HAL open science

A game of quantum advantage: linking verification and simulation

Daniel Stilck França, Raúl García-Patrón

▶ To cite this version:

Daniel Stilck França, Raúl García-Patrón. A game of quantum advantage: linking verification and simulation. Quantum, 2022, 6, pp.753. 10.22331/q-2022-06-30-753 . hal-03967168

HAL Id: hal-03967168

https://hal.science/hal-03967168v1

Submitted on 1 Feb 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Distributed under a Creative Commons Attribution 4.0 International License

A game of quantum advantage: linking verification and simulation

Daniel Stilck França[1,2] and Raul Garcia-Patron[3]

[1]QMATH, Department of Mathematical Sciences, University of Copenhagen, Denmark
[2]Univ Lyon, ENS Lyon, UCBL, CNRS, Inria, LIP, F-69342, Lyon Cedex 07, France
[3]School of Informatics, University of Edinburgh, Edinburgh EH8 9AB, UK

We present a formalism that captures the process of proving quantum superiority to skeptics as an interactive game between two agents, supervised by a referee. The model captures most of the currently existing quantum advantage verification techniques. In this formalism, Bob samples from a distribution on a quantum device that is supposed to demonstrate a quantum advantage. The other player, the skeptical Alice, is then allowed to propose mock distributions supposed to reproduce Bob's device's statistics. Bob then needs to provide witness functions to prove that Alice's proposed mock distributions cannot properly approximate his device. Within this framework, we establish three results. First, for random quantum circuits, Bob being able to efficiently distinguish his distribution from Alice's implies efficient approximate simulation of the distribution. Secondly, finding a polynomial time function to distinguish the output of random circuits from the uniform distribution can also spoof the heavy output generation problem in polynomial time. This pinpoints that exponential resources may be unavoidable for even the most basic verification tasks in the setting of random quantum circuits. Finally, by employing strong data processing inequalities, our framework allows us to analyse the effect of noise on classical simulability and verification of more general near-term quantum advantage proposals.

1 Introduction

The transition from the reign of classical computers to quantum superiority is expected not to be a singular event but rather a process of accumulation of evidence. It will most probably happen through an iterative process of claims of proofs and refutations until a consensus is reached among the scientific community. A few years back the series of claims of the advantage of quantum annealers followed by rebuttals and an intense debate inside the quantum computation community [51, 52, 12, 56] can be seen as an example of that. Similarly, recent claims of quantum advantage [8, 57, 58, 59] were followed by a growing interest in its potential simulation by a classical device [32, 47, 55].

Ideally, one would like to demonstrate the advantage of quantum computers solving a well-established hard computational problem, such as factoring large

numbers or simulating large-sized molecules. Such demonstration will most likely need a fault-tolerant quantum computer, which will not be available in the near future. Thus, a lot of attention has been focused in the last years on quantum advantage proposals based on sampling from the outcomes of random quantum circuits, a task considered achievable. This effort culminated in landmark experiment of [8]. and its more recent followups [57, 58, 59].

The classical hardness of computing the outcome probability of random circuits has been reduced to standard complexity-theoretic assumptions in various settings [30, 14, 39, 41, 16, 3, 17, 27]. For instance, in [14] the authors prove that it is $\sharp P$ hard to compute the exact probability of the outputs of random quantum circuits for a fraction of $\frac{3}{4}+\mathrm{poly}(n)^{-1}$ of instances of random quantum circuits on n qubits. This result can be extended to devices with very low noise by assuming a couple of widely-accepted conjectures. Despite this significant progress in putting the classical hardness of sampling from a distribution close to the outputs of the random circuit on solid grounds, equivalent hardness statements are not known to hold for the levels of noise that affect current quantum computing architectures. Moreover, certifying closeness to the ideal distribution in total variation distance requires an exponential number of samples [29]. Thus, it is both not feasible to verify closeness in total variation, and the actual distance is unlikely to be in the regime of current quantum advantage experiments.

These shortcomings have shifted the interest to benchmarks of advantage that are thought to be more robust against noise and that are known to be verifiable with a feasible number of samples, albeit not computationally efficient. Prominent examples are the heavy output generation problem (XHOG) [3, 5] and the related linear cross-entropy benchmarking (linear XEB) fidelity [45, 13, 8]. The recent quantum advantage experiments used linear XEB as a benchmark for the quantum state generated by 50 to 60 qubit devices. However, these approaches have two main drawbacks. First, they require the computation of the probability of sampled strings under the circuit's ideal distribution, which consumes a running time growing exponentially with the system's size. Secondly, the number of required samples grows exponentially with the size of the system for a constant gate error probability [8]. Thus, the linear XEB verification approach requires us to be in the "sweet spot" where both the number of samples needed given the noise level and the size of the circuit are not too large to render the verification impossible. This approach is not scalable to larger system sizes with current levels of noise. Besides, it is still unknown how to reduce the hardness of the heavy output generation problem (XHOG) [3, 5, 37] to standard complexity-theoretic assumptions.

This difficulty of finding efficient certification protocols and benchmarks for random circuit sampling extends to other proposals of near-term quantum advantage, such as boson sampling [1]. This has sparked interest in simpler sanity checks, such as efficiently distinguishing the output distribution from the uniform and other "easy" distributions [2, 19, 48, 53]. As mentioned in [29], these verification forms are manifestly weaker than certifying the total variation distance and do not preclude the possibility of the device being classically simulable. However,

in the setting of random circuit sampling, not even an efficient verification test that allows us to distinguish the outputs from the uniform distribution is known. Furthermore, although efficient verification protocols for quantum computation exist [40, 20], they are likely to require fault-tolerance and are beyond what can be achieved with near-term implementations.

Independent of the recent quantum advantage experiment, the development of more efficient certification techniques of quantum advantage that can be scaled with the increasing size of the quantum computer is an area of relevance for near-term quantum computation [15]. In parallel, it is important to develop a better understanding of how noise reduces the power of quantum computers and how noise affects quantum advantage proposals or near-term applications of quantum devices.

To these ends, in this work, we envision this certification process as an interactive game between two agents, Alice that uses classical computing resources, and Bob that holds a (noisy) quantum computer and wants to convince Alice of its computational advantage. They are both supervised by the referee Robert. To win, Bob has to find functions that allow him to efficiently distinguish the output of his device from every alternative distribution Alice proposes. In turn, Alice needs to propose alternative distributions that approximate the statistics obtained by Bob's quantum computer; otherwise, she loses the game.

Central to our result is the connection between the mirror descent algorithm [54, 18] and the proposed framework of the certification game. This allows us to connect distinguishing probability distributions from a target quantum distribution and learning an approximate classical description of the latter. Furthermore, as we will see, mirror descent is particularly well-suited to learning distributions of high entropy, which is the case for NISQ devices and current quantum advantage proposals. Our framework is inspired by a recent result of the authors [23]. There, we show how to use mirror descent, strong data processing inequalities and related concepts to analyse the performance of noisy quantum devices performing optimization. In contrast, this article's main result also holds for noiseless circuits and our overarching goal is to formally link the hardness of verification of quantum advantage proposals and their approximate classical simulation.

2 Summary of results

We envision a quantum advantage demonstration as a game played between Bob, who is sampling from a (noisy) quantum circuit and claiming that it demonstrates a quantum advantage, and Alice, who is skeptical of Bob's claim and defends that her classical computer can mimic Bob's behaviour efficiently.

As Bob is the person looking to convince the others that his quantum computer has an advantage over Alice's classical device, we place the burden of the demonstration on him. In addition, Bob publicly discloses a description of the hardware and the quantum algorithm his device is implementing. The game consists of different rounds at which Alice can propose an alternative hypothesis to

the claim that Bob has achieved a quantum advantage. At the beginning of the game, they both agree on a distinguishability parameter $\epsilon > 0$, which captures how close Alice needs to match Bob's result, and confidence probability δ, which captures the probability of the outcome of the game being correct. In what follows, we denote the probability distribution Bob is sampling from by ν.

2.1 The quantum advantage game

In what follows we present a framework that captures most of the existing quantum advantage verification protocols to date as an interactive game between players.

At the beginning of the first round, Alice discloses an alternative hypothesis to quantum advantage in the form of a randomized classical algorithm sampling from a given distribution $\mu_0(x)$, for example the uniform distribution. It is then Bob's role to refute Alice's proposal and show that his distribution is at least ϵ away from Alice's in total variation distance. As we will discuss in more detail later in Section 3.1, certifying this distance constraint is equivalent to Bob providing a function $f_1 : \{0,1\}^n \to [-1,1]$ such that

$$|\mathbb{E}_{\mu_0}(f_1) - \mathbb{E}_{\nu}(f_1)| \geq \epsilon. \tag{1}$$

Alice is then allowed to update her hypothesis to $\mu_1(x)$ given the information gained from the first round of the game and sample from a distribution μ_1 that could potentially satisfy

$$|\mathbb{E}_{\mu_1}(f_1) - \mathbb{E}_{\nu}(f_1)| \leq \epsilon.$$

If so, Bob then needs to refute this mock distribution, providing a new witness f_2. Alice's new distribution then needs to pass both tests for the functions f_1 and f_2. The game continues with Bob proposing new distinguishing functions f_{t+1} and Alice mock distributions μ_t that approximate all previous expectation values up to ϵ.

The game ends if one of two players is declared defeated following a set of previously established rules. For example, Alice could concede defeat if she takes too much time to propose a new candidate or sample from it. Similarly, Bob could be forced to acknowledge his defeat if he takes too long to offer a new witness that challenges Alice. In some sense, the rules should be consistent with the process of building community consensus on the validity of a quantum advantage result.

It is important to remark that the condition in Eq. (1) must be checked by a referee Robert, to whom Alice and Bob provide samples at each round. That is, they give Robert enough samples of their distributions to estimate the expectation values by computing the empirical average for the functions f_i on the samples. The required number of samples required to be confident that the condition is satisfied up to a small failure probability can be estimated by e.g. an application of Höffding's inequality, as we show later. Note that this guarantees that Alice cannot cheat by using Bob's samples to find better distributions. For instance, if Bob's quantum computer is solving an NP problem, then knowing

the samples themselves would give Alice an efficient classical strategy. In order to suppress statistical anomalies, we also make the number of samples depend on the current round. More specifically, the number of samples for round t should be $\mathcal{O}(t\epsilon^{-2}\log(t\delta^{-1}))$. This choice ensures that the overall probability of an error occurring remains $\mathcal{O}(\delta)$.

In order for the verification game to be scalable, we may further request that sampling from Alice's distributions, evaluating Bob's test functions f_i and the size of the messages sent to Robert to be tasks that need to be achieved efficiently with the resources at hand. In what follows, we define efficient as a consumption of resources that scales polynomially with the size of the problem, i.e., the number of qubits of the quantum computer. But a more at hands definition, where we impose constraints on their size and time of computation justified by the state of the art of classical computing hardware, is also compatible with this framework and most probably be the definition used in any real experimental demonstration. Test functions that are not scalable, such as the XEB used in random quantum circuit experiments, are also contained in our game framework after some modifications. However, our no-go results do not directly apply to them as they focus on efficient functions. Indeed, the non-scalability of benchmarks like the XEB makes it impractical for future quantum computers of larger sizes than today's.

The framework presented above is quite general, capturing most of the existing quantum advantage verification proposals to date. It is natural to ask how to phrase some current quantum advantage tests and attempts to spoof them within our framework. In the examples below we illustrate how to use the formalism to describe the scenario in which we benchmark against supposedly better solutions to NP-complete optimization problems, sampling from random quantum circuits and briefly discuss the case of boson sampling machines.

2.2 Verification of NP problems

As an example, let us consider the optimization version of an NP-complete problem, such as MAXCUT on a Δ-regular graph with $\Delta \geq 3$:

Example 2.1 (MAXCUT)**.** *suppose that Bob claims his quantum computer can achieve a better value for an NP optimization problem, say MAXCUT, than Alice's classical computer. Recall that for a graph $G = (V, E)$ with n vertices and maximum degree Δ, MAXCUT of the graph can be cast as finding the maximum over $\{0,1\}^n$ of the function*

$$f_G(x) = \frac{1}{n\Delta} \sum_{(i,j)\in E} (1 - \delta(x_i, x_j)). \tag{2}$$

Here, the normalization $n\Delta$ ensures that $0 \leq f(x) \leq 1$. Thus, in this case, Bob can propose the function f_G to distinguish his distribution from Alice's classical computer. If the average value for MAXCUT he can achieve is indeed at least ϵ better than what Alice can achieve, he wins the game. On the other hand, if classical methods can yield a better cut than Bob's quantum computer, then f_G cannot claim to have achieved a quantum advantage. Note that our choice

of normalization in Eq. (2) implies that an additive error ϵ on approximating the expectation value of f_G implies a multiplicative error of order ϵ for the cut's value.

As exemplified above, for NP optimization problems, there is a clear choice for which function Bob should propose and it can be computed in polynomial time.

At first sight, the possibility of the game always requiring an exponential number of rounds seems a plausible outcome. However, there is an update rule for Alice's distribution that will lead to the game having at most $\mathcal{O}(n\epsilon^{-2})$ rounds, where n is the number of qubits of Bob's device, as we explain below. Note that we do not claim that this update rule will always lead to Alice winning, only that it will define a finite series of probability distributions that converge to the one the quantum device. The key question is whether Alice can sample from those explicit distributions efficiently or not. Interestingly, below we will show that this leads to a successful strategy against random circuits under the condition Bob provides the efficient functions f_t.

This update rule uses the connection between our certification game and mirror descent with the von Neumann entropy as potential [18], a method to learn probability distributions by approximating them by a sequence of Gibbs distributions. In a nutshell, Alice can exploit each test function f_t that Bob provides to improve her guess of the distribution ν. The updates are of the form $\mu_{t+1} \propto e^{\log(\mu_t) - \frac{\epsilon}{4} f_t}$. She can use her method of choice to sample from the Gibbs distribution, such as rejection sampling. One can then show that at every round of the game, Alice's gets closer to the ideal distribution by at least a finite amount, converging to μ_t being ϵ-close to the ideal quantum distribution ν in a finite number of rounds. In fact, regardless of the distribution Bob is sampling from, if Alice chooses to use mirror descent to update her distribution, then it follows from standard properties of mirror descent that the game will end after at most $8n\epsilon^{-2}$ rounds. We explain this in more detail in Section 4 and refer to Appendix A for a discussion and proof of its basic properties. The caveat is that the knowledge from which Gibbs distribution Alice needs to sample does not guarantee that she can do it efficiently. Let us exemplify this again with MAXCUT:

Example 2.2 (MAXCUT continued-mirror descent and simulated annealing)**.** *In the same setting as in Example 2.1, one can show that if Alice uses mirror descent to update her probability distributions, her sequence of probability distributions μ_t is given by:*

$$\mu_t(x) = \frac{e^{\frac{t}{4\epsilon} f_G(x)}}{\mathcal{Z}_t}, \quad \mathcal{Z}_t = \sum_{x \in \{0,1\}^n} e^{\frac{t}{4\epsilon} f_G(x)}.$$

That is, her strategy will be akin to performing simulated annealing to try to obtain a better value of MAXCUT. This is one of the most widely used methods for combinatorial optimization [35]. If one picks t large enough, μ_t is guaranteed to be sharply concentrated around the maximum of f_G, but the complexity of sampling from μ_t increases accordingly and at some point Alice won't be able to sample from it anymore. We refer to e.g. [26, Chapter 28] and references therein for a

discussion of the application of simulated annealing to combinatorial optimization. We can interpret the parameter $t/4\epsilon$ as the inverse temperature β. Thus, if Alice picks the mirror descent strategy, she would win if a classical Monte Carlo algorithm has a performance comparable to that of Bob's quantum computer.

