

Theory of quantum error mitigation for non-Clifford gates

David Layden, Bradley Mitchell, and Karthik Siva

IBM Research

Quantum error mitigation techniques mimic noiseless quantum circuits by running several related noisy circuits and combining their outputs in particular ways. How well such techniques work is thought to depend strongly on how noisy the underlying gates are. Weakly-entangling gates, like $R_{ZZ}(\theta)$ for small angles θ , can be much less noisy than entangling Clifford gates, like CNOT and CZ, and they arise naturally in circuits used to simulate quantum dynamics. However, such weakly-entangling gates are non-Clifford, and are therefore incompatible with two of the most prominent error mitigation techniques to date: probabilistic error cancellation (PEC) and the related form of zero-noise extrapolation (ZNE). This paper generalizes these techniques to non-Clifford gates, and comprises two complementary parts. The first part shows how to effectively transform any given quantum channel into (almost) any desired channel, at the cost of a sampling overhead, by adding random Pauli gates and processing the measurement outcomes. This enables us to cancel or properly amplify noise in non-Clifford gates, provided we can first characterize such gates in detail. The second part therefore introduces techniques to do so for noisy $R_{ZZ}(\theta)$ gates. These techniques are robust to state preparation and measurement (SPAM) errors, and exhibit concentration and sensitivity—crucial statistical properties for many experiments. They are related to randomized benchmarking, and may also be of interest beyond the context of error mitigation. We find that while non-Clifford gates can be less noisy than related Cliffords, their noise is fundamentally more complex, which can lead to surprising and sometimes unwanted effects in error mitigation. Whether this trade-off can be broadly advantageous remains to be seen.

Contents

1	Introduction	2
1.1	Contributions	2
2	Error Mitigation	4
2.1	Mathematical Background	4
2.2	Clifford Gates	5
2.3	Non-Clifford Gates	6
2.3.1	Formalism of Pauli Shaping	7
2.3.2	Examples of Pauli Shaping	10
3	Characterizing Noisy Gates	14
3.1	Clifford Gates	14
3.2	Non-Clifford Gates	17
3.2.1	Type 1 Elements	18
3.2.2	Type 2 and 3 Elements	19

1 Introduction

While fault tolerance is essential for realizing the full potential of quantum computing, error mitigation may unlock some of this potential before the advent of large-scale fault tolerance [1]. The most prominent error mitigation techniques seek to compute noiseless expectation values for a given quantum circuit by running several related circuits on a noisy quantum computer, then combining their measurement outcomes in nontrivial ways [2, 3]. The resulting precision and/or accuracy (within a fixed running time) typically improves with increasing gate quality, but declines—often exponentially—with the size of the target circuit. In other words, better gates can enable these techniques on bigger circuits—a dynamic that may provide a continuous path towards fault tolerance [4].

Due to its strong dependence on circuit size, error mitigation is most promising for problems which admit an exponential quantum speedup in computing expectation values, and where the required gates closely match the connectivity between qubits in hardware. The most evident example is quantum simulation using Trotter/Floquet-type circuits [5, 6]. Such circuits implement repeated unitaries of the form $U_j = \exp(-iH_j \delta t)$ using one- and two-qubit gates, where $\{H_j\}$ are components of the Hamiltonian being simulated and δt represents a timestep. One typically wants a small δt to reduce Trotter error, which makes each U_j (at most) weakly entangling [7]. For example, alternating between U_1 and U_2 layers, generated by $H_1 = g \sum_i X_i$ and $H_2 = -J \sum_{\langle i,j \rangle} Z_i Z_j$ respectively, approximates evolution by the transverse-field Ising model $H_1 + H_2$. Using the notation

$$R_P(\theta) = e^{-iP\theta/2} \quad (1)$$

where P denotes a Pauli operator, U_1 comprises only single-qubit unitaries $R_X(\phi)$ with $\phi = 2g\delta t$, while U_2 comprises two-qubit unitaries $R_{ZZ}(\theta)$ with $\theta = -2J\delta t$, which generate little entanglement per layer when δt is small enough to give a reasonable Trotter error.

Such weakly-entangling unitaries can be realized through two main strategies, which use qualitatively different two-qubit gates [8]. The first uses fixed two-qubit Clifford gates, like CNOTs, regardless of δt , while the second uses weakly-entangling, non-Clifford two-qubit gates that approach I as $\delta t \rightarrow 0$. (Both strategies can also use arbitrary single-qubit gates as needed.) Following the first strategy, for example, one might compile $R_{ZZ}(\theta)$ into two CNOTs and a single-qubit $R_Z(\theta)$ gate, where each CNOT is locally equivalent to $R_{ZZ}(\pi/2)$. Following the second strategy, one would instead implement $R_{ZZ}(\theta)$ directly, up to single-qubit gates, by shortening the control sequence used to perform CNOTs—effectively doing a fraction of a CNOT rather than two. We will call these two strategies *digital* and *semi-analog*, respectively. (The latter is not fully analog as it still discretizes the quantum dynamics of interest into circuits.) Since two-qubit gates are the dominant source of errors in most pre-fault-tolerant devices, these two strategies offer very different advantages. The digital strategy is manifestly compatible with the most prominent error mitigation techniques, which handle errors on two-qubit Clifford gates [9, 10]. On the other hand, the semi-analog strategy can incur substantially fewer errors to begin with, by virtue of having faster, and sometimes also fewer, two-qubit gates [11, 12]. However, because these latter gates are non-Clifford, they have been largely incompatible with said error mitigation techniques to date.

1.1 Contributions

Motivated by this semi-analog strategy for quantum simulation, we introduce a broad error mitigation technique for non-Clifford gates. It comprises two interlocking, yet distinct parts: in Section 2 we develop a method for mitigating noise on arbitrary gates, given a detailed description of the noise. Then in Section 3, we introduce a robust method to characterize said noise on common non-Clifford

gates. Together, these two parts greatly generalize the most prominent error mitigation techniques, thereby making them compatible with the relatively low-noise, non-Clifford gates that arise naturally in many experiments.

More precisely, Section 2 extends probabilistic error cancellation (PEC) and zero-noise extrapolation (ZNE), two of the most successful error mitigation techniques to date, to generic noise on arbitrary gates. In their current form, these techniques can effectively transform one Pauli channel into another by inserting single-qubit gates at random from some appropriate distribution. This suffices to mitigate multi-qubit Clifford gates because generic noise on such gates can always be transformed into a Pauli channel through Pauli-twirling, using single-qubit gates. The main obstacle to generalizing PEC and ZNE is that noise on non-Clifford gates cannot be twirled in this way, so it cannot generally be transformed into a Pauli channel. Our key insight here is to reframe the task: we effectively construct a “super-channel” (a linear map acting on quantum channels) that transforms a noisy non-Clifford gate into a desired channel—usually the ideal gate. We then show that this super-channel can be effectively realized by inserting random single-qubit gates into the circuit, as in Clifford PEC/ZNE, just with a different distribution. As usual, this distribution depends on the precise gate noise, which we assume is already known. Our technique applies to arbitrary gates, with only mild regularity requirements on their noise, and reduces to existing methods when applied to Cliffords. However, it behaves unexpectedly on some non-Clifford gates, in ways that have no clear analog in the Clifford case. Most notably, the sampling overhead required to effectively cancel noise on a non-Clifford gate can be discontinuous in the noise strength, and can remain large for arbitrarily weak noise, depending on subtle properties of the noise. We do not attempt to fully characterize the gates and noise structures where this effect occurs, but instead give three illustrative examples (two analytical and one numerical) with $R_{ZZ}(\theta)$ gates to inspire further study of this new effect.

In Section 3 we turn to the complementary problem: how to thoroughly characterize noise on non-Clifford gates to begin with. The first major hurdle is that the measurement outcomes in any experiment depend not only on gate noise, but also on state preparation and measurement (SPAM) errors, which can be significantly stronger and therefore make gate noise hard to isolate. This rules out quantum process tomography, which is not generally robust to SPAM errors. A potential alternative is gate set tomography, which learns gate and SPAM errors simultaneously and self-consistently [13, 14, 15, 16]. While very general, it can be correspondingly expensive in terms of the experiments and computation required. Cycle benchmarking (CB) is therefore a more widely adopted technique, which avoids learning SPAM errors by channeling their effects into a single parameter that can ultimately be discarded. Our main contribution here is generalizing CB to non-Clifford gates. In doing so, one encounters the second major hurdle: standard CB is very robust to statistical fluctuations from running a finite number of random circuits finitely many times—a property that does not automatically carry over to the non-Clifford case. We show that our generalization is similarly robust by construction, while illustrating the pitfalls of more obvious approaches. How to ensure statistical robustness in the non-Clifford case depends significantly on the gate in question, since there is a wider variety of non-Clifford gates than of Cliffords. So for the sake of concreteness, we focus in this section on characterizing noisy $R_{ZZ}(\theta)$ gates, which are common across many experiments (unlike in Section 2, which is completely general). Our technique learns all the noise parameters needed to effectively cancel generic noise using our method from Section 2. We illustrate this numerically under strong SPAM errors and large statistical fluctuations. Finally, we find that, while not strictly necessary, learning certain noise parameters through CB-like methods seems fundamentally hard. Curiously, these are exactly the parameters that can cause unusual behavior in the sampling overhead in Section 2, which suggests a novel connection between learnability and mitigation.

We make several common mathematical assumptions about noise in this work. Specifically, we assume throughout that all noise is constant in time, or rather, that it drifts slowly enough to be treated as constant. We also assume that the noise on multi-qubit gates is Markovian, meaning that each noisy gate can be described by a fixed quantum channel regardless of when it occurs in a circuit. Moreover, we approximate single-qubit gates as being noiseless compared to multi-qubit gates, to reflect the fact that the latter are typically much noisier on pre-fault-tolerant devices. This is not

strictly necessary, but it simplifies the presentation and keeps the focus on the main topics at hand. Finally, we assume in Section 3 that the gate noise is not so strong as to totally overwhelm the $R_{ZZ}(\theta)$ gate, as described by Eq. (51). From a theoretical perspective, we refer to this condition as the “weak noise regime,” although it encompasses noise that is very strong by experimental standards (e.g., 95% average fidelity, as in our numerics), and is therefore a mild assumption in practice.

2 Error Mitigation

We begin this section by reviewing the essential mathematical background in Sec. 2.1, and the existing formalism of PEC and ZNE for Clifford gates in Sec. 2.2, before finally introducing our more general mitigation scheme in Sec. 2.3.

2.1 Mathematical Background

A generic operation on n qubits, unitary or not, can be described by a completely positive trace-preserving (CPTP¹) map $\mathcal{G}$, also called a quantum channel, which maps an input state ρ to an output state $\rho' = \mathcal{G}(\rho)$. There are several distinct matrix representations for a given quantum channel, but it will be convenient here to describe $\mathcal{G}$ by its Pauli transfer matrix (PTM) $\mathbf{G}$, a $4^n \times 4^n$ real matrix with elements

$$\mathbf{G}_{ij} = \text{tr}[P_i \mathcal{G}(P_j)]/2^n \quad (2)$$

between -1 and 1 , where P_i and P_j are n -qubit Paulis [17]. ($\mathbf{G}$ is also known as a Liouville representation of $\mathcal{G}$.) We will denote PTMs and other matrices of the same size in bold to distinguish them from $2^n \times 2^n$ -dimensional unitaries like P_i . Writing the input/output states above in the Pauli basis, $\rho = 2^{-n} \sum_i s_i P_i$ and $\rho' = 2^{-n} \sum_i s'_i P_i$, the PTM of $\mathcal{G}$ is the matrix relating their generalized Bloch vectors as $\vec{s}' = \mathbf{G} \vec{s}$. For any two channels $\mathcal{G}^{(1)}$ and $\mathcal{G}^{(2)}$ with PTMs $\mathbf{G}^{(1)}$ and $\mathbf{G}^{(2)}$, the PTM of $\mathcal{G}^{(1)}$ followed by $\mathcal{G}^{(2)}$ is simply $\mathbf{G}^{(2)} \mathbf{G}^{(1)}$, and performing $\mathcal{G}^{(1)}$ or $\mathcal{G}^{(2)}$ with respective probabilities p and $1-p$ gives a PTM of $p \mathbf{G}^{(1)} + (1-p) \mathbf{G}^{(2)}$.

A channel $\mathcal{P}$ is said to be a Pauli channel if it acts as $\mathcal{P}(\rho) = \sum_i p_i P_i \rho P_i$ for some probability distribution $\vec{p} = (p_i)_{i=0}^{4^n-1}$, meaning that it can be understood as a process where Pauli errors P_i occur with probabilities p_i . Equivalently, a Pauli channel is one whose PTM is diagonal, $\mathbf{P} = \text{diag}(\vec{f})$, with elements f_i that are sometimes called *Pauli eigenvalues* [18] or *Pauli fidelities* [9]. These are related to the error probabilities p_i by $\vec{f} = \mathbf{W} \vec{p}$, where $\mathbf{W}$ is a $4^n \times 4^n$ Walsh-Hadamard matrix with elements

$$\mathbf{W}_{ij} = \begin{cases} +1, & [P_i, P_j] = 0 \\ -1, & \{P_i, P_j\} = 0, \end{cases} \quad (3)$$

that describes a type of discrete Fourier transform and obeys $\mathbf{W} = \mathbf{W}^\top = 4^n \mathbf{W}^{-1}$.

Finally, a generic channel $\mathcal{G}$ can be transformed into a Pauli channel through Pauli-twirling [19, 20]. Twirling means sampling a random unitary V uniformly from some given set each time $\mathcal{G}$ is implemented, and applying V and $V^\dagger$ before and after $\mathcal{G}$, respectively. When V is sampled from

$$\mathbb{P} = \left\{ P^{(1)} \otimes \dots \otimes P^{(n)} \mid P^{(i)} \in \{I, X, Y, Z\} \right\}, \quad (4)$$

the set of all n -qubit Paulis, i.e., when $V \sim \text{unif}(\mathbb{P})$, the process is called Pauli-twirling. For any $\mathcal{G}$, the resulting, overall channel

$$\bar{\mathcal{G}}(\rho) = \frac{1}{4^n} \sum_{P_i \in \mathbb{P}} P_i \mathcal{G}(P_i \rho P_i) P_i \quad (5)$$

is a Pauli channel, with Pauli fidelities $f_i = \mathbf{G}_{ii}$. All off-diagonal elements of $\mathbf{G}$ are averaged away by the twirling.

¹Many of our results apply more broadly to completely positive trace-non-increasing maps (which can describe leakage), but we will focus on CPTP maps here to simplify the presentation.

2.2 Clifford Gates

Suppose we want to implement a Clifford gate, or a layer of simultaneous Clifford gates, described by an n -qubit unitary U and a corresponding quantum channel $\mathcal{U}(\rho) = U\rho U^\dagger$, but we instead implement a slightly different channel $\mathcal{G}$ due to experimental imperfections. It is customary to factor this noisy gate into $\mathcal{G} = \mathcal{U}\mathcal{N}$, where $\mathcal{N} = \mathcal{U}^{-1}\mathcal{G}$ describes the noise and $\mathcal{U}^{-1}(\rho) = U^\dagger\rho U$. (We could equivalently factor it into $\mathcal{G} = \mathcal{N}'\mathcal{U}$ for noise $\mathcal{N}' = \mathcal{G}\mathcal{U}^{-1}$. Generically $\mathcal{N} \neq \mathcal{N}'$, although the choice of order is inconsequential, at least for PEC, as long as we are consistent.) PEC, and the related version of ZNE, both have two conceptual steps in terms of this factorization, as depicted in Fig. 1:

Step 1 Pauli-twirl the noise channel $\mathcal{N}$ to simplify it, using only single-qubit gates.

Step 2 Transform (i.e., cancel for PEC, or amplify for ZNE) the resulting Pauli noise channel $\tilde{\mathcal{N}}$.

The noise $\mathcal{N}$ will not typically be a Pauli channel, but we can transform it into one through Pauli-twirling in Step 1. To do so while leaving $\mathcal{U}$ intact, we must apply a random Pauli $P_j \sim \text{unif}(\mathbb{P})$ before $\mathcal{G}$ and a corresponding unitary $UP_jU^\dagger$ after, in order to reach $\mathcal{N}$ with P_j from both sides [21, 22]. It is essential that $UP_jU^\dagger$ comprise only single-qubit gates, which are typically much less noisy than multi-qubit gates, so as not to introduce more noise comparable to $\mathcal{N}$ while trying to twirl $\mathcal{N}$. This locality is guaranteed when U is Clifford, in which case $UP_jU^\dagger \propto P_i$ is an n -qubit Pauli (up to a global phase) for every $P_j \in \mathbb{P}$, by definition. Step 1 therefore consists of applying random Paulis P_j and P_i on either side of $\mathcal{G}$ according to the distribution described above. Under the approximation that Pauli gates are noiseless compared to multi-qubit gates, which we will make from now on, the resulting overall channel is $\tilde{\mathcal{G}} = \mathcal{U}\tilde{\mathcal{N}}$, where $\tilde{\mathcal{N}}$ is a Pauli channel. Thus simplified, the noise is easily amenable to mitigation.




Figure 1: The two conceptual steps of PEC, and the related form of ZNE, for a noisy Clifford gate $\mathcal{G}$. Any $\mathcal{G}$ can be described as a noise channel $\mathcal{N}$ followed by an ideal gate U . (We write the unitary U in place of the channel $\mathcal{U}$ in the circuit above, and likewise for other noiseless gates, for simplicity.) In Step 1, one adds random Pauli gates on both sides of $\mathcal{G}$, chosen so as to twirl $\mathcal{N}$ into a Pauli channel $\tilde{\mathcal{N}}$. We denote the resulting average channel as $\tilde{\mathcal{G}}$. In Step 2, one then adds a random Pauli gate P_k before $\tilde{\mathcal{G}}$, sampled from a probability (for ZNE) or quasi-probability (for PEC) distribution $\tilde{q}$, which is chosen to correctly amplify ($\alpha > 0$) or invert ($\alpha = -1$) $\tilde{\mathcal{N}}$, respectively. We denote the resulting aggregate channel as $\mathcal{A}$, and use the notation $\stackrel{\text{avg}}{=}$ to indicate that, for PEC, the relation only holds for expectation values.

In Step 2, both PEC and the related version of ZNE perform operations of the form

$$\mathcal{M}(\rho) = \sum_k q_k P_k \rho P_k \quad (6)$$

before $\bar{\mathcal{G}}$, leading to an aggregate channel of

$$\mathcal{A} = \bar{\mathcal{G}}\mathcal{M} = \mathcal{U}\bar{\mathcal{N}}\mathcal{M} \quad (7)$$

with a PTM of

$$\mathbf{A} = \mathbf{U} \left(\overbrace{\begin{pmatrix} \ddots & & 0 \\ & \vec{f} & \\ 0 & & \ddots \end{pmatrix}}^{\bar{\mathbf{N}}=\text{diag}(\vec{f})} \overbrace{\begin{pmatrix} \ddots & & 0 \\ & \mathbf{W}\vec{q} & \\ 0 & & \ddots \end{pmatrix}}^{M=\text{diag}(\mathbf{W}\vec{q})} \right), \quad (8)$$

expressed in terms of the PTMs for $\mathcal{A}$, $\mathcal{U}$, $\bar{\mathcal{N}}$ and $\mathcal{M}$ respectively, where $f_i = \text{tr}[P_i \mathcal{N}(P_i)]/2^n$. In other words, the aggregate channel behaves like an ideal gate $\mathcal{U}$ preceded by an adjustable Pauli noise channel $\bar{\mathcal{N}}\mathcal{M}$, which depends on $\vec{q} = (q_k)_{k=0}^{4^n-1}$ via $\mathcal{M}$. We now address ZNE and PEC in turn, which differ in their choice and their implementation of $\mathcal{M}$.

The idea of ZNE is to purposely increase the effective strength of gate noise so as to measure an expectation value of interest at multiple noise levels, then predict its zero-noise value by extrapolation. It is a heuristic technique whose performance depends on how exactly one increases the noise level. The most successful approach to date (sometimes called probabilistic error amplification, or PEA) seeks to effectively replace $\bar{\mathcal{N}}$ with $\bar{\mathcal{N}}^{1+\alpha}$ for different noise levels $1+\alpha \geq 1$ by picking $\mathcal{M} = \bar{\mathcal{N}}^\alpha$ [10]. As Eq. (8) shows, this can be done by applying random Paulis P_k before $\bar{\mathcal{G}}$ with probabilities q_k chosen so that $\mathbf{W}\vec{q} = \vec{f}^\alpha$, where $\vec{f}^\alpha$ is defined element-wise. The net effect is to amplify the twirled noise by a tunable amount $1+\alpha$ while preserving its structure.

Rather than amplify the twirled noise, PEC seeks to cancel it by picking $\alpha = -1$ so that $\mathcal{M} = \bar{\mathcal{N}}^{-1}$ and therefore $\mathcal{A} = \mathcal{U}$. As per Eq. (8), this can be done by picking $\vec{q}$ such that $\mathbf{W}\vec{q} = \vec{f}^{-1}$, where $(\vec{f}^{-1})_i = 1/f_i$. This $\vec{q}$, however, generally contains negative elements and is therefore not a valid probability distribution. In turn, $\mathcal{M} = \bar{\mathcal{N}}^{-1}$ is not a valid quantum channel, and cannot be implemented as described above in the context of ZNE. It is nonetheless possible to realize this $\mathcal{M}$ in effect, when measuring expectation values, by treating $\vec{q}$ as a quasi-probability distribution. That is, suppose we aim to measure $\langle P_m \rangle = \text{tr}[P_m \mathcal{U}(\rho)]$ for some n -qubit Pauli P_m , but can only implement the noisy gate $\bar{\mathcal{G}}$ (with twirled noise) in place of $\mathcal{U}$. PEC provably recovers the noiseless expectation value, on average, by applying random Paulis P_k before $\bar{\mathcal{G}}$ with probabilities $|q_k|/\gamma$, where $\gamma = \sum_k |q_k|$, then multiplying the ± 1 measurement outcomes (corresponding to eigenspaces of P_m) by $\gamma \text{sgn}(q_k)$ [2]. Assuming noiseless readout, the expected value of these scaled outcomes is

$$\sum_k \frac{|q_k|}{\gamma} \gamma \text{sgn}(q_k) \text{tr}[P_m \bar{\mathcal{G}}(P_k \rho P_k)] = \text{tr}\{P_m \bar{\mathcal{G}}[\mathcal{N}^{-1}(\rho)]\}, \quad (9)$$

which equals $\langle P_m \rangle$ as desired. However, because each shot returns $\pm\gamma$ rather than ± 1 , and $\gamma \geq 1$, one typically needs γ^2 times more shots to estimate $\langle P_m \rangle$ with a given precision than if the gate were noiseless [2, 9]. Moreover, the γ factors multiply when one does PEC for multiple gate layers within a circuit, resulting in a sampling overhead that (typically) grows exponentially in the number of noisy gates. The silver lining, however, is that γ approaches 1 here as gate noise decreases, so PEC could be compatible with classically hard circuits despite this exponential overhead, provided the gate noise is sufficiently weak.

2.3 Non-Clifford Gates

The error mitigation schemes described above, which we will call *Clifford PEC* and *ZNE*, rely critically on $\bar{\mathcal{N}}$ being a Pauli channel due to the Pauli-twirling in Step 1. In terms of PTMs, as in Eq. (8), both techniques enact a diagonal matrix $\mathbf{M}$ before the noisy gate in Step 2. This suffices to amplify or cancel the twirled noise because the latter's PTM, $\bar{\mathbf{N}}$, is also diagonal, so we can set $\mathbf{M} = \bar{\mathbf{N}}^\alpha$ for any α . (We could equally well perform $\mathcal{M}$ after the noisy gate instead, with minor adjustments,

following the alternate factorization $\mathcal{G} = \mathcal{N}'\mathcal{U}$.) If the effective noise channel were not Pauli, e.g., if we tried to implement $\mathbf{M} = \mathbf{N}^\alpha$ directly for a generic, non-diagonal $\mathbf{N}$, Step 2 would generally require multi-qubit non-Clifford gates, which would introduce more noise that also must be mitigated, leading to a circular situation. But this is exactly the problem we face when U is non-Clifford, as it is not generally possible to Pauli-twirl the associated noise using only single-qubit gates since $UP_jU^\dagger$ can be entangling, meaning Step 1 breaks down [23]. More colloquially, there is no way to reach $\mathcal{N}$ from both sides with arbitrary Paulis without introducing more entangling gates, which are themselves noisy. Or alternatively, in terms of PTMs, there is no apparent way to twirl $\mathbf{N}$ into a diagonal matrix which can, in turn, be amplified or inverted by a diagonal $\mathbf{M}$. That means that Clifford PEC and ZNE cannot correctly amplify or cancel noise on non-Clifford gates.

