

Fast state tomography with optimal error bounds

M. Guta,¹ J. Kahn,² R. Kueng,^{3,*} and J. A. Tropp³

¹*School of Mathematical Sciences, University of Nottingham, United Kingdom*

²*Institut de Mathématiques de Toulouse, Toulouse, France*

³*California Institute of Technology, Pasadena, United States*

(Dated: October 12, 2018)

Projected least squares (PLS) is an intuitive and numerically cheap technique for quantum state tomography. The method first computes the least-squares estimator (or a linear inversion estimator) and then projects the initial estimate onto the space of states. The main result of this paper equips this point estimator with a rigorous, non-asymptotic confidence region expressed in terms of the trace distance. The analysis holds for a variety of measurements, including 2-designs and Pauli measurements. The *sample complexity* of the estimator is comparable to the strongest convergence guarantees available in the literature and—in the case of measuring the uniform POVM—saturates fundamental lower bounds. The results are derived by reinterpreting the least-squares estimator as a sum of random matrices and applying a matrix-valued concentration inequality. The theory is supported by numerical simulations for Pauli observables, Pauli basis measurements and mutually unbiased bases.

I. INTRODUCTION

Quantum state tomography is the task of reconstructing a quantum state from experimental data. Many methods have been proposed for this problem. Maximum-likelihood estimation [1, 2] is a popular universal approach that produces point estimators, but error bars are only available in asymptotic scenarios involving fully mixed states [56]. This shortcoming spurred the development of alternatives, such as Bayesian [3–5] and region estimators [6–8]. These methods have other drawbacks, such as comparatively high computational cost and weak (or implicit) convergence guarantees. In parallel, researchers proposed compressed sensing techniques [9–12] to estimate (approximately) low-rank states from fewer samples.

In this work, we revisit linear inversion, one of the oldest and simplest methods for tomography [13]. We prove that a variant, *projected least squares (PLS)*, supports easy-to-interpret, non-asymptotic error guarantees. Our main result demonstrates that roughly $r^2 d \epsilon^{-2} \log d$ independent samples suffice to reconstruct *any* rank- r state ρ of level d up to accuracy ϵ in trace distance. This *sampling rate* improves upon existing results [14] and is competitive with the most powerful techniques in the literature [11, 15]. It (almost) saturates fundamental lower bounds on the minimal number of independent samples [57] required for tomography [11, 16]. As an added benefit, the PLS method is numerically “cheap” in the sense that its computational cost is dominated by forming the least-squares estimator. Numerical simulations indicate that PLS is much faster and performs well in comparison to prominent alternatives, including maximum-likelihood estimation [17] and compressed sensing [9].

A. Background and estimator

For d -level systems, a tomographically complete measurement [58] is described by $d \times d$ hermitian matrices $M_1, \dots, M_m \in \mathbb{H}_d$ that are positive semidefinite and obey $\sum_{i=1}^m M_i = \mathbb{I}$. Measuring a quantum state ρ results in one of m outcomes, indexed by $i \in [m]$. The probability of observing outcome i depends on ρ and is described by Born’s rule:

$$[p]_i = \Pr [i|\rho] = \text{tr}(M_i \rho) \text{ for } i \in [m]. \quad (1)$$

These probabilities can be estimated by *frequencies*: Prepare n copies of the state, measure each of them separately, and set

$$[f_n]_i = \frac{n_i}{n} \text{ for } i \in [m], \quad (2)$$

where n_i is the number of times outcome i was observed. These frequencies converge to the true probabilities as $n \rightarrow \infty$. The *least-squares estimator* is the solution to the least-squares problem that results from replacing the true probabilities in Born’s rule by frequencies (2):

$$\hat{L}_n = \underset{X \in \mathbb{H}_d}{\text{argmin}} \sum_{i=1}^m ([f_n]_i - \text{tr}(M_i X))^2. \quad (3)$$

This optimization inverts the m linear equations specified by Born’s rule; see Eq. (7) below (*linear inversion*). In general, \hat{L}_n can have negative eigenvalues, so it may *fail* to be a quantum state. Several ways to overcome this drawback have been proposed [18–22]. We follow [14] and compute the quantum state closest to \hat{L}_n in Frobenius distance:

$$\hat{\rho}_n = \underset{\sigma \text{ is a quantum state}}{\text{argmin}} \|\hat{L}_n - \sigma\|_2. \quad (4)$$

Mathematically, this corresponds to *projecting* \hat{L}_n onto the convex set of all quantum states. We term this procedure *Projected Least Squares (PLS)*; see Tab. I. To our

*corresponding author: rkueng@caltech.edu

Projected least squares (PLS) estimator $\hat{\rho}_n$

- Estimate probabilities by frequencies (2),
- Compute the least squares estimator (3),
- Project it onto the set of quantum states (4).

TABLE I: Summary of the estimation technique.

knowledge, PLS was first formulated in [14]. Its simplicity is a key advantage for both the analysis and for the actual computation. The least-squares estimator (3) and the projection onto the set of quantum states (11) admit closed-form expressions. The total cost of the PLS estimator is dominated by forming the least-squares estimate \hat{L}_n . PLS requires considerably less storage and arithmetic than existing techniques that are based on more complicated optimization problems.

II. RESULTS

A. Error bounds and confidence regions for $\hat{\rho}_n$

We analyze several important and practically relevant measurement systems: *structured POVMs* (e.g. SIC-POVMs, MUBs and stabilizer states), *global Pauli observables*, *Pauli basis measurements*, and the *uniform/covariant POVM*. For each of these settings, the PLS estimator $\hat{\rho}_n$ provably converges to the true state ρ in trace distance $\|\cdot\|_1$.

Theorem 1 (Error bound for $\hat{\rho}_n$). *Let $\rho \in \mathbb{H}_d$ be state and fix a number of samples $n \in \mathbb{N}$. Then, for each of the aforementioned measurements, the PLS estimator (Tab. I) obeys*

$$\Pr [\|\hat{\rho}_n - \rho\|_1 \geq \epsilon] \leq de^{-\frac{ne^2}{43g(d)r^2}} \text{ for } \epsilon \in [0, 1],$$

where $r = \min\{\text{rank}(\rho), \text{rank}(\hat{\rho}_n)\}$ and $g(d)$ specifies dependence on the ambient dimension:

$$g(d) = 2d \text{ for structured POVMs; see Eq. (8) for } \hat{L}_n,$$

$$g(d) = d^2 \text{ for Pauli observables; see Eq. (9) for } \hat{L}_n,$$

$$g(d) \simeq d^{1.6} \text{ for Pauli basis measurements; see Eq. (10) for } \hat{L}_n.$$

The following immediate Corollary endows $\hat{\rho}_n$ with rigorous error-bars in trace distance.

Corollary 1 (δ -confidence region). *The trace-norm ball of size $\text{rank}(\hat{\rho}_n)\sqrt{43g(d)n^{-1}\log(d/\delta)}$ around $\hat{\rho}_n$ (intersected with state space) is a δ -confidence region for the true state ρ .*

We emphasize the following aspects of this result:

- (i) (Almost) optimal sampling rate: Theorem 1 highlights that

$$n \geq 43g(d) \frac{\text{rank}(\rho)^2}{\epsilon^2} \log\left(\frac{d}{\delta}\right) \quad (5)$$

samples suffice to ensure $\|\hat{\rho}_n - \rho\|_1 \leq \epsilon$ with probability at least $1 - \delta$. For structured POVMs and Pauli observables, this *sampling rate* is comparable to the *best* theoretical bounds for alternative tomography algorithms [9, 23]. Moreover, fundamental lower bounds in [16] and [11] indicate that this scaling is optimal up to a single $\log(d)$ -factor, so it cannot be improved substantially.

- (ii) *implicit exploitation of (approximate) low rank*: the number of samples required to achieve a good estimator scales quadratically in the rank, rather than the ambient dimension d . This behavior extends to the case where ρ , or $\hat{\rho}_n$, is well-approximated by a rank- r matrix; see Theorem 4 in the appendix. These results are comparable with guarantees for compressed sensing methods [11] that are specifically designed to exploit low-rank. Fig 2 (below) provides numerical confirmation.

Proof sketch for Theorem 1. The least-squares estimator \hat{L}_n can be viewed as a sum of n independent random matrices. To illustrate this, consider a single structured POVM measurement. Then \hat{L}_1 defined in (8) is an instance of the random matrix $X = (d+1)|v_k\rangle\langle v_k| - \mathbb{I}$, where $k \in [m]$ occurs with probability $\frac{d}{m}\langle v_k|\rho|v_k\rangle$ (Born's rule). This generalizes to $\hat{L}_n = \frac{1}{n}\sum_{i=1}^n X_i$, where the matrices X_i are statistically independent. Such sums of random matrices concentrate sharply around their expectation value $\mathbb{E}\hat{L}_n = \rho$, and matrix concentration inequalities [24] quantify this convergence:

$$\Pr [\|\hat{L}_n - \rho\|_\infty \geq \tau] \leq de^{-\frac{3n\tau^2}{8g(d)}} \quad \tau \in [0, 1]. \quad (6)$$

This bound induces a similar operator norm bound for the PLS estimator: the projection is a contraction in operator norm and the shift in eigenvalues is bounded by τ . The resulting bound $\|\hat{\rho}_n - \rho\|_\infty \leq 2\tau$ may be transformed into a trace-norm bound. This comparison only depends on the (approximate) rank of the states involved. We refer to the appendix for details. \square

B. Optimal performance guarantee for the uniform POVM

Theorem 1 involves a factor of the dimension d that may be extraneous. This, in turn, introduces an additional $\log(d)$ -gap between Eq. (5) and existing lower bounds [16]. The dimensional factors emerge because we employ matrix-valued concentration inequalities in the proof. Our second main result shows that we can remove the dimensional factor for the *uniform POVM*, which encompasses all rank-one projectors:

Theorem 2 (Convergence of $\hat{\rho}_n$ for the uniform POVM). *For uniform POVM measurements, the PLS estimator obeys*

$$\Pr [\|\hat{\rho}_n - \rho\|_1 \geq \epsilon] \leq e^{2.2d - \frac{\epsilon^2 n}{480r^2}} \text{ for } \epsilon > 0,$$

where $r = \min \{\text{rank}(\rho), \text{rank}(\hat{\rho}_n)\}$.

This result exactly reproduces the best existing performance guarantees for tomography from independent measurements [15]. The bound follows from standard techniques from high-dimensional probability theory.

