



Theory of versatile fidelity estimation with confidenceAkshay Seshadri *Department of Physics, University of Colorado Boulder, Boulder, Colorado 80309-0390, USA*Martin Ringbauer *Universität Innsbruck, Institut für Experimentalphysik, Technikerstrasse 25, 6020 Innsbruck, Austria*Jacob Spainhour *Department of Applied Mathematics, University of Colorado Boulder, Boulder, Colorado 80309-0526, USA*Thomas Monz *Universität Innsbruck, Institut für Experimentalphysik, Technikerstrasse 25, 6020 Innsbruck, Austria
and Alpine Quantum Technologies GmbH, 6020 Innsbruck, Austria*Stephen Becker *Department of Applied Mathematics, University of Colorado Boulder, Boulder, Colorado 80309-0526, USA*

(Received 18 February 2024; accepted 1 May 2024; published 10 July 2024)

Estimating the fidelity with a target state is important in quantum information tasks. Many fidelity-estimation techniques present a suitable measurement scheme to perform the estimation. In contrast, we present techniques that allow the experimentalist to choose a convenient measurement setting. Our primary focus lies on a method that constructs an estimator with nearly minimax optimal confidence intervals for any specified measurement setting. We demonstrate, through a combination of theoretical and numerical results, various desirable properties of the method: robustness against experimental imperfections, competitive sample complexity, and accurate estimates in practice. We compare this method with maximum likelihood estimation and the associated profile likelihood method, a semi-definite programming based approach, direct fidelity estimation, quantum state verification, and classical shadows. Our method can also be used for estimating the expectation value of any observable with the same guarantees.

DOI: [10.1103/PhysRevA.110.012431](https://doi.org/10.1103/PhysRevA.110.012431)**I. INTRODUCTION**

Quantum information, though a relatively nascent branch of research, has pervaded into many areas of physics and even other subjects like computer science. This prevalence is owed in large part to the wide range of applications that benefit from a quantum approach, such as computation [1], simulation [2], communication [3], and metrology and sensing [4,5]. All of these applications require us to accurately and reliably prepare a desired quantum state. Yet *verifying* that the experimentally prepared state is indeed what we intended is a nontrivial task. A commonly used metric to judge the success of state preparation is the fidelity between the target state ρ and the prepared state σ . When the target state is pure, the fidelity is simply given as $F(\rho, \sigma) = \text{Tr}(\rho\sigma)$ [6]. While we focus on fidelity estimation here, our method can be used to estimate the expectation value of any observable.

A common way to estimate the fidelity is to first reconstruct the experimentally prepared quantum state and then compute the fidelity between the reconstruction and the target state. Quantum state tomography is an active area of research, and several methods have been proposed to reconstruct the quantum state [7–11]. However, these methods suffer from an exponential resource requirement as the Hilbert space grows,

making them infeasible for all but the smallest systems. Particularly, when the task is merely to estimate the fidelity, a method known as direct fidelity estimation (DFE) [12,13] has been shown to achieve the task with exponentially fewer resources. DFE and more recent approaches like classical shadows [14] estimate the fidelity directly from measurement outcomes without going through the intermediate step of reconstructing the state. A related, but slightly different approach, is quantum state verification [15], where one *tests* whether the fidelity is larger than a given threshold. These methods specify measurement protocols, typically involving sampling measurement settings randomly that one needs to follow to estimate the fidelity or perform state verification. In practice, however, experimenters tend to use a fixed set of measurement settings instead of randomly sampling them, which would result in a loss of theoretical guarantees. Furthermore, there can be cases where it is preferable to implement a different measurement scheme than the ones prescribed by these approaches. This is particularly true if one wishes to estimate the expectation value of a single observable, where performing measurements tailored to that observable can give exponential improvements over a method like classical shadows that performs random measurements independent of the observable to be estimated.

However, methods like maximum likelihood estimation (MLE) work for arbitrary measurement schemes, yet, in addition to the need for a costly reconstruction of the quantum state, they do not provide any guarantees on the estimate. Rigorous confidence intervals for estimated fidelities are crucial for applications from quantum error correction [16] to quantum key distribution, where the security of a protocol rests on proving the presence of entanglement [17,18]. Entanglement verification, in turn, is typically achieved using an entanglement witness [19,20], which is intimately tied to fidelity estimation [20], again, highlighting the need for rigorous confidence intervals. Even in cases where state reconstruction with a good confidence interval is possible, such as a bound in trace distance as given in Ref. [21], the number of measurements needed for tomography can be much larger than what is needed for fidelity estimation [12]. One can also use quantum algorithms to estimate the fidelity, for example, Refs. [22,23], which work with a different premise than what we consider in our study. To summarize, we identified two concerns with standard approaches for fidelity estimation: most do not allow for arbitrary measurement settings and/or do not provide rigorous confidence intervals.

Here, we address both of these concerns. We provide a versatile approach for estimating the fidelity directly from raw data for *any* measurement scheme, without the need for state reconstruction. More precisely, given any target state, measurement scheme, and desired confidence level, our method constructs an estimator for the fidelity that takes raw data and gives a fidelity estimate almost instantaneously. Furthermore, the estimator comes equipped with a rigorous confidence interval that is guaranteed to be close to minimax optimal.¹ Turning this around, for a target confidence level $1 - \delta$, our method achieves a confidence interval of size $2\epsilon_*$, where ϵ_* is called the *risk* (or additive error) of the estimate, with a sample complexity of $\approx \ln(2/\delta)/(2\epsilon_*^2)$ independent of the target state ρ when measuring in the basis defined by ρ (see Sec. IID). While such a measurement scheme may often be impractical, we will introduce an alternative scheme using Pauli measurements that achieves a similar sample complexity for stabilizer states. This finding demonstrates that our method can give a scalable sample complexity with a judicious choice of measurement protocol.

We will begin by describing the theory behind our approach in Sec. IIA. Then, we compute the sample complexity of the method in Sec. IID and describe a Pauli measurement scheme similar to DFE (in procedure and sample complexity) in Sec. IIF. In Sec. IIG, we demonstrate the robustness of the estimator generated by our method against noise. Finally, we compare our method with DFE, MLE, profile likelihood, a semi-definite programming approach, quantum state verification, and classical shadows in Sec. III.

II. MINIMAX METHOD

Our approach is based on recent results in statistics by Juditsky and Nemirovski [24–26], who described the esti-

mation of linear functionals in a general setting. The risk associated with the estimate is nearly minimax optimal, and so we call it the minimax method. Roughly speaking, a minimax optimal estimator gives the smallest possible symmetric confidence interval in the worst possible scenario (this will be made precise later).

In Sec. IIA, we elaborate on the main ideas involved in estimating fidelity using the minimax method. An abstract version of the method with the theoretical details is given in Appendix A, while an application-oriented presentation of the key results is given in the accompanying Ref. [27].

A. Fidelity estimation using the minimax method

Suppose that we are working with a quantum system that has a d -dimensional Hilbert space over \mathbb{C} , $d \in \mathbb{N}$. The states of this system can be described by density matrices, which are $d \times d$ positive-semi-definite (complex-valued) matrices with unit trace. We denote the set of density matrices by \mathcal{X} . Our goal is to estimate the fidelity between the state $\sigma \in \mathcal{X}$ prepared in the laboratory and a pure target state $\rho \in \mathcal{X}$, where pure means that ρ is a rank-one matrix. In this case, the fidelity between ρ and σ is given as $F(\rho, \sigma) = \text{Tr}(\rho\sigma)$, which is a linear function of σ for a fixed ρ .

In practice, one needs to estimate the fidelity from partial information about the state obtained through measurements. Any quantum measurement can be described by a positive operator-valued measure (POVM), which comprises a set of positive-semi-definite matrices that sum to the identity. We consider the case where an experimenter uses L different measurement settings, where the POVM for the l th setting ($l = 1, \dots, L$), is described by $\{E_1^{(l)}, \dots, E_{N_l}^{(l)}\}$. The experimenter performs R_l repetitions (shots) of the l th POVM. The probability $p_\sigma^{(l)}(k)$ that a particular outcome $k \in \{1, \dots, N_l\}$ is obtained upon measuring the l th POVM when the state of the system is σ is given by Born's rule as $p_\sigma^{(l)}(k) = \text{Tr}(E_k^{(l)}\sigma)$.

Note that Born's rule can give zero probabilities for some of the outcomes depending on the state and the POVM. However, a technical requirement of Juditsky and Nemirovski's method [24] is that all outcome probabilities are nonzero. For this reason, we include a small positive parameter $\epsilon_o \ll 1$ to make the outcome probabilities positive, i.e.,

$$p_\sigma^{(l)}(k) = \frac{\text{Tr}(E_k^{(l)}\sigma) + \epsilon_o/N_l}{1 + \epsilon_o}.$$

We choose $\epsilon_o = 10^{-5}$ in the numerical simulations so that the Born probabilities are practically unaffected. Because ϵ_o is small with respect to typical experimental imperfections, we will usually drop terms proportional to ϵ_o in the theoretical calculations.

Our goal is to construct an estimator \widehat{F} for the fidelity $F(\rho, \sigma)$ between the experimental state σ and the target state ρ , using the observed outcomes corresponding to the chosen measurement settings. By an estimator, we mean any function that takes the observed outcomes as input and gives an estimate for the fidelity as an output. Constructing a good estimator first requires a measure of error to judge the performance of the estimator. For this purpose, we use the δ risk defined by Juditsky and Nemirovski [24]. Intuitively speaking,

¹Larger, at most, by a factor of order 1 for sufficiently large confidence levels; see Sec. IIA.

having an δ risk of ε means that the error in the estimation is at most ε with a probability larger than $1 - \delta$. Therefore, the δ risk defines a symmetric confidence interval for a confidence level of $1 - \delta$. More precisely, given a fidelity estimator \widehat{F} and a confidence level $1 - \delta$, the δ risk of \widehat{F} is given by

$$\mathcal{R}(\widehat{F}; \delta) = \inf \left\{ \varepsilon' \mid \sup_{\chi \in \mathcal{X}} \text{Prob}_{\text{outcomes} \sim p_{\chi}} [|\widehat{F}(\text{outcomes}) - \text{Tr}(\rho \chi)| > \varepsilon'] < \delta \right\}. \quad (1)$$

Here, “outcomes $\sim p_{\chi}$ ” means that the outcomes for the l th measurement are given by the (ε_{\circ} -modified) Born rule probabilities $p_{\chi}^{(l)}$ for a state χ , for all $l = 1, \dots, L$. The definition of the δ risk says that $\mathcal{R}(\widehat{F}; \delta)$ is the smallest number such that the probability that the estimator \widehat{F} has an error larger than $\mathcal{R}(\widehat{F}; \delta)$ is less than δ , irrespective of the underlying state of the system.

Importantly, the δ risk for any given estimator \widehat{F} can be precomputed, even before performing a measurement. This is possible because it depends only on the target state, the confidence level, and the chosen measurement settings, while implicitly accounting for all possible measurement outcomes in all possible states. We can thus construct estimators that nearly achieve the minimum possible risk before taking any data.

This minimum possible risk $\mathcal{R}_{*}(\delta)$ for a chosen confidence level is called the *minimax optimal risk*, and is obtained by minimizing $\mathcal{R}(\widehat{F}; \delta)$ over all possible estimators \widehat{F} , i.e., $\mathcal{R}_{*}(\delta) = \inf_{\widehat{F}} \mathcal{R}(\widehat{F}; \delta)$. Therefore, the minimax optimal risk gives the smallest possible error by minimizing over all the estimators \widehat{F} , while maximizing over all the states $\chi \in \mathcal{X}$ [see Eq. (1)].

In practice, however, it is computationally very difficult to sift through all possible estimators and choose the optimal one. For this purpose, Juditsky and Nemirovski [24] restricted their attention to a subset \mathcal{F} of possible estimators called affine estimators. Any affine estimator $\phi \in \mathcal{F}$ is of the form $\phi = \sum_{l=1}^L \phi^{(l)}$, where $\phi^{(l)}$ is an estimator for the l th measurement setting. That is, $\phi^{(l)}$ takes an outcome of the l th POVM as input and returns a number as the output. Remarkably, Juditsky and Nemirovski [24] showed how to construct an affine estimator that achieved nearly minimax optimal performance, and therefore restricting attention to affine estimators was not a problem. Specifically, if $\widehat{F}_{*} \in \mathcal{F}$ is the near-optimal affine estimator constructed by Juditsky and Nemirovski’s procedure, then the risk $\varepsilon_{*}(\delta)$ of \widehat{F}_{*} computed by their procedure satisfies

$$\mathcal{R}(\widehat{F}_{*}; \delta) \leq \varepsilon_{*}(\delta) \leq \vartheta(\delta) \mathcal{R}_{*}(\delta), \quad (2)$$

$$\vartheta(\delta) = \frac{2 \ln(2/\delta)}{\ln[1/(4\delta)]} = 2 + \frac{\ln(64)}{\ln(0.25/\delta)} \quad (3)$$

for $\delta \in (0, 0.25)$. We restrict δ to the interval $(0, 0.25)$ (or equivalently, to confidence levels greater than 75%) so that $\vartheta(\delta)$ is well defined. From Eq. (2), we see that the risk $\varepsilon_{*}(\delta)$ computed by Juditsky and Nemirovski’s procedure is actually an upper bound on the (unknown) δ -risk $\mathcal{R}(\widehat{F}_{*}, \delta)$ of

TABLE I. Different types of risk considered in this study. In each case, the true state is unknown, the measurement settings are fixed, and all measurement outcomes consistent with the state and measurement settings are allowed. \widehat{F}_{*} is the near-optimal affine estimator constructed by Juditsky and Nemirovski’s [24] procedure.

$\mathcal{R}_{*}(\delta)$	Minimax optimal risk, which is the minimum possible additive error with a confidence level of $1 - \delta$, accounting for all possible states, measurement outcomes, and all possible estimators.
$\mathcal{R}(\widehat{F}_{*}; \delta)$	δ risk of the estimator \widehat{F}_{*} as defined in Eq. (1). This is the smallest possible additive error for the estimator \widehat{F}_{*} with a confidence level of $1 - \delta$, accounting for all possible states and measurement outcomes.
$\varepsilon_{*}(\delta)$	Risk (or additive error) of the estimator \widehat{F}_{*} for a confidence level of $1 - \delta$ that is computed by Juditsky and Nemirovski’s procedure. This is an upper bound on the δ -risk $\mathcal{R}(\widehat{F}_{*}; \delta)$ and it is nearly minimax optimal in the sense of Eq. (2).

the near-optimal affine estimator \widehat{F}_{*} . Since the probability that the estimator \widehat{F}_{*} fails by an error more than $\mathcal{R}(\widehat{F}_{*}, \delta)$ is less than δ , and because $\mathcal{R}(\widehat{F}_{*}, \delta) \leq \varepsilon_{*}(\delta)$ we can conclude that the probability that \widehat{F}_{*} fails by more than $\varepsilon_{*}(\delta)$ is also less than δ . Therefore, the risk $\varepsilon_{*}(\delta)$ defines a confidence interval corresponding to the chosen confidence level $1 - \delta$.

Equation (2) guarantees that the confidence interval defined by $\varepsilon_{*}(\delta)$ is nearly minimax optimal. This is because the risk $\varepsilon_{*}(\delta)$ of the near-optimal affine estimator \widehat{F}_{*} computed by Juditsky and Nemirovski’s procedure is, at most, a constant times the minimax optimal risk, where the constant factor $\vartheta(\delta)$ depends only on the chosen confidence level. Note that $\vartheta(0.1) < 6.54$ and that $\vartheta(\delta)$ is a decreasing function of δ , converging to 2 as $\delta \rightarrow 0$. Therefore, for confidence levels greater than 90%, ε_{*} is guaranteed to be close to the minimax optimal risk by a factor less than 6.5. In practice, the estimator typically performs better than the theoretically guaranteed bound of Eq. (2). We use the star symbol as a subscript in $\varepsilon_{*}(\delta)$ to emphasize that it is nearly minimax optimal.

In the remainder of this study, we omit the argument δ when writing ε_{*} for the sake of notational simplicity. A short summary of the different types of risk mentioned in this section can be found in Table I.

A practically implementable version [28] of the procedure outlined by Juditsky and Nemirovski [24] for constructing the near-optimal affine estimator \widehat{F}_{*} and its associated risk ε_{*} is as follows.

(1) Find the saddle-point value of the function $\Phi : (\mathcal{X} \times \mathcal{X}) \times (\mathcal{F} \times \mathbb{R}_{+}) \rightarrow \mathbb{R}$ defined as

$$\begin{aligned} \Phi(\chi_1, \chi_2; \phi, \alpha) &= \text{Tr}(\rho \chi_1) - \text{Tr}(\rho \chi_2) + 2\alpha \ln(2/\delta) \\ &+ \alpha \sum_{l=1}^L R_l \left[\ln \left(\frac{\sum_{k=1}^{N_l} e^{-\phi_k^{(l)}/\alpha} \text{Tr}(E_k^{(l)} \chi_1) + \varepsilon_{\circ}/N_l}{1 + \varepsilon_{\circ}} \right) \right. \\ &\left. + \ln \left(\frac{\sum_{k=1}^{N_l} e^{\phi_k^{(l)}/\alpha} \text{Tr}(E_k^{(l)} \chi_2) + \varepsilon_{\circ}/N_l}{1 + \varepsilon_{\circ}} \right) \right] \quad (4) \end{aligned}$$

to a given precision, where ρ is the target state; $\chi_1, \chi_2 \in \mathcal{X}$ are density matrices; $\phi \in \mathcal{F}$ is an affine estimator; and $\alpha > 0$ is a positive number. The number α as such has no intuitive meaning, and enters the function Φ for mathematical reasons. The function Φ is itself a mathematical device designed by Juditsky and Nemirovski to compute the near-optimal affine estimator and the associated risk.

Recall that any affine estimator $\phi \in \mathcal{F}$ can be written as $\phi = \sum_{l=1}^L \phi^{(l)}$. Here, $\phi^{(l)}$ is a function taking outcomes of the l th POVM as input, and giving a real number as an output. The l th POVM has N_l measurement outcomes, and $\phi_k^{(l)}$ denotes the output of $\phi^{(l)}$ for the k th measurement outcome, for each $k = 1, \dots, N_l$.

Juditsky and Nemirovski [24] proved that the function Φ has a well-defined saddle-point value (see Appendix A for more details). We present an algorithm to compute the saddle-point value of Φ to a given precision using convex optimization in Appendix B.

(2) We denote the saddle-point value of Φ by $2\varepsilon_*$, i.e.,

$$\begin{aligned} \varepsilon_* &= \frac{1}{2} \sup_{\chi_1, \chi_2 \in \mathcal{X}} \inf_{\phi \in \mathcal{F}, \alpha > 0} \Phi(\chi_1, \chi_2; \phi, \alpha) \\ &= \frac{1}{2} \inf_{\phi \in \mathcal{F}, \alpha > 0} \max_{\chi_1, \chi_2 \in \mathcal{X}} \Phi(\chi_1, \chi_2; \phi, \alpha). \end{aligned} \quad (5)$$

Suppose that the saddle-point value is attained at $\chi_1^*, \chi_2^* \in \mathcal{X}$, $\phi_* \in \mathcal{F}$, and $\alpha_* > 0$ to within the given precision. Since $\phi_* \in \mathcal{F}$ is an affine estimator, we can write $\phi_* = \sum_{l=1}^L \phi_*^{(l)}$. Suppose that independent and identically distributed outcomes $\{o_1^{(l)}, \dots, o_{R_l}^{(l)}\}$ are observed upon measurement of the l th POVM. Then, the optimal estimator $\widehat{F}_* \in \mathcal{F}$ for estimating the fidelity of the state prepared in the laboratory with the target

state is given as

$$\widehat{F}_*({o_1^{(l)}, \dots, o_{R_l}^{(l)}}_{l=1}^L) = \sum_{l=1}^L \sum_{r=1}^{R_l} \phi_*^{(l)}(o_r^{(l)}) + c, \quad (6)$$

where the constant c is

$$c = \frac{1}{2} [\text{Tr}(\rho \chi_1^*) + \text{Tr}(\rho \chi_2^*)]. \quad (7)$$

The procedure outputs \widehat{F}_* and ε_* , such that $\mathcal{R}(\widehat{F}_*; \delta) \leq \varepsilon_*$. Then, from Eq. (1), we can infer that $|\widehat{F}_*(\text{outcomes}) - \text{Tr}(\rho \sigma)| \leq \varepsilon_*$ with probability greater than or equal to $1 - \delta$, where σ is the actual state prepared in the laboratory. The guarantees given by Eqs. (2) and (3) apply.

Observe that \widehat{F}_* is a function of the outcomes [see Eq. (6)], while, the risk ε_* of the estimator is a function of the measurement protocol. We thus know how well the constructed estimator does even before performing an experiment, which can be used for benchmarking different measurement protocols.

