

# Noise and the frontier of quantum supremacy

Adam Bouland\* Bill Fefferman† Zeph Landau‡ Yunchao Liu§

## Abstract

Noise is the defining feature of the NISQ era, but it remains unclear if noisy quantum devices are capable of quantum speedups. Quantum supremacy experiments have been a major step forward, but gaps remain between the theory behind these experiments and their actual implementations. In this work we initiate the study of the complexity of quantum random circuit sampling experiments with realistic amounts of noise.

Actual quantum supremacy experiments have high levels of uncorrected noise and exponentially decaying fidelities. It is natural to ask if there is any signal of exponential complexity in these highly noisy devices. Surprisingly, we show that it remains hard to compute the output probabilities of noisy random quantum circuits without error correction. More formally, so long as the noise rate of the device is below the error detection threshold, we show it is  $\#P$ -hard to compute the output probabilities of random circuits with a constant rate of noise per gate. This hardness persists even though these probabilities are exponentially close to uniform. Therefore the small deviations away from uniformity are hard to compute, formalizing an important intuition behind Google’s supremacy claim.

Interestingly these hardness results also have implications for the complexity of experiments in a low-noise setting. The issue here is that prior hardness results for computing output probabilities of random circuits are not robust enough to imprecision to connect with the Stockmeyer argument for hardness of sampling from circuits with constant fidelity. We exponentially improve the robustness of prior results to imprecision, both in the cases of Random Circuit Sampling and BosonSampling. In the latter case we bring the proven hardness within a constant factor in the exponent of the robustness required for hardness of sampling for the first time. We then show that our results are in tension with one another – the high-noise result implies the low-noise result is essentially optimal, even with generalizations of our techniques.

---

\*Simons Institute for the Theory of Computing and Department of EECS, UC Berkeley. [abouland@berkeley.edu](mailto:abouland@berkeley.edu)

†Department of Computer Science, University of Chicago. [wjf@uchicago.edu](mailto:wjf@uchicago.edu)

‡Department of EECS, UC Berkeley. [zeph.landau@gmail.com](mailto:zeph.landau@gmail.com)

§Department of EECS, UC Berkeley. [yunchaoliu@berkeley.edu](mailto:yunchaoliu@berkeley.edu)

# 1 Introduction

The last two years have seen the first claimed demonstrations of “quantum supremacy” [AAB<sup>+</sup>19, ZWD<sup>+</sup>20] – the experimental demonstration of an exponential quantum computational speedup. Quantum supremacy is both a necessary step on the path to building useful quantum computers, and also marks the first experimental challenge to the Extended Church-Turing thesis, a bedrock of theoretical computer science. Achieving quantum supremacy requires finding a computational task which is experimentally feasible and has strong complexity-theoretic evidence for classical intractability. Two prominent supremacy proposals have been BosonSampling [AA11] and Random Circuit Sampling [BIS<sup>+</sup>18], both of which involve sampling from the output distribution of a randomly chosen quantum experiment<sup>1</sup>.

Despite having advantages in being near-term implementable, it is a priori unclear why these random quantum experiments should be difficult to classically simulate. While we know of many examples of quantum algorithms that attain exponential speedups over classical computation, they all rely on highly structured circuits (such as quantum Fourier transforms) which are far from typical. Why then should we expect a generic quantum experiment to realize a large computational advantage?

Aaronson and Arkhipov [AA11] (in the setting of BosonSampling), and Bouland *et al.* [BFNV19] and Movassagh [Mov20] (in the setting of Random Circuit Sampling) gave initial evidence that these random experiments might be difficult to simulate classically<sup>2</sup>. These works essentially show that it is  $\#P$ -hard to *exactly compute* the output probabilities of ideal, noiseless, quantum experiments on average<sup>3</sup>. To do this, they establish a worst-to-average-case reduction for computing the output probabilities of random experiments. Since it is a well-established fact that computing the output probability of a worst-case quantum experiment is  $\#P$ -hard (see e.g., [TD02]), these results prove that computing the output probabilities of a random quantum experiment is  $\#P$ -hard as well.

These arguments have two primary limitations which keep them from directly matching the constraints of noisy quantum sampling experiments. First, they are sensitive to additive imprecision, which prevents them from addressing the hardness of classically mimicking *low noise* experiments. For example, in the case of random circuit sampling experiments with  $m$  gates (where we assume  $m$  is greater than the number of qubits,  $n$ ) the prior works are only able to prove hardness of computing an estimate to the output probability that achieves an additive error at most  $2^{-O(m^3)}$  [BFNV19, Mov20]. In contrast, standard reductions [AA11, Sto85] require hardness to additive imprecision  $O(2^{-n})$  in order to show that the more natural sampling task – which is what a low-noise<sup>4</sup> random quantum circuit experiment solves “natively” – cannot be reproduced classically. A similarly large gap exists between the known ( $e^{-O(n^4)}$ ) vs. conjectured ( $O(e^{-n \log n})$ ) robustness of hardness results for BosonSampling. Proving a hardness result which is robust even to this small amount of noise remains open for all quantum supremacy proposals.

The second and more fundamental limitation of these arguments is that they say nothing about the complexity of *high noise* experiments. The prior average-case hardness arguments of [AA11, BMS16, BFN19, Mov20] all explore the complexity of mimicking low-noise experiments performed with a high fidelity that stays constant as the system size grows. This would require improved

<sup>1</sup>Either a random circuit in the case of RCS, or a random linear optical network in the case of BosonSampling.

