Honda Research Institute EU
Honda Research Institute Europe GmbH
https://www.honda-ri.de/

Hierarchically discriminating Haar-randomness in quantum states from a black-box device

Xavier Bonet-Monroig, Hao Wang, Adrián Pérez-Salinas

2026

Preprint:

This is an accepted article published in New Journal of Physics. The final authenticated version is available online at:
https://doi.org/https://doi.org/10.1088/1367-2630/ae868b

New Journal of Physics
The open access journal at the forefront of physics

Deutsche Physikalische Gesellschaft DPG IOP Institute of Physics

PAPER • OPEN ACCESS

Hierarchically discriminating Haar-randomness in quantum states from a black-box device

To cite this article: Xavier Bonet-Monroig *et al* 2026 *New J. Phys.* **28** 074512

View the article online for updates and enhancements.

You may also like

- Device-independent randomness generation from several Bell estimators
Olmo Nieto-Silleras, Cédric Bamps, Jonathan Silman et al.
- Increased certification of semi-device independent random numbers using many inputs and more post-processing
Piotr Mironowicz, Armin Tavakoli, Alley Hameedi et al.
- Loss of information in quantum guessing game
Martin Plesch and Matej Pivoluska

This content was downloaded from IP address 62.80.17.92 on 23/07/2026 at 10:54

New Journal of Physics
The open access journal at the forefront of physics

Deutsche Physikalische Gesellschaft DPG
IOP Institute of Physics

Published in partnership with: Deutsche Physikalische Gesellschaft and the Institute of Physics

Check for updates

PAPER

Hierarchically discriminating Haar-randomness in quantum states from a black-box device

Xavier Bonet-Monroig[1,2,3], **Hao Wang**[2,4] and **Adrián Pérez-Salinas**[2,3,*]

[1] Honda Research Institute Europe GmbH, Carl-Legien-Str. 30, 63073 Offenbach, Germany
[2] Universiteit Leiden, Leiden, The Netherlands
[3] Instituut-Lorentz, Universiteit Leiden, Niels Bohrweg 2, 2333 CA Leiden, The Netherlands
[4] LIACS, Universiteit Leiden, Niels Bohrweg 1, 2333 CA Leiden, The Netherlands
[*] Author to whom any correspondence should be addressed.

E-mail: adperez@phys.ethz.ch, xavier.bonet@honda-ri.de and h.wang@liacs.leidenuniv.nl

Keywords: Haar-random, *t*-designs, state discrimination, Dirichlet distribution

OPEN ACCESS

RECEIVED
19 February 2026

REVISED
27 May 2026

ACCEPTED FOR PUBLICATION
6 July 2026

PUBLISHED
23 July 2026

Original content from this work may be used under the terms of the Creative Commons Attribution 4.0 licence.

Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

Abstract

The concept of randomness in quantum computing has been central to constructing benchmarking tools, cryptographic protocols, as well as a proof of beyond-classical computation. Discerning whether quantum states (or unitaries) are randomly distributed is a computational task that requires an enormous amount of quantum computational resources. This work addresses such a challenge by introducing a hierarchical discrimination algorithm to efficiently test if a set of states S generated from a black-box quantum device with an unknown distribution is (in)compatible with a random distribution. To this end, we reduce the complexity of the problem by selecting an observable with known spectrum to study the statistical properties of its expectation values with respect to the quantum states from an unknown (black-box) quantum device. Concurrently, we use our first technical result, a connection between Haar-randomness and the Dirichlet distribution, to analytically compute Haar-random moments of the observable. Our Haar-random discriminator test is then simply to compare those statistical moments, such that if S fails the test, it is enough to state that the quantum device does not output randomly distributed states. Else, we can not (yet) confirm that the states follow a Haar-random distribution. We further provide an extension to this algorithm by permutation- and unitary-equivalent randomization of the observable at increasing computational resources, which allows us to more accurately state whether S is compatible with Haar-randomness. We envision the use of the discriminator test as a quantum device benchmark, by discriminating whether the states generated are incompatible with Haar-randomness.

1. Introduction

Imagine we are given a black-box quantum device that outputs quantum states on demand, and we are asked to efficiently discriminate if their underlying distribution resembles the Haar-random distribution. Realistically, the black-box might implement randomly chosen sequences of fixed-size quantum circuits, such as a parametrized quantum circuit, or a fixed quantum circuit with random initial states, in all cases possibly without any prior information about their structure. While this question might be seen just as a thought experiment, the problem of efficiently characterizing the distribution of quantum states is at the core of many problems in quantum information and computation; from quantum advantage and its verification with random circuits or boson sampling [1–5] to cryptography [6, 7]. It even plays a central role in the characterization and benchmarking of quantum hardware central [8, 9], or the performance of variational algorithms [10–13], to mention a few.

© 2026 The Author(s). Published by IOP Publishing Ltd on behalf of the Institute of Physics and Deutsche Physikalische Gesellschaft

As a first step to characterize the output distribution, one might be tempted to take the naive approach: perform full tomography of the quantum states and compare their statistical moments to those of the Haar-random distribution. Due to the high symmetry of the Haar-random distribution, full tomography quickly faces the curse of dimensionality, and thus requires an exponential cost in quantum resources.

In this work we present an algorithm that initially avoids the exponential cost by hierarchically discriminating the compatibility of the set of states to be Haar randomly distributed. We base our method on the following observation: a Haar-random distribution must resemble random with respect to all observables, but only one observable showing discrepancies suffices to falsify randomness. Our approach aims to hierarchically allocate computational resources for the analysis of expectation values in such a way that the falsification of t-designs is possible before facing the unavoidable curse-of-dimensionality. The first building block of the method is the use of the Dirichlet distribution [14] to analytically treat expectation values of a known quantum observable w.r.t. Haar-random states as a random variable. Next, we further extend our method by designing permutation- and unitary-equivalent families of observables to address the true verification of a t-design. By randomly selecting observables in an ordered manner, we discriminate between subspaces of several observables of t-designs. More importantly, neither this nor any other method suffices to prove that a set of states is Haar-random without facing the curse-of-dimensionality.

This paper is organized as follows. Section 2 provides a short background on Haar-randomness. Section 3 describes the mathematical foundations of this protocol and its link to the Dirichlet distribution. In section 4 we present the protocol to verify the Haar-randomness of sets of states with respect to a given observable, including extensions for true verification of randomness. We conclude and give an overview of future ideas in section 5.

2. Background

Random unitary operations or quantum states are defined via the Haar measure, that is the only measure over the unitary group that remains invariant under the action of an arbitrary operation $U \in \mathcal{SU}(N)$. For unitary operations, the Haar measure $\mu_{\mathrm{H}}^{(U)}$ satisfies

$$\int_{V\in\mathcal{SU}(N)} \mathrm{d}\mu_{\mathrm{H}}^{(U)}(V) f(V) = \int_{V\in\mathcal{SU}(N)} \mathrm{d}\mu_{\mathrm{H}}^{(U)} f(UV), \qquad \forall U \in \mathcal{SU}(N). \tag{1}$$

With respect to states, the Haar-random distribution satisfies

$$\int_{\substack{\psi\in\mathbb{C}^N(N)\\ \langle\psi|\psi\rangle=1}} \mathrm{d}\mu_{\mathrm{H}}(\psi) f(|\psi\rangle) = \int_{\substack{\psi\in\mathbb{C}^N(N)\\ \langle\psi|\psi\rangle=1}} \mathrm{d}\mu_{\mathrm{H}}(\psi) f(U|\psi\rangle), \qquad \forall U \in \mathcal{SU}(N). \tag{2}$$

Probability distributions can be studied through their generating functions, or equivalently statistical moments, such as mean and variance. In the language of Haar-randomness, statistical moments of order t are related to t-designs, that is, ensembles of states S satisfying

$$\rho_t(S) \equiv \mathbb{E}_{|\psi\rangle\sim S}\left(\left(|\psi\rangle\langle\psi|\right)^{\otimes t}\right) = \int \mathrm{d}\mu_{\mathrm{H}}(\phi)\left(|\phi\rangle\langle\phi|\right)^{\otimes t}, \tag{3}$$

where the integration limits comprising all possible quantum states have been omitted for readability. The generation of t-designs is a necessary ingredient to many tasks in quantum computing, such as benchmarking [9], and its preparation has been extensively studied. In particular, t-designs can be created with $\mathcal{O}(\mathrm{poly}(t))$ quantum gates [15, 16], and the minimal number of states required to construct a t-design increases as $\Omega(N^t)$ for $t \ll N$ [17, 18].

Another interpretation of t-designs is obtained by computing expectation values over quantum observables $\hat{O}$ such that if a set of states S forms a t-design

$$\mu_t\left(\hat{O}, S\right) \equiv \mathbb{E}_{\psi\sim S}\left(\left\langle\psi\middle|\hat{O}\middle|\psi\right\rangle^t\right) = \int \mathrm{d}\mu_{\mathrm{H}}(\phi)\left(\left\langle\phi\middle|\hat{O}\middle|\phi\right\rangle\right)^t \equiv \mu_t\left(\hat{O}\right), \qquad \forall \hat{O}. \tag{4}$$

where the omission of the set of states indicates the Haar-random distribution. This condition holds for any quantum observable. Assume now a $\hat{O}$ is given the by rank-1 projector $U|x\rangle\langle x|U^\dagger$, such that $|x\rangle$ is

an element of the computational basis, and U is chosen so that all projectors form an informationally-complete POVM. By definition, Haar-randomness induces a strong symmetry across all dimensions of the Hilbert space of interest, and the random variables $\langle x|\, U^\dagger\, |\psi\rangle$, for $|\psi\rangle$ Haar-random, are identically distributed. Therefore, assessing if S is a t-design implies verifying its compatibility on a set of linearly independent observables of size $\mathcal{O}(N^2)$. Yet, a single observable that does not meet equation (4) is sufficient to demonstrate an incompatibility between S and t-designs. Therefore, verifying Haar-randomness is a much harder task than falsifying it. The latter fact is the building block of our Haar-random discriminator algorithm; we aim to efficiently discriminate (e.g. falsify) if a set of quantum states is compatible with t-designs.

3. Haar meets Dirichlet

The first contribution of this work is to reinterpret the expectation values of Haar-randomly distributed quantum states as random variables with specific distributions. To this end, we utilize the Dirichlet distribution [14], a mathematical tool that allows us to derive an analytical formula for the statistical moments of expectation values $\mu_t(\hat{O})$. The connection between Haar and Dirichlet has been previously noticed [2]. For example, the first experiments on quantum advantages base their theoretical separation between classical and quantum computers on this observation, under the commonly used name of Porter-Thomas distribution [4].