It is not apparent from the above procedure how finding the saddle point of the function Φ defined in Eq. (4) leads to constructing the fidelity estimator \widehat{F}_* with error ε_* . Fortunately, we can get some mathematical intuition on how this algorithm works by rewriting the function Φ in Eq. (4), following the proof of Juditsky and Nemirovski [24]. For simplicity, we restrict our attention to the case where we record a single outcome of a single POVM (i.e., $L = 1$ and $R_1 = 1$), and the more general case can be handled similarly. Denoting the outcome probability as $p_k^{(1)}(\chi) = [\text{Tr}(E_k^{(1)} \chi) + \epsilon_o] / (1 + \epsilon_o)$, we can rearrange terms in Eq. (4) to obtain

$$\begin{aligned} \frac{\Phi(\chi_1, \chi_2; \phi_*, \alpha_*) - 2\varepsilon_*}{\alpha_*} &= \left[\ln \left(\sum_{k=1}^{N_1} \exp([\text{Tr}(\rho \chi_1) - (\phi_{*k}^{(1)} + c) - \varepsilon_*] / \alpha_*) p_k(\chi_1) \right) + \ln(2/\delta) \right] \\ &+ \left[\ln \left(\sum_{k=1}^{N_1} \exp([-\text{Tr}(\rho \chi_2) + (\phi_{*k}^{(1)} + c) - \varepsilon_*] / \alpha_*) p_k(\chi_2) \right) + \ln(2/\delta) \right]. \end{aligned}$$

Since $\max_{\chi_1, \chi_2} \Phi(\chi_1, \chi_2; \phi_*, \alpha_*) = 2\varepsilon_*$ to within some precision, we have $\Phi(\chi_1, \chi_2; \phi_*, \alpha_*) - 2\varepsilon_* \leq 0$ for all χ_1, χ_2 . Then, noting that $\widehat{F}_*(k) = \phi_{*k}^{(1)} + c$ (for one measurement setting), the choice of c in Eq. (7) ensures that both the terms in square brackets above are less than or equal to zero [24]. In other words, for all χ_1, χ_2 , we have

$$\begin{aligned} \sum_{k=1}^{N_1} \exp([\text{Tr}(\rho \chi_1) - \widehat{F}_*(k) - \varepsilon_*] / \alpha_*) p_k(\chi_1) &\leq \delta/2, \\ \sum_{k=1}^{N_1} \exp([-\text{Tr}(\rho \chi_2) + \widehat{F}_*(k) - \varepsilon_*] / \alpha_*) p_k(\chi_2) &\leq \delta/2, \end{aligned}$$

where by $\widehat{F}_*(k)$ we mean \widehat{F}_* evaluated on the k th outcome. However, since $e^x \geq 1$ for $x \geq 0$ and $\delta \in (0, 0.25)$, for the above inequalities to be true, we must have $\text{Tr}(\rho \chi) - \widehat{F}_* - \varepsilon_* \leq 0$ and $-\text{Tr}(\rho \chi) + \widehat{F}_* - \varepsilon_* \leq 0$ with high probability for any state χ . This is equivalent to saying $|\text{Tr}(\rho \chi) - \widehat{F}_*| \leq \varepsilon_*$ with high probability for all states χ . Note that this idea

can be formalized using Markov's inequality. Therefore, \widehat{F}_* gives a valid estimate for the fidelity $\text{Tr}(\rho \chi)$ within error ε_* with high probability, no matter what the underlying state χ is.

Although the formal procedure above is rather abstract, in Sec. II C, we show that the estimator \widehat{F}_* can be understood as an appropriately weighted sum of the observed frequencies, where the weights depend on the target state and measurement scheme. In Sec. II D, we show that the risk ε_* can be written in terms of classical fidelities determined by the POVMs in the measurement protocol. This is helpful in performing theoretical calculations involving the risk, and potentially useful for benchmarking measurement protocols.

Importantly, the above procedure cannot only be used to construct an estimator for the fidelity, but for the expectation value of any observable. This can be achieved by replacing the target state ρ in the above equations with the Hermitian operator \mathcal{O} corresponding to the observable whose expectation value we wish to estimate.

We discuss the theoretical underpinnings of the minimax method in Appendix A, starting with a brief overview of Juditsky and Nemirovski's general setup [24], followed by details on adapting their method to estimating fidelity and expectation values. Subsequently, details regarding the numerical implementation of the minimax method are given in Appendix B. In particular, we outline a convex optimization algorithm for constructing the optimal estimator \widehat{F}_* and the associated risk ε_* given any target state and measurement settings.

B. Toy problem: One-qubit target state

We explain in detail a simple setting where it is possible to understand the estimator that is constructed by the procedure described above. Let the target state be $\rho = |1\rangle\langle 1|$, where $\{|0\rangle, |1\rangle\}$ are the eigenstates of Pauli Z . Suppose that our experiment consists of performing Pauli Z measurements, and we want to estimate the fidelity of the laboratory state σ with the target state ρ . This problem is essentially classical: given the quantum state σ , we can write it in the computational basis as $\sigma = (1-p)|0\rangle\langle 0| + q|0\rangle\langle 1| + q^*|1\rangle\langle 0| + p|1\rangle\langle 1|$. Since we measure Z , we will obtain the outcome $|0\rangle$ with a probability of $1-p$ and $|1\rangle$ with a probability of p , and the problem is to estimate the fidelity $\text{Tr}(\rho\sigma) = p$. This is the same as having a Bernoulli random variable with parameter p , which we want to estimate. Note that the parameter p corresponds to the probability of observing the outcome 1 for the Bernoulli random variable.

To numerically compute the fidelity estimator, we choose a confidence level of 95% and $R = 100$ repetitions of Pauli Z measurement. Since we measure only one POVM, the estimator given in Eq. (6) can be written as $\widehat{F}_*({o_1, \dots, o_R}) = \sum_{i=1}^R \phi_*(o_i) + c$, where $o_1, \dots, o_R \in \{0, 1\}$ are the measurement outcomes. Here, $\phi_* \in \mathcal{F}$ is an affine estimator corresponding to the saddle point of the function Φ [see Eq. (4)] that accepts a measurement outcome (0 or 1) as input and gives a number as an output, while the constant c is as given in Eq. (7). Numerically, we find that the affine estimator ϕ_* gives the values $\phi_*(0) \approx -0.476 \times 10^{-2}$ and $\phi_*(1) \approx 0.476 \times 10^{-2}$ for the measurement outcomes 0 and 1, respectively. The value of the constant c computed numerically is $c = 0.5$.

Now we show that the estimator \widehat{F}_* constructed by the minimax method is essentially the sample mean estimator for a Bernoulli random variable, but with a small additive constant that pushes the estimate away from the boundary of the parameter space. Note that the sample mean estimator is also the maximum likelihood estimator for estimating the parameter of a Bernoulli random variable. Recall that our estimator for fidelity is given as $\widehat{F}_*({o_1, \dots, o_R}) = \sum_{i=1}^R \phi_*(o_i) + c$ for a given set of outcomes $\{o_i\}_{i=1}^R$. We can then write the affine estimator ϕ_* as $\phi_*(x) = [\phi_*(1) - \phi_*(0)]x + \phi_*(0)$ for any outcome $x \in \{0, 1\}$. Therefore,

$$\begin{aligned} \widehat{F}_*({o_i}_{i=1}^R) &= \sum_{i=1}^R \{[\phi_*(1) - \phi_*(0)]o_i + \phi_*(0)\} + c \\ &= [\phi_*(1) - \phi_*(0)] \sum_{i=1}^R o_i + [R\phi_*(0) + c], \end{aligned}$$

where $o_i \in \{0, 1\}$. Using the values obtained numerically, we find that

$$\widehat{F}_*({o_i}_{i=1}^{100}) = \frac{0.952}{100} \sum_{i=1}^{100} o_i + 0.024.$$

Thus, we can interpret \widehat{F}_* as (approximately) the mean of the sample o_1, \dots, o_{100} , but with a small additive constant of 0.024. For finite sample sizes the constant term is justified because, even upon observing all 0 outcomes, we cannot be certain that the probability p of observing outcome 1 is $p = 0$. For a similar reason, we have $0.952/100$ as the coefficient for the sum instead of $1/100$. We observe numerically that this coefficient approaches $1/R$ and the additive constant goes to 0 as R increases.

What if we decide to measure Pauli X (measurement 2) in addition to Pauli Z (measurement 1)? In this case, the estimator \widehat{F}_* can be written as $\widehat{F}_*({o_1, \dots, o_R}) = \sum_{i=1}^R \phi_*^{(1)}(o_i) + \sum_{i=1}^R \phi_*^{(2)}(o_i) + c$. Here, $\phi_*^{(1)}$ is the affine estimator at the saddle point corresponding to the Pauli Z measurement, while $\phi_*^{(2)}$ is the affine estimator at the saddle point corresponding to the Pauli X measurement. In such a scenario, we find that the saddle-point value corresponding to the Pauli Z measurement is unchanged. That is, $\phi_*^{(1)}(0) \approx -0.476 \times 10^{-2}$, $\phi_*^{(1)}(1) \approx 0.476 \times 10^{-2}$, and the constant $c = 0.5$ is also the same as before. However, the saddle-point value corresponding to Pauli X measurement is $\phi_*^{(2)}(0) = \phi_*^{(2)}(1) = 0$. Therefore, the outcomes from the Z measurement are weighted as before, but those from X measurement are discarded. The reason is simple: for estimating the fidelity with $\rho = |1\rangle\langle 1|$, measurement of Pauli X gives no useful information, irrespective of what the actual state σ is. Indeed, upon measuring X , we get $|+\rangle$ with probability $1/2 + \text{Re}(q)$ and $|-\rangle$ with probability $1/2 - \text{Re}(q)$, both of which are independent of $\text{Tr}(\rho\sigma) = p$ that we want to estimate. Thus, our method properly incorporates the available measurements to give an estimator for fidelity.

C. Affine estimator for fidelity

An affine function is a linear function plus a constant term. In the previous section, we saw that the estimator for the one-qubit toy problem can be expressed as an affine function of the sample mean. In this section, we show that such a result holds more generally. Our estimator can be written as an affine function of the observed frequencies.

To begin, recall that the outcomes of the l th measurement setting is described by the set $\Omega^{(l)} = \{1, \dots, N_l\}$, where the index k corresponds to the POVM element $E_k^{(l)}$. Corresponding to each outcome $k \in \Omega^{(l)}$, we associate a canonical basis element $e_k^{(l)} \in \mathbb{R}^{N_l}$. From Eq. (6), we know that the fidelity estimator is given as

$$\widehat{F}_*({o_1^{(l)}, \dots, o_{R_l}^{(l)}}_{l=1}^L) = \sum_{l=1}^L \sum_{r=1}^{R_l} \phi_*^{(l)}(o_r^{(l)}) + c,$$

where $\{o_1^{(l)}, \dots, o_{R_l}^{(l)}\}$ are the outcomes corresponding to the l th measurement setting, and the constant c is given by Eq. (7). In the form written above, $\phi_*^{(l)} \in \mathcal{F}^{(l)}$ are real-valued functions defined on the set $\Omega^{(l)}$. Regarding $\phi_*^{(l)}$ as an

N_l -dimensional real vector with ‘‘coefficient’’ vector

$$\mathbf{a}^{(l)} = (\phi_*^{(l)}(1), \dots, \phi_*^{(l)}(N_l)) \in \mathbb{R}^{N_l},$$

we can write $\phi_*^{(l)}(k) = \langle \mathbf{a}^{(l)}, \mathbf{e}_k^{(l)} \rangle$ for $k \in \Omega^{(l)}$. Then, the estimator can be written as

$$\widehat{F}_*({o}_1^{(l)}, \dots, {o}_{R_l}^{(l)}) = \sum_{l=1}^L \sum_{r=1}^{R_l} \langle \mathbf{a}^{(l)}, \mathbf{e}_{o_r^{(l)}}^{(l)} \rangle + c,$$

which is affine in $\mathbf{e}_{o_r^{(l)}}^{(l)}$. However, in the above equation, \widehat{F}_* is not an affine function of the input. To remedy this, we define the vector of experimentally observed frequencies $\mathbf{f}^{(l)} \in \mathbb{R}^{N_l}$, obtained by binning the outcomes $o_1^{(l)}, \dots, o_{R_l}^{(l)}$, such that $f_k^{(l)}$ denotes the relative frequency of the outcome $k \in \Omega^{(l)}$ observed in the experiment for the l th measurement setting. We can then write

$$\mathbf{f}^{(l)} = \frac{1}{R_l} \sum_{r=1}^{R_l} \mathbf{e}_{o_r^{(l)}}^{(l)}.$$

We can then express the fidelity estimator in a way that is indeed affine in the observed frequencies as

$$\widehat{F}_*(\mathbf{f}^{(1)}, \dots, \mathbf{f}^{(L)}) = \sum_{l=1}^L R_l \langle \mathbf{a}^{(l)}, \mathbf{f}^{(l)} \rangle + c. \quad (8)$$

Equation (8) gives an intuitive understanding of the estimator constructed by the minimax method (it is simply an appropriate weighting of the relative frequencies observed in the experiment). The weights are obtained from the saddle point of the function Φ defined in Eq. (4).

D. Risk and the best sample complexity

We now learned about the estimator, but what about the risk? We found that the risk ε_* given by Eq. (5) is half the saddle-point value of the function Φ . However, in this form, it is difficult to infer what this quantity is. In Appendix A, we show that one can write the risk ε_* given by the minimax method in a form that is more amenable to interpretation as

$$\varepsilon_* = \max_{\chi_1, \chi_2 \in \mathcal{X}} \frac{1}{2} [\text{Tr}(\rho \chi_1) - \text{Tr}(\rho \chi_2)]$$

$$\text{s.t. } \prod_{l=1}^L [F_C(\chi_1, \chi_2, \{E_k^{(l)}\})]^{R_l/2} \geq \frac{\delta}{2}, \quad (9)$$

where

$$F_C(\chi_1, \chi_2, \{E_k^{(l)}\}) = \left(\sum_{k=1}^{N_l} \sqrt{\text{Tr}(E_k^{(l)} \chi_1) \text{Tr}(E_k^{(l)} \chi_2)} \right)^2 \quad (10)$$

denotes the classical fidelity between the probabilities determined by states χ_1 and χ_2 corresponding to the POVM measurement $\{E_k^{(l)}\}$. Here, the measurement protocol chosen by the experimenter corresponds to measuring L different POVMs, where the l th POVM $\{E_1^{(l)}, \dots, E_{N_l}^{(l)}\}$ is measured R_l times, for $l = 1, \dots, L$. Thus, the risk ε_* can be intuitively interpreted as half the largest deviation $F(\rho, \chi_1) - F(\rho, \chi_2)$ of the fidelity of the target state ρ with the states χ_1, χ_2 compatible with the measurement protocol and confidence level. The factor of half comes from the fact that we are looking for a symmetric confidence interval.




FIG. 1. Variation of the risk with the number of Pauli measurements L and number of repetitions R_L of each Pauli measurement for a four-qubit random target state. All the computed risks correspond to a 95% confidence level.

Note that the fidelity between any two quantum states must lie between 0 and 1, and thus, $0 \leq \text{Tr}(\rho \chi) \leq 1$ for any density matrix χ . Then, we can infer from Eq. (9) that the maximum possible value of risk is 0.5. Indeed, an uncertainty of ± 0.5 gives a confidence interval of length 1, which is also the size of the interval for fidelity.

However, it can be shown from Eq. (9) that the risk decreases when the number of measurement settings or the number of shots for any measurement setting is increased. To that end, note that F_C ranges between 0 and 1, so raising it to a larger power makes it smaller. Thus, increasing the number of shots R_l makes it harder to satisfy the constraint that $\prod_l F_C^{R_l/2}$ must be larger than $\delta/2$. Similarly, if we include another measurement setting (in effect, increasing L), we have one more fraction F_C multiplying the left-hand side of the constraint, thereby making the constraint tighter. Since a tighter constraint implies a more restricted search space for maximization, we can infer that the risk will become smaller (or stay the same) in either case. This also shows that the risk is dependent on the chosen measurement protocol, and protocols that make the constraint tighter will have a smaller risk in estimating the fidelity.

We numerically quantify the variation of the risk with the number of measurement settings and the number of repetitions for a four-qubit randomly chosen target state. We apply 10% depolarizing noise to obtain the actual state and perform Pauli measurements. We see from Fig. 1 that the risk decreases with the number of Pauli measurements as well as the number of repetitions.

A natural question that arises is how low the risk can be. To find a lower bound, recall that we can write the fidelity between any two states as a minimization of the classical fidelity over all POVMs [29]

$$F(\chi_1, \chi_2) = \min_{\text{POVM } \{F_i\}} F_C(\chi_1, \chi_2, \{F_i\}).$$

In particular, we have $F(\chi_1, \chi_2) \leq F_C(\chi_1, \chi_2, \{E_k^{(l)}\})$ for every POVM $\{E_k^{(l)}\}$ that we are using. Thus, we obtain the following lower bound on our risk:

$$\varepsilon_* \geq \max_{\chi_1, \chi_2 \in \mathcal{X}} \frac{1}{2} [\text{Tr}(\rho \chi_1) - \text{Tr}(\rho \chi_2)]$$

$$\text{s.t. } F(\chi_1, \chi_2) \geq \left(\frac{\delta}{2} \right)^{\frac{2}{R}},$$

where $R = \sum_{l=1}^L R_l$ is the total number of shots. Evaluating the right-hand side of the above equation, we obtain a lower bound for the risk. This result is summarized in the following theorem.

Theorem II.1. Let ρ be any pure target state. Suppose that ε_* is the risk associated with the fidelity estimator given by the minimax method corresponding to a confidence level of $1 - \delta \in (0.75, 1)$ and any measurement scheme. Then, the risk is bounded below as

$$\varepsilon_* \geq \frac{1}{2} \sqrt{1 - \left(\frac{\delta}{2}\right)^{2/R}}, \quad (11)$$

where R is the total number of measurement outcomes (all measurement settings combined) to be supplied to the estimator. This lower bound can be achieved by measuring in the basis defined by ρ . That is, R repetitions of the POVM $\{\rho, \mathbb{I} - \rho\}$ achieves the risk ε_* .

Stated differently, the best sample complexity that can be obtained using the minimax method corresponding to a risk of $\varepsilon_* \in (0, 0.5)$ and confidence level $1 - \delta$ is given by

$$\begin{aligned} R &\geq \frac{2 \ln(2/\delta)}{|\ln(1 - 4\varepsilon_*^2)|} \\ &\approx \frac{\ln(2/\delta)}{2\varepsilon_*^2} \text{ when } \varepsilon_*^2 \ll 1. \end{aligned} \quad (12)$$

Proof. See Appendix C. ■

Note that sample complexity refers to the smallest number of measurements outcomes required for achieving the desired risk for the chosen confidence level.

Observe that the lower bound on the risk in Eq. (11) is independent of the system dimension, the target state, and the true state. It only depends on the confidence level and total number of repetitions. Note that it is possible to obtain dimension-independent sample complexity because we assume that the target state ρ is pure. Without this assumption, the sample complexity for fidelity estimation is expected to scale with the dimension (see, for example, Refs. [30,31]). That said, the assumption that the target state is pure is well motivated from a practical standpoint as we usually seek to prepare pure states in experiments [e.g., Bell states or Greenberger—Horne—Zeilinger (GHZ) states]. Indeed, this is a common assumption in many fidelity estimation studies [12–14]. Since the true state is allowed to be mixed, our method is applicable in practical scenarios. Assuming that the target state is pure, one can, in principle, achieve the optimal sample complexity by measuring in a basis defined by the target state ρ . We acknowledge that, in practice, it is often impractical to implement the POVM $\{\rho, \mathbb{I} - \rho\}$ (or an orthonormal basis containing ρ) that achieves this sample complexity. However, next we show that we can achieve something very close to this optimal sample complexity for stabilizer states using a practical measurement scheme.

E. Stabilizer states

Stabilizer states are the cornerstone of numerous applications, ranging from measurement-based quantum computing [32] to quantum error correction [16]. An n -qubit stabilizer state ρ is the unique $+1$ eigenstate of exactly n Pauli operators

that generate a stabilizer group of size $d = 2^n$. The state ρ can thus be written as [33]

$$\rho = \frac{1}{d} \left(\mathbb{I} + \sum_{S \in \mathcal{S}_n \setminus \{\mathbb{I}\}} S \right), \quad (13)$$

where \mathcal{S}_n is the stabilizer subgroup corresponding to the stabilizer state ρ .

To estimate the fidelity with a target stabilizer state efficiently, we implement the minimax optimal strategy [15,33] while restricting to Pauli measurements. The measurement strategy is simple: uniformly sample an element from the stabilizer subgroup (excluding the identity), and record the eigenvalue of the outcome (whether $+1$ or -1) of the stabilizer measurement. This strategy is very similar to DFE, with the exception that we exclude the identity. This measurement scheme can be implemented by an effective POVM with elements $\{\Theta, \Delta_\Theta\}$ given by

$$\begin{aligned} \Theta &= \rho + \frac{d/2 - 1}{d - 1} \Delta_\rho, \\ \Delta_\Theta &= \frac{d/2}{d - 1} \Delta_\rho, \end{aligned} \quad (14)$$

where $\Delta_\rho = \mathbb{I} - \rho$. See the proof of Theorem II.2 for details on how this is obtained (also see Refs. [15,33]). We can see that the effective POVM is a combination of the target state ρ and Δ_ρ (which is orthogonal to ρ). This is similar to measuring in the basis defined by the target state, so from our results in Sec. IID, we can expect the risk to be small. Indeed, we find that for a sufficient number of repetitions of this measurement (independent of the dimension or the stabilizer state), the risk of the estimator is, at most, four times the optimal risk described in Eq. (11). We summarize this result below.

Proposition II.1. Let ρ be an n -qubit stabilizer state. Suppose that we uniformly sample R elements from the stabilizer group (with replacement) and measure them. Then, for any risk $\varepsilon_* \in (0, 0.5)$ and confidence level $1 - \delta \in (0.75, 1)$,

$$R \geq 2 \frac{\ln(2/\delta)}{\left| \ln \left(1 - \left(\frac{d}{d-1}\right)^2 \varepsilon_*^2 \right) \right|} \quad (15)$$

$$\approx 2 \left(\frac{d-1}{d} \right)^2 \frac{\ln(2/\delta)}{\varepsilon_*^2} \text{ when } \varepsilon_*^2 \ll 1 \quad (16)$$

suffices to build an estimator using the minimax method that achieves this risk. Here, $d = 2^n$ is the dimension of the system.

Proof. Take $\xi = d$ in corollary C.1.2 in the Appendix. ■

Since $(d-1)/d < 1$, the dimension dependence actually improves the sample complexity, and for large dimensions one needs to measure only a constant number of stabilizers. This sample complexity is of the same order as that obtained using DFE.