<sup>2</sup>See also the complementary work of [AC17, AG19] for alternative evidence that directly addresses the hardness of scoring well on benchmarking tests such as cross-entropy.

<sup>3</sup>Please see Section 2 for a discussion of the differences between these works.

<sup>4</sup>That is, such a hardness result would imply that no classical algorithm can sample from any distribution within a fixed constant total variation distance of the ideal distribution sampled by the noiseless quantum circuit. Interestingly, we do not even have hardness of *exact* sampling results for random quantum circuits owing to this same imprecision issue!

control over the noise (e.g. error correction) to achieve asymptotically. But in the NISQ era we do not have the quantum resources to implement error correction and so experiments are highly noisy. For example, Google’s recent quantum supremacy experiment estimated that their fidelity was merely  $\sim 0.2\%$  (i.e., the experiment was  $\sim 99.8\%$  noise) [AAB<sup>+</sup>19] and that their fidelity will further decrease as they scale their system. Therefore their quantum supremacy claim hinges on whether or not random circuit sampling is intractable in the high noise regime, in which there is only a small signal of the correct experiment in a sea of noise, and this signal diminishes with system size. In some cases noise can render the simulation of certain quantum circuits substantially easier (e.g. [BMS17, ZSW20]). For this reason, it is critical to examine if these hardness of computation results are robust to noise models that are relevant to near-term quantum experiments. For example, can anything be said about the hardness of computing the output probabilities of such circuits in the high noise regime, without error correction?

In this work we make progress on both these limitations. First, we substantially improve the robustness of prior hardness results. For example, in the case of RCS, we show that it is  $\#P$ -hard to compute the output probabilities of (noiseless) random circuits to additive imprecision  $2^{-O(m \log m)}$ . In the regime of constant depth circuits where  $m = O(n)$ , this robustness nearly reaches that needed to establish the classical hardness of sampling in the low noise regime, though a small gap remains.

**Theorem 1** (Simplified). *For any constant  $\eta < 1/4$ , it is  $\#P$ -hard to compute the output probability of random quantum circuits  $C$  up to additive error  $2^{-O(m \log m)}$ , with probability at least  $1 - \eta$  over the choice of  $C$  with  $m$  gates.*

We note that Kondo, Mori and Movassagh [KMM21] recently claimed the same robustness as our Theorem 1 and arrived at it independently.

We also give an analogous result for BosonSampling, showing that it is hard to compute the output probabilities of (noiseless) random  $n$ -photon  $m = n^2$  mode linear optical networks to additive imprecision  $e^{-6n \log n - O(n)}$  (see Corollary 1). This nearly matches the desired robustness for BosonSampling ( $O(e^{-n \log n})$ ) up to a constant factor in the exponent. To our knowledge, this is the first time that a quantum supremacy proposal exhibits a proven robustness of hardness which is polynomially related to the conjectured robustness in absolute terms.

Second, we show that these hardness results remain true even in the high noise regime – i.e. in the presence of a constant rate of noise, without error correction. For example, consider the case of random circuit sampling, and suppose one fixes a noise model, such as i.i.d. depolarizing noise at rate  $\gamma = \Theta(1)$  after each 2-qubit gate. We show that it is also  $\#P$ -hard to compute the output probabilities of the noisy circuits<sup>5</sup> on average to additive imprecision  $2^{-O(m \log m)}$ . This holds so long as the noise rate  $\gamma$  is below the error detection threshold.

**Theorem 2** (Simplified). *The same average-case hardness result as in Theorem 1 applies to noisy random quantum circuits, assuming the noise is gate-independent and can be error-detected.*

This second hardness result is even more interesting when we consider additional structure present in the uncorrected noise setting: the output probabilities of random circuits converge exponentially quickly to the uniform distribution. More specifically, it is widely conjectured that the output probabilities of such noisy random quantum circuits are  $2^{-O(m)}$  close to uniform [BIS<sup>+</sup>18]. In fact we rigorously prove this convergence in a toy model in which errors are interspersed with global Haar-random unitaries, though only slower exponential convergences have been shown for random circuit sampling [ABOIN96, GD18]. Therefore the output probabilities of the noisy circuits are very close to  $2^{-n}$ , up to some exponentially suppressed deviations.

---

<sup>5</sup>Here we mean the output probabilities averaged over the noise, i.e. the corresponding diagonal entries of the density matrix after evolution under the circuit and noise channels.

This rapid convergence to uniformity has several implications. First, central to Google’s supremacy claim was an assumption that these tiny deviations are in the direction of the ideal distribution<sup>6</sup>, and their experimental data was consistent with this claim [AAB<sup>+</sup>19]. Our result supports their intractability claim, in the sense that even these exponentially suppressed deviations from uniformity are hard to compute classically, though it remains open if *sampling* from these highly noisy distributions is classically intractable.

Second, this convergence to uniformity implies that Theorem 2 is optimal up to the log factor in the exponent! This is because it cannot possibly be hard to compute the output probabilities of high-noise circuits to additive imprecision  $2^{-o(m)}$  – because the trivial guess of  $2^{-n}$  is already within  $2^{-O(m)}$  of the noisy probability (on average). Therefore the robustness of our result to imprecision  $2^{-O(m \log m)}$  in the high noise case is essentially tight, as one cannot have a sublinear dependence on  $m$  in the exponent. Since our low-noise result is proven with very similar techniques, this means it might be very difficult to improve the robustness of the low-noise result as well. In this sense our two results are complementary to one another.