2.3 Random Quantum Circuits

Random quantum circuits use a more sophisticated benchmarking strategy based on the linear cross-entropy.

Example 2.3 (linear cross-entropy, spoofing it and correlators)**.** *The current approach to benchmark quantum advantage experiments based on sampling from random quantum circuits is the linear cross-entropy [45, 8]. Given the outcome distribution of the ideal circuit ν and another distribution μ, its value is given by:*

$$\mathcal{F}_{\mathrm{XEB}}(\mu) = 2^n \mathbb{E}_\mu(\nu) - 1. \tag{3}$$

We discuss this metric thoroughly in Sec. 5, but roughly speaking the goal of this benchmark is to sample from a distribution μ that achieves $\mathcal{F}_{\mathrm{XEB}}(\mu) > 0$. Note that it corresponds to the expectation value of the function

$$f(x) = 2^n \nu(x) - 1. \tag{4}$$

In principle, the function defined in Eq. (4) does not fit our framework, as it could in principle take values larger than 1. However, as we explain in Appendix. B, under some assumptions that are believed to hold for outputs of random quantum circuit it is possible to show that the suitably cut-off function

$$f_r(x) = r^{-1}(\min\{2^n \mathbb{E}_\mu(\nu), r\} - 1) \tag{5}$$

for $r = \mathcal{O}(1)$ can be used to approximate the value Eq. (3). Furthermore, in Prop. B.2 we show that the function in Eq. (5) can also be used to distinguish the output of the circuit from the uniform distribution. Thus, we see that this commonly used benchmark also fits our framework and Bob could propose the variation of the linear cross-entropy in Eq. (5) during the first round, although it is not efficiently computable.

Recently, tensor network contraction techniques have been proposed to spoof this benchmark, which would correspond to Alice passing the first round of the game if Bob proposes the function in Eq. (5). Let us now show how our framework could be used to provide Bob with extra functions to win against the approach championed in [46]. Roughly speaking, in that paper the authors fix the outcome of k out of the n qubits to some fixed output, say $|0\rangle^{\otimes k}$, and then search for strings with higher than average probability under ν in the space of strings with those outcomes fixed. As fixing some outcomes significantly reduces the computational cost of computing outcome probabilities, the authors are then able to generate many samples which have an expectation value for the linear cross-entropy that is close to the value reported in [8]. If the authors of [46] were playing Alice with the strategy outlined in that paper, then Bob could resort to a simple strategy

to distinguish his distribution from Alice's: using correlators. Let i be one of the k qubits always set to $|0\rangle$. We then let the function for the second round be given by $f_2(x) = 1$ if $x_i = 0$ and 0 else. If Bob is sampling from the output distribution ν of a random quantum circuits, it will be the case that $\mathbb{E}_\nu(f_2) \simeq \frac{1}{2}$ up to exponentially small corrections with high probability, whereas for Alice's distribution μ_0 we have $\mathbb{E}_{\mu_0}(f_2) = 1$. Thus, this way Bob can easily distinguish his distribution from Alice's and protect himself from spoofs based on strategies like that of [46].

2.4 Gaussian Boson Sampling

In the context of Boson sampling, correlators have been championed as a benchmark of the quantum advantage certification of the devices [48]. That is, one computes some k-point correlation function of the ideal outcome distribution and compares it to the output of the device. Such tests easily fit into our framework. To see this, suppose we wish to consider a 2-point correlation function on the first two bits. In that case, we could just pick the function $f(x) = \delta_{x_1,x_2}$, which satisfies the conditions discussed before. Applying mirror descent to the case of the correlators then gives rise to a classical Gibbs state that reproduces the local correlators of the ideal distribution. Interestingly, this strategy was recently used in [55], where the authors observed that already fitting to some few-body correlators seems sufficient to obtain a better approximation in total variation to the true distribution than the noisy quantum device of [58].

2.5 Summary and discussion

Thus, we see that our framework is able to recover many of the strategies currently used in the literature to benchmark quantum advantage proposals, besides also giving advice as to how to refute spoofing strategies. Moreover, the mirror descent approach can also give rise to competitive spoofing techniques, as observed in [55].

However, it is important to notice that our framework does not cover the most general efficient procedure to distinguish probability distributions. Indeed, our framework only includes procedures that use the empirical averages of single samples to distinguish distributions. Such a setting is very close in spirit to that of statistical queries [34, 50]. However, a more general efficient procedure could apply an efficient function that depends on a polynomial number of samples to try to distinguish the distributions. In Appendix E we discuss possible extensions and limitations of our results in this direction.

2.6 Main results

As we anticipated, this framework of quantum advantage certification allows us to prove three main results on the impossibility of Bob winning using efficiently computable test functions, a connection between the HOG conjecture and the indistingushability from uniform distribution and an analysis of the effects of hardware errors on a quantum advantage verification protocol.

2.6.1 Bob can not win with efficient distinguishing functions

Our first result is that for random circuits, we are only required to play a number of rounds that scales like $\mathcal{O}(\epsilon^{-2})$. Moreover, Bob never wins if Alice plays mirror descent, $\epsilon = \Omega(\log(n)^{-1})$ and the distinguishing functions can be computed efficiently.

From a high-level perspective, Theorem 2.1 below states that if Bob can always efficiently find distinguishing functions and they can be computed efficiently, then Alice can also find and sample from a high-temperature Gibbs state that is close to the ideal distribution. Note that in this work we do not come up with such a strategy for Bob, but rather explore the consequences of the existence of such a strategy to understand the limitations and connections of verification and simulation of random quantum circuit experiments.

Theorem 2.1 (Alice approximately learns ν after ϵ^{-2} rounds for random circuits, informal version of Thm. 4.1)**.** *Let ν be the probability distribution of the outcome of a random quantum circuit on n qubits stemming from a 2^{-2n-1}-approximate two design and $\epsilon > 0$ be given. Suppose that Bob succeeds in providing functions $f_1, \ldots, f_T$ that can be computed in polynomial time and distinguish ν from a sequence $\mu_1, \ldots, \mu_T$ of distributions that Alice discloses, with T the maximal number of rounds at most $\mathcal{O}(\epsilon^{-2})$. Then there is an algorithm that allows Alice to learn a distribution μ_{T+1} exploiting the revealed information on f_t that can be sampled from in time $e^{\mathcal{O}(\epsilon^{-1})}$. Moreover, μ_T is ϵ close in total variation distance to ν.*

Note that we can efficiently sample from the output distribution as long as $\epsilon = \Omega(\log(n)^{-1})$. We will prove this result in Thm. 4.1, but it is intimately connected to the fact that the output distributions of random quantum circuits are very "flat", as the probability of the outcomes is mostly of the order 2^{-n}. For such distributions mirror descent converges very fast and we will see that they are well-approximated by high temperature Gibbs states. On the other hand, for the optimization problems like MAXCUT we expect good solvers to have outcome distributions that are highly concentrated on low energy strings. In contrast with very flat distributions, mirror descent converges slower for such concentrated distributions.

An important corollary of our theorem is that either the hardness conjectures of random quantum circuits are not valid for distances $\epsilon = \Omega(\log(n)^{-1})$ or a complete certification strategy for Bob, providing a discrimination function for every guess of Alice, is impossible with polynomial resources. Given the wide range of results that establish the hardness of sampling from random quantum circuits requiring slightly stronger assumptions than our result, we believe that our results indicate that efficient and scalable certification of random circuits in terms of empirical averages of functions is not possible.

2.6.2 Distinguishing from uniform would invalidate HOG conjecture

Our second result concerns a connection between the hardness of fooling the XHOG problem and distinguishing the output of a random circuit from the uniform distribution. We refer to Sec. 5 for a precise definition of the XHOG problem.

There we also show that if the conjectures related to the hardness of fooling the XHOG problem are true, not even distinguishing from the uniform distribution in polynomial time should be possible. Thus, although it might be possible that Bob can only efficiently distinguish during the first rounds, before mirror descent converges as in Thm 2.1, this suggests that even the first round might be difficult to win if we restrict to efficient strategies. Thus, Bob will have to resort to verification strategies that take super-polynomial time to demonstrate a quantum advantage within our game. More precisely, in Section 5 we prove the following result.

Theorem 2.2 (Distinguishing from uniform and HOG, informal version of Thm. 5.1)**.** *Let $f : \{0,1\}^n \to \{0,1\}$ be a function that for some $\epsilon > 0$ satisfies: $\mathbb{E}_\nu(f) - \mathbb{E}_\mathcal{U}(f) \geq \epsilon$, where ν is the outcome distribution of a random quantum circuit stemming from a 2^{-2n-1}-approximate two design in n qubits. Then there is an algorithm that samples from a distribution that fools XHOG up to ϵ using $\mathcal{O}(\epsilon^{-2})$ evaluations of f in expectation.*

One possible criticism of the above framework is that it might be in general hard to distinguish the outcome of any circuit stemming from a random ensemble of circuits from the uniform distribution. However, this is not true, as we show in Appendix F that in case Bob is sampling from a randomly generated stabilizer circuit, Alice can easily fool XHOG.

2.6.3 Effects of hardware errors

All the results above concern the outcome distribution of the *ideal* circuit. In Sec. 6 we extend our results to the approximate simulability of the outcome distribution of noisy devices. We show that, under doubly-stochastic noise, the number of rounds of the verification game when Alice uses mirror descent decreases exponentially with the depth of the circuit Bob is implementing. As we believe these results are interesting beyond quantum advantage proposals and apply to the broader topic of classical simulability of noisy circuits, we state them in more general terms. Below we state a specialized version of our main result regarding the complexity of approximating the statistics of outcomes of noisy circuits, Theorem 6.1:

Theorem 2.3 (Informal version of Thm. 6.1)**.** *Let ν be the outcome distribution of a noisy quantum circuit on n qubits of depth D affected by local depolarizing noise with rate p after each gate, measured in the computational basis. Given functions $f_1, \ldots, f_k : \{0,1\}^n \to [-1,1]$ and $\epsilon > 0$, mirror descent will converge to a distribution μ_T satisfying:*

$$|\mathbb{E}_\nu(f_i) - \mathbb{E}_{\mu_T}(f_i)| \leq \epsilon$$

for all $1 \leq i \leq k$ in at most $T = \mathcal{O}(\epsilon^{-2}(1-p)^{2D+2}n)$ iterations. Moreover, we can sample from μ_T by evaluating $f_1, \ldots, f_k$ at most

$$\exp\left(\frac{4(1-p)^{2D+2}n}{\epsilon}\right)$$

times.

As shown in the recent [23], which we discuss in more detail shortly, this restrains the power of noisy quantum computers to demonstrate a significant advantage versus classical methods for more structured problems. Let us exemplify this with the noisy MAXCUT example:

Example 2.4 (MAXCUT continued- noisy circuits)**.** *Let us exemplify the consequences of Theorem 2.3 to approximating MAXCUT on a noisy quantum device. Suppose that Bob's device suffers from local depolarizing noise with rate p and consists of a circuit of depth D. In this scenario, Alice will be able to obtain an expected value of MAXCUT that is ϵ close to Bob's by sampling from the distribution μ_t given by:*

$$\mu_t(x) = \frac{e^{\beta f_G(x)}}{\mathcal{Z}}, \quad \mathcal{Z} = \sum_{x\in\{0,1\}^n} e^{\beta f_G(x)}$$

with $\beta = \epsilon^{-1}(1-p)^{2D+2}n$. That is, the noise decreases the inverse temperature β we have to go when performing classical simulated annealing to obtain comparable results.

As is discussed in more detail in [23], results like the one above can be used to rigorously establish maximal depths before noisy quantum devices are not outperformed by polynomial time classical algorithms. But the main message of the example above in our context of verification is that if there are clear candidate functions to distinguish the output of the noisy quantum circuit, then the noise will make it easier for Alice to simulate the output of the device, as one would expect. We refer to Section 6 for a derivation of this bound and a more detailed discussion of its consequences.

3 Preliminaries

We will now introduce some basic definitions and notation together with the concepts of mirror descent and rejection sampling, which are relevant to our work.

3.1 Notation

Probability distributions on binary strings: we define $F = (\{0,1\}^n)^{[-1,1]}$ to be the set of functions $f : \{0,1\}^n \to [-1,1]$.

Given two probability distributions μ, ν on $\{0,1\}^n$, we define their total variation distance $\|\nu - \mu\|_{TV}$ as:

$$\|\nu - \mu\|_{TV} = \frac{1}{2} \sum_{x\in\{0,1\}^n} |\mu(x) - \nu(x)| \,.$$

Moreover, given a function $f : \{0,1\}^n \to \mathbb{R}$ we define

$$\|f\|_\infty = \sup_{x\in\{0,1\}^n} |f(x)|.$$

Distinguishability measures for quantum states and unitary designs: we are also going to need other distinguishability measures for distributions and quantum states. We will introduce them only for quantum states and note that the corresponding classical definition is obtained by considering the classical probability distribution as a diagonal quantum state. For two quantum states $\rho, \sigma \in \mathcal{M}_{2^n}$ we define their relative entropy to be:

$$S(\rho||\sigma) = \mathrm{tr}\left(\rho\left(\log\rho - \log\sigma\right)\right)$$

if kern $\rho \subset$ kern σ and $+\infty$ otherwise. Moreover, we define the $\alpha-$Rényi entropies S_α for $\alpha > 1$ to be given by:

$$S_\alpha(\rho) = -\frac{1}{\alpha - 1}\log\left(\mathrm{tr}\left(\rho^\alpha\right)\right).$$

and the von Neumann entropy to be $S_1(\rho) = S(\rho) = -\mathrm{tr}\left(\rho\log(\rho)\right)$. Note that we have:

$$n \geq S(\rho) \geq S_\alpha(\rho).$$

Let us also set our notation and terminology for random quantum circuits. Given a distribution τ on the unitary group on n qubits, $U(2^n)$, and $t \in \mathbb{N}$, we define $\mathcal{G}_\tau^{(t)} : \mathcal{M}_{2^{tn}} \to \mathcal{M}_{2^{tn}}$ to be the quantum channel:

$$\mathcal{G}_\tau^{(t)}(X) = \int\limits_{U(2^n)} U^{\otimes t} X \left(U^\dagger\right)^{\otimes t} d\tau.$$

$\mathcal{G}_\tau^{(t)}$ is then said to be an ϵ-approximate t-design [7] if

$$\|\mathcal{G}_\tau^{(t)} - \mathcal{G}_{\mu_G}^{(t)}\|_\diamond \leq \epsilon,$$

where $\|\cdot\|_\diamond$ is the diamond norm and μ_G is the Haar measure on the unitary group. Moreover, given C distributed according to τ, we will always denote by ν the probability measure we obtain by measuring $C\left|0\right\rangle$ in the computational basis, i.e.

$$\nu(x) = \mathrm{tr}\left(\left|x\right\rangle\left\langle x\right|C\left|0\right\rangle\left\langle 0\right|^{\otimes n}C^\dagger\right)$$

for $x \in \{0,1\}^n$.

