mutus, Ofer Naaman, Matthew Neeley, Charles Neill, Hartmut Neven, Michael Newman, Murphy Yuezhen Niu, Thomas E O’Brien, Eric Ostby, Bálint Pató, Andre Petukhov, Harald Putterman, Chris Quintana, Jan-Michael Reiner, Pedram Roushan, Nicholas C Rubin, Daniel Sank, Kevin J Satzinger, Vadim Smelyanskiy, Doug Strain, Kevin J Sung, Peter Schmitteckert, Marco Szalay, Norm M Tubman, Amit Vainsencher, Theodore White, Nicolas Vogt,

- Z Jamie Yao, Ping Yeh, Adam Zalcman, and Sebastian Zanker, “Observation of separated dynamics of charge and spin in the Fermi-Hubbard model,” [arXiv:2010.07965 \(2020\)](#).
- [10] Thomas E O’Brien, Stefano Polla, Nicholas C Rubin, William J Huggins, Sam McArdle, Sergio Boixo, Jarrod R McClean, and Ryan Babbush, “Error mitigation via verified phase estimation,” [arXiv:2010.02538 \(2020\)](#).
- [11] Jarrod R McClean, Mollie E Kimchi-Schwartz, Jonathan Carter, and Wibe 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] Jarrod R McClean, Zhang Jiang, Nicholas C Rubin, Ryan Babbush, and Hartmut 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] Sam McArdle, Xiao Yuan, and Simon 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 (2019).
- [17] William J Huggins, Jarrod McClean, Nicholas Rubin, Zhang Jiang, Nathan Wiebe, K Birgitta Whaley, and Ryan Babbush, “Efficient and noise resilient measurements for quantum chemistry on Near-Term quantum computers,” [arXiv:1907.13117 \(2019\)](#).
- [18] Google AI Quantum and Collaborators, “Hartree-Fock on a superconducting qubit quantum computer,” *Science* **369**, 1084–1089 (2020).
- [19] Yanzhu Chen, Maziar Farahzad, Shinjae Yoo, and Tzu-Chieh Wei, “Detector tomography on IBM quantum computers and mitigation of an imperfect measurement,” *Phys. Rev. A* **100**, 052315 (2019).
- [20] Filip B Maciejewski, Zoltán Zimborás, and Michał Oszmaniec, “Mitigation of readout noise in near-term quantum devices by classical post-processing based on detector tomography,” *Quantum* **4**, 257 (2020).
- [21] A Berthiaume, D Deutsch, and R Jozsa, “The stabilization of quantum computations,” in *Proceedings Workshop on Physics and Computation. PhysComp ’94* (ieeexplore.ieee.org, 1994) pp. 60–62.
- [22] Adriano Barenco, André Berthiaume, David Deutsch, Artur Ekert, Richard Jozsa, and Chiara Macchiavello, “Stabilization of quantum computations by symmetrization,” *SIAM J. Comput.* **26**, 1541–1557 (1997).
- [23] Asher Peres, “Error symmetrization in quantum computers,” *Int. J. Theor. Phys.* **38**, 799–805 (1999).
- [24] Paweł Horodecki and Artur Ekert, “Method for direct detection of quantum entanglement,” *Phys. Rev. Lett.* **89**, 127902 (2002).
- [25] Artur K Ekert, Carolina Moura Alves, Daniel K L Oi, Michał Horodecki, Paweł Horodecki, and L C Kwek, “Direct estimations of linear and nonlinear functionals of a quantum state,” *Phys. Rev. Lett.* **88**, 217901 (2002).
- [26] Todd A Brun, “Measuring polynomial functions of states,” [arXiv:quant-ph/0401067 \(2004\)](#).
- [27] Matthew B Hastings, Iván González, Ann B Kallin, and Roger G Melko, “Measuring renyi entanglement entropy in quantum monte carlo simulations,” *Phys. Rev. Lett.* **104**, 157201 (2010).
- [28] Rajibul Islam, Ruichao Ma, Philipp M Preiss, M Eric Tai, Alexander Lukin, Matthew Rispoli, and Markus Greiner, “Measuring entanglement entropy in a quantum many-body system,” *Nature* **528**, 77–83 (2015).
- [29] Juan Carlos Garcia-Escartin and Pedro Chamorro-Posada, “SWAP test and Hong-Ou-Mandel effect are equivalent,” *Phys. Rev. A* **87**, 052330 (2013).
- [30] Sonika Johri, Damian S Steiger, and Matthias Troyer, “Entanglement spectroscopy on a quantum computer,” *Phys. Rev. B Condens. Matter* **96**, 195136 (2017).
- [31] Yiğit Subaşı, Lukasz Cincio, and Patrick J Coles, “Entanglement spectroscopy with a depth-two quantum circuit,” *J. Phys. A: Math. Theor.* **52**, 044001 (2019).
- [32] Leonardo Banchi, Abolfazl Bayat, and Sougato Bose, “Entanglement entropy scaling in solid-state spin arrays via capacitance measurements,” *Phys. Rev. B Condens. Matter* **94**, 241117 (2016).
- [33] Jordan Cotler, Soonwon Choi, Alexander Lukin, Hrant Gharibyan, Tarun Grover, M Eric Tai, Matthew Rispoli, Robert Schittko, Philipp M Preiss, Adam M Kaufman, Markus Greiner, Hannes Pichler, and Patrick Hayden, “Quantum virtual cooling,” *Phys. Rev. X* **9**, 031013 (2019).
- [34] Sergey Bravyi and Alexei Kitaev, “Universal quantum computation with ideal clifford gates and noisy ancillas,” *Phys. Rev. A* **71**, 022316 (2005).
- [35] E Knill, “Quantum computing with realistically noisy devices,” *Nature* **434**, 39–44 (2005).
- [36] Jeongwan Haah and Matthew B Hastings, “Codes and protocols for distilling T , controlled- S , and toffoli gates,” *Quantum* **2**, 71 (2018).
- [37] J I Cirac, A K Ekert, and C Macchiavello, “Optimal purification of single qubits,” *Phys. Rev. Lett.* **82**, 4344–4347 (1999).
- [38] Jacob Biamonte and Ville Bergholm, “Tensor networks in a nutshell,” [arXiv:1708.00006 \(2017\)](#).
- [39] Jacob C Bridgeman and Christopher T Chubb, “Handwaving and interpretive dance: an introductory course on tensor networks,” *J. Phys. A: Math. Theor.* **50**, 223001 (2017).
- [40] Roman Orus, “A practical introduction to tensor networks: Matrix product states and projected entangled pair states,” *Ann. Phys.* **349**, 117–158 (2014).
- [41] Jordan Cotler and Frank Wilczek, “Quantum overlapping tomography,” *Phys. Rev. Lett.* **124**, 100401 (2020).
- [42] Xavier Bonet-Monroig, Ryan Babbush, and Thomas E O’Brien, “Nearly optimal measurement scheduling for partial tomography of quantum states,” *Phys. Rev. X* **10**, 031064 (2020).
- [43] A Yu. Kitaev, “Quantum measurements and the abelian stabilizer problem,” [arXiv:quant-ph/9511026 \(1995\)](#).
- [44] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C Bardin, Rami Barends, Rupak Biswas, Sergio Boixo, Fernando G S L Brandao, David A Buell, Brian Burkett, Yu Chen, Zijun Chen, Ben Chiaro, Roberto Collins, William Courtney, Andrew Dunsworth, Ed-