## 2 Our results and techniques

Showing both of these results require the development of new techniques to deal with increased amounts of imprecision and noise in the output probabilities. For simplicity we will first describe our results in the context of random circuit sampling. We will later show that the same techniques immediately give the analogous results for BosonSampling as well.

**Previous work on average-case hardness.** The core idea of prior average case hardness proofs, such as that of Lipton [Lip89], is to find an algebraic structure in the problem and use it as an error correcting code. Specifically, for many  $\#P$ -hard problems, the hard function can be written as a low degree polynomial of the input values. One can think of this low degree polynomial as a Reed-Solomon code, which allows one to infer the (possibly incorrect) value of a worst-case instance from the (typically correct) values of many average-case instances. In order to do so one must find a way of relating the value of a worst case instance to the value of several average-case instances in a way which preserves the low degree polynomial structure. For example, in the case of permanent of matrices over finite fields, given an arbitrary matrix  $A \in \mathbb{F}_q^{n \times n}$ , if one picks a random matrix  $B$  and lets  $A(\theta) = A + \theta B$  for  $\theta \in \mathbb{F}_q$ , one obtains a family of related instances  $\{A(\theta)\}_{\theta \in \mathbb{F}_q}$  which are each individually random, but correlated in a way which preserves this low-degree polynomial structure – namely,  $\text{Per}(A(\theta))$  is a low degree polynomial in  $\theta$ . This allows one to infer the value of the worst-case instance  $\text{Per}(A(0))$  from many values at  $\theta \neq 0$  by polynomial interpolation.

For random quantum circuits similar ideas can be made to work, but creating these correlated families of instances is not straightforward, due to the fact that the algebraic structure of quantum circuits is very different than that of permanents. There are currently two different ways to establish this polynomial structure [BFNV19, Mov20]. The original approach of [BFNV19] accomplished this by taking a family of correlated quantum circuits without a low degree polynomial structure, and imposing one by truncating a particular Taylor series. They used this structure to show the output probabilities of “truncated” circuits are hard to compute on average. While at first glance this average-case result might appear unrelated to the average-case hardness of random circuits – as

---

<sup>6</sup>In particular [AAB<sup>+</sup>19] argued for the hardness of their experiment under a highly simplified noise model of global depolarizing noise, since then the output probabilities are a convex combination of the true distribution with the uniform distribution. Our result generalizes this statement in a more realistic, per-gate noise model, and more generally to any error-detectable noise model.

the “input type” to this problem is different than the random circuit itself, due to the truncations – surprisingly Bouland *et al.* showed this average-case result is related to the original conjecture, in that robust hardness with respect to this truncated Taylor distribution is equivalent to robust hardness with respect to random circuits (see SI of [BFNV19]). Interestingly, while Bouland *et al.*’s result provided the first evidence in support of the approximate average-case supremacy conjecture for random circuits, their work left open whether or not the *true* output probabilities of these circuits are hard to compute on average, as pointed out in the SI of [BFNV19] and discussed in more detail by [Mov20, NPD<sup>+</sup>20]. More recently, Movassagh [Mov20] found a clever workaround to avoid the truncation issue. By simply choosing a different path between the worst and average case instances – known as the “Cayley path” – he showed it is hard to compute the exact output probabilities of random circuits, using a similar flow of argument to [BFNV19]. As Movassagh’s proof techniques are simpler, and the differences between these approaches could matter in high noise setting (see later discussion in Section 4.1 and Appendix A), we use Movassagh’s interpolation method throughout this work when discussing random circuit sampling results.

More generally, the strategy of establishing average-case hardness via polynomial interpolation is well studied in the context of problems over finite fields. For example, using the Berlekamp-Welch algorithm [WB86] and list decoding techniques, it has been shown that it is hard to compute the permanent of random matrices in  $\mathbb{F}_q^{n \times n}$  with constant [GS92] or even inverse polynomial [CPS99] probability, so long as the field size  $q$  is sufficiently large (see also [Gur06] for an overview). However, in the context of quantum supremacy, the relevant polynomials are all over the fields  $\mathbb{C}$  or  $\mathbb{R}$ . For example, the real-valued output probabilities of quantum circuits are low degree polynomials in the complex gate entries of the input circuit. This complicates the picture for proving average-case hardness, because first, one needs to consider that some of the values might be incorrect, and second on these correct values, there may be some imprecision in the value, i.e. the correct values may only be within  $\pm\delta$  of the true real value<sup>7</sup>.

## 2.1 The bottleneck in prior works

Existing average-case hardness results over the reals, such as those in [AA11, BFN19, Mov20], can tolerate only a tiny amount ( $\delta = 2^{-\text{poly}(n)}$ ) of such imprecision. This is for two reasons. First, many of the sophisticated tools used to show strong average-case hardness over finite fields, such as Berlekamp-Welch and list decoding techniques, do not immediately port over to the real case. Therefore one is forced to use straightforward Lagrangian interpolation/extrapolation techniques in the average case argument<sup>8</sup>, following the original arguments of Lipton [Lip89]. Second, it is well known that polynomial interpolation/extrapolation over the reals is exponentially sensitive to noise. This is essentially because there exist low degree polynomials (like Chebyshev polynomials) which are very close to one another in some interval but diverge rapidly outside of that interval. In fact very small errors in the value of average-case instances can blow up exponentially when extrapolated to the value of the worst-case instance.