2.3.1 Formalism of Pauli Shaping

We propose a simple generalization of these techniques that applies to both Clifford and non-Clifford gates. To motivate it, we begin with a thought-experiment: consider a noisy non-Clifford gate $\mathcal{G} = \mathcal{N}'\mathcal{U}$ shown in Fig. 2, where $\mathcal{U}(\rho) = U\rho U^\dagger$ is the intended (non-Clifford) unitary and $\mathcal{N}'$ happens to be a Pauli noise channel from the outset. Since there is no need to twirl such noise, we could still do Clifford PEC or ZNE (skipping Step 1) by inserting $\mathcal{M} = (\mathcal{N}')^\alpha$, as in Eq. (6), *after* $\mathcal{G}$. Notice, however, that if we factored the same noisy gate in the order $\mathcal{G} = \mathcal{U}\mathcal{N}$ instead, the resulting noise channel $\mathcal{N} = \mathcal{U}^{-1}\mathcal{N}'\mathcal{U}$ would generally be non-Pauli (since U is non-Clifford), and could not be amplified or inverted by inserting an $\mathcal{M}$ *before* $\mathcal{G}$, as described in the previous section. Moreover, there is no apparent way to twirl $\mathcal{N}$ into a Pauli channel using single-qubit gates since U is not Clifford. This thought-experiment suggests two conclusions about mitigating errors on non-Clifford gates. First, whether we insert $\mathcal{M}$ before or after $\mathcal{G}$ can make a critical difference for PEC and ZNE, unlike in the Clifford case. We should therefore seek a formalism that finds the right placement automatically. Second, while it is always possible to factor a noisy gate $\mathcal{G}$ into noise and an ideal gate, doing so for non-Clifford gates can give qualitatively different noise channels (e.g., Pauli or non-Pauli) depending on the factorization order, which is an arbitrary mathematical choice of no physical significance. It can therefore be more informative to think in terms of the noisy gate $\mathcal{G}$ directly, rather than factoring out a noise channel.




Figure 2: A thought-experiment where a noisy non-Clifford gate $\mathcal{G}$ happens to factorize into an ideal gate U followed by a Pauli noise channel $\mathcal{N}'$, as in the middle circuit. $\mathcal{N}'$ can therefore be amplified or inverted by inserting random Paulis after the noisy gate. However, if we factorized $\mathcal{G}$ as shown on the right, the resulting noise channel $\mathcal{N}$ would generally be non-Pauli, and could not be amplified or inverted by inserting random Paulis before $\mathcal{G}$. This is a fundamental difference between Clifford and non-Clifford gates.

In light of these conclusions, we propose the following scheme. Consider performing P_j , then a noisy gate $\mathcal{G}$, then P_i , as shown in Fig. 3, where P_i and P_j are arbitrary n -qubit Paulis. The resulting channel

$$\mathcal{A}^{(ij)}(\rho) = P_i \mathcal{G}(P_j \rho P_j) P_i \quad (10)$$

has a PTM $\mathbf{A}^{(ij)}$ whose $(k, \ell)^{\text{th}}$ element is

$$\mathbf{A}_{k\ell}^{(ij)} = \text{tr} [P_k \mathcal{A}^{(ij)}(P_\ell)] / 2^n = \mathbf{W}_{ki} \mathbf{W}_{j\ell} \mathbf{G}_{k\ell}, \quad (11)$$

where we have used the trace's cyclic property, the fact that $P_a P_b P_a = \mathbf{W}_{ab} P_b$ for all $P_a, P_b \in \mathbb{P}$, and the definition of the PTM elements $\mathbf{G}_{k\ell}$ of $\mathcal{G}$ from Eq. (2). Now consider the aggregate channel $\mathcal{A}$ that is a linear combination of $\mathcal{A}^{(ij)}$, weighted by real coefficients $\mathbf{Q}_{ij}$ forming a $4^n \times 4^n$ matrix $\mathbf{Q}$ of our




Figure 3: A circuit description of Pauli shaping. To shape a given channel $\mathcal{G}$ into a desired channel $\mathcal{A}$, one can insert random n -qubit Paulis P_j and P_i before and after $\mathcal{G}$, respectively, to form a channel $\mathcal{A}^{(ij)}$. One picks P_i and P_j randomly with probability $|Q_{ij}|/\gamma$ in each shot, then multiplies the measurement outcomes by $\gamma \text{sgn}(Q_{ij})$ to realize $\mathcal{A}$, for $\mathbf{Q}$ and γ from Eqs. (14) and (15) respectively. When all $Q_{ij} \geq 0$, i.e., when $\mathbf{Q}$ is a valid probability distribution of pairs of Paulis, the last step is trivial and the channel $\mathcal{A}$ is actually realized. When some $Q_{ij} < 0$, $\mathbf{Q}$ is instead a quasi-probability distribution, and $\mathcal{A}$ is only realized in terms of expectation values, as indicated by the notation $\stackrel{\text{avg}}{=}$.

choice:

$$\mathcal{A}(\rho) = \sum_{ij} Q_{ij} \mathcal{A}^{(ij)}(\rho) = \sum_{ij} Q_{ij} P_i \mathcal{G}(P_j \rho P_j) P_i. \quad (12)$$

Using Eq. (11), the $(k, \ell)^{\text{th}}$ PTM element of $\mathcal{A}$ is

$$A_{k\ell} = \sum_{ij} W_{ki} Q_{ij} W_{j\ell} G_{k\ell}, \quad (13)$$

so its PTM, in matrix form, is

$$\mathbf{A} = (\mathbf{W} \mathbf{Q} \mathbf{W}) \odot \mathbf{G} = \mathbf{C} \odot \mathbf{G}, \quad (14)$$

where $\odot$ denotes an element-wise (i.e., Hadamard) product. It is convenient to define $\mathbf{C} = \mathbf{W} \mathbf{Q} \mathbf{W}$ in Eq. (14), which we call a *characteristic matrix* in analogy to characteristic functions from probability theory, which are Fourier transforms of probability density functions [24]. Intuitively, $\mathbf{Q}$ and $\mathbf{C}$ can be understood as Fourier transforms of one another (since $\mathbf{W}$ describes a type of discrete Fourier transform), so the element-wise product in Eq. (14) is reminiscent of a convolution in some “frequency domain.”

More concretely, Eq. (14) shows that through an appropriate choice of $\mathbf{Q}$, the aggregate channel $\mathcal{A}$ can be chosen almost arbitrarily, at least in terms of expectation values, at the cost of a potential sampling overhead. That is, to realize a desired $\mathcal{A}$ with PTM elements A_{ij} , it suffices to pick characteristic matrix elements $C_{ij} = A_{ij}/G_{ij}$. (Anytime $A_{ij} = G_{ij} = 0$, the corresponding C_{ij} can be chosen arbitrarily.) The resulting $\mathbf{C}$ then corresponds to a unique $\mathbf{Q} = \mathbf{W} \mathbf{C} \mathbf{W} / 2^{4n}$. If all $Q_{ij} \geq 0$ then $\mathbf{Q}$ can be interpreted as a probability distribution over *pairs* of Paulis $(P_j, P_i) \in \mathbb{P} \times \mathbb{P}$, rather than over individual Paulis like $\vec{q}$ in Eq. (6). In other words, the corresponding $\mathcal{A}$ can be realized, with no sampling overhead, by performing P_j , then $\mathcal{G}$, then P_i with probability Q_{ij} , as in Fig. 3. (Normalization, $\sum_{ij} Q_{ij} = 1$, is guaranteed if $\mathcal{A}$ and $\mathcal{G}$ are both trace-preserving.) Much like in the Clifford case, we can still realize a desired $\mathcal{A}$ in expectation when the corresponding $\mathbf{Q}$ contains negative elements by treating $\mathbf{Q}$ as a quasi-probability distribution. That is, we can insert P_j and P_i before and after $\mathcal{G}$ respectively with probability $|Q_{ij}|/\gamma$, where

$$\gamma = \sum_{ij} |Q_{ij}|, \quad (15)$$

then multiply the measurement outcomes (the measured eigenvalue of the observable in question) by $\gamma \text{sgn}(Q_{ij})$. The proof is almost identical to that of the Clifford case (see Appendix A), and the meaning of γ is the same: γ^2 is a sampling overhead that combines multiplicatively with that from other gate layers, leading to an exponential overhead. We are not aware of any prior name for this technique, which is encapsulated by Eq. (14) and Fig. 3, so we will refer to it here as *Pauli shaping*. Moreover, we will generally refer to $\mathbf{Q}$ as a quasi-probability matrix, and similarly for its elements, even though it can also describe a true probability distribution.

Rather than just invert or amplify Pauli noise, Pauli shaping effectively transforms any implemented channel $\mathcal{G}$ into (almost) any desired channel $\mathcal{A}$. It applies to both Clifford and non-Clifford gates. (The only minor limitation is that it requires $\mathbf{G}_{ij} \neq 0$ in order to achieve $\mathbf{A}_{ij} \neq 0$ for any Paulis $P_i, P_j \in \mathbb{P}$, otherwise no choice of characteristic matrix will satisfy $\mathbf{A}_{ij} = \mathbf{C}_{ij} \mathbf{G}_{ij}$. While it is unclear how to rephrase this condition in more conventional terms, we expect it to be easily satisfied outside of pathological scenarios involving fine-tuned or infinite noise, as illustrated in Sec. 2.3.2.) For instance, one can use Pauli shaping to do PEC for arbitrary gates by choosing $\mathbf{C}$ so that $\mathcal{A} = \mathcal{U}$, i.e., by demanding $\mathbf{A}_{ij} = \mathbf{U}_{ij}$. Similarly, one could do several forms of ZNE for different notions of noise amplification. One possibility would be to mimic the Clifford case by picking $\mathbf{C}$ such that $\mathcal{A} = \mathcal{U}\mathcal{N}^{1+\alpha}$ for different noise levels $1 + \alpha \geq 1$, even though $\mathcal{N} = \mathcal{U}^{-1}\mathcal{G}$ need not be a Pauli channel. (The other noise factorization order works too.) Alternatively, one could aim to find a Lindbladian $\mathcal{L}$ such that $\mathcal{G} = \exp(-i\mathcal{H} + \mathcal{L})$, where $\mathcal{H} = [H, \cdot]$ is an effective Hamiltonian superoperator (with $H = H^\dagger$) that generates the intended gate $\mathcal{U} = e^{-i\mathcal{H}}$, then similarly implement $\mathcal{A} = \exp(-i\mathcal{H} + \alpha\mathcal{L})$ for different α . Variants of these schemes where the noise is twirled over a subset of $\mathbb{P}$, depending on the intended gate, are also possible.

Pauli shaping reduces to Clifford PEC/ZNE when the target gate U is Clifford. This may not be obvious since the latter is typically broken into two conceptual steps (as in Sec. 2.2), which can obscure the full picture: while Step 2 only adds random Paulis on one side of the noisy gate $\mathcal{G}$, Step 1 adds them on both sides to twirl the noise. By combining any adjacent Paulis as shown in Fig. 4, these two steps can be jointly described as inserting random Paulis P_j and P_i before and after a noisy gate $\mathcal{G}$, respectively, with quasi-probability

$$\mathbf{Q}_{ij} = 4^{-n} q_{\sigma(i) \oplus j}, \quad (16)$$

where $\vec{q} = (q_k)_{k=0}^{4^n-1}$ are the coefficients appearing throughout Sec. 2.2, the function σ is defined by $P_{\sigma(i)} \propto U^\dagger P_i U$, and we write $k = i \oplus j$ when $P_k \propto P_i P_j$. These are the same probabilities, or quasi-probabilities, as one gets by starting with the formalism of Pauli shaping (see Appendix B). The key difference between Pauli shaping and Clifford PEC/ZNE, then, is that the former allows more general correlations between the random Paulis flanking $\mathcal{G}$, giving it a much broader scope without requiring deeper circuits. More precisely, Pauli shaping uses a quasi-probability matrix $\mathbf{Q}$ with up to $O(4^{2n})$ distinct elements, whereas Clifford PEC/ZNE implicitly uses a $\mathbf{Q}$ with only $O(4^n)$ distinct elements, repeated according to Eq. (16). Of course, it is not feasible to compute $2^{O(n)}$ quasi-probabilities in either case for modern quantum devices. In Clifford PEC/ZNE, “sparse” noise models, which approximate $\bar{\mathcal{N}}$, or equivalently $\bar{\mathcal{G}}$, using only $\text{poly}(n)$ parameters based on hardware connectivity, have been empirically successful [9, 10]. (See also [18, 25] for related theoretical results.) We expect a similar approach to be possible for non-Clifford gates, on similar grounds, but we leave such modeling to follow-up work.




Figure 4: Two equivalent ways to view Clifford PEC/ZNE. Left: the conceptual steps are shown separately, where the random Paulis $P_i \sim \text{unif}(\mathbb{P})$ and $P_j = P_{\sigma(i)} \propto U^\dagger P_i U$ serve to twirl the factored noise channel $\mathcal{N}$, then a random Pauli P_k is added with quasi-probability $k \sim \vec{q}$ to amplify or invert the twirled noise. Right: P_k and P_j are combined into a single Pauli $P_\ell = P_{k \oplus \sigma(i)}$, so P_ℓ and P_i are added before and after $\mathcal{G}$, respectively, with quasi-probability $q_k/|\mathbb{P}| = 4^{-n} q_{\sigma(i) \oplus \ell}$.

2.3.2 Examples of Pauli Shaping

To illustrate the scope and behavior of Pauli shaping more broadly, we present three concrete examples in this section—two analytical and one numerical. While Pauli shaping can be applied to arbitrary gates, these examples focus on a 2-qubit gate that is common in many experiments [26, 27, 28]:

$$U = R_{ZZ}(\theta), \quad (17)$$

with a generic angle $\theta \notin \{0, \pm\pi/2, \pi\}$ for which the gate is non-Clifford. When constructing PTMs and the like related to this gate, it will be convenient to order the 2-qubit Paulis as

$$\mathbb{P} = (II, ZZ, XX, YY, IZ, ZI, YX, XY, XI, YZ, XZ, YI, IX, ZY, ZX, IY), \quad (18)$$

so that those in the first/second half commute/anti-commute with ZZ , and each Pauli is related to one of its neighbors through multiplication by ZZ . (Of course, the choice of basis ordering is arbitrary and has no impact on the results that follow.) Expressed in the ordered basis (18), the PTM of the channel $\mathcal{U}(\rho) = U\rho U^\dagger$ describing the ideal gate is block-diagonal:

$$U = \begin{pmatrix} I & & & & & & & 0 \\ & I & & & & & & \\ & & I & & & & & \\ & & & I & & & & \\ & & & & R^{(2)} & & & \\ & & & & & R^{(2)} & & \\ & & & & & & R^{(2)} & \\ 0 & & & & & & & R^{(2)} \end{pmatrix}, \quad (19)$$

with 2×2 identity matrices I in the upper-left (acting trivially on Paulis that commute with ZZ) and 2×2 rotation matrices

$$R^{(2)} = \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix} \quad (20)$$

in the lower-right (mixing together Paulis that anti-commute with ZZ).

Example 1- To illustrate the enhanced scope of Pauli shaping, suppose the only error in implementing U is a coherent over- or under-rotation by an angle $\epsilon \in \mathbb{R}$, that is:

$$\mathcal{G}(\rho) = R_{ZZ}(\theta + \epsilon) \rho R_{ZZ}(\theta + \epsilon)^\dagger. \quad (21)$$

This is a paradigmatic non-Pauli error [23], which cannot be twirled into a Pauli error channel because U is non-Clifford. Pauli shaping lets us recover the ideal gate², in terms of expectation values, by using a block-diagonal characteristic matrix (also expressed in the ordered basis (18)):

$$\mathbf{C} = \text{diag}(C_1, C_1, C_1, C_1, C_2, C_2, C_2, C_2), \quad (22)$$

where the four bottom blocks must be

$$C_2 = \begin{pmatrix} \frac{\cos(\theta)}{\cos(\theta+\epsilon)} & \frac{\sin(\theta)}{\sin(\theta+\epsilon)} \\ \frac{\sin(\theta)}{\sin(\theta+\epsilon)} & \frac{\cos(\theta)}{\cos(\theta+\epsilon)} \end{pmatrix} \quad (23)$$

in order to satisfy Eq. (14). Similarly, the four top blocks of $\mathbf{C}$ must have the form

$$C_1 = \begin{pmatrix} 1 & x \\ y & 1 \end{pmatrix}, \quad (24)$$

²The only mathematical requirement for Pauli shaping, namely that $\mathbf{G}_{ij} \neq 0$ wherever $\mathbf{A}_{ij} \neq 0$, is met here for all values of θ and ϵ except those where $2(\theta + \epsilon)$ is an integer multiple of π ; that is, the miscalibrated gate $R_{ZZ}(\theta + \epsilon)$ must not be exactly Clifford. This is more of a theoretical concern than a practical one.

where the parameters x and y can be chosen arbitrarily since both get multiplied by zero in Eq. (14), so their values do not affect the resulting aggregate channel. They do, however, generally affect the required sampling overhead, in potentially complicated ways. A judicious choice of x and y (see Appendix C for details) cancels the coherent error, in expectation, at the cost of a γ^2 sampling overhead, where

$$\gamma = \max \left\{ \left| \frac{\cos(\theta)}{\cos(\theta + \epsilon)} \right|, \left| \frac{\sin(\theta)}{\sin(\theta + \epsilon)} \right| \right\} = 1 + O(\epsilon). \quad (25)$$

In other words, the sampling overhead (per gate) vanishes linearly as the calibration error goes to zero.

Example 2- Unfortunately, while Pauli shaping applies to arbitrary gates, it can be prohibitively expensive on non-Clifford gates depending on the noise in question. To illustrate the issue, consider a different noisy implementation of $U = R_{ZZ}(\theta)$ with a block-diagonal PTM of

$$\mathbf{G} = \begin{pmatrix} \begin{pmatrix} 1 & 0 \\ 2\epsilon & 1-2\epsilon \end{pmatrix} & & & & & & & 0 \\ & \begin{pmatrix} 1-\epsilon & -\epsilon \\ -\epsilon & 1-\epsilon \end{pmatrix} & & & & & & \\ & & \begin{pmatrix} 1-\epsilon & \epsilon \\ \epsilon & 1-\epsilon \end{pmatrix} & & & & & \\ & & & \begin{pmatrix} 1-\epsilon & \epsilon \\ \epsilon & 1-\epsilon \end{pmatrix} & & & & \\ & & & & \sqrt{1-2\epsilon} R^{(2)} & & & \\ & & & & & \sqrt{1-2\epsilon} R^{(2)} & & \\ & & & & & & \sqrt{1-2\epsilon} R^{(2)} & \\ 0 & & & & & & & \sqrt{1-2\epsilon} R^{(2)} \end{pmatrix}, \quad (26)$$

where $\epsilon \in [0, \frac{1}{2})$ describes the noise strength³ ($\epsilon = 0$ for the ideal gate) and $R^{(2)}$ is the 2×2 rotation matrix from Eq. (20). The optimal characteristic matrix (in terms of γ) that recovers the ideal gate, in expectation, is similarly block-diagonal (see Appendix C). To satisfy $\mathbf{U} = \mathbf{C} \odot \mathbf{G}$, it must equal

$$\mathbf{C} = \text{diag}(C_1, C_2, C_2, C_2, C_3, C_3, C_3, C_3), \quad (27)$$

where

$$C_1 = \begin{pmatrix} 1 & x \\ 0 & (1-2\epsilon)^{-1} \end{pmatrix} \quad (28)$$

for some parameter x that we are free to choose, and

$$C_2 = \frac{1}{1-\epsilon} \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \quad C_3 = \frac{1}{\sqrt{1-2\epsilon}} \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}. \quad (29)$$

The resulting γ is a complicated function of ϵ and x in general (see Appendix C), but its noiseless limit is

$$\lim_{\epsilon \rightarrow 0} \gamma = \frac{1}{32} (28|x| + 3|x-8| + |x+24|), \quad (30)$$

which attains a minimum value of $\gamma = 1.5$ at $x = 0$.

While contrived, this second example illustrates a stark difference between Pauli shaping on Clifford and non-Clifford gates. With Clifford gates, the sampling overhead required to cancel any (twirled) noise, in expectation, approaches 1 in the noiseless limit. The same is not true, in general, for non-Clifford gates, where canceling certain types of noise can require a large sampling overhead, regardless of how weak the noise is. (Since the overheads from different gates combine multiplicatively, $\gamma \approx 1.5$ is effectively much larger than $\gamma \approx 1$ since it leads to a total overhead that grows much faster with circuit size.) In the case of Eq. (26), the issue, mathematically, comes from the $O(\epsilon)$ off-diagonal

³This example can be derived from a Lindblad equation (see Appendix C) in which $\epsilon \rightarrow \frac{1}{2}$ describes the limit of infinite noise strength. The mathematical requirement for Pauli shaping that $\mathbf{G}_{ij} \neq 0$ wherever $\mathbf{A}_{ij} \neq 0$, is met here for all for all $\epsilon < \frac{1}{2}$, i.e., all finite noise strengths.




Figure 5: A noisy $R_{ZZ}(\theta)$ gate used as a numerical example. *Left:* The PTM of the channel $\mathcal{G}$ constructed in Example 3 of Sec. 2.3.2, which describes a noisy $U = R_{ZZ}(47^\circ)$ gate. *Center:* The PTM of the associated noise channel $\mathcal{N} = U^{-1}\mathcal{G}$, with the identity matrix subtracted to reduce the range of values and thereby improve visibility. *Right:* The relative error in learning the PTM elements of this $\mathcal{G}$ that are essential for Pauli shaping, using the schemes developed in Sec. 3. The learned elements are all accurate to within $\lesssim 1\%$ (with room for further optimization), even though the underlying simulation includes state preparation and measurement (SPAM) errors, each with only 90% average fidelity [29]. The elements in gray were not learned, as their values are not needed to cancel the noise through Pauli shaping, which requires one to find numbers $\{C_{ij}\}_{i,j=0}^{15}$ such that $U_{ij} = C_{ij}G_{ij}$ (as in Eq. (14), with $A = U$). There is no need to learn $G_{0,0}$, as it equals 1 for any CPTP map, which fixes $C_{0,0} = 1$. For all other grayed out indices, $U_{ij} = 0$ (see Eqs. (19) and (20)), so we can pick $C_{ij} = 0$ regardless of G_{ij} .

elements in the upper-left part of $\mathcal{G}$, which cannot be twirled away without spoiling the action of the gate. Clifford gates have no such elements—rather, all elements of their PTMs that should be zero (for noiseless gates) can be twirled away in noisy gates. In this second non-Clifford example, however, the only way to suppress the aforementioned PTM elements is to set the corresponding elements of $\mathcal{C}$ to zero, leading to a large γ . (This is true whenever these PTM elements are nonzero, regardless of their exact values, i.e., even if they did not all equal $\pm\epsilon$ as in Eq. (26).) Because Eq. (14) involves an element-wise product of $\mathcal{C}$ and $\mathcal{G}$, instead of a typical matrix multiplication, this holds true even for arbitrarily small $\epsilon > 0$. In other words, the only number C_{ij} for which $C_{ij} \times \epsilon = 0$ is $C_{ij} = 0$, no matter how small ϵ might be, which can impose a prohibitive lower bound on γ even when the noise is weak. This issue is illustrated numerically in the next example.