- ward Farhi, Brooks Foxen, Austin Fowler, Craig Gidney, Marissa Giustina, Rob Graff, Keith Guerin, Steve Habegger, Matthew P Harrigan, Michael J Hartmann, Alan Ho, Markus Hoffmann, Trent Huang, Travis S Humble, Sergei V Isakov, Evan Jeffrey, Zhang Jiang, Dvir Kafri, Kostyantyn Kechedzhi, Julian Kelly, Paul V Klimov, Sergey Knysh, Alexander Korotkov, Fedor Kostritsa, David Landhuis, Mike Lindmark, Erik Lucero, Dmitry Lyakh, Salvatore Mandrà, Jarrod R McClean, Matthew McEwen, Anthony Megrant, Xiao Mi, Kristel Michielsen, Masoud Mohseni, Josh Mutus, Ofer Naaman, Matthew Neeley, Charles Neill, Murphy Yuezhen Niu, Eric Ostby, Andre Petukhov, John C Platt, Chris Quintana, Eleanor G Rieffel, Pedram Roushan, Nicholas C Rubin, Daniel Sank, Kevin J Satzinger, Vadim Smelyanskiy, Kevin J Sung, Matthew D Trevithick, Amit Vainsencher, Benjamin Villalonga, Theodore White, Z Jamie Yao, Ping Yeh, Adam Zalcman, Hartmut Neven, and John M Martinis, “Quantum supremacy using a programmable superconducting processor,” *Nature* **574**, 505–510 (2019).
- [45] D V Dmitriev and V Ya Krivnov, “Quasi-one-dimensional anisotropic heisenberg model in a transverse magnetic field,” *Journal of Experimental and Theoretical Physics Letters* **80**, 303–307 (2004).
- [46] Suguru Endo, Qi Zhao, Ying Li, Simon Benjamin, and Xiao Yuan, “Mitigating algorithmic errors in a hamiltonian simulation,” *Phys. Rev. A* **99**, 012334 (2019).
- [47] Minh C Tran, Yuan Su, Daniel Carney, and Jacob M Taylor, “Faster digital quantum simulation by symmetry protection,” [arXiv:2006.16248](https://arxiv.org/abs/2006.16248) (2020).
- [48] Earl Campbell, “Random compiler for fast hamiltonian simulation,” *Phys. Rev. Lett.* **123**, 070503 (2019).
- [49] Andrew M Childs, Aaron Ostrander, and Yuan Su, “Faster quantum simulation by randomization,” *Quantum* **3**, 182 (2019).
- [50] Yingkai Ouyang, David R White, and Earl T Campbell, “Compilation by stochastic Hamiltonian sparsification,” *Quantum* **4**, 235 (2020).
- [51] Chi-Fang Chen, Hsin-Yuan Huang, Richard Kueng, and Joel A Tropp, “Quantum simulation via randomized product formulas: Low gate complexity with accuracy guarantees,” [arXiv:2008.11751](https://arxiv.org/abs/2008.11751) (2020).
- [52] Bálint Koczor, “Exponential error suppression for Near-Term quantum devices,” [arXiv:2011.05942](https://arxiv.org/abs/2011.05942) (2020).
- [53] Maurice George Kendall and Others, *The Advanced Theory of Statistics* (Charles Griffin and Co., Ltd., London, 1946).
- [54] 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**, 456–469 (2010).
- [55] Clare Horsman, Austin G Fowler, Simon Devitt, and Rodney Van Meter, “Surface code quantum computing by lattice surgery,” *New J. Phys.* **14**, 123011 (2012).
- [56] Daniel Litinski, “A Game of Surface Codes: Large-Scale Quantum Computing with Lattice Surgery,” *Quantum* **3**, 128 (2019).