Because of these two limitations, the strongest existing average-case hardness results for computing the output probabilities of random quantum circuits can only tolerate imprecision of order  $2^{-O(m^3)}$  in the output probabilities of  $n$ -qubit  $m$ -gate quantum circuits. More precisely, Movassagh [Mov20], building on prior arguments of Aaronson and Arkhipov [AA11] and Bouland *et al.*

<sup>7</sup>Indeed, some amount of imprecision is inevitable in any digital model of computation, where real numbers are approximated with poly-sized binary representations to  $\delta = 2^{-\text{poly}(n)}$  accuracy.

<sup>8</sup>We note that if one adopts a model of computation which admits exact representations of real numbers, then the Berlekamp-Welch algorithm does work, as noted in [BFNV19], but this is not known to work in a digital model of computation where one cannot exactly represent real numbers.

[BFNV19]<sup>9</sup>, showed the following:

**Theorem** ([Mov20], Simplified). *It is  $\#\mathsf{P}$ -hard to compute the output probability of random quantum circuits  $C$  up to additive error  $2^{-O(m^3)}$ , with probability  $\geq 1 - \frac{1}{\text{poly}(n)}$  over the choice of  $C$  with  $n$  qubits and  $m$  gates.*

The origin of this  $O(m^3)$  factor in the exponent is simple to understand. A well-known result due to Paturi states that if two degree- $d$  polynomials are  $\delta$ -close to one another in the interval  $[0, \Delta]$ , then their values at 1 can differ by at most  $\delta \cdot 2^{O(d\Delta^{-1})}$  [Pat92], which has an exponential dependence in the relative distance  $1/\Delta$  and can be saturated by Chebyshev polynomials in many parameter regimes. The degree of the relevant polynomial in the context of random quantum circuits is  $O(m)$ . The relative distance of interpolation  $1/\Delta$  in existing arguments is  $O(m^2)$ . Therefore, the Paturi argument shows that polynomial extrapolation over this interval incurs a  $2^{O(m^3)}$  blowup in the imprecision when going from the average-case instance to the worst-case instance, which means that a precision of at least  $2^{-O(m^3)}$  is required in the average-case output value to infer the worst-case value (which is bounded in  $[0, 1]$ ) to any nontrivial level of precision.

Let us be slightly more precise as to why the relative distance of interpolation is  $O(m^2)$ . In these average case arguments, given an arbitrary (worst-case) circuit  $C$ , one defines a (randomly chosen) family of circuits  $C(\theta)$  ( $\theta \in [0, 1]$ ) which essentially<sup>10</sup> has the following properties:

- $C(1) = C$ , that is, the worst-case circuit corresponds to  $\theta = 1$ .
- $C(0)$  is random, and  $C(\theta)$  is close to randomly distributed when  $\theta$  is small. More specifically,  $C(\theta)$  is  $O(m\theta)$  close to random in total variation distance for small values of  $\theta$ .
- The probability that the circuit  $C(\theta)$  outputs the all-zero outcome  $0^n$ , denoted as  $\mathsf{p}_0(C(\theta))$ , is a degree  $d = O(m)$  polynomial in  $\theta$ .

The basic idea of the average-case hardness arguments is to perform standard Lagrangian interpolation. Consider circuits  $C(\theta)$  for many small values of  $\theta > 0$  that are less than some maximum value  $\theta_{\max} = \Delta - \text{say } d + 1$  of them. One then uses the algorithm which works in the average case to compute  $\mathsf{p}_0(C(\theta))$  for these circuits. If the algorithm can correctly compute  $\mathsf{p}_0(C(\theta))$  within  $\pm\delta$  imprecision on all of them, then these points uniquely define a polynomial  $q(\theta)$  by Lagrangian interpolation which is close to the polynomial  $\mathsf{p}_0(C(\theta))$  at these evaluated points. By the Paturi error bound, the polynomial  $q$  is  $\delta \cdot 2^{O(m\Delta^{-1})}$  close to  $\mathsf{p}_0(C(\theta))$  in the interval  $[0, 1]$ <sup>11</sup>. Therefore one simply uses  $q(1)$  as an estimate for the worst case output probability  $\mathsf{p}_0(C(1)) = \mathsf{p}_0(C)$ , which is within additive error  $\delta \cdot 2^{O(m\Delta^{-1})}$ . So long as  $\delta \ll 2^{-O(m\Delta^{-1})}$ , this will be a good approximation to the desired value.

The reason that existing arguments set  $\Delta = O(1/m^2)$  – leading to a net imprecision tolerance of  $\delta = 2^{-O(m^3)}$  – is that Lagrangian interpolation requires *all* of the  $O(m)$  evaluated points to be correct (that is, up to imprecision  $\delta$ ). So each evaluation point needs to be correct with probability

<sup>9</sup>More precisely, [BFNV19] showed (using the techniques of Aaronson and Arkhipov [AA11]) that it is hard to compute the “truncated” circuit probabilities to accuracy  $2^{-\text{poly}(m)}$ , but did not quantify the polynomial in their argument as it depends on a number of unspecified parameter settings. Movassagh [Mov20] improved this to show hardness of computing the true output probabilities, and quantified the robustness to  $2^{-O(m^3)}$ .

<sup>10</sup>Technically speaking neither the techniques of [BFNV19] nor [Mov20] have all of these properties, but the authors found workarounds to this limitation.