The intuitive connection between random states and the Dirichlet distribution is straightforward. The Dirichlet distribution ($\mathrm{Dir}(\boldsymbol{\alpha})$) is a model for multidimensional random variables $\vec{x} = (x_1, x_2, \ldots, x_m)$ with the constraints that $x_i \in [0,1]$, and $\sum_i x_i = 1$, and fully determined by a vector of parameters $\vec{\alpha}$ with the same dimensionality as $\vec{x}$. That is, the Dirichlet distribution is a probability distribution over the simplex. On the other hand, Haar-random states correspond to states for which $|\langle i|\psi|\rangle|^2$ are randomly and uniformly distributed, while the simplex constraint is imposed by normalization. A convenient property of the Dirichlet distribution is the existence of a closed form for any statistical moment, for instance $\mathbb{E}_{\mathrm{Dir}(\boldsymbol{\alpha})}(x_i) = \alpha_i / \|\boldsymbol{\alpha}\|_1$. We refer the reader to appendix B for details about the Dirichlet distribution.

Lemma 1. *Let $|\psi\rangle$ be a random state drawn from the Haar-random distribution, and let $\boldsymbol{x}$ be a random variable defined as $x_i = |\langle\psi|i\rangle|^2$, for an arbitrary unitary operation U, where $|i\rangle$ is the ith element of the reference basis. Then*

$$\boldsymbol{x} \sim \mathrm{Dir}\,(\mathbf{1})\,, \tag{5}$$

with $\mathrm{Dir}\,(\boldsymbol{\alpha})$ *being the Dirichlet distribution, and* $\mathbf{1}$ *being a vector where all entries are* 1.

See appendix B for a proof of lemma 1. This result is a direct implication of the identical distributions of random variables $\langle x|\psi\rangle$, together with the unitarity constraint of quantum states. For completeness, we recall that the marginals of this distribution are all symmetric and more commonly known as the Porter-Thomas distribution [2, 19] used to build the so-called *quantum supremacy* experiment [4],

$$\mathrm{Prob}\left(|\langle\psi|i\rangle|^2 \geqslant p\right) = (1-p)^{N-1} \lesssim \exp\left(-pN\right). \tag{6}$$

Using this interpretation of Haar-random states, lemma 1 allows us to analytically compute the expectation value of a Haar-random ensemble of states with respect to an observable $\hat{O}$. For the sake of simplicity, in the remainder of the paper, we consider $\hat{O}$ a positive semi-definite Hermitian observable.

Theorem 1. *Let $|\psi\rangle$ be a random state drawn from the Haar-random distribution. Let $\hat{O}$ be an observable, with G distinct eigenvalues $\boldsymbol{\lambda} = (\lambda_1, \lambda_2, \ldots, \lambda_G), \lambda_i \geqslant 0$, and algebraic multiplicities $\boldsymbol{m} = (m_1, m_2, \ldots, m_G)$. Then, the expectation value is a random variable*

$$\langle\psi|\,\hat{O}\,|\psi\rangle = \boldsymbol{\lambda}\cdot\boldsymbol{x}, \tag{7}$$

where $\boldsymbol{x}$ is a random variable $\boldsymbol{x} \sim \mathrm{Dir}\,(\boldsymbol{m})$.

To show this, it suffices to observe that the Dirichlet-based description from lemma 1 holds in any basis, in particular in the one that diagonalizes $\hat{O}$. Aggregation properties of coordinates of Dirichlet-random variables suffice for this description. A detailed proof can be found in appendix B.

From the results in theorem 1, it is direct to compute the statistical moments of $\langle\psi|\,\hat{O}\,|\psi\rangle$, for Haar-random states.

Lemma 2. *Let $|\psi\rangle$ be a random state drawn from the Haar-random distribution. Let $\hat{O}$ have G eigenvalues $\boldsymbol{\lambda} = (\lambda_1, \lambda_2, \ldots, \lambda_G)$ and multiplicities $\boldsymbol{m} = (m_1, m_2, \ldots, m_G)$. Then, the $\langle\psi|\,\hat{O}\,|\psi\rangle$ is a random variable with statistical moments*

$$\mu_t\left(\hat{O}\right) = \sum_{\substack{\boldsymbol{k} \\ \|\boldsymbol{k}\|_1 = t}} \binom{t}{\boldsymbol{k}} \left(\prod_{i=1}^{G} \lambda_i^{k_i}\right) \frac{\Gamma(N)}{\Gamma(N+t)} \prod_{i=1}^{G} \frac{\Gamma(m_i + k_i)}{\Gamma(m_i)}, \tag{8}$$

with $\Gamma(\cdot)$ being Euler's gamma function. For integers $\Gamma(n+1) = n!$.

These computations follow immediately from the multinomial theorem and the properties of Dirichlet distribution. A detailed proof of this lemma is provided in appendix C. As an example, consider $\hat{O} = (I+Z)/2$, with eigenvalues $\lambda = \{0, 1\}$ and multiplicities $N/2$. In this case, only the eigenvalues $\lambda \neq 0$ contribute, and we obtain

$$\mu_1\left(\hat{O}\right) = \frac{1}{2} = \frac{\Gamma(N)}{\Gamma(N+1)} \frac{\Gamma(N/2+1)}{\Gamma(N/2)} \tag{9}$$

$$\mu_2\left(\hat{O}\right) = \frac{1}{4}\frac{N+2}{N+1} = \frac{\Gamma(N)}{\Gamma(N+2)} \frac{\Gamma(N/2+2)}{\Gamma(N/2)}. \tag{10}$$

Despite the existence of a closed analytical formula to calculate any $\mu_t(\hat{O})$, its classical computational cost is prohibitive for large t. The reason for this bottleneck is the high number of different elements in the multinomial theorem. Additionally, the rapid growth of the $\Gamma(\cdot)$ function prevents accurate estimations of its values. Alternatively, it is possible to efficiently estimate $\mu_t(\hat{O})$ up to precision $\exp\left(\mathcal{O}(t^2/N)\right)$ by applying Stirling's approximation to lemma 2.

Corollary 1. *In the assumptions of lemma 2, we can upper- and lower-bound the statistical moments as*

$$\left(\frac{\mathrm{Tr}\,\hat{O}}{N}\right)^t \exp\left(\frac{-t^2}{2N}\right) \leqslant \mu_t\left(\hat{O}\right) \leqslant \left(\frac{\mathrm{Tr}\,\hat{O}}{N}\right)^t \exp\left(\frac{t^2}{2\bar{m}}\right) \tag{11}$$

with $\bar{m}^{-1} = \sum_{i=1} m_i^{-1}$ being an effective degeneracy of the spectrum.

A proof of this corollary can be found in appendix D.

4. Hierarchical Haar-random discriminator

This section contains our main contribution; an inductive test to discriminate if an unknown set of quantum states S resembles a t-design. The Haar-random discriminator test is described in pseudocode in appendix A. We strongly encourage the reader to consult the appendix for a comprehensive description of the method. A visual representation of the Haar-random discriminator algorithm is shown in figure 1 to help the reader understand it without the need of the algorithmic details. The starting point is a black box with two buttons to push: produce a copy of the previous state, and produce a new state. A given observable $\hat{O}$ is measured with respect to the states output by the black-box, and the compatibility with a t-design according to $\hat{O}$ is measured, up to a maximum t (top line of figure 1). The number of required different random states scales as $M \sim \exp\left(t^2\bar{m}^{-1}\right)$, with $\bar{m}$ being an effective multiplicity of the eigenvalues of $\hat{O}$. If the test is satisfactory, the observable is extended through permutations, and the procedure to increase t is repeated (middle line of figure 1). If the second test is passed, the observable is extended through measurements in other bases (bottom line of figure 1). If this last test does not yield incompatibility, we can only conclude that more tests are needed to verify t-designs. If the ensemble is Haar-random (or pseudorandom), our hierarchical test will never falsify the ensemble, for any t.

Our method is said to be hierarchical in the sense that tighter conditions for compatibility with t-designs are tested at each step, at more stringent computational resources. We give an interpretation in geometrical terms. A Haar-random distribution of states can be interpreted as the uniform distribution over a sphere (over complex numbers). A t-design requires an even distribution of states to saturate the

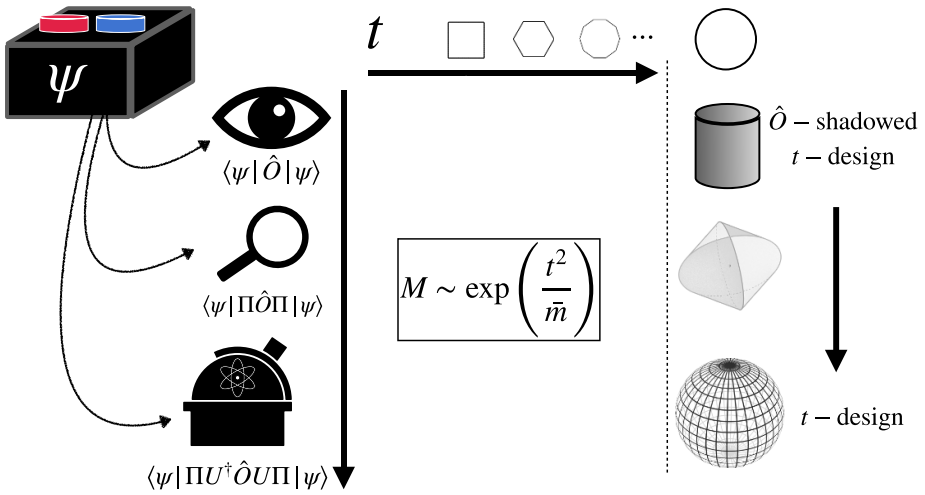

Figure 1. Schematic description of our hierarchical test. The goal is to test the randomness of a set of states generated by a black box with two buttons: *copy* and *new state*. In this picture, randomness is represented by circles and spheres. Our test is built on expectation values of an observable $\hat{O}$. The hierarchy of our test grows in two directions: t (from t-designs) and complexity of the observable. An increasing number of t—polyhedra with many edges—imposes a denser distribution of states. The sampling cost of the test increases exponentially in t. For the observables, $\hat{O}$ cannot verify t-designs due to the high dimension of quantum states. Analogously, the right projection of a cylinder from a sphere are equivalent. To compare states and t-designs, we need to construct increasingly more complex families of observables, represented by more sophisticated observation apparatuses. The algorithms and subroutines of each of the corresponding schemes are given in appendix A: t-design check, algorithm A2, Permutation check, algorithm A3, Mutually Unbiased Bases check, algorithm A4, and the combined hierarchical Haar-random discriminator, algorithm A1.

frame potential [18]. The value t captures the degree of the statistical moment to be matched, or equivalently, how fine the even distribution must be [17, 20, 21]. A direct interpretation is to match a t-design with a polygon such that the number of edges grows with t, as in figure 1. The observable $\hat{O}$ is in the end a weighted projector over the Hilbert space. We interpret that as a geometrical projector, in the right column of figure 1. A t-design is a sphere, but there exist bodies whose projections are circles, i.e. indistinguishable from projections of a sphere, such as the cylinder (circular when projected along the z axis) or the bicylinder (circular when projected along x and y axis). Partial projections, such as $\mu_t(\hat{O}, S)$ can fool us, hence the need to steadily increase the requirements to verify t-designs.

