ARTICLE OPEN

Experimental quantum verification in the presence of temporally correlated noise

S. Mavadia^{1,2}, C. L. Edmunds^{1,2}, C. Hempel^{1,2}, H. Ball¹, F. Roy¹, T. M. Stace³ and M. J. Biercuk^{1,2}

Growth in the capabilities of quantum information hardware mandates access to techniques for performance verification that function under realistic laboratory conditions. Here we experimentally characterise the impact of common temporally correlated noise processes on both randomised benchmarking (RB) and gate-set tomography (GST). Our analysis highlights the role of sequence structure in enhancing or suppressing the sensitivity of quantum verification protocols to either slowly or rapidly varying noise, which we treat in the limiting cases of quasi-DC miscalibration and white noise power spectra. We perform experiments with a single trapped $^{171}\text{Yb}^+$ ion-qubit and inject engineered noise ($\propto \hat{\sigma}_z$) to probe protocol performance. Experiments on RB validate predictions that measured fidelities over sequences are described by a gamma distribution varying between approximately Gaussian, and a broad, highly skewed distribution for rapidly and slowly varying noise, respectively. Similarly we find a strong gate set dependence of default experimental GST procedures in the presence of correlated errors, leading to significant deviations between estimated and calculated diamond distances in the presence of correlated $\hat{\sigma}_z$ errors. Numerical simulations demonstrate that expansion of the gate set to include negative rotations can suppress these discrepancies and increase reported diamond distances by orders of magnitude for the same error processes. Similar effects do not occur for correlated $\hat{\sigma}_x$ or $\hat{\sigma}_y$ errors or depolarising noise processes, highlighting the impact of the critical interplay of selected gate set and the gauge optimisation process on the meaning of the reported diamond norm in correlated noise environments.

npj Quantum Information (2018)4:7; doi:10.1038/s41534-017-0052-0

INTRODUCTION

Quantum characterisation, validation, and verification (QCVV) techniques are broadly used in the quantum information community in order to evaluate the performance of experimental hardware. A variety of techniques have emerged including randomised benchmarking (RB),^{1,2} purity benchmarking,³ process tomography,^{4–7} adaptive methods,⁸ and gate-set tomography (GST).^{9,10} Each protocol has relative strengths and weaknesses; for instance, RB has low experimental overhead but only provides average information about gate performance, while process tomography provides more information at the cost of unfavourable scaling in measurement overhead.¹¹ Despite their differences, these protocols share the common theme that they were originally developed and mathematically formalised assuming that error processes are statistically independent and do not exhibit strong correlations in time.^{1,2,10}

Even in highly controlled laboratory environments there are a range of noise sources that, when applied to a qubit concurrent with logical gate operations, produce effective error models that diverge significantly from the assumptions underlying most QCVV protocols. For example, slow variations in ambient magnetic fields or drifts in amplifier gain can produce temporally correlated noise processes, often characterised through a power spectral density possessing large weight at low frequencies.^{12–14} Moreover, these error processes may exhibit gate-dependent behaviour. So far such processes have been largely ignored in experimental QCVV,

with predominantly phenomenological attempts used to explain deviations from ideal outputs.¹⁵ Understanding that such an approach is untenable when attempting to rigorously compare QCVV results to metrics relevant to quantum error correction has recently led to an expansion of theoretical activity in this space.^{3,16–20}

In this work our objectives are to experimentally characterise and explain the impact of temporally correlated noise processes on the outputs of QCVV protocols, and to identify potential modifications enabling users to improve the utility of the information returned. We perform QCVV experiments using a single trapped $^{171}\text{Yb}^+$ ion as a long-lived, high-stability qubit. Our study implements engineered frequency noise ($\propto \hat{\sigma}_z$) in the control system in order to study the impact of different temporal noise correlations on QCVV results. We apply noise in the two extremes, either quasi-DC offsets or noise with an effective white power spectrum to approximate slowly and rapidly varying noise, respectively. Measurements reveal that QCVV outputs diverge significantly when subject to these different types of noise, highlighting potential circumstances where the information extracted from a given protocol may no longer accurately represent the true error processes experienced by individual gates. Our experiments are compared against analytic calculations linking the underlying structure of the QCVV sequences with the manifestation of specific characteristics associated with the presence of noise correlations.

¹ARC Centre for Engineered Quantum Systems, School of Physics, The University of Sydney, Sydney, NSW, Australia; ²National Measurement Institute, West Lindfield, NSW 2070, Australia and ³ARC Centre for Engineered Quantum Systems, School of Physics and Mathematics, The University of Queensland, St Lucia, QLD, Australia

Correspondence: M. J. Biercuk (michael.biercuk@sydney.edu.au)

S. Mavadia, C. L. Edmunds and C. Hempel contributed equally to this work.

Received: 3 July 2017 Revised: 28 November 2017 Accepted: 5 December 2017

Published online: 06 February 2018

Fig. 1 QCVV sequence construction and mapping to accumulated error. **a** Overview of unitary sequence construction for RB and GST, using Clifford gates, C_i or fiducial operations, $F_{\alpha, \beta}$ and repeated germs $(G)^n$ respectively. **b** Schematic representation of slowly and rapidly varying noise with relevant time scales defined by the sequence where δ represents the instantaneous noise values drawn from a normal distribution with σ^2 variance. Grey lines are other possible noise realisations. For RB, the noise is sampled from this distribution and varies shot-to-shot between noise realisations, while in GST a single value is selected for the entire set of experiments. **c** Sequence-dependent “random walk” calculated for an arbitrary QCVV sequence (here according to the RB prescription) with $J = 100$ in Pauli space. Green dot indicates origin and black triangle indicates sequence terminus. Blue line represents the 3D walk, which can be used to calculate the trace infidelity while the grey line represents the 2D projection, measurable in a standard projective measurement. The green arrow indicates the net walk vector, $\vec{V}_{2D}$, given unit step size

We examine two common QCVV protocols in the experimental quantum information community: RB and GST. The construction of these protocols follows a similar pattern, a series of unitary quantum operations is applied to one or more qubits sequentially in time, followed by a projective measurement (Fig. 1a). Experimental measurements are acquired and combined, then experimental parameters are changed according to some prescription (e.g. changing the sequence length J) and further data are collected. The variation in QCVV protocols predominantly comes from the different constituent operations that are applied and the analysis techniques by which measurement results are post-processed to extract information.