## Appendix A: Measurement by Diagonalization with Three or More Copies

In [Section II A](#) and [Section II B](#) we described protocols for measuring the expectation value of  $O$  with respect to the state  $\frac{\rho^2}{\text{Tr}(\rho^2)}$  by diagonalizing  $S^{(2)}$  and either  $O^1 S^{(2)}$  or  $O^{(2)} S^{(2)}$ . These approaches can be generalized to higher powers of  $\rho$  in a natural way. The cyclic shift operator  $S^{(M)}$ , like the swap operator  $S^{(2)}$ , factorizes into a tensor product over  $N$   $M$ -tuples, one for each of the  $N$  qubits of the system of interest. The symmetrized operator  $Z_k^{(M)} = \frac{1}{M} \sum_{i=1}^M Z_k^i$  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, 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 multi-qubit observables that we discussed above for  $M = 2$  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 multi-qubit Pauli operator  $P$ , we can do so using a generalization of [Eq. 18](#), 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 have numerically optimized the quantum circuit of [Figure 11](#) to simultaneously diagonalize  $Z_k^{(3)} S_k^{(3)}$  and  $S_k^{(3)}$ . We obtained 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

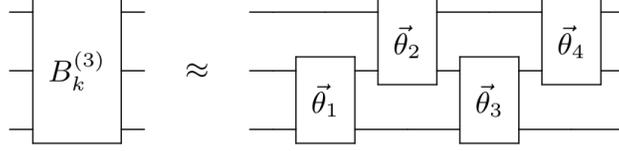


Figure 11. The ansatz that we numerically optimize to approximately diagonalize  $S_k^{(3)}$  and  $Z_k^{(3)}$ . The two-qubit gates parameterized by the  $\vec{\theta}_i$ s are arbitrary two-qubit gates. We performed the numerical optimization using the Julia language.