Example 3- The last two examples considered error models with very particular structures in order to be analytically tractable. In this one we take the opposite approach by choosing a complicated error model, generated at random with no special structure, and analyzing it numerically. Specifically, suppose one intends to perform a $U = R_{ZZ}(\theta)$ gate with $\theta = 47^\circ$, a nondescript angle chosen for illustration so that U is neither Clifford, nor even close to a Clifford gate. Suppose also that, in place of U , one can only implement a noisy version described by the channel $\mathcal{G} = \exp(-i\mathcal{H} + \mathcal{L})$ where the first term, defined by $\mathcal{H}(\rho) = \frac{\theta}{2}[ZZ, \rho]$, describes the ideal gate, and the second term is a perturbation of general Lindblad form (to ensure that $\mathcal{G}$ remains CPTP):

$$\mathcal{L}(\rho) = -i\epsilon_H[H_{\text{err}}, \rho] + \epsilon_D \sum_{j,k=1}^{15} \Gamma_{jk} \left(P_j \rho P_k - \frac{1}{2} \{P_k P_j, \rho\} \right). \quad (31)$$

To form $\mathcal{G}$, we picked generic matrices $H_{\text{err}} = H_{\text{err}}^\dagger$ and $\Gamma = \Gamma^\dagger \geq 0$ at random by sampling elements from standard normal distributions. We then set the strengths of the Hamiltonian and dissipative noise terms, ϵ_H and ϵ_D respectively, through bisection so that $\mathcal{G}$ has an average fidelity [29] of $F_{\text{avg}} = 0.95$ with respect to U , and the noise unitarity [30] is $u = 0.9$. These two numbers were chosen for illustration,




Figure 6: A numerical example of noise cancellation through Pauli shaping. *Top*: An illustrative circuit used to estimate $\langle ZX \rangle$ as a function of ϕ , where $|y_+\rangle$ denotes the $+1$ eigenstate of Y . *Bottom*: Implementing $\mathcal{G}$, the noisy gate depicted in Fig. 5a, in place of the ideal $U = R_{ZZ}(47^\circ)$ gates generally yields incorrect values of $\langle ZX \rangle$, as shown by the two curves. (Both use perfect initial states and measurements.) However, Pauli shaping can recover the correct expectation values despite using $\mathcal{G}$. Every data point is the mean from 10^7 shots, each from an independent noisy circuit drawn from the quasi-probability distribution given by Eq. (14) with $\mathbf{A} = \mathbf{U}$. The gray points use a distribution derived from the true PTM $\mathbf{G}$, whereas the blue points use one derived from an imperfect version thereof, which is learned using the schemes from Sec. 3, and described in Fig. 5c. Error bars show the estimated standard error of the mean.

to give strong noise that is somewhat non-unital, and partway between the fully incoherent and fully coherent. The PTM of the resulting $\mathcal{G}$ is shown in Fig. 5a, and that of the corresponding non-Pauli noise channel $\mathcal{N} = \mathcal{U}^{-1}\mathcal{G}$ is depicted in Fig. 5b.

To demonstrate the broad applicability of Pauli shaping, we used it to mitigate repeated $\mathcal{G}$ gates in the circuit from the top panel of Fig. 6a. As shown in Fig. 6b, this gives the correct observable expectation values⁴. However, because $\mathbf{G}_{i,i+1}$ and $\mathbf{G}_{i+1,i}$ are mostly nonzero for $0 \leq i \leq 7$, as in Example 2, this noise cancellation requires $\gamma \approx 1.67$. Note that while Example 1 features coherent noise and small γ , whereas Examples 2 and 3 feature (at least partly) incoherent noise and large γ , one should not infer a trend here. That is, noise coherence is not what causes γ to be large or small. As a simple counter-example, consider an $R_{ZZ}(\theta)$ gate with weak Pauli noise—the resulting γ is small like in Example 1, even though the noise is purely incoherent. The same is true for a mixture of over/under-rotation and Pauli noise.

We will discuss the phenomenon of large γ despite weak noise more broadly in Sec. 4. For now, however, we turn our attention to the challenge of accurately measuring the PTM $\mathbf{G}$ of a noisy non-Clifford gate. Since the characteristic matrix $\mathbf{C}$, and ultimately, the quasi-probability matrix $\mathbf{Q}$ in Eq. (14), depend on $\mathbf{G}$, addressing this challenge is essential for Pauli shaping.

⁴The mathematical requirement for Pauli shaping that $\mathbf{G}_{ij} \neq 0$ wherever $\mathbf{A}_{ij} \neq 0$ is manifestly met for $\mathbf{G}$ in Fig. 5a when $\mathbf{A} = \mathbf{U}$.

3 Characterizing Noisy Gates

The performance of both PEC and ZNE depends on how accurately one can characterize, or “learn,” the noisy gate in question, or equivalently, the factored noise channel. This is true for both Clifford and non-Clifford gates. A notable obstacle, however, is that the measurement outcomes from any circuit meant to characterize a noisy gate will also depend on state preparation and measurement (SPAM) errors, which can be significant [31]. It is therefore desirable to use learning schemes that can distinguish gate and SPAM errors as much as possible, so as to accurately isolate the former. Here we review such learning schemes for Clifford gates in Sec. 3.1, then introduce new ones suited for non-Clifford $R_{ZZ}(\theta)$ gates in Sec. 3.2. But first we discuss readout twirling, which is a common component throughout both sections [32].

Readout twirling consists of applying an X gate with 50% probability independently on each qubit before readout, then flipping the measured bits wherever such gates were applied. (We use “measurement” and “readout” synonymously.) In principle, these random X gates should be sampled independently in each circuit execution, i.e., each “shot.” Suppose we want to measure an n -qubit Pauli expectation value $\langle P_i \rangle = \text{tr}(P_i \rho)$ for $P_i \in \mathbb{P}$ and some state ρ , and we denote our estimate thereof after a finite number N_{tot} of independent shots as

$$\hat{\mu} = \frac{(\# \text{ of } +1 \text{ outcomes}) - (\# \text{ of } -1 \text{ outcomes})}{N_{\text{tot}}}, \quad (32)$$

where ± 1 outcomes refer to the observed eigenvalues of P_i . Because $\hat{\mu}$ will vary from one experiment to the next, even under identical conditions, we can treat it as a random variable whose expectation value, $\mathbb{E}(\hat{\mu})$, describes an average over many hypothetical experiments. Absent any measurement errors, $\mathbb{E}(\hat{\mu}) = \langle P_i \rangle$, i.e., $\hat{\mu}$ is an unbiased estimate of $\langle P_i \rangle$, which means we can get it arbitrarily close to the true value $\langle P_i \rangle$, with arbitrarily high probability, by simply taking enough shots. However, measurement errors could bias $\hat{\mu}$ in complicated ways, such that $\mathbb{E}(\hat{\mu})$ bears no simple relation to $\langle P_i \rangle$. This issue is partially remedied through readout twirling, which ensures that $\mathbb{E}(\hat{\mu}) = m_i \langle P_i \rangle$ for some coefficient m_i that depends on the statistics of the readout noise but not on the measured quantum state ρ . ($m_i = 1$ for ideal readout—see Appendix D.) In other words, readout twirling still gives a biased estimate for $\langle P_i \rangle$, but the bias has a predictable form that will let us distinguish gate errors from SPAM errors.

3.1 Clifford Gates

Consider a Clifford unitary U whose noisy implementation is described by the channel $\mathcal{G} = \mathcal{U}\mathcal{N}$, where $\mathcal{N}$ describes generic noise and $\mathcal{U}(\rho) = U\rho U^\dagger$, as in Sec. 2.2. We will assume throughout that the noise does not vary with time, and is independent of any previous gates. Since we can Pauli-twirl $\mathcal{N}$ using only single-qubit gates, it suffices to learn the twirled noise channel $\bar{\mathcal{N}}$. In the language of PTMs, we only need to learn $\bar{\mathbf{N}} = \text{diag}(\bar{f})$, since the Pauli fidelities $f_i = \mathbf{N}_{ii}$ are the only components of the noise that figure in Clifford PEC/ZNE. It is possible to learn these (at least in part) in a way that is robust to SPAM errors through cycle benchmarking (CB) [33], a variant of randomized benchmarking.

To introduce CB, we will begin with the simple case where $U = I$. This means that every Pauli $P_i \in \mathbb{P}$ is an eigenvector of the twirled, noisy identity gate $\bar{\mathcal{G}} = \mathcal{U}\bar{\mathcal{N}} = \bar{\mathcal{N}}$ with eigenvalue f_i —that is:

$$\bar{\mathcal{G}}(P_i) = f_i P_i. \quad (33)$$

Applying $\bar{\mathcal{G}}$ d times to an initial state $\rho = 2^{-n} \sum_j s_j P_j$ therefore leads to a final state of

$$\rho' = \bar{\mathcal{G}}^d(\rho) = \frac{1}{2^n} \sum_j s_j f_j^d P_j \quad (34)$$

with Pauli expectation values $\langle P_i \rangle = \text{tr}(P_i \rho') = s_i f_i^d$ that decay exponentially in the circuit depth d at rates f_i . CB exploits this relation by performing the following steps to estimate each Pauli fidelity f_i :

1. Prepare an initial state $\rho = 2^{-n} \sum_j s_j P_j$ for which $s_i = \text{tr}(P_i \rho)$ is as large as possible (to maximize the eventual signal-to-noise ratio). E.g., attempt to prepare $\rho = |\psi\rangle\langle\psi|$ where $|\psi\rangle$ is a separable $+1$ eigenstate of $P_i \in \mathbb{P}$, so $s_i = 1$ ideally.
2. Apply $\bar{\mathcal{G}}$ d times to ρ for varying depths d .
3. Estimate $\langle P_i \rangle$ for the resulting state $\bar{\mathcal{G}}^d(\rho)$ as in Eq. (32) using readout twirling, denoting the result by $\hat{\mu}$.

The expected value of $\hat{\mu}$ (which we will denote as μ), i.e., the average estimate of $\langle P_i \rangle$ for a circuit depth d from noisy experimental data, is

$$\mu := \mathbb{E}(\hat{\mu}) = s_i m_i \times f_i^d, \quad (35)$$

where the coefficients s_i and m_i depend on state preparation and measurement errors, respectively, but not on f_i or d . Therefore, even though $\mu \neq \langle P_i \rangle$ in general, CB obtains an estimate of f_i that is robust to SPAM by fitting the tuples $(d, \hat{\mu})$ to a function $d \mapsto Ar^d$ and extracting r . These steps are then repeated for all desired Pauli fidelities f_i (although it is possible to re-use some of the same experimental data to get different Pauli fidelities, as we discuss later).

While there are more sophisticated alternatives with substantially better scaling [34, 35, 36], CB's relative simplicity and robustness make it especially practical on large pre-fault-tolerant quantum devices. In fact, besides being robust to SPAM, CB is well-behaved in two other ways that are critical in many experiments but less often discussed theoretically. First, it is *sensitive*, in that a small change in f_i produces a large change in the data, so the number of shots needed is reasonably small. Second, it *concentrates*, meaning that the different random circuits that arise due to twirling give similar expectation values, so few of them are needed [37, 38]. We now elaborate on both points in turn.

Sensitivity: Suppose we estimate $\langle P_i \rangle$ for a circuit depth d , as per the CB steps described above. Each ± 1 measurement outcome follows a Bernoulli distribution⁵ with mean μ from Eq. (35). One way to quantify how much information about f_i is conveyed by each such outcome is through its (classical) Fisher information [24]

$$\mathcal{I}(f_i) = \frac{1}{1 - \mu^2} \left(\frac{\partial \mu}{\partial f_i} \right)^2 = \frac{\mu^2}{1 - \mu^2} \frac{d^2}{f_i^2}, \quad (36)$$

which diverges for large d in the noiseless limit (see Appendix E). Why consider this limit? Because we need $f_i \approx 1$ for error mitigation to be feasible in the first place, so while the exact value of f_i , and therefore also of $\mathcal{I}(f_i)$, is unknown *a priori*, a learning scheme that works well as $f_i \rightarrow 1$ should also work well in the relevant regime of f_i , by continuity. (Also, crucially, this limit is analytically tractable.) Concretely, $\mu \rightarrow f_i^d$ in the absence of SPAM errors, so we can re-write Eq. (36) as

$$\mathcal{I}(f_i) \xrightarrow[\text{errors}]{\text{no SPAM}} \frac{1}{f_i^2 \ln(f_i)^2} \frac{x^2}{4(e^x - 1)} \lesssim \frac{0.162}{f_i^2 \ln(f_i)^2}, \quad (37)$$

where we maximized over $x = 2d \ln(1/f_i)$ numerically in the last step, which is a convenient way to maximize over d in effect (see Appendix E). This maximum value of $\mathcal{I}(f_i)$ then diverges as $f_i \rightarrow 1$, thanks to the slow decay of μ versus d when $f_i \approx 1$, wherein a small change in f_i produces a big change in μ at large depths d , as quantified by $\partial \mu / \partial f_i$. Consequently, CB does not require an exorbitant number of shots to precisely estimate large Pauli fidelities, since each shot can be very informative.

Concentration: CB, as described above, uses a new random circuit for each shot because it must twirl the gate and measurement noise. For any given depth, this leads to independent and identically distributed (IID) measurement outcomes, which are relatively simple to analyze. For instance, the

⁵We use the term “Bernoulli” loosely here since this distribution is supported on $\{1, -1\}$ rather than $\{0, 1\}$.

standard error (i.e., the standard deviation of $\hat{\mu}$, a widely-used measure of statistical uncertainty) in estimating $\mu = \mathbb{E}(\hat{\mu})$ with N_{tot} shots has the familiar $O(1/\sqrt{N_{\text{tot}}})$ scaling, specifically:

$$\sqrt{\text{Var}(\hat{\mu})} = \sqrt{\frac{1 - \mu^2}{N_{\text{tot}}}}. \quad (38)$$

In many experiments, however, loading a new circuit into the control electronics used to implement gates is slow compared to running an already-loaded circuit [39]. It is therefore common practice to use a modified version of CB where, for any given depth, N_c different random circuits are chosen, each of which is run $N_{s/c} \geq 1$ times (the subscript stands for “shots per circuit”) for a total of $N_{\text{tot}} = N_c N_{s/c}$ shots. The resulting estimate of $\langle P_i \rangle$ (i.e., the sample average), denoted $\hat{\mu}'$, has the same mean as $\hat{\mu}$, and reduces to $\hat{\mu}$ when $N_{s/c} = 1$ as in proper CB. In general, however, it suffers from larger statistical fluctuations than $\hat{\mu}$, which do not generally scale as $O(1/\sqrt{N_{\text{tot}}})$ since the measurement outcomes are not IID. Rather, $\hat{\mu}'$ has a standard error of

$$\sqrt{\text{Var}(\hat{\mu}')} = \sqrt{\frac{1 - \mu^2}{N_{\text{tot}}} + \left(\frac{N_{s/c} - 1}{N_{s/c}}\right) \frac{\Delta^2}{N_c}}, \quad (39)$$

where

$$\Delta^2 = \text{Var}\{\text{tr}[P_i \mathcal{T}(\rho)]\} = \mathbb{E}\{\text{tr}[P_i \mathcal{T}(\rho)]^2\} - \mu^2 \quad (40)$$

is the variance in expectation values over different random, noisy circuits $\mathcal{T}$. (See Appendix F for details.) That is, CB defines many random noisy circuits $\mathcal{T}$, sometimes called “twirl circuits” or “twirl instances,” comprising d sequential noisy gates $\mathcal{G}$ interleaved by random Paulis, as depicted in Fig. 7a. Each such circuit can have a different expectation value, and Δ^2 is the variance thereof, as shown in Fig. 7b. It is an important, albeit rarely analyzed quantity, since the standard error of $\hat{\mu}'$ approaches $\Delta/\sqrt{N_c}$ as $N_{s/c} \rightarrow \infty$. A large Δ would therefore necessitate many different twirl circuits ($N_c \gg 1$) to precisely estimate μ for any given depth d , which can be very slow (in terms of wall-clock time). Fortunately, $\Delta \rightarrow 0$ in the noiseless limit for CB. This can be seen from Eq. (40), since $\text{tr}[P_i \mathcal{T}(\rho)] \rightarrow 1$ for all $\mathcal{T}$ in this limit. We will not seek a formal bound on Δ more broadly; rather, we simply note that the common practice of reusing a small number of random circuits many times is well-justified for CB when the noise is weak.




Figure 7: Cycle benchmarking, like the non-Clifford learning schemes introduced below, involves running random circuits comprising d noisy gates $\mathcal{G}$ interleaved with random Paulis drawn from some prescribed distribution. These gates are then followed by a readout noise channel flanked by random bit flips [32] (see Appendix D). We describe such a noisy circuit by a channel $\mathcal{T}$, as shown in the top panel. Each of these possible circuits can give a different Pauli expectation value $\text{tr}[P_i \mathcal{T}(\rho)]$, and some values arise more often than others, as illustrated in the bottom panel. (This is not quantum randomness from measurement outcomes, but rather, classical randomness from the choice of circuit.) The mean of this distribution, μ , describes our average estimate of $\langle P_i \rangle$, and the standard deviation Δ determines how many random circuits are needed to do so accurately.

These two properties, namely sensitivity and concentration for weak noise, do not arise automatically. As we will see in the next section, special care must be taken to maintain them when generalizing CB for non-Clifford gates.

Before moving on, however, we must consider non-trivial Clifford gates $U \neq I$. CB works very similarly in this setting—the key difference being that U now maps a generic Pauli P_i to a potentially different one $P_j = \pm U P_i U^\dagger$, so

$$\bar{\mathcal{G}}(P_i) = \pm f_i P_j. \quad (41)$$

(In other words, P_i is now a generalized eigenvector of $\bar{\mathcal{G}}$ [40].) Since the right-hand side is proportional to P_j , another application of $\bar{\mathcal{G}}$ will introduce a factor of f_j rather than another f_i . However, repeated applications of $\bar{\mathcal{G}}$ will eventually yield a term proportional to P_i , ultimately leading to an exponential decay similar to that in Eq. (35). Consider $U = \text{CNOT}$ and $P_i = IZ$ for illustration. This U maps P_i to $P_j = ZZ$ and vice versa, meaning:

$$\bar{\mathcal{G}}^d(P_i) = (f_i f_j)^{d/2} P_i \quad (42)$$

for even depths d . Therefore, by measuring $\langle P_i \rangle$ for various (even) circuit depths, as described earlier, one can obtain an estimate of $f_i f_j$ that is robust to SPAM errors. Unfortunately, this product does not uniquely specify f_i or f_j . One can partially sidestep this “degeneracy” and isolate certain Pauli fidelities by interleaving other single-qubit gates between applications of $\bar{\mathcal{G}}$. However, some f_i fundamentally cannot be learned in a SPAM-robust way for general Clifford gates [41]. Perhaps unsurprisingly, a similar issue will arise with non-Clifford gates as well. For CNOTs, principled guesses for f_i and f_j , given a measured value of $f_i f_j$, have led to very effective error mitigation [9, 10]. We will assume here that similar heuristics may work for non-Clifford gates too.

3.2 Non-Clifford Gates

While there exist some SPAM-robust methods for learning general, noisy non-Clifford gates (at least in part) [13, 14, 15, 16, 42, 43], it is easier to ensure sensitivity and concentration by devising schemes that are tailored for specific gates of interest. In particular, we will consider 2-qubit gates of the form $U = R_{PQ}(\theta)$, where $P, Q \in \{X, Y, Z\}$ and $\theta \notin \{0, \pm\pi/2, \pi\}$ is a generic rotation angle for which the gate is non-Clifford. Such gates arise naturally in several types of qubits [26, 27, 28]. For concreteness, we will focus on $U = R_{ZZ}(\theta)$ gates, from which the more general $R_{PQ}(\theta)$ case follows through a change of basis. In contrast with the CB formalism for Clifford gates from the previous section, it will be more natural here to learn PTM elements of the noisy gate $\mathcal{G} = \mathcal{UN}$ itself, rather than those of the factored noise channel $\mathcal{N}$. Of course, learning $\mathcal{G}$ is equivalent to learning $\mathcal{N}$, although the former approach is especially convenient for Pauli shaping, which deals directly with $\mathcal{G}$, not $\mathcal{N}$ (see, e.g., Eq. (14)).

For mathematical convenience, we will express all PTMs here in the same ordered basis as in Sec. 2, i.e., we will order the set of 2-qubit Paulis $\mathbb{P}$ as in Eq. (18). It will also be convenient to define the subsets

$$\mathbb{P}_C = \{II, ZZ, XX, YY, IZ, ZI, YX, XY\} \quad \mathbb{P}_A = \{XI, YZ, XZ, YI, IX, ZY, ZX, IY\} \quad (43)$$

of $\mathbb{P} = \mathbb{P}_C \cup \mathbb{P}_A$, comprising the eight Paulis that commute and anti-commute with ZZ respectively. The PTM $\mathbf{U}$ of the ideal gate, given in Eqs. (19) and (20), is then block-diagonal with 2×2 blocks. By contrast, the PTM $\mathbf{G}$ of the noisy gate has no such special structure in general; rather, almost all its elements could be nonzero due to noise. Learning $\mathbf{G}$ in full would be a daunting task. To cancel arbitrary noise through Pauli shaping, however, we need only form a characteristic matrix $\mathbf{C}$ such that $\mathbf{U} = \mathbf{C} \odot \mathbf{G}$ (see Eq. (14)). Notice that since $U_{ij} = 0$ for most indices i and j , we can simply set these $C_{ij} = 0$ and avoid learning the corresponding G_{ij} . This argument does not necessarily hold if we instead want to amplify the noise in $\mathbf{G}$ for ZNE, though. Rather, we observe that arbitrary noise can be greatly simplified (though not to the point of being Pauli noise), without affecting the logical

action of the gate, by twirling $\mathcal{G}$ over $\mathbb{P}_c$. The resulting channel

$$\bar{\mathcal{G}}_{\text{C}}(\rho) = \frac{1}{8} \sum_{P_i \in \mathbb{P}_{\text{C}}} P_i \mathcal{G}(P_i \rho P_i) P_i, \quad (44)$$

has a PTM of

$$\bar{G}_c = \begin{pmatrix} G_{0,0} & G_{0,1} & & & & \\ G_{1,0} & G_{1,1} & & & & \\ & & \ddots & & & \\ & & & G_{6,6} & G_{6,7} & \\ & & & G_{7,6} & G_{7,7} & \\ & & & & & G_{8,8} & G_{8,9} \\ & & & & & G_{9,8} & G_{9,9} \\ & & & & & & \ddots \\ & & & & & & & G_{14,14} & G_{14,15} \\ & & & & & & & G_{15,14} & G_{15,15} \\ 0 & & & & & & & & & 0 \end{pmatrix}, \quad (45)$$

which is 2×2 block-diagonal regardless of $\mathbf{G}$, and whose nonzero elements are equal to the corresponding ones of $\mathbf{G}$ (see Appendix C). This “partially” twirled (over $\mathbb{P}_c$ rather than $\mathbb{P}$, hence the subscript c) noisy gate is therefore a simpler starting point for ZNE, and it also includes all the $\mathbf{G}_{ij}$ that are relevant for noise cancellation. Consequently, we will focus here on learning only the elements of $\mathbf{G}_c$, i.e., those on the 2×2 diagonal blocks of $\mathbf{G}$. In principle, they could take (almost) arbitrary values. But in practice, if the noise is weak enough to be amenable to mitigation, then these elements should be close to their ideal values. This leads to four qualitatively different types of PTM elements to be learned, which are color-coded in Eq. (45):

Type 1 ($G_{ii} \approx 1$) These top-left, diagonal elements describe how $\mathcal{G}$ preserves Paulis that commute with ZZ , and should be close to 1. They must be learned in order to cancel the gate noise.

Type 2 ($G_{ii} \approx \cos \theta$) These bottom-right, diagonal elements describe how $\mathcal{G}$ preserves Paulis that anti-commute with ZZ , and should be close to $\cos \theta$. They must be learned in order to cancel the gate noise.

Type 3 ($\mathbf{G}_{ij} \approx \pm \sin \theta$) These bottom-right, off-diagonal elements describe how $\mathcal{G}$ mixes Paulis that anti-commute with ZZ , and should be close to $\pm \sin \theta$. They must be learned in order to cancel the gate noise.

Type 4 ($G_{ij} \approx 0$) These top-left, off-diagonal elements describe how $\mathcal{G}$ mixes Paulis that commute with ZZ , and should be close to 0. They need not necessarily be learned in order to cancel gate noise, since the corresponding elements of $\mathbf{U}$ equal zero. However, they could be necessary for ZNE.

These distinct types stand in contrast with the Clifford case, where all Pauli fidelities approach 1 in the noiseless limit (so there is only one type). We now introduce learning schemes tailored for each different type of PTM element.

3.2.1 Type 1 Elements

Type 1 elements behave much like Pauli fidelities. We can therefore learn them using *modified cycle benchmarking*, in which we Pauli-twirl $\mathcal{G}$ (or equivalently $\bar{\mathcal{G}}_C$), then follow the steps of standard CB for the resulting channel. Note that this is different than standard CB for nontrivial (i.e., $U \neq I$) Clifford gates. There, one inserts random Paulis on either side of $\mathcal{G}$ so as to Pauli-twirl the associated *noise channel* $\mathcal{N}$ without changing the logical effect of the gate, described by $\mathcal{U}$. Here, we instead Pauli-twirl

the noisy gate itself, thus intentionally spoiling its logical effect—in a particular way—and turning it into a Pauli channel with PTM

$$\bar{\mathbf{G}} = \text{diag}(\mathbf{G}_{0,0}, \dots, \mathbf{G}_{7,7}, \mathbf{G}_{8,8}, \dots, \mathbf{G}_{15,15}). \quad (46)$$

This channel is not unitary, even in the noiseless limit. However, because $\bar{\mathbf{G}}$ is diagonal, we can learn the elements $\mathbf{G}_{ii}$, for $1 \leq i \leq 7$, by preparing an eigenstate of P_i , applying this Pauli channel d times for various depths d , estimating $\langle P_i \rangle$ for each using readout twirling, and fitting the results to $d \mapsto Ar^d$, from which r gives a SPAM-robust estimate of $\mathbf{G}_{ii}$. ($\mathbf{G}_{00} = 1$ assuming the noisy gate is CPTP.) Moreover, this scheme is sensitive and it concentrates, much like standard CB; that is, $\mathcal{I}(\mathbf{G}_{ii}) \rightarrow \infty$ and $\Delta \rightarrow 0$ in the noiseless limit. (See Appendices E and F.) It is illustrated numerically in Fig. 8a and summarized in Fig. 9a.