Proof sketch of Theorem 2. The operator norm has a variational formulation: $\|\hat{L}_n - \rho\|_\infty = \max_{z \in \mathbb{S}^d} |\langle z | \hat{L}_n - \rho | z \rangle|$. The optimization over the unit sphere may be replaced by a maximization over a finite point set, called a *covering net*, whose cardinality scales exponentially in d . For any $z \in \mathbb{S}^d$, $\langle z | \hat{L}_n - \rho | z \rangle$ is a sum of n i.i.d. random variables that exhibit subexponential tail decay. (Measuring the uniform POVM allows us to draw this conclusion.) Standard concentration inequalities yield a tail bound that decays exponentially in the number n of samples. Applying a union bound over all points z_i in the net then ensures $\Pr [\|\hat{L}_n - \rho\|_\infty \geq \tau] \leq 2e^{c_1 d - c_2 n \tau^2}$. Subsequently, closeness in operator norm for \hat{L}_n may be converted into closeness in trace-norm for $\hat{\rho}_n$ at the cost of an additional (effective) rank factor. \square

III. ALGORITHMIC CONSIDERATIONS

A. Explicit solutions for the least squares estimator (3)

Tomographically complete measurements can be viewed as injective linear maps $\mathcal{M} : \mathbb{H}_d \rightarrow \mathbb{R}^m$ with components $[\mathcal{M}(X)]_i = \text{tr}(M_i X)$ specified by Born's rule (1). It is well known that the least-squares problem (3) admits the closed-form solution:

$$\hat{L}_n = (\mathcal{M}^\dagger \mathcal{M})^{-1} (\mathcal{M}^\dagger (f_n)). \quad (7)$$

We evaluate this formula for different measurements and content ourselves with sketching key steps and results (see appendix for details).

a. Structured POVMs and the uniform POVM: Also known as 2-designs, these systems include highly structured, rank-one POVMs $\left\{ \frac{d}{m} |v_i\rangle\langle v_i| \right\}_{i=1}^m$, such as symmetric informationally complete POVMs [25], maximal sets of mutually unbiased bases [26], the set of all stabilizer states [27, 28], as well as the uniform POVM. By definition, for $X \in \mathbb{H}_d$, all of the above systems obey

$$\mathcal{M}^\dagger \mathcal{M}(X) = \frac{d^2}{m} \sum_{i=1}^m \langle v_i | X | v_i \rangle |v_i\rangle\langle v_i| = \frac{md}{d+1} (X + \text{tr}(X)\mathbb{I}).$$

These equations can readily be inverted, and Eq. (7) simplifies to

$$\hat{L}_n = (d+1) \sum_{i=1}^m [f_n]_i |v_i\rangle\langle v_i| - \mathbb{I}. \quad (8)$$

b. Pauli observables: Fix $d = 2^k$ (k qubits), and let $W_1, \dots, W_{d^2} \in \mathbb{H}_d$ be the set of Pauli observables, comprising all possible k -fold tensor products of the elementary 2×2 Pauli matrices. We can approximate the expectation value $\text{tr}(W_i \rho)$ of each Pauli observable by the empirical mean $\hat{\mu}_i = \left[f_{n/d^2}^+ \right]_i - \left[f_{n/d^2}^- \right]_i$ of the 2-outcome POVM $P_i^\pm = \frac{1}{2} (\mathbb{I} \pm W_i)$. Pauli matrices form a unitary operator basis, and the evaluation of Eq. (7) is simple:

$$\hat{L}_n = \frac{1}{d} \sum_{i=1}^{d^2} \hat{\mu}_i W_i = \frac{1}{d} \sum_{i=1}^{d^2} \left(\left[f_{n/d^2}^+ \right]_i - \left[f_{n/d^2}^- \right]_i \right) W_i. \quad (9)$$

c. Pauli basis measurements Rather than approximating (global) expectation values, it is possible to perform different combinations of local Pauli measurements. For $d = 2^k$, there are 3^k potential combinations in total. Each of the settings $s \in \{x, y, z\}^k$ corresponds to a basis measurement $|b_o^{(s)}\rangle\langle b_o^{(s)}|$, where $o \in \{\pm 1\}^k$ labels the 2^k potential outcomes. The union \mathcal{M} of all 3^k bases obeys $(\mathcal{M}^\dagger \mathcal{M})(X) = 3^k \mathcal{D}_{1/3}^{\otimes k}(X)$, where $\mathcal{D}_{1/3}(X) = (\rho + \text{tr}(X)\mathbb{I})/3$ denotes a single-qubit depolarizing channel. Evaluating Eq. (7) yields

$$\hat{L}_n = \frac{1}{3^k} \sum_{s,o} [f_{n/3^k}]_o^{(s)} \left(\mathcal{D}_{1/3}^{\otimes k} \right)^{-1} \left(|b_o^{(s)}\rangle\langle b_o^{(s)}| \right). \quad (10)$$

Finally, we point out that all of these explicit solutions are guaranteed to have unit trace: $\text{tr}(\hat{L}_n) = 1$. They are generally not positive semidefinite,

B. Explicit solutions for the projection step (4)

The PLS estimator is defined to be the state closest in Frobenius norm to the least-squares estimator \hat{L}_n . The search (4) admits a simple, *analytic* solution [19]. Define the all-ones vector $\mathbf{1} \in \mathbb{R}^d$ and the thresholding function $[\cdot]^+$ with components $[y]_i^+ = \max\{[y]_i, 0\}$. Let $\hat{L}_n = U \text{diag}(\lambda) U^\dagger$ be an eigenvalue decomposition. Then

$$\hat{\rho}_n = U \text{diag} \left([\lambda - x_0 \mathbf{1}]^+ \right) U^\dagger, \quad (11)$$

where $x_0 \in \mathbb{R}$ is chosen so that $\text{tr}(\hat{\rho}_n) = 1$. The fact that \hat{L}_n itself has unit trace ensures that this solution to Eq. (4) is unique. The number x_0 may be determined by applying a root-finding algorithm to the non-increasing function $f(x) = 2 + \text{tr}(\hat{L}_n) - dx + \sum_{i=1}^d |\lambda_i - x|$.

C. Runtime analysis

The two steps discussed here are inherently *scalable*: just count frequencies to determine the LI estimators (8,9,10) at a total cost of (at most) $\min\{m, n\}$ matrix additions. The subsequent projection onto state-space




FIG. 1: PLS (blue) vs. ML (red) for 4-qubit Pauli basis measurements: boxplots of trace distance error for ML vs. PLS for 100 datasets generated with random states of rank 1,5,10 and 16, and 200 repetitions per setting. *Inset*: trace distance error as a function of sample size for a pure target state.

is a particular type of soft-thresholding. The associated computational cost is dominated by the eigenvalue decomposition and has runtime (at most) $\mathcal{O}(d^3)$.

In summary, forming \hat{L}_n is the dominant cost of a naïve implementation. However, the high degree of structure may allow us to employ techniques from randomized linear algebra [29] to further reduce the cost.

IV. NUMERICAL EXPERIMENTS

We numerically compare the performance of PLS to *maximum likelihood* (ML) and *compressed sensing* (CS), respectively. Additional numerical studies for MUBs can be found in the appendix.

Fig. 1 compares ML and PLS for Pauli basis measurements in dimension $d = 2^4$. The trace-norm error incurred by PLS is within a factor of two of ML for low-rank states. Additional simulations (not included here) indicate that this gap closes for full-rank states.

CS is a natural benchmark for low-rank tomography. The papers [9, 10] apply to Pauli observables, and they show that a random choice of $m \geq Crd \log^6(d)$ Pauli observables is sufficient to reconstruct *any* rank- r state. The actual reconstruction is performed by solving a convex optimization problem, e.g. the least-squares fit over the set of quantum states [30, 31]. Numerical studies from [11] suggest that $m = 256$ is appropriate for $d = 2^5$ and $r = 1$. Fig. 2 shows that PLS consistently outperforms the CS estimator in this regime. Importantly, PLS was also much faster to evaluate than both, ML and CS.




FIG. 2: PLS (blue) vs. CS (red) for m 5-qubit Pauli observables and a pure target state: Trace distance error for CS ($m = 256$) and PLS ($m = 1024$) as a function of (total) sample size.

V. CONCLUSION AND OUTLOOK

Linear inversion is one of the oldest and simplest approaches to solve the practically important task of quantum state tomography. In this work, we focused on a variant called *projected least squares* (PLS) that projects the least-squares estimator onto the set of all quantum states. Not only is this estimator numerically cheap, but it comes with strong, non-asymptotic convergence guarantees. These results are derived using concentration inequalities for sums of random matrices, and they exploit the randomness inherent in quantum experiments.

We show that PLS is competitive, both in theory and in practice. For a variety of measurements, the results match the *best* existing theoretical results for the sampling rate of other tomography methods. In particular, for the uniform POVM, an order of $\frac{r^2 d}{\epsilon^2}$ samples suffice to reconstruct any rank- r state up to accuracy ϵ in trace distance. This result also saturates existing lower bounds [16] on the minimal sampling rate required for *any* tomographic procedure with independent measurements. Numerical studies underline these competitive features.

Outlook: Corollary 1 is not (yet) optimal. *Bootstrapping* could be used to obtain tighter confidence regions, and the low computational cost of PLS may speed up this process considerably. It also seems fruitful to combine the ideas presented here with recent insights from [32]. Finally, the proof of Theorem 1 indicates that PLS is stable with respect to time-dependent state generation (*drift*). We intend to address these points in future work.

Acknowledgements: The authors thank Philippe Faist, Matthias Kleinmann, Anirudh Acharya and Theodore Kypraios for fruitful discussions. RK and JT are supported by ONR Award No. N00014-17-12146. RK also acknowledges funding provided by the Institute of Quantum Information and Matter, an NSF Physics Frontiers Center (NSF Grant PHY-1733907).