the following equations are used,

$$B_k^{(3)} S_k^{(3)} B_k^{(3)\dagger} \rightarrow \frac{1}{8} \left( 2 + \right. \quad (A1)$$

$$\begin{aligned} & (-3 - \sqrt{3}i) Z_k^1 + \\ & (3 - \sqrt{3}i) Z_k^1 Z_k^2 + \\ & (3 + \sqrt{3}i) Z_k^3 + \\ & 2i\sqrt{3} Z_k^1 Z_k^3 + \\ & \left. (3 - \sqrt{3}i) Z_k^2 Z_k^3 \right), \end{aligned}$$

$$B_k^{(3)} Z_k^{(3)} B_k^{(3)\dagger} \rightarrow \frac{1}{3} (Z_k^1 + Z_k^2 + Z_k^3). \quad (A2)$$

## Appendix B: Ancilla-Assisted Measurement with Three or More Copies

In [Section II C](#), we described a protocol for measuring the expectation value of an arbitrary observable  $O$  with respect to the state  $\frac{\rho^2}{\text{Tr}(\rho^2)}$  using the controlled-swap operator between two copies of  $\rho$ .

Now let us consider the case with three or more copies of  $\rho$ .  $S^{(N)}$  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  $\text{Re}(\text{Tr}(S^{(N)} \rho^{\otimes N}))$  [[25](#), [26](#), [33](#)].

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 C: Variance of the Corrected Expectation Value Estimator

In this section, we shall calculate the variance of the estimator for the corrected expectation value obtained by applying [Eq. 11](#) with  $M = 2$ . Specifically, we shall consider the estimation of the expectation value of an observable  $O$  with respect to the state  $\frac{\rho^2}{\text{Tr}(\rho^2)}$  constructed by repeatedly measuring the operators  $S^{(2)} O^{(2)}$  and  $S^{(2)}$ ,

$$\frac{\text{Tr}(O \rho^2)}{\text{Tr}(\rho^2)} = \frac{\text{Tr}(O^{(2)} S^{(2)} \rho^{\otimes 2})}{\text{Tr}(S^{(2)} \rho^{\otimes 2})}. \quad (C1)$$

We shall assume that both operators are simultaneously measured by averaging over  $R$  repetitions of state-preparation and measurement. This assumption applies to the methods of [Section II A](#) when the expectation value of interest is a single-qubit Pauli operator and it applies generally to the methods of [Section II C](#).

The measurements of these operators are classical random variables, and we can proceed by determining the variance

of these two random variables and their covariance. Let's start with the numerator.

$$\text{Var}(S^{(2)}O^{(2)}) = \langle (S^{(2)}O^{(2)})^2 \rangle - \langle S^{(2)}O^{(2)} \rangle^2 \quad (\text{C2})$$

$$= \langle (O^{(2)})^2 \rangle - \text{Tr}(\rho^2 O)^2 \quad (\text{C3})$$

$$= \frac{1}{4} \text{Tr}((\rho \otimes \rho)(O^2 \otimes \mathbb{I} + 2O \otimes O + \mathbb{I} \otimes O^2)) - \text{Tr}(\rho^2 O)^2 \quad (\text{C4})$$

$$= \frac{1}{2} \text{Tr}(\rho O^2) + \frac{1}{2} \text{Tr}(\rho O)^2 - \text{Tr}(\rho^2 O)^2 \quad (\text{C5})$$

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

$$\text{Var}(S^{(2)}) = \langle (S^{(2)})^2 \rangle - \langle S^{(2)} \rangle^2 \quad (\text{C6})$$

$$= 1 - \text{Tr}(\rho^2)^2. \quad (\text{C7})$$

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

$$\text{Cov}(S^{(2)}O^{(2)}, S^{(2)}) = \langle S^{(2)}O^{(2)}S^{(2)} \rangle - \langle S^{(2)}O^{(2)} \rangle \langle S^{(2)} \rangle \quad (\text{C8})$$

$$= \langle O^{(2)} \rangle - \text{Tr}(\rho^2 O) \text{Tr}(\rho^2) \quad (\text{C9})$$

$$= \text{Tr}(\rho O) - \text{Tr}(\rho^2 O) \text{Tr}(\rho^2). \quad (\text{C10})$$