At the core of the hierarchical discriminator lies the estimation of $\mu_t(\hat{O}, S)$ from measurements of the quantum states generated from the black box. The device must be able to generate M random and independent quantum states, with C copies of each of them on demand. We do not require the C copies to be measured simultaneously, depicted in figure 1 as the red and blue push-buttons.

To estimate $\mu_t(\hat{O}, S)$, we compute the unbiased estimator

$$\overline{\mu}_t\left(\hat{O}, S\right) = \frac{1}{M}\sum_{|\psi\rangle \in S}\left(\langle\psi|\,\hat{O}\,|\psi\rangle\right)^t, \tag{12}$$

with numerical error

$$\delta\left(\overline{\mu}_t\left(\hat{O}, S\right)\right) = \frac{\hat{\sigma}_t\left(\hat{O}, S\right)}{\sqrt{M}}. \tag{13}$$

The variance $\hat{\sigma}_t(\hat{O}, S)$ can be empirically approximated as

$$\bar{\sigma}_t^2\left(\hat{O}, S\right) = \frac{1}{M-1}\sum_{|\psi\rangle_i \in S}\left(\left(\langle\psi|\,\hat{O}\,|\psi\rangle\right)^t - \overline{\mu}_t\left(\hat{O}, S\right)\right)^2 \approx \bar{\mu}_{2t}\left(\hat{O}, S\right) - \bar{\mu}_t^2\left(\hat{O}, S\right). \tag{14}$$

As we will see, the computing overhead scales exponentially with t, making procedures for large values unfeasible.

We recall that the expectation value $\langle\psi|\,\hat{O}\,|\psi\rangle$ cannot be exactly determined, hence the need for copies of the same state. The expectation value is then only accessible as an approximation of the form

$$\langle\psi|\,\hat{O}\,|\psi\rangle^t \mapsto \langle\psi|\,\hat{O}\,|\psi\rangle^t\left(1+\frac{\eta}{\sqrt{C}}\right)^t = \langle\psi|\,\hat{O}\,|\psi\rangle^t\left(1+\frac{t\eta}{\sqrt{C}}+\mathcal{O}\left(\frac{t^2}{C}\right)\right), \tag{15}$$

with η a Gaussian random variable with unit variance, and C the number of copies. The last equality is given by Taylor expansion. This implies that the dominant measurement error when estimating $\mu_t(\hat{O}, S)$ comes from the number of different states by a factor $C^{-1/2}$. In the remainder of this manuscript, we neglect the effect of the number of samples by assuming $C \gg t^2$, since typical values of t are small.

4.1. Increasing t-design

The first step for our test is the attempt to falsify the compatibility between S and a t-design measured with respect to a given observable. For S to be a Haar-random ensemble of states, it must pass the test for arbitrarily large values of t. Stricter conditions are required as t increases. Notice that $\mu_{t+1}(\hat{O}, S) = 0 \implies \mu_t(\hat{O}, S) = 0$, hence a satisfactory test at degree $t+1$ provides at least the same guarantees as a test for degree t. If S is not fully Haar-random, our test will falsify the compatibility between S and a t-design for some t.

Example 1 (Clifford circuits). As an example of an ensemble of states for which our hierarchical test works, consider S_{Stab} to be the set of uniformly distributed stabilizer states, i.e. states reached by a random Clifford circuit acting on $|0\rangle$. In this case $\mathcal{D}_t^{(\hat{O})}(S) = 0$, up to statistical fluctuations, for $t = \{1, 2, 3\}$ [22, 23]. Running the hierarchical test for any $\hat{O}$, one might be misled into thinking that S_{Stab} is Haar-randomly distributed if the test is halted in $t \leqslant 3$. However, upon inspection of $t = 4$, S_{Stab} will fail the average randomness verification test, correctly identifying S_{Stab} not being compatible with a Haar-random distribution of states.

The compatibility between S and a t-design is formally defined via $\hat{O}$-shadowed t-design.

Definition 1 ($\hat{O}$-shadowed t-design). Let S be an ensemble of states, $\hat{O}$ a quantum observable, and $\mathcal{D}_t^{(\hat{O})}(S)$ the $\hat{O}$-discriminator of this ensemble such that

$$\mathcal{D}_t^{(\hat{O})}(S) = \mu_t\left(\hat{O}, S\right) - \mu_t\left(\hat{O}\right). \tag{16}$$

Then, S is a $\hat{O}$-shadowed t-design, to precision ϵ, if

$$\left|\mathcal{D}_t^{(\hat{O})}(S)\right| \leqslant \epsilon. \tag{17}$$

To falsify compatibility, it is sufficient to obtain an estimation of $\mathcal{D}_t^{(\hat{O})}(S)$ such that with high probability

$$|\mathcal{D}_t^{(\hat{O})}(S)| > \delta\left(\mathcal{D}_t^{(\hat{O})}(S)\right), \tag{18}$$

and hence we just need to consider a sufficient number of samples to reduce the error below this threshold.

Theorem 2 (Samples to falsify $\hat{O}$-shadowed t-designs). *Let S be a set of N-dimensional states. A number of different states M suffices to falsify the compatibility of S with a $\hat{O}$-shadowed t-design, with*

$$M \geqslant \frac{\exp\left(\frac{t^2}{2\bar{m}}\right)}{\epsilon^2}\left(\frac{\text{Tr}\left(\hat{O}\right)}{N}\right)^{2t}, \tag{19}$$

with $\bar{m}^{-1} = \sum_{i=1}^{G} m_i^{-1}$, and m_i the multiplicities of the eigenvalues of $\hat{O}$.

A proof of the theorem can be found in appendix E. The scaling of M dominates the cost of the randomness verification, being $\mathcal{O}\left(\exp\left(\frac{t^2}{\bar{m}}\right)\right)$ in relative error.

In this analysis, we assume knowledge of the spectrum of $\hat{O}$ to compute the analytical values $\mu_t(\hat{O})$. One might be tempted to optimize the eigenvalue distribution to tune the obtained values of $\mu_t(\hat{O})$, and the number of Monte Carlo samples required to verify $\hat{O}$-shadowed Haar-randomness to a certain precision. Optimal spectral distributions must exist to minimize this bound. Additionally, sets of (non-commuting) observables $\{\hat{O}_1, \hat{O}_2, \cdots\}$ can be investigated. From equation (4) it is clear that product of expectation values, e.g. $\langle\psi|\hat{O}_1|\psi\rangle\langle\psi|\hat{O}_2|\psi\rangle$ must also match their distributions on Haar-random states.

The above procedure can be extended at the expense of classical post-processing via numerical estimation of high-order moments. These two questions are left for future research.

Thus far, we have been able to discern whether S is a $\hat{O}$-shadowed t-design, but this is not sufficient to unequivocally disprove that S is Haar-random. In our protocol, randomness is lost at two points: (i) The aggregation properties used in theorem 1 imply that measurements via $\hat{O}$ are only capable of distinguishing between eigenspaces, but any properties of the states within the subspace are ignored. (ii) lemma 1 implies that any measurement is sensitive to the amplitudes of the coefficients in the eigenbasis of $\hat{O}$, yet the relative phases of these quantum states are completely neglected. In the rest of the manuscript we provide means to overcome these limitations, and its computational cost.

4.2. Permutations

The first step towards improving the Haar-random discrimination capabilities is to ensure that states within the same subspace are also tested. This can be done by randomizing the spectrum of the $\hat{O}$ by applying permutations. Furthermore, this can be achieved without the need of additional quantum resources, re-using the already measured expectation values in classical post-processing. This step corresponds to going from the naked eye to a magnifying glass in figure 1, a slight increase in cost for a significant improvement in visual capacity.

Example (Relevance of permutations). Consider the family of random states

$$|\psi\rangle = \sum_{k=0}^{n} \sqrt{p_k}\,|0\rangle^{\otimes k}|1\rangle^{\otimes n-k}, \tag{20}$$

with $\boldsymbol{p} = \{p_1, \ldots, p_n\}$ sampled from a Dirichlet distribution defined by the parameters $\boldsymbol{\alpha}_k = \binom{n}{k}$, for all $k \in \{0, 1, \ldots, n\}$. To discriminate the Haar-randomness of the ensemble, let us consider the observable

$$\hat{O} = \sum_{q=1}^{n} \frac{1 - Z_q}{2}, \tag{21}$$

with Z_q being the Z-Pauli matrix acting on the qth qubit. This observable is diagonal in the computational basis, $\hat{O}|k\rangle = \lambda_k |k\rangle$, with eigenvalues λ_k being the number of 1s in the binary representation of the state. The spectral distribution of $\hat{O}$ is $\lambda = \{0, 1, \ldots n\}$, with multiplicity $\binom{n}{k}$. In this example, the set of states is specifically designed to accumulate all the relative sizes of the eigenspaces (defined by $\hat{O}$) in a one-dimensional linear subspace. One can see that such a set of quantum states is not random under the Haar-random measure, yet it appears to be when measured with respect to $\hat{O}$. By simply applying a unitary transformation to the states $\hat{O}$-shadowed is no longer Haar-random, showcasing the limitations of the protocol in section 4.

To overcome the hurdle, we can average over permutations on the spectrum of $\hat{O}$ via Monte Carlo, a process that requires only classical post-processing. We extend definition 1 to the family of observables defined as $\hat{O}_\Pi = \Pi \hat{O} \Pi$, where Π is any permutations of N elements in the eigenbasis of $\hat{O}$. We define the set S being a $(\hat{O}_\mathcal{P})$-shadowed t-design if

$$\mathcal{D}_t^{(\hat{O},\Pi)}(S) = \mathbb{E}_\Pi \left(\mu_t \left(\hat{O}_\Pi, S \right) \right) - \mu_t \left(\hat{O} \right) \tag{22}$$

$$\left| \mathcal{D}_t^{(\hat{O},\Pi)}(S) \right| < \epsilon, \tag{23}$$

and we extend theorem 2 to include permutations.