- [1] Z. Hradil, *Phys. Rev. A* **55**, R1561 (1997).
- [2] J. Řeháček, Z. Hradil, and M. Ježek, *Phys. Rev. A* **63**, 040303 (2001).
- [3] R. Blume-Kohout, *New J. Phys.* **12**, 043034 (2010).
- [4] C. Granade, J. Combes, and D. G. Corrie, *New J. Phys* **18**, 033024 (2016).
- [5] C. Granade, C. Ferrie, and S. T. Flammia, *New J. Phys.* **19**, 113017 (2017).
- [6] M. Christandl and R. Renner, *Phys. Rev. Lett.* **109**, 120403 (2012).
- [7] R. Blume-Kohout, *arXiv preprint arXiv:1202.5270* (2012).
- [8] P. Faist and R. Renner, *Phys. Rev. Lett.* **117**, 010404 (2016).
- [9] D. Gross, Y.-K. Liu, S. T. Flammia, S. Becker, and J. Eisert, *Phys. Rev. Lett.* **105**, 150401 (2010).
- [10] Y.-K. Liu, in *Adv. Neural Inf. Process. Syst.* **24**, edited by J. Shawe-Taylor, R. S. Zemel, P. L. Bartlett, F. Pereira, and K. Q. Weinberger (Curran Associates, Inc., 2011) pp. 1638–1646.
- [11] S. T. Flammia, D. Gross, Y.-K. Liu, and J. Eisert, *New J. Phys.* **14**, 095022 (2012).
- [12] C. Riofrio, D. Gross, S. Flammia, T. Monz, D. Nigg, R. Blatt, and J. Eisert, *Nat. Commun.* **8**, 15305 (2017).
- [13] G. M. D’Ariano and P. Perinotti, *Phys. Rev. Lett.* **98**, 020403 (2007).
- [14] T. Sugiyama, P. S. Turner, and M. Muraio, *Phys. Rev. Lett.* **111**, 160406 (2013).
- [15] R. Kueng, H. Rauhut, and U. Terstiege, *Appl. Comput. Harmon. Anal.* **42**, 88 (2017).
- [16] J. Haah, A. W. Harrow, Z. Ji, X. Wu, and N. Yu, *IEEE Trans. Inform. Theory* **63**, 5628 (2017).
- [17] H. Häffner, W. Hänsel, C. F. Roos, J. Benhelm, D. Chek-al kar, M. Chwalla, T. Körber, U. D. Rapol, M. Riebe, P. O. Schmidt, C. Becher, O. Gühne, W. Dür, and R. Blatt, *Nature* **438**, 643 EP (2005).
- [18] M. S. Kaznady and D. F. V. James, *Phys. Rev. A* **79**, 022109 (2009).
- [19] J. A. Smolin, J. M. Gambetta, and G. Smith, *Phys. Rev. Lett.* **108**, 070502 (2012).
- [20] P. Alquier, C. Butucea, M. Hebiri, K. Meziani, and T. Morimae, *Phys. Rev. A* **88**, 032113 (2013).
- [21] C. Butucea, M. Guta, and T. Kypraios, *New Journal of Physics* **17**, 113050 (2015).
- [22] V. Koltchinskii and D. Xia, *J. Mach. Learn. Res.* **16**, 1757 (2015).
- [23] R. Kueng, H. Zhu, and D. Gross, *arXiv preprint arXiv:1610.08070* (2016).
- [24] J. A. Tropp, *Found. Comput. Math.* **12**, 389 (2012).
- [25] J. M. Renes, R. Blume-Kohout, A. J. Scott, and C. M. Caves, *J. Math. Phys.* **45**, 2171 (2004).
- [26] A. Klappenecker and M. Rötteler, in *International Symposium on Information Theory, 2005. ISIT 2005. Proceedings* (2005) pp. 1740–1744.
- [27] C. Dankert, R. Cleve, J. Emerson, and E. Livine, *Phys. Rev. A* **80**, 012304 (2009).
- [28] D. Gross, K. Audenaert, and J. Eisert, *J. Math. Phys.* **48**, 052104 (2007).
- [29] N. Halko, P. Martinsson, and J. Tropp, *SIAM Review* **53**, 217 (2011), <https://doi.org/10.1137/090771806>.
- [30] A. Kaley, R. L. Kosut, and I. H. Deutsch, *NPJ Quantum Inf.* **1**, 15018 (2015).
- [31] M. Kabanava, R. Kueng, H. Rauhut, and U. Terstiege, *Inf. Inference* **5**, 405 (2016).
- [32] T. L. Scholten and R. Blume-Kohout, *New J. Phys.* **20**, 023050 (2018).
- [33] R. O’Donnell and J. Wright, in *Proceedings of the Forty-eighth Annual ACM Symposium on Theory of Computing, STOC ’16* (ACM, New York, NY, USA, 2016) pp. 899–912.
- [34] A. J. Scott, *J. Phys. A* **39**, 13507 (2006).
- [35] D. Gross, F. Kraemer, and R. Kueng, *J. Fourier Anal. Appl.* **21**, 229 (2015).
- [36] J. Schwinger, *Proc. Natl. Acad. Sci. USA* **46**, 570 (1960).
- [37] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information: 10th Anniversary Edition*, 10th ed. (Cambridge University Press, New York, NY, USA, 2011).
- [38] R. Kueng and D. Gross, *preprint arXiv:1510.02767* (2015).
- [39] H. Zhu, *Phys. Rev. A* **96**, 062336 (2017).
- [40] Z. Webb, *arXiv preprint arXiv:1510.02769* (2015).
- [41] N. Tomczak-Jaegermann, *Stud. Math.* **50**, 163 (1974).
- [42] F. Lust-Piquard, *CR Acad. Sci. Paris* **303**, 289 (1986).
- [43] G. Pisier and Q. Xu, *Commun. Math. Phys.* **189**, 667 (1997).
- [44] M. Rudelson, *J. Funct. Anal.* **164**, 60 (1999).
- [45] R. Ahlswede and A. Winter, *IEEE Trans. Inform. Theory* **48**, 569 (2002).
- [46] D. Gross, *IEEE Trans. Inform. Theory* **57**, 1548 (2011).
- [47] R. I. Oliveira *et al.*, *Electron. Commun. Probab* **15**, 26 (2010).
- [48] J. A. Tropp, *Applied and Computational Harmonic Analysis* **44**, 700 (2018).
- [49] S. Foucart and H. Rauhut, *A Mathematical Introduction to Compressive Sensing* (Birkhäuser Basel, 2013).
- [50] C. Lancien and A. Winter, *preprint arXiv:1711.00697* (2017).
- [51] M. Talagrand, *The generic chaining: upper and lower bounds of stochastic processes* (Springer Science & Business Media, 2006).
- [52] R. Vershynin, in *Compressed Sensing: Theory and Practice*, edited by Y. C. Eldar and G. Kutyniok (Cambridge University Press, 2012) pp. 210–268.
- [53] W. K. Wootters and B. D. Fields, *Ann. Phys.* **191**, 363 (1989).
- [54] Bandyopadhyay, Boykin, Roychowdhury, and Vatan, *Algorithmica* **34**, 512 (2002).
- [55] A. Klappenecker and M. Rötteler, in *Finite Fields and Applications*, edited by G. L. Mullen, A. Poli, and H. Stichtenoth (Springer, Berlin, Heidelberg, 2004) pp. 137–144.
- [56] Experimental works [17] do provide bootstrap error bars, but these methods are not theoretically grounded for rank-deficient states. (General statistical theory provides asymptotic error bars and bootstrap theory, but only for fully mixed states where local asymptotic normality holds.)
- [57] Correlated measurements over several copies of the state can achieve even lower sampling rates [16, 33], but are much more challenging to implement. Here we focus on the practically more relevant case of measuring each copy of ρ independently.
- [58] A measurement is tomographically complete if, for every pair of distinct states $\rho \neq \sigma$, there exists $i \in [m]$ such that $\text{tr}(M_i \rho) \neq \text{tr}(M_i \sigma)$.

Appendix

At the heart of this work is *projected least squares* (PLS) – a simple point estimator for quantum state tomography from tomographically complete measurements $\{M_1, \dots, M_m\} \subset \mathbb{H}_d$. PLS is a three-step procedure, see also Tab. I:

1. Estimate outcome probabilities by frequencies.
2. Construct the least squares (linear inversion) estimator:

$$\hat{L}_n = \operatorname{argmin}_{X \in \mathbb{H}_d} \sum_{i=1}^m (f_i - \operatorname{tr}(M_i X))^2. \quad (12)$$

3. Project onto the set of all quantum states:

$$\hat{\rho}_n = \operatorname{argmin}_{\sigma \text{ is a quantum state}} \|\hat{L}_n - \sigma\|_2. \quad (13)$$

We analyze the performance of PLS for a variety of concrete measurement scenarios: *structured POVMs*, *Pauli observables*, *Pauli basis measurements* and the *uniform POVM*. For each of them, $\hat{\rho}_n$ may be equipped with rigorous non-asymptotic confidence regions in trace distance. In this appendix, we complement the rather succinct presentation in the main text with additional explanations, motivations and more detailed arguments.

Outline: In Section VI we provide explicit least squares solutions (12) for the different measurements. We also review essential features and properties of the individual scenarios to provide context.

Section VII contains the main conceptual insight of this work: least squares estimators may be interpreted as sums of independent random matrices – the randomness is due to the fundamental laws of quantum mechanics (Born’s rule). This allows us to apply strong matrix-valued concentration inequalities to show that, with high probability, \hat{L}_n is close to the true target state in operator norm.

Section VIII is devoted to showing that closeness of \hat{L}_n in operator norm implies closeness of $\hat{\rho}_n$ in trace norm.

We combine these two insights in Section IX to arrive at the main result of this work: convergence guarantees for the PLS estimator in trace norm. The result derived there is a strict generalization of Theorem 1 quoted in the main text. It extends to the notion of *effective rank* which may be beneficial in concrete applications. We illustrate this potential benefit with a caricature of a faulty state preparation apparatus.

In Section X yet stronger convergence guarantees for the uniform POVM are derived. The proof technique is completely different and we believe that it may be of independent interest to the community.

Finally, we present additional numerical experiments in Section XI.

VI. CLOSED-FORM EXPRESSIONS FOR LEAST SQUARES ESTIMATORS

As outlined in the main text, any POVM measurement can be viewed as a linear map $\mathcal{M} : \mathbb{H}_d \rightarrow \mathbb{R}^m$, defined component-wise as $[\mathcal{M}(X)]_i = \operatorname{tr}(M_i X)$ for $i \in [m]$. This map is injective if and only if the measurement is tomographically complete. Provided that this is the case, the least squares estimator (12) admits a unique solution:

$$\hat{L}_n = \left(\mathcal{M}^\dagger \mathcal{M} \right)^{-1} \mathcal{M}^\dagger (f_n),$$

where $f_n \in \mathbb{R}^m$ subsumes the individual frequency estimates. In this section, we evaluate this formula explicitly for different types of prominent measurements.

A. The uniform POVM and 2-designs

The *uniform/covariant* POVM in d -dimensions corresponds to the union of all (properly re-normalized) rank-one projectors: $\{d|v\rangle\langle v|d\}_{v \in \mathbb{S}^d}$. Here, $d\nu$ denotes the unique, unitarily invariant, measure on the complex unit sphere induced by the Haar measure (over the unitary group $U(d)$). Its high degree of symmetry allows for analyzing this POVM by means of powerful tools from representation theory. This is widely known, see e.g. [34, 35], but we include a short presentation here to be self-contained. Define the *frame operator* of order k : $F_{(k)} = \int_{\mathbb{S}^d} (|v\rangle\langle v|)^{\otimes k} d\nu \in \mathbb{H}_d^{\otimes k}$. Unitary invariance of $d\nu$ implies that this frame operator commutes with every k -fold tensor product of a unitary matrix $U \in U(d)$:

$$\begin{aligned} U^{\otimes k} F_{(k)} &= \int_{\mathbb{S}^d} (U|v\rangle\langle v|)^{\otimes k} d\nu = \int_{\mathbb{S}^d} (|\tilde{v}\rangle\langle \tilde{v}|U)^{\otimes k} d\tilde{\nu} \\ &= F_{(k)} U^{\otimes k}. \end{aligned}$$

Here, we have used a change of variables ($\tilde{v} = Uv$) together with the fact that $d\nu$ is unitarily invariant ($d\tilde{\nu} = d\nu$). Schur’s Lemma – one of the most fundamental tools in representation theory – states that any matrix that commutes with every element of a given group representation must be proportional to a sum of the projectors onto the associated irreducible representations (irreps). For the task at hand, the representation of interest is the diagonal representation of the unitary group: $U \mapsto U^{\otimes k}$ for all $U \in U(d)$. This representation affords, in general, many irreps that may be characterized using Schur–Weyl duality. The *symmetric subspace* $\operatorname{Sym}^{(k)} \subset \left(\mathbb{C}^d\right)^{\otimes k}$ is one of them and corresponds to the subspace of all vectors that are invariant under permuting tensor factors. Crucially, $F_{(k)}$ is an average over rank-one projectors onto vectors $|v\rangle^{\otimes k} \in \operatorname{Sym}^{(k)}$ and, therefore, its range must be contained entirely within $\operatorname{Sym}^{(k)}$. Combining

this with the assertion of Schur's lemma then yields

$$F_{(k)} = \int_{\mathbb{S}^d} (|v\rangle\langle v|)^{\otimes k} dv = \binom{d+k-1}{k}^{-1} P_{\text{Sym}^{(k)}} \quad k \in \mathbb{N}, \quad (14)$$

The pre-factor $\binom{d+k-1}{k}^{-1} = \dim(\text{Sym}^{(k)})^{-1}$ follows from the fact that $F_{(k)}$ has unit trace.