There isn't a closed-form expression for the variance of the ratio of two random variables [53], but we can take the standard approximation based on a Taylor series expansion,

$$\text{Var}\left(\frac{A}{B}\right) \approx \frac{1}{\langle B \rangle^2} \text{Var}(A) - 2 \frac{\langle A \rangle}{\langle B \rangle^3} \text{Cov}(A, B) + \frac{\langle A \rangle^2}{\langle B \rangle^4} \text{Var}(B). \quad (\text{C11})$$

We estimate the expectation values for the numerator and denominator of Eq. C1 by averaging over a series of  $R$  experiments. This scales the variances calculated above by a factor of  $\frac{1}{R}$ . If  $R$  is sufficiently large, then the approximation presented in Eq. C11 will be a good one. Applying this expression to determine the variance of the estimator from Eq. C1 yields

$$\begin{aligned} \text{Var}(\text{Estimator}) &\approx \frac{1}{R} \left( \frac{1}{\text{Tr}(\rho^2)^2} \left( \frac{1}{2} \text{Tr}(\rho O^2) + \frac{1}{2} \text{Tr}(\rho O)^2 - \text{Tr}(\rho^2 O)^2 \right) \right. \\ &\quad \left. - 2 \frac{\text{Tr}(\rho^2 O)}{\text{Tr}(\rho^2)^3} (\text{Tr}(\rho O) - \text{Tr}(\rho^2 O) \text{Tr}(\rho^2)) + \frac{\text{Tr}(\rho^2 O)^2}{\text{Tr}(\rho^2)^4} (1 - \text{Tr}(\rho^2)^2) \right). \end{aligned} \quad (\text{C12})$$

#### Appendix D: Variance of the Proposed Collective Measurement

In Section IID, we defined the operator

$$\tilde{O} = \frac{1}{\binom{K}{2}} \sum_{i=1}^K \sum_{j>i} \frac{1}{2} (O^i + O^j) S^{(i,j)}, \quad (\text{D1})$$

where we use  $S^{(i,j)}$  to denote the swap operator between subsystems  $i$  and  $j$  and  $O$  is an arbitrary Pauli operator. We shall now calculate the variance of measurements of this operator with respect to the state  $\rho^{\otimes K}$ .

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

$$= \frac{1}{4 \binom{K}{2}^2} \sum_{i=1}^K \sum_{j>i} \sum_{a=1}^K \sum_{b>a} \langle (O^i + O^j) S^{(i,j)} (O^a + O^b) S^{(a,b)} \rangle - \text{Tr}(O \rho^2)^2. \quad (\text{D3})$$

We shall bound the variance by breaking the sum up into three cases. In the first case,  $i = a$  and  $j = b$ . In the second case, there are only three distinct values amongst the indices  $i, j, a, b$ . In the fourth case, all four of the indices take distinct values.

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

$$\frac{1}{4\binom{K}{2}^2} \sum_{i=1}^K \sum_{j>i} \langle (O^i + O^j)S^{(i,j)}(O^i + O^j)S^{(i,j)} \rangle = \frac{1}{4\binom{K}{2}^2} \sum_{i=1}^K \sum_{j>i} \langle 2 + O^i O^j + O^j O^i \rangle \quad (\text{D4})$$

$$\leq \frac{1}{\binom{K}{2}}. \quad (\text{D5})$$

Here we have used the properties that  $S^{(i,j)}$  is self-inverse that that it commutes with  $(O^i + O^j)$ .