Corollary 2. *Let $\mathcal{P}$ be a set of permutations. Fix a permutation Π. For this Π we estimate $\mathcal{D}_t^{(\hat{O},\Pi)}(S)$ via Monte Carlo with M samples. We repeat the process with M_Π different permutations. We can verify that S is a $(\hat{O}_\mathcal{P})$-shadowed t-design if*

$$\left| D_t^{(\hat{O},\Pi)}(S) \right| \leq \left(\frac{\mathrm{Tr}\,\hat{O}}{N} \right)^t \exp\left(\frac{t^2}{4\bar{m}} \right) \left(\frac{1}{\sqrt{M_\Pi}} + \frac{1}{\sqrt{M}} \right). \tag{24}$$

The proof is a standard error propagation of Monte Carlo errors.

Now, we connect the arguments of the Dirichlet distribution with the permutations Π. The action of Π, as seen from the Dirichlet distribution from which $\boldsymbol{x}$ is sampled, is to permute the coefficients

$\boldsymbol{\alpha}$. Consider the case where S is a $(\Pi\hat{O}\Pi)$-shadowed t-design for some Π. Then, a subset of $\boldsymbol{\alpha}$ need not be affected by such permutation, and subsequently, the underlying Dirichlet distribution must be symmetric. We can argue that if $\mathcal{D}_t^{(\hat{O})}(S) = \mathcal{R}_t^{(\Pi\hat{O}\Pi)}(S) \approx 0$ for sufficiently many permutations Π, then exchanging any set of coefficients leads to the same random variable, and in turn, implies a fully symmetric Dirichlet distribution as in lemma 1.

One can compute $\mu_t(\hat{O}, S)$ by solving a linear system of equations with the partial statistical moments $\mathbb{E}\left(\prod_i x_i^{k_i}\right)$ as the unknown variables. Clearly, this vector has exponential size, and thus one must run over exponentially many permutations to solve the linear system. Alternatively, we can leverage our sampling strategy to statistically estimate how close a set of states is from the symmetric Dirichlet distribution, which is invariant under permutation.

Proposition 1 (Non-symmetry of Dirichlet distribution). *Let S be an ensemble whose sampling provides underlying probability distributions of the form* $\mathrm{Dir}(\boldsymbol{\alpha})$*, with unknown $\boldsymbol{\alpha}$. Then, the estimation of $\mu_t(\hat{O}_{\mathcal{P}}, S)$ can be used to estimate the average deviations between the coordinates of $\boldsymbol{\alpha}$ by*

$$\mathbb{E}_{i,j}\left(\left(\alpha_i - \alpha_j\right)^2\right) \in \mathcal{O}\left(\left(\frac{\mathrm{Tr}\left(\hat{O}^2\right)}{N^2}\right)^{-t} \mathrm{Var}_{\Pi}\left(\mu_t\left(\hat{O}_{\mathcal{P}}, S\right) - \mu_t\left(\hat{O}\right)\right)\right). \tag{25}$$

This proposition is not a formal proof, but rather an educated guess to estimate the average values. The full derivation is available in appendix F.

4.3. Mutually unbiased bases (MUB)

Even after one has sufficiently permuted over the eigenbasis of $\hat{O}$ without signaling any incompatibility with Haar randomness, it is not sufficient to guarantee that the states form a t-design. The reason for such insufficiency lies in the fact that the information retrieval of measuring $\langle\psi|\hat{O}|\psi\rangle$ is necessarily incomplete. By simple counting, we observe that permutations limit the number of accessible degrees of freedom to 2^n. Yet, full tomography of a n-qubit quantum state requires at least 2^{2n} independent measurements [24].

To illustrate the relevance of the measurement basis, we recall an example based on separations between classical and quantum computers.

Example 3 (Relevance of measurement basis). Instantaneous quantum polynomial (IQP) circuits are defined as

$$U_{\mathrm{IQP}} = H^{\otimes n} D H^{\otimes n}, \tag{26}$$

for D a non-Clifford diagonal gate, that is with T or CCZ gates. The task of sampling from these circuits cannot be simulated by any classical algorithm (under commonly accepted conjectures in computational complexity) [25], and it has been observed that its sampling has a similar underlying distribution to that of random circuits [2]. Consider again the observable equation (21), then the family of IQP states seems Haar-random. Permutations of the computational basis do not significantly modify this state, since they can be absorbed into D. However, IQP circuits cannot be Haar-random due to a lack of degrees of freedom (e.g. sampling in the X basis returns uniformly distributed bitstrings). In this case, sampling in the X basis suffices to falsify the Haar-randomness of IQP circuits.

Our next step is to show that the Haar-random hierarchical test can be extended to an informationally complete set of measurements. This additional step allows us to properly test compatibility with t-designs, going beyond its shadowed version. To achieve information completeness, we resort to complete sets of MUB [26], only requiring Clifford circuits [27]. Then, we only need to transform $\hat{O}$ through a quantum circuit and repeat the previous steps in the algorithm (see appendix A for the pseudo-code description).

Lemma 3. *Let S be a set of states of dimension N. Let $\hat{O}$ be an observable, let $\mathcal{P}$ be the set of all permutations over the eigenbasis of $\hat{O}$, and let $\mathcal{U}$ be a complete set of MUB. Then, the set of operators $\hat{O}_{\Pi,U} = \{\Pi U^\dagger \hat{O} U \Pi\}$ for $\Pi \in \mathcal{P}, U \in \mathcal{U}$ is tomographically complete.*

The operators of the form $\Pi_1\hat{O}\Pi_1 - \Pi_2\hat{O}\Pi_2$, can be used to construct observables as $\|\hat{O}\|\left(|i\rangle\langle i| - |j\rangle\langle j|\right)$, for any pair (i,j). These operators, plus the identity, suffice to obtain all relevant information in the eigenbasis of $\hat{O}$, which combined with the set of MUB, are sufficient to construct a tomographically complete measurement [26]. See appendix G for a detailed proof.

This observation motivates the use of MUB for extending $\hat{O}$-shadowed Haar-randomness. We define a set of states S to be ϵ-close to a $(\hat{O}_{\mathcal{P},\mathcal{U}})$-shadowed t-design if

$$\mathcal{D}_t^{(\hat{O},\Pi,U)}(S) = \mathbb{E}_{\Pi,U}\left(\mu_t\left(\hat{O}_{\Pi,U},S\right)\right) - \mu_t\left(\hat{O}\right) \tag{27}$$

$$\left|\mathcal{D}_t^{(\hat{O},\Pi,U)}(S)\right| < \epsilon, \tag{28}$$

where $\mathcal{U}$ is complete set of MUB and $\mathcal{P}$ is the set of all permutations in the eigenbasis of $\hat{O}$. Notice that U can absorb the permutations. We choose to keep permutations separately to lower the demands on quantum resources.

As in previous sections, we can estimate the statistical moments $\overline{\mu}_t(\hat{O}_{\Pi,U})$ via Monte Carlo to approximately obtain $\mathcal{D}_t^{(\hat{O},\Pi,U)}(S)$. This allows us to statistically verify the Haar-randomness of the set of states, irrespective of the observable $\hat{O}$.

Corollary 3. *Let S be a set of states of dimension N. Let $\hat{O}$ be a $N\times N$ Hermitian matrix, and let $\{\Pi\}$ be the set of all permutations, and let $\mathcal{U}$ be a complete set of MUB. Fix a basis $U\in\mathcal{U}$. For this U we estimate $\mathcal{R}_t^{(U^\dagger\Pi\hat{O}\Pi U)}(S)$ via Monte Carlo, with M_Π different permutations, each with M samples, and repeat the process with M_U different bases. We can verify that S is compatible with a $(\hat{O}_{\mathcal{P},\mathcal{U}})$-shadowed t-design if*

$$\left|\mathcal{D}_t^{(\hat{O},\Pi,U)}(S)\right| \leqslant \left(\frac{\mathrm{Tr}\hat{O}}{N}\right)^t \exp\left(\frac{t^2}{4\bar{m}}\right)\left(\frac{1}{\sqrt{M}}+\frac{1}{\sqrt{M_\Pi}}+\frac{1}{\sqrt{M_U}}\right) \tag{29}$$

The tomographical completeness of the measurements stated in lemma 3, implies that low values of $\mathcal{D}_t^{(\hat{O},\Pi,U)}(S)$ are only compatible with S being approximately a t-design. This step inevitably requires additional quantum resources that potentially grow exponentially with the order of the t. In our pictorial representation this step corresponds to going from a magnifying glass to a radio telescope, a simile that perfectly captures the essence of our method; gaining immense understanding of our surroundings at a very high cost.

4.4. Mixed states and frame potential

Finally, we can go beyond pure states and extend the algorithm to a set S of mixed states. In the Schmidt decomposition, ρ is given by

$$\rho = \sum_{i=1}^{N} p_i\, U|i\rangle\langle i|\, U^\dagger. \tag{30}$$

The set S is now defined by p_i and U, where both elements are drawn from two different probability distributions. Consider the case in which U is drawn from the unitary Haar distribution, then we can reabsorb any matrix diagonalizing $\hat{O}$ into U. Analogously to theorem 1, the expectation value is a random variable specified by

$$\mathrm{Tr}\left(\rho\ \hat{O}\right) = \sum_{i,j} p_i\lambda_j\left|\langle j|\, U|i\rangle\right|^2. \tag{31}$$

Random matrices are accessible by applying the QR decomposition to an $N\times N$ matrix filled with random Gaussian variables, or by sampling random N-dimensional vectors from spaces of decreasing dimensions [28]. Unfortunately, $\left|\langle j|\, U|i\rangle\right|^2$ does not admit a simple representation as in the case of pure states [29]. A detailed description of $\mathrm{Tr}\left(\rho\hat{O}\right)$ variable is out of the scope of this work and left for future research.

Another connection of the Haar-random discriminator algorithm is that of frame potentials [18, 30]. A Monte Carlo procedure sampling relative overlaps between states from an unknown set allows us to numerically estimate the frame potential, which can be compared to existing analytical bounds [31]. This procedure will yield results similar to those of this manuscript with different requirements. In the case of frame potentials one would require simultaneous copies of different states in S, as well as SWAP operations for a fidelity test. In contrast, our hierarchical discriminator algorithm only requires measurements of individual expectation values and classical post-processing. A very interesting future research question is to compare which of these two approaches yields more robust results with the same quantum resources.