3.2.2 Type 2 and 3 Elements

It may seem from Eq. (46) that we could learn the Type 2 elements ($\mathbf{G}_{ii}$ for $8 \leq i \leq 15$) in the same way. Indeed, this approach would formally give SPAM-robust estimates of said elements—but it would not be sensitive nor would it concentrate, thus making it of limited practical use. These issues are illustrated numerically in Fig. 8b. In more mathematical terms, Eq. (37) implies that

$$\mathcal{I}(\mathbf{G}_{ii}) \lesssim \frac{0.162}{\cos(\theta)^2 \ln[\cos(\theta)]^2} < \infty \quad (47)$$

in the noiseless limit, meaning that each shot would give relatively little information about $\mathbf{G}_{ii}$, so far more shots would be needed than for Type 1 elements (or Pauli fidelities in the Clifford case). Intuitively, the problem is that $\mathbf{G}_{ii}^d \approx \cos(\theta)^d$ generally decays quickly with d , so the measured expectation values $\langle P_i \rangle$ would quickly approach zero regardless of $\mathbf{G}_{ii}$'s exact value, and resolving them to within a reasonable relative error would require many shots. To make matters worse, the expectation values of different random circuits would not concentrate; rather, they would have a variance of

$$\Delta^2 \rightarrow \frac{1}{2} \left[1 + \cos(2\theta)^d - 2 \cos(\theta)^{2d} \right] \quad (48)$$

in the noiseless limit, which quickly asymptotes to $1/2$ with growing depth d . (See Appendix F.) The issue is that the lower-right blocks of the ideal PTM $\mathbf{U}$ (see Eqs. (19) and (20)) are 2-dimensional rotation matrices, so Pauli-twirling the ideal gate implements rotations over $\mathbb{P}_A$ by a uniformly random angle of $\pm\theta$. Repeating such twirled gates therefore produces a random walk with a rapidly growing variance given by Eq. (48). Ultimately, this means that modified CB is impractical for Type 2 elements, since it would require many more shots from many more random circuits (compared to Type 1 elements). These issues highlight the importance of grouping PTM elements into distinct types based on their approximate values—Type 1 and Type 2 elements may appear similarly in Eq. (46), but they can behave very differently since the latter can be much smaller.

Instead, we introduce two other learning schemes which, together, satisfy all of our desiderata. The first of these schemes (called *partial-twirl benchmarking*) yields some information about the Type 3 elements, which we then use, together with the second scheme (called *correlated-twirl benchmarking*), to get the Type 2 elements.

Partial-twirl benchmarking: The main idea of this scheme is to apply $\bar{\mathcal{G}}_C^d$ for various depths d , i.e., to apply the noisy gate $\mathcal{G}$ d times, twirling each one independently over $\mathbb{P}_C$, the set of Paulis that commute with ZZ , rather than over all Paulis. Estimating $\langle P_i \rangle$ for $P_i \in \mathbb{P}_A$ at different depths and fitting the results will then give SPAM-robust estimates of certain PTM elements.

Since the PTM of the partially-twirled gate, $\bar{\mathbf{G}}_C$, is block-diagonal (see Eq. (45)), we can find $\bar{\mathbf{G}}_C^d$ by simply taking the d^{th} power of each 2×2 block. Consider one such block from the bottom-right of $\bar{\mathbf{G}}_C$, which we will denote as B :

$$B := \begin{pmatrix} \mathbf{G}_{ii} & \mathbf{G}_{ij} \\ \mathbf{G}_{ji} & \mathbf{G}_{jj} \end{pmatrix}, \quad (49)$$




Figure 8: A numerical example of how modified CB learns the Type 1 elements of G from Fig. 5a, but fails on the Type 2 elements. *Top*: To illustrate how modified CB concentrates for Type 1 elements, we estimated $\langle ZZ \rangle$ at each depth d using only $N_c = 10$ random circuits, repeated for $N_{s/c} = 10^3$ shots each, and denote the resulting means by $\hat{\mu}'$. The error bars, which show the standard error of $\hat{\mu}'$ estimated through Eq. (39), are small because the random circuits for a given d have similar ZZ expectation values (i.e., Δ is small). The slow decay enables a good fit and ultimately a good estimate of $G_{1,1}$, despite strong SPAM errors, which we model by inserting random error channels with average fidelity $F_{\text{avg}} = 0.9$ at the beginning and end of each circuit. The relative error in learning this, and other Type 1 elements, is shown in Fig. 5c. *Bottom*: Analogous data for a Type 2 element of this G using the same settings (N_c , $N_{s/c}$, and SPAM error strength). The learning error is much worse because the decay is faster, and because the error bars/scatter are larger due to different twirl circuits having very different XZ expectation values (i.e., Δ is large). This data is not included in Fig. 5c; rather, it only motivates the two other learning schemes that follow.




Figure 9: A summary of the non-Clifford learning schemes introduced in this paper. All three schemes work by repeating the noisy gate $\mathcal{G}$ d times for various depths d , twirling each one in a prescribed way, then estimating Pauli expectation values for the resulting quantum state. Fitting these estimated values, $\hat{\mu}$, versus d to a prescribed function then gives SPAM-robust estimates of the various PTM elements of $\mathcal{G}$ (or products thereof). Each row above depicts a different, complementary, learning scheme. The first column (*Twirl sequence*) shows the sets of Paulis over which each $\mathcal{G}$ is twirled. The branching in the bottom row indicates that either path is taken with 50% probability. The middle column (*Expectation values*) illustrates typical results, with orange dots depicting empirical estimates $\hat{\mu}$ of the true expectation values, and solid gray curves depicting the resulting fits. (The dashed curves in the middle plot show the exponential envelope $d \mapsto \pm r^d$, and are meant to guide the eye.) Finally, the right column (*Output*) lists the PTM elements, or products thereof, that can be extracted by fitting the data.

where $i \in \{8, 10, 12, 14\}$ and $j = i + 1$. The form of B^d depends qualitatively on the eigenvalues of B , which are:

$$\lambda_{\pm} = \frac{1}{2} \left[\mathbf{G}_{ii} + \mathbf{G}_{jj} \pm \sqrt{(\mathbf{G}_{ii} - \mathbf{G}_{jj})^2 + 4 \mathbf{G}_{ij} \mathbf{G}_{ji}} \right]. \quad (50)$$

Intuitively, the elements of B^d are functions of $\lambda_{\pm}^d$, so there can be two distinct cases: If $\lambda_{\pm}$ are real, they will produce exponential decays. If $\lambda_{\pm} = r e^{\pm i\omega}$ are complex, because the term in the square root is negative, they will instead produce exponentially-damped oscillations with some frequency ω and decay rate r . We will call these two cases, namely when $\text{Im}(\lambda_{\pm}) = 0$ and $\text{Im}(\lambda_{\pm}) \neq 0$, the strong and weak noise regimes respectively. (They are analogous to over/critically-damped and under-damped classical harmonic oscillators, respectively.) In the strong noise regime, applying $\mathcal{G}_c$ repeatedly to a generic initial state ρ and measuring $\langle P_i \rangle$ for $P_i \in \mathbb{P}_A$ will give expectation values that decay steadily towards their asymptotic values with growing circuit depth d . In the weak noise regime, these expectation values will instead oscillate with d as they gradually decay, much like Rabi oscillations. (In the noiseless limit there is no decay and the oscillations persist as $d \rightarrow \infty$.) The two regimes should therefore typically be easy to distinguish experimentally. This scheme assumes that the gate is in the weak noise regime, i.e., that:

$$(\mathbf{G}_{ii} - \mathbf{G}_{jj})^2 \stackrel{\text{assumed}}{<} -4 \mathbf{G}_{ij} \mathbf{G}_{ji} \quad (51)$$

for all 2×2 blocks of Type 2 and 3 elements. This is our only assumption about these PTM elements, and it amounts to assuming that the gate's logical effect is not overwhelmed by noise. In the noiseless limit, the left and right hand sides of (51) approach 0 and $4 \sin(\theta)^2$ respectively. More broadly, we expect this condition to be easily satisfied on modern quantum processors, provided the chosen θ is not unreasonably small.

Assuming condition (51), we can write B^d in the form

$$B^d = \begin{pmatrix} ar^d \cos(\omega d - \delta) & \cdots \\ \cdots & \cdots \end{pmatrix}, \quad (52)$$

where

$$\begin{aligned} a &= 2\sqrt{\frac{\mathbf{G}_{ij}\mathbf{G}_{ji}}{(\mathbf{G}_{ii} - \mathbf{G}_{jj})^2 + 4\mathbf{G}_{ij}\mathbf{G}_{ji}}} \\ r &= \sqrt{\mathbf{G}_{ii}\mathbf{G}_{jj} - \mathbf{G}_{ij}\mathbf{G}_{ji}} \\ \omega &= \arctan2\left(\sqrt{-(\mathbf{G}_{ii} - \mathbf{G}_{jj})^2 - 4\mathbf{G}_{ij}\mathbf{G}_{ji}}, \mathbf{G}_{ii} + \mathbf{G}_{jj}\right) \\ \delta &= \arctan2\left(\mathbf{G}_{ii} - \mathbf{G}_{jj}, \sqrt{-(\mathbf{G}_{ii} - \mathbf{G}_{jj})^2 - 4\mathbf{G}_{ij}\mathbf{G}_{ji}}\right), \end{aligned} \quad (53)$$

and ellipses denote different matrix elements that are generally nonzero (see Appendix G). We also use the notation $\arctan2(y, x)$ to denote $\arctan(y/x)$ with an appropriate quadrant correction, as in many programming languages. In the noiseless limit⁶ $(a, r, \omega, \delta) \rightarrow (1, 1, \theta, 0)$. More generally, Eqs. (52) and (53) suggest that one could perhaps follow steps akin to CB, but with measured expectation values at different d forming a decaying sinusoid $\propto r^d \cos(\omega d - \delta)$ rather than a pure exponential decay. Fitting the data to this curve would then give a decay rate r , a frequency ω , and a phase δ , from which one could learn about the relevant PTM elements by inverting Eq. (53).

There remains one problem, however. In the other learning schemes discussed so far, the PTM of interest was diagonal, so an expectation value $\langle P_i \rangle$ at depth d only depended on one component $s_i = \text{tr}(\rho P_i)$ of the initial state $\rho = \frac{1}{4} \sum_k s_k P_k$. Here, however, because B^d is not diagonal, $\langle P_i \rangle$ will be a linear combination of s_i and s_j , with weights that vary with d . That is, assuming ideal measurements for the moment (for simplicity):

$$\langle P_i \rangle = \text{tr}[\bar{\mathcal{G}}_c^d(\rho) P_i] = s_i ar^d \cos(\omega d - \delta) + s_j (B^d)_{01}, \quad (54)$$

where $(B^d)_{01}$ is the top-right element of B^d from Eq. (52), which is generally nonzero and depends on d . This means that state preparation errors, which can cause s_i and s_j to deviate from their intended values independently, can impact our estimates of $\langle P_i \rangle$ in more pernicious ways than before, i.e., not just as a constant scale factor that can be absorbed into the amplitude of the fitted curve and ignored. To sidestep this issue, we propose applying P_i to ρ with 50% probability before applying $\bar{\mathcal{G}}_c$. The resulting state

$$\rho' = \frac{1}{2}(\rho + P_i \rho P_i) = \frac{1}{4} \sum_k s'_k P_k \quad (55)$$

has

$$s'_i = \text{tr}(\rho' P_i) = s_i \quad s'_j = \text{tr}(\rho' P_j) = 0, \quad (56)$$

i.e., it has the same P_i component as ρ but no P_j component, because $\{P_i, P_j\} = 0$. We will refer to this step as *state-prep twirling*, in analogy to readout twirling, since it uses randomization to make state preparation errors better behaved. By suppressing the second term in Eq. (54) (effectively replacing s_j with $s'_j = 0$), state-prep twirling ensures that state preparation errors only contribute a constant scale factor in our estimates of $\langle P_i \rangle$, as in CB. This allows us to fit our estimates of $\langle P_i \rangle$ versus d and extract values of r , ω , and δ that are robust to SPAM errors.

Ultimately, then, partial-twirl benchmarking comprises the following steps. For each $i \in \{8, 10, 12, 14\}$ and $j = i + 1$ (indices that label 2-qubit Paulis according to Eq. (18)):

⁶This can be a useful initial guess for curve fitting.

1. Prepare an initial state $\rho = \frac{1}{4} \sum_k s_k P_k$ for which $s_i = \text{tr}(P_i \rho)$ is as large as possible. E.g., attempt to prepare $\rho = |\psi\rangle\langle\psi|$ where $|\psi\rangle$ is a separable +1 eigenstate of P_i , so $s_i = 1$ ideally.
2. (State-prep twirling.) Apply P_i with 50% probability to ρ , independently in each shot, to produce an average state ρ' .
3. Apply $\bar{\mathcal{G}}_C$ d times to ρ' for varying depths d , where $\bar{\mathcal{G}}_C$ denotes the noisy $R_{ZZ}(\theta)$ gate $\mathcal{G}$ twirled over $\mathbb{P}_C$, the eight Paulis that commute with ZZ .
4. Estimate $\langle P_i \rangle$ for the resulting state $\bar{\mathcal{G}}^d(\rho')$ as in Eq. (32) using readout twirling, denoting the result by $\hat{\mu}$.

The expected value of $\hat{\mu}$ for a circuit depth d is then

$$\mu := \mathbb{E}(\hat{\mu}) = s_i m_i a \times r^d \cos(\omega d - \delta), \quad (57)$$

where a , r , ω and δ are given by Eq. (53), and the coefficient m_i describes the measurement errors, as introduced at the start of Sec. 3. (Cf. the equivalent expression for cycle benchmarking in Eq. (35).) Note that SPAM errors only affect the amplitude of this decaying sinusoid. We can therefore estimate its decay rate, frequency and phase in a SPAM-robust way by fitting the tuples $(d, \hat{\mu})$ to $d \mapsto A r^d \cos(\omega d - \delta)$ and extracting r , ω and δ , respectively. We can then solve Eq. (53) for the underlying PTM elements to get:

$$\mathbf{G}_{ij} \mathbf{G}_{ji} = - \left(\frac{r \sin(\omega)}{\cos(\delta)} \right)^2 \quad (58)$$

$$\mathbf{G}_{ii} = r \left[\cos(\omega) + \sin(\omega) \tan(\delta) \right] \quad (59)$$

$$\mathbf{G}_{jj} = r \left[\cos(\omega) - \sin(\omega) \tan(\delta) \right], \quad (60)$$

which do not depend on the amplitude A . The procedure is illustrated numerically in Fig. 10 and summarized in Fig. 9b.

The expectation values from this scheme concentrate as desired, i.e., $\Delta \rightarrow 0$ in the noiseless limit (see Appendix F). Unfortunately, the steps above only give us products $\mathbf{G}_{ij} \mathbf{G}_{ji}$ of Type 3 elements, rather than their isolated values. While there are partial workarounds⁷, we suspect this is a fundamental limitation like the CB “degeneracy” arising in the Clifford case [41]. If so, one could rely on similar approximations here to isolate the Type 3 elements, e.g., assume that $\mathbf{G}_{ij} = -\mathbf{G}_{ji}$ as in the noiseless limit.

There is one remaining issue with this scheme: the measurement results are highly sensitive to the decay rate and oscillation frequency, r and ω respectively, but not to the phase δ (see Appendix E). Intuitively, a small change in r or ω leads to a big change in μ at large depths, as quantified by $\partial\mu/\partial r$ and $\partial\mu/\partial\omega$. In contrast, a small change in δ only produces a small offset in μ regardless of the depth. In other words, the phase is typically harder to fit precisely than the other two parameters. This is a minor issue for Type 3 elements, since $\delta \approx 0$ for weak noise, and Eq. (58) only depends on δ to order $O(\delta^2)$. However, Eqs. (59) and (60) both depend on it more strongly, namely to order $O(\delta)$, so our inability to precisely fit δ can lead to poor estimates of Type 2 elements ($\mathbf{G}_{ii}$ and $\mathbf{G}_{jj}$) using this method. We therefore introduce one final scheme to more accurately estimate these latter elements.

Correlated-twirl benchmarking: Due to the above difficulty in fitting δ , partial-twirl benchmarking should only be used to learn the Type 3 elements—a different scheme can then be used to learn the Type 2 elements. The key insight underpinning this final scheme is that twirling $\mathcal{G}$ over $\mathbb{P}_A$, the 8 Paulis that anti-commute with ZZ , leads to a block-diagonal PTM $\bar{\mathbf{G}}_A$ that resembles $\bar{\mathbf{G}}_C$ in Eq. (45), but with all the off-diagonal components negated⁸. Suppose we apply $\mathcal{G}$ twice, and with equal probability

⁷One can interleave single-qubit Cliffords between the noisy gates, for instance [41].

⁸In other words, $\frac{1}{2} \bar{\mathbf{G}}_C + \frac{1}{2} \bar{\mathbf{G}}_A = \bar{\mathbf{G}}$, since twirling over $\mathbb{P}_C$ or $\mathbb{P}_A$, each with 50% probability, amounts to a full Pauli-twirl.




Figure 10: A numerical example of how partial-twirl benchmarking learns the Type 3 elements of $\mathbf{G}$ from Fig. 5a. We estimated $\langle XZ \rangle$ at each depth d using the same settings as in Fig. 8 ($N_c = 10$, $N_{s/c} = 10^3$, and state preparation and measurement errors with $F_{\text{avg}} = 0.9$ each), and denote the resulting means by $\hat{\mu}'$. By design, the decay is slow and the error bars, which show the standard error of $\hat{\mu}'$ estimated through Eq. (39), are small for the same reason as in Fig. 8a (small Δ). The resulting fit therefore gives a good estimate of $\mathbf{G}_{10,11}\mathbf{G}_{11,10}$. To fully isolate these and other Type 3 elements, we assumed $\mathbf{G}_{11,10} = -\mathbf{G}_{10,11} > 0$ as in the noiseless limit, leading to the sub-1% errors shown in Fig. 5c. Notice from the oscillations that this example is clearly in the required “weak noise regime,” defined by Eq. (51), despite $\mathbf{G}$ describing much stronger noise than is in most experiments.

we either twirl the first instance over $\mathbb{P}_c$ then the second over $\mathbb{P}_A$, or twirl the first over $\mathbb{P}_A$ then the second over $\mathbb{P}_c$. We refer to this as *correlated twirling*, since the second gate is twirled in a manner that depends on how the first gate was twirled. It results in a Pauli channel, but not the same one as if we had simply Pauli-twirled $\mathcal{G}$ directly. As in the previous section, all PTMs in question are block-diagonal, so it suffices to consider a generic 2×2 PTM block. Specifically, the overall PTM that describes correlated twirling has blocks:

$$\underbrace{\frac{1}{2} \begin{pmatrix} \mathbf{G}_{ii} & \mathbf{G}_{ij} \\ \mathbf{G}_{ji} & \mathbf{G}_{jj} \end{pmatrix}}_{\text{block from } \bar{\mathbf{G}}_c} \underbrace{\begin{pmatrix} \mathbf{G}_{ii} & -\mathbf{G}_{ij} \\ -\mathbf{G}_{ji} & \mathbf{G}_{jj} \end{pmatrix}}_{\text{block from } \bar{\mathbf{G}}_A} + \underbrace{\frac{1}{2} \begin{pmatrix} \mathbf{G}_{ii} & -\mathbf{G}_{ij} \\ -\mathbf{G}_{ji} & \mathbf{G}_{jj} \end{pmatrix}}_{\text{block from } \bar{\mathbf{G}}_A} \underbrace{\begin{pmatrix} \mathbf{G}_{ii} & \mathbf{G}_{ij} \\ \mathbf{G}_{ji} & \mathbf{G}_{jj} \end{pmatrix}}_{\text{block from } \bar{\mathbf{G}}_c} = \begin{pmatrix} \mathbf{G}_{ii}^2 - \mathbf{G}_{ij}\mathbf{G}_{ji} & 0 \\ 0 & \mathbf{G}_{jj}^2 - \mathbf{G}_{ij}\mathbf{G}_{ji} \end{pmatrix}. \quad (61)$$

The resulting PTM is therefore diagonal, but with elements that depend non-trivially on the elements of $\mathbf{G}$. (We use the colors for Type 2 and 3 elements in Eq. (61), but it applies also to Type 1 and Type 4 elements.) Therefore, by repeating this sequence for varying depths, as in cycle benchmarking, we can learn $\mathbf{G}_{ii}^2 - \mathbf{G}_{ij}\mathbf{G}_{ji}$ and $\mathbf{G}_{jj}^2 - \mathbf{G}_{ij}\mathbf{G}_{ji}$ in a SPAM-robust way. And since we have already learned $\mathbf{G}_{ij}\mathbf{G}_{ji}$ through partial-twirl benchmarking, we can therefore isolate the Type 2 elements $\mathbf{G}_{ii}$ and $\mathbf{G}_{jj}$.

We refer to this scheme as *correlated-twirl benchmarking*. It uses the same steps as CB, but with different—and to the best of our knowledge, unusual—twirling. That is, for each $i \in \{8, 10, 12, 14\}$ and $j = i + 1$ (indices that label 2-qubit Paulis according to Eq. (18)):

1. Prepare an initial state $\rho = \frac{1}{4} \sum_k s_k P_k$ for which $s_i = \text{tr}(P_i \rho)$ is as large as possible. E.g., attempt to prepare $\rho = |\psi\rangle\langle\psi|$ where $|\psi\rangle$ is a separable +1 eigenstate of P_i , so $s_i = 1$ ideally.
2. Apply $\bar{\mathcal{G}}_c$ then $\bar{\mathcal{G}}_A$, or $\bar{\mathcal{G}}_A$ then $\bar{\mathcal{G}}_c$, each with 50% probability, where $\bar{\mathcal{G}}_c$ and $\bar{\mathcal{G}}_A$ denote the noisy gate $\mathcal{G}$ twirled over $\mathbb{P}_c$ or $\mathbb{P}_A$ respectively. Repeat this process independently $d/2$ times for varying (even) depths d , as shown in Fig. 9c.
3. Estimate $\langle P_i \rangle$ for the resulting state as in Eq. (32) using readout twirling, denoting the result $\hat{\mu}$.




Figure 11: A numerical example of how correlated-twirl benchmarking learns the Type 2 elements of G from Fig. 5a. We estimated $\langle XZ \rangle$ at each depth d using the same settings as in Figs. 8 and 10 ($N_c = 10$, $N_{s/c} = 10^3$, and state preparation and measurement errors with $F_{\text{avg}} = 0.9$ each), and denote the resulting means by $\hat{\mu}'$. This data contrasts sharply with Fig. 8b by design: the decay is slow and the error bars, which show the standard error of $\hat{\mu}'$ estimated through Eq. (39), are small for the same reason as in Fig. 8a (small Δ). The resulting fit therefore gives a good estimate of $G_{10,10}^2 - G_{10,11}G_{11,10}$, and in turn of $G_{10,10}$. The relative error in learning this, and other Type 2 elements, is shown in Fig. 5c.

The resulting expectation values decay exponentially with depth, with no oscillations, as in CB. Concretely, the expected value of $\hat{\mu}$ is

$$\mu := \mathbb{E}(\hat{\mu}) = s_i m_i \times (G_{ii}^2 - G_{ij}G_{ji})^{d/2}, \quad (62)$$

where the coefficients s_i and m_i depend on state preparation and measurement errors respectively, as in CB, but not on the noisy gate in question. So by fitting the tuples $(d, \hat{\mu})$ to $d \mapsto Ar^{d/2}$, the resulting r gives a SPAM-robust estimate of $G_{ii}^2 - G_{ij}G_{ji}$. Like CB, this scheme is sensitive and it concentrates (see Appendices E and F). Finally, by adding the learned value of $G_{ij}G_{ji}$ from partial-twirl benchmarking, we obtain an estimate of the Type 2 element G_{ii} . These same steps can be repeated with $i \leftrightarrow j$ to learn G_{jj} . The procedure is illustrated numerically in Fig. 11 and summarized in Fig. 9c.