Next, let's 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  $\rho^{\otimes K}$  can be evaluated separately and multiplied together. We can use this fact to simplify and bound this component of the sum,

$$\frac{1}{4\binom{K}{2}^2} \sum_{i=1}^K \sum_{j>i} \sum_{a=1, a \neq i, a \neq j}^K \sum_{b>a, b \neq i, b \neq j} \langle (O^i + O^j)S^{(i,j)}(O^a + O^b)S^{(a,b)} \rangle \quad (\text{D6})$$

$$= \frac{1}{4\binom{K}{2}^2} \sum_{i=1}^K \sum_{j>i} \sum_{a=1, a \neq i, a \neq j}^K \sum_{b>a, b \neq i, b \neq j} \langle (O^i + O^j)S^{(i,j)} \rangle \langle (O^a + O^b)S^{(a,b)} \rangle \quad (\text{D7})$$

$$= \frac{1}{\binom{K}{2}^2} \sum_{i=1}^K \sum_{j>i} \sum_{a=1, a \neq i, a \neq j}^K \sum_{b>a, b \neq i, b \neq j} \text{Tr}(O\rho^2)^2 \quad (\text{D8})$$

$$= \frac{1}{\binom{K}{2}} \binom{K-2}{2} \text{Tr}(O\rho^2)^2 < \text{Tr}(O\rho^2)^2. \quad (\text{D9})$$

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

$$\langle (O^i + O^j)S^{(i,j)}(O^i + O^b)S^{(i,b)} \rangle \quad (\text{D10})$$

$$= \langle (O^i + O^j)(O^j + O^b)S^{(i,j)}S^{(i,b)} \rangle \quad (\text{D11})$$

$$= \langle (O^i O^j + O^i O^b + O^j O^b + 1)S^{(i,j)}S^{(i,b)} \rangle. \quad (\text{D12})$$

Here we have used the property that each  $O^k$  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 [Figure 12](#)) establishes that

$$\langle (O^i O^j + O^i O^b + O^j O^b + 1)S^{(i,j)}S^{(i,b)} \rangle \quad (\text{D13})$$

$$= 3\text{Tr}(O\rho O\rho^2) + \text{Tr}(\rho^3). \quad (\text{D14})$$

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. [D14](#) in terms of these projectors, yielding

$$3\text{Tr}(O\rho O\rho^2) + \text{Tr}(\rho^3) \quad (\text{D15})$$

$$= 3\text{Tr}(P_+\rho P_+\rho^2) - 3\text{Tr}(P_-\rho P_+\rho^2) - 3\text{Tr}(P_+\rho P_-\rho^2) + 3\text{Tr}(P_-\rho P_-\rho^2) + \text{Tr}(\rho^3) \quad (\text{D16})$$

$$\leq 7\text{Tr}(\rho^3). \quad (\text{D17})$$

For simplicity, let us define the indicator function

$$W(i, j, a, b) = \begin{cases} 1 & \text{if } i, j, a, b \text{ take exactly three distinct values} \\ 0 & \text{otherwise} \end{cases} \quad (\text{D18})$$