3.2 Distinguishing distributions

The total variation is widely accepted as one of the most natural and operationally relevant measures of distinguishability for two probability distributions. One of the reasons for that is its dual formulation. One can show that:

$$\|\nu - \mu\|_{TV} = \frac{1}{2}\sup_{f\in F}(\mathbb{E}_\mu(f) - \mathbb{E}_\nu(f)). \tag{6}$$

Thus, it quantifies by how much the expectation values of two functions can differ on the two distributions. Moreover, defining $S = \{x \in \{0,1\}^n : \mu(x) \geq \nu(x)\}$ and letting χ_S be the indicator function of S, it is easy to see that:

$$\|\nu - \mu\|_{TV} = \frac{1}{2}(\mathbb{E}_\mu(\chi_S - \chi_{S^c}) - \mathbb{E}_\nu(\chi_S - \chi_{S^c})).$$

That is, we can restrict to differences of indicator functions in Eq. (6).

The total variation distance also has an operational interpretation in terms of distinguishability of two distributions. Indeed, consider the scenario in which with probability $\frac{1}{2}$ we are given a sample from μ and with probability $\frac{1}{2}$ we are given a sample from ν. Then once can show that the optimal probability of guessing correctly from which distribution the sample came from is given by

$$p_{\text{guess}} = \frac{1}{2}[1 + \|\mu - \nu\|_{TV}].$$

Furthermore, the optimal strategy consists of responding μ if the sample was in S and ν otherwise. Thus, we see that the total variation distance naturally allows us to quantify the distinguishability of two distributions in the one-shot setting. However, if we have access to m samples of the distribution instead of one and have to distinguish them, then the success probability is then $\|\mu^{\otimes m} - \nu^{\otimes m}\|_{TV}$.

The characterization given in Eq. (6) can also be used in yet another way to distinguish probability distributions given access to multiple samples. Suppose we have a witness function f that the total variation distance between μ and ν is at least ϵ, i.e.

$$|\mathbb{E}_\mu(f) - \mathbb{E}_\nu(f))| \geq \epsilon. \tag{7}$$

We can then use the empirical average w.r.t. to f to distinguish the distributions. To see why, given samples $X_1, \ldots, X_s$ from μ and $Y_1, \ldots, Y_s$ from ν, it follows from Hoeffding's inequality that:

$$\left| s^{-1} \sum_{i=1}^{s} f(X_i) - \mathbb{E}_\mu(f) \right| \leq \frac{\epsilon}{2}, \quad \left| s^{-1} \sum_{i=1}^{s} f(Y_i) - \mathbb{E}_\nu(f) \right| \leq \frac{\epsilon}{2} \tag{8}$$

with probability of at least $1 - \delta$ as long as

$$s = \mathcal{O}(\epsilon^{-2} \log(\delta^{-1})). \tag{9}$$

Thus, we can check if the empirical average of the samples is closer to $\mathbb{E}_\mu(f)$ or $\mathbb{E}_\nu(f)$ and use this as criterium to chose from which distribution we think the samples came from. A simple application of the triangle inequality demonstrates that this strategy will succeed with probability at least $1 - \delta$. Thus, we conclude from Eq. (9) and the discussion above that as long as $\epsilon^{-2} \log(\delta^{-1}) = \mathcal{O}(\text{poly}(n))$, polynomially many samples and evaluations of the function f are sufficient to certify that two distributions are at least at a certain distance ϵ in total variation and distinguish them. Of course, this in no sense discards the possibility that

finding the distinguishing function f itself or evaluating it may not be possible in polynomial time.

The discussion above allows us to estimate the number of samples we need to provide at each round of the game to ensure that the probability of a deviation greater than ϵ from the target is upper bounded by $1 - \mathcal{O}(\delta)$ for some δ. As at each round t we have to estimate t expectation values up to ϵ, obtaining $\mathcal{O}(\epsilon^{-2}\log(t\delta^{-t}))$ samples for each round ensures that the probability one of them deviates by more than ϵ is at most δ^t. By a union bound, the probability that there was a deviation after T rounds is at most

$$\sum_{t=1}^{T} \delta^t \leq \frac{\delta}{1-\delta} = \mathcal{O}(\delta)$$

for $\delta \leq \frac{1}{3}$. Thus, letting the number of samples per round grow like $t\log(t\delta^{-1})$ is enough to ensure that the probability of an error occurring at some point remains of order δ.

3.2.1 Discussion on generality of the model

This set of strategies to distinguish probability distributions is closely related to the statistical queries model [34, 50]. In this model to learn or distinguish distributions, one is not given access to samples from a distribution ν. Rather, one is given access to an oracle that is also specified by a precision parameter $\epsilon > 0$. When queried with a function $f \in F$, the oracle returns an estimate e_f satisfying $|e_f - \mathbb{E}_\nu(f)| \leq \epsilon$. Thus, in some sense we can say that in our game it is Bob's task to distinguish his distribution from Alice's in the statistical query model. However, as discussed in more detail in Appendix E, some of our results generalize to the case where the distinguishing functions f do not act on one sample, but rather a block of samples.

Note, however, that this is not the most general model to distinguish two probability distributions efficiently given samples. Indeed, one could consider more generally the scenario where we are given polynomially many samples of the distribution and can act on all of them simultaneously with a function that can be computed in polynomial time. Proving the impossibility of distinguishing two distributions in such a scenario is a daunting task, as discussed in more detail in Appendix E, and is out of reach of the results of this manuscript. Nevertheless, we stress that our results do apply to the strategies encountered in the literature.

3.3 Mirror descent

Mirror descent with the von Neumann entropy as potential [18, 54] is an optimization and learning algorithm to approximate probability distributions efficiently and in a structured way through a series of Gibbs probability measures. It allows us to formally connect the problem of distinguishing probability distributions and learning them. That is, given some target distribution ν on $\{0,1\}^n$ we wish to learn, say the output distribution of a given random quantum circuit, mirror descent is an iterative procedure that proposes a sequence of $\mu_0, \ldots, \mu_T$ of guesses

Algorithm 1 *Mirror descent for learning probability distributions.*

```
function MIRROR DESCENT(T, ϵ)
    Set μ_0 = U                                   ▷ initialize to the uniform distribution
    for t = 1, ..., T = ⌈8S(ν‖U)ϵ^{-2}⌉ do
        Demand function f_{t+1} such that E_{μ_t}(f_{t+1}) − E_ν(f_{t+1}) ≥ ϵ
        if Given f_t then
            Set μ_{t+1}(x) = exp(−(ϵ/4) Σ_{i=1}^{t+1} f_i(x))/Z_{t+1}.   ▷ Update the guess.
        end if
        if no such function exists then
            Return μ_t
            break loop
        end if
    end for
    Return μ_T and exit function                  ▷ Current guess is ϵ indistinguishable from ν
end function
```

for ν. Furthermore, the initial distribution μ_0 is the uniform distribution $\mathcal{U}$. The algorithm requires us to find functions $f_1, \ldots, f_T : \{0,1\}^n \to [-1,1]$ that allow us to distinguish μ_t from ν, i.e.

$$\mathbb{E}_{\mu_t}(f_{t+1}) - \mathbb{E}_{\nu}(f_{t+1}) \geq \epsilon \tag{10}$$

for some given distinguishability parameter $\epsilon > 0$. Note that we may assume without loss of generality that the equation in (10) holds without the absolute value, as if the inequality holds in the reverse direction we can just pick $-f_t$ instead. One can now appreciate the direct connection between mirror descent and our verification game. Of course, it is a priori not clear how to find such functions in a traditional mirror descent application. In the certification game, this problem is overcome by having the responsibility to provide f on Bob's side. Also note that if no function f_t exists that satisfies (10), then $\|\nu - \mu_t\|_{TV} \leq \epsilon$ by the dual formulation of the total variation distance in eq. (6).

As outlined in Algorithm 1, mirror descent works by updating the probability measure as:

$$\mu_{t+1} = \exp\left(-\frac{\epsilon}{4}\sum_{i=1}^{t+1} f_i\right)/\mathcal{Z}_{t+1}, \tag{11}$$

where

$$\mathcal{Z}_{t+1} = \sum_{x\in\{0,1\}^n} \exp\left(-\frac{\epsilon}{4}\sum_{i=1}^{t+1} f_i(x)\right)$$

is the partition function. As we update the distributions, the candidate distributions μ_t become closer and closer to the target distribution, as made precise by the following lemma:

Lemma 3.1. *The distributions μ_t of the algorithm 1 satisfy:*

$$S\left(\nu\|\mu_t\right) \leq -t\frac{\epsilon^2}{8} + S\left(\nu\|\mathcal{U}\right), \tag{12}$$

where $\mathcal{U}$ is the uniform distribution.

Eq. (12) is a standard property of mirror descent [18]. We give a simplified proof and discuss basic properties of this algorithm in Appendix A. Also note that in principle we can "recycle" distinguishing functions. That is, if Alice updates her guess a few times, it could be the case that her distribution μ_t does not satisfy

$$\mathbb{E}_{\mu_t}(f_i) - \mathbb{E}_{\nu}(f_i) \leq \epsilon \tag{13}$$

for some previously disclosed f_i. In this case, she can update in terms of f_i again until all previous expectation values also agree. This version of the algorithm is given in Algorithm 3 of Appendix A.

Exploiting the direct connection between the verification protocol and mirror descent, we can directly use lemma 3.1 to bound the number of rounds of the game in terms of $S\left(\nu\|\mathcal{U}\right)$. We then immediately obtain:

Theorem 3.1. *The output of algorithm 1 satisfies:*

$$\|\mu_t - \nu\|_{TV} \leq \epsilon \tag{14}$$

after at most $T \leq \lfloor 8\epsilon^{-2}S\left(\nu\|\mathcal{U}\right)\rfloor + 1$ iterations.

Proof. If we break the algorithm at Line (9), then, by the variational formulation of the total variation distance we have that Eq. (14) holds.

To see that this must happen after at most $\lfloor 8\epsilon^{-2}S\left(\nu\|\mathcal{U}\right)\rfloor + 1$ steps, note that by Eq. (12) we have the relation

$$0 \leq S\left(\nu\|\mu_t\right) \leq S\left(\nu\|\mathcal{U}\right) - t\frac{\epsilon^2}{8}. \tag{15}$$

Thus, a total number of iterations T that is larger than $\lfloor 8\epsilon^{-2}S\left(\nu\|\mathcal{U}\right)\rfloor + 1$ would contradict the positivity of the relative entropy. □

Note that Eq. (15) ensures that we make constant progress in relative entropy at each iteration of the algorithm. Theorem 3.1 implies that if Bob provides a sequence of functions $f_1, \ldots, f_{T+1}$ that allow distinguishing ν from the sequence $\mu_0, \ldots, \mu_T$ of at most $\mathcal{O}(\epsilon^{-2}S\left(\nu\|\mathcal{U}\right))$ distributions up to an error ϵ, then we can also find a distribution that is ϵ close to it in total variation distance. Furthermore, as we will show later, it is possible to use this connection to the relative entropy to quantify the effect of noise on the complexity of learning the distribution.

3.4 Rejection sampling

Let us now show one way how to generate samples from μ_t and the underlying complexity. We will use the standard technique of rejection sampling described in Algorithm 2.

We refer to [Appendix B.5][38] for a brief review of its properties. In rejection sampling we sample indirectly from a target distribution μ_t by first generating a sample x from an easy to sample distribution $\gamma(x)$ and accepting the sample with probability $\mu_t(x)/(M\gamma(x))$, where M is a constant such that the ratio is bounded by 1. It is a standard fact that rejection sampling will output a sample from μ_t after M runs in expectation, as the probability of rejection follows a geometric distribution with parameter M^{-1}. In the case of Gibbs distributions $\mu_t = \exp(-H_t)/\mathcal{Z}_t$, where

$$H_t(x) = \frac{\epsilon}{4}\sum_{i=1}^{t} f_i(x),$$

a common choice for γ is the uniform distribution and $M_t = \frac{2^n}{\mathcal{Z}_t}$, where $\mathcal{Z}_t$ is once again the partition function. Note that for this choice of M_t, we have that:

$$\frac{\mu_t(x)}{M_t\mathcal{U}(x)} = e^{-H_t(x)} \leq 1,$$

as we may assume without loss of generality that $H_t(x) \geq 0$ for all $x \in \{0,1\}^n$. In particular, note that with this choice, we never have to compute the partition function $\mathcal{Z}_t$ to run rejection sampling, only $H_t(x)$. Thus, we conclude that the complexity of running one round of rejection sampling is the same as computing $H_t(x)$. Let us now estimate how many rounds are required in expectation before we accept a sample:

Lemma 3.2 (Sampling from μ_t)**.** *Let μ_t be the guess at iteration t of Algorithm 1. Running rejection sampling with the uniform distribution as γ and $M_t = \frac{2^n}{\mathcal{Z}_t}$ returns a sample from μ_t after at most $e^{\frac{\epsilon}{4}t}$ trials and evaluations of H_t, in expectation.*

Algorithm 2 *Rejection sampling.*

Require: ability to generate samples from distribution γ on $\{0,1\}^n$, distribution μ_t on $\{0,1\}^n$, constant M such that $\frac{\mu_t(x)}{M\gamma(x)} \leq 1$, ability to compute $\frac{\mu_t(x)}{M\gamma(x)}$ and samples from $\mathcal{U}([0,1])$.

```
function REJECTION SAMPLING(M)
    Sample u distributed according to U([0,1]) and x distributed acoording to γ.
    if u ≤ μ_t(x)/(Mγ(x)) then
        Output x
    end if
end function
```

Proof. Note that by our previous discussion the expected number of trials is $M_t = \frac{2^n}{\mathcal{Z}_t}$. By construction, f_t are functions with image $[-1, 1]$. Thus, it follows from a triangle inequality that:

$$\|H_t\|_\infty \leq \frac{\epsilon}{4} \sum_{i=1}^{t} \|f_i\|_\infty \leq \frac{t\epsilon}{4}.$$

This implies that

$$\mathcal{Z}_t = \sum_{x \in \{0,1\}^n} \exp(-H_t(x)) \geq 2^n e^{-\frac{\epsilon}{4}t}.$$

We conclude from the last inequality that

$$M_t = \frac{2^n}{\mathcal{Z}_t} \leq e^{\frac{\epsilon}{4}t}$$

which yields the claim. □

We see that as long as $\epsilon t = \mathcal{O}(\log(n))$, then we can sample from μ_t in a polynomial expected number of trials and evaluations of H_t.

In practice, rejection sampling is not necessarily the most efficient way of simulating probability distributions and other techniques to sample from a Gibbs distribution such as Glauber dynamics or simulated annealing [38] perform better. However, rejection sampling allows for a simple analytical analysis, which is more challenging for more refined techniques.