Furthermore, following the ideas given in the proof of Theorem C.1, we can simplify the algorithm given in Appendix B to compute the estimator. In the simplified algorithm, we only need to perform a two-dimensional optimization regardless of the dimension of the system. Consequently, building the estimator is both time and memory efficient.

TABLE II. True fidelity, the fidelity estimate, and the number of stabilizers sampled corresponding to a risk of 0.05 and 95% confidence level for n -qubit stabilizer states.

		n		
		2	3	4
True fidelity	$F(\rho, \sigma)$	0.925	0.912	0.906
Estimate of fidelity	$\widehat{F}_*(\text{outcomes})$	0.933	0.911	0.898
Number samples	R	1657	2256	2591

In the following we present simulated results for two-, three-, and four-qubit stabilizer states, for a risk of 0.05 and a confidence level of 95%. We implement the measurement strategy discussed above, namely, randomly sampling R stabilizers and recording the outcome eigenvalue. The estimated fidelity in each case is summarized in Table II. Even for the four-qubit stabilizer state, constructing the estimator and generating measurements to compute the estimate together took only a few seconds on a personal computer.

Now, on the other extreme, we consider the case where an insufficient number of stabilizer measurements are provided. Specifically, we consider a measurement protocol where we measure only $n - 1$ generators of an n -qubit stabilizer state. We assume these to be subspace measurements, i.e., the measurements correspond to projecting on the eigenspaces of the generators with eigenvalue ± 1 . This is an insufficient measurement protocol because the measurement of $n - 1$ stabilizers cannot uniquely identify the stabilizer state, as there are two orthogonal states that will be consistent with the measurement statistics. As a consequence, any reasonable fidelity estimation protocol should give complete uncertainty for the estimated fidelity. Indeed, the following result shows that the minimax method gives a risk of 0.5 (implying total uncertainty) for this measurement protocol.

Proposition II.2. Let ρ be an n -qubit stabilizer state generated by S_1, \dots, S_n . Suppose that the measurement protocol consists of measuring the first $n - 1$ generators S_1, \dots, S_{n-1} , such that the measurements correspond to projecting on eigenspaces of S_i with eigenvalue ± 1 for $i = 1, \dots, n$. Then, irrespective of the number of shots, the minimax method gives a risk of 0.5.

Proof. See the end of Appendix C. \blacksquare

We remark that if we consider eigenbasis measurements (in contrast to subspace measurements considered above), computing the risk analytically is more complicated. By eigenbasis measurement, we mean that the measurement of the stabilizer S has the POVM $\{E_1, \dots, E_d\}$ where E_i is the projection on i th eigenvector of S . Note that the eigenbasis for S is not unique since S has a degenerate spectrum, and therefore, the risk can depend on the choice of eigenbasis used for the measurement. For example, if the target state ρ happens to coincide with one of the eigenvectors used in the measurement, then that measurement has sufficient information to accurately estimate the fidelity. However, when this does not happen, one can expect the risk to be 0.5; see Sec. III B for an example.

F. Randomized Pauli measurement scheme

In the spirit of the randomized measurement strategy for stabilizer states we discussed above, we present a generalized version of such a Pauli measurement scheme for arbitrary (pure) states. The idea comes from the simple observation that any density matrix χ can be written as

$$\chi = \frac{\mathbb{I}}{d} + \sum_{i=1}^{d^2-1} \frac{\text{Tr}(W_i \chi)}{d} W_i \quad (17)$$

$$= \frac{\mathbb{I}}{d} + \sum_{i=1}^{d^2-1} \frac{|\text{Tr}(W_i \chi)|}{d} S_i, \quad (18)$$

where $S_i = \text{sign}[\text{Tr}(W_i \chi)] W_i$ are the Pauli operators appended with a sign. This is, in particular, valid for the target state ρ . Then, consider the probability distribution

$$p_i = \frac{|\text{Tr}(W_i \rho)|}{\sum_{i=1}^{d^2-1} |\text{Tr}(W_i \rho)|}, \quad i = 1, \dots, d^2 - 1, \quad (19)$$

where ρ is the target state. The measurement scheme is as follows.

Scheme II.1. Pauli Measurement Scheme

(1) Sample a (nonidentity) Pauli operator W_i with probability p_i ($i = 1, \dots, d^2 - 1$) given in Eq. (19) and record the outcome (± 1) of the measurement.

(2) Flip the measurement outcome $\pm 1 \rightarrow \mp 1$ if $\text{Tr}(\rho W_i) < 0$, or else retain the original measurement outcome.

(3) Repeat this procedure R times and feed the outcomes into the estimator given by the minimax method.

Because we exclude the measurement of identity, i runs from 1 to $d^2 - 1$. We flip the outcomes because we need to measure $S_i = \text{sign}[\text{Tr}(W_i \rho)] W_i$. In Theorem II.2, we show how to choose the number of repetitions R so as to obtain a desired value of the risk. Note that the Pauli measurement can either be a projection on the subspace with eigenvalue ± 1 or a projection on the eigenvectors. For the estimator given by the minimax method, the values $+1, -1$ are inconsequential; all that matters is how many times $+1$ and -1 are observed.

Note that a very similar measurement scheme has been considered for verifying the ground state of a class of Hamiltonians [34]. Further, the random sampling scheme described above is very similar to the measurement strategy used in DFE, except that we sample the Pauli operators using a different probability distribution.

For the proposed measurement strategy, the minimax method gives the following sample complexity.

Theorem II.2. Let ρ be an n -qubit pure target state. Suppose that we perform R Pauli measurements as described in Scheme II.1. Then, for a given risk $\varepsilon_* \in (0, 0.5)$ and a confidence level $1 - \delta \in (0.75, 1)$,

$$R \geq 2 \frac{\ln(2/\delta)}{\left| \ln \left(1 - \frac{d^2}{\mathcal{N}^2} \varepsilon_*^2 \right) \right|} \quad (20)$$

$$\approx 2 \left(\frac{\mathcal{N}}{d} \right)^2 \frac{\ln(2/\delta)}{\varepsilon_*^2} \text{ when } \varepsilon_*^2 \ll 1 \quad (21)$$

measurements are sufficient to achieve the risk. Here,

$$\mathcal{N} = \sum_{i=1}^{d^2-1} |\text{Tr}(W_i \rho)|$$

and for any pure target state ρ , we have

$$\mathcal{N} \leq (d-1)\sqrt{d+1}.$$

Therefore

$$2\left(\frac{\mathcal{N}}{d}\right)^2 \frac{\ln(2/\delta)}{\varepsilon_*^2} \leq 2(d+1)\left(1 - \frac{1}{d}\right)^2 \frac{\ln(2/\delta)}{\varepsilon_*^2}.$$

Proof. See the end of Appendix C. \blacksquare

Notably, the sample complexity is upper bounded by $O(d)\ln(2/\delta)/\varepsilon_*^2$, which is comparable to the upper bound on (the expected value of) the sample complexity in DFE (see Eq. (10) in Ref. [12]). We can show that the sample complexity indeed improves when considering a subset of well-conditioned states defined by Flammia and Liu [12]. To that end, suppose that for any given i , the state ρ satisfies either $|\text{Tr}(\rho W_i)| = \alpha$ or $\text{Tr}(\rho W_i) = 0$. Then, using Eqs. (21) and (C9), we obtain the bound of

$$\frac{2}{\alpha^2} \left(\frac{d-1}{d}\right)^2 \frac{\ln(2/\delta)}{\varepsilon_*^2} \quad (22)$$

on the sample complexity. As before, this is comparable to DFE. In addition to giving a good sample complexity, we can obtain the estimator efficiently by reducing the optimization to a two-dimensional problem, following the ideas in Theorem C.1. Thus, if \mathcal{N} can be computed efficiently, we can efficiently construct the estimator even for large dimensions.

G. Robustness against errors

For any fidelity estimator to be useful in practice, it is crucial for it to be robust to common types of experimental errors. Systematic errors in the measurements will ultimately perturb the observed outcome frequencies. Since the estimator given by the minimax method is affine, we can ensure that small changes in the observed frequencies only lead to small changes in the fidelity estimate. These changes can be quantified as follows.

Let $\mathbf{f}^{(1)}, \dots, \mathbf{f}^{(L)}$ denote the ‘‘ideal’’ observed frequencies, i.e., when there are no measurement errors in the measurement settings $l = 1, \dots, L$. Let $\tilde{\mathbf{f}}^{(1)}, \dots, \tilde{\mathbf{f}}^{(L)}$ denote the actual observed frequencies due to the presence of measurement errors. From Eq. (8), we know that the minimax estimator can be written as

$$\hat{F}_*(\mathbf{f}^{(1)}, \dots, \mathbf{f}^{(L)}) = \sum_{l=1}^L R_l \langle \mathbf{a}^{(l)}, \mathbf{f}^{(l)} \rangle + c.$$

For simplicity, denote $\mathbf{f} = (\mathbf{f}^{(1)}, \dots, \mathbf{f}^{(L)})$ and $\mathbf{a} = (R_1 \mathbf{a}^{(1)}, \dots, R_L \mathbf{a}^{(L)})$ a larger vector obtained by concatenating all vectors from $l = 1, \dots, L$. Denote $\Delta \mathbf{f} = \tilde{\mathbf{f}} - \mathbf{f}$ the deviation of the observed frequencies in the erroneous measurement case from the observed frequencies in the ideal measurement case. Using Hölder’s inequality, we can infer that

$$|\hat{F}_*(\mathbf{f}) - \hat{F}_*(\tilde{\mathbf{f}})| \leq \|\mathbf{a}\|_1 \|\Delta \mathbf{f}\|_\infty, \quad (23)$$

where for any vector \mathbf{x} , we have $\|\mathbf{x}\|_1 = \sum_i |x_i|$ and $\|\mathbf{x}\|_\infty = \max_i |x_i|$. We obtain

$$\|\mathbf{a}\|_1 = \sum_{l=1}^L R_l \sum_{k=1}^{N_l} |a_k^{(l)}|$$

and

$$\|\Delta \mathbf{f}\|_\infty = \max_{l=1, \dots, L} \max_{k=1, \dots, N_l} |\tilde{f}_k^{(l)} - f_k^{(l)}|.$$

The quantity $\|\mathbf{a}\|_1$ depends only on the minimax estimator and is finite. Thus, for small-enough perturbations, we can ensure that the fidelity estimates in the erroneous measurement case are close to the fidelity estimates in the error-free case. This argument qualitatively shows that the estimator given by the minimax method is robust against perturbations.

We illustrate this property by considering a simple example of an erroneous measurement setup that one could encounter in practice. Suppose we wish to perform Pauli measurements, but due to some issue with the experimental setup, each qubit undergoes a systematic rotation at the beginning of the measurement process, changing the POVM. The ideal measurements correspond to measuring Pauli operators and recording the observed eigenvalues. This procedure amounts to choosing POVMs that project onto the ± 1 eigenspace of the respective Pauli operator (we call this subspace projection). To simulate erroneous measurements, we suppose that each qubit undergoes an erroneous rotation along a given axis by a fixed angle. For the purposes of our example, we generate a randomly chosen axis (unit vector) \mathbf{n}_j for the j th qubit. The rotation angle θ_j for the j th qubit is chosen from a normal distribution with zero mean and standard deviation σ_θ that controls the amount of noise. Thus, the unitary describing this erroneous rotation is given by $U_{\text{error}} = \otimes_{j=1}^n \exp(i\theta_j \mathbf{n}_j \cdot \boldsymbol{\sigma})$, where $\boldsymbol{\sigma} = (X, Y, Z)$ denotes the one-qubit Pauli operators. We assume that this error is systematic, i.e., the same error acts on all the POVM elements. In other words, if the ideal POVMs are $E_k^{(l)}$, the actual POVMs being implemented are $U_{\text{error}} E_k^{(l)} U_{\text{error}}^\dagger$. Since we choose $E_\pm^{(l)} = (\mathbb{I} \pm P_l)/2$ as projections onto eigenspaces of the Pauli operator P_l , the erroneous POVM corresponds to $(\mathbb{I} + U_{\text{error}} P_l U_{\text{error}}^\dagger)/2$. This is just a slight rotation of the Pauli observable to be measured due to calibration issues, for example.

To study the performance of our fidelity estimation method under such a systematic error, we consider a three-qubit GHZ state as the target state, and the actual state is obtained by applying 10% depolarizing noise to the target state. We measure $L = 7$ (nonidentity) Pauli operators that have nonzero expectation with the target state. Each of these operators is measured $R_l = 300$ times. We choose a confidence level of 95%. In Table III, we list the fidelity estimate given by our method using ideal Pauli measurements, the fidelity estimate due to erroneous Pauli measurements, as well as the bound on the deviation of the estimates given by Eq. (23). We find that the deviations in the fidelity estimate always lie within this bound. In particular, for small rotation angles, the fidelity estimates with and without error are nearly the same. This substantiates our claim that our method is robust to small amounts of measurement errors without any modifications to the estimation procedure.

TABLE III. Analysis of fidelity estimates given by the minimax method under systematic measurement error. The true fidelity with a three-qubit GHZ target state is 0.912. Ideal Pauli measurements give a fidelity estimate of 0.915. A systematic measurement error is present due to erroneous rotation of Pauli operators, with the size of the rotation controlled by σ_θ . The calculations in the table correspond to one randomly drawn instance of these erroneous rotations. Fidelity estimates using erroneous measurements, the absolute value of deviation of estimates from erroneous and ideal measurements, and the bound given by Eq. (23) are listed as a function of σ_θ .

σ_θ	Estimates from		
	erroneous measurements	Deviation	Bound
0.05 rad (2.86°)	0.911	0.004	0.004
0.1 rad (5.73°)	0.901	0.014	0.058
0.15 rad (8.59°)	0.878	0.037	0.069
0.2 rad (11.45°)	0.849	0.066	0.111

Depending on the experimental setup, the measurement noise encountered in the system can be different from the one we have considered above. However, even for other forms of noise, our method is guaranteed to be robust, although the exact dependence on the noise may vary. If one is interested in precise quantitative bounds for a specific error model, these can be computationally estimated following the procedure outlined above.

III. COMPARISON WITH OTHER METHODS

We now compare the minimax method with two commonly used techniques to estimate the fidelity: direct fidelity estimation (DFE) and maximum likelihood estimation (MLE) [and a related approach called profile likelihood (PL)]. We further compare it with a simple semidefinite programming (SDP) based approach.

A. Direct fidelity estimation

Flammia and Liu [12] and Silva *et al.* [13] constructed a fidelity estimator by judiciously sampling Pauli measurements. If the target state ρ is well conditioned, i.e., $|\text{Tr}(\rho W)| \geq \alpha$ or $\text{Tr}(\rho W) = 0$ for each Pauli operator W and some fixed $\alpha > 0$, then DFE gives a good sample complexity. Specifically, their method used $O[\ln(1/\delta)/\alpha^2 \varepsilon^2]$ measurement outcomes to obtain an estimate for fidelity within an additive error of ε and a confidence level of $1 - \delta$ [12]. Their rigid measurement scheme, however, can have disadvantages in practice, compared to a more flexible approach that works for arbitrary POVMs.

Consider a random four-qubit state as the target state, with the actual state obtained by applying 10% depolarizing noise to the target state. We choose a risk of 0.05 in DFE and a confidence level of 95%, and obtain the Pauli measurements that need to be performed as prescribed by DFE. However, to perform these Pauli measurements, we choose two different types of POVM: (1) projection on +1 and -1 eigenvalue subspaces of the Pauli operator, and (2) projection on each eigenstate of the Pauli operator, which is common in experiments. For DFE, it does not matter much which of the two

TABLE IV. Comparison of DFE method with the minimax method for a four-qubit random state and four-qubit GHZ state as target states. The Pauli measurements are performed as per the prescription of DFE, but two different types of POVMs are used: projection on subspace with +1 and -1 eigenvalue, and projection on each eigenstate of the Pauli operator. We find that the minimax method has lower risk in most cases, while the estimates are comparable to those of DFE method.

	Random state			
	Subspace projection		Eigenbasis projection	
	DFE	Minimax	DFE	Minimax
True fidelity	0.906		0.906	
Estimate	0.895	0.895	0.900	0.902
Risk	0.050	0.023	0.050	0.012
	GHZ state			
	Subspace projection		Eigenbasis projection	
	DFE	Minimax	DFE	Minimax
True fidelity	0.906		0.906	
Estimate	0.907	0.907	0.904	0.906
Risk	0.050	0.022	0.050	0.018

POVMs is implemented because the estimator there only uses the expectation value, while the minimax method makes a distinction between these POVMs. From Table IV, we can see that the risk (≈ 0.023) for the minimax method is less than half the DFE risk (0.05) when using projection on subspaces (1). When projecting on each eigenbasis element (2), we obtain a risk (≈ 0.012), which is even lower than before, and clearly much better than the DFE risk. That is, we are able to use the larger expressive power of POVM (2) compared to POVM (1) to lower the risk. Since the minimax method has already computed an estimator with a low risk (as per the prescription of the DFE method), all subsequent experiments can use the *same* measurement settings to estimate the fidelity. Random sampling of Pauli measurements is not necessary.

Next, we consider a four-qubit GHZ state as the target state, which is a well-conditioned state as per DFE. As before, we perform measurements as prescribed by DFE by choosing a risk of 0.05 for DFE method and a confidence level of 95%. We summarize the estimated fidelity and risks in the second half of Table IV. Similarly to the case of the random state, we observe that the risk given by the minimax method for the GHZ state is always lower than the DFE risk. Changing from POVM (1) to POVM (2) again leads to an improvement in the risk.

Finally, we argue that the Pauli measurement scheme prescribed by DFE can be far from optimal for certain target states. To demonstrate this, we choose a rather extreme example of a random four-qubit target state with the following measurement scheme: perform as many measurements and repetitions as prescribed by DFE for achieving a risk of 0.05 at a confidence level of 95%, but instead of performing Pauli measurements, we choose random POVMs, generated as per Ref. [35]. In this case, we obtain a risk of ≈ 0.023 from the minimax method, which is larger than the risk for Pauli measurements using the minimax method (see Table IV), but

still a factor of 2 lower than the DFE risk. This suggests that, when the target state is random, performing random measurements is almost as good as performing Pauli measurements. Hence there is no real advantage in using DFE in terms of sample complexity. If, instead, we used the optimal (though impractical) measurement scheme described in Sec. II D, we can obtain a risk of 0.05 using just ≈ 700 total measurements (independent of the target state and dimension) as opposed to $\approx 248\,000$ total measurements required by DFE for the random target state. While this random state example is rather atypical, it stands to demonstrate that there are cases where the DFE measurement scheme is outperformed by a more tailored scheme, which the minimax method can benefit from.

A potential disadvantage for the minimax method is that once the dimension of the system becomes very large, depending on the measurement scheme, constructing the estimator can become inefficient. One reason for this inefficiency is that the algorithm we use for optimization requires projection on the set of density matrices. Because we perform this projection through diagonalization, the projection can cost up to $O(d^3)$ iterations, where d is the dimension of the system. In contrast, DFE can handle large dimensions well because the classical computations involved are simple and most of the complexity is absorbed into performing the experiments. This drawback of the minimax method is alleviated if we use the Pauli measurement scheme described in Scheme II.1, which is similar to DFE. We have written an efficient algorithm to construct the estimator given by the minimax method for this measurement scheme, as noted in Sec. II F. Moreover, the estimator can be precomputed once the measurement scheme has been defined.

B. MLE and profile likelihood

Maximum likelihood estimation (MLE) is a popular approach used for quantum tomography [8]. One can think of MLE as minimizing the Kullback-Liebler divergence between the observed frequencies and the Born probabilities [36]. This naturally leads to the (negative) log-likelihood function

$$\ell(\{f^{(l)}\}_{l=1}^L | \chi) = - \sum_{l=1}^L \frac{R_l}{R} \sum_{k=1}^{N_l} f_k^{(l)} \ln [\text{Tr}(E_k^{(l)} \chi)], \quad (24)$$

where $R = \sum_{l=1}^L R_l$. One then minimizes the function ℓ (or equivalently, maximizes the likelihood function) to obtain the MLE estimate $\hat{\sigma}$ for the experimental quantum state. We calculate $\text{Tr}(\rho \hat{\sigma})$ to estimate the fidelity with a target state ρ . While the estimates are usually good enough when a sufficient number of measurements is used for the reconstruction, a major disadvantage is that the method provides no confidence intervals for the estimated fidelity.

A common approach to compute an uncertainty of the MLE estimate is Monte Carlo (MC) resampling. Herein, one numerically generates (artificial) outcomes based on the observed frequencies and the expected statistical noise distribution, referred to as resampling. For each set of outcomes one reconstructs the state and estimates the fidelity as before. This process is repeated many times to obtain a large number of MLE fidelity estimates which can be used to calculate an (asymmetric) interval around the median that corresponds to

the chosen confidence level. If necessary, hedging [37] can be implemented to deal with zero probabilities. Note that the MC approach is similar to the nonparametric bootstrap method used for finding confidence intervals [38].

A disadvantage of such an MLE-based approach is that we need to reconstruct the state, which can be costly. Moreover, there are certain freedoms in the resampling and hedging definitions, which can affect the results. We argue that the MC approach can give overconfident uncertainty bounds. This issue can be seen in the following example: take a two-qubit Bell state, stabilized by XX and ZZ , as the target state. We measure only XX with 500 repetitions, which is not enough to uniquely determine the state. The true state is obtained by applying 10% depolarizing noise to the Bell state, so that the actual fidelity is 0.925. We generate 100 MLE fidelity estimates (with new measurement outcomes generated at every repetition), and we find that the average MLE fidelity estimate is 0.44 when the measurements correspond to projection onto an eigenbasis of XX .