5. Conclusions

We have introduced a hierarchical discriminator to falsify the compatibility between a set of states S and a $\hat{O}$-shadowed t-design, i.e. a t-designed projected on a fixed observable $\hat{O}$. The test is hierarchical in two directions: larger t values impose an exponentially-in-t number of evenly distributed states in S, as seen by a fixed $\hat{O}$. Second, the observable $\hat{O}$ can be modified via random permutations and change of bases to extend the test from $\hat{O}$-shadowed to actual t-designs in a statistical sense. The input of this method requires only the expectation values of the observable $\hat{O}$. Our first main result is to show that the expectation values of an observable measured with Haar-random states output values that can be related to a Dirichlet distribution. The second main result shows that if the spectrum of the observable is known, then the statistical moments of the expectation values of Haar-random states can be calculated exactly, albeit computationally costly. To overcome this cost, we provide an efficient method to lower- and upper-bound the Haar-random moments. We can extend these results by artificially increasing the measurements on the ensemble. Permutations in the eigenbasis of $\hat{O}$ allow us to verify the symmetry of the underlying Dirichlet distributions. MUB allow us to perform tomographically complete measurements. At the end of the procedure, if the measurements are on average compatible with a t-design, then the ensemble of states must also be (approximately) a t-design.

We foresee applications of the discriminator method presented here in multiple directions. First, it provides a rapid and efficient Haar-randomness check, potentially preventing the squandering of quantum computational resources. In addition, a test of $\hat{O}$-shadowed t-designs with respect to a single observable might be useful in the context of trainability of parameterized quantum circuits where an observable $\hat{O}$ is used to define a cost function. Our Haar-random discriminator method will indicate the presence of vanishing gradients or barren plateaus [12], or the lack thereof as hinted in section 4. Our formulation clearly exposes that t-designs are sufficient yet not necessary for them to arise. Moreover, this method can be used as an offline characterization tool for variational algorithms to explore heuristic approaches in areas such as optimization or quantum machine learning, in alignment with recent research [32]. Finally, the hierarchical discriminator protocol can also serve as a statistical alternative to tomographical verification of spherical t-designs if $\mathcal{D}_t^{(\hat{O},\Pi,U)}(S) \approx 0$, under the condition that $t \in \mathcal{O}(1)$ to be computationally affordable. One can envision the creation of a benchmark metric to characterize quantum hardware quality by how accurate it generates t-design quantum states [15].

Acknowledgments

The authors would like to thank Carlo Beenakker, Jordi Tura-Brugués and Vedran Dunjko for their support on this project, and Andrei Udriste, Stefano Polla, Patrick Emonts, Tim Coopmans and Berta Casas for useful comments. The authors would like to extend their gratitude to all members of aQa Leiden for fruitful discussions. This work was supported by the Dutch National Growth Fund (NGF), as part of the Quantum Delta NL programme.

Data availability statement

No new data were created or analysed in this study.

Appendix A. Haar-random discriminator algorithm and subroutines

We summarize our Haar-random discriminator method as pseudo-code algorithms. The first step is to measure the expectation values of an observable with known spectrum w.r.t. Haar-randomly distributed quantum states. In parallel, one computes classically the exact moments of the t-design with $\hat{O}$ following equation (8), lemma 2. The underlying assumption is then that the random states in the set S live in a Hilbert space of dimension N. The algorithm described here works for any t without extra measurements.

Algorithm 1. Haar-random Discriminator.

Procedure: HaarRandDiscriminator($\hat{O}$, S, $t_{\max}$, $\mathcal{P}, \mathcal{U}$);
Set $\mu_t(\hat{O})$ from equation (8);
Set M from equation (19);
$\overline{\mu}_t(\hat{O}, S) \leftarrow$ tDesignCheck($\hat{O}$, S, $t_{\max}$);
Compute $\mathcal{D}_t^{(\hat{O})}(S)$ from equation (16);
if $|\mathcal{D}_t^{(\hat{O})}(S)| > \delta(x)$ **then**
  Output: S is NOT COMPATIBLE with $\hat{O}$-shadowed t-design;
  STOP;
else
  $\overline{\mu}_t(\hat{O}_\Pi, S) \leftarrow$ PermutationTest($\hat{O}$, S, $t_{\max}$, $\mathcal{P}$);
  Compute $\mathcal{D}_t^{(\hat{O},\Pi)}(S)$ from equation (16);
  if $|\mathcal{D}_t^{(\hat{O},\Pi)}(S)| > \delta(x)$ **then**
    Output: S is NOT COMPATIBLE with $\hat{O}_{\mathcal{P}}$-shadowed t-design;
    STOP;
  else
    $\overline{\mu}_t(\hat{O}_{\Pi,U}, S) \leftarrow$ MUBTest($\hat{O}$, S, $t_{\max}$, $\mathcal{P}, \mathcal{U}$);
    Compute $\mathcal{D}_t^{(\hat{O},\Pi,U)}(S)$ from equation (16);
    if $|\mathcal{D}_t^{(\hat{O},\Pi,U)}(S)| > \delta(x)$ **then**
      Output: S is NOT COMPATIBLE with $\hat{O}_{\mathcal{P},\mathcal{U}}$-shadowed t-design;
      STOP;
    else
      Output: S is COMPATIBLE with a t-design;

Algorithm 2. $\hat{O}$-shadow t-design.

Procedure: tDesignCheck($\hat{O}$, S, $t_{\max}$);
Input: a quantum observable $\hat{O}$, a black-box quantum device that outputs quantum states $|\Psi_i\rangle \in S$, maximum t-moment to estimate $t_{\max}$;
Sample M quantum states from S;
for $i \in [1 \ldots M]$ **do**
  Measure $\langle \hat{O} \rangle_i = \langle \Psi_i | \hat{O} | \Psi_i \rangle$;
for $t \in [1 \ldots t_{\max}]$ **do**
  Estimate $\overline{\mu}_t(\hat{O}, S)$ from equation (12);
Output: $\overline{\mu}_t(\hat{O}, S)$

Algorithm 3. Add Permutations.

Procedure: PermutationCheck($\hat{O}$, S, $t_{\max}$, $\mathcal{P}$);
Input: a quantum observable $\hat{O}$, a black-box quantum device that outputs quantum states $|\Psi_i\rangle \in S$, maximum t-moment to estimate $t_{\max} > 1$, and a set of permutations $\mathcal{P}$;
for $\Pi \in \mathcal{P}$ **do**
  Set $\hat{O}_\Pi = \Pi \hat{O} \Pi$;
  Call tDesignTest($\hat{O}_\Pi$, S, $t_{\max}$);
Output: $\overline{\mu}_t(\hat{O}_\Pi, S)$

Algorithm 4. Add Mutually Unbiased Bases.

Procedure: MUBCheck($\hat{O}$, S, $t_{\max}$, $\mathcal{P}, \mathcal{U}$);
Input: a quantum observable $\hat{O}$, a black-box quantum device that outputs quantum states $|\Psi_i\rangle \in S$, maximum t-moment to estimate $t_{\max} > 1$, a set of permutations $\mathcal{P}$, and a complete set of Mutually Unbiased Bases $\mathcal{U}$;
for $U \in \mathcal{U}$ **do**
  Set $\hat{O}_U = U^\dagger \hat{O} U$;
  Call PermutationTest($\hat{O}_U$, S, $t_{\max}$, $\mathcal{P}$);
Output: $\overline{\mu}_t(\hat{O}_{\Pi,U}, S)$

Appendix B. Proof of lemma 1 and theorem 1

Consider a set of states $\{|\psi\rangle\} \sim \mathcal{H}$, where $\mathcal{H}$ is the Haar-random distribution over the complex projective space $\mathbb{C}P^{N-1}$, defined as

$$S^N = \left\{ |\psi\rangle \,:\, |\psi\rangle \in \mathbb{C}^N, \| \, |\psi\rangle \, \| = 1 \right\}. \tag{32}$$

We begin by defining the Dirichlet distribution.

Definition 2 (Dirichlet distribution [14]). The Dirichlet distribution $\boldsymbol{x} \sim \mathrm{Dir}(\boldsymbol{\alpha})$ parameterized by $\alpha \in \mathbb{R}^N_{>0}$ is supported on the $(N-1)$-standard simplex, i.e. $\boldsymbol{x} = (x_1, x_2, \ldots, x_N), \|\boldsymbol{x}\|_1 = 1$. It has the following probability density function with respect to the Lebesgue measure on $\mathbb{R}^{N-1}$:

$$f_{\mathrm{Dir}}(\boldsymbol{x}, \boldsymbol{\alpha}) = \frac{\Gamma(\alpha_0)}{\prod_{i=1}^N \Gamma(\alpha_i)} \prod_{i=1}^N x_i^{\alpha_i - 1}, \tag{33}$$

where $\alpha_0 = \|\boldsymbol{\alpha}\|_1$. It is possible to exactly compute the statistical moments of arbitrary order $\boldsymbol{k} = (k_1, k_2, \ldots, k_N)$,

$$\mathbb{E}_{\boldsymbol{x} \sim \mathrm{Dir}(\boldsymbol{\alpha})} \left(\prod_{i=1}^N x_i^{k_i} \right) = \frac{\Gamma(\alpha_0)}{\Gamma(\alpha_0 + k_0)} \prod_{i=1}^N \frac{\Gamma(\alpha_i + k_i)}{\Gamma(\alpha_i)}. \tag{34}$$

We now show the relationship between Haar-random states and the Dirichlet distribution. By assumption, a Haar-random ensemble $\{|\psi\rangle\}$ and $\{U|\psi\rangle\}$ are statistically equivalent, for any unitary matrix $U \in \mathcal{SU}(N)$. The projections $|\langle\psi| U |i\rangle|^2$ are now a set of random variables, subject to

$$\sum_{i=0}^{N-1} |\langle\psi| U |i\rangle|^2 = 1. \tag{35}$$

Therefore, the random variables $x_i = |\langle\psi| U |i\rangle|^2$ must follow a Dirichlet distribution. This must hold for any U. In particular, it must hold for any permutation between elements in the computational basis, and therefore the random variables x_i and x_j must be indistinguishable for any pair (i,j). This yields

$$\boldsymbol{x} \sim \mathrm{Dir}(\alpha \mathbf{1}), \tag{36}$$

where $\mathbf{1}$ is a vector with all entries equal to 1, and α is a normalization constant, and it is the only missing part. For the normalization, we can recall the known result that coordinates (squared) in a multidimensional sphere can be expressed as Dirichlet distributions with parameters $1/2$ [33]. Notice that our quantum states have N independent complex values, thus each coordinate has the double contribution. The additivity of Dirichlet distribution over its coordinates suffices to state $\alpha = 1/2 + 1/2 = 1$. This completes the proof. □

We proceed now to prove theorem 1. The sets of quantum states are Haar-random. We choose U as the matrix that diagonalizes the observable $\hat{O}$. This allows us to express the quantity $\langle\psi| \hat{O} |\psi\rangle$ in the basis of $\hat{O}$ as

$$\langle\psi| \hat{O} |\psi\rangle = \sum_{i=1}^N \lambda_i |\langle\lambda_i|\psi\rangle|^2, \tag{37}$$

where λ_i are the eigenvalues of H. Trivially, we obtain that

$$\langle\psi| \hat{O} |\psi\rangle = \sum_{i=1}^{2^n} \lambda_i x_i, \tag{38}$$

with $\boldsymbol{x} \sim \mathrm{Dir}(\mathbf{1}/2)$.