4 Random quantum circuits

Let us now discuss the implications to the verification of quantum advantage proposals based on sampling the output distribution of random circuits. The key technical assumption behind various state-of-the-art classical hardness proofs for quantum advantage proposals is the property that the underlying ensemble is an approximate two design [28]. Thus, we will also depart from this assumption. We then have:

Lemma 4.1. *Let ν be the probability distribution of the outcome of a random quantum circuit on n qubits stemming from a 2^{-2n-1}-approximate two design. Then, with probability at least $1 - \delta$, we have:*

$$S(\nu) \geq n - [\log(1/\delta) + \log(3)].$$

Proof. We refer to Prop. C.1 in Appendix C for a proof. □

Similar statements were shown in [29, 2]. From this, we have:

Theorem 4.1 (Distinguishing output distributions and classical simulability)**.** *Let ν be the probability distribution of the outcome of a random quantum circuit on n qubits stemming from a 2^{-2n-1}-approximate two design and $\epsilon > 0$ be given.*

Suppose that we can distinguish ν from a sequence $\mu_1, \ldots, \mu_T$ of distributions that can be sampled from in polynomial time. Moreover, we can distinguish them by functions $f_1, \ldots, f_T$ that can be evaluated in polynomial time. That is:

$$\forall 1 \leq t \leq T : \mathbb{E}_{\mu_t}(f_t) - \mathbb{E}_{\nu}(f_t) \geq \epsilon. \tag{16}$$

with f_t computable in polynomial time. Then, with probability at least $1-\delta$, we can find distributions μ_t satisfying:

$$\|\nu - \mu_t\|_{TV} = \sqrt{2\left(\log(3) + \log(1/\delta) - t\frac{\epsilon^2}{8}\right)} \tag{17}$$

and sample from it in time $\mathcal{O}(e^{\frac{c}{\epsilon}}\,\mathrm{poly}(n)) = \mathcal{O}(\mathrm{poly}(n))$. In particular, if $T = \mathcal{O}(\epsilon^{-2})$, then the output distribution μ_T will also be ϵ close in total variation distance to the target.

Proof. As stated in Lemma 4.1, we have that with probability at least $1-\delta$

$$S(\nu) \geq n - [\log(3) + \log(1/\delta)]\,.$$

Conditioned on the event above, it follows from Thm. 3.1 that mirror descent outputs a distribution satisfying Eq. (17) after at most

$$8\epsilon^{-2}\left(\log(3) + \log(1/\delta)\right)$$

iterations, or equivalently after that many game rounds. Now, at each iteration t of mirror descent, we need a function satisfying Eq. (16). Moreover, we have that $\mu_t \propto \exp\left(-\epsilon/4\sum_i f_i\right)$. Thus, if all the f_i can be computed in polynomial time, then it follows from Lemma 3.2 that we can also sample from μ_t using rejection sampling in polynomial time. This is because Lemma 3.2 implies that we need at most

$$\exp\left(\frac{2(\log(3) + \log(1/\delta))}{\epsilon} + \tfrac{\epsilon}{4}\right)$$

rejection sampling rounds, in expectation. As each round of rejection sampling requires us to evaluate the functions f_i once and we suppose that they can be evaluated in polynomial time, this gives the claim. □

Therefore, if Bob provides for every proposed distribution μ_t of Alice a poly-time computable function f_t that distinguishes it from ν, after at most a constant number of rounds Alice will be sampling efficiently from an approximate distribution. It is interesting to point out that the certification game and the sampling of Alice remains efficient, even if we relax the condition of Lemma 4.1 to $S(\nu) \geq n - \mathcal{O}(\log(n))$ or request ϵ to decrease with the size n of the quantum device with scaling $\epsilon = \mathcal{O}(\log(n)^{-1})$.

A direct consequence of our result is that if the hardness conjecture of random quantum circuits is true, then Bob must fail to provide a certification function

that is efficiently computable at some stage of the certification game. A natural question would then be to ask at which stage Bob will fail to provide such a function. In the following section, we will prove that if the XHOG conjecture [5] is correct, Bob must fail at the first round of the game, i.e., even distinguishing the output distributions from the uniform distribution cannot be done in polynomial time.

We note that these results have important differences from the results in [2, Appendix 11]. There the authors show the existence of a high min-entropy distribution that can be sampled from classically and is indistinguishable from the random quantum circuit by classical circuits of polynomial size. This is because in our case we have the guarantee of a good approximation in total variation distance, i.e. the distributions are indistinguishable under any function after a couple of iterations. Another difference is that, given the distinguishing functions, our framework allows for finding the probability distribution that approximates the random circuit. Moreover, if the distinguishing functions are given and can be computed efficiently, then the outcome distribution can also be sampled from efficiently. However, to the best of our knowledge, the aforementioned result does not give an algorithm to find such an approximate distribution. Finally, our framework allows us to work with the Shannon entropy instead of the min-entropy. The min-entropy is notoriously more difficult to bound and always smaller than the Shannon entropy.

5 Distinguishing from the uniform distribution

To the best of our knowledge, the state-of-the-art approach for the verification of quantum advantage proposals based on random circuit sampling is the linear cross-entropy heavy output generation problem (XHOG) [5], which is closely related to the linear cross-entropy benchmarking (linear XEB) fidelity $\mathcal{F}_{\mathrm{XEB}}$ [45, 13, 8]. The XHOG refers to the problem of, given some circuit C, generating distinct samples $z_1, \ldots, z_k$ such that:

$$\mathbb{E}_i \left[|\langle z_i | C | 0^n \rangle|^2 \right] \geq b/2^n \tag{18}$$

for some $b > 1$ with probability at least $s = \frac{1}{2} + \frac{1}{2b}$ and k satisfying:

$$k \geq \frac{1}{((2s-1)b-1)(b-1)}. \tag{19}$$

Note that, given the samples $z_1, \ldots, z_k$, verifying that they indeed satisfy eq. (18) requires us to compute the probability of the outcomes under the ideal circuit. In turn, the linear cross-entropy fidelity for a distribution μ, $\mathcal{F}_{\mathrm{XEB}}(\mu)$, as defined in [45, 13, 8], is given by:

$$\mathcal{F}_{\mathrm{XEB}}(\mu) = 2^n \mathbb{E}_\mu (\nu) - 1, \tag{20}$$

where we interpreted the probability distribution ν as a function on $\{0,1\}^n$ that outputs the corresponding probability $\nu(x)$. The linear cross-entropy can also be

formulated as

$$\mathcal{F}_{\mathrm{XEB}}(\mu) = 2^n \mathbb{E}_\mu \left(f_\nu\right) - 1,$$

where f_ν is given by $f_\nu(x) = 2^n \nu(x)$. Although such a function does not immediately fit our framework, as it may take values higher than 1, in Appendix B we show how to approximate it by a bounded function. The underlying intuition is that as ν is very flat for random quantum circuits, for very few inputs the function f will take values that are not of constant order. Thus, as long as the measure μ is not too concentrated on strings of high value, it is possible to cut-off the function f without significantly changing the expectation value.

A simple manipulation then shows that samples z_i from μ satisfy

$$\mathbb{E}_i \left[|\langle z_i | C | 0^n \rangle|^2 \right] \geq \frac{1 + \mathcal{F}_{\mathrm{XEB}}(\mu)}{2^n}.$$

In [5], the authors relate the complexity of computing outcome probabilities of random quantum circuits to the XHOG problem. More precisely, the authors start by assuming that there is no polynomial-time classical algorithm that takes as input a (random) quantum circuit C and produces an estimate p of $p_0 = \mathbb{P}[C \,\text{outputs}\, 0]$ such that

$$\mathbb{E}\left[\left(p_0 - 2^{-n}\right)^2\right] = \mathbb{E}\left[\left(p_0 - p\right)^2\right] + \Omega\left(2^{-3n}\right).$$

where the expectations are taken over circuits C as well as the algorithm's internal randomness. They then show that this conjecture implies that there is no polynomial time algorithm that solves the XHOG problem. As the verification of XHOG requires us to estimate the outcome probabilities, this path to proving and verifying the hardness of the sampling task also implies that XHOG is not verifiable in polynomial time.

But the hardness of XHOG imposes barriers to even more basic verification tasks. As we will see now, the hardness of XHOG would imply that it is impossible to efficiently distinguish the output from the circuit from the uniform distributions. In turn, efficient distinguishability implies a polynomial time algorithm to spoof XHOG. Therefore, Bob will have to resort to verification strategies that take superpolynomial time to demonstrate a quantum advantage in our game. It should not be surprising that fooling XHOG is related to distinguishing from the uniform distribution, as both tasks require us to identify higher-than-average probability strings. What is particular to the case of random circuits is the fact that distinguishing implies we can also sample from a distribution that fools XHOG. The proof of the statement will once again rely on the fact that the output distribution is essentially flat. We will start by showing that any distribution with large 2−Rényi entropy cannot have small sets of large mass in the following sense:

Lemma 5.1. *Let ν be a probability distribution on $\{0,1\}^n$ such that $S_2(\nu) \geq n - \log(c)$ for some constant $c > 0$. Then, for any $\epsilon > 0$ and subset $L \subset \{0,1\}^n$*

we have that

$$\nu(L) = \sum_{x \in L} \nu(x) \geq \epsilon$$

implies that

$$|L| \geq \epsilon^2 c^{-1} 2^n.$$

Proof. Note that the condition $S_2(\nu) \geq n - \log(c)$ is equivalent to

$$\sum_{x \in \{0,1\}^n} \nu(x)^2 \leq c 2^{-n}. \tag{21}$$

Moreover, we have

$$\sum_{x \in \{0,1\}^n} \nu(x)^2 \geq \sum_{x \in L} \nu(x)^2 \geq \frac{\epsilon^2}{|L|}. \tag{22}$$

To see the last inequality, note that, by the concavity of the function $x \mapsto x^2$,

$$\frac{1}{|L|} \sum_{x \in L} \nu(x)^2 \geq \left(\frac{1}{|L|} \sum_{x \in L} \nu(x) \right)^2 \geq \frac{\epsilon^2}{|L|^2}.$$

Combining (22) with (21) we conclude that:

$$\frac{\epsilon^2}{|L|} \leq c 2^{-n},$$

which yields the claim after rearranging the terms. □

It then immediately follows that:

Lemma 5.2. *Let $f : \{0,1\}^n \to \{0,1\}$ be a function that for some $\epsilon > 0$ satisfies:*

$$\mathbb{E}_\nu(f) - \mathbb{E}_{\mathcal{U}}(f) \geq \epsilon, \tag{23}$$

where ν is the outcome distribution of a random quantum circuit stemming from a 2^{-2n-1}-approximate two design on n qubits. Let

$$L = \{x \in \{0,1\}^n : f(x) = 1\}.$$

Then, with probability at least $1 - \delta$,

$$|L| = \Omega\left(\epsilon^2 \delta 2^n\right). \tag{24}$$

and

$$\frac{1}{|L|} \nu(L) \geq \frac{1}{2^n} + \frac{\epsilon}{|L|} \geq \frac{1+\epsilon}{2^n}. \tag{25}$$

Proof. First note that Eq. (23) is equivalent to

$$\nu(L) \geq \frac{|L|}{2^n} + \epsilon$$

and, in particular, $\nu(L) \geq \epsilon$. Moreover, Eq. (25) readily follows by dividing the equation above by $|L|$. As we show in Prop. C.1 in Eq. (47) that with probability at least $1 - \delta$

$$S_2(\nu) \geq n - \log(3) + \log(\delta). \tag{26}$$

Conditioned on Eq. (26), it follows from Lemma 5.1 that

$$|L| \geq c\epsilon^2 \delta 2^n \tag{27}$$

for some constant $c > 0$, which yields Eq. (24). □

We restricted the result above to functions with binary outputs to simplify the arguments, but we show in Appendix D that this can be done without loss of generality. That is, given a function f that distinguished the distributions for ϵ and range $[-1, 1]$, there always exists some f' that is binary and has the same properties and distinguishes the distribution up to $\frac{\epsilon^2}{17}$.

If the function f in Lemma 5.2 can computed in polynomial time, then we can use it to fool XHOG in polynomial time:

Theorem 5.1 (From distinguishing to fooling XHOG)**.** *Let f as in Lemma 5.2 for some $\epsilon > 0$. Moreover, let $\mathcal{U}(L)$ be the uniform distribution on L. Then we can sample from $\mathcal{U}(L)$ by evaluating f a total of $\mathcal{O}(\epsilon^{-2})$ many times, in expectation. Moreover, samples from $\mathcal{U}(L)$ violate HOG up to ϵ.*

Proof. Let us start with the statement that samples from $\mathcal{U}(L)$ violate XHOG up to at least ϵ. To see this, note that:

$$\mathbb{E}_\nu\left(\mathcal{U}(L)\right) = \sum_{x \in L} \frac{\nu(x)}{|L|} \geq \frac{1+\epsilon}{2^n}$$

by Eq. (25), where with some abuse of notation we see $\mathcal{U}(L)$ as a function that outputs the probability of x under $\mathcal{U}(L)$ given x. To sample from $\mathcal{U}(L)$, we can once again resort to a variation of rejection sampling. We sample a point $x_1 \in \{0,1\}^n$ from the uniform distribution and compute $f(x_1)$. If $f(x_1) = 1$, we output x_1. If not, we rerun this procedure with a new sample x_2. It is easy to see that when we accept, x_i is uniformly distributed on L. Moreover, the probability of accepting is $\frac{L}{2^n}$. By Eq. (27), it then follows that the probability of accepting is at least $\Omega(\epsilon^2)$, from which we obtain that the expected number of rounds is $\mathcal{O}(\epsilon^{-2})$. As for each round we have to evaluate f once, the claim follows. □

It follows that if we can efficiently distinguish the output from the uniform distribution, then we can fool XHOG. In particular, it would follow from the conjectures of [5] that it is not possible to distinguish the output of random

circuits from the uniform distribution for ϵ at least inverse polynomial in n in polynomial time if XHOG cannot be solved in polynomial time.

Note that the results of this section only required the property that the underlying random circuit ensemble is a two design. However, it is well-known that random Clifford unitaries also are two designs and can still be simulated efficiently classically. It is then natural to ask if finding a function distinguishing the output of a random Clifford from the uniform distribution can be found in polynomial time. As we show in Appendix F, for random Cliffords, the function can be found and be computed in polynomial time.

6 Noisy devices

Most of the current implementations of quantum circuits have high levels of noise. In principle, highly correlated noise can make simulating the quantum device classically even more complex [36]. However, in other contexts, as demonstrated for boson sampling [49], noise can render the simulation significantly less complex. Here, we show that in the scenario of doubly stochastic Markovian noise, i.e., quantum channels that map the maximally mixed state to itself, the noise diminishes the complexity of approximating the underlying distribution being sampled from. To see what we mean, let us go back the game described in Section 2. There we considered Bob to have a quantum advantage and win the game if he could find functions $f_1, \ldots, f_t$ whose expectation values under his distribution Alice cannot reproduce classically. However, suppose now that Bob's device is affected by noise, say a global depolarizing channel with parameter $0 \leq p \leq 1$. One would then expect that as $p \to 1$, it should be easier for Alice to reproduce the statistics of Bob's device and win the game. As we show below, if Alice once again uses mirror descent to find her candidate distributions, then the algorithm will converge faster to a distribution that reproduces the statistics of Bob's distribution as the level of noise increases. For example, for MAXCUT, this will also make it easier for her to sample from the distribution mirror descent proposes, as a smaller number of iterations implies not having to go down to lower temperatures.