<sup>11</sup>This is not entirely straightforward, because Paturi’s bound on error blowup is for polynomials which are uniformly close in some interval, and these polynomials are only close at particular points, but this turns out not to be an issue as we will describe later.

at least  $1 - O(1/m)$ , so that all points are correct with high probability by a union bound<sup>12</sup>. However, if one has an algorithm which works on random instances, it will only work with probability  $\geq 1 - O(m\Delta)$  on each evaluated points, because these points are not actually distributed randomly, but merely  $O(m\Delta)$  close to random, as given in the second property defined above. Therefore one must set  $\Delta = O(1/m^2)$  for this argument, to ensure that all of the evaluated points are correct with high probability.

## 2.2 Going beyond the imprecision bottleneck

One can easily see that the value of  $\Delta$ , as well as the degree of the base polynomial  $d = O(m)$ , are the controlling factors of the error in this argument. A natural question is whether or not either of these can be improved. Increasing the value of  $\Delta$  in the interpolation argument would cause some of the evaluated points to be incorrect, which means standard Lagrangian interpolation cannot be used. In the finite field case, it is well-known how to handle a constant fraction of incorrect points via the Berlekamp-Welch algorithm [WB86, GS92]. It is natural to ask if this algorithm, or a variant of it, could be made to work over the Reals. Critically, in order for such an algorithm to be relevant to our context, we would need the algorithm itself to *not blow up errors too badly* – or else the gains we make from choosing closer evaluation points could be swamped by the error blowup of the Real version of Berlekamp-Welch.

### 2.2.1 First step: robust Berlekamp-Welch over the Reals

Our first result is to fill in this gap, and show that a variant of Berlekamp-Welch argument still holds over  $\mathbb{R}$  or  $\mathbb{C}$ , without too much blowup in error. That is, suppose there exists a degree- $d$  polynomial  $p(x)$  over the reals, as well as  $\text{poly}(d)$  evaluation points  $\{x_i\}$  which are equally spaced over the interval  $[0, 1]$  and claimed values of the polynomial  $y_i$  at those points, of which only  $1 - \eta$  fraction are correct within additive error  $\pm\delta$ , where  $\eta < \frac{1}{4}$  is a constant. Then our results show that any other polynomial  $q(x)$  which also  $\delta$ -agrees with (possibly different)  $1 - \eta$  fraction of these evaluation points can diverge from  $p(x)$  by at most  $\delta \cdot 2^{O(d)}$  at all points in the interval.

Fortunately this error blowup happens to be less than that incurred by our interpolation argument, so is irrelevant to the asymptotics of our result. This statement is relatively easy to prove if all of the points the polynomials agree on are “bunched” together on one side of the interval, but extending this to the case where the correct/incorrect points are interspersed with one another requires a very careful argument. In particular our argument inductively shows that one can slowly “bunch” the correct and incorrect points together one by one without loss of generality, without blowing up the error. From there one can then appeal to the bunched case which follows from more standard arguments combining well-known statements about the maximality of Chebyshev polynomial blowups with Markov’s inequality for the maximum derivative of polynomials (see Section 5 for details).

The astute reader will notice that this argument is information theoretic in nature – that is, if one were to be handed a polynomial which agrees with a large constant fraction of the points, then that polynomial can’t diverge “too badly” from the true polynomial, but it provides no way of *finding* such a polynomial (which is part of the result of the Berlekamp-Welch algorithm). Interestingly, in the context of proving average-case  $\#P$  hardness, this is all that we need! That is to say, we only need a form of Berlekamp-Welch error correction with an NP decoder. This is because for the purposes of the complexity-theoretic arguments relevant to quantum supremacy, it suffices to

---

<sup>12</sup>These points are highly correlated with one another, so unfortunately a union bound is the best tool we know how to use here.

prove that computing the output probabilities of random circuits lies outside of the polynomial hierarchy<sup>13</sup>. Therefore using an  $\text{NP}$  decoder in our argument simply means we are showing these probabilities are  $\#\text{P}$  hard to compute under  $\text{BPP}^{\text{NP}}$  reductions, which is still outside of  $\text{PH}$ , unless  $\text{PH}$  collapses to a finite level. Whether an algorithmic version of this result holds over the Reals remains unclear, and is left as an intriguing open problem.

### 2.2.2 The next barrier at $2^{-O(m^2)}$ precision

With this in hand, we have now shown that it is  $\#\text{P}$ -hard (under  $\text{BPP}^{\text{NP}}$  reductions) to compute the output probabilities of random circuits to additive imprecision<sup>14</sup>  $2^{-O(m^2)}$ . We did so by decreasing the relative distance  $1/\Delta$  we needed to interpolate in our argument from  $O(m^2)$  to  $O(m)$  in relative terms, by allowing a constant fraction of the interpolation points to be incorrect. This proof proceeded by generalizing techniques developed for finite fields to the Reals, but with significant modifications to accommodate the imprecision of real values.