This closed-form expression is very useful. In particular, it implies that the uniform POVM $\{|v\rangle\langle v|\}_{v \in \mathbb{S}^d}$ is almost an isometry. Fix $X \in \mathbb{H}_d$ and compute

$$\begin{aligned} & (d+1) \int_{\mathbb{S}^d} d\langle v|X|v\rangle |v\rangle\langle v| dv \\ &= (d+1) d\text{tr}_1 \left(X \otimes \mathbb{I} \int_{\mathbb{S}^d} (|v\rangle\langle v|)^{\otimes 2} dv \right) \\ &= 2\text{tr}_1 \left(X \otimes \mathbb{I} P_{\text{Sym}^{(2)}} \right), \end{aligned} \quad (15)$$

where $\text{tr}_1(A \otimes B) = \text{tr}(A)B$ denotes the partial trace over the first tensor factor. The projector onto the totally symmetric subspace of two parties has an explicit representation: $P_{\text{Sym}^{(2)}} = \frac{1}{2}(\mathbb{I} + \mathbb{F})$, where \mathbb{F} denotes the *flip operator*, i.e. $\mathbb{F}|x\rangle \otimes |y\rangle = |y\rangle \otimes |x\rangle$ for all $|x\rangle, |y\rangle \in \mathbb{C}^d$ and extend it linearly to the entire tensor product. Inserting this explicit characterization into Eq. (33) yields

$$\begin{aligned} (d+1) \int_{\mathbb{S}^d} d\langle v|X|v\rangle |v\rangle\langle v| dv &= \text{tr}_1(X \otimes \mathbb{I}(\mathbb{I} + \mathbb{F})) \\ &= X + \text{tr}(X)\mathbb{I}. \end{aligned} \quad (16)$$

We emphasize that the full symmetry of the uniform POVM is not required to derive this formula: Eq. (14) for $k=2$ is sufficient. This motivates the following definition:

Definition 1 (2-design). *A (finite) set of m rank-one projectors $\{|v_i\rangle\langle v_i|\}_{i=1}^m$ is called a (complex-projective) 2-design if*

$$\frac{1}{m} \sum_{i=1}^m (|v_i\rangle\langle v_i|)^{\otimes 2} = \binom{d+1}{2}^{-1} P_{\text{Sym}^{(2)}}.$$

Taking the partial trace of this expression yields

$$\frac{1}{m} \sum_{i=1}^m |v_i\rangle\langle v_i| = \frac{1}{d}\mathbb{I},$$

highlighting that each 2-design is proportional to a POVM $\mathcal{M} = \left\{ \frac{d}{m} |v_i\rangle\langle v_i| \right\}_{i=1}^m$. Moreover, viewed as a map $\mathcal{M} : \mathbb{H}_d \rightarrow \mathbb{R}^m$, every such POVM obeys

$$\mathcal{M}^\dagger \mathcal{M}(X) = \frac{d^2}{m^2} \sum_{i=1}^m \langle v_i|X|v_i\rangle |v_i\rangle\langle v_i| = \frac{d(X + \text{tr}(X)\mathbb{I})}{(d+1)m}$$

for any $X \in \mathbb{H}_d$, which can be readily inverted:

$$\left(\mathcal{M}^\dagger \mathcal{M} \right)^{-1}(X) = \frac{m}{d} \left((d+1)X - \text{tr}(X)\mathbb{I} \right). \quad (17)$$

Inserting this formula into the closed-form expression of the linear-inversion estimator yields

$$\begin{aligned} \hat{L}_n &= \left(\mathcal{M}^\dagger \mathcal{M} \right)^{-1} \left(\mathcal{M}^\dagger(f) \right) \\ &= \left(\mathcal{M}^\dagger \mathcal{M} \right)^{-1} \left(\frac{d}{m} \sum_{i=1}^m f_i |v_i\rangle\langle v_i| \right) \\ &= \sum_{i=1}^m f_i \left((d+1) |v_i\rangle\langle v_i| - \text{tr}(|v_i\rangle\langle v_i|)\mathbb{I} \right) \\ &= (d+1) \sum_{i=1}^m f_i |v_i\rangle\langle v_i| - \mathbb{I} \end{aligned} \quad (18)$$

for any frequency vector $f_n \in \mathbb{R}^m$. Mathematically, this is a consequence of the fact that 2-design POVMs "almost" form a tight frame on \mathbb{H}_d . The close connection to well-behaved, tomographically complete, rank-one POVMs has spurred considerable interest in the identification of 2-designs. Over the past decades, the following concrete examples have been identified:

- (i) *Equiangular lines (SIC POVMs)*: a family of m unit vectors $|v_1\rangle, \dots, |v_m\rangle \in \mathbb{S}^d$ is equiangular, if $|\langle v_i, v_j \rangle|^2$ is constant for all $i \neq j$. The maximal cardinality of such a set is $m = d^2$ in which case the angle must be fixed: $|\langle v_i, v_j \rangle|^2 = \frac{1}{d+1}$. Such maximal sets of equiangular lines are known to form 2-designs [25] and have been termed *symmetric, informationally complete (SIC) POVMs*. This nomenclature underlines the importance of Eq. (16) for the original quantum motivation of the study of equiangular lines. While several explicit constructions of SIC POVMs exist, the general question of their existence remains an intriguing open problem.
- (ii) *Mutually unbiased bases (MUBs)*: Two orthonormal bases $\{|b_i\rangle\}_{i=1}^d$ and $\{|c_i\rangle\}_{i=1}^d$ of \mathbb{C}^d are *mutually unbiased* if $|\langle b_i, c_j \rangle|^2 = \frac{1}{d}$ for all $1 \leq i, j \leq d$. The study of such mutually unbiased bases (MUBs) has a rich history in quantum mechanics that dates back to Schwinger [36]. It is known that at most $(d+1)$ pairwise mutually unbiased bases can exist in dimension d and explicit algebraic constructions are known for prime power dimensions ($d = p^k$). Klappenecker and Roettler [26] showed that maximal sets of MUBs are guaranteed to form 2-designs.
- (iii) *stabilizer states (STABs)*: the stabilizer formalism is one of the cornerstones of quantum computation, fault tolerance and error correction, see e.g. [37]. Let \mathcal{P}_k be the Pauli group on k qubits ($d = 2^k$), i.e. the group generated by k -fold tensor products of the elementary Pauli matrices. It is then possible to find maximal abelian subgroups $\mathcal{S} \subset \mathcal{P}_k$

of size $d = 2^k$. Since all matrices $W \in \mathcal{S}$ commute, they can be simultaneously diagonalized and determine a single unit vector which is the joint eigenvector with eigenvalue $+1$ of all the matrices in \mathcal{S} (provided that $-\mathbb{I} \notin \mathcal{S}$). Such vectors are called *stabilizer states* (STAB) and the group $\mathcal{S} \subset \mathcal{P}_k$ is its associated *stabilizer group*. A total of $m = 2^k \prod_{i=0}^k (d^i + 1) = 2^{\frac{1}{2}k^2 + o(k)}$ different stabilizer states can be generated this way. The union of all of them is actually known to form a 3-design [38–40] and, therefore, also a 2-design. The latter is also a consequence of earlier results [27, 28]

B. Pauli observables

For $d = 2^k$, the Pauli matrices $W_1, \dots, W_{d^2} \in \mathbb{H}_d$ arise from all possible k -fold tensor products of elementary Pauli matrices $\{\mathbb{I}, \sigma_x, \sigma_y, \sigma_z\} \subset \mathbb{H}_2$. They are well-known to form a unitary operator basis:

$$X = \frac{1}{d} \sum_{i=1}^{d^2} \text{tr}(W_i X) W_i, \quad (20)$$

for all $X \in \mathbb{H}_d$. While they do constitute observables, Pauli matrices by themselves are not POVMs. However, every observable W_i may be associated with a two-outcome POVM $\mathcal{M}_i = \{P_i^\pm\} = \left\{\frac{1}{2}(\mathbb{I} \pm W_i)\right\}$. The union $\bigcup_{i=1}^{d^2} \mathcal{M}_i$ of all these 2-outcome POVMs constitutes a linear map $\mathcal{M} : \mathbb{H}_d \rightarrow \mathbb{R}^{2m}$ that obeys

$$\begin{aligned} \mathcal{M}^\dagger \mathcal{M}(X) &= \sum_{i=1}^{d^2} (\text{tr}(P_i^+ X) P_i^+ + \text{tr}(P_i^- X) P_i^-) \\ &= \sum_{i=1}^{d^2} \frac{1}{2} (\text{tr}(X) \mathbb{I} + \text{tr}(W_i X) W_i) \\ &= \frac{d}{2} (d \text{tr}(X) \mathbb{I} + X), \end{aligned}$$

where the last line is due to Eq. (20). Once more, this expression can be readily inverted:

$$(\mathcal{M}^\dagger \mathcal{M})^{-1}(X) = \frac{2}{d} X - \frac{2 \text{tr}(X)}{d^2 + 1} \mathbb{I}.$$

Before we continue, we note that one Pauli matrix is equal to the identity, say $W_1 = \mathbb{I}$, and the associated POVM is trivial. Hence, we suppose that n copies of ρ are distributed equally among all $d^2 - 1$ non-trivial 2-Outcome POVMs \mathcal{M}_i . We denote the resulting frequencies by $[f]_i^\pm$ and suppress the dependence on the number of samples. Then, the explicit solution to the least

squares problem becomes

$$\begin{aligned} \hat{L}_n &= (\mathcal{M}^\dagger \mathcal{M})^{-1} (\mathcal{M}^\dagger (f_n)) \\ &= (\mathcal{M}^\dagger \mathcal{M})^{-1} (\mathbb{I}) + \sum_{i=2}^{d^2} \sum_{o=\pm} [f]_i^o (\mathcal{M}^\dagger \mathcal{M})^{-1} (P_i^o) \\ &= \frac{2}{d(d^2 + 1)} \mathbb{I} + \sum_{i=2}^{d^2} [f]_i^+ \left(\frac{1}{d} (\mathbb{I} + W_i) - \frac{d}{d^2 + 1} \mathbb{I} \right) \\ &\quad + \sum_{i=2}^{d^2} [f]_i^- \left(\frac{1}{d} (\mathbb{I} - W_i) - \frac{d}{d^2 + 1} \mathbb{I} \right) \\ &= \frac{1}{d} \sum_{i=2}^{d^2} ([f]_i^+ - [f]_i^-) W_i + \frac{2 + \sum_{i=2}^{d^2} ([f]_i^+ + [f]_i^-)}{d(d^2 + 1)} \mathbb{I}. \end{aligned}$$

We can simplify this expression further by noticing that each 2-outcome POVM is dichotomic: either $+$ or $-$ is observed for every run. This implies $[f]_i^+ + [f]_i^- = 1$ and, by extension, $\sum_{i=2}^{d^2} ([f]_i^+ + [f]_i^-) = d^2 - 1$. Hence,

$$\begin{aligned} \hat{L}_n &= \frac{1}{d} \sum_{i=2}^{d^2} ([f]_i^+ - [f]_i^-) W_i + \frac{1}{d} \mathbb{I} \\ &= \frac{1}{d} \sum_{i=1}^{d^2} ([f]_i^+ - [f]_i^-) W_i, \end{aligned} \quad (21)$$

because $[f]_1^+ = 1$. This is the formula from the main text and has a compelling interpretation: the difference $\hat{\mu}_i = [f]_i^+ - [f]_i^-$ is an empirical estimate for the expectation value $\mu_i = \text{tr}(W_i X)$ of the i -th Pauli observable. Finally, note that this estimator is again unbiased with respect to random fluctuations in the sample statistics:

$$\mathbb{E}[\hat{L}_n] = \frac{1}{d} \sum_{i=1}^{d^2} \text{tr}(W_i \rho) W_i = \rho. \quad (22)$$

C. Pauli basis measurements

Before considering the general case, we find it instructive to consider the single qubit case in more detail. For now, fix $d = 2$ and note that there are three non-trivial Pauli matrices σ_s with $s \in \{x, y, z\}$. We may associate each σ_s with a 2-outcome POVM that is also a basis measurement: $\frac{1}{2}(\mathbb{I} \pm \sigma_s) = |b_\pm^{(s)}\rangle\langle b_\pm^{(s)}|$. For $s \neq s'$

$$\left| \langle b_\pm^{(s)}, b_{o'}^{(\pm)} \rangle \right|^2 = \frac{1}{4} \text{tr}((\mathbb{I} \pm \sigma_s)(\mathbb{I} \pm \sigma_{s'})) = \frac{1}{2},$$

because $\sigma_s, \sigma_{s'}$ and $\sigma_s \sigma_{s'}$ have vanishing trace. This implies that the six vectors $|b_o^{(s)}\rangle\langle b_o^{(s)}|$ with $o \in \{\pm 1\}$ form a maximal set of 3 = $(d + 1)$ mutually unbiased bases. Such vector sets form spherical 2-designs and Eq. (16) ensures for any $X \in \mathbb{H}_d$

$$\sum_{s,o} \langle b_o^{(s)} | X | b_o^{(s)} \rangle | b_o^{(s)} \rangle \langle b_o^{(s)} | = (X + \text{tr}(X) \mathbb{I}) = 3\mathcal{D}_{1/3}(X). \quad (23)$$

Here, $\mathcal{D}_{1/3} : \mathbb{H}_d \rightarrow \mathbb{H}_d$ denotes a single-qubit depolarizing channel with loss parameter $p = \frac{1}{3}$.