6.1 Strong data processing inequalities and mirror descent

Let us first recall the following definition:

Definition 6.1 (Strong data processing inequality)**.** *A doubly stochastic quantum channel $T : \mathcal{M}_d \to \mathcal{M}_d$ (i.e. a CPTP map satisfying $T(I) = T^*(I) = I$) is said to satisfy a strong data processing inequality with contraction $\alpha > 0$ w.r.t. $\frac{I}{d}$ if for all states ρ we have:*

$$S\left(T(\rho)\|\frac{I}{d}\right) \leq (1-\alpha)S\left(\rho\|\frac{I}{d}\right).$$

Strong data processing inequalities for doubly stochastic can be derived using the framework of hypercontractivity and logarithmic Sobolev inequalities [33, 43,

44, 31, 9] and are known explicitly in some cases. See also [11] for another approach. As shown by other works [6, 10, 42, 23], strong data processing can be used to quantify how useful a noisy quantum device is and for how long it can sustain interesting computations.

Let us illustrate the strong data processing inequality with an example. Let $\Phi_p : \mathcal{M}_2 \to \mathcal{M}_2$ be the depolarizing channel on one qubit with depolarizing parameter p, i.e.

$$\Phi_p(\rho) = (1-p)\rho + p\frac{I}{2}.$$

The authors of [43] show that:

$$S\left(\Phi_p^{\otimes n}(\rho)\|\frac{I}{2^n}\right) \leq \left(1-2p+p^2\right) S\left(\rho\|\frac{I}{2^n}\right)$$

and similar results are available for other relevant noise models. In particular, optimal inequalities have been derived in [33, 44] for tensor products of single qubit, doubly stochastic channels.

Suppose now that the noisy circuit of interest consists of n qubits initialized to $|0\rangle^{\otimes n}$, D layers of unitaries $U_1, U_2, \ldots, U_D$ and measurement in the computational basis. However, due to imperfections in the implementation, the state initialization, the measurements and the unitaries are noisy. We will model this by assuming that every layer of unitaries is proceeded by a layer of doubly stochastic quantum channel Φ that satisfies a strong data processing inequality α. Moreover, we will assume that the measurement is also affected by an extra noisy channel Φ. We only make these assumptions to simplify the notation and argument, but it is easy to adapt the argument to different channels at different times and different noise rates. Under the assumptions above, the probability distribution ν describing the outcomes of the noisy device is given by:

$$\nu(i) = \operatorname{tr}\left(|i\rangle\langle i| T(|0\rangle\langle 0|^{\otimes n})\right) = \operatorname{tr}\left(|i\rangle\langle i| (M \circ \Phi \circ \mathcal{U}_D \circ \Phi \circ \ldots \circ \mathcal{U}_1 \circ \Phi(|0\rangle\langle 0|^{\otimes n})\right), \tag{28}$$

where $\mathcal{U}_i$ is the channel given by conjugations with U_i and M is the q.c. channel

$$M(\rho) = \sum_{i=0}^{2^n-1} \operatorname{tr}\left(\rho|i\rangle\langle i|\right)|i\rangle\langle i|$$

and

$$T(|0\rangle\langle 0|^{\otimes n}) = (M \circ \Phi \circ \mathcal{U}_D \circ \Phi \circ \ldots \circ \mathcal{U}_1 \circ \Phi(|0\rangle\langle 0|^{\otimes n}).$$

We then have that if we want to approximate the values of k functions, then the number of iterations of mirror descent T required to achieve that decreases exponentially with the noise level. More precisely:

Theorem 6.1. *Let ν be the distribution defined in Eq. (28), $\epsilon > 0$ and assume Φ satisfies a strong data processing inequality with parameter α. Given functions*

$f_1, \ldots, f_k : \{0,1\}^n \to [-1,1]$, mirror descent will converge to a distribution μ_t satisfying:

$$|\mathbb{E}_\nu(f_i) - \mathbb{E}_{\mu_t}(f_i)| \leq \epsilon$$

for all $1 \leq i \leq k$ in at most $T = \mathcal{O}(\epsilon^{-2}(1-\alpha)^{D+1}n)$ iterations. Moreover, we can sample from μ_t by evaluating $f_1, \ldots, f_k$ at most

$$\exp\left(\frac{4(1-\alpha)^{D+1}n}{\epsilon}\right)$$

times.

Proof. Mirror descent will converge to a distribution with the desired properties after $8\epsilon^{-2}S(\nu\|\mathcal{U})$ iterations, see Prop. A.1 for a proof. Moreover, by Lemma 3.2, the complexity of sampling from μ is bounded by

$$\exp\left(\frac{4S(\nu\|\mathcal{U})}{\epsilon}\right)$$

evaluations of the functions f_i. Thus, the statement follows if we can bound the relative entropy of the outcome. Note that:

$$S(\nu\|\mathcal{U}) = S\left(T(|0\rangle\langle 0|^{\otimes n})\|\frac{I}{2^n}\right),$$

as the maximally mixed state gives rise to the uniform distribution when measured in the computational basis. By the data processing inequality:

$$S\left(T(|0\rangle\langle 0|^{\otimes n})\|\frac{I}{2^n}\right) \leq S\left(\Phi \circ \mathcal{U}_D \circ \Phi \circ \ldots \circ \mathcal{U}_1 \circ \Phi(|0\rangle\langle 0|^{\otimes n})\|\frac{I}{2^n}\right)$$

and by our assumption on Φ:

$$S\left(\Phi \circ \mathcal{U}_D \circ \Phi \circ \ldots \circ \mathcal{U}_1 \circ \Phi(|0\rangle\langle 0|^{\otimes n})\|\frac{I}{2^n}\right) \leq (1-\alpha)S\left(\mathcal{U}_D \circ \Phi \circ \ldots \circ \mathcal{U}_1 \circ \Phi(|0\rangle\langle 0|^{\otimes n})\|\frac{I}{2^n}\right).$$

The relative entropy is unitarily invariant, thus:

$$(1-\alpha)S\left(\mathcal{U}_D \circ \Phi \circ \ldots \circ \mathcal{U}_1 \circ \Phi(|0\rangle\langle 0|^{\otimes n})\|\frac{I}{2^n}\right) = (1-\alpha)S\left(\Phi \circ \ldots \circ \mathcal{U}_1 \circ \Phi(|0\rangle\langle 0|^{\otimes n})\|\frac{I}{2^n}\right).$$

Applying the chain of arguments above another D times we conclude that:

$$S\left(T(|0\rangle\langle 0|^{\otimes n})\|\frac{I}{2^n}\right) \leq (1-\alpha)^{D+1}S\left(|0\rangle\langle 0|^{\otimes n}\|\frac{I}{2^n}\right) = (1-\alpha)^{D+1}n,$$

which yields the claim. □

Thus, we see that the complexity of approximate sampling from the output of noisy circuits decreases exponentially with the noise level and depth. For

instance, for circuits with local depolarizing noise with parameter p and depth D, the theorem above gives a complexity of:

$$\exp\left(\frac{4(1-p)^{2D+2}n}{\epsilon}\right). \tag{29}$$

evaluations of the distinguishing functions. Thus, we see that our framework has the desirable feature that it becomes easier for Alice to mock Bob's device as the noise increases. Unfortunately, the scaling with n in the bound above is undesirable for high-entropy distributions. Thus, we plan to derive more specialized bounds in upcoming work.

7 Acknowledgments

DSF was supported by VILLUM FONDEN via the QMATH Centre of Excellence under Grant No. 10059. RGP was supported by the Quantum Computing and Simulation Hub, an EPSRC-funded project, part of the UK National Quantum Technologies Programme. We thank Anthony Leverrier and Juani Bermejo-Vega for helpful comments and discussions.

References

[1] Scott Aaronson and Alex Arkhipov. The computational complexity of linear optics. In *Research in Optical Sciences*. OSA, 2014. DOI: 10.1364/qim.2014.qth1a.2.

[2] Scott Aaronson and Alex Arkhipov. Boson sampling is far from uniform. *Quantum Info. Comput.*, 14(15–16):1383–1423, November 2014. ISSN 1533-7146. DOI: https://doi.org/10.26421/qic14.15-16-7.

[3] Scott Aaronson and Lijie Chen. Complexity-theoretic foundations of quantum supremacy experiments. In *Proceedings of the 32nd Computational Complexity Conference*, 2017. ISBN 9783959770408. DOI: https://doi.org/10.48550/arXiv.1612.05903.

[4] Scott Aaronson and Daniel Gottesman. Improved simulation of stabilizer circuits. *Physical Review A*, 70(5), Nov 2004. ISSN 1094-1622. DOI: 10.1103/physreva.70.052328.

[5] Scott Aaronson and Sam Gunn. On the classical hardness of spoofing linear cross-entropy benchmarking. *Theory of Computing*, 16(11):1–8, 2020. DOI: 10.4086/toc.2020.v016a011.

[6] Dorit Aharonov, Michael Ben-Or, Russell Impagliazzo, and Noam Nisan. Limitations of noisy reversible computation. *arXiv preprint quant-ph/9611028*, 1996.

[7] Andris Ambainis and Joseph Emerson. Quantum t-designs: t-wise independence in the quantum world. In *Twenty-Second Annual IEEE Conference on Computational Complexity 07)*. IEEE, jun 2007. DOI: 10.1109/ccc.2007.26.

[8] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C Bardin, Rami Barends, Rupak Biswas, Sergio Boixo, Fernando G S L Brandao, David A Buell, Brian Burkett, Yu Chen, Zijun Chen, Ben Chiaro, Roberto Collins, William Courtney, Andrew Dunsworth, Edward Farhi, Brooks Foxen, Austin Fowler, Craig Gidney, Marissa Giustina, Rob Graff, Keith Guerin, Steve Habegger, Matthew P Harrigan, Michael J Hartmann, Alan Ho, Markus Hoffmann, Trent Huang, Travis S Humble, Sergei V Isakov, Evan Jeffrey, Zhang Jiang, Dvir Kafri, Kostyantyn Kechedzhi, Julian Kelly, Paul V Klimov, Sergey Knysh, Alexander Korotkov, Fedor Kostritsa, David Landhuis, Mike Lindmark, Erik Lucero, Dmitry Lyakh, Salvatore Mandrà, Jarrod R McClean, Matthew McEwen, Anthony Megrant, Xiao Mi, Kristel Michielsen, Masoud Mohseni, Josh Mutus, Ofer Naaman, Matthew Neeley, Charles Neill, Murphy Yuezhen Niu, Eric Ostby, Andre Petukhov, John C Platt, Chris Quintana, Eleanor G Rieffel, Pedram Roushan, Nicholas C Rubin, Daniel Sank, Kevin J Satzinger, Vadim Smelyanskiy, Kevin J Sung, Matthew D Trevithick, Amit Vainsencher, Benjamin Villalonga, Theodore White, Z Jamie Yao, Ping Yeh, Adam Zalcman, Hartmut Neven, and John M Martinis. Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779):505–510, 2019. ISSN 1476-4687. DOI: 10.1038/s41586-019-1666-5.

[9] Salman Beigi, Nilanjana Datta, and Cambyse Rouzé. Quantum reverse hypercontractivity: Its tensorization and application to strong converses. *Communications in Mathematical Physics*, 376(2):753–794, may 2020. DOI: 10.1007/s00220-020-03750-z.

[10] Michael Ben-Or, Daniel Gottesman, and Avinatan Hassidim. Quantum refrigerator. *arXiv preprint arXiv:1301.1995*, 2013.

[11] Mario Berta, David Sutter, and Michael Walter. Quantum Brascamp-Lieb Dualities, 2019. arXiv:1909.02383v2.

[12] Sergio Boixo, Troels F. Rønnow, Sergei V. Isakov, Zhihui Wang, David Wecker, Daniel A. Lidar, John M. Martinis, and Matthias Troyer. Evidence for quantum annealing with more than one hundred qubits. *Nature Physics*, 10(3):218–224, feb 2014. DOI: 10.1038/nphys2900.

[13] Sergio Boixo, Sergei V. Isakov, Vadim N. Smelyanskiy, Ryan Babbush, Nan Ding, Zhang Jiang, Michael J. Bremner, John M. Martinis, and Hartmut Neven. Characterizing quantum supremacy in near-term devices. *Nature Physics*, 14(6):595–600, apr 2018. DOI: 10.1038/s41567-018-0124-x.

[14] Adam Bouland, Bill Fefferman, Chinmay Nirkhe, and Umesh Vazirani. On the complexity and verification of quantum random circuit sampling. *Nature Physics*, 15(2):159, 2019. DOI: https://doi.org/10.1038/s41567-018-0318-2.

[15] Zvika Brakerski, Venkata Koppula, Umesh Vazirani, and Thomas Vidick. Simpler Proofs of Quantumness. In Steven T. Flammia, editor, *15th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2020)*, volume 158 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 8:1–8:14, Dagstuhl, Germany, 2020. Schloss

Dagstuhl–Leibniz-Zentrum für Informatik. ISBN 978-3-95977-146-7. DOI: 10.4230/LIPIcs.TQC.2020.8.

[16] Michael J Bremner, Richard Jozsa, and Dan J Shepherd. Classical simulation of commuting quantum computations implies collapse of the polynomial hierarchy. In *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, volume 467, pages 459–472. The Royal Society, 2011. DOI: https://doi.org/10.1098/rspa.2010.0301.

[17] Michael J. Bremner, Ashley Montanaro, and Dan J. Shepherd. Achieving quantum supremacy with sparse and noisy commuting quantum computations. *Quantum*, 1:8, apr 2017. DOI: 10.22331/q-2017-04-25-8.

[18] Sébastien Bubeck. Convex Optimization: Algorithms and Complexity. *Foundations and Trends® in Machine Learning*, 8(3-4):231–357, 2015. ISSN 1935-8237. DOI: 10.1561/2200000050.

[19] Jacques Carolan, Jasmin D. A. Meinecke, Peter J. Shadbolt, Nicholas J. Russell, Nur Ismail, Kerstin Wörhoff, Terry Rudolph, Mark G. Thompson, Jeremy L. Brien, Jonathan C. F. Matthews, and Anthony Laing. On the experimental verification of quantum complexity in linear optics. *Nature Photonics*, 8(8):621–626, jul 2014. DOI: 10.1038/nphoton.2014.152.

[20] Kai-Min Chung, Yi Lee, Han-Hsuan Lin, and Xiaodi Wu. Constant-round Blind Classical Verification of Quantum Sampling. *arXiv:2012.04848 [quant-ph]*, December 2020. arXiv: 2012.04848.

[21] Christoph Dankert, Richard Cleve, Joseph Emerson, and Etera Livine. Exact and approximate unitary 2-designs and their application to fidelity estimation. *Physical Review A*, 80(1), jul 2009. DOI: 10.1103/physreva.80.012304.

[22] D.P. DiVincenzo, D.W. Leung, and B.M. Terhal. Quantum data hiding. *IEEE Transactions on Information Theory*, 48(3):580–598, Mar 2002. ISSN 0018-9448. DOI: 10.1109/18.985948.

[23] Daniel Stilck França and Raul Garcia-Patrón. Limitations of optimization algorithms on noisy quantum devices. *Nature Physics*, 17(11):1221–1227, oct 2021. DOI: 10.1038/s41567-021-01356-3.

[24] Xun Gao, Marcin Kalinowski, Chi-Ning Chou, Mikhail D. Lukin, Boaz Barak, and Soonwon Choi. Limitations of linear cross-entropy as a measure for quantum advantage, 2021. URL `https://arxiv.org/abs/2112.01657`.

[25] Daniel Gottesman. The heisenberg representation of quantum computers, 1998. arXiv:quant-ph/9807006.

[26] Martin Grötschel, László Lovász, and Alexander Schrijver. *Geometric algorithms and combinatorial optimization*, volume 2. Springer Science & Business Media, 2012.