In RB, sequences are constructed by concatenating unitary operations U_i selected at random from the 24 Clifford operations C_i . The final operation in a sequence of length J is selected to invert the net rotation $U_J = \left(\prod_{i=1}^{J-1} C_i\right)^{-1}$, such that the sequence implements a net identity $\prod_{i=1}^J C_i = \hat{I}$. In experimental GST as defined by the pyGSTi package, by contrast, operations are selected deterministically according to a tabulated routine comprising specifically crafted sequences that are designed to maximise overall sensitivity to all detectable error types. These operations are constructed by concatenating so-called “germs”, short sequences implementing predefined unitary rotations, which, in our case, are constructed from a subset of Clifford gates. The first and last unitaries $U_{1,J} \in \{F_{\alpha}, F_{\beta}\}$, termed the “fiducial” operations, effectively set the reference frame for state-preparation and measurement (Fig. 1a); see “Methods” for further detail.

In our experiments we engineer noise in order to permit quantitative analysis of QCVV outputs under known conditions. We compare the outputs obtained from both RB and GST for two distinct noise-correlation regimes. Firstly, where the engineered noise is implemented as a constant miscalibration over the entire sequence, which is the extreme case for slowly varying noise and produces temporally correlated errors. Secondly, where the engineered noise is rapidly varying (yielding an approximately white power spectrum), which leads to errors that are uncorrelated between gates (Fig. 1b). We now introduce a framework for interpreting the impact of sequence structure and noise correlations on measurement outcomes to facilitate an analysis of our results.

RESULTS

Mapping noise to measured error in RB

The key analytic tool for our study is a formalism mapping an applied noise model to an output error for a given Clifford sequence, following a procedure derived in ref. 17. Error accumulation over a given Clifford sequence maps to a “random walk” in a three-dimensional vector-space representing the action of sequential error unitaries in the operator space spanned by the Pauli operators, $\hat{\sigma}_{\{x,y,z\}}$ (Fig. 1c). For $\hat{\sigma}_z$ noise, the l th step of the walk is calculated by conjugating $\hat{\sigma}_z$ with the entire operator subsequence $K_{l-1} \equiv \prod_{q=1}^{l-1} U_q$ up to the $(l-1)$ th gate, with multiplication performed from the left. This conjugation always results in a member of the Pauli group, allowing us to compactly write $\mathbf{P}_l \equiv K_{l-1}^\dagger \hat{\sigma}_z K_{l-1} = \hat{\mathbf{r}}_l \cdot \vec{\sigma}$, where $\vec{\sigma} = (\hat{\sigma}_x, \hat{\sigma}_y, \hat{\sigma}_z)$ and $\hat{\mathbf{r}}_l \in \{\pm \hat{\mathbf{x}}, \pm \hat{\mathbf{y}}, \pm \hat{\mathbf{z}}\}$. The direction of $\mathbf{P}_l$ in Pauli space therefore maps to the Cartesian unit vector $\hat{\mathbf{r}}_l$ associated with the l th step of a J -step walk $\vec{\mathbf{R}} \equiv \sum_{l=1}^J \delta_l \hat{\mathbf{r}}_l$. For our chosen error model, the step length, δ_l , captures the integrated phase between the driving field and qubit during execution of the single gate U_l . In terms of experimental parameters, $\delta_l = \Delta/\Omega$, where Δ/Ω is the detuning expressed in terms of the experimental Rabi frequency Ω (see “Methods”).

The overall form of the walk is a statistical measure of how the sequence itself interacts with the noise process to produce a net, measurable accumulation of error. Sequences that are highly susceptible to error accumulation produce walks that migrate far from the origin, while sequences exhibiting error suppression produce walks that meander back towards the origin. The net walk length is captured in the mean-squared distance from the origin $\langle \|\vec{\mathbf{R}}\|^2 \rangle$, averaged over noise realisations. This links to the

trace fidelity, defined as $\mathcal{F}_{\text{trace}} = \left\langle \left| \text{Tr} \left(\prod_{l=1}^J \tilde{U}_l \right) \right|^2 \right\rangle / 4$, where $\tilde{U}$ are modified unitary operations to take into account the effect of the $\hat{\sigma}_z$ noise. We then define the infidelity $\mathcal{I}_{\text{trace}} = 1 - \mathcal{F}_{\text{trace}} \simeq \langle \|\vec{\mathbf{R}}\|^2 \rangle$.

Appropriately linking this picture of error accumulation to standard laboratory measurements requires consideration of the measurement routine itself. In typical measurements the qubit Bloch vector at the end of the sequence is projected onto the quantisation axis, z , with basis states $|0\rangle$ and $|1\rangle$. A measurement of this type is therefore insensitive to net rotations around that

axis of the Bloch sphere, meaning that it only probes a 2D projection of the 3D walk onto the xy -plane. Our preferred metric is the survival probability, $\mathcal{P}$, that may be linked directly to such a 2D projection (grey line, Fig. 1c) as $\langle \|\vec{R}_{2D}\|^2 \rangle = 1 - \mathcal{P}$ where $\mathcal{P} = \left\langle \left| \langle 0 | \prod_{i=1}^J U_i | 0 \rangle \right|^2 \right\rangle$, $\langle \|\vec{R}_{2D}\|^2 \rangle = \langle \|\vec{R}\|^2 \rangle - \langle \|\vec{R}_z\|^2 \rangle$, and $\langle \|\vec{R}_z\|^2 \rangle$ is the mean-squared walk length along the quantisation axis (see Supplementary Material for details).

At this stage we must link the correlation properties of the noise to the form of the walk for a specific sequence. Considering only the underlying properties of the sequence, we may assume unit-length steps, resulting in a deterministic sequence-dependent walk with length $\vec{V} \equiv \sum_{i=1}^J \vec{r}_i$. The presence or absence of temporal noise correlations is now captured through a rescaling of the individual steps in the deterministic walk for a specific sequence. In the case of slowly varying noise, and to first-order approximation, the net error can be separated into two independent parts, $\|\vec{R}\|^2 = \delta^2 \|\vec{V}\|^2$, where δ is the value of the noise and $\|\vec{V}\|$ is the net unit-step walk specific to a particular sequence.¹⁷ However, in the case of rapidly varying noise these two terms are no longer separable and the net error must be calculated as the convolution of the noise value at each timestep and each individual step in the random walk, $\|\vec{R}\|^2 = \|\sum_{i=1}^J \delta_i \vec{r}_i\|^2$.