The large error (an estimate of 0.44 compared to the true value of 0.925) is expected since the measurements are chosen poorly with respect to the state. This problem, however, is not detected by the MC approach, where we find that the average uncertainty corresponding to a confidence level of 95% is (0.32, 0.38). This indicates that MC uncertainty is overconfident (because $0.44 + 0.38 = 0.82 < 0.925$). Indeed, this is seen through the empirical coverage probability which turns out to be 0, in stark contrast with the high confidence level of 95% that was chosen. By empirical coverage probability, we mean the fraction of cases where the true fidelity lies inside the computed confidence interval. We obtain very similar results when using subspace measurements [average fidelity of 0.45, average uncertainty of (0.31, 0.37), and empirical coverage probability of 0]. Therefore, we cannot always trust the uncertainty given by the MC approach. In contrast, the minimax approach gives a risk of 0.5 for this example (implying total uncertainty) when insufficient number of stabilizer measurements are provided irrespective of whether eigenbasis or subspace measurements are used.

We note that such a problem with the MC approach is also present in more realistic scenarios. For example, consider a three-qubit W state as the target state, and suppose that the true state has a fidelity of 99.1% with the target state. We measure $N = 19$ Pauli operators (subspace measurements), choosing all the nonidentity Pauli operators that have nonzero overlap with the W state. Each Pauli measurement is repeated 100 times, and the outcomes together are used to estimate the fidelity corresponding to a confidence level of 95%. Using a total of 150 such estimates, we find that the MC method gives an empirical coverage probability of 72%, which is much smaller than the chosen confidence level of 95%. Thus, the MC approach is overconfident and gives tighter bounds than warranted in practical situations of interest.

In addition to the statistical problems noted above, we remark that MC resampling is also computationally costly. This is because MC resampling requires performing MLE for each (artificial) fidelity estimate generated, which is costly. Moreover, this costly computation needs to be performed after the experimental data has been collected. In contrast, our method only needs to compute the estimator once (potentially




FIG. 2. A schematic of a typical profile likelihood (PL) plotted against the parameter F (representing fidelity). The MLE estimate (green) corresponds to the minimum of the PL curve. A cutoff for PL (red dashed horizontal line) gives a bound $[F_1, F_2]$ for the estimated fidelity.

costly computation), and this can be done before collecting measurement data. Once the estimator has been constructed, it runs very efficiently on the experimental data. Thus, we not only obtain statistical but also computational benefits over MLE with MC resampling.

An alternate approach to obtain a bound for the MLE estimate is calculating the PL function [39]. Given any value $F \in [0, 1]$, PL corresponds to the solution of the optimization problem

$$\begin{aligned} \text{PL}(F) &= \min_{\chi \in \mathcal{X}} \ell(\{f^{(l)}\}_{l=1}^L | \chi) \\ \text{s.t. } & \text{Tr}(\rho \chi) = F. \end{aligned} \quad (25)$$

Note that we define the profile likelihood in terms of negative log-likelihood instead of the likelihood function. Our definition of PL amounts to solving the MLE optimization problem, except for the added constraint that the fidelity with the target state must be equal to F . The MLE solution can be obtained from PL by adding an additional layer of optimization: $\text{MLE} = \min_{F \in [0,1]} \text{PL}(F)$. It can be shown that $\text{PL}(F)$ is convex in F . The advantage of calculating PL is that given a cutoff value, one can obtain a bound on the fidelity similar to an error bar. To this end, we draw a horizontal line (the cutoff) on the PL versus F plot, and the locations along the F axis where this line intersects the curve gives a bound on estimated fidelity. The MLE estimate lies inside this interval because it corresponds to the minimum. A schematic of a typical PL plot is shown in Fig. 2. A somewhat similar idea for obtaining confidence regions was proposed by Faist and Renner [40].

Using PL is a natural way of providing a bound for MLE since it returns an interval of fidelity estimates that correspond to large-enough likelihood. However, it does not provide a true confidence interval as the location of the cutoff value is unknown. The issue is that likelihoods can be considered a generic notion of “plausibility” but not, in the technical sense, “uncertainty;” and that PL gives a *likelihood region*, which is not the same as a coverage probability [41].

What we want is that given many MLE estimates (from actual experimental data), the true fidelity must lie $1 - \delta$ fraction of times inside the bound calculated from PL, where $1 - \delta$ is the chosen confidence level. One general heuristic that is used to give such a cutoff for PL is based on Wilks’ theorem [42]. However, this theorem is only valid asymptotically as the




FIG. 3. Plot of empirical coverage probability against the average width of the bound given by the PL and SDP methods for a three-qubit random target state and 48 Pauli measurements. It can be seen that PL gives tighter bounds than the SDP method. For reference, the minimax bound corresponding to a confidence level of 95% is shown as well.

sample size becomes arbitrarily large, and furthermore, it does not work well for quantum states, especially those that are low rank [43,44]. Note that Scholten and Blume-Kohout [43] gave an alternative to Wilks’ theorem for quantum states, but like Wilks’ theorem, their alternative was only exact asymptotically. Furthermore, it used models different from what we need for PL (they considered nested Hilbert spaces of increasing dimension, while we require density matrices of the same dimension but with an added constraint). Therefore, as of now, we do not know how to obtain a confidence interval using PL.

Hence, for the purpose of demonstration, we choose different cutoff values for the PL and compute the bounds (as described below) on the MLE estimate corresponding to each chosen value. We take a three-qubit random target state and apply 10% depolarizing noise to obtain the actual state. We measure $L = 48 = 0.75 \times 4^3$ Pauli operators, chosen in the decreasing order of weights given by DFE, repeating each measurement 100 times. Using the knowledge of the true state, we check in which of these bounds the true fidelity lies. This process is repeated 10^3 times by generating different observed frequencies using the true state, and subsequently, we obtain the coverage probability corresponding to each value of the cutoff. We also find the average width of the bound for each cutoff value. We plot the coverage probability against this average width in Fig. 3, finding that the bounds are reasonably tight. Note that the true fidelity is usually not known, so the coverage probability cannot be computed in practice.

C. An SDP-based approach

Finally, we compare with a simple semi-definite programming (SDP) method to obtain a bound on fidelity. This involves solving the following intuitive optimization problem to obtain bounds on fidelity

$$\begin{aligned} F_{\min} &= \min_{\chi \in \mathcal{X}} \text{Tr}(\rho \chi) \\ \text{s.t. } & \sum_{i=1}^L \sum_{k=1}^{N_i} (\text{Tr}(E_k^{(i)} \chi) - f_k^{(i)})^2 \leq \epsilon_m, \end{aligned} \quad (26)$$

$$\begin{aligned} F_{\max} &= \max_{\chi \in \mathcal{X}} \text{Tr}(\rho \chi) \\ \text{s.t. } & \sum_{i=1}^L \sum_{k=1}^{N_i} (\text{Tr}(E_k^{(i)} \chi) - f_k^{(i)})^2 \leq \epsilon_m. \end{aligned} \quad (27)$$

TABLE V. Comparison of different methods for estimation of fidelity: minimax method, maximum likelihood estimation (MLE) with Monte Carlo (MC) sampling for estimating the fidelity and uncertainty, profile likelihood (PL) and semi-definite programming (SDP) methods for calculating bounds on fidelity and the direct fidelity estimation (DFE) method.

	Minimax	MLE, MC	PL	SDP	DFE
No unknown parameters required?	✓ ^a	✓	✗	✗	✓
Provides a rigorous confidence interval (never overconfident)?	✓	✗	✗ ^b	✗ ^b	✓
Risk level known before seeing the outcomes?	✓	✗	✗	✗	✓
Applies to any measurement setting?	✓	✓	✓	✓	✗
No significant computation required in practice?	✗ ^c	✗	✗	✗	✓

^aTechnically, ϵ_o can be considered as a free parameter, but we fix it at $\epsilon_o = 10^{-5}$. Since we do not have to tune ϵ_o , we do not list it as an unknown parameter.

^bBecause no systematic method of obtaining a confidence interval is known.

^cThe computational time required for the minimax method depends on the system dimension, the target state, the measurement settings, and the algorithm used. For the Pauli measurement scheme in Sec. II F, we have a fast algorithm irrespective of the target state or the dimension.

In essence, we find the minimum and maximum fidelity with the target state over density matrices that satisfy the measurement statistics up to an error of ϵ_m . The advantage of such an approach is that the bounds obtained are independent of the method used to estimate fidelity. The drawback, however, is that the parameter ϵ_m needs to be chosen “by hand.” Since we need to tune the parameter ϵ_m similar to the cutoff in the PL method, we plot the coverage probability against average width of the bound as before in the Fig. 3. We use the same state and measurement settings as for PL, and 10^3 estimates to compute the coverage probability. We can see that the average width of the bound given by SDP method is typically larger than both PL and the minimax method.

We remark that the minimax method gives a wider confidence interval than PL in Fig. 3 due to its conservative definition. In turn, the minimax method’s confidence interval is guaranteed to hold irrespective of what the true state of the system is. Indeed, we can construct the fidelity estimator using the minimax method even before taking any data. In contrast, the PL and SDP methods compute bounds on the fidelity after the experimental data are collected, so they can be tighter in principle. However, the PL method requires choosing a cutoff while the SDP method requires choosing the parameter ϵ_m . For experimental data, we have no systematic way to choose these quantities to ensure that the bound corresponds to a genuine confidence interval. Therefore, we cannot use PL and SDP methods to generate confidence intervals in practice. In contrast, the minimax method gives rigorous confidence intervals without free parameters.

Table V provides a quick overview of comparison of these different methods.

D. Quantum state verification

Quantum state verification (QSV) is a procedure to certify that the fidelity of the quantum state prepared in the laboratory with a pure target state is greater than some prespecified value with high probability [15,45–47]. QSV corresponds to a hypothesis testing problem, which is related to, but different from, our problem of estimation of fidelity. Nevertheless, the approach adopted by Pallister *et al.* [15] and related and subsequent work [45–48] shares similarities with our fidelity

estimation method, and therefore, a comparison of these methods is helpful.

Suppose that $\rho = |\psi\rangle\langle\psi|$ is the pure target state one wishes to prepare, but one prepares the states $\sigma_1, \sigma_2, \dots, \sigma_R$ in the laboratory. For comparison with our method, we assume that $\sigma_i = \sigma$ are identically prepared states and that measurements are independent. One is given the promise that either $F(\rho, \sigma) = 1$ or $F(\rho, \sigma) \leq 1 - \epsilon_q$ for some $\epsilon_q \in (0, 1)$. If one can reject the hypothesis $F(\rho, \sigma) \leq 1 - \epsilon_q$ with probability $1 - \delta$, then we can infer that $F(\rho, \sigma) = 1$ with a confidence level of $1 - \delta$. The error ϵ_q denotes (a bound on) the deviation of the true fidelity from the maximum fidelity of 1. This is different from the additive error ϵ_* in fidelity estimation, which bounds the deviation of the true fidelity from the estimated fidelity.

The goal of QSV [15] is to obtain a minimax optimal protocol for rejecting the hypothesis $F(\rho, \sigma) \leq 1 - \epsilon_q$. Pallister *et al.* [15] considered random measurement strategies to achieve this goal, and showed that one can obtain the minimax optimal strategy by solving an optimization problem. They demonstrated their method by obtaining a minimax optimal measurement strategy for arbitrary two-qubit states. In particular, to reject $F(\rho, \sigma) \leq 1 - \epsilon_q$ with a confidence level of $1 - \delta$, their strategy needs

$$O\left(\frac{\ln(1/\delta)}{\epsilon_q}\right)$$

samples. The scaling of $1/\epsilon_q$ is optimal for QSV under the aforementioned assumptions. We note that quantum state certification [30] also achieves $O(1/\epsilon_q)$ scaling but using joint measurements, which are difficult to implement with current technology. In contrast, Pallister *et al.* [15] presented a local, single-copy measurement protocol for achieving this scaling for two-qubit states.

We show in Appendix D that by using the same measurement strategy, our method can estimate fidelity to an additive error ϵ_* with a confidence level of $1 - \delta$ by using

$$O\left(\frac{\ln(2/\delta)}{\epsilon_*^2}\right)$$

samples. Note that the scaling of $1/\epsilon_*^2$ is optimal for any fidelity estimation protocol using only independent single-copy

measurements [14]. Thus, we demonstrate for the two-qubit example, that the optimal QSV measurement protocol is also optimal for fidelity estimation using our method.

Notice, however, that the scaling of the optimal sample complexity is different for QSV [$O(1/\epsilon_q)$] and fidelity estimation [$O(1/\epsilon_*^2)$]. Such a difference boils down to the differences in problem setup for QSV and fidelity estimation (for example, the different meanings of ϵ_q and ϵ_*) and the underlying assumptions. For instance, if we remove the assumption of QSV that either $F(\rho, \sigma) = 1$ or $F(\rho, \sigma) \leq 1 - \epsilon_q$, as one would in an actual experiment, the scaling of QSV can be worse than $1/\epsilon_q$ (see Ref. [47] and its Supplementary Material). Despite these differences, the approach adopted in QSV can help with designing optimal measurement protocols that our method can benefit from for fidelity estimation, as demonstrated by the two-qubit example above. Additional details on QSV and how it compares with our method can be found in Appendix D.

E. Classical shadows

Recently, a method called classical shadows was introduced for the estimation of expectation values of observables [14]. We begin by remarking that the philosophy of classical shadows differs from our method. In classical shadows, one seeks to find an optimal estimation procedure, assuming that there is no *a priori* information about the observables whose expectation values need to be computed. In contrast, we seek to find the optimal estimation procedure, given *a priori* knowledge of the observable, which is often known in practice. In the context of fidelity estimation, we know the target state ρ with which we wish to estimate the fidelity.

To obtain a classical shadow, one randomly selects a unitary U from a given ensemble \mathcal{U} according to some probability distribution, rotates the state by U (i.e., $\sigma \rightarrow U\sigma U^\dagger$), and performs a computational basis measurement. If $|b\rangle$ is the computation basis element observed after one such measurement, then in expectation, we obtain a quantum channel \mathcal{M}_c , defined by $\mathbb{E}[U^\dagger|b\rangle\langle b|U] = \mathcal{M}_c(\sigma)$, where σ is the underlying state. If the unitary ensemble is tomographically complete, then the channel \mathcal{M}_c is invertible [14], and we call $\mathcal{M}_c^{-1}(U^\dagger|b\rangle\langle b|U)$ a classical snapshot of σ . After R measurements, the array of classical snapshots $\{\mathcal{M}_c^{-1}(U_1^\dagger|b_1\rangle\langle b_1|U_1), \dots, \mathcal{M}_c^{-1}(U_R^\dagger|b_R\rangle\langle b_R|U_R)\}$ (where U_i is the randomly sampled unitary and $|b_i\rangle$ is the computational basis element observed at the i th time step) is called a classical shadow. Then, given an observable \mathcal{O} , an estimate of the expectation value $\text{Tr}(\mathcal{O}\sigma)$ using a single classical snapshot is given by $\text{Tr}[\mathcal{O}\mathcal{M}_c^{-1}(U_1^\dagger|b_1\rangle\langle b_1|U_1)]$. Multiple estimates are combined using the median of means to give the final estimate [14].

Classical shadows are minimax optimal up to a constant factor in the worst case over all observables, and thus, are suitable to use when the observable is not known *a priori*. Our method is minimax optimal up to a constant factor for a given observable and measurement protocol. Thus, given the observable whose expectation needs to be estimated and the measurement protocol being implemented, our method can match the performance (sample complexity) of classical shadows up to a constant factor.

Both these approaches have their respective advantages. Since classical shadows use a fixed randomized measurement protocol, they can simultaneously estimate N observables to a precision of ϵ with $O[S \ln(N)/\epsilon^2]$ samples, where S is the maximum shadow norm of the observables [14]. However, since the observables are not assumed to be known beforehand, the measurements are not tailored to the observables being estimated. In contrast, for a given observable, we are allowed to choose an appropriate measurement protocol to estimate the expectation value of that observable. As a result, our method, in principle, can use exponentially fewer samples than classical shadows for estimating the expectation value. A simple example is estimating the expectation with a Pauli observable acting nontrivially on all qubits. The prefactor (shadow norm) S scales exponentially using randomized global Clifford measurements as well as randomized Pauli measurements, as was pointed out in Huang *et al.* [14]. As a result, one has to take exponentially many samples to estimate this expectation value using classical shadows. In contrast, measuring in the eigenbasis of the given Pauli observables only needs $O(1/\epsilon^2)$ samples to estimate the expectation to an error of ϵ .

For fidelity estimation with a pure target state ρ , we begin by studying the case when the unitary ensemble \mathcal{U} is the Clifford group \mathcal{C}_n on n qubits, and we sample uniformly at random from \mathcal{C}_n . It was shown [14] that classical shadows only need $O(1/\epsilon^2)$ measurements to estimate the fidelity with ρ to an error of ϵ , which matches the optimal sample complexity up to a constant factor. Owing to minimax optimality, our method can match the sample complexity of classical shadows using global Clifford measurements using $O(1/\epsilon^2)$ measurements. However, since the size of the Clifford group scales as $|\mathcal{C}_n| = 2^{n^2+2n} \prod_{i=1}^n (4^i - 1)$ with the number of qubits n [49], it is computationally challenging to compute our estimate for global Clifford measurements by directly optimizing over POVM elements. To circumvent this problem, we need to simplify the optimization problem involved in computing our estimator analytically, as done, for example, for the randomized Pauli measurement scheme described in Sec. II F. We leave the problem of simplifying the optimization to a future study.

We remark that the computational difficulty with computing estimates is not specific to our method, but also shared by classical shadows. For Clifford measurements, the classical shadow stores the stabilizer states $U_i^\dagger|b_k\rangle\langle b_k|U_i$, which can be done efficiently. However, the target state (or more generally, the observable \mathcal{O}) with which the expectation needs to be computed need not be a stabilizer state (or Clifford unitary). For efficient computation of the estimate $\text{Tr}(\mathcal{O}U_i^\dagger|b_k\rangle\langle b_k|U_i)$, one needs to be able to write the observable as a linear combination of polynomially many Cliffords (or stabilizer states). If this is not possible, there may be no efficient way of storing the observable in memory or computing the estimate. Thus, for many observables and target states of interest, even classical shadows suffers from exponential classical computation time, albeit in a different manner than our method. We remark that for stabilizer states, we have an efficient algorithm to compute our estimator that also achieves optimal sample complexity as described in Sec. II E. We conjecture this algorithm can be generalized to states that can be expressed as a linear

combinations of polynomially many stabilizer states, with, at most, a polynomial overhead.

Note, however, that Clifford measurements can be considered as an exceptional case, especially in the near term. Not only does the size of the Clifford group grow extremely quickly with qubit number, even the experimental complexity for realizing n -qubit Clifford gates grows rapidly, making this a rather impractical approach. In practice, experiments almost always rely on local measurements, such as Pauli measurements.

For fidelity estimation using Pauli measurements, we can get an exponential advantage over classical shadows. A simple example is estimating the fidelity with the target state $|0 \dots 0\rangle$ (computational basis state with all zeros). Since the shadow norm using Pauli measurements scales exponentially with the locality of the observable [14], and because $|0 \dots 0\rangle$ has a nontrivial support on all n qubits, classical shadows need exponentially many samples to estimate the fidelity with this state using random Pauli measurements. In contrast, our method (and even DFE [12]) can estimate the fidelity with $|0 \dots 0\rangle$, using computational basis measurements that does not scale with the number of qubits. This is because the measurement protocol is tailored to the target state, which is known *a priori*.

IV. CONCLUDING REMARKS AND FUTURE RESEARCH

The minimax method can be used to obtain an estimator for the fidelity with a pure target state for any measurement scheme. For a given setting, the estimator only needs to be computed once, and can subsequently be evaluated on raw measurement outcomes instantaneously. This gives our method a practical advantage over other estimation protocols which require random sampling of measurement settings.

Crucially, the minimax method not only constructs an estimator but also provides rigorous confidence intervals that are nearly minimax optimal. We showed that this property translates to practical sample complexity when the measurement scheme is carefully chosen. Notably, the risk is a property of the chosen measurement scheme (including the number of repetitions) and target state and is thus computed before seeing any data. Because the risk is precomputed, it is taken to be symmetric around the fidelity estimate. As a consequence, it might be suboptimal in practice, as the confidence interval might include unphysical values when the estimate is close to 0 or 1.

However, the fact that the risk is known beforehand allows us to use the method for benchmarking experimental protocols without having to take any data. This can, therefore, be useful in guiding the design of experiments. Further, when extending the method to quantum channels using the Choi-Jamiołkowski isomorphism [50–52], such benchmarking can also be done for protocols that estimate gate fidelity.

The computation involved in finding the fidelity estimator is practical for relatively small dimensions, but can become computationally intensive for larger systems because intermediate steps in our algorithm need $O(d^3)$ iterations. Therefore, finding more efficient ways to do the optimization needed to find the estimator would prove very useful in practice. We show that this is possible for a specific measurement

protocol involving Pauli measurements (see Sec. II F), where the optimization is reduced to two dimensions irrespective of the dimension of the system. The only challenge is efficiently computing the Pauli weights for the measurement protocol, but that is not a drawback of the algorithm itself. For example, when these weights can be efficiently computed, as is possible for well-conditioned states, the estimator can be efficiently computed in very large dimensions. It would be interesting to extend such an approach to more general measurement settings to efficiently compute the fidelity estimator. Similarly, approaches to compute the estimator efficiently for randomized Clifford measurements need to be investigated. This problem essentially reduces to efficiently computing the classical fidelity between two probability distributions determined by randomized Clifford measurements on two different states. Furthermore, it would be interesting to see if using conjugate gradient methods for the initial descent and then using accelerated projected gradient methods like Nesterov's method [53] can give significant improvements in performing the optimization for target states and measurement settings. Such an approach was proposed by the authors of Ref. [54] to give a fast algorithm for MLE.