[27] J. Haferkamp, D. Hangleiter, A. Bouland, B. Fefferman, J. Eisert, and J. Bermejo-Vega. Closing gaps of a quantum advantage with short-time hamiltonian dynamics. *Physical Review Letters*, 125(25):250501, dec 2020. DOI: 10.1103/physrevlett.125.250501.

[28] Dominik Hangleiter, Juani Bermejo-Vega, Martin Schwarz, and Jens Eisert. Anticoncentration theorems for schemes showing a quantum speedup. *Quantum*, 2:65, may 2018. DOI: 10.22331/q-2018-05-22-65.

[29] Dominik Hangleiter, Martin Kliesch, Jens Eisert, and Christian Gogolin. Sample complexity of device-independently certified "quantum supremacy". *Phys. Rev. Lett.*, 122:210502, May 2019. DOI: 10.1103/PhysRevLett.122.210502.

[30] Aram W Harrow and Ashley Montanaro. Quantum computational supremacy. *Nature*, 549(7671):203, 2017. DOI: https://doi.org/10.1038/nature23458.

[31] Christoph Hirche, Cambyse Rouzé, and Daniel Stilck França. On contraction coefficients, partial orders and approximation of capacities for quantum channels, 2020. arXiv:2011.05949v1.

[32] Cupjin Huang, Fang Zhang, Michael Newman, Junjie Cai, Xun Gao, Zhengxiong Tian, Junyin Wu, Haihong Xu, Huanjun Yu, Bo Yuan, Mario Szegedy, Yaoyun Shi, and Jianxin Chen. Classical simulation of quantum supremacy circuits, 2020. arXiv:2005.06787.

[33] Michael J. Kastoryano and Kristan Temme. Quantum logarithmic sobolev inequalities and rapid mixing. *Journal of Mathematical Physics*, 54(5):052202, may 2013. DOI: 10.1063/1.4804995.

[34] Michael Kearns. Efficient noise-tolerant learning from statistical queries. *Journal of the ACM*, 45(6):983–1006, nov 1998. DOI: 10.1145/293347.293351.

[35] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi. Optimization by simulated annealing. *Science*, 220(4598):671–680, may 1983. DOI: 10.1126/science.220.4598.671.

[36] M. Kliesch, T. Barthel, C. Gogolin, M. Kastoryano, and J. Eisert. Dissipative quantum church-turing theorem. *Physical Review Letters*, 107(12), sep 2011. DOI: 10.1103/physrevlett.107.120501.

[37] William Kretschmer. The Quantum Supremacy Tsirelson Inequality. In James R. Lee, editor, *12th Innovations in Theoretical Computer Science Conference (ITCS 2021)*, volume 185 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 13:1–13:13, Dagstuhl, Germany, 2021. Schloss Dagstuhl–Leibniz-Zentrum für Informatik. ISBN 978-3-95977-177-1. DOI: 10.4230/LIPIcs.ITCS.2021.13.

[38] David A Levin and Yuval Peres. *Markov chains and mixing times*, volume 107. American Mathematical Soc., 2017.

[39] AP Lund, Michael J Bremner, and TC Ralph. Quantum sampling problems, Boson sampling and quantum supremacy. *npj Quantum Information*, 3(1): 15, 2017. DOI: https://doi.org/10.1038/s41534-017-0018-2.

[40] Urmila Mahadev. Classical Verification of Quantum Computations. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 259–267, Paris, October 2018. IEEE. ISBN 978-1-5386-4230-6. DOI: 10.1109/FOCS.2018.00033.

[41] Ramis Movassagh. Efficient unitary paths and quantum computational supremacy: A proof of average-case hardness of random circuit sampling. *arXiv preprint arXiv:1810.04681*, 2018.

[42] Alexander Müller-Hermes, David Reeb, and Michael M. Wolf. Quantum subdivision capacities and continuous-time quantum coding. *IEEE*

Transactions on Information Theory, 61(1):565–581, jan 2015. DOI: 10.1109/tit.2014.2366456.

[43] Alexander Müller-Hermes, Daniel Stilck França, and Michael M. Wolf. Relative entropy convergence for depolarizing channels. *Journal of Mathematical Physics*, 57(2):022202, feb 2016. DOI: 10.1063/1.4939560.

[44] Alexander Müller-Hermes, Daniel Stilck França, and Michael M. Wolf. Entropy production of doubly stochastic quantum channels. *Journal of Mathematical Physics*, 57(2):022203, feb 2016. DOI: 10.1063/1.4941136.

[45] C. Neill, P. Roushan, K. Kechedzhi, S. Boixo, S. V. Isakov, V. Smelyanskiy, A. Megrant, B. Chiaro, A. Dunsworth, K. Arya, R. Barends, B. Burkett, Y. Chen, Z. Chen, A. Fowler, B. Foxen, M. Giustina, R. Graff, E. Jeffrey, T. Huang, J. Kelly, P. Klimov, E. Lucero, J. Mutus, M. Neeley, C. Quintana, D. Sank, A. Vainsencher, J. Wenner, T. C. White, H. Neven, and J. M. Martinis. A blueprint for demonstrating quantum supremacy with superconducting qubits. *Science*, 360(6385):195–199, apr 2018. DOI: 10.1126/science.aao4309.

[46] Feng Pan and Pan Zhang. Simulation of quantum circuits using the big-batch tensor network method. *Physical Review Letters*, 128(3):030501, jan 2022. DOI: 10.1103/physrevlett.128.030501.

[47] Edwin Pednault, John A. Gunnels, Giacomo Nannicini, Lior Horesh, and Robert Wisnieff. Leveraging secondary storage to simulate deep 54-qubit sycamore circuits, 2019. arXiv 1910.09534.

[48] D. S. Phillips, M. Walschaers, J. J. Renema, I. A. Walmsley, N. Treps, and J. Sperling. Benchmarking of Gaussian boson sampling using two-point correlators. *Physical Review A*, 99(2):023836, February 2019. ISSN 2469-9926, 2469-9934. DOI: 10.1103/PhysRevA.99.023836.

[49] Haoyu Qi, Daniel J. Brod, Nicolás Quesada, and Raul Garcia-Patron. Regimes of classical simulability for noisy gaussian boson sampling. *Physical Review Letters*, 124(10), mar 2020. DOI: 10.1103/physrevlett.124.100502.

[50] Lev Reyzin. Statistical queries and statistical algorithms: Foundations and applications, 2020.

[51] Seung Woo Shin, Graeme Smith, John A. Smolin, and Umesh Vazirani. How "quantum" is the d-wave machine?, 2014. arXiv 1401.7087.

[52] John A. Smolin and Graeme Smith. Classical signature of quantum annealing. *Frontiers in Physics*, 2, sep 2014. DOI: 10.3389/fphy.2014.00052.

[53] Nicolò Spagnolo, Chiara Vitelli, Marco Bentivegna, Daniel J. Brod, Andrea Crespi, Fulvio Flamini, Sandro Giacomini, Giorgio Milani, Roberta Ramponi, Paolo Mataloni, Roberto Osellame, Ernesto F. Galvão, and Fabio Sciarrino. Experimental validation of photonic boson sampling. *Nature Photonics*, 8 (8):615–620, jun 2014. DOI: 10.1038/nphoton.2014.135.

[54] Koji Tsuda, Gunnar Rätsch, and Manfred K Warmuth. Matrix exponentiated gradient updates for on-line learning and bregman projection. *J. Mach. Learn. Res.*, 6(Jun):995–1018, 2005.

[55] Benjamin Villalonga, Murphy Yuezhen Niu, Li Li, Hartmut Neven, John C. Platt, Vadim N. Smelyanskiy, and Sergio Boixo. Efficient approximation of experimental Gaussian boson sampling, 2021. arXiv:2109.11525v1.

[56] Lei Wang, Troels F. Rønnow, Sergio Boixo, Sergei V. Isakov, Zhihui Wang, David Wecker, Daniel A. Lidar, John M. Martinis, and Matthias Troyer. Comment on: "classical signature of quantum annealing", 2013. arXiv 1305.5837.

[57] Yulin Wu, Wan-Su Bao, Sirui Cao, Fusheng Chen, Ming-Cheng Chen, Xiawei Chen, Tung-Hsun Chung, Hui Deng, Yajie Du, Daojin Fan, Ming Gong, Cheng Guo, Chu Guo, Shaojun Guo, Lianchen Han, Linyin Hong, He-Liang Huang, Yong-Heng Huo, Liping Li, Na Li, Shaowei Li, Yuan Li, Futian Liang, Chun Lin, Jin Lin, Haoran Qian, Dan Qiao, Hao Rong, Hong Su, Lihua Sun, Liangyuan Wang, Shiyu Wang, Dachao Wu, Yu Xu, Kai Yan, Weifeng Yang, Yang Yang, Yangsen Ye, Jianghan Yin, Chong Ying, Jiale Yu, Chen Zha, Cha Zhang, Haibin Zhang, Kaili Zhang, Yiming Zhang, Han Zhao, Youwei Zhao, Liang Zhou, Qingling Zhu, Chao-Yang Lu, Cheng-Zhi Peng, Xiaobo Zhu, and Jian-Wei Pan. Strong quantum computational advantage using a superconducting quantum processor. *Physical Review Letters*, 127(18): 180501, oct 2021. DOI: 10.1103/physrevlett.127.180501.

[58] Han-Sen Zhong, Hui Wang, Yu-Hao Deng, Ming-Cheng Chen, Li-Chao Peng, Yi-Han Luo, Jian Qin, Dian Wu, Xing Ding, Yi Hu, Peng Hu, Xiao-Yan Yang, Wei-Jun Zhang, Hao Li, Yuxuan Li, Xiao Jiang, Lin Gan, Guangwen Yang, Lixing You, Zhen Wang, Li Li, Nai-Le Liu, Chao-Yang Lu, and Jian-Wei Pan. Quantum computational advantage using photons. *Science*, 370(6523): 1460–1463, December 2020. DOI: 10.1126/science.abe8770.

[59] Qingling Zhu, Sirui Cao, Fusheng Chen, Ming-Cheng Chen, Xiawei Chen, Tung-Hsun Chung, Hui Deng, Yajie Du, Daojin Fan, Ming Gong, Cheng Guo, Chu Guo, Shaojun Guo, Lianchen Han, Linyin Hong, He-Liang Huang, Yong-Heng Huo, Liping Li, Na Li, Shaowei Li, Yuan Li, Futian Liang, Chun Lin, Jin Lin, Haoran Qian, Dan Qiao, Hao Rong, Hong Su, Lihua Sun, Liangyuan Wang, Shiyu Wang, Dachao Wu, Yulin Wu, Yu Xu, Kai Yan, Weifeng Yang, Yang Yang, Yangsen Ye, Jianghan Yin, Chong Ying, Jiale Yu, Chen Zha, Cha Zhang, Haibin Zhang, Kaili Zhang, Yiming Zhang, Han Zhao, Youwei Zhao, Liang Zhou, Chao-Yang Lu, Cheng-Zhi Peng, Xiaobo Zhu, and Jian-Wei Pan. Quantum computational advantage via 60-qubit 24-cycle random circuit sampling. *Science Bulletin*, 67(3):240–245, feb 2022. DOI: 10.1016/j.scib.2021.10.017.

A Basic properties of mirror descent with the Shannon entropy as potential

In this section we review some basic properties of the mirror descent algorithm with the Shannon entropy as potential. We start with a simple proof of the update rule behind mirror descent in Lemma A.1. It shows how to obtain a Gibbs state τ_1 that is closer in relative entropy to a target distribution ν departing from a Gibbs state τ_0 and a function that distinguishes the latter from the target distribution.

Lemma A.1. *Let ν be a probability distribution on n bits and $\epsilon > 0$. Fix a function $H_0 : \{0,1\}^n \to \mathbb{R}$ and let $\tau_0 = e^{-H_0}/\mathcal{Z}_0$. Suppose that for some other (bounded) function $f : \{0,1\}^n \to [-1,1]$ we have:*

$$\mathbb{E}_{\tau_0}(f) - \mathbb{E}_{\nu}(f) \geq \epsilon. \tag{30}$$

Set $H_1 = H_0 + \frac{\epsilon}{4} f$. Then, the Gibbs state $\tau_1 = e^{-H_1}/\mathcal{Z}_1$ obeys

$$S(\nu\|\tau_1) - S(\nu\|\tau_0) \leq -\frac{\epsilon^2}{8}$$

Proof. We have:

$$S(\nu\|\tau_1) - S(\nu\|\tau_0) = \mathbb{E}_{\nu}\left[H_1 - H_0\right] + \ln\left(\frac{\mathcal{Z}_1}{\mathcal{Z}_0}\right) \tag{31}$$

By construction, $H_1 - H_0 = \frac{\epsilon}{4} f$ and the first term equals $\frac{\epsilon}{4}\mathbb{E}_{\nu}(f)$. The logarithmic ratio can be bounded using Jensen's inequality:

$$\ln\left(\frac{\mathcal{Z}_1}{\mathcal{Z}_0}\right) = -\ln\left(\mathbb{E}_{\tau_1}\left[e^{\frac{\epsilon}{4}f}\right]\right) \leq -\mathbb{E}_{\tau_1}\left[\frac{\epsilon}{4} f\right],$$

as the function $-\log(x)$ is convex. It then follows that

$$\mathbb{E}_{\nu}\left[H_1 - H_0\right] + \ln\left(\frac{\mathcal{Z}_1}{\mathcal{Z}_0}\right) \leq \frac{\epsilon}{4}\left(\mathbb{E}_{\nu}[\,f\,] - \mathbb{E}_{\tau_1}[f]\right). \tag{32}$$

We will now show that the expectation values of f on the distributions τ_1 and τ_0 are $\mathcal{O}(\epsilon)$ close. We can then replace the expectation value over τ_1 by only paying a small price and then use our hypothesis in Eq. (30).