Experimental platform and engineered noise

We perform experiments using the hyperfine qubit in a single trapped $^{171}\text{Yb}^+$ ion driven by microwaves near 12.64 GHz, with basis states $|0\rangle \equiv |^2S_{1/2}, F=0, m_F=0\rangle$ and $|1\rangle \equiv |^2S_{1/2}, F=1, m_F=0\rangle$. Our calibration process permits accurate determination of the (first-order magnetic-field-insensitive) qubit transition frequency to within approximately 1 Hz. In our laboratory, this qubit and the associated control system have been demonstrated to possess a coherence time of $T_2 \sim 1$ s, measurement fidelity of $\sim 99.7\%$ limited by photon collection efficiency, and error rates from intrinsic system noise of $p_{\text{RB}} \approx 6 \times 10^{-5}$ using “baseline” RB experiments (see Supplementary Figures). Details of the control system and experimental protocols for QCVV techniques used here are presented in the *Methods*, and information about various detection procedures in use for estimating $\mathcal{P}$ (including a Bayesian method) are found in the Supplementary Materials.

We engineer $\hat{\sigma}_z$ noise applied concurrently with Clifford operations through the application of a detuning, Δ , of the qubit driving field from resonance using an externally modulated vector signal generator. As the detuning is applied concurrently with driven qubit rotations about x and y axes, rotation errors arise along multiple directions on the Bloch sphere, rather than being purely $\hat{\sigma}_z$ in character. An additional violation of typical assumptions employed in RB is that different Clifford gates are physically decomposed into base rotations with different durations, which means that our formal error model will also be gate-dependent.¹⁹

For each of our two limiting noise cases we engineer N different noise “realisations” in order to average over an appropriate ensemble. In our experiments, we set the distribution of noise $\Delta/\Omega \sim \mathcal{N}(0, \sigma^2)$, where σ^2 is the variance of the distribution, such that the root-mean-square value is approximately equivalent in both cases once averaged over all noise realisations. The specific implementation of noise engineering and its impact on the conduct of RB and GST is described in “Methods”, and additional details on the error model are provided in the Supplementary Materials.

Experiments involve state preparation in the $|0\rangle$ state, application of a unitary sequence appropriate for a QCVV protocol while subject to noise, and projective measurement of the qubit along

the quantisation axis. The sequence of operations applied and the measurement procedure are determined by the protocol in use.

RB survival probability distributions

In the limit of rapidly varying noise, all sequences of randomly ordered Clifford gates with length J are equivalent under noise averaging, and all sequence survival probabilities tend towards the mean. Recent theoretical studies have demonstrated that measurements on RB sequences in the presence of temporal noise correlations, can produce a divergence between average and worst-case reported trace fidelities.^{17,20} Thus we find that measurement outcomes for different RB sequences are characterised by distributions with distinctly different shapes depending on the temporal correlations in the noise. The standard practice of combining all measurements to extract an RB error rate, p_{RB} , from the decay of the mean over all J -gate sequences as a function of J , results in a global ensemble average and does not take advantage of this information (formally, as the noise we implement exhibits temporal correlations, the value of p_{RB} one extracts may not be meaningful as a measure of average Clifford gate error). Our analysis takes advantage of the additional information which is always present in an RB experiment in order to evaluate the impact of noise correlations and deduce useful information about the underlying error process.

In our experimental study we measure the noise-averaged survival probabilities for a set of sequences $\{\eta_i\}_J$, indexed by i and of length J , for different lengths $25 \leq J \leq 200$ (Fig. 2a), where we implement the same set of J -gate sequences under application of either slowly or rapidly varying detuning noise. For an arbitrary individual sequence, η_i and a single noise realisation, n , we perform r nominally identical repetitions of the experiment. We combine the information from the outcomes of these individual repetitions to produce a maximum-likelihood estimate of survival probability, $\mathcal{P}_{i,n}$ (see Supplementary Materials). The use of multiple repetitions under identical conditions reduces quantum projection noise in the qubit measurement and assists in isolating specific quantitative contributions to the distribution of survival probabilities, though this is not possible without noise engineering. In general, we average measured outcomes over a fixed number of noise realisations to yield $\mathcal{P}_{i,\langle \cdot \rangle}$ for a fixed sequence η_i . From here on, we will refer to this noise-averaged survival probability as $\mathcal{P}$.

In the case of rapidly varying noise we observe the distribution of sequence outcomes is symmetrically spread around the sequence-averaged mean survival probability, $\overline{\mathcal{P}}(J)$, and the entire distribution shifts away from zero error with increasing J (red data, Fig. 2a). The presence of slowly varying noise, by contrast, produces a broad distribution of measured $\mathcal{P}$ over each set $\{\eta_i\}_J$, demonstrating a positively skewed set of outcomes and the persistence of a long tail at higher error rates (lower survival probabilities). In this case, as J increases the distribution broadens but remains skewed. Under both noise correlation cases, the measured $\overline{\mathcal{P}}(J)$ remain approximately the same. The differences in the distribution of measured survival probabilities over sequences under these two noise models reproduces the central predictions of ref.¹⁷

We compare the characteristics of the distributions themselves against analytic predictions for both slowly and rapidly varying noise, beginning with the measured expectation, $\mathbb{E}(\mathcal{I})$, and variance, $\mathbb{V}(\mathcal{I})$ (Fig. 2b, c), finding good agreement by taking only the applied noise strength as an input into a theoretical model (see Supplementary Materials). More specifically, theoretical predictions suggest that the distribution of outcomes under both noise models – as well as intermediate models described by coloured power spectra – should be well described by a gamma distribution.¹⁷ The general gamma distribution probability density

Fig. 2 RB distributions over sequences in the presence of different noise correlations. **a** Standard RB protocol showing survival probability as a function of J for the same set of sequences implemented under slowly varying (grey) or rapidly varying (red) noise with $\Delta_{\text{RMS}} = 1$ kHz. In these experiments the Rabi frequency, $\Omega = 22.5$ kHz. Each experiment is repeated $r = 25$ to $r = 30$ times under fixed conditions, and each sequence fidelity is averaged over 200 noise realisations. Lines represent exponential fits to the sequence-averaged survival probability $1 - \bar{P}(J) = 0.5 - (0.5 - \kappa)e^{-p_{\text{RB}}J}$, weighted by the variance over sequences for each J , and are used to extract the error-rate p_{RB} . Here $\kappa = 3 \times 10^{-3}$ represents state preparation and measurement error. **b, c** Scaling of $\mathbb{E}(\mathcal{I})$ and $\mathbb{V}(\mathcal{I})$ against sequence length J , comparing experimental values (markers) against first-principles theory (lines) as per¹⁷ modified to state fidelity (2D walk) and noise applied concurrently with gate implementation. See Supplementary Material for details. **d–g** Histograms for data in panel **a** in the presence of slowly varying noise. Green line: fitted gamma distribution with shape parameter fixed, $\alpha = 1$. Black line: gamma distribution using input parameters calculated from first principles (see text). χ^2 values for calculated (fitted) gamma distributions are {0.354(0.091), 0.212(0.078), 0.241(0.204), 0.348(0.348)}

function is given by

$$\Gamma(a, \beta): f(x) = \frac{x^{a-1}}{\Gamma(a)\beta^a} \exp\left[-\frac{x}{\beta}\right], \quad (1)$$

where a and β are the shape and scale parameters and $\Gamma(x)$ is the gamma function. The form of the gamma distribution will vary significantly between the limiting noise cases treated here, tending towards a symmetric Gaussian for rapidly varying noise and a broader positively skewed distribution in the presence of slowly varying noise, as determined by the values of a and β .