While the three learning schemes we have introduced may seem more complicated than CB, they actually have similar experimental requirements. Not only do they involve the same kinds of circuits (just drawn from different distributions), but they require a similar number of distinct experiments. In particular, a 2-qubit Pauli channel has 15 non-trivial Pauli fidelities⁹. One might therefore expect that it takes 15 distinct experiments to learn these with CB, where each “experiment” consists of estimating some Pauli expectation value $\langle P_i \rangle$ for various depths d . However, it is possible to recycle data and use the same measurement outcomes to estimate $\langle ZX \rangle$ and $\langle IX \rangle$ simultaneously, for instance, since $[I, Z] = 0$. In fact, 6 different expectation values (for the weight-1 Paulis, namely $IX, IY, \dots, ZI$) can be found for free in this way, reducing the required number of distinct experiments to 9. The same trick applies to all the schemes we have introduced, which therefore require only 11 distinct experiments in total. Specifically, the Type 1 elements require 5 distinct experiments (modified CB), the Type 3 elements require just 2 experiments (partial-twirl benchmarking), and the Type 2 elements require 4 experiments (correlated-twirl benchmarking).

⁹There are 16 diagonal elements, but the top-left PTM element of a CPTP map always equals 1.

3.2.3 Type 4 Elements

The only potentially nonzero PTM elements in Eq. (45) not learned so far are those of Type 4. Unfortunately, we do not yet know of a good way to learn these. The issue is that, like the Type 3 elements, all SPAM-robust learning schemes we have found¹⁰ let us measure products $G_{ij}G_{ji}$, rather than isolated elements G_{ij} and G_{ji} . We suspect this to be a fundamental limitation, analogous to that for Clifford gates [41]. And since these Type 4 elements should be close to zero for low-noise gates, we expect their products to be extremely small in practice, and therefore very challenging to resolve. For now, we propose to simply bound them using our knowledge of the other PTM elements, by demanding that the learned channel be CPTP [17].

Of course, it is possible to cancel arbitrary noise through Pauli shaping without learning these Type 4 elements, as mentioned earlier, by simply setting the corresponding characteristic matrix elements C_{ij} to be zero. This is what we did in Examples 2 and 3 of Sec. 2.3.2 after all. However, this is what led to a pathologically large values γ . If we could learn these Type 4 elements, we could perhaps amplify them correctly and do ZNE, or simply ignore them (and allow $C_{ij} \neq 0$) with confidence, and thereby get a much lower γ . As it stands, however, these small—but potentially nonzero—elements seem both hard to learn and expensive to mitigate (specifically, to cancel). Such elements are unique to non-Clifford gates like $R_{ZZ}(\theta)$, since any PTM element of a noisy Clifford that should be zero (ideally) can be made zero through twirling. The same is not true of non-Clifford gates, whose noise generally cannot be twirled over the full Pauli group without spoiling the effect of the gate. It is interesting to note that the average fidelity between the ideal gate $\mathcal{U}$ and its noisy realization $\mathcal{G}$ depends on the Type 1–3 elements, but not on the Type 4 elements.

4 Discussion & Outlook

This work was motivated by the prospect of using error mitigation to simulate quantum dynamics on pre-fault-tolerant quantum computers in a semi-analog way. More specifically, the Trotter/Floquet circuits arising in quantum simulation can be realized using weakly-entangling gates (e.g., $R_{ZZ}(\theta)$ for small angles θ), which can be performed faster and with higher fidelity in some experiments than can entangling Clifford gates (e.g., CNOTs). However, the current prevailing machinery of error mitigation—including both the noise learning and noise cancellation or amplification components—relies critically on the gate(s) of interest being Clifford, and is therefore incompatible with such non-Clifford, weakly-entangling gates. We have shown how to generalize both components of error mitigation to non-Clifford gates. Specifically, we introduced the framework of *Pauli shaping*, which transforms any quantum channel into almost any other channel in expectation, at the cost of a sampling overhead, and which reduces to earlier methods when applied to Clifford gates. As this technique relies on detailed knowledge of the former channel, we also introduced three schemes to characterize noisy $R_{ZZ}(\theta)$ gates, which are natural in many experiments, in a SPAM-robust way. In doing so, however, we uncovered several new challenges that do not arise with Clifford gates.

Clifford gates have a simple structure by definition, in that they map every Pauli operator (in a density matrix) to some other Pauli, rather than to a mixture thereof. This makes it possible to twirl the noise in an imperfect Clifford gate over the full set of Paulis using only single-qubit gates, leading to relatively simple noise in effect. In this sense, non-Clifford gates have more complicated structures that admit less twirling, so mitigating them presents a trade-off: the associated noise is potentially weaker, but more complex. While this upside is substantial, the cost can be serious, leading to unwanted effects in error mitigation that have no analogue in Clifford gates. The main examples we encountered involve PTM elements of noisy $R_{ZZ}(\theta)$ gates that describe how a Pauli P_i gets mapped to a different one P_j , where $[P_i, ZZ] = 0$ and $P_j \propto P_i ZZ$ (e.g., $P_i = ZI$ and $P_j = IZ$). We called these Type 4 elements in Sec. 3. They equal zero in noiseless gates, but can be slightly nonzero in practice due to experimental imperfections. Because they belong to a non-Clifford gate, we know of no good way to eliminate them

¹⁰For instance, modified CB followed by correlated-twirl benchmarking.

through twirling. And while it is possible to do so through Pauli shaping, the resulting sampling overhead is impractically large, no matter how weak the noise. This would not be an issue if these Type 4 PTM elements happened to be negligibly small and could simply be ignored. However, they also seem particularly difficult to measure in a SPAM-robust way, making it hard to know precisely how small they are in experiments. There are no such troublesome PTM elements in Clifford gates, because these gates' simpler structure enables more twirling, which leads to simpler noise.

We do not yet know how common such small-but-not-easily-eliminated PTM elements are for other non-Clifford gates, although they appear to be quite generic. However, we can imagine several tentative ways to sidestep them. One way could be at the device physics level, by designing gates whose errors overwhelmingly affect the larger PTM elements. Another would be to synthesize a noiseless non-Clifford gate by Pauli-shaping a probabilistic mixture of Cliffords in which entangling gates rarely arise. For instance, rather than mitigate a physical $R_{ZZ}(\theta)$ gate, one could instead perform I , ZZ , or a noisy $R_{ZZ}(\pi/2)$ gate (which is Clifford, so the Type 4 elements can be twirled away) with appropriate probabilities, then use Pauli shaping to transform the resulting channel into a noiseless $R_{ZZ}(\theta)$ in the spirit of [44, 45, 46, 47]. If θ is small, $R_{ZZ}(\pi/2)$ need only be performed with low probability, so the overall channel would contain little gate noise, and the resulting overhead could be reasonable. Another, more speculative approach, could be to approximately amplify non-Clifford gate noise (for ZNE) without ever learning the troublesome PTM elements. Consider, for example, a noisy $R_{ZZ}(\theta)$ gate twirled over the Paulis that commute with ZZ , whose PTM is therefore block-diagonal with 2×2 blocks as in Eq. (45). Consider one such block

$$B = \begin{pmatrix} \mathbf{G}_{ii} & \mathbf{G}_{ij} \\ \mathbf{G}_{ji} & \mathbf{G}_{jj} \end{pmatrix}, \quad (63)$$

where $\mathbf{G}_{ii}$ and $\mathbf{G}_{jj}$ (called Type 1 elements in Sec. 3) are known, and the Type 4 elements $\mathbf{G}_{ij}$ and $\mathbf{G}_{ji}$ are unknown but small. There are several possible notions of noise amplification in such a gate, some of which involve replacing B with $B^{1+\alpha}$ through Pauli shaping, for different noise levels $1 + \alpha \geq 1$. To first order in $\mathbf{G}_{ij}$ and $\mathbf{G}_{ji}$:

$$B^{1+\alpha} \approx \begin{pmatrix} \mathbf{G}_{ii}^{1+\alpha} & \eta \mathbf{G}_{ij} \\ \eta \mathbf{G}_{ji} & \mathbf{G}_{jj}^{1+\alpha} \end{pmatrix} = B \odot \overbrace{\begin{pmatrix} \mathbf{G}_{ii}^\alpha & \eta \\ \eta & \mathbf{G}_{jj}^\alpha \end{pmatrix}}^C \quad (64)$$

for

$$\eta = \frac{\mathbf{G}_{ii}^{1+\alpha} - \mathbf{G}_{jj}^{1+\alpha}}{\mathbf{G}_{ii} - \mathbf{G}_{jj}}. \quad (65)$$

Since C depends only on $\mathbf{G}_{ii}$ and $\mathbf{G}_{jj}$, one could amplify the noise, up to an approximation error of order $O(\mathbf{G}_{ij}^2) + O(\mathbf{G}_{ij}\mathbf{G}_{ji}) + O(\mathbf{G}_{ji}^2)$, without knowing the exact value of $\mathbf{G}_{ij}$ or $\mathbf{G}_{ji}$. Doing so would entail a sampling overhead, although a potentially much smaller one than is required to cancel the noise (see Example 2 of Appendix C).

An additional practical complication with mitigating non-Clifford gates is that the noise must be learned separately for different angles θ . But since noisy Cliffords on different qubits are already learned separately (because the noise can vary between qubits), and regularly (because of noise drift), the angle dependence does not seem to add much difficulty to any given quantum simulation. However, it makes the first-order Trotter formula preferable to higher-order product formulas [48], which typically involve multiple different angles for any pair of qubits.

Similarly, while our mitigation techniques in Sec. 2 are very general, our noise learning techniques in Sec. 3 are deliberately specific to R_{ZZ} gates (or rather R_{PQ} gates, for Paulis $P, Q \in \{X, Y, Z\}$). Other 2-qubit gates, such as weakly-entangling iSWAP-like gates, can have a different structure which admits less twirling, leading to PTMs with larger nonzero blocks than in Eq. (45). Whether and how our noise learning techniques generalize to such gates is an important open question to be answered case by case.

Finally, noise channels describing multiple 2-qubit gates performed in parallel have exponentially many parameters in general, for both Clifford and non-Clifford gates. Fortunately, the underlying device physics often allows one to approximate said noise using polynomially many parameters, thereby making the learning and mitigation tractable. The simplest approach is to approximate each noisy gate as being independent of (though not necessarily identical to) the others, meaning the techniques from Sections 2 and 3 can be applied to each gate in parallel. In the Clifford case, sparse Pauli-Lindblad models are a generalization that has been empirically successful [9, 10]. They model the twirled noise as arising from an effective Lindbladian $\mathcal{L}_{\text{eff}}$ that comprises only geometrically local terms matching the device connectivity (e.g., with 2-body terms only between coupled qubits), and learn its parameters through cycle benchmarking. We expect our learning techniques can be similarly generalized for parallel R_{ZZ} gates twirled over $\mathbb{P}_C$. The main difference would be in allowing $\mathcal{L}_{\text{eff}}$ to include a slightly broader range of local terms (e.g., ZZ Hamiltonian terms, and certain dissipative cross-terms) because the noise, while simplified from twirling, would remain non-Pauli. This extension of sparse Pauli-Lindblad models to non-Clifford gates seems conceptually simple (though perhaps tedious to work out), and would be a natural direction for subsequent research.

Whether or not non-Clifford error mitigation can outperform the Clifford variety remains to be seen. Ultimately, this comes down to whether reduced noise strength can outweigh increased noise complexity, which in turn, depends on specific techniques to handle this complexity, like those mentioned above. We expect such techniques to be a fruitful area for future research.

Acknowledgements

Acknowledgments. We wish to thank Lev Bishop, Andrew Eddins, Luke Govia, Seth Merkel, Kristan Temme, and Ewout van den Berg for helpful discussions.

References

- [1] Zhenyu Cai, Ryan Babbush, Simon C. Benjamin, Suguru Endo, William J. Huggins, Ying Li, Jarrod R. McClean, and Thomas E. O’Brien. “Quantum error mitigation”. *Rev. Mod. Phys.* **95**, 045005 (2023).
- [2] Kristan Temme, Sergey Bravyi, and Jay M. Gambetta. “Error mitigation for short-depth quantum circuits”. *Phys. Rev. Lett.* **119**, 180509 (2017).
- [3] Ying Li and Simon C. Benjamin. “Efficient variational quantum simulator incorporating active error minimization”. *Phys. Rev. X* **7**, 021050 (2017).
- [4] Sergey Bravyi, Oliver Dial, Jay M. Gambetta, Darío Gil, and Zaira Nazario. “The future of quantum computing with superconducting qubits”. *Journal of Applied Physics* **132**, 160902 (2022).
- [5] Seth Lloyd. “Universal quantum simulators”. *Science* **273**, 1073–1078 (1996).
- [6] Andrew M. Childs, Dmitri Maslov, Yunseong Nam, Neil J. Ross, and Yuan Su. “Toward the first quantum simulation with quantum speedup”. *Proceedings of the National Academy of Sciences* **115**, 9456–9461 (2018).
- [7] Andrew M. Childs, Yuan Su, Minh C. Tran, Nathan Wiebe, and Shuchen Zhu. “Theory of Trotter error with commutator scaling”. *Phys. Rev. X* **11**, 011020 (2021).
- [8] Laura Clinton, Johannes Bausch, and Toby Cubitt. “Hamiltonian simulation algorithms for near-term quantum hardware”. *Nature Communications* **12**, 4989 (2021).
- [9] Ewout Van Den Berg, Zlatko K Mineev, Abhinav Kandala, and Kristan Temme. “Probabilistic error cancellation with sparse Pauli-Lindblad models on noisy quantum processors”. *Nature Physics* (2023).
- [10] Youngseok Kim, Andrew Eddins, Sajant Anand, Ken Xuan Wei, Ewout Van Den Berg, Sami Rosenblatt, Hasan Nayfeh, Yantao Wu, Michael Zaletel, Kristan Temme, et al. “Evidence for the utility of quantum computing before fault tolerance”. *Nature* **618**, 500–505 (2023).

- [11] Nathan Earnest, Caroline Tornow, and Daniel J. Egger. “Pulse-efficient circuit transpilation for quantum applications on cross-resonance-based hardware”. *Phys. Rev. Res.* **3**, 043088 (2021).
- [12] John P. T. Stenger, Nicholas T. Bronn, Daniel J. Egger, and David Pekker. “Simulating the dynamics of braiding of Majorana zero modes using an IBM quantum computer”. *Phys. Rev. Res.* **3**, 033171 (2021).
- [13] Seth T. Merkel, Jay M. Gambetta, John A. Smolin, Stefano Poletto, Antonio D. Córcoles, Blake R. Johnson, Colm A. Ryan, and Matthias Steffen. “Self-consistent quantum process tomography”. *Phys. Rev. A* **87**, 062119 (2013).
- [14] Robin Blume-Kohout, John King Gamble, Erik Nielsen, Jonathan Mizrahi, Jonathan D. Sterk, and Peter Maunz. “Robust, self-consistent, closed-form tomography of quantum logic gates on a trapped ion qubit” (2013). [arXiv:1310.4492](#).
- [15] Erik Nielsen, John King Gamble, Kenneth Rudinger, Travis Scholten, Kevin Young, and Robin Blume-Kohout. “Gate Set Tomography”. *Quantum* **5**, 557 (2021).
- [16] Suguru Endo, Simon C. Benjamin, and Ying Li. “Practical quantum error mitigation for near-future applications”. *Phys. Rev. X* **8**, 031027 (2018).
- [17] Daniel Greenbaum. “Introduction to quantum gate set tomography” (2015). [arXiv:1509.02921](#).
- [18] Steven T Flammia and Joel J Wallman. “Efficient estimation of Pauli channels”. *ACM Transactions on Quantum Computing* **1**, 1–32 (2020).
- [19] W. Dür, M. Hein, J. I. Cirac, and H.-J. Briegel. “Standard forms of noisy quantum operations via depolarization”. *Phys. Rev. A* **72**, 052326 (2005).
- [20] Christoph Dankert, Richard Cleve, Joseph Emerson, and Etera Livine. “Exact and approximate unitary 2-designs and their application to fidelity estimation”. *Phys. Rev. A* **80**, 012304 (2009).
- [21] E. Knill. “Fault-tolerant postselected quantum computation: Threshold analysis” (2004). [arXiv:quant-ph/0404104](#).
- [22] Joel J. Wallman and Joseph Emerson. “Noise tailoring for scalable quantum computation via randomized compiling”. *Phys. Rev. A* **94**, 052325 (2016).
- [23] Jader P. Santos, Ben Bar, and Raam Uzdin. “Pseudo twirling mitigation of coherent errors in non-clifford gates”. *npj Quantum Information* **10**, 100 (2024).
- [24] G. Casella and R.L. Berger. “Statistical inference”. *Duxbury advanced series in statistics and decision sciences*. Thomson Learning. (2002).
- [25] Robin Harper, Wenjun Yu, and Steven T. Flammia. “Fast estimation of sparse quantum noise”. *PRX Quantum* **2**, 010322 (2021).
- [26] P. Krantz, M. Kjaergaard, F. Yan, T. P. Orlando, S. Gustavsson, and W. D. Oliver. “A quantum engineer’s guide to superconducting qubits”. *Applied Physics Reviews* **6**, 021318 (2019).
- [27] S. A. Moses, C. H. Baldwin, M. S. Allman, R. Ancona, L. Ascarrunz, C. Barnes, J. Bartolotta, B. Bjork, P. Blanchard, M. Bohn, et al. “A race-track trapped-ion quantum processor”. *Phys. Rev. X* **13**, 041052 (2023).
- [28] Simon J Evered, Dolev Bluvstein, Marcin Kalinowski, Sepehr Ebadi, Tom Manovitz, Hengyun Zhou, Sophie H Li, Alexandra A Geim, Tout T Wang, Nishad Maskara, et al. “High-fidelity parallel entangling gates on a neutral atom quantum computer”. *Nature* **622**, 268–272 (2023).
- [29] Michael A Nielsen. “A simple formula for the average gate fidelity of a quantum dynamical operation”. *Physics Letters A* **303**, 249–252 (2002).
- [30] Joel Wallman, Chris Granade, Robin Harper, and Steven T Flammia. “Estimating the coherence of noise”. *New Journal of Physics* **17**, 113020 (2015).
- [31] Sergey Bravyi, Sarah Sheldon, Abhinav Kandala, David C. McKay, and Jay M. Gambetta. “Mitigating measurement errors in multiqubit experiments”. *Phys. Rev. A* **103**, 042605 (2021).
- [32] Ewout van den Berg, Zlatko K. Mineev, and Kristan Temme. “Model-free readout-error mitigation for quantum expectation values”. *Phys. Rev. A* **105**, 032620 (2022).
- [33] Alexander Erhard, Joel J Wallman, Lukas Postler, Michael Meth, Roman Stricker, Esteban A Martinez, Philipp Schindler, Thomas Monz, Joseph Emerson, and Rainer Blatt. “Characterizing large-scale quantum computers via cycle benchmarking”. *Nature Communications* **10**, 5347 (2019).


- [34] Senrui Chen, Sisi Zhou, Alireza Seif, and Liang Jiang. “Quantum advantages for Pauli channel estimation”. *Phys. Rev. A* **105**, 032435 (2022).
- [35] Sitan Chen, Weiyuan Gong, and Qi Ye. “Optimal tradeoffs for estimating Pauli observables”. In 2024 IEEE 65th Annual Symposium on Foundations of Computer Science (FOCS). Pages 1086–1105. (2024).
- [36] Robbie King, David Gosset, Robin Kothari, and Ryan Babbush. “Triply efficient shadow tomography”. *PRX Quantum* **6**, 010336 (2025).
- [37] Joel J Wallman and Steven T Flammia. “Randomized benchmarking with confidence”. *New Journal of Physics* **16**, 103032 (2014).
- [38] Jonas Helsen, Joel J. Wallman, Steven T. Flammia, and Stephanie Wehner. “Multiqubit randomized benchmarking using few samples”. *Phys. Rev. A* **100**, 032304 (2019).
- [39] Andrew Wack, Hanhee Paik, Ali Javadi-Abhari, Petar Jurcevic, Ismael Faro, Jay M. Gambetta, and Blake R. Johnson. “Quality, speed, and scale: three key attributes to measure the performance of near-term quantum computers” (2021). [arXiv:2110.14108](#).
- [40] Steven T. Flammia. “Averaged Circuit Eigenvalue Sampling”. In 17th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2022). Volume 232 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 4:1–4:10. Dagstuhl, Germany (2022). Schloss Dagstuhl – Leibniz-Zentrum für Informatik.
- [41] Senrui Chen, Yunchao Liu, Matthew Otten, Alireza Seif, Bill Fefferman, and Liang Jiang. “The learnability of Pauli noise”. *Nature Communications* **14**, 52 (2023).
- [42] Shelby Kimmel, Marcus P. da Silva, Colm A. Ryan, Blake R. Johnson, and Thomas Ohki. “Robust extraction of tomographic information via randomized benchmarking”. *Phys. Rev. X* **4**, 011050 (2014).
- [43] Jonas Helsen, Francesco Battistel, and Barbara M Terhal. “Spectral quantum tomography”. *npj Quantum Information* **5**, 74 (2019).
- [44] Earl Campbell. “Random compiler for fast Hamiltonian simulation”. *Phys. Rev. Lett.* **123**, 070503 (2019).
- [45] Bálint Koczor, John J. L. Morton, and Simon C. Benjamin. “Probabilistic interpolation of quantum rotation angles”. *Phys. Rev. Lett.* **132**, 130602 (2024).
- [46] Etienne Granet and Henrik Dreyer. “Hamiltonian dynamics on digital quantum computers without discretization error”. *npj Quantum Information* **10**, 82 (2024).
- [47] Bálint Koczor. “Sparse probabilistic synthesis of quantum operations”. *PRX Quantum* **5**, 040352 (2024).
- [48] Masuo Suzuki. “General theory of fractal path integrals with applications to many-body theories and statistical physics”. *Journal of Mathematical Physics* **32**, 400–407 (1991).

A Pauli Shaping Produces the Intended Expectation Values


Suppose we wish to apply a channel $\mathcal{A}$ to an initial state ρ on n qubits, then estimate the expectation value

$$\langle O \rangle_{\mathcal{A}} = \text{tr} [O\mathcal{A}(\rho)] \quad (66)$$

of an observable $O = \sum_{\lambda} \lambda |\lambda\rangle\langle\lambda|$ for the resulting state. (Any gates preceding or following $\mathcal{A}$ in a quantum circuit can be absorbed into the definition of ρ and O respectively.) Suppose, however, that we can only implement a different channel $\mathcal{G}$ in place of $\mathcal{A}$, e.g., due to experimental imperfections. We will show that through Pauli shaping—that is, by inserting random n -qubit Paulis on either side of $\mathcal{G}$ with an appropriate distribution, then scaling the measurement outcomes $\{\lambda\}$ —one can recover $\langle O \rangle_{\mathcal{A}}$ without needing to implement $\mathcal{A}$. The situation is summarized below, with the circuit on the left showing what we would like to implement, and that on the right showing what we will implement instead.



Desired circuit involving $\mathcal{A}$, giving $\langle O \rangle_{\mathcal{A}}$



Pauli shaping circuits involving $\mathcal{G}$, giving $\langle O \rangle_{\text{ps}}$

We begin by defining the PTMs $\mathbf{A}$ and $\mathbf{G}$ of the channels $\mathcal{A}$ and $\mathcal{G}$, respectively, by

$$\mathbf{A}_{ij} = 2^{-n} \text{tr}[P_i \mathcal{A}(P_j)] \quad \text{and} \quad \mathbf{G}_{ij} = 2^{-n} \text{tr}[P_i \mathcal{G}(P_j)] \quad (67)$$

for n -qubit Paulis P_i and P_j . We then take the characteristic matrix $\mathbf{C}$ to be any $4^n \times 4^n$ real matrix satisfying $\mathbf{A} = \mathbf{C} \odot \mathbf{G}$ as in Eq. (14) of the main text, where $\odot$ denotes a Hadamard/element-wise product. Finally, we define the corresponding quasi-probability matrix

$$\mathbf{Q} = 2^{-4n} \mathbf{W} \mathbf{C} \mathbf{W}, \quad (68)$$

where $\mathbf{W}$ is the Walsh matrix defined in Eq. (3) of the main text, as well as the normalizing factor

$$\gamma = \sum_{ij} |\mathbf{Q}_{ij}|. \quad (69)$$

Claim: By inserting Paulis P_j and P_i before and after $\mathcal{G}$, respectively, with probability $|\mathbf{Q}_{ij}|/\gamma$ in each shot independently, and multiplying the measurement outcomes (i.e., the recorded eigenvalues of O) by $\gamma \text{sgn}(\mathbf{Q}_{ij})$, the resulting expectation value $\langle O \rangle_{\text{ps}}$ equals the desired one $\langle O \rangle_{\mathcal{A}}$.