To that end, a direct computation shows that

$$S\left(\tau_0\|\tau_1\right) = \frac{\epsilon}{4}\mathbb{E}_{\tau_0}(f) + \ln\left(\frac{\mathcal{Z}_1}{\mathcal{Z}_0}\right). \tag{33}$$

As in Eq. (32), we then can estimate the second term by:

$$S\left(\tau_0\|\tau_1\right) = \frac{\epsilon}{4}\mathbb{E}_{\tau_0}(f) + \ln\left(\frac{\mathcal{Z}_1}{\mathcal{Z}_0}\right) \leq \frac{\epsilon}{4}\left(\mathbb{E}_{\tau_0}(f) - \mathbb{E}_{\tau_1}(f)\right). \tag{34}$$

By Pinsker's inequality:

$$\left(\mathbb{E}_{\tau_0}(f) - \mathbb{E}_{\tau_1}(f)\right)^2 \leq 2S\left(\tau_0\|\tau_1\right) \leq \frac{\epsilon}{2}\left(\mathbb{E}_{\tau_0}(f) - \mathbb{E}_{\tau_1}(f)\right).$$

As $\left(\mathbb{E}_{\tau_0}(f) - \mathbb{E}_{\tau_1}(f)\right) \geq 0$ (see Lemma A.2 below for a proof) this yields

$$\left(\mathbb{E}_{\tau_0}(f) - \mathbb{E}_{\tau_1}(f)\right) \leq \frac{\epsilon}{2}.$$

We then see that:

$$\frac{\epsilon}{4}\left(\mathbb{E}_{\nu}[\,f\,] - \mathbb{E}_{\tau_1}[f]\right) \leq \frac{\epsilon}{4}\left(\mathbb{E}_{\nu}[f] - \mathbb{E}_{\tau_0}[f] + \frac{\epsilon}{2}\right).$$

Algorithm 3 *Mirror descent for reproducing expectation values.*

Require: Expectation value of functions $f_1, \ldots, f_k : \{0,1\}^n \to [-1,1]$ with respect to probability measure ν on n bits.

```
function MIRROR DESCENT(T, ε)
    Set μ_0 = U                                   ▷ initialize to the uniform distribution
    for t = 1, ..., T = ⌈8S(ν‖U) ε^{-2}⌉ do
        Check if |E_{μ_t}(f_i) − E_ν(f_t)| ≤ ε for all 1 ≤ i ≤ k.
        if Given that for a f_i we have |E_{μ_t}(f_i) − E_ν(f_t)| ≥ ε then
            if E_{μ_t}(f_i) − E_ν(f_t) ≥ ε then
                Set μ_{t+1}(x) = exp(−(ε/4) f_i(x) + log(μ_t))/Z_{t+1}.   ▷ Update the guess.
            else if E_{μ_t}(f_i) − E_ν(f_t) ≤ −ε then
                Set μ_{t+1}(x) = exp((ε/4) f_i(x) + log(μ_t))/Z_{t+1}.   ▷ Update the guess.
            end if
        else if For all |E_{μ_t}(f_i) − E_ν(f_t)| ≤ ε then
            Return μ_t
            break loop
        end if
    end for
    Return μ_T and exit function                  ▷ Current guess is ε indistinguishable from ν
end function
```

By our assumption in Eq. (30) we may then bound the right hand side in Eq. (32) by $-\frac{\epsilon^2}{8}$ and finally obtain:

$$S(\nu\|\tau_1) - S(\nu\|\tau_0) = \mathbb{E}_\nu\left[H_1 - H_0\right] + \ln\left(\frac{\mathcal{Z}_1}{\mathcal{Z}_0}\right) \leq -\frac{\epsilon^2}{8}.$$

The claim follows. □

With this Lemma at hand, we can then show that mirror descent will converge to a probability distribution approximating the expectation values of a given set of functions with respect to a probability measure. We will now show that mirror descent allows for recovering the expectation values of a set of functions and also give guarantees as to how well we approximate the distribution globally.

We then have:

Proposition A.1 (Mirror descent converges)**.** *Algorithm 3 returns a probability measure μ_t that satisfies*

$$|\mathbb{E}_{\mu_t}(f_i) - \mathbb{E}_\nu(f_t)| \leq \epsilon \tag{35}$$

for all $1 \leq i \leq k$ after at most $t \leq \lceil 8S(\nu\|\mathcal{U})\,\epsilon^{-2}\rceil$ iterations. Moreover, μ_t satisfies

$$\|\mu_t - \nu\|_{TV} \leq \sqrt{2\left(S(\nu\|\mathcal{U}) - \frac{t\epsilon^2}{8}\right)}. \tag{36}$$

Proof. If we exit the algorithm at line 12, then the output satisfies Eq. (35) by definition. Thus, it only remains to prove that this is indeed the case after $8\lceil S(\nu\|\mathcal{U})\,\epsilon^{-2}\rceil$ iterations. But note that Lemma A.1 and the update rules of Algorithm 3 ensure that we have:

$$S(\nu\|\mu_{t+1}) - S(\nu\|\mu_t) \leq -\frac{\epsilon^2}{8}$$

for all t. Applying a telescopic sum and our initial choice $\mu_0 = \mathcal{U}$ we see that

$$S(\nu\|\mu_t) \leq S(\nu\|\mathcal{U}) - t\frac{\epsilon^2}{8} \tag{37}$$

and the claim on the number of iterations follows from the positivity of the relative entropy. The claim in Eq. (36) follows from Eq. (37). □

Finally, let us prove for completeness the following standard fact that we used in the proof of Lemma A.1:

Lemma A.2. *Let $\tau_0 = \frac{e^{-H_0}}{\mathcal{Z}_0}$ be an arbitrary Gibbs probability measure on n bits and for a function $f : \{0,1\}^n \to [-1,1]$ define for $\lambda > 0$ the Gibbs probability measure*

$$\tau_\lambda = \frac{e^{-H_0-\lambda f}}{\mathcal{Z}_\lambda}.$$

Then $g : \lambda \mapsto \mathbb{E}_{\tau_\lambda}(f)$ is a monotone decreasing function.

Proof. The proof is quite standard and simple. It is easy to check that:

$$\frac{d}{d\lambda}g(\lambda) = -(\mathbb{E}_{\tau_\lambda}(f^2) - \mathbb{E}_{\tau_\lambda}(f)^2),$$

which is the negative of the variance of the function f under τ_λ. Thus, we clearly have that $\frac{d}{d\lambda}g(\lambda) \leq 0$ and the claim follows.

□

B Distinguishing function from the linear cross-entropy

At first sight, the current verification procedures for random circuits, the linear cross-entropy benchmark or the heavy output generation problem, do not readily fit into our framework. As we will see now, this is not defined as the expectation value of a bounded function. Indeed, define the function

$$f(x) = 2^n\nu(x) - 1$$

where $\nu(x)$ is the probability of string x under ν, the output of the ideal circuit. We have that the linear cross-entropy is given by

$$\mathcal{F}_{\mathrm{XEB}}(\mu) = \mathbb{E}_\mu(f) - 1. \tag{38}$$

In principle, the function f could take values between $[-1, 2^n - 1]$, whereas our framework required the distinguishing functions to take values in $[-1, 1]$. However, we will now show that by suitably discarding high values of f and restricting μ to a suitable set of distributions, we can massage the linear cross-entropy into our framework. That is, we will find a bounded function f_r that approximates $\mathcal{F}_{\mathrm{XEB}}$. Once again, the main property required to show this is the fact that random quantum circuits have very flat outcome distributions.

To prove our claims, we will first assume that the distribution of probabilities of the outcomes is well-approximated by a Porter-Thomas distribution, as explained in detail below. We refer to [13] for a justification of this assumption and numerical evidence of its validity. It is possible to obtain similar but weaker results departing from the assumption that the output is an approximate 3-design. However, for the sake of conciseness, we will restrict to the Porter-Thomas distribution.

The Porter-Thomas assumption is an approximation of the probability that a given outcome string x will have for a family of quantum circuits. More specifically, it assumes that for all strings x, the random variable corresponding to the value of $\nu(x)$ under this family of random quantum circuits follows the density κ given by

$$\kappa(p) = 2^n e^{-p2^n}. \tag{39}$$

It is not difficult to see that this distribution is highly concentrated around its mean, 2^{-n}, and that its variance is 2^{-2n}, as it corresponds to an exponential distribution with parameter 2^{-n}.

We then have:

Proposition B.1. *Let ν be the output of a random quantum circuit on n qubits and assume that the density of outcomes is given by a Porter-Thomas distribution with parameter 2^{-n}, as in Eq. (39). For a parameter $r \geq 1$ define $f_r : \{0,1\}^n \to [-1, 1]$ as*

$$f_r(x) = r^{-1}(\min\{2^n \nu(x), r\} - 1).$$

Then for another distribution μ satisfying for some constant $C > 0$

$$\mu(x) \leq C\nu(x) \tag{40}$$

almost surely for all x such that $\nu(x) \geq r2^{-n}$ we have:

$$\mathbb{E}\left[|\mathcal{F}_{\mathrm{XEB}}(\mu) - r\mathbb{E}_\mu(f_r)|\right] \leq Ce^{-r}(2 + r). \tag{41}$$

where the expectation is taken over the circuits.

Proof. By the definition of $f_r(x)$ we have that

$$\mathbb{E}\left[|\mathcal{F}_{\mathrm{XEB}}(\mu) - r\mathbb{E}_\mu(f_r)|\right] = \mathbb{E}\left[\sum_{x:\nu(x)\geq \frac{r}{2^n}} (2^n\nu(x) - r)\mu(x)\right].$$

By our assumption on the distribution in Eq. (40) we have that

$$\mathbb{E}\left[\sum_{x:\nu(x)\geq\frac{r}{2^n}}(2^n\nu(x)-r)\mu(x)\right]\leq C\mathbb{E}\left[\sum_{x:\nu(x)\geq\frac{r}{2^n}}(2^n\nu(x)^2-r\nu(x))\right].$$

Furthermore, by the assumption that the probability of the output strings follows a Porter-Thomas distribution, we conclude that:

$$\mathbb{E}\left[\sum_{x:\nu(x)\geq\frac{r}{2^n}}(2^n\nu(x)^2-r\nu(x))\right]=\int_{\frac{r}{2^n}}^{+\infty}2^{2n}e^{-2^nx}(2^nx^2-rx)dx=e^{-r}(2+r),$$

which yields the claim. □

Thus, we see that for distributions that do not differ too much from the distribution of the outcome of the random circuit in the sense of Eq. (40), $\mathcal{F}_{\text{XEB}}$ can be approximated in our framework. By picking a cut off at $\log(\epsilon^{-2})$ we ensure that the difference between the truncated f_r and $\mathcal{F}_{\text{XEB}}$ only differ by $\mathcal{O}(\epsilon)$.

Let us now discuss the condition in Eq. (40) in more detail. The condition in Eq. (40) has a natural interpretation for the problem at hand. The goal of Prop. B.1 is to identify conditions under which $\mathcal{F}_{\text{XEB}}$ is well-approximated by a bounded function. However, if a probability measure μ only satisfies Eq. (40) for large values of C, it means it assigns high probability outcomes of ν even more weight than ν. This in turn will yield higher values for $\mathcal{F}_{\text{XEB}}(\mu)$. However, for outcome distributions that are not strongly concentrated on heavy outcomes, we expect Eq. (40) to hold for moderate values of C. For instance, for the uniform distribution the condition holds with $C=1$.

However, it is possible to construct distributions that converge to the true distribution in total variation and for which $\mathcal{F}_{\text{XEB}}(\mu)$ diverges. At the same time, it is possible to construct distributions that are a constant distance away from the ideal distribution in total variation, do not satisfy Eq. (40) and nevertheless satisfy $\mathcal{F}_{\text{XEB}}(\mu)=\mathcal{F}_{\text{XEB}}(\nu)$. We will give the explicit constructions of these distributions shortly. But they showcase that in principle there is no connection between $\mathcal{F}_{\text{XEB}}(\mu)$ and the total variation distance between μ and ν.

Both constructions will exploit the fact that the $\mathcal{F}_{\text{XEB}}$ is unbounded, as expected. Indeed, if the benchmark we were using were bounded, then at least we can always conclude from a convergence in total variation distance that the expectation values also have to converge.

Thus, we believe that these examples showcase why we cannot expect that $\mathcal{F}_{\text{XEB}}$ can always be captured in our framework. Whereas our framework is intimately connected to the the two distributions being close in total variation distance, this is not the case for similar linear cross entropy. This was also observed in [24], where the authors give additional arguments why the linear cross entropy is not connected with the total variation distance or fidelity in general.

Example B.1 (Distributions close in total variation distance but diverging $\mathcal{F}_{\text{XEB}}$)**.** *To construct our examples, observe that it follows from the Porter-Thomas assumption in Eq. (39) that for some given $c_1<1$, we expect $2^{(1-c_1)n}$ strings to have*

probability at least $c_1 n 2^{-n}$. Indeed, the expected number of strings with probability at least $c_2 n 2^{-1}$ is:

$$2^n \int_{c_1 n 2^{-n}}^{\infty} 2^n e^{-p 2^n} dp = 2^{(1-c_1)n}.$$

Now define

$$B_{c_1} = \{x \in \{0,1\}^n : \nu(x) \geq c_1 n 2^{-n}\}$$

and let $\mathcal{U}_{B_{c_1}}$ be the uniform distribution on B_{c_1}. Further define the distribution

$$\mu_{c_1} = \left(1 - \frac{1}{\sqrt{n}}\right)\nu + \frac{1}{\sqrt{n}}\mathcal{U}_{B_{c_1}}.$$

Clearly, $\|\mu_{c_1} - \nu\|_{TV} = \mathcal{O}(n^{-\frac{1}{2}})$. However, $\mathcal{F}_{\mathrm{XEB}}(\mu_{c_1}) = \Omega(\sqrt{n})$, as $\mathcal{F}_{\mathrm{XEB}}(\mathcal{U}_{B_{c_1}}) = \Omega(n)$ by definition.

Example B.2 (Distributions far away in total variation distance but $\mathcal{F}_{\mathrm{XEB}}$ is similar)**.** *To construct this example, we will resort to the same distribution as above. Let $a = \mathcal{F}_{\mathrm{XEB}}(\mathcal{U}_{B_{c_1}})$. Again, by the definiition of B_{c_1}, $a \geq c_1 n - 1$. Now let μ'_{c_1} be defined as*

$$\mu'_{c_1} = (1 - 1/a)\mathcal{U} + \frac{1}{a}\mathcal{U}_{B_{c_1}}. \tag{42}$$

Using the fact that $\mathcal{F}_{\mathrm{XEB}}(\mathcal{U}) = 0$, by the linearity of $\mathcal{F}_{\mathrm{XEB}}$ we get that $\mathcal{F}_{\mathrm{XEB}}(\mu'_{c_1}) = 1$, which is the expected value for ν. Thus the value of $\mathcal{F}_{\mathrm{XEB}}$ coincides for both distributions. But a reverse triangle inequality together with the fact that $\|\nu - \mathcal{U}\|_{TV} = \Omega(1)$ shows that

$$\|\mu'_{c_1} - \nu\|_{TV} = \Omega(1). \tag{43}$$

In spite of the limitations of the $\mathcal{F}_{\mathrm{XEB}}$ showcased above, for the uniform distribution the situation is less complicated. As the uniform distribution corresponds to the guess in the first round of the game and deserves a detailed analysis, we will now directly compute by how much a suitably cut-off and normalized linear cross-entropy allows for distinguishing the output of the random quantum circuit from the uniform distribution.

Proposition B.2. *Let f_r and ν as in the statement of Prop. B.1 and $\mathcal{U}$ be the uniform distribution on n bits. Then we have for $r \geq 1$:*

$$\mathbb{E}\left[\mathbb{E}_\nu(f_r) - \mathbb{E}_{\mathcal{U}}(f_r)\right] = \frac{1 - e^{-r}(1 + 2r)}{r} \tag{44}$$

where the first expectation value is taken over the random quantum circuits.