Figure 2d–g shows histograms of RB sequence survival probabilities in the presence of the extreme case of slowly varying noise, quasi-DC miscalibration. We overlay gamma distributions calculated from first principles using no free parameters (black lines) as $\Gamma(1, (2J\sigma^2/3)(1/2 + \pi^2/96))$, and fixing $\alpha = 1$ while allowing β to vary as a fit parameter (green lines). The theoretical prediction captures both the measured skew towards high survival probabilities and the approximate “length” of the tail at low survival probabilities. We believe that residual

disagreement between data and first-principles calculations arises due to both limited sequence sampling and contributions from higher-order analytic error terms when the approximation $J\sigma^2 \ll 1$ is no longer valid. Importantly, data and theory show the mode of the distribution is close to unity survival probability ($P = 1$) and therefore corresponds to a lower error than the mean. For details on modifications to the theory presented in¹⁷ accounting for the specific noise and gate-dependent error model employed in our experiments, contributions from higher-order terms, and expanded data sets including larger sequence numbers, see Supplementary Material.

Modification of RB for identification of model violation

The fact that the distribution of sequence survival probabilities under slowly varying noise does not converge to the mean indicates sequence-dependence in the resulting error accumulation. The emergence of this phenomenology is elucidated through an examination of the walks for different sequences. Under this type of noise certain sequences possess walks with large $\|\vec{V}_{2D}\|^2$, hence amplifying the accumulation of error, while others tend back towards the origin and show reduced accumulated error (Fig. 3a, b). We classify sequences as “long-walk” if they possess a 2D projection beyond the diffusive mean-squared limit for an unbiased random walk, $\|\vec{V}_{2D}\|^2 > \frac{2}{3}J$.

We link between the sequence walk in Pauli space and the noise-averaged survival probability by displaying the experimentally measured $1 - \bar{P}$ for sequences of fixed length $J = 200$ against the calculated 2D walk length, $\|\vec{V}_{2D}\|^2$ (Fig. 3c). Data are presented for both rapidly varying (red open markers) and slowly varying (grey solid markers) noise, where the same set of sequences is used between the noise models. Measurements for rapidly varying noise are fit with a line possessing a slope approximately consistent with zero, while for the same sequences under slowly varying noise, the measurements show a positive dependence on $\|\vec{V}_{2D}\|^2$ as expected. We believe the significant scatter in the plot is partially due to a concurrently acting noise source and higher-order contributions to error, neither of which are incorporated in the first principles calculation of the walk, $\|\vec{V}_{2D}\|^2$ (see Supplementary Material and Appendix C of ref. ¹⁷). Nonetheless, the effect of sequence structure on measured survival probability is clearly visible for the case of slowly varying noise.

In aggregate, this phenomenology gives rise to the skewed gamma distribution under slowly varying noise described above, and the convergence of all noise-averaged survival probabilities for individual sequences to the ensemble average when the noise is rapidly varying. However, preselection of RB sequences possessing large calculated, unit-step walks also provides a mechanism to both identify the presence of temporally correlated errors and extract an RB outcome that more closely approximates worst-case errors. In Fig. 3d we plot $1 - \bar{P}$ vs. J for a subset of sequences preselected to possess long walks as in Fig. 3a, whose error rates we denote $p_{\text{RB}}^{\text{LW}}(J)$. We choose that the preselection of long walks is based on the condition $\|\vec{V}_{2D}\|^2 > 2 \times \frac{2}{3}J$.

When these long-walk sequences are subjected to rapidly varying noise, the distribution of survival probabilities over sequences remains approximately Gaussian about the mean, and the expectation value over this subset closely approximates the expectation value over an unbiased random sampling of the 24^J possible J -gate sequences, $\bar{P}_{\text{LW}}^{\text{rapid}} \approx \bar{P}^{\text{rapid}}$, (Fig. 3d, red solid line and blue dashed line). However, in the presence of slowly varying noise we observe a larger spread in $\bar{P}_{\text{LW}}^{\text{slow}}$ than that achieved with unbiased sampling. The difference between the sequence-averaged survival probabilities in these noise cases arises solely because of the intrinsic properties of the sequences in use.

Fig. 3 RB using long-walk sequences. **a, b** Schematic representations of long **a** and short **b** length walks in 3D (coloured lines) and 2D (black lines), defined relative to a limit deduced from diffusive behaviour, as indicated by the blue circle. **c** Noise-averaged fidelity distributions of the same sequences as a function of walk length in the 2D plane. Measured infidelity vs. 2D walk length, $\|\vec{V}_{2D}\|^2$, when subject to slowly varying (grey) and rapidly varying (red) noise with linear fit overlaid. The slope of this fit is $(0.8 \pm 1) \times 10^{-5}$, consistent with zero. **d** RB using long-walk sequences. Solid red line corresponds to RB performed using 20 long-walk sequences and rapidly varying noise. Extracted $p_{RB}^{(LW)}$ matches that extracted under the same conditions using an unbiased sampling of all sequences (dashed line). Grey line corresponds to RB using the same long-walk sequences and slowly varying noise. For the exponential fits, state-preparation and measurement error, κ , is fixed to 3×10^{-3} .

Extracting an RB gate-error-rate, $p_{RB}^{(LW)}$ from $\overline{P}_{LW}(J)$ in the presence of slowly varying noise, we typically find an increase $p_{RB}^{(LW)} \sim 2 - 5 \times p_{RB}$ relative to standard sequence sampling, depending on the number of long-walk sequences employed, and the threshold value of $\|\vec{V}_{2D}\|^2$ used to define a “long walk” (Fig. 3c). This approach effectively constitutes the construction of an RB protocol that increases the reported error rate by enhancing sensitivity to a particular noise type, which in our case is $\propto \hat{\sigma}_z$. Alternative sequences may also be calculated that are more sensitive to $\hat{\sigma}_x$ or $\hat{\sigma}_y$ noise than randomly selected RB sequences. These error enhancing sequences give a clear, qualitative signature of the violation of the assumption that the error process is uncorrelated in time, although we do not claim that such a signature is in general uniquely associated with the presence of temporal noise correlations. Furthermore, because calculation of $\|\vec{V}_{2D}\|^2$ and sequence preselection is performed numerically in advance, this approach alleviates the requirement to average extensively in experiment over sequences in order to reveal the skewed fidelity distribution.