Proof: We begin by finding $\langle O \rangle_{\text{ps}}$. There are two types of randomness involved in Pauli shaping: that of the measurement outcomes $\{\lambda\}$ for a given quantum circuit (i.e., for fixed Paulis P_i and P_j), and that from the choice of circuit. For a fixed circuit in which $\mathcal{G}$ is flanked by Paulis P_j and P_i as shown above, the probability of observing λ when measuring O is

$$\Pr(\lambda | i, j) = \text{tr} \left[|\lambda\rangle\langle\lambda| P_i \mathcal{G}(P_j \rho P_j) P_i \right]. \quad (70)$$

One then records $\lambda \times \gamma \text{sgn}(\mathbf{Q}_{ij})$ in place of λ , as described above. The probability of running this random circuit in the first place is

$$\Pr(i, j) = |\mathbf{Q}_{ij}|/\gamma. \quad (71)$$

Therefore, the overall expectation value from Pauli shaping, combining both sources of randomness and invoking the chain rule (for probabilities), is:

$$\begin{aligned} \langle O \rangle_{\text{ps}} &= \sum_{\lambda ij} \lambda \gamma \text{sgn}(\mathbf{Q}_{ij}) \underbrace{\Pr(\lambda | i, j) \Pr(i, j)}_{\Pr(\lambda, i, j)} = \sum_{ij} \gamma \text{sgn}(\mathbf{Q}_{ij}) \text{tr} [O P_i \mathcal{G}(P_j \rho P_j) P_i] \frac{|\mathbf{Q}_{ij}|}{\gamma} \\ &= \sum_{ij} \mathbf{Q}_{ij} \text{tr} [O P_i \mathcal{G}(P_j \rho P_j) P_i]. \end{aligned} \quad (72)$$

To show that this expression equals $\langle O \rangle_{\mathcal{A}}$, we then decompose both $\rho = 2^{-n} \sum_{\ell} s_{\ell} P_{\ell}$ and $O = \sum_k r_k P_k$ in the Pauli basis using appropriate coefficients $\{s_{\ell}\}$ and $\{r_k\}$ to get

$$\begin{aligned} \langle O \rangle_{\text{ps}} &= 2^{-n} \sum_{ijk\ell} \mathbf{Q}_{ij} r_k s_{\ell} \text{tr} [P_k P_i \mathcal{G}(P_j P_{\ell} P_j) P_i] = 2^{-n} \sum_{k\ell} r_k s_{\ell} \left(\sum_{ij} \mathbf{W}_{ki} \mathbf{Q}_{ij} \mathbf{W}_{j\ell} \right) \text{tr} [P_k \mathcal{G}(P_{\ell})] \\ &= \sum_{k\ell} r_k s_{\ell} \mathbf{C}_{k\ell} \mathbf{G}_{k\ell}, \end{aligned} \quad (73)$$

where we've used Eq. (67) together with the facts that $P_i P_k P_i = \mathbf{W}_{ki} P_k$ and $\mathbf{W} \mathbf{Q} \mathbf{W} = \mathbf{C}$. Finally, since $\mathbf{A}_{k\ell} = \mathbf{C}_{k\ell} \mathbf{G}_{k\ell}$ by definition of $\mathbf{C}$, we can use Eq. (67) again to conclude that:

$$\langle O \rangle_{\text{PS}} = \sum_{k\ell} r_k s_\ell \mathbf{A}_{k\ell} = 2^{-n} \sum_{k\ell} r_k s_\ell \text{tr}[P_k \mathcal{A}(P_\ell)] = \text{tr}[O \mathcal{A}(\rho)] = \langle O \rangle_{\mathcal{A}}. \quad \square \quad (74)$$

Of course, the variance of the recorded outcomes is generally larger with Pauli shaping than it would be if we could implement $\mathcal{A}$ rather than $\mathcal{G}$. Given access to $\mathcal{A}$, the variance would be

$$\Delta O_{\mathcal{A}}^2 = \text{tr}[O^2 \mathcal{A}(\rho)] - \langle O \rangle_{\mathcal{A}}^2, \quad (75)$$

which depends on ρ and $\mathcal{A}$, of course, but can be upper-bounded as

$$\begin{aligned} \Delta O_{\mathcal{A}}^2 &\leq \text{tr}[O^2 \mathcal{A}(\rho)] \leq \sum_i \sigma_i(O^2) \sigma_i[\mathcal{A}(\rho)] \\ &\leq \sigma_1(O^2) \sum_i \sigma_i[\mathcal{A}(\rho)] = \|O^2\| \text{tr}[\mathcal{A}(\rho)] = \|O\|^2 \end{aligned} \quad (76)$$

using the von Neumann trace inequality, where $\sigma_i(M)$ denotes the i^{th} largest singular value of a matrix M and $\|\cdot\|$ denotes the operator/spectral/2 norm. (We use the standard notation ΔO^2 to denote the variance in measurement outcomes for a quantum observable O . It has no relation to the quantity Δ in Eq. (39) of the main text. The letter σ is similarly overloaded: its meaning above has no relation to the function in Eq. (16), which describes the action of a generic Clifford gate on Paulis.) Due to the extra randomness inherent in Pauli shaping, its recorded outcomes are generally less concentrated, having a variance of

$$\begin{aligned} \Delta O_{\text{PS}}^2 &= \sum_{\lambda ij} \left[\lambda \gamma \text{sgn}(\mathbf{Q}_{ij}) \right]^2 \text{Pr}(\lambda | i, j) \text{Pr}(i, j) - \langle O \rangle_{\text{PS}}^2 \\ &= \gamma \sum_{ij} |\mathbf{Q}_{ij}| \text{tr}[O^2 P_i \mathcal{G}(P_j \rho P_j) P_i] - \langle O \rangle_{\mathcal{A}}^2. \end{aligned} \quad (77)$$

This quantity also depends on ρ , $\mathcal{A}$, and $\mathcal{G}$, but it can be similarly upper-bounded as

$$\Delta O_{\text{PS}}^2 \leq \gamma \left(\sum_{ij} |\mathbf{Q}_{ij}| \right) \max_{k\ell} \text{tr}[O^2 P_k \mathcal{G}(P_\ell \rho P_\ell) P_k] \leq \gamma^2 \|O\|^2. \quad (78)$$

Since the standard error of the mean in estimating $\langle O \rangle_{\mathcal{A}}$ using N shots (i.e., N circuit executions) with access to $\mathcal{A}$ is $\Delta O_{\mathcal{A}}/\sqrt{N}$, whereas that from Pauli shaping is $\Delta O_{\text{PS}}/\sqrt{N}$, Pauli shaping incurs roughly a γ^2 sampling overhead for estimating expectation values to within a given statistical error.

B Pauli Shaping Reduces to Clifford PEC/ZNE

Consider a $2^n \times 2^n$ Clifford unitary U . Using the same notation as in the main text, we define an invertible function $\sigma : \{0, 4^n - 1\} \rightarrow \{0, 4^n - 1\}$ such that $P_{\sigma(i)} \propto U^\dagger P_i U$ for every n -qubit Pauli P_i . Moreover, we use the notation $k = i \oplus j$ when $P_k \propto P_i P_j$. We begin with two lemmas about the Walsh matrix elements $\mathbf{W}_{ij}$ defined in Eq. (3) of the main text.

Lemma B1: $\mathbf{W}_{ij} \mathbf{W}_{ik} = \mathbf{W}_{i, j \oplus k}$.

Proof: By definition, $P_{j \oplus k} = z P_j P_k$ for some $z \in \mathbb{C}$. Then:

$$\mathbf{W}_{i, j \oplus k} P_{j \oplus k} = P_i P_{j \oplus k} P_i = z P_i P_j P_k P_i = z \mathbf{W}_{ij} \mathbf{W}_{ik} P_j (P_i)^2 P_k = \mathbf{W}_{ij} \mathbf{W}_{ik} P_{j \oplus k}. \quad \square$$

Lemma B2: $\mathbf{W}_{ij} = \mathbf{W}_{\sigma(i), \sigma(j)}$.

Proof: By definition, $P_{\sigma(i)} = v_i U^\dagger P_i U$ and $P_{\sigma(j)} = v_j U^\dagger P_j U$ for $v_i, v_j \in \{-1, 1\}$. Then:

$$\mathbf{W}_{\sigma(i), \sigma(j)} P_{\sigma(j)} = P_{\sigma(i)} P_{\sigma(j)} P_{\sigma(i)} = v_i^2 v_j (U^\dagger P_i U)(U^\dagger P_j U)(U^\dagger P_i U) = v_j U^\dagger P_i P_j P_i U = \mathbf{W}_{ij} P_{\sigma(j)}. \quad \square$$

Suppose we implement a channel $\mathcal{G}$ in place of an ideal gate $\mathcal{U}(\rho) = U\rho U^\dagger$, which we want to transform (in expectation) into $\mathcal{A}$ from Eq. (7) of the main text through Pauli shaping—without invoking the formalism of Clifford PEC/ZNE. That is, we want to realize an aggregate PTM of

$$\mathbf{A} = \mathbf{U} \bar{\mathbf{N}}^{1+\alpha} \quad (79)$$

as in Eq. (8), for some desired α , where $\mathbf{U}$ is the PTM of $\mathcal{U}$ and $\bar{\mathbf{N}} = \text{diag}(\vec{f})$ is the PTM of the twirled noise channel $\bar{\mathcal{N}}$ (which comes from Pauli-twirling $\mathcal{N} = \mathcal{U}^{-1}\mathcal{G}$) with Pauli fidelities $f_i = \text{tr}[P_i \mathcal{N}(P_i)]/2^n$. The elements of $\mathbf{U}$ are give by

$$U_{ij} = \text{tr}(P_i U P_j U^\dagger)/2^n = v_i \text{tr}(P_{\sigma(i)} P_j)/2^n = v_i \delta_{\sigma(i), j}, \quad (80)$$

using the same notation of $P_{\sigma(i)} = v_i U^\dagger P_i U$ for $v_i = \pm 1$ as in the proof of Lemma B2. The elements of $\mathbf{A}$ are therefore

$$A_{ij} = \sum_k U_{ik} \left(\bar{\mathbf{N}}^{1+\alpha} \right)_{kj} = \sum_k v_i \delta_{\sigma(i), k} f_k^{1+\alpha} \delta_{kj} = v_i f_j^{1+\alpha} \delta_{\sigma(i), j}. \quad (81)$$

It follows that, for generic noise, we need a characteristic matrix $\mathbf{C}$ with elements

$$C_{ij} = f_j^\alpha \delta_{\sigma(i), j} \quad (82)$$

to satisfy Eq. (14) of the main text (that is, to get $\mathbf{A} = \mathbf{C} \odot \mathbf{G}$, where $\mathbf{G}$ is the PTM of $\mathcal{G}$). The associated quasi-probability matrix $\mathbf{Q} = \mathbf{W} \mathbf{C} \mathbf{W}/4^{2n}$ has elements

$$\begin{aligned} Q_{ij} &= 4^{-2n} \sum_{k\ell} \mathbf{W}_{ik} C_{k\ell} \mathbf{W}_{\ell j} = 4^{-2n} \sum_{k\ell} \mathbf{W}_{ik} \left(f_\ell^\alpha \delta_{\sigma(k), \ell} \right) \mathbf{W}_{\ell j} = 4^{-2n} \sum_k \mathbf{W}_{ik} \mathbf{W}_{\sigma(k), j} f_{\sigma(k)}^\alpha \\ &\stackrel{\text{Lemma B2}}{=} 4^{-2n} \sum_k \mathbf{W}_{\sigma(i), \sigma(k)} \mathbf{W}_{\sigma(k), j} f_{\sigma(k)}^\alpha \\ &\stackrel{\text{Lemma B1}}{=} 4^{-2n} \sum_k \mathbf{W}_{\sigma(i) \oplus j, \sigma(k)} f_{\sigma(k)}^\alpha \\ &= 4^{-2n} (\mathbf{W} \vec{f}^\alpha)_{\sigma(i) \oplus j}, \end{aligned} \quad (83)$$

or, written in terms of the vector of quasi-probabilities $\vec{q}$ from Sec. 2.2 of the main text, defined by $\mathbf{W} \vec{q} = \vec{f}^\alpha$:

$$Q_{ij} = 4^{-n} \vec{q}_{\sigma(i) \oplus j}, \quad (84)$$

which is precisely Eq. (16) from the main text. That is, Step 1 of Clifford PEC/ZNE (from Sec. 2.2) can be described as inserting Paulis $P_{\sigma(i)}$ and P_i before and after $\mathcal{G}$ respectively, where $P_i \sim \text{unif}(\mathbb{P})$. Step 2 can be described as inserting an extra Pauli P_k before $P_{\sigma(i)}$ with quasi-probability q_k . Both steps are shown separately in the left circuit below. Combining the two adjacent Paulis (as in the middle circuit), and relabeling $k \oplus \sigma(i)$ as j (as in the right circuit), we arrive at the previous equation. In other words, Pauli shaping is operationally identical to Clifford PEC/ZNE when the target gate U is Clifford, as claimed in the main text.

$$\underbrace{\boxed{P_k}}_{\text{quasi-prob.} = q_k} \underbrace{\boxed{P_{\sigma(i)}} \boxed{\mathcal{G}} \boxed{P_i}}_{\text{prob.} = 4^{-n}} = \underbrace{\boxed{P_{k \oplus \sigma(i)}} \boxed{\mathcal{G}} \boxed{P_i}}_{\text{quasi-prob.} = 4^{-n} q_k} = \underbrace{\boxed{P_j} \boxed{\mathcal{G}} \boxed{P_i}}_{\text{quasi-prob.} = 4^{-n} q_{\sigma(i) \oplus j} \text{ (setting } j = k \oplus \sigma(i))}$$

C Mathematical Details of Pauli Shaping Examples

Before getting into the details of the Pauli shaping examples from Sec. 2.3.2 of the main text, we begin this section with a useful lemma. Consider two characteristic matrices $\mathbf{C}^{(1)}$ and $\mathbf{C}^{(2)}$, which correspond respectively to quasi-probability matrices

$$\mathbf{Q}^{(1)} = 2^{-4n} \mathbf{W} \mathbf{C}^{(1)} \mathbf{W} \quad \mathbf{Q}^{(2)} = 2^{-4n} \mathbf{W} \mathbf{C}^{(2)} \mathbf{W}. \quad (85)$$

Suppose we apply Pauli shaping to a channel $\mathcal{G}$ (with PTM $\mathbf{G}$) according to $\mathbf{C}^{(1)}$ to create an aggregate channel with PTM $\mathbf{C}^{(1)} \odot \mathbf{G}$ at the cost of a sampling overhead γ_1^2 , where $\gamma_1 = \sum_{ij} |\mathbf{Q}_{ij}^{(1)}|$. Then, treating this aggregate channel as a black box, suppose we do a second, outer layer of Pauli shaping according to $\mathbf{C}^{(2)}$, thus producing an aggregate channel with PTM

$$\mathbf{C}^{(2)} \odot (\mathbf{C}^{(1)} \odot \mathbf{G}) = (\mathbf{C}^{(2)} \odot \mathbf{C}^{(1)}) \odot \mathbf{G}, \quad (86)$$

at the cost of a total sampling overhead $(\gamma_1 \gamma_2)^2$, where $\gamma_2 = \sum_{ij} |\mathbf{Q}_{ij}^{(2)}|$. Alternatively, we could realize the same aggregate channel by applying a single layer of Pauli shaping with characteristic matrix $\mathbf{C} = \mathbf{C}^{(2)} \odot \mathbf{C}^{(1)}$, with a corresponding quasi-probability matrix $\mathbf{Q} = \mathbf{W} \mathbf{C} \mathbf{W} / 4^{2n}$, incurring a potentially different sampling overhead of γ^2 where $\gamma = \sum_{ij} |\mathbf{Q}_{ij}|$.

Lemma C1: $\gamma \leq \gamma_1 \gamma_2$.

Proof: We begin by rewriting the relation $\mathbf{C} = \mathbf{C}^{(2)} \odot \mathbf{C}^{(1)}$ in terms of the quasi-probability matrices associated with each characteristic matrix:

$$\mathbf{W} \mathbf{Q} \mathbf{W} = (\mathbf{W} \mathbf{Q}^{(2)} \mathbf{W}) \odot (\mathbf{W} \mathbf{Q}^{(1)} \mathbf{W}), \quad (87)$$

or equivalently:

$$\mathbf{Q} = 4^{-2n} \mathbf{W} \left[(\mathbf{W} \mathbf{Q}^{(2)} \mathbf{W}) \odot (\mathbf{W} \mathbf{Q}^{(1)} \mathbf{W}) \right] \mathbf{W}. \quad (88)$$

The elements of $\mathbf{Q}$ can therefore be expressed as

$$\begin{aligned} \mathbf{Q}_{ij} &= 4^{-2n} \sum_{abcdk\ell} \mathbf{w}_{ia} \left[(\mathbf{w}_{ak} \mathbf{Q}_{k\ell}^{(2)} \mathbf{w}_{\ell b}) (\mathbf{w}_{ac} \mathbf{Q}_{cd}^{(1)} \mathbf{w}_{db}) \right] \mathbf{w}_{bj} \\ &\stackrel{\text{Lemma B1}}{=} 4^{-2n} \sum_{abcdk\ell} (\mathbf{w}_{i \oplus k, a} \mathbf{w}_{ac}) (\mathbf{w}_{j \oplus \ell, b} \mathbf{w}_{bd}) \mathbf{Q}_{k\ell}^{(2)} \mathbf{Q}_{cd}^{(1)} \\ &= \sum_{cdk\ell} \delta_{i \oplus k, c} \delta_{j \oplus \ell, d} \mathbf{Q}_{k\ell}^{(2)} \mathbf{Q}_{cd}^{(1)} \\ &= \sum_{k\ell} \mathbf{Q}_{k\ell}^{(2)} \mathbf{Q}_{i \oplus k, j \oplus \ell}^{(1)}. \end{aligned} \quad (89)$$

In other words, $\mathbf{Q}$ is a convolution of $\mathbf{Q}^{(1)}$ and $\mathbf{Q}^{(2)}$. It follows immediately that

$$\gamma = \sum_{ij} |\mathbf{Q}_{ij}| = \sum_{ij} \left| \sum_{k\ell} \mathbf{Q}_{k\ell}^{(2)} \mathbf{Q}_{i \oplus k, j \oplus \ell}^{(1)} \right| \leq \sum_{k\ell} |\mathbf{Q}_{k\ell}^{(2)}| \sum_{ij} |\mathbf{Q}_{i \oplus k, j \oplus \ell}^{(1)}| = \sum_{k\ell} |\mathbf{Q}_{k\ell}^{(2)}| \gamma_1 = \gamma_1 \gamma_2. \quad \square \quad (90)$$

An important consequence of this lemma is that, in Examples 1 and 2 from Sec. 2.3.2 of the main text, we only need to consider block-diagonal characteristic matrices. More precisely, suppose we want to transform a noisy gate $\mathcal{G}$ into an a channel $\mathcal{A}$ through Pauli shaping, and that we do so using a generic characteristic matrix $\mathbf{C}^{(2)}$ such that $\mathbf{A} = \mathbf{C}^{(2)} \odot \mathbf{G}$. Now consider the operation of twirling $\mathcal{G}$ over the 8 Paulis that commute with ZZ . In the ordered basis (18) from the main text, this twirling is described by

$$\mathbf{Q}^{(1)} = \text{diag} \left(\frac{1}{8}, \frac{1}{8}, \frac{1}{8}, \frac{1}{8}, \frac{1}{8}, \frac{1}{8}, \frac{1}{8}, \frac{1}{8}, 0, 0, 0, 0, 0, 0, 0, 0 \right), \quad (91)$$

whose corresponding characteristic matrix is block-diagonal with identical 2×2 blocks:

$$\mathbf{C}^{(1)} = \mathbf{W}\mathbf{Q}^{(1)}\mathbf{W} = \begin{pmatrix} 1 & 1 & & & & 0 \\ 1 & 1 & & & & \\ & & 1 & 1 & & \\ & & 1 & 1 & & \\ & & & & \ddots & \\ & & & & & 1 & 1 \\ 0 & & & & & 1 & 1 \end{pmatrix}. \quad (92)$$

In both examples, $\mathbf{G}$ is block-diagonal with 2×2 blocks, so $\mathbf{G} = \mathbf{C}^{(1)} \odot \mathbf{G}$. Therefore,

$$\mathbf{A} = \mathbf{C}^{(2)} \odot \mathbf{G} = \mathbf{C}^{(2)} \odot (\mathbf{C}^{(1)} \odot \mathbf{G}) = \mathbf{C} \odot \mathbf{G}, \quad (93)$$

where $\mathbf{C} = \mathbf{C}^{(2)} \odot \mathbf{C}^{(1)}$ is also block-diagonal with 2×2 blocks, and $\gamma \leq \gamma_1\gamma_2$ from Lemma C1. Therefore, in both examples from Sec. 2.3.2 of the main text, the optimal (i.e., lowest possible) sampling overhead is achieved by a block-diagonal characteristic matrix with 2×2 blocks.

Example 1

To correct a coherent over- or under-rotation by ϵ , as in the first example, we consider a block-diagonal characteristic matrix $\mathbf{C}$ given by Eqs. (22)–(24) of the main text. The bottom four blocks are fixed by Eq. (14), whereas we choose the top four to be identical in order to simplify the calculations. The associated $\gamma = \frac{1}{256} \sum_{ij} |(\mathbf{WCW})_{ij}|$ is given by

$$8\gamma = 8|x-y| + |2+x+y+2c+2s| + |2+x+y-2c-2s| + |2-x-y-2c+2s| + |2-x-y+2c-2s|, \quad (94)$$

where x and y are the free coefficients introduced in the main text, and

$$c = \frac{\cos(\theta)}{\cos(\theta + \epsilon)} \quad s = \frac{\sin(\theta)}{\sin(\theta + \epsilon)}. \quad (95)$$

By inspection, Eq. (94) is minimized when $x = y$, which we will consider from now on. Notice also that Eq. (94) is invariant under the transformations $(c, s) \mapsto (s, c)$ and $(c, s) \mapsto (-c, -s)$. We can simplify the equation by removing these symmetries and expressing γ as

$$4\gamma = |x - \underbrace{(-1 - M - m)}_{x_1}| + |x - \underbrace{(-1 + M + m)}_{x_2}| + |x - \underbrace{(1 - M + m)}_{x_3}| + |x - \underbrace{(1 + M - m)}_{x_4}| \quad (96)$$

for

$$M = \max\{|c|, |s|\} \quad m = \text{sgn}(cs) \min\{|c|, |s|\}. \quad (97)$$

Note that

$$\sin^2(\theta) - \sin^2(\theta + \epsilon) = -[\cos^2(\theta) - \cos^2(\theta + \epsilon)] \quad (98)$$

for any θ and ϵ , and that the sign of the left (and right) hand side indicates whether $|s|$ (and $|c|$) is greater or smaller than 1. It follows that $M \geq 1$ and $|m| \leq 1$, which implies $x_1 \leq x_3 \leq x_2 \leq x_4$. It is helpful to visualize γ as a sum of two functions, one comprising the first two terms in Eq. (96) and one comprising the last two, as shown in Fig. 12. It is then clear, geometrically, that the minimum value of γ is achieved at

$$x = \frac{1}{2}(x_2 + x_3) = m, \quad (99)$$

which gives $\gamma = M$, as in Eq. (25) of the main text.




Figure 12: The components of Eq. (96) grouped into two functions which sum to 4γ .