An obvious advantage of the minimax method over tomographic methods like MLE is that the state need not be reconstructed. We showed that the uncertainty given by the Monte Carlo method (for MLE), although typically tighter, need not correspond to a genuine confidence interval as it can be overconfident. PL also gives tighter bounds than the minimax method because it is computed after seeing the data, but we are not aware of any systematic method to obtain a confidence interval using PL. Finding an approach to obtain confidence intervals from PL and similarly for the SDP approach outlined above are interesting directions for future research. Finally, the confidence intervals from the minimax method, though guaranteed to be correct, are often not as tight as they could be. This could potentially be improved by generalizing to an asymmetric risk or by computing the risk after seeing the data, or both.

An open source implementation of the minimax method can be found in Ref. [55].

ACKNOWLEDGMENTS

The authors thank Emanuel Knill, Scott Glancy, Yanbao Zhang, and Rainer Blatt for helpful discussions on the manuscript. This material is based upon work supported by the National Science Foundation under Grants No. 1819251 and No. 2112901. This work utilized the Summit supercomputer, which is supported by the National Science Foundation (Awards No. ACI-1532235 and No. ACI-1532236), the University of Colorado Boulder, and Colorado State University. The Summit supercomputer is a joint effort of the University of Colorado Boulder and Colorado State University. We gratefully acknowledge support by the Austrian Science Fund (FWF Grant-DOI 10.55776/F71) (SFB BeyondC) and the Institut für Quanteninformation GmbH. We also acknowledge funding from the EU H2020-FETFLAG-2018-03 under Grant Agreement No. 820495, by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research

Projects Activity (IARPA), via the U.S. Army Research Office (ARO) Grants No. W911NF-16-1-0070 and No. W911NF-21-1-0007, and the US Air Force Office of Scientific Research (AFOSR) via IOE Grant No. FA9550-19-1-7044 LASCEM. This project received funding from the European Union's Horizon 2020 research and innovation program under the Marie Skłodowska-Curie Grant Agreement No. 840450.

APPENDIX A: MINIMAX METHOD: THEORY

We discuss here the theory surrounding the minimax method for fidelity estimation. We begin by giving a short overview of Juditsky and Nemirovski's framework [24] in Appendix A 1. Then, in Appendix A 2, we describe how we adapt their general method for the purpose of fidelity estimation so as to obtain the procedure described in Sec. II A.

1. Juditsky and Nemirovski's premise

Suppose we are given a set of "states" $\mathcal{X} \subseteq \mathbb{R}^d$ that is a compact and convex subset of \mathbb{R}^d . We wish to estimate the linear functional $\langle g, x \rangle$, where $g \in \mathbb{R}^d$ is some fixed vector, while the state $x \in \mathcal{X}$ of the system is unknown to us. We do not have direct access to the state x . However, we have access to a single measurement outcome determined by the state x . Measurements are modeled using random variables that assign probabilities to the possible outcomes depending on the state.

To that end, Juditsky and Nemirovski [24] consider a family of random variables Z_μ parametrized by $\mu \in \mathcal{M}$, where $\mathcal{M} \subseteq \mathbb{R}^m$ is some set of parameters. These random variables take values in a Polish space² (Ω, Σ) equipped with a σ -finite Borel measure \mathbb{P} [24]. We assume that Z_μ has a probability density p_μ with respect to this reference measure \mathbb{P} . The state $x \in \mathcal{X}$ determines the random variable $Z_{A(x)}$ through an affine function $A : \mathcal{X} \rightarrow \mathcal{M}$, and we are given one outcome of this random variable for the purpose of estimation. Looking ahead to the specialization of this framework to the quantum setting, one can think of the set of observations Ω as a finite set, Borel measurable functions on Ω as functions on Ω , and the integrals $\int_\Omega f(\omega) d\mathbb{P}$ as the finite sum $\sum_{\omega \in \Omega} f(\omega)$.

Our goal is to construct an estimator for $\langle g, x \rangle$ that uses an outcome of the random variable $Z_{A(x)}$ to give an estimate. An estimator is a real-valued Borel measurable function on Ω . The set of estimators \mathcal{F} we are allowed to work with is any finite-dimensional vector space comprised of real-valued Borel measurable functions on Ω as long as it contains constant functions [24]. The mapping $\mathcal{D}(\mu) = p_\mu$ between the parameter μ and the corresponding probability density p_μ is called a parametric density family [24]. To be able to choose an appropriate estimator from \mathcal{F} given that the probability density of the random variable is $p_{A(x)}$, we want that the set of estimators \mathcal{F} and the parametric density family \mathcal{D} interact well with each other. This gives rise to the notion of a good pair defined by Juditsky and Nemirovski [24].

Definition A.1 (Good pair). We call a given pair $(\mathcal{D}, \mathcal{F})$ of parametric density family \mathcal{D} and finite-dimensional space \mathcal{F} of Borel functions on Ω a good pair if the following conditions hold.

- (1) \mathcal{M} is a relatively open convex set in \mathbb{R}^m . By relatively open, we mean $\mathcal{M} = \text{relint}(\mathcal{M}) \equiv \{x \in \mathcal{M} \mid \exists r > 0 \text{ with } B(x, r) \cap \text{aff}(\mathcal{M}) \subseteq \mathcal{M}\}$, where $\text{aff}(\mathcal{M})$ is the affine hull of \mathcal{M} .
- (2) Whenever $\mu \in \mathcal{M}$, we have $p_\mu(\omega) > 0$ for all $\omega \in \Omega$.
- (3) Whenever $\mu, \nu \in \mathcal{M}$, $\phi(\omega) = \ln[p_\mu(\omega)/p_\nu(\omega)] \in \mathcal{F}$.
- (4) Whenever $\phi \in \mathcal{F}$, the function

$$F_\phi(\mu) = \ln \left(\int_\Omega \exp[\phi(\omega)] p_\mu(\omega) d\mathbb{P} \right)$$

is well defined and concave in $\mu \in \mathcal{M}$.

Any estimator in $\widehat{\mathcal{G}} \in \mathcal{F}$ is called an affine estimator (note, however, that $\widehat{\mathcal{G}}$ need not be an affine function).

To judge the performance of an arbitrary estimator $\widehat{\mathcal{G}}$, we define the δ risk as follows [24].

Definition A.2 (δ risk). Given a confidence level $1 - \delta \in (0, 1)$, we define the δ risk associated with an estimator $\widehat{\mathcal{G}}$ as

$$\mathcal{R}(\widehat{\mathcal{G}}; \delta) = \inf \left\{ \epsilon : \sup_{x \in \mathcal{X}} \text{Prob}_{\omega \sim p_{A(x)}} \{ \omega : |\widehat{\mathcal{G}}(\omega) - \langle g, x \rangle| > \epsilon \} < \delta \right\},$$

where $\omega \sim p_{A(x)}$ means that ω is sampled according to $p_{A(x)}$. The corresponding minimax optimal risk is defined as

$$\mathcal{R}_*(\delta) = \inf_{\widehat{\mathcal{G}}} \mathcal{R}(\widehat{\mathcal{G}}; \delta),$$

where the infimum is taken over all Borel functions $\widehat{\mathcal{G}}$ on Ω . Restricting to just the affine estimators, the affine risk is defined as

$$\mathcal{R}_{\text{aff}}(\delta) = \inf_{\widehat{\mathcal{G}} \in \mathcal{F}} \mathcal{R}(\widehat{\mathcal{G}}; \delta).$$

It turns out that we do not lose much by restricting our attention to affine estimators. Indeed, Juditsky and Nemirovski [24] proved that if $(\mathcal{D}, \mathcal{F})$ is a good pair, there is an estimator $\widehat{F}_* \in \mathcal{F}$ with δ risk at most ε_* , such that

$$\mathcal{R}_{\text{aff}}(\delta) \leq \varepsilon_* \leq \vartheta(\delta) \mathcal{R}_*(\delta),$$

$$\vartheta(\delta) = 2 + \frac{\ln(64)}{\ln(0.25/\delta)},$$

for $\delta \in (0, 0.25)$.

The estimator \widehat{F}_* and the risk ε_* are constructed as follows [24,28].

(1) Find the saddle-point value of the function $\Phi : (\mathcal{X} \times \mathcal{X}) \times (\mathcal{F} \times \mathbb{R}_+) \rightarrow \mathbb{R}$ defined as

$$\Phi(x, y; \phi, \alpha) = \langle g, x \rangle - \langle g, y \rangle + 2\alpha \ln(2/\delta)$$

$$+ \alpha \left[\ln \left(\int_\Omega \exp(-\phi(\omega)/\alpha) p_{A(x)}(\omega) d\mathbb{P} \right) + \ln \left(\int_\Omega \exp(\phi(\omega)/\alpha) p_{A(y)}(\omega) d\mathbb{P} \right) \right] \quad (\text{A1})$$

to a given precision. Juditsky and Nemirovski [24] showed that Φ has the following properties. Φ is concave in (x, y)

²A Polish space is a topological space that is homeomorphic to a separable complete metric space. We endow this with the Borel σ algebra.

and convex in (ϕ, α) , and also $\Phi \geq 0$. Further, Φ has a well-defined saddle point. See Ref. [24] for other properties and a more general treatment of the problem.

(2) Denote the saddle-point value of Φ by $2\varepsilon_*$:

$$\begin{aligned} \varepsilon_* &= \frac{1}{2} \sup_{x, y \in \mathcal{X}} \inf_{\phi \in \mathcal{F}, \alpha > 0} \Phi(x, y; \phi, \alpha) \\ &= \frac{1}{2} \inf_{\phi \in \mathcal{F}, \alpha > 0} \max_{x, y \in \mathcal{X}} \Phi(x, y; \phi, \alpha). \end{aligned} \quad (\text{A2})$$

Say the saddle-point value is achieved at some points $x^*, y^* \in \mathcal{X}$, $\phi_* \in \mathcal{F}$, and $\alpha_* > 0$ to a precision $\epsilon > 0$. Suppose that an outcome $\omega \in \Omega$ is observed upon measurement of the random variable $Z_{A(x)}$. Then, the estimator $\widehat{F}_* \in \mathcal{F}$ is given as

$$\widehat{F}_*(\omega) = \phi_*(\omega) + c,$$

where the affine estimator ϕ_* is given by

$$\frac{\phi_*}{\alpha_*} = \frac{1}{2} \ln \left(\frac{P_{A(x^*)}}{P_{A(y^*)}} \right), \quad (\text{A3})$$

and the constant c is

$$c = \frac{1}{2} (\langle g, x^* \rangle + \langle g, y^* \rangle). \quad (\text{A4})$$

The δ risk associated with this estimator satisfies $\mathcal{R}(\widehat{F}_*; \delta) \leq \varepsilon_* + \epsilon$, so the final output of the procedure is $\widehat{F}_*(\omega) \pm (\varepsilon_* + \epsilon)$. See Refs. [24, 28] for details.

Importantly, the estimator \widehat{F}_* is a function that can accept any outcome $\omega \in \Omega$. In other words, the estimate is provided depending on the outcome, but the risk ε_* is computed before seeing any outcome.

So far we described the one-shot scenario, i.e., producing an estimate for $\langle g, x \rangle$ from one outcome of a single random variable $Z_{A(x)}$. In practice, we will need to consider outcomes of different random variables $Z_{A^{(l)}(x)}$, which corresponds to $l = 1, \dots, L$ different measurement settings. More precisely, we are given Polish spaces $(\Omega^{(l)}, \Sigma^{(l)})$ equipped with a σ -finite Borel measure $\mathbb{P}^{(l)}$ for $l = 1, \dots, L$. We are also given a set of parameters $\mathcal{M}^{(l)}$ for $l = 1, \dots, L$. For each $l = 1, \dots, L$, we are given a family of random variables Z_{μ_l} taking values in $\Omega^{(l)}$, where $\mu_l \in \mathcal{M}^{(l)}$. The random variable Z_{μ_l} has a probability density $p_{\mu_l}^{(l)}$ with respect to the reference measure $\mathbb{P}^{(l)}$. As before, we are given affine mappings $A^{(l)}: \mathcal{X} \rightarrow \mathcal{M}^{(l)}$ for $l = 1, \dots, L$ that map the state $x \in \mathcal{X}$ of the system to a corresponding parameter. For each $l = 1, \dots, L$, we can choose estimators for the l th measurement from the set $\mathcal{F}^{(l)}$, which is a finite-dimensional vector space of real-valued Borel measurable functions on $\Omega^{(l)}$ that contains constant functions. To incorporate the outcomes of these different random variables, Juditsky and Nemirovski [24] defined the direct product of good pairs, which essentially constructed one large good pair from many smaller ones.

Definition A.3 (Direct product of good pairs). Considering the following quantities for $l = 1, \dots, L$. Let $(\Omega^{(l)}, \Sigma^{(l)})$ be a Polish space endowed with a Borel σ -finite measure $\mathbb{P}^{(l)}$. Let $\mathcal{D}^{(l)}(\mu_l) = p_{\mu_l}^{(l)}$ be the parametric density family for $\mu_l \in \mathcal{M}^{(l)}$. Let $\mathcal{F}^{(l)}$ be a finite-dimensional linear space of Borel functions on $\Omega^{(l)}$ containing constants, such that the pair $(\mathcal{D}^{(l)}, \mathcal{F}^{(l)})$ is good. Then the direct product of these good pairs $(\mathcal{D}, \mathcal{F}) = \bigotimes_{l=1}^L (\mathcal{D}^{(l)}, \mathcal{F}^{(l)})$ is defined as follows.

TABLE VI. A dictionary specifying the meaning of each quantity appearing in the text for the purpose of fidelity estimation. The index l varies from 1 to L , where L denotes the number of measurement settings. The l th measurement setting is described by the POVM $\{E_1^{(l)}, \dots, E_{N_l}^{(l)}\}$.

\mathcal{X}	Set of density matrices
$\Omega^{(l)}$	Measurement outcomes $\{1, \dots, N_l\}$
$\mathbb{P}^{(l)}$	Counting measure on $(\Omega^{(l)}, \Sigma^{(l)})$ with $\Sigma^{(l)} = 2^{\Omega^{(l)}}$
$\mathcal{M}^{(l)}$	Relatively open simplex $\{x \in \mathbb{R}^{N_l} \mid x_i > 0, \sum_i x_i = 1\}$
p_{μ}	$p_{\mu} = (\mu_1, \dots, \mu_{N_l}), \mu \in \mathcal{M}^{(l)}$
$A^{(l)}$	$A^{(l)}(x)_k = \frac{\text{Tr}(E_k^{(l)} \chi) + \epsilon_0 / N_l}{1 + \epsilon_0}, k = 1, \dots, N_l, \epsilon_0 > 0$
$\mathcal{F}^{(l)}$	Set of estimators: real-valued functions on $\Omega^{(l)}$
g	Pure target state ρ

(1) The large space is $\Omega = \Omega^{(1)} \times \dots \times \Omega^{(L)}$ endowed with the product measure $\mathbb{P} = \mathbb{P}^{(1)} \times \dots \times \mathbb{P}^{(L)}$.

(2) The set of parameters is $\mathcal{M} = \mathcal{M}^{(1)} \times \dots \times \mathcal{M}^{(L)}$, and the associated parametric density family is $\mathcal{D}(\mu) = p_{\mu} \equiv \prod_{l=1}^L p_{\mu_l}^{(l)}$ for $\mu = (\mu_1, \dots, \mu_L) \in \mathcal{M}$.

(3) The linear space \mathcal{F} comprises of all functions ϕ defined as $\phi(\omega_1, \omega_2, \dots, \omega_L) = \sum_{l=1}^L \phi^{(l)}(\omega_l)$, where $\phi^{(l)} \in \mathcal{F}^{(l)}$ and $\omega_l \in \Omega^{(l)}$ for $l = 1, \dots, L$.

It can be verified that the direct product of good pairs is a good pair [24]. Therefore, we can apply the above procedure for constructing an optimal estimator to the direct product of good pairs to obtain an optimal estimator that accounts for all the given measurement outcomes.

We now note a simplification of the risk ε_* obtained using results of Juditsky and Nemirovski [24]. Recall that we defined the risk ε_* as half the saddle-point value of the function Φ [see Eq. (A2)]. However, this definition is not very amenable to theoretical calculations. Therefore, we note an alternate expression for the risk, given by Proposition 3.1 of Juditsky and Nemirovski [24]

$$\varepsilon_* = \frac{1}{2} \max_{x, y \in \mathcal{X}} \left\{ \langle g, x \rangle - \langle g, y \rangle \mid \text{AffH}[A(x), A(y)] \geq \frac{\delta}{2} \right\}. \quad (\text{A5})$$

Juditsky and Nemirovski [24] referred to the quantity

$$\text{AffH}(\mu, \nu) = \int \sqrt{p_{\mu} p_{\nu}} d\mathbb{P} \quad (\text{A6})$$

as the Hellinger affinity (this quantity is sometimes referred to as the Bhattacharyya coefficient in the discrete case; see, for example, Ref. [56]). Juditsky and Nemirovski [24] proved that the Hellinger affinity is continuous and log-concave in $(\mu, \nu) \in \mathcal{M} \times \mathcal{M}$ when $(\mathcal{D}, \mathcal{F})$ is a good pair.

2. Fidelity estimation

Now we give the details on how we adapt Juditsky and Nemirovski's general framework to fidelity estimation. In Table VI, we provide a dictionary relating the general quantities defined in Appendix A1 to our scenario of fidelity estimation. We assume that the quantum system has a d -dimensional Hilbert space over \mathbb{C} , $d \in \mathbb{N}$. The set of $d \times d$

complex-valued Hermitian matrices forms a real-vector space that is isomorphic to \mathbb{R}^{d^2} , and thus, we can write $\mathcal{X} \subseteq \mathbb{R}^{d^2}$, where \mathcal{X} is the set of density matrices. Note that \mathcal{X} is a compact and convex set. For the l th measurement setting, we consider the positive operator-valued measure (POVM) $\{E_1^{(l)}, \dots, E_{N_l}^{(l)}\}$, where $l = 1, \dots, L$. Since the definition of a good pair requires that the probability of each outcome is nonzero, we add a small parameter $0 < \epsilon_o \ll 1$ to make the outcome probabilities positive, as noted in Sec. II A. These outcome probabilities are represented by the affine map $A^{(l)} : \mathcal{X} \rightarrow \mathcal{M}^{(l)}$ given in Table VI.

Note that the set of outcomes $\Omega^{(l)}$ is a finite set for each $l = 1, \dots, L$. We consider the discrete topology $\Sigma^{(l)} = 2^{\Omega^{(l)}}$ on $\Omega^{(l)}$, so that $(\Omega^{(l)}, \Sigma^{(l)})$ forms a Polish space. The Borel σ algebra coincides with the topology, so we use the same symbol for both. Because $\Sigma^{(l)}$ is discrete, any real-valued function on $\Omega^{(l)}$ is Borel measurable, and therefore, we omit the phrase ‘‘Borel measurable’’ when talking about functions (or estimators) on $\Omega^{(l)}$. Furthermore, since each $\Omega^{(l)}$ is a finite set, real-valued functions defined on it can be considered as $|\Omega^{(l)}| = N_l$ -dimensional real vectors. Thus, we treat elements of $\mathcal{F}^{(l)}$ as vectors.

Using these facts, we show that $\mathcal{D}^{(l)}$ and $\mathcal{F}^{(l)}$ as defined in Table VI form a good pair. By definition of $\mathcal{M}^{(l)}$ and p_μ in Table VI, it is easy to see the first and second conditions for a good pair given in Definition A.1 hold. We check that

$$\begin{aligned} \Phi(\chi_1, \chi_2; \phi, \alpha) &= \text{Tr}(\rho \chi_1) - \text{Tr}(\rho \chi_2) + 2\alpha \ln(2/\delta) \\ &+ \alpha \sum_{l=1}^L \sum_{r=1}^{R_l} \left[\ln \left(\sum_{k=1}^{N_l} e^{-\phi_k^{(l,r)}/\alpha} \frac{\text{Tr}(E_k^{(l)} \chi_1) + \epsilon_o/N_l}{1 + \epsilon_o} \right) + \ln \left(\sum_{k=1}^{N_l} e^{\phi_k^{(l,r)}/\alpha} \frac{\text{Tr}(E_k^{(l)} \chi_2) + \epsilon_o/N_l}{1 + \epsilon_o} \right) \right], \end{aligned}$$

where we use the fact that $\exp(\sum_{l,r} \phi^{(l,r)}) = \prod_{l,r} \exp(\phi^{(l,r)})$, $p_\mu = \prod_{l,r} p_{\mu_{l,r}}^{(l)}$, $A(\chi) = \bigoplus_{l,r} A^{(l)}(\chi)$, and that $\mathbb{P} = (\mathbb{P}^{(1)})^{R_1} \times \dots \times (\mathbb{P}^{(L)})^{R_L}$ is a product measure. By Remark 3.2 of Juditsky and Nemirovski [24], we can just use R_l copies of $\phi^{(l)}$ instead of $\phi^{(l,r)}$ in finding the saddle point of the above function, and thus we obtain Eq. (4).

Suppose that the (χ_1, χ_2) and the α components of the saddle point of the function Φ are attained at $\chi_1^*, \chi_2^* \in \mathcal{X}$, and $\alpha_* > 0$, respectively, to a given precision. Then, from Eq. (A3), dDefinition A.3, and the preceding remarks, we can infer that the ϕ component of the saddle point can be described by the function $\phi_* = \sum_{l=1}^L \sum_{r=1}^{R_l} \phi_*^{(l)}$, where for $l = 1, \dots, L$, we have

$$\phi_*^{(l)} = \frac{\alpha_*}{2} \ln \left(\frac{p_{A^{(l)}(\chi_1^*)}}{p_{A^{(l)}(\chi_2^*)}} \right). \quad (\text{A7})$$

Note that replacing $g = \rho$ with $g = \mathcal{O}$ for any Hermitian operator (observable) \mathcal{O} , we can obtain an estimator for the expectation value of that observable.