Experimental GST in the presence of correlated noise

We now apply the sequence-dependent Pauli walk framework to the default experimental GST gate set in order to understand the interplay of sequence structure and temporal noise correlations in the experimental GST estimation procedures. We begin by collating all standard experimental GST sequences up to 256

gates in length using gates $G_i \equiv \hat{I}$, the identity, G_x , a $\pi/2$ $\hat{\sigma}_x$ rotation and G_y defined similarly. We define sequences to include fiducial operations and germs (see “Methods” and ref.¹⁰), and calculate the corresponding walk lengths. Here we assume unit step size under application of either a constant $\hat{\sigma}_z$ or $\hat{\sigma}_x$ unitary error process (Fig. 4a, b) such that $\|\vec{R}\|^2 = \delta^2 \|\vec{V}\|^2$, and plot $\|\vec{V}_{2D}\|^2$ as a proxy for projected sequence error vs. J . We overlay the results on the calculated probability distribution of unit-step walks for RB sequences, presented as a colour scale for comparison. Points appear clumped due to the experimental GST prescription using different fiducials (leading to different sequence lengths) surrounding a reported germ, as highlighted in Fig. 4b.

Examining these data indicates that GST sequences used in the default package broadly sample the range of expected fidelities in the presence of strongly correlated $\hat{\sigma}_x$ errors, more effectively so than RB. However, their structure appears to systematically suppress measured errors in the presence of correlated $\hat{\sigma}_z$ errors. This mimics the positive skew of RB sequence survival probabilities in the presence of slowly varying noise, as observed in the colour scale. In the presence of correlated $\hat{\sigma}_z$ errors, only GST sequences consisting of repeated G_i germs, formally equivalent to Ramsey experiments,²¹ show sensitivity to this kind of error. We now explore the impact of these observations in further detail by both numerical investigations and experiments involving engineered unitary $\hat{\sigma}_z$ errors.

Given measurement outcomes (experimental or simulated) for the prescribed sequences, the open-source analysis package pyGSTi²² is used to extract a large set of results characterising the performance of the gate set. One important metric calculated by the protocol for each gate is the diamond distance, $\|G_{\text{ideal}} - G_{\text{err}}\|_{\diamond}$, which is meant to provide a worst-case bound on the distance to the ideal gate operation. Experimental GST has found wide adoption in part because of its ability to calculate this metric, which is postulated to be important for formal analyses of fault-tolerance in the context of quantum error correction.

In our first test, we numerically probe the sensitivity of the experimental GST analysis procedure to correlated error using the aforementioned pyGSTi toolkit. We introduce constant $\hat{\sigma}_x$, $\hat{\sigma}_y$, or $\hat{\sigma}_z$ errors via concurrent unitary rotations added to the formerly ideal operations. Therefore the exact mathematical representation of each gate ($G_{i,x,y}$) is known from analytical transformations and we have two paths to evaluate gate performance (Fig. 4c). First, we directly calculate the diamond distance ($\|G_{\text{ideal}} - G_{\text{err}}\|_{\diamond}$) using the matrix representation of G_{err} , maintaining the initial frame of reference. Second, we estimate it by employing pyGSTi to simulate data using G_{err} and determine the diamond distance ($\|G_{\text{ideal}} - G_{\text{err}}^{(\text{est})}\|_{\diamond}$) of the estimate $G_{\text{err}}^{(\text{est})}$ obtained by the toolkit’s fitting routines.

As a self-consistent QCV implementation, the experimental GST estimation procedure incorporates a gauge optimisation by construction, as it makes no assumptions in regard to the qubit and its measurement basis. It performs two rounds of gauge optimisation, allowing identification of a frame in which to minimise the distance of the entire set of estimated gates in relation to the target gates. The relevance of this gauge freedom on RB-derived estimates of gate performance was highlighted recently in²³ To illustrate how gauge freedom affects the results, we separately calculate the diamond distance with and without gauge optimising our analytic gate set G_{err} using routines included in the pyGSTi toolkit.

We plot the calculated and estimated diamond norms for $G_{i,x,y}$ subject to processes similar to either a constant overrotation (i.e. proportional to $\hat{\sigma}_x$ or $\hat{\sigma}_y$ depending on the gate in question, with no error on G_i operations), or a constant detuning error (i.e. proportional to $\hat{\sigma}_z$), as shown Fig. 4d, e. Here we see that the

Fig. 4 Demonstration of GST sensitivity to correlated error models. **a, b** Sensitivity of GST sequences to $\hat{\sigma}_x$, $\hat{\sigma}_z$ errors using the length of the sequence-dependent walk vector $\vec{V}_{2D}$. GST sequence walks are shown as red crosses on a background colour scale illustrating the distribution over 10^6 RB walks and their average (yellow line). Here gates are defined as constituent Clifford operations of length $\tau_{pi/2}$. **c** Flow diagram for the numerical analysis of the diamond norm estimation under correlated errors concurrent with gates G . **d, e** Results of the analysis for the standard gate set G_x, G_y, G_z with the calculated diamond distance shown as solid lines (dashed lines) without (with) gauge optimisation on all graphs, and GST estimation depicted as symbols. Both overrotation errors on the G_x, G_y gates **d** and concurrent detuning errors **e** are studied. For overrotation errors the ideal rotation angle, $\theta \rightarrow (1 + \epsilon)\theta$. **f, g** Analysis is repeated by extending the gate set to include $-G_x, -G_y$. In panels (d) and (f) which employ only overrotation errors, the calculated diamond distance for G_z vanishes and we do not show the noise floor for visual clarity. **h** Experimental investigation of concurrent detuning $\hat{\sigma}_z$ errors via a deliberately engineered detuning Δ . Markers indicate GST estimates from experimental data and solid lines represent analytical calculations performed within the pyGSTi toolkit

estimated diamond distance for operators $G_{i,x,y}$ closely matches the calculated value in the presence of numerical overrotation errors. When used with its standard gate set $\{G_x, G_y, G_z\}$, pyGSTi's estimate of G_x and G_y errors arising from constant unitary $\hat{\sigma}_z$ errors differs significantly, however, and only the diamond norm

estimate for G_z appears similar to the directly calculated value. Other estimated quantities such as process infidelity and the associated Choi matrices are affected in a similar way (see Supplementary Material). However, performing gauge optimisation on the analytically calculated matrices G_{err} as well (within the pyGSTi package) reduces the difference in the reported diamond distance for $\hat{\sigma}_z$ errors, and produces agreement with the much lower $G_{x,y}$ diamond distance reported by the GST estimation procedure (Fig. 4e). Among the error models we have tested, for this gate set such behaviour is only manifested in the presence of temporally correlated $\hat{\sigma}_z$ errors and does not appear using various other error processes built into the pyGSTi analysis package (see Supplementary Material for details). We note that full gauge optimisation is a requirement for self-consistency of results within GST.