This behavior extends to multi-qubit systems, i.e. $d = 2^k$. Suppose that we perform k local (single-qubit) Pauli measurements on a k -qubit state $\rho \in \mathbb{H}_d$. Then, there are a total of k potential combinations that we label by a string $\mathbf{s} = (s_1, \dots, s_k) \in \{x, y, z\}^k$. Each of them corresponds to a POVM $\mathcal{M}^{(\mathbf{s})}$ with $2^k = d$ outcomes that we label by $\mathbf{o} = (o_1, \dots, o_k) \in \{\pm 1\}^k$. The POVM element associated with index \mathbf{s} and outcome \mathbf{o} has an appealing tensor-product structure: $|b_{\mathbf{o}}^{(\mathbf{s})}\rangle\langle b_{\mathbf{o}}^{(\mathbf{s})}| = \bigotimes_{i=1}^k |b_{o_i}^{(s_i)}\rangle\langle b_{o_i}^{(s_i)}|$. Let $\mathcal{M} = \bigcup_{\mathbf{s}} \mathcal{M}^{(\mathbf{s})} : \mathbb{H}_d \rightarrow \mathbb{R}^{3^k} \times \mathbb{R}^{2^k}$ denote the union of all such basis measurements. Then, the following formula is true for tensor product matrices $X = \bigotimes_{i=1}^k X_i$ and $X_i \in \mathbb{H}_2$:

$$\begin{aligned} \mathcal{M}^\dagger \mathcal{M}(X) &= \sum_{\mathbf{s}, \mathbf{o}} \langle b_{\mathbf{o}}^{(\mathbf{s})} | X | b_{\mathbf{o}}^{(\mathbf{s})} \rangle | b_{\mathbf{o}}^{(\mathbf{s})} \rangle \langle b_{\mathbf{o}}^{(\mathbf{s})} | \\ &= \bigotimes_{i=1}^k \left(\sum_{s_i, o_i} \langle b_{o_i}^{(s_i)} | X_i | b_{o_i}^{(s_i)} \rangle | b_{o_i}^{(s_i)} \rangle \langle b_{o_i}^{(s_i)} | \right) \\ &= 3^k \bigotimes_{i=1}^k \mathcal{D}(X_i) = 3^k \mathcal{D}_{1/3}^{\otimes k}(X), \end{aligned} \quad (24)$$

where we have used Eq.(23). Linear extension ensures that this formula remains valid for arbitrary matrices $X \in \mathbb{H}_d$. Since the single qubit depolarizing channel is invertible, the same is true for its k -fold tensor product and we conclude

$$\left(\mathcal{M}^\dagger \mathcal{M} \right)^{-1}(X) = \frac{1}{3^k} \left(\mathcal{D}_{1/3}^{\otimes k} \right)^{-1}(X).$$

Inserting this explicit expression into the closed-form expression for the least squares estimator yields

$$\begin{aligned} \hat{L}_n &= \left(\mathcal{M}^\dagger \mathcal{M} \right)^{-1} \left(\mathcal{M}^\dagger(f) \right) \\ &= \frac{1}{3^k} \sum_{\mathbf{s}, \mathbf{o}} [f]_{\mathbf{o}}^{(\mathbf{s})} \left(\mathcal{D}_{1/3}^{\otimes k} \right)^{-1} \left(| b_{\mathbf{o}}^{(\mathbf{s})} \rangle \langle b_{\mathbf{o}}^{(\mathbf{s})} | \right), \end{aligned}$$

as advertised in the main text. Here, $[f]_{\mathbf{o}}^{(\mathbf{s})}$ is assumed to be a frequency approximation to $p_{\mathbf{o}}^{(\mathbf{s})} = \langle b_{\mathbf{o}}^{(\mathbf{s})} | \rho | b_{\mathbf{o}}^{(\mathbf{s})} \rangle$.

We conclude this section with a single-qubit observations that allows for characterizing this expression in a more explicit fashion. Note that one may rewrite $\mathcal{D}_{1/3}(X)$ as $\frac{\text{tr}(X)}{2} \mathbb{I} + \frac{1}{6} \sum_s \text{tr}(\sigma_s X) \sigma_s$. This facilitates the computation of the single-qubit inverse:

$$\mathcal{D}_{1/3}^{-1}(X) = \frac{\text{tr}(X)}{2} \mathbb{I} + \frac{3}{2} \sum_s \text{tr}(\sigma_s X) \sigma_s$$

and, in particular

$$\begin{aligned} \mathcal{D}_{1/3}^{-1} \left(| b_{\pm}^{(s)} \rangle \langle b_{\pm}^{(s)} | \right) &= \frac{1}{2} \left(\mathcal{D}^{-1}(\mathbb{I}) \pm \mathcal{D}^{-1}(\sigma_s) \right) \\ &= \frac{1}{2} \left(\mathbb{I} \pm 3\sigma_s \right) = 3 | b_{\pm}^{(s)} \rangle \langle b_{\pm}^{(s)} | - \mathbb{I}. \end{aligned}$$

This in turn implies

$$\hat{L}_n = \frac{1}{3^k} \sum_{\mathbf{s}, \mathbf{o}} [f]_{\mathbf{o}}^{(\mathbf{s})} \bigotimes_{i=1}^k \left(3 | b_{o_i}^{(s_i)} \rangle \langle b_{o_i}^{(s_i)} | - \mathbb{I} \right). \quad (25)$$

which, again, is an unbiased estimator with respect to random fluctuations in the sample statistics.

Finally, we also point out another consequence that will be important later on:

$$\left(\mathcal{D}_{1/3}^{-1} \left(| b_{\mathbf{o}}^{(\mathbf{s})} \rangle \langle b_{\mathbf{o}}^{(\mathbf{s})} | \right) \right)^2 = 5 \mathcal{D}_{3/5} \left(| b_{\mathbf{o}}^{(\mathbf{s})} \rangle \langle b_{\mathbf{o}}^{(\mathbf{s})} | \right), \quad (26)$$

where $\mathcal{D}_{3/5}$ is another single-qubit depolarizing channel.

VII. THE MATRIX BERNSTEIN INEQUALITY AND CONCENTRATION IN OPERATOR NORM

Scalar concentration inequalities provide sharp bounds on the probability of a sum of independent random variables deviating from their mean. Classical examples include Hoeffding's, Chernoff's and Bernstein's inequality – all of which have found widespread use in a variety of scientific disciplines. The main results of this work are based on a matrix generalizations of these classical statements – in particular the *matrix Bernstein inequality* developed by one of the authors, see [24, Theorem 1.4].

Theorem 3 (Matrix Bernstein inequality). *Consider a sequence of n independent, hermitian random matrices $A_1, \dots, A_n \in \mathbb{H}_d$. Assume that each A_i satisfies*

$$\mathbb{E}[A_i] = 0 \quad \text{and} \quad \|A_i\|_{\infty} \leq R \text{ almost surely.}$$

Then, for any $t > 0$

$$\Pr \left[\left\| \sum_{i=1}^n (A_i - \mathbb{E}[A_i]) \right\|_{\infty} \geq t \right] \leq \begin{cases} d \exp \left(-\frac{3t^2}{8\sigma^2} \right) & t \leq \frac{\sigma^2}{R}, \\ d \exp \left(-\frac{3t}{8R} \right) & t \geq \frac{\sigma^2}{R}, \end{cases}$$

where $\sigma^2 = \left\| \sum_{i=1}^n \mathbb{E}[A_i^2] \right\|_{\infty}$.

First results of this kind originate in Banach space theory [41–44] and were later independently developed in quantum information theory [45, 46]. Further advances by Oliveira [47] and one of the authors [24] led to the result that we employ here. We refer to the monograph [24] for a detailed exposition of related work and history.

Similar to the scalar Bernstein inequality, the tail behavior in Theorem 3 consists of two regimes. Small deviations are suppressed in a subgaussian fashion, while larger deviations follow a subexponential decay. The ratio $\frac{\sigma^2}{R}$ marks the transition from one regime into the other. We also note in passing that this result recovers the scalar Bernstein inequality for $d = 1$ ($\mathbb{H}_1 \simeq \mathbb{R}$).

A. Concentration for structured POVM measurements

For structured measurements (2-designs) we may rewrite the (plain) least squares estimator (19) as

$$\hat{L}_n = (d+1) \sum_{i=1}^m [f_n]_i |v_i\rangle\langle v_i| - \mathbb{I} = \frac{1}{n} \sum_{i=1}^m X_i,$$

where each X_i is an i.i.d. copy of the random matrix $X \in \mathbb{H}_d$ that assumes $(d+1)|v_k\rangle\langle v_k| - \mathbb{I}$ with probability $\frac{d}{m} \langle v_k | \rho | v_k \rangle$ for all $k \in [m]$. Unbiasedness with respect to the sample statistics ensures $\mathbb{E}[\hat{L}_n] = \mathbb{E}[X] = \rho$. Hence, $\hat{L}_n - \rho$ is a sum of iid, centered random matrices $\frac{1}{n}(X_i - \mathbb{E}[X_i])$. These obey

$$\|X_i - \mathbb{E}[X_i]\|_\infty = \frac{1}{n} \|(d+1)|v_k\rangle\langle v_k| - \mathbb{I} - \rho\|_\infty \leq \frac{d}{n} =: R,$$

where $k \in [m]$ is arbitrary. Next, note that the random matrix X obeys

$$\mathbb{E}[(X - \mathbb{E}[X])^2] = \mathbb{E}[X^2] - \mathbb{E}[X]^2 = \mathbb{E}[X^2] - \rho^2$$

and also

$$\begin{aligned} \mathbb{E}[X^2] &= \sum_{k=1}^m \frac{d}{m} \langle v_k | \rho | v_k \rangle ((d+1)|v_k\rangle\langle v_k| - \mathbb{I})^2 \\ &= \frac{d(d^2-1)}{N} \sum_{k=1}^N \langle v_k | \rho | v_k \rangle |v_k\rangle\langle v_k| + \mathbb{I} \\ &= (d-1)(\rho + \mathbb{I}) + \mathbb{I}, \end{aligned}$$

according to Eq. (16). This allows us to bound the variance parameter:

$$\begin{aligned} \left\| \frac{1}{n} \sum_{i=1}^n \left(\frac{1}{n} (X - \mathbb{E}[X]) \right)^2 \right\|_\infty &= \frac{1}{n} \|(d+1)\rho + d\mathbb{I} - \rho^2\|_\infty \\ &\leq \frac{2d}{n} =: \sigma^2. \end{aligned}$$

The ratio $\frac{\sigma^2}{R} = 2$ indicates that any choice of $\tau \in [0, 2]$ will fall into the subgaussian regime of the matrix Bernstein inequality and Theorem 3 yields

$$\begin{aligned} \Pr[\|\hat{L}_n - \rho\|_\infty \geq \tau] &= \Pr\left[\left\| \frac{1}{n} \sum_{i=1}^n (X_i - \mathbb{E}[X_i]) \right\|_\infty \geq \tau\right] \\ &\leq de^{-\frac{3\tau^2 n}{16d}}. \end{aligned} \quad (27)$$

B. Concentration for (global) Pauli observables

We assume that the total number of samples n is distributed equally among the d^2 different Pauli measurements. Similar to before, unbiasedness (22) and the explicit characterization of the LI estimator (21) allow us

to write

$$\hat{L}_n - \rho = \sum_{k=1}^{d^2} \frac{1}{n} \sum_{i=1}^{n/d^2} \left(X_i^{(k)} - \mathbb{E}[X_i^{(k)}] \right).$$