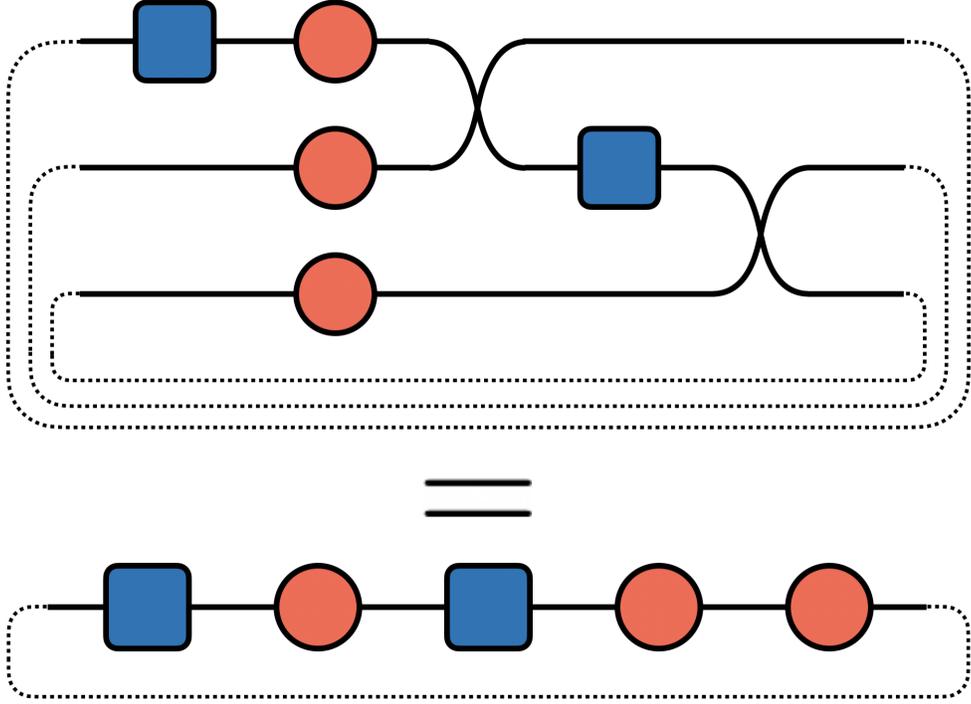


Figure 12. A diagrammatic proof of Eq. D14.

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

$$\frac{1}{4\binom{K}{2}^2} \sum_{i=1}^K \sum_{j>i}^K \sum_{a=1}^K \sum_{b>a}^K \langle (O^i + O^j) S^{(i,j)} (O^a + O^b) S^{(a,b)} \rangle I(i, j, a, b) \quad (\text{D19})$$

$$\leq \frac{7}{4\binom{K}{2}^2} \sum_{i=1}^K \sum_{j>i}^K \sum_{a=1}^K \sum_{b>a}^K I(i, j, a, b) \text{Tr}(\rho^3). \quad (\text{D20})$$

$$\leq \frac{7(2K - 4)}{4\binom{K}{2}} \text{Tr}(\rho^3) \quad (\text{D21})$$

$$\leq \frac{7(K - 2)}{K(K - 1)} \text{Tr}(\rho^3). \quad (\text{D22})$$

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

$$\text{Var}(\tilde{O}) \leq \frac{2 + 7(K - 2)\text{Tr}(\rho^3)}{K(K - 1)}. \quad (\text{D23})$$

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

## Appendix E: Details of Numerical Simulations

### 1. Scrambling Circuits

In Section IV A we briefly described the random circuits that we simulated to produce Figure 4, Figure 5, and Figure 6. Here we expand upon that description.

The first class of circuits are essentially the one-dimensional analogues 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,

$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & -i & 0 \\ 0 & -i & 0 & 0 \\ 0 & 0 & 0 & e^{-\frac{i\pi}{6}} \end{bmatrix}. \quad (\text{E1})$$

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

$$\{X, Y, Z, \sqrt{X}, \sqrt{Y}, \sqrt{Z}\}, \quad (\text{E2})$$

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 exactly 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 together.

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,

$$\Delta(\rho) = (1-p)\rho + \frac{p}{3} + \frac{p}{3}(X\rho X + Y\rho Y + Z\rho Z). \quad (\text{E3})$$

An equivalent formulation which will be useful for our purposes is

$$\Delta(\rho) = (1 - \frac{4}{3}p)\rho + \frac{4}{3}p\frac{\mathbb{I}}{2}. \quad (\text{E4})$$

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

$$\Delta(|\phi\rangle\langle\phi|) = (1 - \frac{2}{3}p)|\phi\rangle\langle\phi| + \frac{2}{3}p|\phi^\perp\rangle\langle\phi^\perp|, \quad (\text{E5})$$

where  $|\phi^\perp\rangle$  is orthogonal to  $|\phi\rangle$ . Eq. E4 allows us to easily analyze  $D$  repeated applications of the channel,

$$\Delta(\rho)^D = (1 - \frac{4}{3}p)^D\rho + (1 - (1 - \frac{4}{3}p)^D)\frac{\mathbb{I}}{2}, \quad (\text{E6})$$

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

$$\tilde{p} = \frac{3}{4} - \frac{3}{4}(1 - \frac{4}{3}p)^D. \quad (\text{E7})$$

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

$$\bigotimes_{i=1}^N (1 - \frac{2}{3}\tilde{p}_{(i)})|\phi_i\rangle\langle\phi_i| + \frac{2}{3}\tilde{p}_{(i)}|\phi_i^\perp\rangle\langle\phi_i^\perp|. \quad (\text{E8})$$