We now wish to push this result closer to achieving hardness to  $2^{-O(m)}$  imprecision. Once again the degree of our polynomial is  $O(m)$ , so unless we find a different polynomial to work with, the only way forward seems to be to increase the interpolation distance further – namely to set  $\Delta = O(1)$ . A natural question is if more sophisticated techniques developed for finite fields, such as list decoding, might help us to do this, by allowing us to handle a larger fraction of incorrect points, and hence handle a larger interpolation distance. Unfortunately these sort of ideas will not help us here by more than constant factors. The issue is that the distribution of circuits  $C(\theta)$  are only  $O(m\theta)$  close to random in total variation distance. Therefore, we cannot guarantee we can correctly evaluate *any* points beyond  $\Delta = \Theta(1/m)$  with any substantial probability<sup>15</sup>. In fact even proving a variant of list decoding over the Reals would only improve our result by a factor of two in the constant prefactor in the exponent. As a result, new ideas will be required to tackle this barrier which are not derived from analogues in the finite field case.

### 2.2.3 Second step: improved robustness of interpolation over the Reals

Therefore, our second step is to show improved robustness for real-valued polynomial interpolation. We will first show this using a novel error suppression technique based on variable rescaling which is specific to the Reals. We will then show that this technique proves that in prior robustness results, the Paturi bound was “overly pessimistic” as it was being applied in a parameter regime in which it is not tight! Therefore one can accurately interpolate degree  $d$  polynomials to  $2^{-O(d \log \Delta^{-1})}$  additive error, and moreover significantly simplify prior proofs of robustness. We provide a pedagogical explanation of this proof below.

To see this, we first note that a variable rescaling trick already shows improved robustness using the standard Paturi bound. The basic idea is to consider “rescaling” the variables in our problem by replacing our interpolation variable  $\theta$  with the variable  $x^k$  for some large integer  $k \in \mathbb{N}$ . We then perform our polynomial interpolation algorithm over the variable  $x$  instead of the variable  $\theta$ .

---

<sup>13</sup>This is because the whole point of these arguments is that the output probabilities of classical sampling algorithms can be approximated in the polynomial hierarchy - see Appendix A for details.

<sup>14</sup>We note that Movassagh [Mov20] independently claimed a similar result, but was missing the robust Berlekamp-Welch theorem necessary for it to be rigorously proven (as also noted in [ODMZ20]). Our robust Berlekamp-Welch argument (derived independently from Movassagh’s work) completes his proof of this claim, but with the same weakening to hardness under  $\text{BPP}^{\text{NP}}$  reductions.

<sup>15</sup>For example, if one considers setting  $\Delta = \Theta(1)$ , the total variation distance between the distribution of  $C(\theta)$  and random becomes  $1 - 1/\exp$ , and none of the evaluated points will necessarily be correct with non-negligible probability.

That is, we still consider using the average-case algorithm to compute the output probabilities of many circuits  $C(\theta)$  for many values of  $\theta < \Delta = O(1/m)$ . But when running the robust Berlekamp-Welch argument, instead of asking the NP oracle for the degree  $O(m)$  polynomial in  $\theta$  which  $\delta$ -approximates the correct values on  $1 - \eta$  fraction of the evaluating points, and using that to extrapolate our value at  $\theta = 1$ , we instead ask for the same, degree  $O(mk)$  polynomial in the variable  $x$ , and use that polynomial in  $x$  to extrapolate our value at  $x = 1$ .

At first glance this sounds strange – how could simple renaming variables improve the robustness of interpolation? And why would we be using higher degree polynomials, when they increase the error of extrapolation in general? What is going on here is that we are performing a continuous transformation of the interval  $[0, 1]$  over which we are interpolating. That is, consider the map  $\Phi : [0, 1] \rightarrow [0, 1]$  defined as  $\Phi(x) = x^k$ . This transformation is bijective, continuous, and preserves the endpoints of the interval. Critically, this transformation “stretches out” the part of the interval near 0, and compresses it near 1. Therefore, for a fixed value of  $\Delta = O(1/m)$ , the corresponding value of  $x_{\max}$  is substantially increased – it becomes  $1/m^{1/k}$ . In other words, the interpolation distance is much shorter for this new variable, which decreases the blowup in error from interpolation. Of course we pay a price for performing this transformation, as the degree of our polynomial in  $x$  is a factor of  $k$  larger than the degree of the polynomial in  $\theta$ . However this blowup due to an increased polynomial degree is less important than the savings from decreased extrapolation distance, and in net this results in a substantially reduced error in our extrapolation.

To see this more concretely, recall that the extrapolation error of a degree  $d$  polynomial interpolated over distance  $\Delta$  is  $\delta \cdot 2^{O(d\Delta^{-1})}$  [Pat92]. Our robust Berlekamp-Welch based algorithm performs interpolation with  $d = O(m)$  and  $1/\Delta = O(m)$ . After applying our variable rescaling, the degree becomes  $O(mk)$ , while  $\Delta$  becomes  $1/m^{1/k}$ . Therefore after this transformation this extrapolation error becomes only  $\delta \cdot 2^{k \cdot O(m) \cdot m^{1/k}}$ . By choosing  $k = \log m$  this becomes merely a blowup of  $2^{O(m \log m)}$ . Crucially the error  $\delta$  in the *output* probability of the quantum circuit is not affected by this rescaling of the *input* variables of the polynomial (namely  $\theta$ ) – and therefore the value of  $\delta$  is the same in the rescaled polynomial. This therefore allows us to show our main result:

**Theorem 1** (Restated, simplified). *For any constant  $\eta < 1/4$ , it is  $\#P$ -hard under a  $\text{BPP}^{\text{NP}}$  reduction to compute the output probability of random quantum circuits  $C$  up to additive error  $2^{-O(m \log m)}$ , with probability at least  $1 - \eta$  over the choice of  $C$  with  $m$  gates.*

It also shows an analogous hardness result in the case of BosonSampling, as we will describe in Section 4 (see Corollary 1).