Here, each $X_i^{(k)}$ is an independent instance of the random matrix $X^{(k)} = \pm dW_k$ with probability $\frac{1}{2}(1 \pm \text{tr}(W_k \rho))$ each. This is a sum of centered random matrices that are independent, but in general not identically distributed. However, independence alone suffices for applying Theorem 3. We note in passing that this would not be the case for earlier (weaker) versions of the matrix Bernstein inequality. Bound

$$\frac{1}{n} \|X_i^{(k)} - \mathbb{E}[X_i^{(k)}]\|_\infty = \frac{d}{n} \|(1 \pm \text{tr}(W_i \rho))\mathbb{I}\|_\infty \leq \frac{2d}{n} =: R$$

and use the fact that

$$\begin{aligned} \mathbb{E}\left[\left(X_i^{(k)} - \mathbb{E}[X_i^{(k)}]\right)^2\right] &= \mathbb{E}\left[\left(X^{(k)}\right)^2\right] - \mathbb{E}\left[X^{(k)}\right]^2 \\ &\leq \mathbb{E}\left[\left(X^{(k)}\right)^2\right] \end{aligned} \quad (28)$$

(in the positive semidefinite order) to considerably simplify the variance computation:

$$\frac{d^2}{n} \left\| \sum_{k=1}^{d^2} \mathbb{E}\left[\left(X^{(k)}\right)^2\right] \right\|_\infty = \frac{1}{n} \left\| \sum_{k=1}^{d^2} \mathbb{I} \right\|_\infty = \frac{d^2}{n} =: \sigma^2,$$

because $\left(X^{(k)}\right)^2 = \frac{1}{d^2}\mathbb{I}$. Applying Theorem 3 yields

$$\Pr[\|\hat{L}_n - \rho\|_\infty \geq \tau] \leq de^{-\frac{3\tau^2 n}{8}} \quad \tau \in [0, d/2].$$

C. Concentration for Pauli-basis measurements

Once more, we assume that the total budget of samples n is distributed equally among all 3^k Pauli basis choices. Unbiasedness of the LI estimator together with the explicit description (25) allows us to once more interpret $\hat{L}_n - \rho$ as a sum of independent, centered random matrices:

$$\hat{L}_n - \rho = \sum_{\mathbf{s}} \frac{1}{n} \sum_{i=1}^{n/3^k} \left(X_i^{(\mathbf{s})} - \mathbb{E}[X_i^{(\mathbf{s})}] \right).$$

For each $\mathbf{s} \in \{x, y, z\}^k$, $X_i^{(\mathbf{s})}$ is an independent copy of the random matrix

$$X^{(\mathbf{s})} = \bigotimes_{i=1}^k \left(3|b_{o_i}^{(s_i)}\rangle\langle b_{o_i}^{(s_i)}| - \mathbb{I} \right)$$

with probability $\langle b_o^{(s)} | \rho | b_o^{(s)} \rangle$ for each $o \in \{\pm 1\}^k$. Jensen's inequality implies

$$\begin{aligned} \frac{1}{n} \left\| X^{(s)} - \mathbb{E} \left[X^{(s)} \right] \right\|_{\infty} &\leq \frac{2}{n} \left\| X^{(s)} \right\|_{\infty} \\ &= \frac{2}{n} \prod_{i=1}^k \left\| 3 |b_{o_i}^{(s_i)}\rangle \langle b_{o_i}^{(s_i)}| - \mathbb{I} \right\|_{\infty} = \frac{2^{k+1}}{n} =: R. \end{aligned}$$

For the variance, we once more use (28) and compute

$$\begin{aligned} \frac{1}{n} \sum_s \mathbb{E} \left(X^{(s)} \right)^2 &= \sum_{s,o} \langle b_o^{(s)} | \rho | b_o^{(s)} \rangle \bigotimes_{i=1}^k \left(\mathcal{D}^{-1} \left(|b_o^{(s)}\rangle \langle b_o^{(s)}| \right) \right)^2 \\ &= \sum_s \sum_o \langle b_o^{(s)} | \rho | b_o^{(s)} \rangle \bigotimes_{i=1}^k \frac{5}{3} \mathcal{D}_{3/5} \left(|b_o^{(s)}\rangle \langle b_o^{(s)}| \right) \\ &= 5^k \mathcal{D}_{3/5}^{\otimes k} \left(\frac{1}{3^k} \sum_{s,o} \langle b_o^{(s)} | \rho | b_o^{(s)} \rangle |b_o^{(s)}\rangle \langle b_o^{(s)}| \right) \\ &= 5^k \mathcal{D}_{3/5}^{\otimes k} \left(\mathcal{D}_{1/3}^{\otimes k}(\rho) \right) = \frac{5^k}{n} \mathcal{D}_{1/5}^{\otimes k}(\rho), \end{aligned}$$

where we have used Eq. (26) and the fact that the combination of two depolarizing channels is again a depolarizing channel. This expression can be evaluated explicitly. For $\alpha \subset [k]$, let $\text{tr}_{\alpha}(\rho)$ denote the partial trace over all indices contained in α . Then, for $X = \bigotimes_{j=1}^k X_j \in \mathbb{H}_d$

$$\begin{aligned} 5^k \mathcal{D}_{1/5}^{\otimes k} \left(\bigotimes_{j=1}^k X_j \right) &= \bigotimes_{j=1}^k (\text{tr}(X_j) \mathbb{I} + X_j) \\ &= \sum_{\alpha \subset [k]} 2^{|\alpha|} \text{tr}_{\alpha}(X) \otimes \mathbb{I}^{\otimes \alpha} \end{aligned}$$

and this extends linearly to all of $\mathbb{H}_d \simeq \mathbb{H}_2^{\otimes k}$. Consequently,

$$\begin{aligned} \left\| \frac{1}{n} \sum_s \mathbb{E} \left[\left(X^{(s)} \right)^2 \right] \right\|_{\infty} &= \frac{1}{n} \left\| \sum_{\alpha \subset [k]} 2^{|\alpha|} \text{tr}_{\alpha}(\rho) \otimes \mathbb{I}^{\otimes \alpha} \right\|_{\infty} \\ &\leq \frac{1}{n} \sum_{\alpha \subset [k]} 2^{|\alpha|} \|\text{tr}_{\alpha}(\rho)\|_{\infty} \|\mathbb{I}^{\otimes \alpha}\|_{\infty} \\ &\leq \frac{1}{n} \sum_{\alpha \subset [k]} 2^{|\alpha|} = \frac{1}{n} \sum_{j=0}^k \binom{k}{j} 2^j \\ &= \frac{(2+1)^k}{n} = \frac{3^k}{n} =: \sigma^2. \end{aligned}$$

This estimate is actually tight for pure product states of the form $\rho = (|\psi\rangle\langle\psi|)^{\otimes k}$. We may now apply Theorem 3 to conclude

$$\Pr \left[\|\hat{L}_n - \rho\|_{\infty} \geq \tau \right] \leq d \exp \left(-\frac{3n\tau^2}{8 \times 3^k} \right) \quad \tau \in [0, 1].$$

VIII. CONVERSION OF CONFIDENCE REGIONS FROM OPERATOR NORM TO TRACE NORM

The final ingredient for the framework presented in this manuscript is a reliable way to transform operator-norm closeness of the (plain) least squares estimator \hat{L}_n into a statement about closeness of the PLS estimator $\hat{\rho}_n$ in trace distance. Recall that the optimization problem (13) admits an analytic solution [19]. Let $U \text{diag}(\lambda) U^{\dagger}$ be an eigenvalue decomposition of \hat{L}_n . Then,

$$\hat{\rho}_n = U \text{diag} \left([\lambda - x_0 \mathbf{1}]^+ \right) U^{\dagger}, \quad (29)$$

where x_0 is chosen such that $\text{tr}(\hat{\rho}_n) = 1$ and $[y]_i^+ = \max\{[y]_i, 0\}$ denotes thresholding on non-negative components. This solution is *unique*, provided that $\text{tr}(\hat{L}_n) = 1$, which is the case for all the least squares estimators we consider.

The conversion from closeness in operator norm to closeness in trace norm will introduce a factor that is proportional to the *effective rank* of the density matrix ρ , rather than a full dimensional factor. For $r \in \mathbb{N}$, we define the best rank- r approximation ρ_r of a quantum state $\rho \in \mathbb{H}_d$ as the optimal feasible point of

$$\sigma_r(\rho) = \underset{\text{rank}(Z) \leq r}{\text{minimize}} \quad \|\rho - Z\|_1. \quad (30)$$

This problem can be solved analytically. Let $\rho = \sum_{i=1}^d \lambda_i |x_i\rangle \langle x_i|$ be an eigenvalue decomposition with eigenvalues arranged in non-increasing order. Then,

$$Z^{\sharp} = \sum_{i=1}^r \lambda_i |x_i\rangle \langle x_i|, \quad \text{and} \quad \sigma_r(\rho) = \sum_{i=r+1}^d \lambda_i = 1 - \text{tr}(\rho_r),$$

highlighting that the best rank- r approximation is simply a truncation onto the r largest contributions in the eigenvalue decomposition. This truncated description is accurate if the residual error $\sigma_r(\rho)$ is small. If this is the case it is reasonable to say that ρ is well approximated by a rank- r matrix Z^{\sharp} and has *effective rank* r .

Proposition 1. *Suppose that $\hat{L}_n \in \mathbb{H}_d$ obeys $\text{tr}(\hat{L}_n) = 1$ and $\|\hat{L}_n - \rho\|_{\infty} \leq \tau$ for some quantum state $\rho \in \mathbb{H}_d$ and $\tau \geq 0$. Then, for any $r \in \mathbb{N}$ the PLS estimator $\hat{\rho}_n$ obeys*

$$\|\hat{\rho}_n - \rho\|_1 \leq 4r\tau + 2 \min \{ \sigma_r(\rho), \sigma_r(\hat{\rho}_n) \},$$

where $\sigma_r(\rho)$ is defined in Eq. (30).

The statement readily follows from combining two auxiliary results. The first one states that the threshold value x_0 in the analytic solution of $\hat{\rho}_n$ must be small if \hat{L}_n is operator-norm close to a quantum state.

Lemma 1. *Instantiate the assumptions from Proposition 1. Then, the threshold value in Eq. (29) obeys $x_0 \in [0, \tau]$.*

Proof. By assumption \hat{L}_n has unit trace. If it is in addition psd, $\hat{\rho}_n = \hat{L}_n$, because \hat{L}_n is already a quantum state and the projection is trivial ($x_0 = 0$). Otherwise, \hat{L}_n is indefinite and unit trace ensures that the positive part dominates. Hence, x_0 must be strictly positive to enforce $\text{tr}(\hat{\rho}_n) = 1$.

For the upper bound, let $P \in \mathbb{H}_d$ denote the orthogonal projection onto the range of $\hat{\rho}_n$. Then, $\hat{\rho}_n = P(\hat{L}_n - x_0\mathbb{I})P = P\hat{L}_nP - x_0P$, according to Eq. (29). In semidefinite order, this implies

$$\begin{aligned}\hat{\rho}_n &= P(\hat{L}_n - \rho)P - x_0P + P\rho P \\ &\leq (\|\hat{L}_n - \rho\|_\infty - x_0)P + P\rho P \\ &\leq (\tau - x_0)P + P\rho P,\end{aligned}$$

where the last line follows from the assumption $\|\hat{L}_n - \rho\|_\infty \leq \tau$. The trace preserves semidefinite order and we conclude

$$0 \leq \text{tr}(\hat{\rho}_n) - \text{tr}(P\rho P) \leq (\tau - x_0)\text{tr}(P)$$

which implies an upper bound of τ ($\text{tr}(P) > 0$). \square

The second technical lemma generalizes a result that is somewhat folklore in quantum information theory: the ‘‘effective rank’’ of a difference of two quantum states is proportional to the minimal rank of the two density operators involved.