To further investigate the influence of the gauge degree of freedom, we repeat our numerical analysis under the application of identical unitary errors, but extend the gate set by adding negative rotations $-G_x, -G_y$ corresponding to $-\pi/2$ $\hat{\sigma}_x$ and $\hat{\sigma}_y$ rotations and incorporating a number of associated compound germs (Fig. 4f, g). The resulting gauge-optimised calculated and estimated diamond-distance values now increase, moving closer to the analytic calculation obtained without gauge optimisation. The behaviour of estimated diamond distances for operations $-G_x$ and $-G_y$ are indistinguishable from those presented to within numerical uncertainty. This simple change in the gate set directly reveals the role of gauge optimisation in the discrepancies we noted above. The additional information now available to experimental GST via the extended gate and germ set effectively constrains the optimisation procedure, allowing it to detect errors that could previously be absorbed in a gauge transformation.

We follow up on these numerical investigations by performing experiments using experimental GST sequences subjected to engineered unitary $\hat{\sigma}_z$ -errors of varying strength. As before, we generate an operation with known error magnitude and form, allowing us to directly produce a matrix representation for the gate and hence calculate the diamond distance for the (deliberately) imperfect gates we apply to our trapped-ion system. Again the experimental GST procedure produces an estimate of the diamond distance that matches the calculation for G_z , but yields estimates of the diamond distance from experimental data approximately an order of magnitude below the (unoptimised) calculated value for $G_{x,y}$ (Fig. 4h). Allowing gauge optimisation on the calculated diamond distance changes its scaling with error magnitude as in simulations above. We do not find strong agreement between data for $G_{x,y}$ and this gauge-optimised scaling, but cannot exclude the possibility that other finite sampling effects may cause saturation of small reported diamond distances.

In addition to the cases presented above we have also performed experimental GST with a wide variety of engineered, time-varying errors. These include detuning and amplitude noise exhibiting 50 Hz fluctuations and slow drifts (*i.e.* varying in time during individual sequences), constant overrotations, and added state-preparation and measurement (SPAM) errors. While these do not form part of this manuscript, they might help inform further studies by other authors in the future. All data sets, corresponding pyGSTi analysis files and resultant reports are included as part of the Supplementary Material.

DISCUSSION

In our studies we have employed a simple analytic framework - a formalism mapping noise to error accumulation in sequences of Clifford operations - to explore the sensitivity of RB and GST to slowly varying noise processes. Theoretical predictions derived from this framework match RB experiments employing engineered noise with known characteristics: either slowly varying or

rapidly varying on the sequence timescale. This highlights the utility of the random-walk analysis in determining sequence-dependent sensitivities of QCVV protocols in the presence of temporally correlated noise.

We have compared RB survival probabilities over sequences to a gamma distribution $\Gamma(\alpha = 1, \beta)$, where β is determined by the type of error model employed in the experiment, and shown good agreement using no free parameters. In addition we have demonstrated that in the presence of slowly varying noise, the mode of the distribution of survival probabilities over sequences is shifted towards lower error rates than the mean and that a long tail of high-error outcomes appears as predicted in.¹⁷

Overall, the experiments reported here give a clear experimental signature of the violation of the assumption that errors between gates are independent. While we do not claim that the features we observe are in general uniquely derived from this interpretation, we hope these results may help experimentalists seeking to interpret complex RB data sets. We believe that more detailed reporting of RB outcomes including the publication of distributions of the survival probability $\mathcal{P}$, as well as the sequences employed, will facilitate more meaningful comparisons between RB data sets derived from different physical systems, as the relevance of p_{RB} is diminished when error processes exhibit temporal correlations.

Through a combination of analytic calculations, numerics, and experiments with engineered errors we have found a similar bias towards lower estimates of diamond distance in experimental GST procedures when using the standard $G_{I,x,y}$ gate set subjected to strongly correlated, unitary $\hat{\sigma}_z$ errors. The asymmetry we observe between the manifestation of correlated $\hat{\sigma}_x/\hat{\sigma}_y$ and $\hat{\sigma}_z$ error-sensitivity has previously only been reported in the context of RB.²³ We have shown explicitly how the low diamond-distance estimates under this kind of noise are related to the gauge optimisation performed as part of the protocol; limiting the gauge freedom by extending the gate set under application of an identical error process dramatically changed the estimated diamond distance of the very same gates in numerical simulations. This highlights that estimates are always reported up to an implicit gauge degree of freedom, making absolute comparisons of diamond norms challenging.

These observations are commensurate with a simple physical interpretation of the effect of an optimised gauge transformation in the circumstances we examine. In the presence of correlated $\hat{\sigma}_z$ errors, when the gate set is limited to $G_{I,x,y}$ gates, the reconstructed operator includes an extra error component along the z-axis. The effect of gauge optimisation is to rotate the axis of rotation of the G_x and G_y operators back to the equatorial plane, effectively cancelling this error. Under this circumstance the magnitude of rotation of these gates is smaller than expected in a fixed lab frame, and the second-order nature of the residual errors result in a steeper gradient of the dotted line in Fig. 4e. In contrast the G_z rotation should have no net rotation and therefore this error will not be cancelled by a simple gauge transformation.

Gauge optimisation is designed to produce the best estimate for errors over the entire gate set in relation to a given target, and in a sense acts to “distribute” nominal errors over all constituent rotations in the gate set. The validity of such a gauge transformation in the presence of independent protocols for establishing a measurement basis remains an open question and has been highlighted recently by Rudnicki et al.²⁴ The variation of calculated and estimated diamond distances under correlated $\hat{\sigma}_z$ errors when subjected to seemingly small modifications of the default gate set has again not been reported previously in the context of GST, and indicates an important dependence of its output on the specific gate set employed, the characteristics of the underlying error source, and the gauge optimisation procedure.