A deeper examination of the above results reveals that this argument proves the Paturi bound itself it not tight in the parameter regime in which it was being applied to quantum supremacy arguments. The standard justification of the Paturi bound – that the error of a degree  $d$  polynomial interpolated over distance  $1/\Delta$  blows up by  $2^{O(d\Delta^{-1})}$  – is that Chebyshev polynomials saturate the bound. This is true in the regime of  $\Delta = \Theta(1)$  – essentially because Chebyshev polynomials have coefficients of order  $2^{O(d)}$ , and the individual monomials happen to “cancel out” in the unit interval, but this carefully orchestrated cancellation quickly falls apart as one leaves the unit interval, and the value of the polynomial approaches  $2^{O(d)}$ . This behavior is what the Paturi bound tightly captures.

However our variable rescaling trick illustrates that in the regime of asymptotically decreasing  $\Delta$ , this bound cannot possibly be tight, because nothing in our argument invalidated Chebyshev polynomials<sup>16</sup>. In other words, our variable rescaling argument proved that Chebyshev polynomials

---

<sup>16</sup>That is, if  $T_d$  is the  $d$ th Chebyshev polynomial, our errors could uniquely specify  $T_d(x^k)$  as our error polynomial, as all we have done is replaced  $x$  with  $x^k$ .

cannot blow up faster than  $2^{O(d \log d)}$  when extrapolated to distance  $1/\Delta = O(d)$ . In retrospect this has a much simpler proof, and it follows because Chebyshev polynomials are after all only polynomials, and not exponential functions! That is, while Chebyshev polynomials appear to grow exponentially as they “leap out” of the small interval, this exponential growth cannot be sustained over longer ranges – because the function can only grow polynomially quickly in the large distance limit. To put this more formally, if one has a degree- $d$  polynomial  $p(x) = \sum_{i=0}^d c_i x^i$ , then the value of the polynomial at  $\Delta^{-1}$  must obey  $|p(\Delta^{-1})| \leq \sum_i |c_i| \Delta^{-d}$ . So if  $\Delta^{-1} = d^c$  for some constant  $c > 0$ , and we know  $c_i = 2^{O(d)}$ , we immediately get that Chebyshev polynomials – and hence our interpolation error – is bounded by  $d 2^{O(d)} (d^c)^d = 2^{O(d \log d)}$ . That is to say, over longer extrapolation distances, Chebyshev polynomials (and in fact any polynomials with bounded coefficients, which is guaranteed in our settings) only have a  $2^{O(d \log d)}$  blowup.

Armed with this insight, we provide a significantly simpler proof of robustness of polynomial interpolation to  $2^{-O(d \log d)}$  additive imprecision starting from first principles. This proof is both simpler than our variable rescaling trick, and results in better constants in the exponent. Moreover, it substantially simplifies prior robustness of interpolation results, as once it is combined with our robust Berlekamp-Welch arguments, it removes the need to invoke heavy-duty techniques of Rakhmanov [Rak07] in the technical portion of the result and instead only requires applying the much simpler Markov’s inequality<sup>17</sup>. Instead our proof is self-contained and based on fundamental facts about real-valued polynomials.

In the end our final result for random circuit sampling with  $m$  2-qubit gates shows robustness to additive imprecision  $\delta = e^{-16m \log m - O(m)}$ . For an  $n$ -photon  $m = O(n^c)$ -mode BosonSampling experiment, our result shows robustness to additive imprecision  $\delta = e^{-(c+4)n \log n - O(n)}$ . In contrast the prior leading constants in the exponent of our results for RCS and BosonSampling were 147 and  $(c+70)$ , respectively, when using the variable rescaling error suppression technique (see Section 5.1 for a detailed discussion).

### 2.3 Extending these arguments to the high-noise regime

Our second result is to extend these average-case hardness arguments to the high noise regime. That is, suppose we fix a noise model, such as i.i.d. depolarizing noise at rate  $\gamma$ . We show that it is also  $\#P$ -hard (under  $\mathsf{BPP}^{\mathsf{NP}}$  reductions) to estimate the output probabilities of noisy random circuits under this noise model to additive error  $2^{-O(m \log m)}$  as well. This holds so long as the noise rate  $\gamma$  is below the error *detection* (not correction!) threshold. More generally, our result holds for any stochastic noise model which admits error detection, and which is gate-independent.

There are two key ideas which enable this strengthening of our result. The first is to note that our worst to average-case reduction techniques apply essentially immediately to the noisy case! This is because for stochastic noise models like i.i.d. depolarizing noise, the output probability of the noisy circuit is a convex combination of the output probabilities of many different noiseless circuits. To be more precise, recall that depolarizing noise with Pauli noise rate  $\gamma$  after a two-qubit gate is equivalent to applying one of the 15 different nontrivial Pauli operators (such as  $I \otimes X$  or  $Y \otimes Z$ ) with probability  $\gamma/15$  after that gate, or else applying the identity operator with probability  $1 - \gamma$ . Therefore, applying depolarizing noise after every gate of an  $m$ -gate circuit is equivalent to taking