Finally, to adapt the simplified expression for risk given in Eq. (A5) to the quantum case, we compute the Hellinger affinity [see Eq. (A6)]. The Hellinger affinity for the quantum

the last two conditions given in Definition A.1 hold. Since $\mathcal{F}^{(l)}$ contains all functions on $\Omega^{(l)}$, it contains $\ln(p_\mu/p_\nu)$ in particular. Next, we see that for $\phi^{(l)} \in \mathcal{F}^{(l)}$, we can write

$$F_{\phi^{(l)}}(\mu) = \ln \left(\sum_{k=1}^{N_l} \exp(\phi_k^{(l)}) \mu_k \right)$$

because $\mathbb{P}^{(l)}$ is the counting measure. As noted above, we consider $\phi^{(l)} \in \mathcal{F}^{(l)}$ as an N_l -dimensional real vector. Then, since $\nabla_\mu^2 F_{\phi^{(l)}} = -\mathbf{e}\mathbf{e}^T / ((\mathbf{e}, \mu))^2$, where $\mathbf{e} = (e^{\phi_1^{(l)}}, \dots, e^{\phi_{N_l}^{(l)}})$, $F_{\phi^{(l)}}$ is concave in μ . Thus, $(\mathcal{D}^{(l)}, \mathcal{F}^{(l)})$ forms a good pair.

Now suppose that we perform R_l repetitions (shots) of the l th measurement setting. Then, the space to be considered for all measurement settings put together is $\Omega = (\Omega^{(1)})^{R_1} \times \dots \times (\Omega^{(L)})^{R_L}$, the parameter space for probability distributions is $\mathcal{M} = (\mathcal{M}^{(1)})^{R_1} \times \dots \times (\mathcal{M}^{(L)})^{R_L}$, and the set of estimators $\mathcal{F} \subseteq (\mathcal{F}^{(1)})^{R_1} \times \dots \times (\mathcal{F}^{(L)})^{R_L}$ is chosen as per Definition A.3. The mapping $A : \mathcal{X} \rightarrow \mathcal{M}$ is given by $A(\chi) = \bigoplus_{l=1}^L \bigoplus_{r=1}^{R_l} A^{(l)}(\chi) \equiv (A^{(1)}(\chi), A^{(1)}(\chi), \dots, A^{(L)}(\chi), A^{(L)}(\chi))$, where $A^{(l)}(\chi)$ is repeated R_l times. Then, we can use the direct product of good pairs (Definition A.3) to compute the function Φ defined in Eq. (A1) when all the measurement settings are considered together.

Note that $\phi \in \mathcal{F} \subseteq (\mathcal{F}^{(1)})^{R_1} \times \dots \times (\mathcal{F}^{(L)})^{R_L}$ implies $\phi = \sum_{l=1}^L \sum_{r=1}^{R_l} \phi^{(l,r)}$, where $\phi^{(l,r)}$ belongs to the r th copy of $\mathcal{F}^{(l)}$. Then, using Eq. (A1), we obtain

case is given as

$$\begin{aligned} \text{AffH}[A(\chi_1), A(\chi_2)] &= \prod_{l=1}^L \left[\sum_{k=1}^{N_l} \left(\frac{\text{Tr}(E_k^{(l)} \chi_1) + \epsilon_o/N_l}{1 + \epsilon_o} \right)^{1/2} \right. \\ &\quad \times \left. \left(\frac{\text{Tr}(E_k^{(l)} \chi_2) + \epsilon_o/N_l}{1 + \epsilon_o} \right)^{1/2} \right]^{R_l} \\ &\approx \prod_{l=1}^L [F_C(\chi_1, \chi_2; \{E_k^{(l)}\})]^{R_l/2}, \quad (\text{A8}) \end{aligned}$$

where, in the last step, we neglect $\epsilon_o \ll 1$ to simplify the equations (we include ϵ_o in the numerical simulations). Here F_C is the classical fidelity defined in Eq. (10). Then substituting the expression for the Hellinger affinity in Eq. (A5), we obtain Eq. (9).

APPENDIX B: MINIMAX METHOD: NUMERICAL IMPLEMENTATION

We outline the procedure followed to find the saddle point of the function Φ defined in Eq. (4), from which we can compute the fidelity estimator.

We present two algorithms to compute the saddle point. The first algorithm is based on Nesterov's method (accelerated projected gradient), and the second algorithm uses standard CVXPY software. The implementation based on Nesterov's method is more memory efficient, whereas the implementation based on CVXPY is faster in practice. A detailed comparison of the computation time required by each method can be found in the accompanying Ref. [27]. We now describe both our algorithms and show that they converge to the saddle point.

1. Implementation using Nesterov's method

This is done in two steps: find the $\chi_1^*, \chi_2^* \in \mathcal{X}$, and $\alpha_* > 0$ components of the saddle point and then compute ϕ_*/α_* to obtain the $\phi_* \in \mathcal{F}$ component of the saddle point.

To execute the first step, we resort to the following expression for the saddle-point value noted in Ref. [28], which is obtained by appropriately rewriting the expression given by Juditsky and Nemirovski [24]

$$2\varepsilon_* = \inf_{\alpha > 0} \left\{ 2\alpha \ln(2/\delta) + \max_{\chi_1, \chi_2 \in \mathcal{X}} [\langle \rho, \chi_1 \rangle - \langle \rho, \chi_2 \rangle + 2\alpha \ln\{\text{AffH}[A(\chi_1), A(\chi_2)]\}] \right\}, \quad (\text{B1})$$

where the Hellinger affinity $\text{AffH}[A(\chi_1), A(\chi_2)]$ is given in Eq. (A8).

Now we note that, as per the definitions used in Table VI, the function $\text{AffH}(\mu, \nu) = \prod_{l=1}^L [\sum_{i=1}^{N_l} \sqrt{\mu_i^{(l)} \nu_i^{(l)}}]^{R_l}$ is smooth on its domain because $\mu, \nu > 0$ (component-wise) for each $\mu, \nu \in \mathcal{M}$. Since $\text{AffH}(\mu, \nu) > 0$, the function $\ln[\text{AffH}(\mu, \nu)]$ is well defined and smooth on its domain. Further, since $A : \mathcal{X} \rightarrow \mathcal{M}$ is affine, $\ln\{\text{AffH}[A(\chi_1), A(\chi_2)]\}$ is smooth on $\mathcal{X} \times \mathcal{X}$. In particular, the derivatives of $\ln\{\text{AffH}[A(\chi_1), A(\chi_2)]\}$ are continuous. Since $\mathcal{X} \times \mathcal{X}$ is compact, the Hessian of $\ln\{\text{AffH}[A(\chi_1), A(\chi_2)]\}$ is bounded, and therefore, the gradient of $\ln\{\text{AffH}[A(\chi_1), A(\chi_2)]\}$ is Lipschitz continuous. Moreover, $\ln\{\text{AffH}[A(\chi_1), A(\chi_2)]\}$ is a jointly concave function of the density matrices (see Appendix A 1).

With this in mind, we use the following procedure to find a saddle point of the function Φ defined in Eq. (4) to any given precision.

(1) For any fixed $\alpha > 0$, we solve the ‘‘inner’’ convex optimization problem in Eq. (B1)

$$\begin{aligned} & \max_{\chi_1, \chi_2 \in \mathcal{X}} \left[\underbrace{\text{Tr}(\rho\chi_1) - \text{Tr}(\rho\chi_2) + 2\alpha \ln\{\text{AffH}[A(\chi_1), A(\chi_2)]\}}_{f(\chi_1, \chi_2)} \right] \\ & = \max_{\chi \in \bar{\mathcal{X}}} f(\chi), \end{aligned}$$

by using the version of Nesterov's second method [53] given in Ref. [57]. For the second equation above, we define $\chi = (\chi_1, \chi_2)$ and $\bar{\mathcal{X}} = \mathcal{X} \times \mathcal{X}$. Nesterov's second method is suited to problems where the objective f is a convex function³ with a Lipschitz continuous gradient (see Theorem 1(c) in Ref. [57] for the convergence guarantee). In such scenarios, Nesterov's second method gives an accelerated version of projected gradient ascent or descent, such that each iterate lies

in the domain $\bar{\mathcal{X}}$. This is a useful method to optimize convex functions of density matrices. When the Lipschitz constant is not known, a backtracking scheme can be used [57].

(2) We perform the ‘‘outer’’ convex optimization over α in Eq. (B1) using SCIPY's MINIMIZE_SCALAR routine. Through this optimization, we obtain the $\chi_1^*, \chi_2^* \in \mathcal{X}$, and $\alpha_* > 0$ components of the saddle point.

(3) Using the so-obtained $\chi_1^*, \chi_2^* \in \mathcal{X}$, and $\alpha_* > 0$, we find $\phi_* \in \mathcal{F}$ using Eq. (A7).

Once we have a saddle point, the estimator can be easily computed using Eqs. (6) and (7).

Note that we embed the Hermitian matrices into a real vector space before performing the above optimizations. This is possible because an isometric isomorphism exists between the set of Hermitian operators of a fixed size and a real vector space. We use $\epsilon_o = 10^{-5}$ in the numerical implementation.

2. Implementation with CVXPY

The optimal risk ε_* as defined by the constrained optimization problem in Eq. (A5) can be evaluated directly by using the CVXPY PYTHON library, provided we take a logarithm of the constraint to make the problem convex and multiply it by 2 to match the theory of Ref. [24]. This approach is faster than the above saddle-point algorithm, at the cost of increased memory usage. A CVXPY implementation with a more balanced trade-off for larger-dimensional problems is possible in the case of Pauli measurements by applying each POVM with a matrix-free algorithm.

Regardless of such implementation details, however, one needs access to the optimal α_* from the saddle-point problem in Eq. (4) to construct the estimator \hat{F}_* . To do this without solving the saddle-point problem directly, we show that the parameter α of the saddle-point problem is actually the dual variable associated with our constraint. The optimal value of this dual variable is automatically returned by CVXPY alongside the optimal primal variables χ_1^* and χ_2^* as a consequence of the primal-dual algorithm with which it solves the problem.

Written in a standard form, Eq. (A5) becomes

$$\begin{aligned} & \underset{\chi_1, \chi_2 \in \mathcal{X}}{\text{minimize}} && \langle \rho, \chi_2 - \chi_1 \rangle, \\ & \text{subject to} && 2 \ln(\epsilon/2) - 2 \ln\{\text{AffH}[A(\chi_1), A(\chi_2)]\} \leq 0, \end{aligned} \quad (\text{B2})$$

where the primal optimal value p^* is related to the risk ε_* by $-\frac{1}{2}p^* = \varepsilon_*$. Introducing $\tilde{\alpha}$ as a dual variable, the Lagrangian for our nonlinear programming problem is given by

$$\begin{aligned} \mathcal{L}(\chi_1, \chi_2; \tilde{\alpha}) &= \langle \rho, \chi_2 - \chi_1 \rangle + \tilde{\alpha} (2 \ln(\epsilon/2) \\ &\quad - 2 \ln\{\text{AffH}[A(\chi_1), A(\chi_2)]\}). \end{aligned}$$

This defines a concave Lagrange dual function $g : \mathbb{R} \rightarrow \mathbb{R}$ defined by

$$g(\tilde{\alpha}) = \min_{\chi_1, \chi_2 \in \mathcal{X}} \mathcal{L}(\chi_1, \chi_2; \tilde{\alpha}).$$

Let d^* denote the maximum value of the function g . Because our original problem in Eq. (A5) is convex and Slater's condition holds (as strict feasibility is achieved whenever $\chi_1 = \chi_2$), we have strong duality between the two problems. This means

³Or a concave function in the case of maximization.

that $p^* = d^*$, and it follows that

$$\begin{aligned} 2\varepsilon_* &= -p^* = -d^* = -\max_{\tilde{\alpha} \geq 0} g(\tilde{\alpha}) \\ &= \min_{\tilde{\alpha} \geq 0} \max_{\chi_1, \chi_2 \in \mathcal{X}} -\mathcal{L}(\chi_1, \chi_2; \tilde{\alpha}) \\ &= \min_{\tilde{\alpha} \geq 0} \max_{\chi_1, \chi_2 \in \mathcal{X}} \{ \langle \rho, \chi_1 - \chi_2 \rangle \\ &\quad + 2\tilde{\alpha} [\ln\{\text{AffH}[A(\chi_1), A(\chi_2)]\} - \ln(\epsilon/2)] \} \\ &= \inf_{\tilde{\alpha} > 0} \left\{ 2\tilde{\alpha} \ln(2/\epsilon) + \max_{\chi_1, \chi_2 \in \mathcal{X}} [\langle \rho, \chi_1 \rangle - \langle \rho, \chi_2 \rangle \right. \\ &\quad \left. + 2\tilde{\alpha} \ln\{\text{AffH}[A(\chi_1), A(\chi_2)]\}] \right\}. \end{aligned}$$

With $\tilde{\alpha} = \alpha$, this is exactly the expression in Eq. (B1), showing the equivalence of the two values. Importantly, this means that the estimator found by solving the saddle-point problem in Eq. (B1) can be equivalently constructed by finding both the primal and dual optima of Eq. (B2), which we compute using CVXPY.

APPENDIX C: MINIMAX METHOD: SAMPLE COMPLEXITY

We begin by computing the best sample complexity that can be achieved by the minimax method. A detailed statement of this result is given in Theorem II.1. Below, we present a proof of this result.

Proof of Theorem II.1. From Eq. (9), we know that the risk can be written as

$$\varepsilon_* = \frac{1}{2} \max_{\chi_1, \chi_2 \in \mathcal{X}} \left\{ \left| \text{Tr}(\rho\chi_1) - \text{Tr}(\rho\chi_2) \right| \right. \\ \left. \prod_{l=1}^L [F_C(\chi_1, \chi_2, \{E_k^{(l)}\})]^{R_l/2} \geq \frac{\delta}{2} \right\},$$

where

$$F_C(\chi_1, \chi_2, \{E_k^{(l)}\}) = \left(\sum_{k=1}^{N_l} \sqrt{\text{Tr}(E_k^{(l)}\chi_1)\text{Tr}(E_k^{(l)}\chi_2)} \right)^2$$

is the classical fidelity between χ_1 and χ_2 determined by the POVM $\{E_k^{(l)}\}$. As noted in Sec. IID, we can write the fidelity between any two states as follows [29]:

$$F(\chi_1, \chi_2) = \min_{\text{POVM}\{F_i\}} F_C(\chi_1, \chi_2, \{F_i\}).$$

In particular, we have $F(\chi_1, \chi_2) \leq F_C(\chi_1, \chi_2, \{E_k^{(l)}\})$ for every POVM $\{E_k^{(l)}\}$ that we are using. Thus, we obtain the following lower bound on our risk:

$$\varepsilon_* \geq \frac{1}{2} \max_{\chi_1, \chi_2 \in \mathcal{X}} \left\{ \left| \text{Tr}(\rho\chi_1) - \text{Tr}(\rho\chi_2) \right| \right. \\ \left. F(\chi_1, \chi_2) \geq \left(\frac{\delta}{2} \right)^{\frac{2}{R}} \right\},$$

where $R = \sum_{l=1}^L R_l$ is the total number of shots. We now proceed to evaluating the lower bound. For convenience, we denote $\gamma = (\delta/2)^{2/R}$ so that the constraint becomes

$F(\chi_1, \chi_2) \geq \gamma$. Next we note that, for a pure state ρ and possibly mixed states χ_1 and χ_2 , we have the following inequality for the fidelity in terms of the trace distance (see Chap. 9 in Ref. [58]):

$$\text{Tr}(\rho\chi_1) \leq \text{Tr}(\rho\chi_2) + \frac{1}{2} \|\chi_1 - \chi_2\|_1, \quad (\text{C1})$$

where $\|\chi\|_1$ is the Schatten 1-norm of χ . Then, using the Fuchs-van de Graaf inequality $(1/2)\|\chi_1 - \chi_2\|_1 \leq \sqrt{1 - F(\chi_1, \chi_2)}$ [56], we can write

$$\begin{aligned} \text{Tr}(\rho\chi_1) - \text{Tr}(\rho\chi_2) &\leq \sqrt{1 - F(\chi_1, \chi_2)} \\ &\leq \sqrt{1 - \gamma}, \end{aligned} \quad (\text{C2})$$

where the second line holds when $F(\chi_1, \chi_2) \geq \gamma$.

We show that the upper bound in Eq. (C2) can be achieved by explicitly constructing the density matrices χ_1^* and χ_2^* achieving the maximum. For this purpose, we define $\Delta_\rho = \mathbb{I} - \rho$, and suppose that the dimension of the system is d (i.e., ρ is an $d \times d$ matrix). Then, let

$$\begin{aligned} \chi_1^* &= \frac{1 + \sqrt{1 - \gamma}}{2} \rho + \frac{1 - \sqrt{1 - \gamma}}{2} \frac{\Delta_\rho}{d - 1}, \\ \chi_2^* &= \frac{1 - \sqrt{1 - \gamma}}{2} \rho + \frac{1 + \sqrt{1 - \gamma}}{2} \frac{\Delta_\rho}{d - 1}. \end{aligned}$$

Since ρ is pure, there is some (normalized) vector $|v_1\rangle$ such that $\rho = |v_1\rangle\langle v_1|$. Let $\{|v_2\rangle, \dots, |v_d\rangle\}$ be any orthonormal basis for the orthogonal complement of the subspace spanned by $|v_1\rangle$. Then, we can write $\Delta_\rho = \sum_{i=2}^d |v_i\rangle\langle v_i|$ using the resolution of identity. Therefore, in the basis $\{|v_1\rangle, \dots, |v_d\rangle\}$, the matrices χ_1^* and χ_2^* are diagonal. Since $\gamma < 1$, the diagonal entries of these matrices are real (and positive), and it is easy to check that they sum to 1, showing that χ_1^* and χ_2^* are density matrices. Since they are diagonal, it is easy to compute the fidelity between them. We find that $F(\chi_1^*, \chi_2^*) = \gamma$, and therefore, these density matrices satisfy the constraint $F(\chi_1^*, \chi_2^*) \geq \gamma$. Further, we can see that these density matrices saturate the upper bound in Eq. (C2), i.e., $\text{Tr}(\rho\chi_1^*) - \text{Tr}(\rho\chi_2^*) = \sqrt{1 - \gamma}$. Thus, we find that the lower bound on the risk is

$$\varepsilon_* \geq \frac{1}{2} \sqrt{1 - \left(\frac{\delta}{2} \right)^{2/R}},$$

where we use $\gamma = (\delta/2)^{2/R}$.

The inequality given in the above equation is tight: the POVM defined by $\{\rho, \Delta_\rho\}$ achieves the lower bound (see Corollary C.1.1). Thus, the best sample complexity given by the minimax method corresponding to a risk of $\varepsilon_* < 0.5$ and confidence level $1 - \delta \in (0.75, 1)$ is

$$\begin{aligned} R &\geq \frac{2 \ln(2/\delta)}{|\ln(1 - 4\varepsilon_*^2)|} \\ &\approx \frac{\ln(2/\delta)}{2\varepsilon_*^2} \text{ when } \varepsilon_*^2 \ll 1, \end{aligned}$$

as noted. ■

We now consider a family of two-outcome POVM measurements, and show that the risk given by the minimax

method can be obtained by solving a one-dimensional optimization problem. Using this, we obtain a simple formula for an upper bound on the risk, and consequently, also a good bound on the sample complexity. In particular, this provides an upper bound on the sample complexity for the randomized Pauli measurement scheme given in Scheme II.1.

Theorem C.1. Suppose we are given a pure target state ρ , and we perform R repetitions of the POVM $\{\Theta, \Delta_\Theta\}$ defined as

$$\begin{aligned}\Theta &= \omega_1 \rho + \omega_2 \Delta_\rho, \\ \Delta_\Theta &= (1 - \omega_1) \rho + (1 - \omega_2) \Delta_\rho,\end{aligned}$$

where $\Delta_\rho = \mathbb{I} - \rho$ and $\omega_1, \omega_2 \in [0, 1]$ are parameters satisfying $\omega_1 > \omega_2$. Also, define

$$\gamma = \left(\frac{\delta}{2}\right)^{2/R},$$

and

$$R_o = \frac{\ln(2/\delta)}{|\ln(\sqrt{\omega_1 \omega_2} + \sqrt{(1 - \omega_1)(1 - \omega_2)})|}.$$

Then, if $R > R_o$, the risk of the estimator given by the minimax method can be obtained by solving the one-dimensional optimization problem

$$\varepsilon_* = \frac{\sqrt{1 - \gamma}}{2(\omega_1 - \omega_2)} \max_{a \in \mathcal{A}_a} \sqrt{1 - (2a - 1)^2 \gamma}, \quad (\text{C3})$$

where the set of allowed values for a is given as

$$\begin{aligned}\mathcal{A}_a &= [0, 1] \cap ((-\infty, a_-^{(1)}] \cup [a_+^{(1)}, \infty)) \\ &\cap ((-\infty, a_-^{(2)}] \cup [a_+^{(2)}, \infty)),\end{aligned}$$

$$\begin{aligned}\text{AffH}(\alpha_1, \alpha_2) &= \frac{1}{1 + \epsilon_o} \left(\omega_1 \alpha_1 + \omega_2 (1 - \alpha_1) + \frac{\epsilon_o}{2} \right)^{1/2} \left(\omega_1 \alpha_2 + \omega_2 (1 - \alpha_2) + \frac{\epsilon_o}{2} \right)^{1/2} \\ &+ \frac{1}{1 + \epsilon_o} \left((1 - \omega_1) \alpha_1 + (1 - \omega_2) (1 - \alpha_1) + \frac{\epsilon_o}{2} \right)^{1/2} \left((1 - \omega_1) \alpha_2 + (1 - \omega_2) (1 - \alpha_2) + \frac{\epsilon_o}{2} \right)^{1/2} \\ &\approx \sqrt{(\omega_2 + (\omega_1 - \omega_2) \alpha_1)(\omega_2 + (\omega_1 - \omega_2) \alpha_2)} + \sqrt{[(1 - \omega_2) + (\omega_2 - \omega_1) \alpha_1][(1 - \omega_2) + (\omega_2 - \omega_1) \alpha_2]}.\end{aligned} \quad (\text{C5})$$

Note that because of the parameter $\epsilon_o > 0$, the Hellinger affinity is differentiable. Since $\epsilon_o \ll 1$, we neglect it in Eq. (C5) to prevent the equations from becoming cumbersome later. We can write the risk as

$$\begin{aligned}2\varepsilon_* &= \max_{\alpha_1, \alpha_2 \in [0, 1]} (\alpha_1 - \alpha_2) \\ \text{s.t. } &-\ln[\text{AffH}(\alpha_1, \alpha_2)] \leq -\ln(\sqrt{\gamma}).\end{aligned} \quad (\text{C6})$$

We take a logarithm to make the optimization problem convex (see Proposition 3.1 in Ref. [24]).