Clearly the observed performance of experimental GST in the presence of correlated $\hat{\sigma}_x$ noise, such as resulting from experimental overrotations, can make GST a valuable tool in debugging an experimental system,²⁵ although precise calibrations can also be carried out efficiently using a subset of the full experimental GST protocol.²⁶ The effect of gauge optimisation in the presence of $\hat{\sigma}_z$ errors and with use of the default gate set, however, is concerning as a key implied benefit of experimental GST is its ability to provide a rigorous upper bound on gate errors using a fully self-contained analysis package. Recent experimental work¹⁰ on the topic claimed such upper bounds on gate errors using experimental GST and compared these to the fault-tolerance threshold with high reported confidence and tight uncertainties. The results above and observations made²⁴ suggest that there may be residual uncertainty in interpretation of such data due to the potential unresolved conflict between full gauge freedom and the nominal existence of a measurement basis constraining that freedom. Furthermore, when acquiring and evaluating data, care has to be taken to suppress any form of model violations reported by the GST toolkit in its likelihood analysis, as otherwise the extracted performance metrics may become unreliable. These deviations are currently not reflected in the uncertainties (*i.e.* error bars) calculated for those metrics by the toolkit and discussions with its authors suggest that a connection between the two is a non-trivial process.

In light of the investigations reported here, we believe that there is a need for greater awareness of the subtleties of the use of both RB and experimental GST in the presence of temporally correlated noise environments. In order to enhance the meaning and utility of reported results we advocate that QCVV benchmarks such as p_{RB} and experimental GST diamond distances should be reported together with a quantitative measure of violation from a purely Markovian, temporally uncorrelated model. In the case of RB, this could be the difference between the extracted p_{RB} of long and short walk sequences; in experimental GST the deviation is already being reported as part of the routine, yet the question about the impact of gauge optimisation that we identified remains. Similarly, if using experimental GST as a standalone gate evaluation procedure one cannot know a priori the form of the underlying noise - and hence any associated experimental GST insensitivities. Increasing the rigour of resultant upper bounds on diamond distances could require performing experimental GST using multiple different gate sets in order to identify potential “blind spots”, owing to the implicitly required gauge transformations. Given the experimental overhead, however, this brute force approach is not necessarily attractive and further modifications to experimental GST could resolve the issue with considerably greater efficiency. Overall, we hope that these observations will assist in both the interpretation of QCVV experiments when model violation may occur, and the development of new techniques with improved rigour and efficiency for larger scale systems.

METHODS

Experimental gate implementation

Quantum gates are implemented on a single $^{171}\text{Yb}^+$ ion by driving its qubit transition at 12.6 GHz with microwave pulses produced by a vector signal generator (VSG, model Keysight E8267D). The phase of the driving field is adjustable via I-Q modulation allowing us to implement rotations around any axis lying in the xy-plane of the Bloch sphere. Rotations around the z-axis are carried out as frame-updates, *i.e.* pre-calculated, instantaneous changes of the generator I-Q values. Identity operations are realised as idle periods, whereby no signals are applied for a time equivalent to that of a π or $\pi/2$ rotation. We additionally implement pulse modulation (RF blanking) to suppress transients in microwave power at pulse edges. In this way, we implement the full set of Clifford gates as listed in supplementary materials.

All RB and GST sequences are uploaded to the VSG prior to the experiments and selected when required. When the number of implemented sequences is large, as is the case with GST, the latter step is the bottleneck in our experiments as sequence selection, depending on the constituent number of gates J , can take up to tens of seconds using our signal generator due to the use of the in-built, high-suppression, RF blanking switch which adds significant overhead.

Experimental noise implementation

In RB experiments correlated noise is implemented by shifting the VSG drive frequency by a fixed amount based on a list of $N = 200$ samples from a Gaussian noise distribution (see Supplementary Material). The same list of noise realisations is repeated for each RB sequence in a set of given length J , yielding sets of noise-averaged fidelities. In GST experiments we implement constant noise of the same strength over all the sequences. In addition to a baseline experiment, only a small set of noise detunings are implemented due to the large overhead imposed by sequence selection prior to execution.

Rapidly varying noise in RB is implemented via the VSG's external frequency modulation, whereby the frequency offset is encoded as a series of calibrated offset voltages on an arbitrary waveform generator (Keysight 33622A) and supplied time-synchronous to each gate within a sequence. Again, $N = 200$ different realisations, each consisting of J samples are applied to each RB sequence to extract a noise-averaged fidelity. Further details can be found in the Supplementary Material.

Concurrent noise model and gate-dependent errors

Deliberately induced $\hat{\sigma}_z$ errors are implemented via a fixed detuning Δ from the qubit's transition frequency, which is tracked by regularly spaced Ramsey experiments to better than 1 Hz accuracy relative to a Rabi frequency of $\Omega = 22.5$ kHz. We apply noisy gates in which a concurrent $\hat{\sigma}_z$ rotation modifies the unitary evolution of our physically implemented gates (only $\hat{\sigma}_I \equiv \hat{I}$, $\hat{\sigma}_x$ and $\hat{\sigma}_y$) given by matrix-exponentiation of the corresponding Pauli-matrices $\hat{\sigma}_{\{I,x,y\}}$ as

$$\tilde{U}(\theta, \Delta, \Omega) = \exp \left\{ -i \left(\frac{\theta}{2} \hat{\sigma}_{\{I,x,y\}} + \frac{|\theta| \Delta}{2 \Omega} \hat{\sigma}_z \right) \right\}. \quad (2)$$

The first term in the exponential corresponds to the unperturbed unitary where the rotation angle θ is chosen to be either $\theta = \pm\pi$ or $\theta = \pm\pi/2$. Here the effective error magnitude scales in relation to the Rabi frequency Ω , and the absolute value of θ ensures that the sign of the detuning term is preserved under positive and negative gate rotations.

This implementation leads to gate-dependent errors. Hence π rotations accumulate twice the phase in the presence of a nonzero Δ as $\pi/2$ rotations.

Gate Set Tomography

Initial F_α and final F_β fiducial operations are taken from the set $\{\emptyset, G_x, G_y, G_x G_x, G_x G_x G_x, G_y G_y G_y\}$, where $\emptyset$ stands for no gate operation, and G_x and G_y stand for $\pi/2$ rotations around the x and y -axes of the Bloch sphere. They are chosen to form an informationally complete set of input states and measurement bases akin to quantum process tomography. The germs used in our experiments are

$$\begin{aligned} &(G_x), (G_y), (G_I), (G_x G_y), \\ &(G_x G_y G_I), (G_x G_I G_y), (G_x G_I G_I), (G_y G_I G_I), \\ &(G_x G_x G_I G_y), (G_x G_y G_y G_I), \\ &(G_x G_x G_y G_x G_y G_y), \end{aligned}$$

identical to those used in reference¹⁰ and recommended as standard GST in the pyGSTi tutorials. In our numerical analysis, we extend the standard gate set from $\{G_I, G_x, G_y\} \rightarrow \{G_I, G_x, G_y, -G_x, -G_y\}$ while also expanding the germ set from 11 to 39 elements to maintain amplification completeness (see Supplementary Material for details). Each of these germs is concatenated with itself up to a maximum length that successively increases as $L = \{1, 2, 4, 8, 16, 32, 64, 128, 256\}$ and measured in all 36 combinations of the fiducials F_α and F_β . In the experimental implementation, we first record a baseline measurement without added error and then step through the cases of added detunings $\Delta = \{75, 500, 1000, 1400\}$ Hz for all 2737 sequences of the standard set. Due to overhead associated with switching between sequences, we recorded 220 repetitions for each sequence in consecutive order. The toolkit's authors advise to instead

interleave sequences and repetitions to spread slow drifts across the data set in order to reduce model violations in the fitting routines (Erik Nielsen, private communications 2017).