Example 2

The second example of Pauli shaping from Sec. 2.3.2 of the main text can be derived from a Lindblad equation. A general Lindblad equation on n qubits can be written as

$$\frac{d\rho}{dt} = \mathcal{L}(\rho) = -i[H, \rho] + \sum_{j,k=1}^{4^n-1} \Gamma_{jk} \left(P_j \rho P_k - \frac{1}{2} \{P_k P_j, \rho\} \right), \quad (100)$$

where $H = H^\dagger$ is a Hamiltonian, $\{P_j\}$ are n -qubit Paulis (excluding the identity), and $\Gamma = (\Gamma_{jk})_{j,k=1}^{4^n-1} = \Gamma^\dagger$ is a positive semi-definite matrix. When all terms are time-independent, the solution to Eq. (100) for an initial state ρ is

$$\rho' = e^{\mathcal{L}t}(\rho); \quad (101)$$

in other words, it is described by the channel $\mathcal{G} = e^{\mathcal{L}t}$. To simplify the notation, we will absorb t into the definition of $\mathcal{L}$, i.e., set $t = 1$ and consider $\mathcal{G} = e^{\mathcal{L}}$. The second example in the main text comes from picking $H = \frac{\theta}{2} ZZ$, which describes the effect of the ideal gate, together with a nonzero dissipative term with a coefficient matrix

$$\Gamma = -\frac{1}{2} \ln(1 - 2\epsilon) \begin{pmatrix} 0 & & & & & & 0 \\ & \ddots & & & & & \\ & & 0 & & & & \\ & & & 1 & i & & \\ & & & -i & 1 & & \\ & & & & & 1 & i \\ & & & & & -i & 1 \\ & & & & & & 1 & i \\ & & & & & & -i & 1 \\ 0 & & & & & & & & 1 & i \\ & & & & & & & & -i & 1 \end{pmatrix}, \quad (102)$$

expressed in the ordered basis (18) from the main text. Computing the PTM of $\mathcal{L}$ and taking its matrix exponential then gives Eq. (26). We can cancel this noise through Pauli shaping using the

characteristic matrix $\mathbf{C}$ from Eq. (27) of the main text, which corresponds to a γ of

$$\begin{aligned} \gamma = & \frac{1}{16} \left| 1 - x + \frac{1}{1-2\epsilon} + \frac{6}{1-\epsilon} \right| \\ & + \frac{1}{32} \left(\left| 1 + x + \frac{1}{1-2\epsilon} - \frac{16}{\sqrt{1-2\epsilon}} + \frac{6}{1-\epsilon} \right| + \left| 1 + x + \frac{1}{1-2\epsilon} + \frac{16}{\sqrt{1-2\epsilon}} + \frac{6}{1-\epsilon} \right| \right) \\ & + \frac{1}{4} \left(\left| 1 - x + \frac{1}{2\epsilon-1} \right| + \left| 1 + x + \frac{1}{2\epsilon-1} \right| \right) + \frac{3}{16} \left(\left| x - \frac{2\epsilon^2}{1-3\epsilon+2\epsilon^2} \right| + \left| x + \frac{2\epsilon^2}{1-3\epsilon+2\epsilon^2} \right| \right), \end{aligned} \quad (103)$$

where x is a free parameter that we can choose so as to minimize γ . Taking the $\epsilon \rightarrow 0$ limit of this expression gives Eq. (30) of the main text, which is lower-bounded by $3/2$.

It is possible, however, to approximately amplify the noise at much lower cost. Specifically, notice that we can factor the PTM $\mathbf{G}$ of the noisy gate from Eq. (26) as either $\mathbf{G} = \mathbf{U}\mathbf{N}$ or $\mathbf{G} = \mathbf{N}\mathbf{U}$ here, for the same noise channel

$$\mathbf{N} = \begin{pmatrix} \begin{pmatrix} 1 & 0 \\ 2\epsilon & 1-2\epsilon \end{pmatrix} & & & & & & 0 \\ & \begin{pmatrix} 1-\epsilon & -\epsilon \\ -\epsilon & 1-\epsilon \end{pmatrix} & & & & & \\ & & \begin{pmatrix} 1-\epsilon & \epsilon \\ \epsilon & 1-\epsilon \end{pmatrix} & & & & \\ & & & \begin{pmatrix} 1-\epsilon & \epsilon \\ \epsilon & 1-\epsilon \end{pmatrix} & & & \\ & & & & \sqrt{1-2\epsilon} I & & \\ & & & & & \sqrt{1-2\epsilon} I & \\ & & & & & & \sqrt{1-2\epsilon} I \\ 0 & & & & & & & \sqrt{1-2\epsilon} I \end{pmatrix} \quad (104)$$

either way, where I is a 2×2 matrix and $\mathbf{U}$ is the PTM of the ideal $R_{ZZ}(\theta)$ gate from Eq. (19). There are several reasonable notions of what it means to amplify the noise in $\mathbf{G}$ by a factor of $1 + \alpha$, as discussed in the main text, but one is to use Pauli shaping to replace $\mathbf{G}$ with

$$\mathbf{U}\mathbf{N}^{1+\alpha} = \mathbf{N}^{1+\alpha}\mathbf{U} = \begin{pmatrix} \begin{pmatrix} 1 & 0 \\ 2\epsilon & 1-2\epsilon \end{pmatrix}^{1+\alpha} & & & & & & 0 \\ & \begin{pmatrix} 1-\epsilon & -\epsilon \\ -\epsilon & 1-\epsilon \end{pmatrix}^{1+\alpha} & & & & & \\ & & \begin{pmatrix} 1-\epsilon & \epsilon \\ \epsilon & 1-\epsilon \end{pmatrix}^{1+\alpha} & & & & \\ & & & \begin{pmatrix} 1-\epsilon & \epsilon \\ \epsilon & 1-\epsilon \end{pmatrix}^{1+\alpha} & & & \\ & & & & (1-2\epsilon)^{\frac{1+\alpha}{2}} R^{(2)} & & \\ & & & & & \ddots & \\ 0 & & & & & & (1-2\epsilon)^{\frac{1+\alpha}{2}} R^{(2)} \end{pmatrix}. \quad (105)$$

Consider a 2×2 block from the top-left of $\mathbf{G}$:

$$B = \begin{pmatrix} \mathbf{G}_{ii} & \mathbf{G}_{ij} \\ \mathbf{G}_{ji} & \mathbf{G}_{jj} \end{pmatrix}, \quad (106)$$

where $i \in \{0, 2, 4, 6\}$ and $j = i + 1$. In order to Pauli-shape $\mathbf{G}$ into $\mathbf{U}\mathbf{N}^{1+\alpha}$, we need to effectively replace each such block B by $B^{1+\alpha}$. In general, if the off-diagonal elements of B are small (as they are in this example) we can hope to instead approximate $B^{1+\alpha}$ with a simple expression that differs from the exact $B^{1+\alpha}$ by errors of order $O(\mathbf{G}_{ij}^2)$, $O(\mathbf{G}_{ji}^2)$ and $O(\mathbf{G}_{ij}\mathbf{G}_{ji})$. To keep track of error orders, we will rewrite the off-diagonal terms of B as $\mathbf{G}_{ij} = \zeta \mathbf{G}'_{ij}$ and $\mathbf{G}_{ji} = \zeta \mathbf{G}'_{ji}$ for some small bookkeeping parameter ζ that we will use for series expansions. To (quickly) approximate $B^{1+\alpha}$ to first order in $O(\zeta)$, we will use the standard identity

$$\frac{d}{d\zeta} e^{F(\zeta)} = \int_0^1 e^{zF(\zeta)} F'(\zeta) e^{(1-z)F(\zeta)} dz \quad (107)$$

for the derivative of a matrix exponential, where F is a suitably well-behaved matrix-valued function. Using Eq. (107), one can formally show that

$$B = \underbrace{\begin{pmatrix} \mathbf{G}_{ii} & 0 \\ 0 & \mathbf{G}_{jj} \end{pmatrix}}_{B_D} + \zeta \underbrace{\begin{pmatrix} 0 & \mathbf{G}'_{ij} \\ \mathbf{G}'_{ji} & 0 \end{pmatrix}}_{B_A} = e^{F(\zeta)} \quad (108)$$

for some function

$$F(\zeta) = \ln(B_D) + \zeta \left[\frac{\ln(\mathbf{G}_{ii}/\mathbf{G}_{jj})}{\mathbf{G}_{ii} - \mathbf{G}_{jj}} \right] B_A + O(\zeta^2), \quad (109)$$

assuming $\mathbf{G}_{ii}, \mathbf{G}_{jj} > 0$. Then, using Eq. (107) again, we can formally expand $B^{1+\alpha} = e^{(1+\alpha)F(\zeta)}$ in powers of ζ as

$$\begin{aligned} B^{1+\alpha} &= e^{(1+\alpha)F(0)} + \zeta \left. \frac{d}{d\zeta} \right|_{\zeta=0} e^{(1+\alpha)F(\zeta)} + O(\zeta^2) \\ &= B_D^{1+\alpha} + \zeta \left[\frac{(1+\alpha) \ln(\mathbf{G}_{ii}/\mathbf{G}_{jj})}{\mathbf{G}_{ii} - \mathbf{G}_{jj}} \right] \int_0^1 B_D^{z(1+\alpha)} B_A B_D^{(1-z)(1+\alpha)} dz + O(\zeta^2) \\ &= B_D^{1+\alpha} + \zeta \eta B_A + O(\zeta^2) \\ &= \begin{pmatrix} \mathbf{G}_{ii}^{1+\alpha} & \eta \mathbf{G}_{ij} \\ \eta \mathbf{G}_{ji} & \mathbf{G}_{jj}^{1+\alpha} \end{pmatrix} + O(\mathbf{G}_{ij}^2) + O(\mathbf{G}_{ij} \mathbf{G}_{ji}) + O(\mathbf{G}_{ji}^2) \end{aligned} \quad (110)$$

for η in Eq. (65) of the main text. (Note that $\eta \rightarrow (1+\alpha)\mathbf{G}_{ii}^\alpha$ in the limit of $\mathbf{G}_{jj} \rightarrow \mathbf{G}_{ii}$.) Therefore, to first order in the off-diagonal elements, $B^{1+\alpha}$ can be written as the Hadamard product of B with a matrix that only depends on the diagonal terms of B , as in Eq. (64).

This means that by Pauli-shaping $\mathbf{G}$ with a characteristic matrix of

$$\mathbf{C} = \text{diag}(C_1, C_2, C_2, C_2, C_3, C_3, C_3, C_3), \quad (111)$$

where

$$C_1 = \begin{pmatrix} 1 & x \\ 1+\alpha & 1-2\alpha\epsilon \end{pmatrix} \quad C_2 = \begin{pmatrix} 1-\alpha\epsilon & 1+\alpha \\ 1+\alpha & 1-\alpha\epsilon \end{pmatrix} \quad C_3 = (1-\alpha\epsilon) \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix} \quad (112)$$

for any x of our choice, we get an aggregate channel of $\mathbf{A} = \mathbf{C} \odot \mathbf{G} = \mathbf{U} \mathbf{N}^{1+\alpha} + O(\epsilon^2)$. In other words, using a characteristic matrix $\mathbf{C}$ that does not require knowledge of the off-diagonal elements $\mathbf{G}_{ij}$ in the top-left of $\mathbf{G}$ (which are called Type 4 elements in Sec. 3 of the main text), we can effectively implement a channel that is very close to the desired $\mathbf{U} \mathbf{N}^{1+\alpha} = \mathbf{N}^{1+\alpha} \mathbf{U}$ when the gate noise is weak. (That is, we constructed C_1 and C_2 without needing to know the values of the Type 4 elements.) Unlike in Clifford ZNE, this noise amplification still comes at a sampling overhead, albeit a much smaller one than is required to cancel the noise. Specifically, picking $x = 1 + \alpha$ for convenience in C_1 leads to

$$\gamma = \frac{1}{256} \sum_{ij} |(\mathbf{W} \mathbf{C} \mathbf{W})_{ij}| = \frac{(3+7\epsilon)|\alpha|}{4} + \frac{1}{4} |4 + \alpha(1-3\epsilon)|, \quad (113)$$

which reduces to $\gamma = 1 + \alpha(1+\epsilon)$ for $\alpha > 0$ and small ϵ . That is, we can approximately amplify weak noise by a factor of $1 + \alpha$ at the cost of a sampling overhead that is approximately linear in α and that vanishes as $\alpha \rightarrow 0$.

This example is unusual in that the associated noise channel does not depend on the factorization order: if we write $\mathcal{G} = \mathcal{U} \mathcal{N} = \mathcal{N}' \mathcal{U}$ as in the main text, then $\mathcal{N} = \mathcal{N}'$. This is because the noise contributes only an overall damping factor to Paulis that anti-commute with ZZ (i.e., in the bottom-right of $\mathbf{G}$). While this will not generally be the case in experiments, the trick above with $B^{1+\alpha}$ nonetheless applies more broadly. To see how, let $\mathcal{G}$ denote an $R_{ZZ}(\theta)$ gate with arbitrary noise. If

we twirl $\mathcal{G}$ over $\mathbb{P}_c$, the set of Paulis that commute with ZZ , the resulting PTM $\bar{\mathcal{G}}_c$ will be 2×2 block-diagonal, as in Eq. (45). Because the ideal PTM $\mathcal{U}$ also has this block structure, so does the PTM of the resulting (partially twirled) noise channel $\bar{\mathcal{N}}_c = \mathcal{U}^{-1} \bar{\mathcal{G}}_c$. And since the top-left blocks of $\mathcal{U}$ are identity matrices, the top-left blocks of the noise-amplified PTM $\bar{\mathcal{N}}_c^{1+\alpha} \mathcal{U}$ are just those of $\bar{\mathcal{G}}_c$ raised to the power of α , as in the example above. This may make it possible to approximately amplify generic noise on $R_{ZZ}(\theta)$ gates without learning the troublesome Type 4 PTM elements. In general, of course, the factorizations $\bar{\mathcal{G}}_c = \mathcal{U} \bar{\mathcal{N}}_c = \bar{\mathcal{N}}'_c \mathcal{U}$ will lead to different noise channels $\bar{\mathcal{N}}_c \neq \bar{\mathcal{N}}'_c$, and it is not clear at present which one—if either—provides a useful notion of noise amplification.

D Readout Twirling Details

Consider estimating the expectation value $\langle P_i \rangle = \text{tr}(\rho P_i)$ of an n -qubit Pauli P_i with respect to a state ρ , subject to measurement/readout errors. In this section, we show that readout twirling [32] gives an estimate that is proportional to the ground truth $\langle P_i \rangle$, on average, with a proportionality coefficient that depends on the statistics of the readout errors, but not on ρ .

We begin by defining some notation. Let

$$W = \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}^{\otimes n} \quad (114)$$

be a 2^n -dimensional Walsh-Hadamard matrix, then $W^\top = W$ and $W^2 = 2^n I$. Also, for any $j \in \{0, 1\}^n$ let

$$\mathcal{X}_j = X^{j_1} \otimes \dots \otimes X^{j_n} \quad \mathcal{Z}_j = Z^{j_1} \otimes \dots \otimes Z^{j_n} \quad (115)$$

be n -qubit Paulis comprising X 's and I 's, or Z 's and I 's, respectively, as labeled by j . Define the vector $\vec{z}_j \in \mathbb{R}^{2^n}$ such that $\mathcal{Z}_j = \text{diag}(\vec{z}_j)$, then evidently $\vec{z}_j = W|j\rangle$, where $|j\rangle = |j_1\rangle \otimes \dots \otimes |j_n\rangle$ is a standard basis vector. Also, note that $\mathcal{X}_j = 2^{-n} W \mathcal{Z}_j W$. Finally, define the vectors $\vec{p}, \vec{v} \in \mathbb{R}^{2^n}$ of ideal probabilities and ideal expectation values, respectively, with elements

$$p_j = \langle j | \rho | j \rangle \quad v_j = \text{tr}(\mathcal{Z}_j \rho) = \langle \mathcal{Z}_j \rangle. \quad (116)$$

Then $\vec{v} = W\vec{p}$, since

$$\sum_k W_{jk} p_k = \sum_k W_{jk} \text{tr}(|k\rangle\langle k| \rho) = \text{tr} \left[\left(\sum_k W_{jk} |k\rangle\langle k| \right) \rho \right] = \text{tr} \left[\text{diag}(W|j\rangle) \rho \right] = \text{tr}(\mathcal{Z}_j \rho) = v_j. \quad (117)$$

The typical way to estimate $\langle P_i \rangle$ is to apply single-qubit gates to rotate P_i into a matrix that is diagonal in the computational basis, then measure each qubit in this same basis (i.e., the Z eigenbasis). Since we are modeling single-qubit gates as being noiseless, it therefore suffices to consider a diagonal P_i from the start, i.e., we can take $P_i = \mathcal{Z}_j$ for some $j \in \{0, 1\}^n$ in this analysis. Each measurement returns a random n -bit string $k \in \{0, 1\}^n$ distributed according to $\vec{p}$, in the absence of readout error. From each measured k , we can record $\langle k | \mathcal{Z}_j | k \rangle = \pm 1$, the corresponding eigenvalue of $\mathcal{Z}_j$. If we perform N_{tot} such measurements on identical copies of ρ , and estimate $\langle \mathcal{Z}_j \rangle$ by $\hat{\mu}$ defined in Eq. (32) of the main text, then

$$\mathbb{E}(\hat{\mu}) = \frac{N_{\text{tot}} \text{Pr}(+1) - N_{\text{tot}} \text{Pr}(-1)}{N_{\text{tot}}} = \sum_k W_{jk} p_k = (W\vec{p})_j = v_j = \langle \mathcal{Z}_j \rangle \quad (118)$$

as expected. That is, without readout errors, $\hat{\mu}$ is an unbiased estimate of $\langle \mathcal{Z}_j \rangle$.

Readout errors can be described by a stochastic matrix A whose elements $A_{\ell k}$ give the probability that a k measurement outcome gets misreported as an ℓ , for any $k, \ell \in \{0, 1\}^n$ [31]. Under such errors,

the distribution of measured bit-strings becomes $\vec{p}' = A\vec{p}$ (according to the law of total probability), rather than $\vec{p}$. Therefore:

$$\mathbb{E}(\hat{\mu}) = (W\vec{p}')_j = (WA\vec{p})_j = 2^{-n}(WAW\vec{v})_j = 2^{-n} \sum_{\ell} (WAW)_{j\ell} \langle \mathcal{Z}_{\ell} \rangle, \quad (119)$$

which generally does not equal $\langle \mathcal{Z}_j \rangle$, nor does it have any simple functional relation to it.

Readout twirling implements a random $\mathcal{X}_m$ operation both before and after the noisy readout, for m drawn uniformly from $\{0, 1\}^n$ in each shot. In effect, this replaces A with

$$\begin{aligned} A' &= 2^{-n} \sum_m \mathcal{X}_m A \mathcal{X}_m = 2^{-3n} \sum_m (W \mathcal{Z}_m W) A (W \mathcal{Z}_m W) = 2^{-3n} W \left[\left(\sum_m \vec{z}_m \vec{z}_m^{\top} \right) \odot (WAW) \right] W \\ &= 2^{-2n} W [I \odot (WAW)] W, \end{aligned} \quad (120)$$

where $\odot$ denotes a Hadamard/element-wise product. To derive this expression, we used the identity

$$\text{diag}(\vec{x}) M \text{diag}(\vec{x}) = (\vec{x} \vec{x}^{\top}) \odot M \quad (121)$$

for any matrix M and vector $\vec{x}$ of compatible sizes, and the fact that

$$\sum_m \vec{z}_m \vec{z}_m^{\top} = \sum_m W|m\rangle\langle m|W^{\top} = 2^n I, \quad (122)$$

where $\{|m\rangle\}$ are standard basis vectors. This means that with noisy, twirled readout, the distribution of measured bit-strings is $\vec{p}'' = A'\vec{p}$. The average value of $\hat{\mu}$ therefore becomes

$$\mathbb{E}(\hat{\mu}) = (W\vec{p}'')_j = 2^{-n}(WA'W\vec{v})_j = 2^{-n} \left([I \odot (WAW)] \vec{v} \right)_j = m_j \langle \mathcal{Z}_j \rangle, \quad (123)$$

where $m_j = 2^{-n}(WAW)_{jj}$ is independent of the measured state ρ , as claimed in the main text.

E Sensitivity of Learning Schemes

Suppose we want to estimate the expectation value of a Pauli observable P_m with respect to a state ρ by repeating many identical measurements in the eigenbasis of P_m on copies of ρ . Let the random variable $Y = \pm 1$ denote the outcome of one such measurement, where the value of Y indicates the observed eigenvalue. Following the notation from the main text, we will denote the expectation value of Y by $\mu = \mathbb{E}(Y)$. In the absence of readout error, μ is equal to the ideal Pauli expectation value of $\langle P_m \rangle = \text{tr}(P_m \rho)$, but with readout error, μ need not equal $\langle P_m \rangle$ in general, as explained in Appendix D.

The probability mass function of Y is

$$p(y) = \Pr(Y = y) = \begin{cases} (1 + \mu)/2, & y = 1 \\ (1 - \mu)/2, & y = -1. \end{cases} \quad (124)$$

Suppose the state ρ is produced by a noise learning scheme, and that μ depends on some parameter ϕ that we wish to learn. The corresponding (classical) Fisher information is

$$\mathcal{I}(\phi) = \mathbb{E} \left[\left(\frac{\partial}{\partial \phi} \ln p(Y) \right)^2 \right] = \sum_{y \in \{1, -1\}} p(y) \left(\frac{y}{2p(y)} \frac{\partial \mu}{\partial \phi} \right)^2 = \frac{1}{1 - \mu^2} \left(\frac{\partial \mu}{\partial \phi} \right)^2, \quad (125)$$

as in Eq. (36) of the main text. Now consider the case where μ decays exponentially with circuit depth d , as in cycle benchmarking (CB), modified CB, and correlated-twirl benchmarking (see Eqs. (35) and (62)). That is, suppose

$$\mu = Ar^d, \quad (126)$$




Figure 13: The function $g(x)$ from Eq. (128).

where we wish to learn the decay rate $0 < r \leq 1$ (so we will write r instead of the general placeholder ϕ in the equations below), and where the amplitude A depends on SPAM errors. The corresponding Fisher information is

$$\mathcal{I}(r) = \frac{\mu^2}{1 - \mu^2} \frac{d^2}{r^2} = \frac{A^2 r^{2d}}{1 - A^2 r^{2d}} \frac{d^2}{r^2}, \quad (127)$$

as in Eq. (36). Since we aim to characterize the Fisher information in the noiseless limit, we begin by taking $A \rightarrow 1$ in the expression above, which corresponds to the limit of ideal state preparation and measurement. The result still depends on the circuit depth d . To find the maximum Fisher information, we will optimize over d , treating it as a continuous parameter for the moment and demanding that $\frac{\partial}{\partial d} \mathcal{I}(r)|_{A=1} = 0$. We know of no closed-form expression for this optimum, so it is convenient to parameterize the Fisher information in terms of $x = 2d \ln(1/r)$ instead:

$$\mathcal{I}(r)|_{A=1} = \frac{1}{r^{-2d} - 1} \frac{d^2}{r^2} = \frac{1}{r^2 \ln(r)^2} \underbrace{\frac{x^2}{4(e^x - 1)}}_{g(x)}, \quad (128)$$

as in Eq. (37) of the main text. This way all dependence on d is captured by $g(x)$, plotted in Fig. 13, which can be shown numerically to achieve a maximum value of $g(x_*) \approx 0.162$ at $x_* \approx 1.59$. This means the maximum Fisher information is approximately

$$\mathcal{I}(r)|_{A=1, d=d_*} \approx \frac{0.162}{r^2 \ln(r)^2} \quad (129)$$

at $d_* \approx \frac{1.59}{2 \ln(1/r)}$, as in Eq. (37). Of course, d_* is not generally an integer, so the true maximum Fisher information can be slightly lower due to rounding. (The impact of this rounding on $\mathcal{I}(r)$ is negligible for $r \approx 1$ since $\frac{\partial x}{\partial d} \approx 0$ in this regime, but it can be substantial when r is small. Eq. (129) should therefore not be used in the latter regime.) If $r \rightarrow 1$ in the noiseless limit then $\mathcal{I}(r)|_{A=1, d=d_*} \rightarrow \infty$, meaning that the measurement outcomes can be informative about the value of r , so relatively few shots are needed. This is the case in Clifford CB, modified CB (for Type 1 elements) and correlated-twirl benchmarking (since $\mathbf{G}_{ii}^2 - \mathbf{G}_{ij} \mathbf{G}_{ji} \rightarrow \sin(\theta)^2 + \cos(\theta)^2 = 1$ in the noiseless limit). However, if we were to use modified CB to learn the Type 2 elements, we would have $r \rightarrow \cos(\theta)$ in the noiseless limit, leading to a much lower Fisher information when $\theta \approx 0$, and therefore requiring substantially more shots. As a result, we recommend using partial-twirl benchmarking and correlated-twirl benchmarking to learn the Type 2 elements instead, in general.