Notice that the symmetric Dirichlet distribution implies that $\langle\psi| \hat{O} |\psi\rangle$ is insensitive to any permutation of coordinates (i,j) if $\lambda_i = \lambda_j$. This can be related to the aggregation property of the Dirichlet distribution, namely

$$(x_1, x_2, \ldots, x_N) \sim \mathrm{Dir}(\alpha_1, \alpha_2, \ldots, \alpha_N) \to (x_1, x_2, \ldots, x_i + x_j, \ldots) \sim \mathrm{Dir}(\alpha_1, \alpha_2, \ldots, \alpha_i + \alpha_j, \ldots). \tag{39}$$

Therefore, we can group all coordinates with the same eigenvalue and aggregate them under a single random variable, with the corresponding parameter proportional to the multiplicity of the eigenvalue. This yields the desired result. □

Appendix C. Proof of lemma 2

The proof is as follows. In order to compute $\mu_t(\hat{O})$ we recall the multinomial theorem as

$$\mathbb{E}_{\mathrm{Dir}(\boldsymbol{\alpha})}\left(\left(\sum_{i=G}\lambda_i x_i\right)^t\right)=\sum_{\substack{\boldsymbol{k}\in\mathbb{N}^G\\ \|\boldsymbol{k}\|_1=t}}\binom{t}{\boldsymbol{k}}\left(\prod_{i=1}^{G}\lambda_i^{k_i}\right)\mathbb{E}_{\mathrm{Dir}(\boldsymbol{\alpha})}\left(\prod_{i=1}^{G}x_i^{k_i}\right), \tag{40}$$

where the multinomial coefficient is defined as

$$\binom{t}{\boldsymbol{k}}=\frac{t!}{\prod_{i=1}^{G}k_i!}. \tag{41}$$

We recall now equation (34) for computing the expectation values over products of x_i. This yields the exact formula

$$\mathbb{E}_{\mathrm{Dir}(\boldsymbol{\alpha})}\left(\left(\sum_{i=G}\lambda_i x_i\right)^t\right)=\sum_{\substack{\boldsymbol{k}\in\mathbb{N}^G\\ \|\boldsymbol{k}\|_1=t}}\binom{t}{\boldsymbol{k}}\left(\prod_{i=1}^{G}\lambda_i^{k_i}\right)\frac{\Gamma(\alpha_0)}{\Gamma(\alpha_0+t)}\prod_{i=1}^{N}\frac{\Gamma(\alpha_i+k_i)}{\Gamma(\alpha_i)}, \tag{42}$$

with $\alpha_0=\|\boldsymbol{\alpha}\|_1$. □

Appendix D. Proof of corollary 1

For the proof, we resume the proof of lemma 2. As a first approximation, we recall

$$\Gamma(x+a)\approx x^a\Gamma(x). \tag{43}$$

This simple approximation allows us to write

$$\mathbb{E}_{\mathrm{Dir}(\boldsymbol{\alpha})}\left(\left(\sum_{i=G}\lambda_i x_i\right)^t\right)\approx\frac{1}{\alpha_0^t}\sum_{\substack{\boldsymbol{k}\in\mathbb{N}^G\\ \|\boldsymbol{k}\|_1=t}}\binom{t}{\boldsymbol{k}}\prod_{i=1}^{G}(\lambda_i\alpha_i)^{k_i}. \tag{44}$$

By definition, the terms in the sum match the multinomial theorem for $\alpha_i=m_i$. This yields the easy approximation

$$\mathbb{E}_{\mathrm{Dir}(\boldsymbol{\alpha})}\left(\left(\sum_{i=G}\lambda_i x_i\right)^t\right)\approx\left(\frac{\mathrm{Tr}\left(\hat{O}\right)}{N}\right)^t. \tag{45}$$

This first approximation implies that $\mu_t(\hat{O})\approx\mu_1(\hat{O})^t$, and it is not enough to carry out faithful comparisons. To this end, we simplify the statistical moments by applying Stirling's approximation with more terms.

$$\log\Gamma(x)=x\log x-x+\mathcal{O}(\log(x)), \tag{46}$$

thus

$$\log\left(\frac{\Gamma(\alpha+t)}{\Gamma(\alpha)}\right)=(\alpha+t)\log(\alpha+t)-\alpha\log(\alpha)-t. \tag{47}$$

Then

$$\log\left(\mathbb{E}_{\mathrm{Dir}(\boldsymbol{\alpha})}\left(\prod_{i=1}^{G}x_i^{k_i}\right)\right)=\log\left(\frac{\Gamma(\alpha_0)}{\Gamma(\alpha_0+t)}\right)+\sum_{i=1}^{G}\log\left(\frac{\Gamma(\alpha_i+k_i)}{\alpha_i}\right)= \tag{48}$$

$$\alpha_0\log\alpha_0+t-(\alpha_0+t)\log(\alpha_0+t)+\mathcal{O}\log(\alpha+t)$$

$$\sum_{i=1}^{G}-\alpha_i\log\alpha_i-k_i+(\alpha_i+k_i)\log(\alpha_i+k_i)= \tag{49}$$

$$\alpha_0 \log \alpha_0 + t - (\alpha_0 + t)\log(\alpha_0 + t) - \sum_{i=1}^{G} \alpha_i \log \alpha_i - (\alpha_i + k_i)\log(\alpha_i + k_i) + \mathcal{O}\log(\alpha + t). \tag{50}$$

Therefore, neglecting the errors,

$$\mathbb{E}_{\mathrm{Dir}(\boldsymbol{\alpha})}\left(\prod_{i=1}^{G} x_i^{k_i}\right) = \frac{\alpha_0^{\alpha_0}}{(\alpha_0 + t)^{\alpha_0 + t}} \prod_{i=1}^{G} \frac{(\alpha_i + k_i)^{\alpha_i + k_i}}{\alpha_i^{\alpha_i}}. \tag{51}$$

We now pull the bounds

$$\alpha_i^{k_i} \exp(k_i) \leqslant \frac{(\alpha_i + k_i)^{\alpha_i + k_i}}{\alpha_i^{\alpha_i}} \leqslant \alpha_i^{k_i} \exp\left(k_i\left(1 + \frac{k_i}{2\alpha_i}\right)\right) \tag{52}$$

$$\alpha_0^{-t} \exp\left(-t\left(1 + \frac{t}{2\alpha_0}\right)\right) \leqslant \frac{\alpha_0^{\alpha_0}}{(\alpha_0 + t)^{\alpha_0 + t}} \leqslant \alpha_0^{-t} \exp(-t). \tag{53}$$

We now plug these expansions into equation (51), again dropping the error, to obtain

$$\exp\left(-\frac{t^2}{2\alpha}\right) \leqslant \frac{\mathbb{E}_{\mathrm{Dir}(\boldsymbol{\alpha})}\left(\prod_{i=1}^{G} x_i^{k_i}\right)}{\alpha_0^{-t} \prod_{i=1}^{G} \alpha_i^{k_i}} \leqslant \exp\left(\sum_{i=1}^{G} \frac{k_i^2}{2\alpha_i}\right). \tag{54}$$

Now, we recall equation (42) to rewrite it as

$$\mathbb{E}_{\mathrm{Dir}(\boldsymbol{\alpha})}\left(\left(\sum_{i=G} \lambda_i x_i\right)^t\right) = \sum_{\substack{\boldsymbol{k} \in \mathbb{N}^G \\ \|\boldsymbol{k}\|_1 = t}} \binom{t}{\boldsymbol{k}} \left(\prod_{i=1}^{G} \lambda_i^{k_i}\right) \alpha_0^{-t} \prod_{i=1}^{N} \alpha_i^{k_i} f(\boldsymbol{k}_i, t), \tag{55}$$

with

$$\exp\left(-\frac{-t^2}{2N}\right) \leqslant f(\boldsymbol{k}_i, t) \leqslant \exp\left(\sum_{i=1}^{G} \frac{k_i^2}{2\alpha_i}\right). \tag{56}$$

These bounds give a multiplicative factor difference with respect to the averages, which injected into the calculations for $\mu_t(\hat{O})$ yields

$$\mu_t\left(\hat{O}\right) \leqslant \left(\frac{\mathrm{Tr}\,\hat{O}}{N}\right)^t \exp\left(\frac{t^2}{2\bar{\alpha}}\right) \tag{57}$$

$$\mu_t\left(\hat{O}\right) \geqslant \left(\frac{\mathrm{Tr}\,\hat{O}}{N}\right)^t \exp\left(-\frac{t^2}{N}\right), \tag{58}$$

with $\frac{1}{\bar{\alpha}} = \sum_{i=1}^{G} \alpha^{-1}$. □

Appendix E. Proof of theorem 2

We start by considering the ensemble of states $S = \{|\phi\rangle\}$, with size M. By Monte Carlo method we obtain the estimate $\bar{\mu}_t(\hat{O}, S)$ with error

$$\delta\left(\bar{\mu}_t\left(\hat{O}, S\right)\right) = \frac{\hat{\sigma}_t\left(\hat{O}, S\right)}{\sqrt{M}}. \tag{59}$$

We compare now the estimate $\bar{\mu}_t(\hat{O}, S)$ to the Haar-random theoretical value $\mu(\hat{O})$ available in lemma 2, or corollary 1. This quantity is designed as

$$\mathcal{D}_t^{(\hat{O})}(S) = \left|\bar{\mu}_t\left(\hat{O}, S\right) - \mu_t\left(\hat{O}\right)\right|. \tag{60}$$

Assuming that Δ_t is small, we can argue that S forms a t-design w.r.t. the observable $\hat{O}$. Due to the Monte Carlo error, this claim is only true with a certain confidence, which depends on the variance

$\hat{\sigma}_t(\hat{O}, S)$, to be computed numerically. In the affirmative case of t-designs, we can estimate the variance easily as

$$\sigma_t^2\left(\hat{O}\right) = \mu_t\left(\hat{O}^2\right) - \mu_t\left(\hat{O}\right)^2 = \mu_{2t}\left(\hat{O}\right) - \mu_t\left(\hat{O}\right)^2. \tag{61}$$

This variance can be exactly computed from lemma 2. We can easily upper-bound it through corollary 1 and obtain

$$\sigma_t^2\left(\hat{O}\right) \leqslant \left(\frac{\mathrm{Tr}\left(\hat{O}\right)}{N}\right)^{2t} \exp\left(\frac{2t^2}{\bar{m}}\right), \tag{62}$$

which yields the condition

$$\delta^2\left(\mathcal{D}_t^{(\hat{O})}(S)\right) \leqslant \left(\frac{\mathrm{Tr}\left(\hat{O}\right)}{N}\right)^{2t} \frac{\exp\left(\frac{2t^2}{\bar{m}}\right)}{M}. \tag{63}$$

Therefore, we can only ensure ϵ accuracy in detecting t-designs if the number of samples scales as

$$M \geqslant \frac{\exp\left(\frac{t^2}{\bar{m}}\right)}{\epsilon^2}\left(\frac{\mathrm{Tr}\left(\hat{O}\right)}{N}\right)^{2t}. \tag{64}$$

□

Appendix F. Proof of proposition 1

To prove this result, we need to state first some auxiliary elements based on the Dirichlet distribution.