Data availability

Data and analysis files in addition to those included in the Supplementary Material are available from the authors on request.

ACKNOWLEDGEMENTS

The authors acknowledge discussions with R.-Blume Kohout, P. Maunz, K. C. Young, and E. Nielsen on GST, and R. Harper, C. Ferrie, and C. Granade for discussions on data analysis. We are grateful to J. Emerson to pointing out the potential utility of adding negative rotations to the GST gate set. Work partially supported by the ARC Centre of Excellence for Engineered Quantum Systems CE110001013, the Intelligence Advanced Research Projects Activity (IARPA) through the US Army Research Office, and a private grant from H. & A. Harley.

AUTHOR CONTRIBUTIONS

S.M. and C.E. led experimental implementation, data collection, and data analysis for RB. C.H. led experimental design, simulations, and analysis of GST. S.M., C.E. and C.H. jointly produced the figures. H.B., F.R., and T.M.S. performed theoretical analyses and calculations. M.J.B. conceived the general direction of this study, oversaw experimental design, and led writing of the manuscript and production of figures.

ADDITIONAL INFORMATION

Supplementary Information accompanies the paper on the *npj Quantum Information* website (<https://doi.org/10.1038/s41534-017-0052-0>).

Competing interests: The authors declare that they have no competing financial interests.

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

REFERENCES

- Emerson, J., Alicki, R. & Życzkowski, K. Scalable noise estimation with random unitary operators. *J. Opt. B Quantum Semiclassical Opt.* **7**, S347–S352 (2005).
- Knill, E. et al. Randomized benchmarking of quantum gates. *Phys. Rev. A* **77**, 012307 (2008).
- Wallman, J., Granade, C., Harper, R. & Flammia, S. T. Estimating the coherence of noise. *New J. Phys.* **17**, 113020 (2015).
- Poyatos, J., Cirac, J. & Zoller, P. Complete characterization of a quantum process: the two-bit quantum gate. *Phys. Rev. Lett.* **78**, 390–393 (1997).
- Chuang, I. L. & Nielsen, M. A. Prescription for experimental determination of the dynamics of a quantum black box. *J. Mod. Opt.* **44**, 2455–2467 (1997).
- Holzäpfel, M., Baumgratz, T., Cramer, M. & Plenio, M. B. Scalable reconstruction of unitary processes and Hamiltonians. *Phys. Rev. A* **91**, 042129 (2015).
- Flammia, S., Gross, D., Liu, Y. K. & Eisert, J. Quantum tomography via compressed sensing: error bounds, sample complexity and efficient estimators. *New J. Phys.* **14**, 095022 (2012).
- Granade, C., Ferrie, C. & Flammia, S. T. Practical adaptive quantum tomography. *New J. Phys.* **19**, 113017, <https://doi.org/10.1088/1367-2630/aa8fe6> (2017).
- Merkel, S. T. et al. Self-consistent quantum process tomography. *Phys. Rev. A* **87**, 062119 (2013).
- Blume-Kohout, R. et al. Demonstration of qubit operations below a rigorous fault tolerance threshold with gate set tomography. *Nat. Commun.* **8**, 14485 (2017).
- Riebe, M. et al. Process tomography of ion trap quantum gates. *Phys. Rev. Lett.* **97**, 220407 (2006).
- Rutman, J. Characterization of phase and frequency instabilities in precision frequency sources: fifteen years of progress. *Proc. IEEE* **66**, 1048–1075 (1978).
- Hooge, F. N., Klempnering, T. G. M. & Vandamme, L. K. J. Experimental studies on 1/f noise. *Rep. Prog. Phys.* **44**, 479 (1981).
- Harlingen, D. J. V., Plourde, B. L. T., Robertson, T. L., Reichardt, P. A. & Clarke, J. In *Quantum Computing and Quantum Bits in Mesoscopic Systems* (eds Leggett, A., Ruggiero, B. & Silvestini, P.) 171–184 (New York, NY, Kluwer Academic Press, 2004).
- Fogarty, M. A. et al. Nonexponential fidelity decay in randomized benchmarking with low-frequency noise. *Phys. Rev. A* **92**, 022326 (2015).
- Wallman, J. J. & Flammia, S. T. Randomized benchmarking with confidence. *New J. Phys.* **16**, 103032 (2014).

17. Ball, H., Stace, T. M., Flammia, S. T. & Biercuk, M. J. Effect of noise correlations on randomized benchmarking. *Phys. Rev. A* **93**, 022303 (2016).
18. Fong, B. H. & Merkel, S. T. Randomized benchmarking, correlated noise, and ising models. Preprint at <https://arxiv.org/abs/1703.09747> (2017).
19. Wallman, J. J. Randomized benchmarking with gate-dependent noise. Preprint at <https://arxiv.org/abs/1703.09835> (2017).
20. Kueng, R., Long, D. M., Doherty, A. C. & Flammia, S. T. Comparing experiments to the fault-tolerance threshold. *Phys. Rev. Lett.* **117**, 170502 (2016).
21. Ramsey, N. F. A molecular beam resonance method with separated oscillating fields. *Phys. Rev.* **78**, 695–699 (1950).
22. Nielsen, E., Rudinger, K., Gamble, J. K. & Blume-Kohout, R. *Pygsti: a Python Implementation of Gate Set Tomography* <http://github.com/pygsti> (2016).
23. Proctor, T., Rudinger, K., Young, K., Sarovar, M. & Blume-Kohout, R. What randomized benchmarking actually measures. *Phys. Rev. Lett.* **119**, 58 (2017).
24. Rudnicki, Ł., Puchała, Z. & Życzkowski, K. *Gauge Invariant Information Concerning Quantum Channels* <https://arxiv.org/abs/1707.06926> (2017).
25. Dehollain, J. P. et al. Optimization of a solid-state electron spin qubit using gate set tomography. *New J. Phys.* **18**, 1–9 (2016).
26. Rudinger, K., Kimmel, S., Lobser, D. & Maunz, P. Experimental demonstration of a cheap and accurate phase estimation. *Phys. Rev. Lett.* **118**, 190502 (2017).

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

© The Author(s) 2018