Now, consider the case $R > R_o$, where R_o is as defined in the statement of the theorem. Then, we argue that at the optimum, $\text{AffH} = \sqrt{\gamma}$. To see this, we first convert the above maximization to a minimization problem, and write its

with

$$\begin{aligned}a_\pm^{(1)} &= \omega_1 \pm \sqrt{\omega_1(1 - \omega_1) \frac{(1 - \gamma)}{\gamma}}, \\ a_\pm^{(2)} &= \omega_2 \pm \sqrt{\omega_2(1 - \omega_2) \frac{(1 - \gamma)}{\gamma}}.\end{aligned}$$

For $R \leq R_o$, the risk is $\varepsilon_* = 0.5$.

In particular, for any risk $\varepsilon_* \in (0, 0.5)$,

$$\begin{aligned}R &\geq 2 \frac{\ln(2/\delta)}{|\ln(1 - 4(\omega_1 - \omega_2)^2 \varepsilon_*^2)|} \\ &\approx \frac{1}{2(\omega_1 - \omega_2)^2} \frac{\ln(2/\delta)}{\varepsilon_*^2}\end{aligned} \quad (\text{C4})$$

repetitions of the measurement are sufficient to achieve that risk with a confidence level of $1 - \delta \in (3/4, 1)$.

Proof. For the case that we have R repetitions of a single POVM $\{\Theta, \Delta_\Theta\}$, the risk can be written as (see Appendix A)

$$\varepsilon_* = \frac{1}{2} \max_{\chi_1, \chi_2 \in \mathcal{X}} \{|\text{Tr}(\rho \chi_1) - \text{Tr}(\rho \chi_2)|$$

$$\text{AffH}[A(\chi_1), A(\chi_2)] \geq \sqrt{\gamma}\},$$

where $A(\chi) = ([\text{Tr}(\Theta \chi) + \epsilon_o/2]/(1 + \epsilon_o), [\text{Tr}(\Delta_\Theta \chi) + \epsilon_o/2]/(1 + \epsilon_o))$. To begin with, we simplify this to a two-dimensional optimization problem. For this purpose, we write the density matrices χ_1, χ_2 as a convex combination of the target state ρ and some other trace-one Hermitian operator in the orthogonal complement of the subspace generated by ρ :

$$\chi_1 = \alpha_1 \rho + (1 - \alpha_1) \rho_1^\perp,$$

$$\chi_2 = \alpha_2 \rho + (1 - \alpha_2) \rho_2^\perp,$$

with $\text{Tr}(\rho \rho_1^\perp) = \text{Tr}(\rho \rho_2^\perp) = 0$ and $0 \leq \alpha_1, \alpha_2 \leq 1$. Using this, the Hellinger affinity can be written as

Lagrangian as

$$\begin{aligned}\mathcal{L} &= -\alpha_1 + \alpha_2 - \lambda \ln[\text{AffH}(\alpha_1, \alpha_2)] + \lambda \ln(\sqrt{\gamma}) \\ &- v_0^1 \alpha_1 + v_1^1 (\alpha_1 - 1) - v_0^2 \alpha_2 + v_1^2 (\alpha_2 - 1),\end{aligned}$$

where $\lambda, v_0^1, v_1^1, v_0^2, v_1^2$ are dual variables. At the optimum, the Karush-Kuhn-Tucker (KKT) conditions must be satisfied [59], which we list below for convenience.

(1) (Primal feasibility) The ‘‘primal’’ variables α_1, α_2 must lie in the domain $[0, 1]$.

(2) (Dual feasibility) The dual variables (corresponding to inequality constraints) $\lambda, v_0^1, v_1^1, v_0^2, v_1^2$ must be nonnegative.

(3) (Complementary slackness) Either the dual variable must vanish or the constraint must be tight.

(4) (Stationarity) The gradient of the Lagrangian with respect to the primal variables must vanish.

Since the gradient of \mathcal{L} with respect to α_1 and α_2 must vanish at the optimum, we have

$$\begin{aligned}\frac{\partial \mathcal{L}}{\partial \alpha_1} &= -1 - \lambda \frac{\partial \ln(\text{AffH})}{\partial \alpha_1} - v_0^1 + v_1^1 = 0, \\ \frac{\partial \mathcal{L}}{\partial \alpha_2} &= 1 - \lambda \frac{\partial \ln(\text{AffH})}{\partial \alpha_2} - v_0^2 + v_1^2 = 0.\end{aligned}$$

$$\begin{aligned}\sqrt{[\omega_2 + (\omega_1 - \omega_2)\alpha_1][\omega_2 + (\omega_1 - \omega_2)\alpha_2]} &\equiv a\sqrt{\gamma}, \\ \sqrt{[(1 - \omega_2) + (\omega_2 - \omega_1)\alpha_1][(1 - \omega_2) + (\omega_2 - \omega_1)\alpha_2]} &\equiv b\sqrt{\gamma}.\end{aligned}$$

Then, $\text{AffH} = \sqrt{\gamma}$ implies

$$a + b = 1,$$

where $a, b \geq 0$. From the above equations, we can deduce that

$$\begin{aligned}\alpha_1 + \alpha_2 &= \frac{(a^2 - b^2)}{\omega_1 - \omega_2} \gamma + \frac{(1 - 2\omega_2)}{\omega_1 - \omega_2}, \\ \alpha_1 \alpha_2 &= \frac{[(1 - \omega_2)a^2 + \omega_2 b^2]}{(\omega_1 - \omega_2)^2} \gamma - \frac{\omega_2(1 - \omega_2)}{(\omega_1 - \omega_2)^2}.\end{aligned}$$

These equations are well defined because $\omega_1 > \omega_2$. Solving these simultaneous equations and applying the constraint $a + b = 1$, we obtain

$$\begin{aligned}\alpha_1 &= \frac{(2a - 1)\gamma + (1 - 2\omega_2)}{2(\omega_1 - \omega_2)} \\ &\quad + \frac{\sqrt{1 - \gamma}}{2(\omega_1 - \omega_2)} \sqrt{1 - (2a - 1)^2 \gamma}, \\ \alpha_2 &= \frac{(2a - 1)\gamma + (1 - 2\omega_2)}{2(\omega_1 - \omega_2)} \\ &\quad - \frac{\sqrt{1 - \gamma}}{2(\omega_1 - \omega_2)} \sqrt{1 - (2a - 1)^2 \gamma}.\end{aligned}$$

Since $a \in [0, 1]$, $(2a - 1)^2 \in [0, 1]$, and $\gamma \in (0, 1)$, the term in the square-root is nonnegative, so α_1, α_2 are real. Furthermore, we use the fact that $\alpha_1 \geq \alpha_2$ at the optimum because the risk involves maximization of $\alpha_1 - \alpha_2$; see Eq. (C6).

Next, we need to impose the constraints $\alpha_1, \alpha_2 \in [0, 1]$. Requiring $\alpha_2 \geq 0$ (and thus $\alpha_1 \geq 0$) gives

$$\begin{aligned}(a - a_+^{(2)})(a - a_-^{(2)}) &\geq 0, \\ a_{\pm}^{(2)} &= \omega_2 \pm \sqrt{\omega_2(1 - \omega_2) \frac{(1 - \gamma)}{\gamma}},\end{aligned}$$

which means a must lie in the region $(-\infty, a_-^{(2)}) \cup [a_+^{(2)}, \infty)$. Similarly, requiring $\alpha_1 \leq 1$ (and thus $\alpha_2 \leq 1$) gives

$$(a - a_+^{(1)})(a - a_-^{(1)}) \geq 0,$$

Dual feasibility implies $\lambda, v_0^1, v_1^1, v_0^2, v_1^2 \geq 0$ at the optimum. If $\lambda = 0$ at the optimum, we must have $v_1^1 = 1 + v_0^1 > 0$ and $v_0^2 = 1 + v_1^2 > 0$. Then complementary slackness implies that $\text{AffH} \geq \sqrt{\gamma}$, $\alpha_1 = 1$ and $\alpha_2 = 0$ at the optimum. However, for $R > R_o$, we have

$$\begin{aligned}\text{AffH}(\alpha_1 = 1, \alpha_2 = 0) &= \sqrt{\omega_1 \omega_2} + \sqrt{(1 - \omega_1)(1 - \omega_2)} \\ &< \sqrt{\gamma},\end{aligned}$$

contradicting with the constraint. Thus, we must have $\lambda > 0$, implying that $\text{AffH} = \sqrt{\gamma}$ as claimed. Using this, we can reduce the problem to a one-dimensional problem that will eventually help perform the optimization. We do this by appropriately parametrizing each term in AffH :

$$a_{\pm}^{(1)} = \omega_1 \pm \sqrt{\omega_1(1 - \omega_1) \frac{(1 - \gamma)}{\gamma}},$$

which implies that a must lie in the region $(-\infty, a_-^{(1)}) \cup [a_+^{(2)}, \infty)$. Therefore, the allowed values of a are

$$\begin{aligned}\mathcal{A}_a &= [0, 1] \cap ((-\infty, a_-^{(2)}) \cup [a_+^{(2)}, \infty)) \\ &\quad \cap ((-\infty, a_-^{(1)}) \cup [a_+^{(1)}, \infty)).\end{aligned}$$

Note that the optimization problem defined by Eq. (C6) has a solution for all $R > 0$ [i.e., $\gamma \in (0, 1)$] because any $\alpha_1, \alpha_2 \in [0, 1]$ with $\alpha_1 = \alpha_2$ satisfies the constraints. Therefore, we must have $\mathcal{A}_a \neq \emptyset$.

Thus, the risk is given as

$$\begin{aligned}\varepsilon_* &= \max_{a \in \mathcal{A}_a} \frac{1}{2} (\alpha_1 - \alpha_2) \\ &= \frac{\sqrt{1 - \gamma}}{2(\omega_1 - \omega_2)} \max_{a \in \mathcal{A}_a} \sqrt{1 - (2a - 1)^2 \gamma}.\end{aligned}$$

Now, for the case when $R \leq R_o$, we have $\sqrt{\gamma} \leq \sqrt{\omega_1 \omega_2} + \sqrt{(1 - \omega_1)(1 - \omega_2)} = \text{AffH}(\alpha_1 = 1, \alpha_2 = 0)$. Therefore, $\alpha_1 = 1$ and $\alpha_2 = 0$ satisfy the constraint of the optimization in Eq. (C6), giving $\varepsilon_* = 0.5$.

The last part of the statement of the theorem follows from the observation that Eq. (C4) implies $R > R_o$ when $\varepsilon_* < 0.5$, and for $R > R_o$, we have the inequality $\varepsilon_* \leq \sqrt{1 - \gamma} / [2(\omega_1 - \omega_2)]$.

Note that there is no loss of generality in requiring that $\omega_1 > \omega_2$, for if $\omega_2 < \omega_1$, we can simply swap Θ and Δ_{Θ} . When $\omega_1 = \omega_2$, we have $\Theta = \Delta_{\Theta} = \mathbb{I}/2$, which means we learn nothing about the state. Indeed $R_o \rightarrow \infty$ as $\omega_1 \rightarrow \omega_2$, alluding to this fact. However, the best we can do is when ω_1 and ω_2 are farthest from each other, and this leads to the following result.

Corollary C.1.1. Let ρ be any pure target state, and $\Delta_{\rho} = \mathbb{I} - \rho$. Then, for R repetitions of the POVM $\{\rho, \Delta_{\rho}\}$, the

estimator given by minimax method achieves the risk

$$\varepsilon_* = \frac{1}{2} \sqrt{1 - \left(\frac{\delta}{2}\right)^{2/R}}.$$

Proof. We have $\omega_1 = 1$ and $\omega_2 = 0$. Substituting this in Theorem C.1, we can see that $a_{\pm}^{(1)} = 1$ and $a_{\pm}^{(2)} = 0$. Therefore, the allowed values of a are $\mathcal{A}_a = [0, 1]$. Subsequently, the risk is given as

$$\begin{aligned} \varepsilon_* &= \frac{\sqrt{1-\gamma}}{2} \max_{a \in [0,1]} \sqrt{1 - (2a-1)^2 \gamma} \\ &= \frac{\sqrt{1-\gamma}}{2}, \end{aligned}$$

as claimed.

We also consider a more restricted family of POVMs that are relevant to the stabilizer measurements described in Sec. II E.

Corollary C.1.2. Suppose we are given a pure target state ρ , and we perform R repetitions of the POVM $\{\Theta, \Delta_{\Theta}\}$ defined as

$$\begin{aligned} \Theta &= \rho + \frac{\xi/2-1}{\xi-1} \Delta_{\rho}, \\ \Delta_{\Theta} &= \frac{\xi/2}{\xi-1} \Delta_{\rho}, \end{aligned}$$

where $\Delta_{\rho} = \mathbb{I} - \rho$ and $\xi \geq 2$ is a parameter. Also, define

$$R_o = 2 \frac{\ln(2/\delta)}{\ln\left(\frac{\xi-1}{\xi/2-1}\right)}.$$

Then, if $R > R_o$, the risk of the estimator given by the minimax method is

$$\varepsilon_* = \begin{cases} \left(\frac{\xi-1}{\xi}\right) \sqrt{1-\gamma}, & b_- \geq 1, \\ \left(\frac{\xi-1}{\xi}\right) (1-\gamma) \sqrt{1+b_-(2-b_-)\left(\frac{\gamma}{1-\gamma}\right)}, & b_- < 1, \\ \left(\frac{\xi-1}{\xi}\right) (1-\gamma) \sqrt{1+b_+(2-b_+)\left(\frac{\gamma}{1-\gamma}\right)}, & b_- < 1, \\ \left(\frac{\xi-1}{\xi}\right) (1-\gamma) \sqrt{1+b_+(2-b_+)\left(\frac{\gamma}{1-\gamma}\right)}, & |b_- - 1| > |b_+ - 1|, \end{cases} \quad (C7)$$

where

$$\gamma = \left(\frac{\delta}{2}\right)^{2/R},$$

and

$$b_{\pm} = \left(\frac{\xi}{\xi-1}\right) \left[1 \pm \sqrt{\left(\frac{1-\gamma}{\gamma}\right) \left(\frac{\xi-2}{\xi}\right)} \right].$$

For $R \leq R_o$, the risk is $\varepsilon_* = 0.5$.

In particular, for any risk $\varepsilon_* \in (0, 0.5)$,

$$\begin{aligned} R &\geq 2 \frac{\ln(2/\delta)}{|\ln(1 - (\frac{\xi}{\xi-1})^2 \varepsilon_*^2)|} \\ &\approx 2 \left(\frac{\xi-1}{\xi}\right)^2 \frac{\ln(2/\delta)}{\varepsilon_*^2} \end{aligned} \quad (C8)$$

repetitions of the measurement are sufficient to achieve that risk with a confidence level of $1 - \delta \in (3/4, 1)$.

Proof. In the context of Theorem C.1, we have

$$\omega_1 = 1 \quad \text{and} \quad \omega_2 = \frac{\xi/2-1}{\xi-1}.$$

This implies $a_{\pm}^{(1)} = 1$, and also

$$a_{\pm}^{(2)} = \frac{\xi/2-1}{\xi-1} \pm \frac{\xi/2}{\xi-1} \sqrt{\frac{(1-\gamma)\xi-2}{\gamma\xi}}.$$

For convenience, we perform the change of variables $b = 2(1-a)$, and define

$$\begin{aligned} b_{\pm} &\equiv 2(1-a_{\mp}^{(2)}) \\ &= \left(\frac{\xi}{\xi-1}\right) \left[1 \pm \sqrt{\left(\frac{1-\gamma}{\gamma}\right) \left(\frac{\xi-2}{\xi}\right)} \right]. \end{aligned}$$

Clearly, $b_- \leq b_+$. Note that $R > R_o$ implies $b_- > 0$, which means $a_+ < 1$. Thus, we have that $\mathcal{A}_a = [0, 1] \cap [(-\infty, a_+^{(2)}) \cup [a_+^{(2)}, \infty)] = [0, a_-^{(2)}] \cup [a_+^{(2)}, 1]$ is nonempty. With respect to the variable b , these allowed values can be expressed as $\mathcal{A}_b = [0, b_-] \cup [b_+, 2]$. The risk is then given as

$$\begin{aligned} \varepsilon_* &= \left(\frac{\xi-1}{\xi}\right) \sqrt{1-\gamma} \max_{b \in \mathcal{A}_b} \sqrt{1 - (1-b)^2 \gamma} \\ &= \left(\frac{\xi-1}{\xi}\right) (1-\gamma) \max_{b \in \mathcal{A}_b} \sqrt{1 + b(2-b) \left(\frac{\gamma}{1-\gamma}\right)}. \end{aligned}$$

Since the objective of maximization is symmetric about $b = 1$ and the maximum is achieved at $b = 1$, the allowed value of b closest to 1 achieves the maximum. Noting that $b_+ > 1$, we obtain the expression given in the statement of the corollary. ■

Finally, we obtain a bound on the sample complexity of the randomized Pauli measurement scheme described in Scheme II.1. The statement is given in Theorem II.2, so we simply give a proof.

Proof of Theorem II.2. For convenience, we reproduce below the measurement strategy given in Scheme II.1. Given a target state ρ , do the following:

(1) Sample a (nonidentity) Pauli operator W_i with probability

$$p_i = \frac{|\text{Tr}(W_i \rho)|}{\sum_{i=1}^{d^2-1} |\text{Tr}(W_i \rho)|}, \quad i = 1, \dots, d^2-1,$$

and record the outcome (± 1) of the measurement.

(2) Flip the measurement outcome (i.e., $\pm 1 \rightarrow \mp 1$) if $\text{Tr}(\rho W_i) < 0$, or else retain the original measurement outcome.

(3) Repeat this procedure R times.

Note that measuring the Pauli W_i and flipping the measurement outcome is equivalent to measuring the operator $S_i = \text{sign}[\text{Tr}(W_i \rho)] W_i$.

We can describe the above measurement strategy using the effective POVM described below, which we obtain by finding a positive-semi-definite operator that reproduces the measurement statistics. The probability of obtaining +1 outcome can

be written as

$$\begin{aligned} \Pr(+1) &= \sum_{i=1}^{d^2-1} (\text{Pr. choosing } S_i) \\ &\quad (\text{Pr. outcome 1 upon measuring } S_i) \\ &\equiv \text{Tr}(\Theta\sigma), \\ \Theta &= \sum_{i=1}^{d^2-1} p_i \mathbb{P}_i^+, \\ \mathbb{P}_i^+ &= \frac{\mathbb{I} + S_i}{2}, \end{aligned}$$

where $S_i = \text{sign}[\text{Tr}(W_i\rho)]W_i$. Substituting for p_i and a simple rearrangement of terms gives

$$\begin{aligned} \Theta &= \frac{\mathbb{I}}{2} + \frac{d}{2\mathcal{N}} \sum_{i=1}^{d^2-1} \frac{|\text{Tr}(W_i\rho)|}{d} S_i \\ &= \frac{[d + (\mathcal{N} - 1)]}{2\mathcal{N}} \rho + \frac{(\mathcal{N} - 1)}{2\mathcal{N}} \Delta_\rho. \end{aligned}$$

$\Delta_\Theta = \mathbb{I} - \Theta$ is given as

$$\Delta_\Theta = \frac{[(\mathcal{N} + 1) - d]}{2\mathcal{N}} \rho + \frac{(\mathcal{N} + 1)}{2\mathcal{N}} \Delta_\rho.$$

The effective POVM is then $\{\Theta, \Delta_\Theta\}$. Substituting

$$\omega_1 = \frac{[d + (\mathcal{N} - 1)]}{2\mathcal{N}} \quad \text{and} \quad \omega_2 = \frac{(\mathcal{N} - 1)}{2\mathcal{N}}$$

in Theorem C.1, we obtain Eq. (20).

Now, note that for any pure state ρ , we have $\text{Tr}(\rho^2) = 1$ and this gives us

$$\sum_{i=1}^{d^2-1} [\text{Tr}(W_i\rho)]^2 = d - 1. \quad (\text{C9})$$

Therefore, to obtain an upper bound on \mathcal{N} , we solve the following optimization problem. Let M be a positive integer and $\beta > 0$ be any real number such that $M > \beta$. We solve

$$\begin{aligned} \max \quad & \sum_{i=1}^M x_i \\ \text{s.t.} \quad & x_i \in [0, 1] \quad \forall i = 1, \dots, M \\ & \sum_{i=1}^M x_i^2 \leq \beta. \end{aligned}$$

The above problem gives a bound for the special case of $x_i = |\text{Tr}(W_i\rho)|$, $M = d^2 - 1$ and $\beta = d - 1$, while saving the trouble of optimizing over all density matrices. Note that we consider the relaxation $\sum_{i=1}^M x_i^2 \leq \beta$, instead of the original equality constraint $\sum_{i=1}^M x_i^2 = \beta$ because quadratic equality constraints do not define a convex set in general. Such a relaxation is inconsequential because equality is attained at the optimum. It can be shown using the KKT conditions that

the optimum corresponds to

$$\begin{aligned} x_i &= \sqrt{\frac{\beta}{M}} \quad \forall i \in \{1, \dots, M\} \\ \Rightarrow \sum_{i=1}^M x_i &= \sqrt{M\beta}. \end{aligned}$$

Returning to the problem of Pauli weights, since $M = d^2 - 1$ and $\beta = d - 1$, we obtain $\mathcal{N} \leq \sqrt{(d^2 - 1)(d - 1)} = \sqrt{d + 1}(d - 1)$ as claimed. Substituting in the expression for sample complexity gives the desired upper bound.