Proof. The proof is similar to the last proposition. We have that:

$$\mathbb{E}\left[\mathbb{E}_\nu(f_r) - \mathbb{E}_{\mathcal{U}}(f_r)\right] = \sum_{x:\nu(x)\leq\frac{r}{2^n}} r^{-1}\nu(x)2^n\left(\nu(x)-\frac{1}{2^n}\right) + (1-r^{-1})\sum_{x:\nu(x)>\frac{r}{2^n}}\left(\nu(x)-\frac{1}{2^n}\right).$$

Taking the expectation, the first sum above translates to the integral

$$r^{-1}\int_0^{\frac{r}{2^n}} 2^{2n}e^{-2^n x}2^n x\left(x-\frac{1}{2^n}\right)dx = r^{-1}\left(1-e^{-r}(1+r^2+r)\right), \tag{45}$$

whereas the second translates to

$$(1-r^{-1})\int_{\frac{r}{2^n}}^{+\infty} 2^{2n}e^{-2^n x}\left(x-\frac{1}{2^n}\right)dx = (1-r^{-1})re^{-r}. \tag{46}$$

Summing the two expressions yields the claim. □

It is not immediately obvious how to maximize the expression in Eq. (44) analytically, but numerically solving it we see that it is around $r \simeq 3.21$, for which we obtain a violation of $\simeq 0.22$. As the expected total variation distance is $1/e \simeq 0.36$ [13] under the Porter-Thomas assumption, we see that this function is not far from the optimal distinguishing function. Thus, Bob could propose the function $f_{3.21}$ to distinguish the distribution of his device and the uniform distribution in the ideal case. Of course, as evaluating $f_{3.21}$ requires us to compute outcome probabilities, this is not an efficient distinguishing function. But the results of this section showcase that the linear cross-entropy fits into our framework by introducing a suitable cut-off as long as the underlying distribution does not put too much additional weight on heavy outputs.

C Bound on the Shannon entropy from design property

We will now show that the Shannon entropy of the output distributions of approximate two designs when measured in the computational basis is essentially maximal. This result is similar in spirit to those of [2, 29].

Proposition C.1. *Let U be a $(2, \epsilon 2^{-2n-1})$ approximate unitary design and define ν as before. Then, with probability at least $1-\delta$:*

$$S(\nu) \geq n - \log(2+\epsilon) - \log(\delta^{-1})$$

Proof. For a Haar random unitary and $x \in \{0,1\}^n$ we have that:

$$\mathbb{E}\left(|\langle x|\, U\, |0\rangle|^2\right) = \frac{1}{2^n}, \quad \mathbb{E}\left(|\langle x|\, U\, |0\rangle|^4\right) = \frac{2}{2^n(2^n+1)}.$$

Thus, for an approximate two design as above, we have that:

$$\mathbb{E}\left(|\langle x| U |0\rangle|^4\right) \leq \frac{2}{2^n(2^n+1)} + \frac{\epsilon}{2^n(2^n+1)}.$$

Recall that the 2-Renyi entropy S_2 is defined as:

$$S_2(\nu) = -\log\left(\sum_x \nu(x)^2\right)$$

and that $S(\nu) \geq S_2(\nu)$. Moreover, the function $-\log$ is convex. Thus, it follows from Jensen's inequality that:

$$\mathbb{E}\left(-\log\left(\sum_x \nu(x)^2\right)\right) \geq -\log\left(\mathbb{E}\left[\sum_x \nu(x)^2\right]\right).$$

From the computations above, we have that:

$$\frac{2+\epsilon}{(2^n+1)} \geq \mathbb{E}\left[\sum_x \nu(x)^2\right],$$

from which we readily obtain that:

$$\mathbb{E}\left(-\log\left(\sum_x \nu(x)^2\right)\right) \geq n - \log(2+\epsilon).$$

It follows from Markov's inequality that

$$\mathbb{P}\left(\sum_x \nu(x)^2 \geq \frac{2+\epsilon}{\delta 2^n}\right) \leq \delta, \tag{47}$$

from which the claim follows. □

D Auxiliary results for Section 5

In Section 5 we showed that being able to efficiently distinguish the probability distributions arising from sampling from random circuits from the uniform distribution implies the ability to fool the XHOG problem associated to the same class of circuits to a certain level.

But we assumed that the function that distinguished the distributions has binary outputs, although our framework for distinguishability allows for functions with outputs in $[-1, 1]$. We now show that it is always possible to obtain an efficiently computable function with binary outputs from an efficiently function with image $[-1, 1]$ that distinguishes the two probability distributions, at the expense of a smaller distinguishability power.

Lemma D.1. *Let $f : \{0,1\}^n \to [-1, 1]$ be a function that can be computed in polynomial time such that for some probability measure ν and $\epsilon > 0$ we have that:*

$$\mathbb{E}_\nu(f) - \mathbb{E}_{\mathcal{U}}(f) \geq \epsilon.$$

Then there exists a function $f' : \{0,1\}^n \to \{0,1\}$ that can be computed in polynomial time such that:

$$\mathbb{E}_\nu(f') - \mathbb{E}_{\mathcal{U}}(f') \geq \frac{\epsilon^2}{17}.$$

Proof. First, we will consider instead of f the shifted and normalized function

$$\tilde{f} = \frac{f+1}{2},$$

as it then has image in $[0,1]$ and clearly

$$\mathbb{E}_\nu(\tilde{f}) - \mathbb{E}_{\mathcal{U}}(\tilde{f}) \geq \frac{\epsilon}{2}.$$

Assume w.l.o.g. that $\epsilon = m^{-1}$ for some integer m and consider the discretization of f_1 of f given by:

$$f_1(x) = (8m)^{-1}\lceil 8m\tilde{f}(x)\rceil.$$

Note that $\|\tilde{f} - f_1\|_\infty \leq (8m)^{-1}$ and, thus,

$$\mathbb{E}_\nu(f_1) - \mathbb{E}_{\mathcal{U}}(f_1) \geq \frac{\epsilon}{4}. \tag{48}$$

Recall that for every real valued random variable X we have:

$$\mathbb{E}(X) = \int \mathbb{P}(X \geq x)dx, \tag{49}$$

Moreover, note that f_1 only takes the $8m+1$ possible values $\{0, (8m)^{-1}, \ldots, 1\}$. Thus. combining this observation with Eq. (48) and the identity in Eq. (49) we see that

$$\mathbb{E}_\nu(f_1) = \sum_{k=0}^{8m} \nu\left(f_1(x) \geq \frac{k}{8m}\right) \geq \frac{\epsilon}{4} + \sum_{k=0}^{8m} \frac{\left|\{f_1(x) \geq \frac{k}{8m}\}\right|}{2^n}. \tag{50}$$

It then follows that for at least one $0 \leq k_0 \leq 8m$ we have that:

$$\nu\left(f_1(x) \geq \frac{k_0}{8m}\right) \overset{(1)}{\geq} \frac{\left|\{f_1(x) \geq \frac{k_0}{8m}\}\right|}{2^n} + \frac{\epsilon}{2(8m+1)} \geq \frac{\left|\{f_1(x) \geq \frac{k_0}{8m}\}\right|}{2^n} + \frac{\epsilon^2}{17}, \tag{51}$$

because if the opposite inequality would hold in (1) for all k_0 we would obtain a contradiction to Eq. (50) by summing over all k_0. Thus, by setting $f'(x) = 1$ if $f_1(x) \geq \frac{k}{8m}$ and 0 else and recalling that $m^{-1} = \epsilon$, we see that Eq. (51) immediately implies

$$\mathbb{E}_\nu(f') - \mathbb{E}_{\mathcal{U}}(f') \geq \frac{\epsilon^2}{17},$$

which yields the claim. □

E Generalizations and limitations of our results

Let us discuss more precisely to what class of verification and distinguishability algorithms our results apply to, how it can be generalized and how it relates to the statistical query model well-known in statistical learning theory [34, 50].

E.1 The one-shot model

Most current verification and distinguishability proposals we are aware of in the random circuit literature have a simple structure. They consist of defining a (not necessarily bounded) function like the cross-entropy benchmark and evaluating its empirical average. We call this scenario one-shot as, despite involving an estimation over many samples, its theoretical analysis involve the distribution of a single realization.

As explained in Sec. 3.2, the optimal probability of success for correctly distinguishing two distributions from one sample is bounded for *any* algorithm by the total variation distance. Thus, if we have that the trace distance between two distributions μ, ν is small, then we can conclude that no algorithm will be able to perform significantly better than random guessing in the one-shot setting.

E.1.1 Comparison to statistical query model

This approach can be naturally cast in the statistical query model [34, 50]. In that model, one is not given access to samples of a distribution μ, but one is allowed to query $\mathbb{E}_\mu(f)$ up to some additive error tolerance $\epsilon > 0$ for arbitrary $f : \{0,1\}^n \to [-1,1]$. We can see that it is possible to formalize our game in this model, as it is Bob's job to distinguish his distribution from Alice's through he expectaion values of such functions. And, as discussed before, this has a natural interpretation in terms of the succes probability of one-shot distinguishing algorithms.

E.2 A more general framework: multiple copies discrimination

However, there is no a-priori reason why we should limit ourselves to the one-shot scenario. We could define more generally efficient distinguishing algorithms that take as input a polynomial number of samples m from a distribution, perform a polynomial-time postprocessing of the data and then outputs a guess.

In this general framework one considers the success probability of arbitrary distinguishing algorithms that take as an input a polynomial number of samples, i.e., $\mu^{\otimes m}$ instead of μ for $m = \text{poly}(n)$. Unfortunately, our techniques cannot discard the existence of such an efficient distinguishing procedure. Indeed, if there is a polynomial time function f such that for some $m \geq 1$ we have that:

$$\left|\mathbb{E}_{\mu^{\otimes m}}(f) - \mathbb{E}_{\nu^{\otimes m}}(f)\right| = \Omega(n^{-m}), \tag{52}$$

then we can take $\mathcal{O}(n^{2m}\log(\delta^{-1}))$ samples from the distribution, compute the empirical average of f on them and distinguish the two with probability of success

at least $1-\delta$. As f can be computed in polynomial time, the empirical average can also be computed efficiently and this yields and efficient distinguishing procedure. And our techniques do not discard the existence of an efficient f as in Eq. (52).

Although in the main text we only considered the case of $m = 1$ for our results, it should be noted that our results naturally extend to the settig in which $m\epsilon^{-1} = \mathcal{O}(\log(n))$. That is, if we are in regime $\epsilon = \Omega(1)$, we can also consider the case in which the distinguishing functions act on a logarithmic number of samples. To see why, note that the proof of Thm. 4.1 relied solely on the fact that with probability at least $1 - \delta$ we have

$$S(\nu\|\mathcal{U}) = \mathcal{O}(1 + \log(\delta^{-1})) \tag{53}$$

for the output distributions of approximate 2-designs.

If we consider instead m copies i.i.d. samples of the distribution ν, the joint output distribution satisfies

$$S(\nu^{\otimes m}\|\mathcal{U}^{\otimes m}) = mS(\nu\|\mathcal{U}) = \mathcal{O}(m(1 + \log(\delta^{-1})))$$

by the additivity of the relative entropy. Thus, as in Thm. 4.1, if we set our error tolerance to be ϵ, mirror descent will converge after $\mathcal{O}(m^2\epsilon^{-2})$ iterations to a distribution μ that approximates $\nu^{\otimes m}$ up trace distance ϵ. Moreover, sampling from μ using rejection sampling takes $\mathcal{O}(e^{m\epsilon^{-1}})$ evluations of f on average. And from this we obtain that as long as the distinguishing functions are efficient and $m\epsilon^{-1} = \mathcal{O}(\log(n))$, the whole procedure is efficient and our results still apply.

F Distinguishing the output distribution of stabilizer states

Note that we only assumed in the proofs of Sec. 5 that the underlying circuit ensemble is an approximate two design. It is well-known that circuits being approximate two designs does not imply that one cannot sample from their output distribution efficiently, as is prominently exemplified by Clifford circuits. Random Cliffords are two designs [21, 22] and it is possible to simulate measurements in the computational basis efficiently for them [25, 4]. We now show that in this case, it is also possible to easily distinguish the outcome distribution from the uniform one, if this is possible at all.

It is not difficult to see that for a Pauli string $P = \otimes_{i=1}^n \sigma_i$ we have for a Clifford C that

$$\mathrm{tr}\left(PC|0\rangle\langle 0|^{\otimes n}C^\dagger\right) \in \{-1, 0, 1\}.$$

This is because, as Cliffords stabilize the Pauli group, we have that $\tilde{P} = C^\dagger PC$ is again, up to a global sign, a Pauli string. And for a Pauli string $\mathrm{tr}\left(\tilde{P}|0\rangle\langle 0|^{\otimes n}\right) \in \{0, 1\}$. Now assume that there exists a Pauli string P consisting only of I and Z Pauli matrices that differs from the identity and such that

$$\mathrm{tr}\left(PC|0\rangle\langle 0|^{\otimes n}C^\dagger\right) \neq 0.$$

Then, interpreting diagonal operators as functions, we have that the function $f = \frac{P+I}{2}$ is a binary function that satisfies:

$$|\mathbb{E}_\nu(f) - \mathbb{E}_\mathcal{U}(f)| = \frac{1}{2}.$$

Thus, in this case, we have found a function that efficiently distinguishes the outcome from uniform. But note that in some cases such a Pauli string does not exist, such as if the Clifford is $H^{\otimes n}$, as then the outcome distribution is uniform.

Let us now discuss how to efficiently find the appropriate distinguishing Pauli string. We refer to [4] for a review of the basics of the stabilizer formalism. First, we recall that a stabilizer state $|\psi\rangle$ on n qubits can always be described by n generators $g_1, \ldots, g_n$ of its stabilizer group $SG(|\psi\rangle)$. Moreover, note that for a Pauli string

$$\mathrm{tr}\left(PC|0\rangle\langle 0|^{\otimes n}C^\dagger\right) \in \{-1, 1\}.$$

is equivalent to $P \in SG(|\psi\rangle)$ or $-P \in SG(|\psi\rangle)$. Thus, by our previous discussion, the problem of finding a function to distinguish the output of the Clifford circuit from the uniform distribution is equivalent to finding a stabilizer of the state consisting solely of Z and I Pauli operators.

Let us now discuss how to achieve this. First, decompose each generator g_i as the product of a string of Pauli X and Pauli Z matrices plus a global ± 1 phase and represent each one of these as vectors (x_i, z_i, s_i) in $\mathbb{F}_2^{2n+1}$. In order to simplify the presentation, we are not going to keep track of the global phase of the elements of the stabilizer group for now. Thus, we restrict to the vectors (x_i, z_i) corresponding to the first $2n$ entries. It is easy to see that if we do not keep track of the global phase, then multiplying two generators is equivalent to adding the corresponding vectors in $\mathbb{F}_2^{2n}$. Now define the $n \times (2n+1)$ binary matrix A with the vectors x_i in its rows. We then have:

Proposition F.1. *Let $|\psi\rangle$ be a stabilizer state with generators $g_1, \ldots, g_n$ and corresponding vectors $(x_i, z_i) \in \mathbb{Z}^{2n}$. Then there exists a Z string $P \in SG(|\psi\rangle)$ if and only if:*

$$\mathrm{span}\{(x_1, z_1), \ldots, (x_n, z_n)\} \cap (\{0\} \times \mathbb{F}_2^n) \neq \{0\} \tag{54}$$

Proof. Note that

$$\mathrm{span}\{(x_1, z_1), \ldots, (x_n, z_n)\}$$

corresponds to the elements of the stabilizer group of the state, up to a global phase. This is because, as discussed before, multiplication in the Pauli group just corresponds to a sum of the vectors, up to the global phase. Thus, if we find a string of Z in the stabilizer group, then it is also in the intersection in eq. (54). □

Thus, we can find the distinguishing Pauli operator by a nonzero element of the subspace

$$\mathrm{span}\{(x_1, z_1), \ldots, (x_n, z_n)\} \cap (\{0\} \times \mathbb{F}_2^n),$$

which can be done by Gaussian elimination.