Lemma 4 (Almost symmetric Dirichlet distribution). *Let $\boldsymbol{x}$ be a Dirichlet-random variable stemming from a Dirichlet distribution with parameters*

$$\boldsymbol{\alpha} = \frac{\alpha_0}{N}(\mathbf{1} + \boldsymbol{\epsilon}), \tag{65}$$

with $\sum_{i=0}^{N} \epsilon_i = 0$. Then

$$\left|\mathbb{E}\left((\boldsymbol{\lambda}\cdot\boldsymbol{x})^t\right) - \mathbb{E}\left((\boldsymbol{\lambda}\cdot\boldsymbol{x}_0)^t\right)\right| \leqslant \left(\frac{\mathrm{Tr}\left(\hat{O}\right)}{N}\right)^t t\|\boldsymbol{\epsilon}\|_\infty, \tag{66}$$

$$\left|\mathbb{E}\left((\boldsymbol{\lambda}\cdot\boldsymbol{x})^t\right) - \mathbb{E}\left((\boldsymbol{\lambda}\cdot\boldsymbol{x}_0)^t\right)\right| \leqslant \left(\frac{\|\hat{O}\|_\infty}{N}\right)^t t\|\boldsymbol{\epsilon}\|_1, \tag{67}$$

with $\boldsymbol{x}_0$ being the Dirichlet random variable corresponding to $\boldsymbol{\epsilon} = \mathbf{0}$.

To show this, we just need to, again, apply the multinomial theorem to the expectation value, to obtain

$$\mathbb{E}_{\mathrm{Dir}(\boldsymbol{\alpha})}\left(\left(\sum_{i=N}\lambda_i x_i\right)^t\right) = \sum_{\substack{\boldsymbol{k}\in\mathbb{N}^G \\ \|\boldsymbol{k}\|_1 = t}} \binom{t}{\boldsymbol{k}}\left(\prod_{i=1}^{N}\lambda_i^{k_i}\right)\frac{\Gamma(\alpha_0)}{\Gamma(\alpha_0+t)}\prod_{i=1}^{N}\frac{\Gamma(\alpha+\epsilon_i+k_i)}{\Gamma(\alpha+\epsilon_i)}, \tag{68}$$

with $\alpha = \alpha_0/N$. We focus on the last term and find bounds, that is $\frac{\Gamma(\alpha+\epsilon+k)}{\Gamma(\alpha+\epsilon)}$.

To bound this element, we can use a Taylor expansion. For small values of ϵ,

$$\frac{\Gamma(\alpha(1+\epsilon)+k)}{\Gamma(\alpha(1+\epsilon))} = \left(\frac{\Gamma(\alpha+k)}{\Gamma(\alpha)}\right)\left(1+\epsilon(\varphi(\alpha+k)-\varphi(\alpha)) + \mathcal{O}\left(\epsilon^2\right)\right), \tag{69}$$

where $\varphi(z)$ is the logarithmic derivative of the Gamma function, that is,

$$\varphi(x) = \frac{d}{dz}\log(\Gamma(z)) = \frac{\Gamma'(z)}{\Gamma(z)}. \tag{70}$$

This function has the property that

$$\varphi(1+z) = \varphi(z) + \frac{1}{z}, \tag{71}$$

therefore

$$\varphi(\alpha+k) - \varphi(\alpha) = \sum_{j=0}^{k-1} \frac{1}{\alpha+j} \equiv \frac{C_{k,\alpha}}{\alpha}. \tag{72}$$

For $\epsilon \leqslant 0$, this constant suffices to have an upper-bound, and for $k=1$, the approximation is an equality. In the case of $\epsilon > 0$, and $k \geqslant 2$, there exists another constant satisfying a linear upper bound. We plug this result into the analytical calculation.

$$\mathbb{E}_{\mathrm{Dir}(\boldsymbol{\alpha})}\left(\left(\sum_{i=1}^{N}\lambda_i x_i\right)^t\right) = \sum_{\substack{\boldsymbol{k}\in\mathbb{N}^N \\ \|\boldsymbol{k}\|_1=t}} \binom{t}{\boldsymbol{k}} \left(\prod_{i=1}^{N}\lambda_i^{k_i}\right) \frac{\Gamma(\alpha_0)}{\Gamma(\alpha_0+t)} \prod_{i=1}^{N} \frac{\Gamma(\alpha+k_i)}{\Gamma(\alpha)} \left(1 + C_{k,\alpha}\epsilon_i + \mathcal{O}\left(\epsilon_i^2\right)\right), \tag{73}$$

and obtain, through Hölder's inequality,

$$\left|\mathbb{E}_{\mathrm{Dir}(\boldsymbol{\alpha})}\left((\boldsymbol{\lambda}\cdot\boldsymbol{x})^t\right) - \mathbb{E}_{\mathrm{Dir}(\boldsymbol{\alpha})}\left((\boldsymbol{\lambda}\cdot\boldsymbol{x}_0)^t\right)\right| \leqslant \mathbb{E}\left((\boldsymbol{\lambda}\cdot\boldsymbol{x}_0)^t\right)\left(C_{t,\alpha}\|\boldsymbol{\epsilon}\|_\infty + \mathcal{O}\left(\|\boldsymbol{\epsilon}\|_\infty^2\right)\right), \tag{74}$$

with $C_{t,\alpha} \in \mathcal{O}(t)$. Alternatively,

$$\left|\mathbb{E}_{\mathrm{Dir}(\boldsymbol{\alpha})}\left((\boldsymbol{\lambda}\cdot\boldsymbol{x})^t\right) - \mathbb{E}_{\mathrm{Dir}(\boldsymbol{\alpha})}\left((\boldsymbol{\lambda}\cdot\boldsymbol{x}_0)^t\right)\right| \leqslant \frac{\|\hat{O}\|_\infty^t}{(N\alpha)^t}\frac{\Gamma(\alpha+t)}{\Gamma(\alpha)} C_{t,\alpha}\|\boldsymbol{\epsilon}\|_1 \tag{75}$$

□

To arrive to proposition 1, we can rely on the previous results. We will be interested in computing

$$\begin{aligned}&\mathrm{Var}\left(\mathbb{E}_{\mathrm{Dir}(\boldsymbol{\alpha})}\left((\boldsymbol{\lambda}\cdot\boldsymbol{x})^t\right) - \mathbb{E}_{\mathrm{Dir}(\alpha\mathbf{1})}\left((\boldsymbol{\lambda}\cdot\boldsymbol{x})^t\right)\right) \\ &= \mathrm{Var}_\Pi\left(\sum_{\substack{\boldsymbol{k}\in\mathbb{N}^N \\ \|\boldsymbol{k}\|_1=t}} \binom{t}{\boldsymbol{k}} \left(\prod_{i=1}^{N}\lambda_i^{k_i}\right) \frac{\Gamma(\alpha_0)}{\Gamma(\alpha_0+t)} \prod_{i=1}^{N} \frac{\Gamma(\alpha+k_i)}{\Gamma(\alpha)} C_{k,\alpha}\epsilon_{\Pi(i)} + \mathcal{O}\left(\epsilon_{\Pi(i)}^2\right)\right).\end{aligned} \tag{76}$$

The variable $\epsilon_{\Pi(i)}$ encodes the differences between the different coordinates of $\boldsymbol{\epsilon}$ with respect to the average values. We consider this a random variable which is equal for all i, due to symmetry of permutations, and uncorrelated, that is

$$\mathrm{Cov}_\Pi\left(\epsilon_{\Pi(i),\Pi(j)}\right) = 0, \tag{77}$$

since for every permutation there will be another one with the opposite effect. Hence,

$$\begin{aligned}&\mathrm{Var}\left(\mathbb{E}_{\mathrm{Dir}(\boldsymbol{\alpha})}\left((\boldsymbol{\lambda}\cdot\boldsymbol{x})^t\right) - \mathbb{E}_{\mathrm{Dir}(\alpha\mathbf{1})}\left((\boldsymbol{\lambda}\cdot\boldsymbol{x})^t\right)\right) \\ &= \mathrm{Var}\left(\epsilon_i\right)\left(\sum_{\substack{\boldsymbol{k}\in\mathbb{N}^N \\ \|\boldsymbol{k}\|_1=t}} \binom{t}{\boldsymbol{k}}^2 \left(\prod_{i=1}^{N}\lambda_i^{2k_i}\right) \left(\frac{\Gamma(\alpha_0)}{\Gamma(\alpha_0+t)}\right)^2 \prod_{i=1}^{N} \left(\frac{\Gamma(\alpha+k_i)}{\Gamma(\alpha)}\right)^2 C_{k,\alpha}^2\right).\end{aligned} \tag{78}$$

We can provide lower bounds by saying $C_{k,\alpha} \geqslant 1$ for $k \geqslant 1$, and 0 otherwise. For small values of t, typical contributions are for $k=1$, and only t/N of them appear, with constant multiplicity. On the other hand, $\Gamma(\alpha_0+t)\Gamma(\alpha)^{-1} \approx (\alpha N)^t$, and $\Gamma(\alpha+1) = \alpha\Gamma(\alpha)$. On the multinomial side, any multinomial factor is smaller than its square. Again making the connection with the multinomial theorem,

$$\mathrm{Var}\left(\mathbb{E}_{\mathrm{Dir}(\boldsymbol{\alpha})}\left((\boldsymbol{\lambda}\cdot\boldsymbol{x})^t\right) - \mathbb{E}_{\mathrm{Dir}(\alpha\mathbf{1})}\left((\boldsymbol{\lambda}\cdot\boldsymbol{x})^t\right)\right) \in \tilde{\Omega}\left(\left(\frac{\mathrm{Tr}\hat{O}^2}{N^2}\right)^t\right)\mathrm{Var}\left(\epsilon\right). \tag{79}$$

The immediate interpretation is that

$$\mathbb{E}_{i,j}\left(\left(\alpha_i - \alpha_j\right)^2\right) \in \mathcal{O}\left(\left(\frac{\mathrm{Tr}\left(\hat{O}^2\right)}{N^2}\right)^{-t}\mathrm{Var}_\Pi\left(\mu_t\left(\hat{O}_{\mathcal{P}}, S\right)\right)\right). \tag{80}$$

Notice that $\mu_t(\hat{O})$ scales at first order as $(\mathrm{Tr}\hat{O}/N)^t$.