Lemma 2. Fix $r \in \mathbb{N}$ and let $\rho, \sigma \in \mathbb{H}_d$ be quantum states. Then,

$$\|\rho - \sigma\|_1 \leq 2r\|\rho - \sigma\|_\infty + 2\min\{\sigma_r(\rho), \sigma_r(\sigma)\},$$

where the residual error $\sigma_r(\cdot)$ was defined in Eq. (30).

Proof. We can without loss of generality assume $\sigma_r(\rho) \leq \sigma_r(\sigma)$. Decompose ρ into $\rho_r + \rho_c$, where ρ_r is the best rank- r approximation (30) and $\rho_c = \rho - \rho_r$ denotes the ‘‘tail’’. By construction, both ρ_r and ρ_c are positive semidefinite matrices that obey $\sigma_r(\rho) = \text{tr}(\rho_c) = 1 - \text{tr}(\rho_r)$. The triangle inequality then implies

$$\|\rho - \sigma\|_1 \leq \|\rho_r - \sigma\|_1 + \sigma_r(\rho),$$

because $\|\rho_c\|_1 = \sigma_r(\rho)$. Next, let $P_+, P_- \in \mathbb{H}_d$ be the projections onto the positive and non-positive ranges of $\rho_r - \sigma$. By construction, P_+ has rank at most $\text{rank}(\rho_r) = r$ and the trace norm equals

$$\|\rho_r - \sigma\|_1 = \text{tr}(P_+(\rho_r - \sigma)) - \text{tr}(P_-(\rho_r - \sigma)).$$

On the other hand,

$$\sigma_r(\rho) = \text{tr}(\sigma - \rho_r) = -\text{tr}(P_+(\rho_r - \sigma)) - \text{tr}(P_-(\rho_r - \sigma)),$$

because $P_+ + P_- = \mathbb{I}$. Combining both relations yields

$$\begin{aligned}\|\rho_r - \sigma\|_1 &= 2\text{tr}(P_+(\rho_r - \sigma)) + \sigma_r(\rho) \\ &\leq 2\text{tr}(P_+(\rho - \sigma)) + \sigma_r(\rho) \\ &\leq 2\|P_+\|_\infty\|\rho - \sigma\|_1 + \sigma_r(\rho),\end{aligned}$$

where we have used $\text{tr}(P_+\rho_c) \geq 0$ and Hoelder’s inequality. Finally, note that $\|P_+\|_1 = \text{rank}(P_+) = r$ by construction and the claim follows. \square

The main result of this section is a rather straightforward combination of these two technical statements.

Proof of Proposition 1. Fix $r \in \mathbb{N}$ and use, Lemma 2 to conclude

$$\|\hat{\rho}_n - \rho\|_1 \leq 2r\|\hat{\rho}_n - \rho\|_\infty + 2\min\{\sigma_r(\rho), \sigma_r(\hat{\rho}_n)\}.$$

Next, note that according to (29), $\hat{\rho}_n$ may be viewed as the positive definite part of the matrix $\hat{L}_n - x_0\mathbb{I}$. Such a restriction to the positive part can never increase the operator norm distance to another positive semidefinite matrix. Hence,

$$\|\hat{\rho}_n - \rho\|_\infty \leq \|\hat{L}_n - \rho\|_\infty + |x_0|\|\mathbb{I}\|_\infty \leq 2\tau,$$

where the last inequality follows from Lemma 1. \square

IX. PROOF OF THE MAIN RESULT

By now we have everything in place to provide a complete proof of the main result of this work.

Theorem 4. Let $\rho \in \mathbb{H}_d$ be a state. Suppose that we either perform n structured POVM measurements (set $g(d) = 2d$), n Pauli observable measurements (set $g(d) = d^2$), or n Pauli basis measurements (set $g(d) = d^{1.6}$). Then, for any $r \in \mathbb{N}$ and $\epsilon \in [0, 1]$, the PLS estimator $\hat{\rho}_n$ (13) obeys

$$\Pr[\|\hat{\rho}_n - \rho\|_1 \geq \epsilon + 2\min\{\sigma_r(\rho), \sigma_r(\hat{\rho}_n)\}] \leq de^{-\frac{n\epsilon^2}{43g(d)r^2}},$$

where $\sigma_r(\rho), \sigma_r(\hat{\rho}_n)$ denote the residual error of approximating ρ and $\hat{\rho}_n$ by a rank- r matrix (30).

Note that Theorem 1 is an immediate consequence of this more general result: simply set $r = \min\{\text{rank}(\rho), \text{rank}(\hat{\rho}_n)\}$ which in turn ensures $\min\{\sigma_r(\rho), \sigma_r(\hat{\rho}_n)\} = 0$.

However, unlike this specification, Theorem 4 does feature an additional degree of freedom. The parameter $r \in \mathbb{N}$ allows for interpolating between small values (small sampling rate, but a potentially large reconstruction error) and large values (high sampling rate, but low reconstruction error). This tradeoff is particularly benign for quantum states that are approximately low-rank. Due to experimental imperfections, such states arise naturally in many experiments that aim at generating a pure quantum state. We illustrate this by means of the following caricature of a faulty state preparation protocol. Suppose that an apparatus either produces a target state $|\psi\rangle\langle\psi|$ perfectly, or fails completely, in the sense that it outputs a maximally mixed state. Then, the resulting state is

$$\rho = (1 - p)|\psi\rangle\langle\psi| + \frac{p}{d}\mathbb{I},$$

where $p \in [0, 1]$ denotes the probability of failure. This state has clearly full rank and Corollary 1 requires at

least $n \geq 43 \frac{g(d)d^2}{\epsilon^2} \log(d/\delta)$ samples to estimate it up to trace-norm accuracy ϵ with high probability. In contrast, Theorem 4 ensures that already $n \geq 43 \frac{g(d)}{\epsilon^2} \log(d/\delta)$ samples suffice to ensure that, with high probability, the PLS estimator obeys $\|\hat{\rho}_n - \rho\|_1 \leq \epsilon + 2p$. For sufficiently high success probabilities/low accuracy ($\epsilon \geq 2p/d$) this clearly outperforms the original statement.

Proof of Theorem 4. We illustrate the proof for structured POVMs – the other settings are completely analogous. Fix $\epsilon \in [0, 1]$, $r \in \mathbb{N}$ and set $\tau = \frac{\epsilon}{4r}$. Then, the main result of Sec. VII A – Equation (27) – ensures that the least squares estimator \hat{L}_n obeys

$$\|\hat{L}_n - \rho\|_\infty \leq \tau = \frac{\epsilon}{4r} \quad (31)$$

with probability of failure bounded by $de^{-\frac{\epsilon^2 n}{86dr^2}}$. Assuming that this condition is true, Proposition 1 readily yields $\|\hat{\rho}_n - \rho\|_1 \leq \epsilon + 2 \min\{\sigma_r(\rho), \sigma_r(\hat{\rho}_n)\}$. \square

X. IMPROVED CONVERGENCE GUARANTEES FOR THE UNIFORM POVM

All the convergence results derived so far feature an additional $\log(d)$ -factor. This is a consequence of the matrix Bernstein inequality that proved instrumental in deriving these results. One can show that such an additional factor necessarily features in *all* matrix concentration inequalities that are based on exclusively first and second moments of the random matrices in question [48].

However, the following question remains: is this $\log(d)$ -factor in Theorem 1 an artifact of the proof technique, or is it an intrinsic feature of tomography via projected least squares?

In this section we rule out the second possibility: a different proof technique allows for avoiding this $\log(d)$ -factor, provided that the POVM is sufficiently symmetric and well-behaved. More precisely, we re-visit the uniform POVM $\{d|v\rangle\langle v|dv\}_{v \in S^d}$ and exploit the fact that Eq. (14) completely characterizes *all* moments of the resulting outcome distribution. This opens the door for applying very strong proof techniques from large dimensional probability theory that found wide-spread applications in a variety of subjects, including compressed sensing [49] and, more recently, quantum information theory [50]. We believe that this technique may be of independent interest and find it therefore worthwhile to present it in a self-contained fashion. Roughly speaking, it is based on the following steps:

- (o) *Reformulation:* the operator norm of a random hermitian matrix $A \in \mathbb{H}_d$ admits a variational definition:

$$\|A\|_\infty = \max_{y \in S^d} |\langle y|A|y\rangle|. \quad (32)$$

- (i) *Discretization:* replace the maximization over the entire complex unit sphere S^d by a maximization over a finite point set \mathcal{N} that covers S^d to sufficiently high accuracy (*covering net*).

- (ii) *Concentration:* fix $y \in \mathcal{N}$ and show that the scalar random variable $s_y = \langle y|A|y\rangle$ concentrates sharply around its expectation value.

- (iii) *Union bound:* apply a union bound over all $|\mathcal{N}|$ random variables s_y to obtain an upper bound on the operator norm $\|A\|_\infty$.

Ideally the tail bound from (iii) is sharp enough to “counter-balance” the $|\mathcal{N}|$ -pre-factor that results from the union bound in step (iv). Should this be not the case, more sophisticated methods, like generic chaining [51], may still allow for drawing non-trivial conclusions. Fortunately, for the task at hand, this turns out to not be necessary and the rather naive strategy sketched above suffices to achieve a result that is (provably) optimal up to a constant factors:

Theorem 5. *Suppose that we perform n independent uniform POVM measurements on a quantum state $\rho \in \mathbb{H}_d$. Then, the associated least squares estimator \hat{L}_n obeys*

$$\Pr[\|\hat{L}_n - \rho\|_\infty \geq \tau] \leq 2 \exp(c_1 d - c_2 n \tau^2),$$

In particular, $n \geq C \frac{d}{\tau^2} \log(1/\delta)$ suffices to ensure $\|\hat{L}_n - \rho\|_\infty \leq \tau$ with probability at least $1 - \delta$. Here $c_1, c_2, C > 0$ denote constants of sufficient size.

No effort has been made to optimize the constants. The proof presented here yields $c_1 = 2 \log(3)$ and $c_2 = \frac{1}{480}$ which could be further improved by a more careful analysis. Importantly, the second part of this statement can be combined with Proposition 1 to readily deduce the last technical result of the main text:

Corollary 2 (Re-statement of Theorem 5). *For any rank- r state ρ , a number of $n \geq C \frac{r^2 d}{\epsilon^2} \log(1/\delta)$ uniform POVM measurements suffice to ensure $\|\hat{\rho}_{(n)}^\# - \rho\|_1 \leq \epsilon$ with probability at least $1 - \delta$.*

Not only does this statement reproduce the best known sampling rates for tomography with independent measurements [15], it also exactly matches lower bounds on the minimal sample complexity associated with *any* tomographic procedure that may apply in this setting [16, Table I].

The remainder of this section is dedicated to proving Theorem 5. For the sake of accessibility, we will divide this proof into three subsections that contain the steps summarized above.