In contrast to the above result, which shows that a good sample complexity can be obtained for the randomized Pauli measurement scheme, we consider the case of a bad measurement protocol. Namely, we are given an n -qubit stabilizer state and we measure $n - 1$ of its generators, where the measurements are subspace measurements. Then, as noted in Proposition II.2, the minimax method gives a risk of 0.5. Here, we present a proof for this statement.

Proof of Proposition II.2. Let ρ be an n -qubit stabilizer state generated by S_1, \dots, S_n . The measurement protocol corresponds to measuring only the first $n - 1$ generators S_1, \dots, S_{n-1} . The measurement of S_l has the POVM $\{E_1^{(l)}, E_2^{(l)}\}$ where $E_1^{(l)}$ is the projection on the $+1$ eigenspace of S_l while $E_2^{(l)}$ is the projection on -1 eigenspace of S_l , for $l = 1, \dots, n - 1$. Suppose that the l th measurement is repeated R_l times.

From Eq. (9), we know that the risk of the minimax method can be written as

$$\begin{aligned} \varepsilon_* &= \frac{1}{2} \max_{\chi_1, \chi_2 \in \mathcal{X}} \left\{ \left| \text{Tr}(\rho\chi_1) - \text{Tr}(\rho\chi_2) \right| \right. \\ &\quad \left. \prod_{l=1}^{n-1} [F_C(\chi_1, \chi_2, \{E_k^{(l)}\})]^{R_l/2} \geq \frac{\delta}{2} \right\}, \end{aligned}$$

where

$$F_C(\chi_1, \chi_2, \{E_k^{(l)}\}) = \left(\sum_{k=1}^2 \sqrt{\text{Tr}(E_k^{(l)}\chi_1)\text{Tr}(E_k^{(l)}\chi_2)} \right)^2$$

is the classical fidelity corresponding to the POVM $\{E_1^{(l)}, E_2^{(l)}\}$ for $l = 1, \dots, n - 1$. Our strategy is to construct two density matrices χ_1 and χ_2 that satisfy the constraints of the optimization defining the risk, such that the value of the risk is 0.5.

To that end, let $\tilde{\rho}$ be the stabilizer state generated by $S_1, \dots, -S_n$, where, the last generator of $\tilde{\rho}$ differs from that of ρ by a negative sign. Note that the states ρ and $\tilde{\rho}$ are orthogonal to each other. Observe that the classical fidelity between the states ρ and $\tilde{\rho}$ corresponding to the POVM $\{E_1^{(l)}, E_2^{(l)}\}$ is $F_C(\rho, \tilde{\rho}, \{E_1^{(l)}, E_2^{(l)}\}) = 1$ for all measured stabilizers since $\text{Tr}(E_1^{(l)}\rho) = \text{Tr}(E_1^{(l)}\tilde{\rho}) = 1$ while $\text{Tr}(E_2^{(l)}\rho) = \text{Tr}(E_2^{(l)}\tilde{\rho}) = 0$ for all $l = 1, \dots, n - 1$.

Thus, taking $\chi_1 = \rho$ and $\chi_2 = \tilde{\rho}$, we find that the risk is $\varepsilon_* = 0.5$, which is the maximum possible value for the risk. In other words, when an insufficient number of stabilizer measurements are provided, the minimax method infers that the fidelity cannot be estimated accurately. ■

APPENDIX D: COMPARISON WITH QSV

We begin by giving a brief overview of the measurement strategy used in QSV [15]. QSV assumes that either $F(\rho, \sigma) = 1$ or $F(\rho, \sigma) \leq 1 - \varepsilon_q$ holds for a fixed $\varepsilon_q \in (0, 1)$. The goal is to reject the hypothesis $F(\rho, \sigma) \leq 1 - \varepsilon_q$. To that end, QSV protocol assumes that we have access to a set of operators $\mathcal{S} = \{P_1, \dots, P_L\}$ with $0 \leq P_i \leq \mathbb{I}$ describing two-outcome POVM $\{P_i, \mathbb{I} - P_i\}$. Pallister *et al.* [15] assumed in addition that $\text{Tr}(P_i \rho) = 1$ for all i . The protocol proceeds by randomly sampling the operator P_i with probability μ_i and measuring the POVM $\{P_i, \mathbb{I} - P_i\}$. If the outcome corresponding to P_i is observed, then one repeats this procedure. However, if one observes the outcome corresponding to $\mathbb{I} - P_i$, then one stops the procedure and declares “fail.” If, after many repetitions of this procedure the protocols does not fail, one can conclude with high probability that $\text{Tr}(\rho\sigma) = 1$.

The above measurement protocol can be described by the effective POVM $\{\Omega, \mathbb{I} - \Omega\}$ with $\Omega = \sum_{i=1}^L \mu_i P_i$ which satisfies $\text{Tr}(\Omega\rho) = 1$. QSV searches for an optimal measurement strategy Ω by minimizing the probability of wrongly declaring “pass” in the worst-case scenario:

$$\min_{\Omega} \max_{\sigma} \text{Tr}(\Omega\sigma) = 1 - \Delta_{\varepsilon_q}. \quad (\text{D1})$$

Consequently, when running the protocol with R states, the false acceptance probability for such an optimal Ω is bounded above by $(1 - \Delta_{\varepsilon_q})^R$. Since the probability of wrongly declaring “fail” is zero by the assumption that $\text{Tr}(\rho\Omega) = 1$, we can conclude

$$R \geq \frac{\ln(1/\delta)}{|\ln(1 - \Delta_{\varepsilon_q})|} \approx \frac{\ln(1/\delta)}{\Delta_{\varepsilon_q}} \quad (\text{D2})$$

measurement outcomes are sufficient to certify that $F(\rho, \sigma) > 1 - \varepsilon_q$ with probability $1 - \delta$. Note that such a measurement protocol is minimax optimal in the sense that one finds the best measurement strategy Ω for the worst-possible state σ satisfying $F(\rho, \sigma) \leq 1 - \varepsilon_q$, under the POVM options \mathcal{S} and the QSV assumption that $F(\rho, \sigma) \notin (1 - \varepsilon_q, 1)$.

The above description shows that the minimax approach of QSV bears similarities with our approach towards fidelity estimation. Below, we discuss two differences between QSV and our method (in addition to those already pointed out in Sec. III D).

First, we note that the assumption that the state σ prepared in the laboratory either satisfies $F(\rho, \sigma) = 1$ or $F(\rho, \sigma) \leq 1 - \varepsilon_q$ might be too stringent in practice. For example, if we reject $F(\rho, \sigma) \leq 1 - \varepsilon_q$, we cannot automatically conclude that the fidelity is 1. Our study makes no such assumption, which makes it more practically amenable. As mentioned earlier, some studies on QSV have relaxed this assumption, but at the cost of increasing the sample complexity compared to the original QSV guarantees [47].

Second, the assumption that the measurement operators P_i satisfy $\text{Tr}(\Omega P_i) = 1$ can be restrictive in practice, even though they allow for designing optimal measurement protocols in theory. For example, this assumption does not hold if one wishes to estimate (or certify) the fidelity of a W state using Pauli measurements. This is a practically relevant problem which can be tackled using our approach [27]. Moreover,

our fidelity estimation method is not restricted to protocols involving random sampling of measurement settings, which is again important for practical applications. That said, when one presents an optimal protocol involving random sampling, we can usually adjust the sampling probabilities so that our fidelity estimation method provides similar optimality results. This was shown, for example, in Sec. II F, which matches the performance guarantees of DFE.

We show that a similar result can be obtained with our method for estimating the fidelity of two-qubit states using the measurement protocol given by Pallister *et al.* [15] without changing the sampling probabilities. Using the equivalence of states under rotation by local unitaries, any two-qubit state can be expressed in the form $|\psi\rangle = \sin(\theta)|00\rangle + \cos(\theta)|11\rangle$ for $\theta \in [0, \pi/2]$ [15]. Then, the optimal measurement strategy for QSV can be expressed as sampling the projectors $\{P_1, \dots, P_4\}$ with probability $\mu_1 = \alpha(\theta)$, $\mu_2 = \mu_3 = \mu_4 = [1 - \alpha(\theta)]/3$ [$\theta \in (0, \pi/2) \setminus \{\pi/4\}$] and measuring them [15]. The projectors are given by $P_1 = |00\rangle\langle 00| + |11\rangle\langle 11|$ and $P_i = \mathbb{I} - |\phi_i\rangle\langle \phi_i|$ for $i = 2, 3, 4$, where the states $|\phi_i\rangle$ and the number $\alpha(\theta)$ are given in Eq. (D4). Pallister *et al.* [15] showed that

$$R \approx [2 + \sin(\theta)\cos(\theta)] \frac{\ln(1/\delta)}{\varepsilon_q} \quad (\text{D3})$$

measurements of the operator $\Omega = \sum_{i=1}^4 \mu_i P_i$ suffice to reject $F(\rho, \sigma) \leq 1 - \varepsilon_q$ with a confidence level of $1 - \delta$. The following result shows that by using this measurement protocol, our method performs fidelity estimation in an optimal manner.

Proposition D.1. Let $|\psi\rangle = \sin(\theta)|00\rangle + \cos(\theta)|11\rangle$, where $\theta \in (0, \pi/2)$, $\theta \neq \pi/4$, denote any two-qubit state up to rotation by local unitaries. Consider the measurement protocol where the projectors P_1, P_2, P_3, P_4 are sampled as per probability $\mu_1 = \alpha(\theta)$, $\mu_2 = \mu_3 = \mu_4 = [1 - \alpha(\theta)]/3$ and their outcomes are recorded. Here, $P_1 = |00\rangle\langle 00| + |11\rangle\langle 11|$, $P_i = \mathbb{I} - |\phi_i\rangle\langle \phi_i|$ for $i = 2, 3, 4$, with

$$\begin{aligned} |\phi_1\rangle &= \left(\frac{1}{\sqrt{1 + \tan\theta}}|0\rangle + \frac{e^{\frac{2\pi i}{3}}}{\sqrt{1 + \cot\theta}}|1\rangle \right) \\ &\otimes \left(\frac{1}{\sqrt{1 + \tan\theta}}|0\rangle + \frac{e^{\frac{\pi i}{3}}}{\sqrt{1 + \cot\theta}}|1\rangle \right), \\ |\phi_2\rangle &= \left(\frac{1}{\sqrt{1 + \tan\theta}}|0\rangle + \frac{e^{\frac{4\pi i}{3}}}{\sqrt{1 + \cot\theta}}|1\rangle \right) \\ &\otimes \left(\frac{1}{\sqrt{1 + \tan\theta}}|0\rangle + \frac{e^{\frac{5\pi i}{3}}}{\sqrt{1 + \cot\theta}}|1\rangle \right), \\ |\phi_3\rangle &= \left(\frac{1}{\sqrt{1 + \tan\theta}}|0\rangle + \frac{1}{\sqrt{1 + \cot\theta}}|1\rangle \right) \\ &\otimes \left(\frac{1}{\sqrt{1 + \tan\theta}}|0\rangle - \frac{1}{\sqrt{1 + \cot\theta}}|1\rangle \right), \text{ and} \\ \alpha(\theta) &= \frac{2 - \sin(2\theta)}{4 + \sin(2\theta)}. \end{aligned} \quad (\text{D4})$$

Then,

$$R \geq 2 \frac{\ln(2/\delta)}{\left| \ln \left(1 - \frac{4}{(2 + \sin(\theta) \cos(\theta))^2 \varepsilon_*^2} \right) \right|} \approx \frac{(2 + \sin(\theta) \cos(\theta))^2 \ln(2/\delta)}{2 \varepsilon_*^2} \text{ when } \varepsilon_* \ll 1 \quad (\text{D5})$$

repetitions of the measurement protocol is sufficient to reach attain a risk of ε_* with a confidence level of $1 - \delta \in (0.75, 1)$.

Proof. The effective measurement operator is given by [15]

$$\begin{aligned} \Omega &= \alpha P_1 + \frac{(1 - \alpha)}{3} \sum_{i=2}^4 P_i \\ &= \alpha P_1 + (1 - \alpha) \Omega_3, \end{aligned}$$

where

$$\Omega_3 = \mathbb{I} - \frac{1}{(1 + t)^2} \begin{pmatrix} 1 & 0 & 0 & -t \\ 0 & t & 0 & 0 \\ 0 & 0 & t & 0 \\ -t & 0 & 0 & t^2 \end{pmatrix}$$

and $t = \tan(\theta)$. By diagonalizing Ω , we find that

$$\Omega = \rho + \frac{2 + \sin(2\theta)}{4 + \sin(2\theta)} (\mathbb{I} - \rho).$$

Then, from Theorem C.1, Eq. (D5) follows. We remark that our risk is invariant under unitary rotations of the target state $\rho \rightarrow U \rho U^\dagger$ and the POVM $E_k \rightarrow U E_k U^\dagger$ as seen from Eq. (9). Therefore, the above result essentially holds for any two-qubit state. ■

We exclude the points $\theta = 0, \pi/4, \pi/2$ in the above result as either $\tan(\theta)$ or $\cot(\theta)$ become undefined at these points. For $\theta = 0, \pi/2$, we obtain a product state, and therefore, one can directly measure Pauli ZZ. For $\theta = \pi/4$, we obtain a stabilizer state, for which we have an optimal measurement strategy (see Sec. II E). These observations were also made by Pallister *et al.* [15].

Note that the number of measurements given in Eq. (D5) required by our method is similar to Eq. (D3), except for a prefactor of order 1 and scaling of $1/\varepsilon_*^2$ with the risk ε_* . This scaling is essentially optimal, as seen from Theorem II.1, and thus, we can use their measurement protocol along with our fidelity estimation method to optimally estimate the fidelity of two-qubit states. This reinforces the expectation that good measurement schemes lead to good sample complexity using our method. Conversely, if our method cannot give a good sample complexity for some measurement protocol, then owing to the definition of minimax optimal risk and Eq. (2), we can infer that no estimation method can give a good sample complexity for that measurement protocol under the same assumptions as our method.

-
- [1] A. Ekert and R. Jozsa, *Rev. Mod. Phys.* **68**, 733 (1996).
[2] I. M. Georgescu, S. Ashhab, and F. Nori, *Rev. Mod. Phys.* **86**, 153 (2014).
[3] N. Gisin and R. Thew, *Nat. Photonics* **1**, 165 (2007).
[4] C. L. Degen, F. Reinhard, and P. Cappellaro, *Rev. Mod. Phys.* **89**, 035002 (2017).
[5] V. Giovannetti, S. Lloyd, and L. Maccone, *Nat. Photonics* **5**, 222 (2011).
[6] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, England, 2010).
[7] K. Vogel and H. Risken, *Phys. Rev. A* **40**, 2847 (1989).
[8] Z. Hradil, *Phys. Rev. A* **55**, R1561(R) (1997).
[9] R. Blume-Kohout, *New J. Phys.* **12**, 043034 (2010).
[10] D. Gross, Y.-K. Liu, S. T. Flammia, S. Becker, and J. Eisert, *Phys. Rev. Lett.* **105**, 150401 (2010).
[11] M. Cramer, M. B. Plenio, S. T. Flammia, R. Somma, D. Gross, S. D. Bartlett, O. Landon-Cardinal, D. Poulin, and Y.-K. Liu, *Nat. Commun.* **1**, 149 (2010).
[12] S. T. Flammia and Y.-K. Liu, *Phys. Rev. Lett.* **106**, 230501 (2011).
[13] M. P. da Silva, O. Landon-Cardinal, and D. Poulin, *Phys. Rev. Lett.* **107**, 210404 (2011).
[14] H.-Y. Huang, R. Kueng, and J. Preskill, *Nat. Phys.* **16**, 1050 (2020).
[15] S. Pallister, N. Linden, and A. Montanaro, *Phys. Rev. Lett.* **120**, 170502 (2018).
[16] D. Gottesman, Ph.D. thesis, California Institute of Technology, Pasadena, CA, 1997.
[17] M. Curty, M. Lewenstein, and N. Lütkenhaus, *Phys. Rev. Lett.* **92**, 217903 (2004).
[18] S. J. van Enk, N. Lütkenhaus, and H. J. Kimble, *Phys. Rev. A* **75**, 052318 (2007).
[19] F. G. S. L. Brandao, *Phys. Rev. A* **72**, 022310 (2005).
[20] M. Bourennane, M. Eibl, C. Kurtsiefer, S. Gaertner, H. Weinfurter, O. Gühne, P. Hyllus, D. Bruß, M. Lewenstein, and A. Sanpera, *Phys. Rev. Lett.* **92**, 087902 (2004).
[21] M. Guță, J. Kahn, R. Kueng, and J. A. Tropp, *J. Phys. A: Math. Theor.* **53**, 204001 (2020).
[22] M. Cerezo, A. Poremba, L. Cincio, and P. J. Coles, *Quantum* **4**, 248 (2020).
[23] A. Gilyén and A. Poremba, [arXiv:2203.15993](https://arxiv.org/abs/2203.15993).
[24] A. B. Juditsky and A. S. Nemirovski, *Ann. Math. Stat.* **37**, 2278 (2009).
[25] A. Gendenshluger, A. Juditsky, and A. Nemirovski, *Electron. J. Statist.* **9**, 1645 (2015).
[26] A. Juditsky and A. Nemirovski, *IMA Inf. Inference* **9**, 423 (2019).
[27] A. Seshadri, M. Ringbauer, J. Spainhour, R. Blatt, T. Monz, and S. Becker, companion paper, *Phys. Rev. Lett.* **133**, 020402 (2024).
[28] A. Seshadri and S. Becker, [arXiv:2112.03390](https://arxiv.org/abs/2112.03390).
[29] C. A. Fuchs, Distinguishability and accessible information in quantum theory, Ph.D. thesis, University of New Mexico, Albuquerque, 1995.
[30] C. Bădescu, R. O'Donnell, and J. Wright, in *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing* (ACM, New York, 2019), pp. 503–514.
[31] J. Haah, A. W. Harrow, Z. Ji, X. Wu, and N. Yu, in *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing* (ACM, New York, 2016), pp. 913–925.

- [32] H.-J. Briegel, D. E. Browne, D. W. R. Raussendorf, and M. Van den Nest, *Nat. Phys.* **5**, 19 (2009).
- [33] M. Kliesch, *Lecture Notes: Characterization, Certification, and Validation of Quantum Systems* (Heinrich-Heine-Universität Düsseldorf, Düsseldorf, 2020).
- [34] Y. Takeuchi and T. Morimae, *Phys. Rev. X* **8**, 021060 (2018).
- [35] T. Heinosaari, M. A. Jivulescu, and I. Nechita, *J. Math. Phys.* **61**, 042202 (2020).
- [36] J. Řeháček, Z. Hradil, and M. Ježek, *Phys. Rev. A* **63**, 040303(R) (2001).
- [37] R. Blume-Kohout, *Phys. Rev. Lett.* **105**, 200504 (2010).
- [38] T. J. DiCiccio and B. Efron, *Stat. Sci.* **11**, 189 (1996).
- [39] S. A. Murphy and A. W. Van der Vaart, *J. Am. Stat. Assoc.* **95**, 449 (2000).
- [40] P. Faist and R. Renner, *Phys. Rev. Lett.* **117**, 010404 (2016).
- [41] D. A. Sprott, *The Likelihood Function* (Springer, New York, 2000), pp. 7–39.
- [42] S. S. Wilks, *Ann. Math. Stat.* **9**, 60 (1938).
- [43] T. L. Scholten and R. Blume-Kohout, *New J. Phys.* **20**, 023050 (2018).
- [44] S. Glancy, E. Knill, and M. Girard, *New J. Phys.* **14**, 095017 (2012).
- [45] H. Zhu and M. Hayashi, *Phys. Rev. Appl.* **12**, 054047 (2019).
- [46] K. Wang and M. Hayashi, *Phys. Rev. A* **100**, 032315 (2019).
- [47] X. Jiang, K. Wang, K. Qian, Z. Chen, Z. Chen, L. Lu, L. Xia, F. Song, S. Zhu, and X. Ma, *npj Quantum Inf.* **6**, 90 (2020).
- [48] H. Zhu and M. Hayashi, *Phys. Rev. A* **100**, 062335 (2019).
- [49] R. Koenig and J. A. Smolin, *J. Math. Phys.* **55**, 122202 (2014).
- [50] M.-D. Choi, *Linear Algebra Appl.* **10**, 285 (1975).
- [51] A. Jamiołkowski, *Rep. Math. Phys.* **3**, 275 (1972).
- [52] M. Jiang, S. Luo, and S. Fu, *Phys. Rev. A* **87**, 022310 (2013).
- [53] Y. Nesterov, *Ekonomika i Mateaticheskie Metody* **24**, 509 (1988).
- [54] J. Shang, Z. Zhang, and H. K. Ng, *Phys. Rev. A* **95**, 062336 (2017).
- [55] A. Seshadri, M. Ringbauer, J. Spainhour, R. Blatt, T. Monz, and S. Becker, minimax-fidelity-estimation, version v2.0 (2024), <https://github.com/akshayseshadri/minimax-fidelity-estimation>.
- [56] C. A. Fuchs and J. Van De Graaf, *IEEE Trans. Inf. Theory* **45**, 1216 (1999).
- [57] P. Tseng, *Math. Program.* **125**, 263 (2010).
- [58] M. M. Wilde, *Quantum Information Theory*, 2nd ed. (Cambridge University Press, Cambridge, England, 2017).
- [59] S. Boyd and L. Vandenberghe, *Convex Optimization* (Cambridge University Press, Cambridge, England, 2004).