To arrive to the final result, we just need to show that the average with respect to permutations is similar to the average of the true symmetric dirichlet distribution. The following corollary supports this claim.

Corollary 4 (Average over permutation is close to true Dirichlet). *Let $\boldsymbol{\epsilon}$ and $\boldsymbol{x}$ be defined as in lemma 4. Let $\Pi \in \mathcal{P}$ be the set of all possible permutations of the elements of $\boldsymbol{\epsilon}$. Then,*

$$\mathrm{Prob}\left(\left|\mathbb{E}_{\Pi\sim P}\left(\mathbb{E}_{\mathrm{Dir}(\boldsymbol{\alpha})}\left((\boldsymbol{\lambda}\cdot\boldsymbol{x})^t\right)\right)-\mathbb{E}_{\mathrm{Dir}(\alpha)}\left((\boldsymbol{\lambda}\cdot\boldsymbol{x_0})^t\right)\right|\geq\Delta\,\mathbb{E}_{\mathrm{Dir}(\boldsymbol{\alpha})}\left((\boldsymbol{\lambda}\cdot\boldsymbol{x_0})^t\right)\right)\leqslant 2\exp\left(-\frac{2^{2N}\Delta^2}{C_{t,\alpha}\|\boldsymbol{\epsilon}\|_\infty^2}\right) \tag{81}$$

The proof is an immediate consequence of Hoeffding's inequality and lemma 4. We put an upper-bound in the value of the differences between the non-symmetric distributions of $\boldsymbol{x}$ that does not depend on the value of $\boldsymbol{\epsilon}$, or the permutation. By applying Hoeffding's inequality, and knowing that the number of possible permutations is 2^N we can infer that the average with respect to permutations is almost indistinguishable from the true symmetric Dirichlet distributions, for any expectation values, that is $\boldsymbol{\epsilon}=\mathbf{0}$. □

To our level of precision, the average over all permutations and the average of true Dirichlet is then indistinguishable, as long as the perturbation is moderate. This yields the final result. □

Appendix G. Proof of lemma 3

To obtain the full tomographic characterization of a quantum state, one needs to conduct experiments over an informationally complete set of positive-operator valued measurements (IC-POVM). Such sets of operators can perfectly distinguish between two different states, providing a sufficient number of copies. While infinitely many ways exist to construct sets of IC-POVM, some of these sets are larger than others. In particular, there exist recipes to construct tight IC-POVM, i.e. sets of operators with a minimal number of projectors. The number of elements of such tight systems scale as N^2.

An almost tight set of IC-POVM can be obtained by combining a reference basis with a set of mutually unbiased bases (MUB). Two bases U, V are MUB if

$$\left|\langle i|\,V^\dagger U|j\rangle\,\right|^2=\frac{1}{N}, \tag{82}$$

where N is the dimensionality of the system. A complete set of MUB is a set of bases that are pairwise unbiased. The number of MUB in a complete set scales as $(N+1)$. The existence of a complete set of MUB is an open problem in general [34]. However, for prime-power dimensions $N=p^n$, it is possible to construct complete sets of MUB [35, 36].

Assume now the set of observables $|i\rangle\langle i|$, for any reference basis, are available to measure. Then the set

$$M_{U,i}=U|i\rangle\langle i|\,U^\dagger, \tag{83}$$

where U runs over a complete set of MUB is a set of IC-POVMs [26].

The only step missing is constructing all measurements corresponding to a single basis through the permutations. We consider the observable $\hat{O}$ in the diagonal basis of the form

$$\hat{O}=\mathrm{diag}\left(\underbrace{0,\ldots,0}_{m_0},\underbrace{\lambda_1,\ldots,\lambda_1}_{m_1},\ldots\underbrace{\lambda_G,\ldots,\lambda_G}_{m_G},\right). \tag{84}$$

We want to construct operators of the form $\|\hat{O}\|\left(|i\rangle\langle i|-|j\rangle\langle j|\right)$, with $\|\hat{O}\|=\lambda_G$. If $|i\rangle$ correspond to the 0th eigenspace, and $|j\rangle$ corresponds to the Gth eigenspace, then

$$\|\hat{O}\|\left(|i\rangle\langle i|-|j\rangle\langle j|\right)=\hat{O}-\Pi\left(i\leftrightarrow j\right)\;\hat{O}\;\Pi\left(i\leftrightarrow j\right). \tag{85}$$

The prefactor of $\|\hat{O}\|$ is independent of the eigenstates (i,j) because the permutations allow the coefficients to be freely moved. In the case $|i\rangle, |j\rangle$ belong to eigenspaces in the middle zone of the spectrum, we just need to apply an extra permutation exchanging $|i\rangle$ with $|0\rangle$, and $|j\rangle$ with $|N-1\rangle$. The new set of operators is given by

$$M_{i,j}=|i\rangle\langle i|-|j\rangle\langle j|. \tag{86}$$

These sets of operations, together with the resolution of the identity

$$I = \sum_{i=0}^{N-1} |i\rangle \langle i| \tag{87}$$

allows us to compute any observable of the form

$$M_i = |i\rangle \langle i| . \tag{88}$$

Together with the sets of MUB [26], this leads to a set of IC-POVMs. □

ORCID iDs

Xavier Bonet-Monroig 0000-0003-4895-4180
Hao Wang 0000-0002-4933-5181
Adrián Pérez-Salinas 0000-0001-5430-6468

References

[1] Aaronson S and Chen L 2016 Complexity-theoretic foundations of quantum supremacy experiments (arXiv:1612.05903)
[2] Boixo S, Isakov S V, Smelyanskiy V N, Babbush R, Ding N, Jiang Z, Bremner M J, Martinis J M and Neven H 2018 *Nat. Phys.* **14** 595–600
[3] Aaronson S and Hung S H 2023 Certified randomness from quantum supremacy
[4] Arute F *et al* 2019 *Nature* **574** 505–10
[5] Zhong H S *et al* 2020 *Science* **370** 1460–3
[6] Ji Z, Liu Y K and Song F 2018 Pseudorandom quantum states (available at: https://eprint.iacr.org/2018/544)
[7] Chen B, Coladangelo A and Sattath O 2024 The power of a single Haar random state: constructing and separating quantum pseudorandomness (arXiv:2404.03295)
[8] Elben A, Flammia S T, Huang H Y, Kueng R, Preskill J, Vermersch B and Zoller P 2023 *Nat. Rev. Phys.* **5** 9–24
[9] Eisert J, Hangleiter D, Walk N, Roth I, Markham D, Parekh R, Chabaud U and Kashefi E 2020 *Nat. Rev. Phys.* **2** 382–90
[10] Cerezo M, Sone A, Volkoff T, Cincio L and Coles P J 2021 *Nat. Commun.* **12** 1791
[11] Holmes Z, Sharma K, Cerezo M and Coles P J 2022 *PRX Quantum* **3** 010313
[12] McClean J R, Boixo S, Smelyanskiy V N, Babbush R and Neven H 2018 *Nat. Commun.* **9** 4812
[13] Ragone M, Bakalov B N, Sauvage F, Kemper A F, Marrero C O, Larocca M and Cerezo M 2023 A unified theory of barren plateaus for deep parametrized quantum circuits (available at: www.nature.com/articles/s41467-024-49909-3)
[14] Johnson N L, Kotz S and Balakrishnan N 1994 *Continuous Univariate Distributions. 1* 2nd edn (Wiley)
[15] Harrow A W and Mehraban S 2023 *Commun. Math. Phys.* **401** 1531–626
[16] Brandão F G S L, Harrow A W and Horodecki M 2016 *Commun. Math. Phys.* **346** 397–434
[17] Hughes D and Waldron S 2021 *Linear Algebr. Appl.* **608** 84–106
[18] Gross D, Audenaert K and Eisert J 2007 *J. Math. Phys.* **48** 052104
[19] Porter C E and Thomas R G 1956 *Phys. Rev.* **104** 483–91
[20] Solé P 1991 *Eur. J. Comb.* **12** 423–31
[21] Waldron S 2003 *IEEE Trans. Inf. Theory* **49** 2307–9
[22] Kueng R and Gross D 2015 Qubit stabilizer states are complex projective 3-designs (arXiv:1510.02767)
[23] Zhu H 2017 *Phys. Rev. A* **96** 062336
[24] Banaszek K, Cramer M and Gross D 2013 *New J. Phys.* **15** 125020
[25] Bremner M J, Montanaro A and Shepherd D J 2016 *Phys. Rev. Lett.* **117** 080501
[26] Scott A J 2006 *J. Phys. A: Math. Gen.* **39** 13507–30
[27] Kern O, Ranade K S and Seyfarth U 2010 *J. Phys. A: Math. Theor.* **43** 275305
[28] Meckes E S 2019 *The Random Matrix Theory of the Classical Compact Groups* Cambridge Tracts in Mathematics (Cambridge University Press)
[29] Zyczkowski K and Kus M 1994 *J. Phys. A: Math. Gen.* **27** 4235
[30] Hunter-Jones N 2019 Unitary designs from statistical mechanics in random quantum circuits (arXiv:1905.12053)
[31] Welch L 1974 *IEEE Trans. Inf. Theory* **20** 397–9
[32] Pérez-Salinas A, Wang H and Bonet-Monroig X 2024 *npj Quantum Inf.* **10** 27
[33] Bailey R W 1992 *Am. Stat.* **46** 117–20
[34] Grassl M 2009 On SIC-POVMs and MUBs in Dimension 6 (arXiv:quant-ph/0406175)
[35] Appleby D M, Bengtsson I and Dang H B 2014 Galois Unitaries, Mutually Unbiased Bases, and MUB-balanced states (arXiv:1409.7987)
[36] Tavakoli A, Farkas M, Rosset D, Bancal J D and Kaniewski J 2021 *Sci. Adv.* **7** eabc3847