Finally, consider the case where

$$\mu = A r^d \cos(\omega d - \delta), \quad (130)$$

as in partial-twirl benchmarking (see Eq. (57) of the main text). Before proceeding, we note that for any $x \in (0, 1)$ and $\varphi \in \mathbb{R}$,

$$\frac{\cos(\varphi)^2}{1 - x \cos(\varphi)^2} \leq \frac{1}{1 - x} \quad \frac{\sin(\varphi)^2}{1 - x \cos(\varphi)^2} \leq 1. \quad (131)$$

The Fisher information for the decay rate r from Eq. (130) is

$$\mathcal{I}(r) = \left(\frac{Ar^d d}{r} \right)^2 \frac{\cos(\omega d - \delta)^2}{1 - A^2 r^{2d} \cos(\omega d - \delta)^2}, \quad (132)$$

which follows from plugging Eq. (130) into Eq. (125). It is difficult to maximize this expression over d in general, since it can behave differently depending on how the oscillations line up with the exponential decay. (E.g., the optimal depth d_* used in Eq. (129) could happen to give $\cos(\omega d_* - \delta) = 0$.) In the noiseless limit, however, the exponential decay will be slow compared to the oscillations, which motivates the upper bound

$$\mathcal{I}(r) \leq \left(\frac{Ar^d d}{r} \right)^2 \max_{\varphi} \frac{\cos(\varphi)^2}{1 - A^2 r^{2d} \cos(\varphi)^2}. \quad (133)$$

We can bound this expression further using Eq. (131) provided $Ar^d < 1$, which is guaranteed whenever

$$d > d_0 := -\frac{\ln(A)}{\ln(r)}. \quad (134)$$

Consider a 2×2 PTM block B (from the diagonal of $\mathbf{G}_c$) which deviates from the ideal expression in Eq. (20) of the main text by an arbitrary perturbation:

$$B = \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix} + \varepsilon \begin{pmatrix} x_{00} & x_{01} \\ x_{10} & x_{11} \end{pmatrix}, \quad (135)$$

for some small $\varepsilon > 0$, then one can show that $d_0 = O(\varepsilon)$ through Eq. (53). In other words, we can use Eq. (131) for any non-trivial depth $d > d_0 \rightarrow 0$ in the noiseless limit to find

$$\mathcal{I}(r) \leq \left(\frac{Ar^d d}{r} \right)^2 \frac{1}{1 - A^2 r^{2d}}, \quad (136)$$

which is identical to Eq. (127), and therefore also diverges at d_* as $r \rightarrow 1$. Similarly, now using ω in place of the generic parameter ϕ :

$$\mathcal{I}(\omega) = (Ar^d d)^2 \frac{\sin(\omega d - \delta)^2}{1 - A^2 r^{2d} \cos(\omega d - \delta)^2} \leq (Ar^d d)^2 \leq \frac{A^2}{e^2 \ln(r)^2} \quad (137)$$

in the noiseless limit, which also diverges as $r \rightarrow 1$. In other words, the measurement outcomes from partial-twirl benchmarking are informative about both r and ω . The phase δ , however, is harder to learn—while $\partial\mu/\partial r$ and $\partial\mu/\partial\omega$ both pick up a factor of d , $\partial\mu/\partial\delta$ does not. More precisely, the second inequality in (131) gives

$$\mathcal{I}(\delta) = (Ar^d)^2 \frac{\sin(\omega d - \delta)^2}{1 - A^2 r^{2d} \cos(\omega d - \delta)^2} \leq (Ar^d)^2 \leq A^2 \quad (138)$$

in the noiseless limit, which is much smaller than the maximum values of both $\mathcal{I}(r)$ and $\mathcal{I}(\omega)$.

F Concentration in Learning Schemes

Suppose we would like to execute random quantum circuits from some particular family (e.g., defined by one of the learning or mitigation schemes discussed in the main text) then estimate some Pauli expectation value $\langle P_m \rangle$. Ideally, we would pick a new random circuit for every shot and record ± 1 based on the observed eigenvalue of P_m , repeating this process N_{tot} times and averaging the results. (We denote this average as $\hat{\mu}$ in the main text.) However, it is convenient in many experiments to

instead pick a small number N_c of random circuits and to run each one $N_{s/c}$ times, leading to a total number of shots $N_{\text{tot}} = N_c N_{s/c}$, as described in the main text.

Consider a random quantum channel $\mathcal{T}$ describing a random (potentially noisy) circuit which is chosen from some prescribed family of circuits with probability $\Pr(\mathcal{T})$. Let the random variable $Y(\mathcal{T}) = \pm 1$ denote a measurement outcome (more precisely, the observed eigenvalue of P_m) from running $\mathcal{T}$. Define the conditional expectation

$$\mu(\mathcal{T}) = \mathbb{E}[Y(\mathcal{T})|\mathcal{T}], \quad (139)$$

which is the expected value of the measurement outcomes from a given circuit $\mathcal{T}$ (i.e., the average outcome in the limit of infinitely many executions of $\mathcal{T}$). Since $\mathcal{T}$ is random, $\mu(\mathcal{T})$ is also a random variable, with some mean

$$\mu = \mathbb{E}[\mu(\mathcal{T})] \quad (140)$$

and some variance Δ^2 , as illustrated in Fig. 7 of the main text. We are interested in the case where there are N_c different random circuits, all drawn independently from the same distribution, and each one is executed $N_{s/c}$ times. Let the random variable $\mathcal{T}_j$ denote the j^{th} random circuit, and the random variable $Y_i(\mathcal{T}_j) = \pm 1$ denote the measurement outcome from the i^{th} time circuit $\mathcal{T}_j$ is run, for $1 \leq j \leq N_c$ and $1 \leq i \leq N_{s/c}$. As described in the main text, we will use the random variable

$$\hat{\mu}' = \frac{1}{N_{\text{tot}}} \sum_{i=1}^{N_{s/c}} \sum_{j=1}^{N_c} Y_i(\mathcal{T}_j) \quad (141)$$

as an estimator for μ . It is unbiased, since

$$\mathbb{E}(\hat{\mu}') = \frac{1}{N_{\text{tot}}} \sum_{i=1}^{N_{s/c}} \sum_{j=1}^{N_c} \mathbb{E}\left\{\mathbb{E}[Y_i(\mathcal{T}_j)|\mathcal{T}_j]\right\} = \frac{1}{N_{\text{tot}}} \sum_{i=1}^{N_{s/c}} \sum_{j=1}^{N_c} \mathbb{E}[\mu(\mathcal{T})] = \mu \quad (142)$$

using the law of total expectation (we follow the standard notation in which the inner $\mathbb{E}$ is over Y_i for fixed $\mathcal{T}_j$, and the outer one is over $\mathcal{T}_j$), and it has a variance of

$$\text{Var}(\hat{\mu}') = \mathbb{E}(\hat{\mu}'^2) - \mathbb{E}(\hat{\mu}')^2 = \mathbb{E}(\hat{\mu}'^2) - \mu^2. \quad (143)$$

There are three distinct types of terms in the sum

$$\mathbb{E}(\hat{\mu}'^2) = \frac{1}{N_{\text{tot}}^2} \sum_{i=1}^{N_{s/c}} \sum_{j=1}^{N_c} \sum_{k=1}^{N_{s/c}} \sum_{\ell=1}^{N_c} \mathbb{E}[Y_i(\mathcal{T}_j)Y_k(\mathcal{T}_\ell)], \quad (144)$$

which we will analyze separately.

Same shot from the same circuit ($i = k$ and $j = \ell$) There are N_{tot} such terms in Eq. (144), for which

$$\mathbb{E}[Y_i(\mathcal{T}_j)Y_i(\mathcal{T}_j)] = \mathbb{E}[(\pm 1)^2] = 1. \quad (145)$$

Different circuits ($j \neq \ell$) There are $N_c(N_c - 1)N_{s/c}^2$ such terms, for which $Y_i(\mathcal{T}_j)$ and $Y_k(\mathcal{T}_\ell)$ are independent since they come from different random circuits $\mathcal{T}_j$ and $\mathcal{T}_\ell$, which are themselves independent:

$$\mathbb{E}[Y_i(\mathcal{T}_j)Y_k(\mathcal{T}_\ell)] = \mathbb{E}\left\{\mathbb{E}[Y_i(\mathcal{T}_j)Y_k(\mathcal{T}_\ell)|\mathcal{T}_j, \mathcal{T}_\ell]\right\} = \mathbb{E}[\mu(\mathcal{T}_j)\mu(\mathcal{T}_\ell)] = \mu^2. \quad (146)$$

Different shots from the same circuit ($i \neq k$ and $j = \ell$) There are $N_{s/c}(N_{s/c} - 1)N_c$ such terms, which are conditionally independent given $\mathcal{T}_j$:

$$\mathbb{E}[Y_i(\mathcal{T}_j)Y_k(\mathcal{T}_j)] = \mathbb{E}\left\{\mathbb{E}[Y_i(\mathcal{T}_j)Y_k(\mathcal{T}_j)|\mathcal{T}_j]\right\} = \mathbb{E}[\mu(\mathcal{T}_j)^2] = \mu^2 + \Delta^2. \quad (147)$$

Therefore

$$\mathbb{E}(\hat{\mu}'^2) = \frac{1}{N_{\text{tot}}^2} \left[N_{\text{tot}} + N_c(N_c - 1)N_{s/c}^2\mu^2 + N_{s/c}(N_{s/c} - 1)N_c(\mu^2 + \Delta^2) \right] \quad (148)$$

$$= \mu^2 + \frac{1 - \mu^2}{N_{\text{tot}}} + \left(\frac{N_{s/c} - 1}{N_{s/c}} \right) \frac{\Delta^2}{N_c}, \quad (149)$$

so the standard error $\sqrt{\text{Var}(\hat{\mu}')}$ is given by Eq. (39) of the main text, which reduces to Eq. (38) when $N_{s/c} = 1$ (in which case we denote the estimator as $\hat{\mu}$ rather than $\hat{\mu}'$).

We now examine the quantity Δ for the different learning schemes discussed in the main text, which are meant to extract certain elements of a PTM $\mathbf{G}$ describing a noisy $R_{ZZ}(\theta)$ gate. In general, Δ will depend on the very elements of $\mathbf{G}$ we wish to measure, so its exact value cannot be known *a priori*. Instead, we approximate it by calculating Δ in the noiseless limit (where it is tractable) for various learning schemes, as discussed in the main text. We use the expression for Δ in Eq. (40) as a starting point. In general, the random channel $\mathcal{T}$ in that equation should describe not just the quantum gates being characterized and the random single-qubit gates surrounding them, but also any state-prep twirling, readout error, and readout twirling. (Readout error can be described by a stochastic matrix A that acts after an ideal measurement, as in Appendix D, or equivalently here, as a quantum channel that acts before an ideal measurement, as in Fig. 7.) However, these latter elements vanish in the noiseless limit, so we need only consider $\mathcal{T}$ comprising repeated ideal $R_{ZZ}(\theta)$ gates (described by a unitary channel $\mathcal{U}$ with PTM $\mathbf{U}$ from Eqs. (19) and (20)) surrounded by random Pauli gates.

We begin with modified cycle benchmarking, where $\mathcal{U}$ is repeated d times, and each occurrence is twirled independently over the full set of 2-qubit Paulis. This means

$$\mathcal{T} = \mathcal{T}^{(d)} \mathcal{T}^{(d-1)} \dots \mathcal{T}^{(1)} \quad (150)$$

for

$$\mathcal{T}^{(k)}(\rho) = P_\ell \mathcal{U}(P_\ell \rho P_\ell) P_\ell = \begin{cases} \mathcal{U}(\rho) = U \rho U^\dagger, & [P_\ell, Z \otimes Z] = 0 \\ \mathcal{U}^\dagger(\rho) = U^\dagger \rho U, & \{P_\ell, Z \otimes Z\} = 0, \end{cases} \quad (151)$$

where $P_\ell \sim \text{unif}(\mathbb{P})$ is sampled independently in each layer k . Therefore, when estimating the expectation value of some Pauli $P_i \in \mathbb{P}$:

$$\Delta^2 = \mathbb{E}\{\text{tr}[P_i \mathcal{T}(\rho)]^2\} - \mu^2 = \text{tr}\left[P_i^{\otimes 2} \mathbb{E}(\mathcal{T}^{\otimes 2})(\rho^{\otimes 2})\right] - \mu^2 = \text{tr}\left\{P_i^{\otimes 2} \left[\mathbb{E}(\mathcal{T}^{(k)\otimes 2})\right]^d (\rho^{\otimes 2})\right\} - \mu^2, \quad (152)$$

where we used the identity $\text{tr}(M)^2 = \text{tr}(M \otimes M)$ for any square matrix M to bring the expectation into the trace, then the fact that each layer is twirled independently to distribute the overall expectation into each layer. Since P_ℓ in Eq. (151) commutes or anti-commutes with $Z \otimes Z$ with equal probabilities,

$$\mathbb{E}(\mathcal{T}^{(k)\otimes 2}) = \frac{1}{2} \mathcal{U}^{\otimes 2} + \frac{1}{2} \mathcal{U}^{\dagger \otimes 2} =: \mathcal{V}. \quad (153)$$

So for a generic initial state $\rho = \frac{1}{4} \sum_k s_k P_k$,

$$\Delta^2 = \text{tr}\left[P_i^{\otimes 2} (\mathcal{V}^d)(\rho)\right] - \mu^2 = \sum_{k\ell} s_k s_\ell \underbrace{\frac{1}{16} \text{tr}\left[P_i^{\otimes 2} \mathcal{V}^d(P_k \otimes P_\ell)\right]}_{(\mathbf{V}^d)_{ii,k\ell}} - \mu^2, \quad (154)$$

where $(\mathbf{V}^d)_{ii,k\ell}$ are the PTM elements of $\mathcal{V}^d$. Moreover, because $\text{span}\{P_i, P_j\}$ is an invariant subspace of $\mathcal{U}$ for $P_j \propto (Z \otimes Z)P_i$, $\text{span}(\mathbb{S})$ is invariant under $\mathcal{V}$ for

$$\mathbb{S} = (P_i \otimes P_i, P_i \otimes P_j, P_j \otimes P_i, P_j \otimes P_j). \quad (155)$$

In other words, because the PTM of $\mathcal{U}$ is block-diagonal with 2×2 blocks, the PTM of $\mathcal{V}$ is block-diagonal with 4×4 blocks (in an appropriate ordered basis). We can therefore evaluate Δ by calculating one of these blocks, denoted $\mathbf{V}_{\mathbb{S}}$, and taking the d^{th} power of it. Concretely,

$$\begin{aligned} \mathbf{V}_{\mathbb{S}} &= \left(\frac{1}{16} \text{tr} [S_k \mathcal{V}(S_\ell)] \right)_{1 \leq k, \ell \leq 4} = \frac{1}{2} \begin{pmatrix} U_{ii} & U_{ij} \\ U_{ji} & U_{jj} \end{pmatrix}^{\otimes 2} + \frac{1}{2} \begin{pmatrix} U_{ii} & -U_{ij} \\ -U_{ji} & U_{jj} \end{pmatrix}^{\otimes 2} \\ &= \begin{pmatrix} U_{ii}^2 & 0 & 0 & U_{ij}^2 \\ 0 & U_{ii}U_{jj} & U_{ij}U_{ji} & 0 \\ 0 & U_{ij}U_{ji} & U_{ii}U_{jj} & 0 \\ U_{ji}^2 & 0 & 0 & U_{jj}^2 \end{pmatrix}, \end{aligned} \quad (156)$$

since $U_{ij} = -U_{ji}$, where S_k denotes the k^{th} element of $\mathbb{S}$. Then

$$\Delta^2 = (1 \ 0 \ 0 \ 0) \begin{pmatrix} U_{ii}^2 & 0 & 0 & U_{ij}^2 \\ 0 & U_{ii}U_{jj} & U_{ij}U_{ji} & 0 \\ 0 & U_{ij}U_{ji} & U_{ii}U_{jj} & 0 \\ U_{ji}^2 & 0 & 0 & U_{jj}^2 \end{pmatrix}^d \begin{pmatrix} s_i s_i \\ s_i s_j \\ s_j s_i \\ s_j s_j \end{pmatrix} - \mu^2, \quad (157)$$

where

$$\mu = \text{tr} \left[P_i \left(\frac{1}{2} \mathcal{U} + \frac{1}{2} \mathcal{U}^\dagger \right)^d (\rho) \right] = U_{ii}^d s_i. \quad (158)$$

If $[P_i, Z \otimes Z]$, which is the case when learning Type 1 elements through modified CB, then $U_{ii} = U_{jj} = 1$ and $U_{ij} = U_{ji} = 0$, so $\mathbf{V}_{\mathbb{S}} = I$ and therefore $\Delta = 0$ in the noiseless limit. The same is true for standard CB on Clifford gates. Note however that this property does not arise automatically: if instead $\{P_i, Z \otimes Z\} = 0$, then $U_{ii} = U_{jj} = \cos(\theta)$ and $U_{ij} = -U_{ji} = \pm \sin(\theta)$, so

$$\Delta^2 = \frac{1}{2} (1, 0, 0, 0) \begin{pmatrix} 1 + \cos(2\theta)^d & 0 & 0 & 1 - \cos(2\theta)^d \\ 0 & 1 + \cos(2\theta)^d & -1 + \cos(2\theta)^d & 0 \\ 0 & -1 + \cos(2\theta)^d & 1 + \cos(2\theta)^d & 0 \\ 1 - \cos(2\theta)^d & 0 & 0 & 1 + \cos(2\theta)^d \end{pmatrix} \begin{pmatrix} s_i s_i \\ s_i s_j \\ s_j s_i \\ s_j s_j \end{pmatrix} - \cos(\theta)^{2d} s_i^2. \quad (159)$$

Since $(s_i, s_j) \rightarrow (1, 0)$ in the noiseless limit, the resulting Δ is given by Eq. (48) of the main text. In other words, the random circuits arising in modified CB lead to expectation values that are highly spread out when measuring Paulis that anti-commute with $Z \otimes Z$. This is another reason why modified CB is generally impractical for learning Type 2 elements.

In contrast, partial-twirl benchmarking (PTB) and correlated-twirl benchmarking (CTB) both have $\Delta = 0$ in the noiseless limit since, for any given depth d , the random circuits they use are all logically equivalent (i.e., they describe identical unitaries). More precisely, in PTB each layer $\mathcal{T}^{(k)}$ has the same form as in Eq. (151), but $P_\ell \sim \text{unif}(\mathbb{P}_C)$, so $\mathcal{T}^{(k)} = \mathcal{U}$ in the noiseless limit, which immediately gives $\Delta = 0$ using Eq. (152). The situation is similar with CTB, although to analyze that scheme it is convenient to decompose $\mathcal{T} = \mathcal{T}^{(d/2)} \dots \mathcal{T}^{(1)}$ with (even) depth d into $d/2$ independent random layers

$$\mathcal{T}^{(k)}(\rho) = \mathcal{P}_m \mathcal{U} \mathcal{P}_m \mathcal{P}_\ell \mathcal{U} \mathcal{P}_\ell, \quad (160)$$

each of which comprises two $R_{ZZ}(\theta)$ gates, where we use the notation $\mathcal{P}_\ell(\rho) = P_\ell \rho P_\ell$. Then $P_\ell \sim \text{unif}(\mathbb{P})$ and $P_m \sim \text{unif}(\mathbb{P}_A)$ if $P_\ell \in \mathbb{P}_C$, or $P_m \sim \text{unif}(\mathbb{P}_C)$ if $P_\ell \in \mathbb{P}_A$. In either case, $\mathcal{T}^{(k)}$ reduces to the identity channel in the noiseless limit, which immediately implies $\Delta = 0$.

G Partial-Twirl Benchmarking Details

In this section we derive Eqs. (52), (53), and (58)–(60) from the main text. Consider the 2×2 matrix

$$B = \begin{pmatrix} \mathbf{G}_{ii} & \mathbf{G}_{ij} \\ \mathbf{G}_{ji} & \mathbf{G}_{jj} \end{pmatrix}, \quad (161)$$

which forms a block on the bottom-right of the diagonal of $\mathbf{G}_c$, the PTM describing a noisy $R_{ZZ}(\theta)$ gate twirled over $\mathbb{P}_c$, the set of Paulis that commute with ZZ . The blue and green colors serve to denote Type 2 and Type 3 elements respectively. As in the main text, we will assume that B has complex eigenvalues, i.e., that its eigenvalues satisfy condition (51). To avoid any ambiguity with complex square roots being multi-valued (e.g., $+i$ and $-i$ are both square roots of -1), we will write these eigenvalues as

$$\lambda_{\pm} = x \pm iy = re^{\pm i\omega} \quad (162)$$

here, where the real and imaginary parts

$$x = \frac{1}{2}(\mathbf{G}_{ii} + \mathbf{G}_{jj}) \quad y = \frac{1}{2}\sqrt{-(\mathbf{G}_{ii} - \mathbf{G}_{jj})^2 - 4\mathbf{G}_{ij}\mathbf{G}_{ji}} \quad (163)$$

are both well-defined, and the magnitude r and argument ω are defined in the usual way in Eq. (53) of the main text. Likewise, we will write the corresponding (unnormalized) eigenvectors of B as

$$\vec{v}_{\pm} = \begin{pmatrix} \mathbf{G}_{ii} - \mathbf{G}_{jj} \pm 2iy \\ 2\mathbf{G}_{ji} \end{pmatrix}. \quad (164)$$

It follows that

$$\begin{pmatrix} 1 \\ 0 \end{pmatrix} = \frac{-i}{4y}(\vec{v}_+ - \vec{v}_-), \quad (165)$$

so for any circuit depth d , we can express the top-left element of B^d as

$$\begin{aligned} (1 \ 0) \begin{pmatrix} \mathbf{G}_{ii} & \mathbf{G}_{ij} \\ \mathbf{G}_{ji} & \mathbf{G}_{jj} \end{pmatrix}^d \begin{pmatrix} 1 \\ 0 \end{pmatrix} &= \frac{-i}{4y} \begin{pmatrix} 1 \\ 0 \end{pmatrix} \cdot (\lambda_+^d \vec{v}_+ - \lambda_-^d \vec{v}_-) = \frac{-i}{4y} [(\mathbf{G}_{ii} - \mathbf{G}_{jj})(\lambda_+^d - \lambda_-^d) + 2iy(\lambda_+^d + \lambda_-^d)] \\ &= r^d \left[\left(\frac{\mathbf{G}_{ii} - \mathbf{G}_{jj}}{2y} \right) \sin(\omega d) + \cos(\omega d) \right]. \end{aligned} \quad (166)$$

We can combine both terms in the square brackets using the identity

$$a \cos(\omega d - \delta) = a [\sin(\delta) \sin(\omega d) + \cos(\delta) \cos(\omega d)] \quad (167)$$

by demanding that the amplitude a and phase δ satisfy

$$a \sin(\delta) = \frac{\mathbf{G}_{ii} - \mathbf{G}_{jj}}{2y} \quad a \cos(\delta) = 1, \quad (168)$$

which immediately gives the expressions for a and δ in Eq. (53) of the main text. These definitions then yield

$$(1 \ 0) \begin{pmatrix} \mathbf{G}_{ii} & \mathbf{G}_{ij} \\ \mathbf{G}_{ji} & \mathbf{G}_{jj} \end{pmatrix}^d \begin{pmatrix} 1 \\ 0 \end{pmatrix} = ar^d \cos(\omega d - \delta), \quad (169)$$

as in Eq. (52).

We now derive Eqs. (58)–(60) of the main text, which give $\mathbf{G}_{ii}$, $\mathbf{G}_{jj}$ and $\mathbf{G}_{ij}\mathbf{G}_{ji}$ as functions of r , ω and δ , which can be estimated experimentally. Note, from the definitions above, that

$$r \cos \omega = x \quad r \sin \omega = y \quad ay = \sqrt{-\mathbf{G}_{ij}\mathbf{G}_{ji}}. \quad (170)$$

Combining these equations with Eq. (168) gives

$$\cos(\delta) = \frac{1}{a} = \frac{y}{\sqrt{-\mathbf{G}_{ij}\mathbf{G}_{ji}}} = \frac{r \sin(\omega)}{\sqrt{-\mathbf{G}_{ij}\mathbf{G}_{ji}}}, \quad (171)$$

which immediately leads to Eq. (58) of the main text. Note that since $\cos(\delta)^{-2} = 1 + O(\delta^2)$ for δ near 0, $\mathbf{G}_{ij}\mathbf{G}_{ji}$ depends only weakly on the fitted value of δ when the noise is weak (meaning δ is small). Similarly,

$$\tan(\delta) = \frac{\mathbf{G}_{ii} - \mathbf{G}_{jj}}{2y} = \frac{\mathbf{G}_{ii} - \mathbf{G}_{jj}}{2r \sin(\omega)}, \quad (172)$$

so

$$\frac{\mathbf{G}_{ii} - \mathbf{G}_{jj}}{2} = r \sin(\omega) \tan(\delta). \quad (173)$$

Adding this quantity to x , or subtracting it from x , gives the expressions for $\mathbf{G}_{ii}$ and $\mathbf{G}_{jj}$ from Eqs. (59) and (60) respectively. However, since $\tan(\delta) = O(\delta)$, these expressions depend strongly on the fitted values of δ , so we do not advise using them. Instead, we recommend correlated-twirl benchmarking.