Here the effective single-qubit depolarizing probability  $\tilde{p}_{(i)}$  depends on  $i$  because the qubits at the end of the circuit are only subject to  $\frac{D}{2}$  applications of the single-qubit depolarizing channel instead of the  $D$  that are applied to qubits in the bulk. Therefore, employing Eq. E7, we have

$$\tilde{p}_{(i)} = \begin{cases} \frac{3}{4} - \frac{3}{4}(1 - \frac{4}{3}p)^{\frac{D}{2}} & \text{for } i = 1 \text{ or } i = N, \\ \frac{3}{4} - \frac{3}{4}(1 - \frac{4}{3}p)^D & \text{for } 2 \leq i \leq N - 1. \end{cases} \quad (\text{E9})$$

We can observe a few things from this equation. 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  $(|\phi_1\rangle \otimes |\phi_2\rangle \otimes \dots \otimes |\phi_N\rangle)$ . Secondly, 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  $\approx \frac{2}{3}P$ . After these states, there will be  $\binom{N}{2}$  eigenvectors corresponding two states with two-errors. This distribution doesn't exactly match the phenomenological noise model we assumed in [Section III A](#), but it is qualitatively similar.

### Appendix F: Interplay with the surface code

While the majority of this work has focused on the NISQ regime, one interesting question to ask is what role this approach can 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 [\[54\]](#) that in conjunction with a minimum-weight perfect matching decoder, the error rate of a surface code cycle is roughly

$$\epsilon_c = 10^{-(d+3)/2} \quad (\text{F1})$$

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 time-like) 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 [\[55, 56\]](#). 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 10 T gates and 10 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

$$f_1 = \left(1 - 10^{-(\sqrt{n/2}+3)/2}\right)^{100\sqrt{n/2}G} \quad (\text{F2})$$

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, consider the error rate of  $G$  gates in the surface code using twice the qubits

$$f_2 = \left(1 - 10^{-(\sqrt{2n}+3)/2}\right)^{100\sqrt{2n}G}. \quad (\text{F3})$$

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 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 is about 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

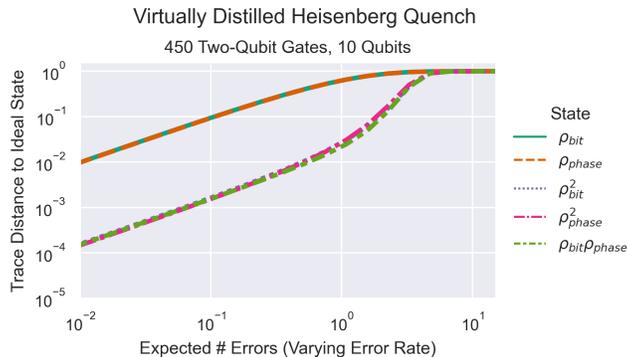


Figure 13. The error in the unmitigated noisy states, and the states accessed by virtual distillation, for the same 10 qubit Heisenberg evolution under two different noise models. We consider a bit-flip error model ( $\rho_{bit}$ , teal curve) and phase-flip error model ( $\rho_{phase}$ , orange dashed curve). We show the error for virtual distillation applied in the usual way to two identical copies of each noisy state ( $\rho_{bit}^2$  and  $\rho_{phase}^2$ , purple dotted curve and pink dashed curve). We also consider the error when virtual distillation is applied to the two different states ( $\rho_{bit}\rho_{phase}$ , green dotted curve). We quantify the error using the trace distance to the state obtained from noiseless evolution as a function of the expected number of single-qubit 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.

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 will 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 G: Virtual Distillation Applied to Distinct States

In the main body of this paper, we applied virtual distillation to a variety of systems under the assumption that we had 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 experience by each copy will not be identical. Consider the case where we apply virtual distillation to two distinct state,  $\rho_A$  and  $\rho_B$ . It's straightforward to show that we then effectively measure expectation values with respect to the state

$$\frac{\rho_A \rho_B}{\text{Tr}(\rho_A \rho_B)}. \quad (\text{G1})$$

Note that we still rely on the important assumption that the two copies are unentangled prior to virtual distillation.

In order to explore the impact of virtually distilling two different states together, we present an additional simulation of the Heisenberg evolution that we considered in [Section IV B](#). In [Figure 13](#), rather than employing the single-qubit depolarizing noise model we used previously, we calculate  $\rho_A = \rho_{bit}$  using an analogous application of a bit-flip error channel, while using a phase-flip channel for  $\rho_B = \rho_{phase}$ . 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.