A. Step I: Reformulation and discretization

Suppose that we perform n uniform POVM measurements on a fixed quantum state $\rho \in \mathbb{H}_d$. Then, the least squares estimator is equivalent to a sum of i.i.d. random matrices:

$$\hat{L}_n = \frac{1}{n} \sum_{i=1}^n X_i.$$

Each X_i is an independent copy of the random matrix X that assumes the value $(d+1)|v\rangle\langle v| - \mathbb{I}$ with probability $d\langle v|\rho|v\rangle dv$ and v may range over the entire complex unit sphere. Unbiasedness of this estimator in turn implies

$$\begin{aligned} \|\hat{L}_n - \rho\|_\infty &= \left\| \frac{1}{n} \sum_{i=1}^n (X_i - \mathbb{E}[X_i]) \right\|_\infty \\ &= \max_{y \in \mathbb{S}^d} \left| \langle y | \frac{1}{n} \sum_{i=1}^n (X_i - \mathbb{E}[X_i]) | y \rangle \right|. \end{aligned}$$

Next, we employ a result that is somewhat folklore in random matrix theory, see e.g. [52, Lemma 5.3]. It states that the maximum over the entire unit sphere may be replaced by a maximum over certain finite point sets, called *covering nets*: A covering-net of \mathbb{S}^d with fineness $\theta > 0$ is a finite set of unit vectors $\{z_j\}_{j=1}^N \subseteq \mathbb{S}^d$ that covers the entire (complex) unit sphere in the sense that every $y \in \mathbb{S}^d$ is at least θ -close to a point in the net.

Lemma 3. *Let $\mathcal{N}_\theta = \{z_j\}_{j=1}^N$ be a covering net of \mathbb{S}^d with fineness θ . Then, for any matrix $A \in \mathbb{H}_d$:*

$$\max_{j \in [N]} |\langle z_j | A | z_j \rangle| \leq \|A\|_\infty \leq \frac{1}{1-2\theta} \max_{j \in [N]} |\langle z_j | A | z_j \rangle|.$$

This result highlights that already a rather coarse net suffices to get reasonable approximations to the operator norm. Here, we choose $\theta = \frac{1}{4}$ which, while certainly not optimal, simplifies exposition. In particular,

$$\|\hat{L}_n - \rho\|_\infty \leq 2 \max_{j \in [N]} \left| \frac{1}{n} \sum_{i=1}^n \langle z_j | X_i - \mathbb{E}[X_i] | z_j \rangle \right|, \quad (33)$$

where the maximization is over a covering net of fineness $\theta = \frac{1}{4}$.

B. Step II: concentration

Note that the right hand side of Eq. (33) corresponds to a maximum over N different random variables – each of them labeled by a unit vector z_j in the net. Let $z \in \mathbb{S}^d$ be such a vector. Then, the associated random variable itself corresponds to an empirical average of n i.i.d. variables:

$$s_z = \langle z | X - \mathbb{E}[X] | z \rangle.$$

Clearly, s_z obeys $\mathbb{E}[s_z] = 0$ and, more importantly, has sub-exponential moment growth. While this follows directly from the fact that s_z is bounded, the following result highlights that this tail-behavior is actually independent of the ambient dimension.

Lemma 4. *Fix $z \in \mathbb{S}^d$. Then for any integer $p \geq 2$, the random variable s_z obeys*

$$\mathbb{E}[|s_z|^p] \leq 27 \times 6^{p-2} p!$$

We divert the proof of this statement to the end of this section and content ourselves with emphasizing that the closed form expression of the frame operator (14) is essential for bounding all moments simultaneously. More relevant to the task at hand is that such a moment behavior ensures that the tails of the distribution of s_z follow an exponential decay: $\Pr[|s_z| \geq t] \leq e^{-ct}$, where c is a constant independent of the dimension d . Strong classical concentration inequalities apply for sums of i.i.d. random variables that exhibit such sub-exponential behavior. We choose to apply a rather general version of the classical Bernstein inequality, see e.g. [49, Theorem 7.30].

Theorem 6. *Let $s_1, \dots, s_n \in \mathbb{R}$ i.i.d. copies of a mean-zero random variable s that obeys $\mathbb{E}[|s|^p] \leq p! R^{p-2} \sigma^2 / 2$ for all integers $p \geq 2$, where $R, \sigma^2 > 0$ are constants. Then, for all $t > 0$,*

$$\Pr \left[\left| \sum_{i=1}^n s_i \right| \geq t \right] \leq 2 \exp \left(- \frac{t^2/2}{n\sigma^2 + Rt} \right).$$

Lemma 4 ensures that the random variable s_z meets this requirement with $\sigma^2 = 54$ and $R = 6$. Hence, the following Corollary is an immediate consequence of Theorem 6.

Corollary 3. *Fix $z \in \mathbb{S}^d$. Then, for any $t \in [0, 1]$*

$$\Pr \left[\left| \frac{1}{n} \sum_{i=1}^n \langle z | X_i - \mathbb{E}[X_i] | z \rangle \right| \geq t \right] \leq 2e^{-\frac{nt^2}{120}}.$$

C. Step III: union bound

Recall that Eq. (33) upper-bounds $\|\hat{L}_n - \rho\|_\infty$ by a maximum over finitely many random variables, each of which is controlled by the strong exponential tail inequality from Corollary 3. To exploit this, we fix $\tau \in [0, 1]$ and apply a union bound (also known as Boole's inequality) over all these different random variables to obtain

$$\begin{aligned} &\Pr \left[\|\hat{L}_n - \rho\|_\infty \geq \tau \right] \\ &\leq \Pr \left[\max_{j \in [N]} \left| \frac{1}{n} \sum_{i=1}^n \langle z_j | X_i - \mathbb{E}[X_i] | z_j \rangle \right| \geq \frac{\tau}{2} \right] \\ &\leq N \max_{j \in [N]} \Pr \left[\left| \frac{1}{n} \sum_{i=1}^n \langle z_j | X_i - \mathbb{E}[X_i] | z_j \rangle \right| \geq \frac{\tau}{2} \right] \leq 2Ne^{-\frac{n\tau^2}{480}}, \end{aligned}$$

where the last line is due to Corollary 3. Here, $N = |\mathcal{N}_{\frac{1}{4}}|$ denotes the cardinality of a covering net for the complex unit sphere S^d with fineness $\theta = \frac{1}{4}$. The complex unit sphere admits an isometric embedding into the real-valued unit sphere in $2d$ -dimensions: Map real and imaginary parts of each complex vector component onto two distinct real parameters. This map preserves Euclidean lengths and, by extension, also the geometry of the unit sphere. Volumetric upper bounds on the cardinality of covering nets for the $2d$ -dimensional real-valued unit sphere are widely known, see e.g. [49, Proposition C.3] and [52, Lemma 5.2]: $|\mathcal{N}_{\theta}| \leq (1 + \frac{2}{\theta})^{2d}$. Since a fineness of $\theta = \frac{1}{4}$ suffices for our purpose, we can conclude $N \leq 3^{2d}$ and consequently,

$$\Pr [\|\hat{L}_n - \rho\|_{\infty} \geq \tau] \leq 2 \times 3^{2d} e^{-\frac{n\tau^2}{480}} = 2e^{2\log(3)d - \frac{n\tau^2}{480}}$$

This concludes the proof of Theorem 5.

D. Proof of Lemma 4

Recall that, by assumption, the random matrix X assumes the value $X = (d+1)|v\rangle\langle v| - \mathbb{I}$ with probability $\langle v|\rho|v\rangle dv$, where v may range over the entire complex unit sphere S^d . Moreover, $\mathbb{E}[X] = \rho$. For fixed $z \in S^d$, we may therefore write

$$s_z = \langle z|X - \mathbb{E}[X]|z\rangle = (d+1)\langle v|B|v\rangle,$$

where $B = |z\rangle\langle z| - \frac{1+\langle z|\rho|z\rangle}{d+1}\mathbb{I} \in \mathbb{H}_d$ has bounded trace norm

$$\|B\|_1 \leq 1 + (1 + \langle z|\rho|z\rangle) \leq 3. \quad (34)$$

Next, recall a basic identity from matrix analysis that states

$$|\langle v|B|v\rangle| = |\text{tr}(|v\rangle\langle v|B)| \leq \text{tr}(|v\rangle\langle v| |B|),$$

where $|B| = \sqrt{B^2}$ denotes the absolute value of the matrix B . Also, the Schatten- p norms of matrices and their absolute values coincides, in particular $\|B\|_1 = \text{tr}(|B|) = \| |B| \|_1$. We can use this trick to absorb the absolute value in the moment computation. More precisely, fix an integer $p \geq 2$ and note that $\mathbb{E}[|s_z|^p]$ obeys

$$\mathbb{E}[|(d+1)\langle v|B|v\rangle|^p] \leq (d+1)^p \mathbb{E}[\text{tr}(|v\rangle\langle v| |B|)^p].$$

We can now include the distribution of the random matrices X , and – by extension – $|v\rangle\langle v|$ – to compute

$$\begin{aligned} \mathbb{E}[|s_z|^p] &\leq (d+1)^p \mathbb{E}[\text{tr}(|v\rangle\langle v| |B|)^p] \\ &= d(d+1)^p \int_{S^d} \langle v|\rho|v\rangle \text{tr}(|v\rangle\langle v| |B|)^p dv \\ &= d(d+1)^p \text{tr} \left(\int_{S^d} (|v\rangle\langle v|)^{\otimes (p+1)} \rho \otimes |B|^{\otimes p} \right) \\ &= d(d+1)^p \binom{d+p}{p+1}^{-1} \text{tr} \left(P_{\text{Sym}^{(p+1)}} \rho \otimes |B|^{\otimes p} \right), \end{aligned}$$

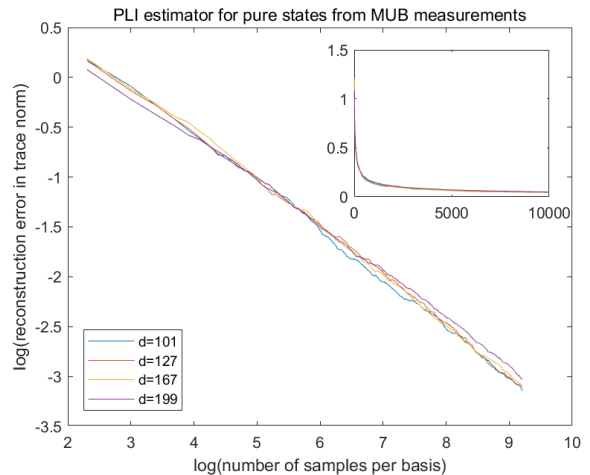


FIG. 3: log trace distance error vs. log sample size for different prime dimensions d . Inset: ordinary plot of the same data.

where the last equation is due to Eq. (14). Next, we note that Hoelder's inequality implies

$$\text{tr} \left(P_{\text{Sym}^{(p+1)}} \rho \otimes |B|^{\otimes p} \right) \leq \|P_{\text{Sym}^{p+1}}\|_{\infty} \|\rho\|_1 \| |B| \|_1^p \leq 3^p,$$

because $P_{\text{Sym}^{(p+1)}}$ is an orthogonal projector, ρ is a quantum state and B is bounded in trace norm (34). For the remaining pre-factor we use the crude bound

$$d(d+1)^p \binom{d+p}{p+1}^{-1} \leq (p+1)! \leq 3 \times 2^{p-2} p!$$

to establish the statement.

XI. ADDITIONAL NUMERICAL EXPERIMENTS

Maximal sets of mutually unbiased bases (MUBs) form a structured POVM (2-design) that lends itself to numerical investigation. Efficient algebraic constructions of MUBs exist in prime power dimensions $d = p^k$ [53–55]. To further underline the implicit advantage of low-rank we fix a prime dimension d and choose a pure state uniformly from the Haar measure on the complex unit sphere in d dimensions. We compute the outcome probabilities for each of the $d+1$ different MUB measurements. We then sample outcomes from each distribution a total of $\frac{n}{d+1}$ times and compute the estimator $\hat{\rho}_n$ associated with the total frequency statistics. Figure 3 shows the relation between reconstruction error (in trace distance) and the number of samples per basis on a log–log-scale for different prime dimensions between $d = 100$ and $d = 200$. This figure suggests that the rate of convergence only depends linearly on the ambient dimension – the additional $\log(d)$ -factor in the main result for structured POVMs is barely visible.