<sup>17</sup>More formally, the previous arguments based on Lagrangian interpolation used a union bound to argue all evaluated points were correct, then invoked Rakhmanov’s argument to argue that the polynomial was uniformly bounded in that interval, then invoked Paturi’s bound to show that the polynomial remained bounded when extrapolated. Our new proof first argues (using our robust Berlekamp-Welch argument) that the polynomial is uniformly bounded, then applies a simple argument to obtain a bound on the coefficients of the difference polynomial, then directly bounds the error at extrapolation, without the need to invoke Rakhmanov’s results.

the original circuit  $C$ , and inserting into the circuit one of the  $16^m$  different “noise patterns”  $\xi$  that could have occurred in the circuit, with appropriately weighted probabilities – call these circuits  $C_\xi$ . The output probability of the noisy circuit is the appropriately weighted convex combination of the output probabilities of  $C_\xi$ . In other words, if  $\mathsf{po}(C, \mathcal{N})$  is the probability that the noisy circuit outputs  $0^n$  under a noise model  $\mathcal{N}$ , then we have that

$$\mathsf{po}(C, \mathcal{N}) = \mathbb{E}_{\xi \sim \mathcal{T}_{\mathcal{N}}} [\mathsf{po}(C_\xi)] \quad (1)$$

where  $\mathcal{T}_{\mathcal{N}}$  is the distribution over error patterns arising from the particular noise model  $\mathcal{N}$ . This is true in general for stochastic noise models.

The key point to note is that our average-case hardness arguments are invariant under taking convex combinations of circuits. This is because the key fact we use in our reduction is that, for some appropriately defined family of circuits  $C(\theta)$ ,  $\mathsf{po}(C(\theta))$  is a low degree polynomial in  $\theta$ . The set of low degree polynomials is convex. Therefore, if one simply defines  $C_\xi(\theta)$  to be the same circuits  $C(\theta)$  with the noise pattern  $\xi$  inserted into the circuit, then we have that

$$\mathsf{po}(C(\theta), \mathcal{N}) = \mathbb{E}_{\xi \sim \mathcal{T}_{\mathcal{N}}} [\mathsf{po}(C_\xi(\theta))] \quad (2)$$

is also a low degree polynomial in  $\theta$ . As a result the exact same worst-to-average case reduction arguments work in the noisy setting (see Section 4.2 and 4.3 for more details). Interestingly, for this to work there is no need to know the noise rate  $\gamma$  or noise model to apply this argument – because the low degree polynomial structure is preserved by *any* noise model that is stochastic and gate-independent.

We have therefore established that average-case noisy output probabilities are as hard to estimate as worst-case ones. The second step required is to establish that these noisy worst-case probabilities are indeed  $\#P$ -hard to compute<sup>18</sup>. To do so we follow the observations first developed by Fujii [Fuj16]. In particular, Fujii observed that no polynomial time algorithm can exactly<sup>19</sup> sample from the output distribution of noisy quantum circuits in the worst case unless  $\mathsf{PH}$  collapses, so long as the noise rate  $\gamma$  is below the error detection threshold. The basic idea is that quantum error detection is highly efficient. With polynomial overhead, one can create a quantum error detecting code which is exponentially sensitive to errors. That is, if no error is detected, then one is certain with probability  $1 - 1/\exp$  that no error indeed occurred. Using this, Fujii showed that one can post-select the “no error detected” event so that the post-selected distribution is exponentially close to the ideal output distribution. By standard post-selection arguments along the lines of [Aar05, BJS11] this suffices to show worst-case hardness. A modification of this argument can be used to show that computing these output probabilities is essentially  $\#P$ -hard. This establishes the worst-case hardness of noisy circuits, which allows us to show the following theorem:

**Theorem 2** (Restated, simplified). *The same average-case hardness result as in Theorem 1 applies to noisy random quantum circuits, assuming the noise is gate-independent and can be error-detected.*

Note that there is one restriction which applies to this theorem which is not present in the statement of Theorem 1. Namely, this theorem assumes that the family of random circuits considered is deep enough to encode a quantum error detecting code which works against the chosen noise model and noise rate. This is because the worst-case circuits used in the argument of Theorem 2 are error detecting circuits, while constant depth circuits suffice for Theorem 1. Therefore our result essentially says that the output probability of noisy random circuits are as hard to compute as the *largest error detected circuit one could have fit* into that particular random circuit architecture.

<sup>18</sup>More formally, we prove  $\mathsf{coC=P}$  hardness, which is hard enough to suffice for quantum supremacy arguments as  $\mathsf{coC=P}$  is outside of the polynomial hierarchy unless  $\mathsf{PH}$  collapses.

<sup>19</sup>That is, up to constant multiplicative error on each of the exponentially small output probabilities.

### 3 Discussion

In this work we have sought to resolve a fundamental disconnect between the theoretical basis of quantum supremacy – based on very *low noise* random quantum circuits [TD02, BJS11, AA11, BMS16, BFN19, Mov20] – and the highly noisy experimental implementations performed so far [AAB<sup>+</sup>19, ZWD<sup>+</sup>20]. Namely, the initial theoretical attempts at proving “quantum supremacy” for sampling from random quantum circuits sought to do so assuming one can maintain a large constant fidelity as the system size grows – which would require error correction to achieve scalably. Even establishing this is classically intractable remains a major theoretical challenge, due to the sensitivity of the average-case hardness of computing results to small amounts of imprecision – a limitation that we made substantial progress on in our first result (Theorem 1) but did not resolve. However the initial experimental attempts at demonstrating supremacy showed a fidelity that *decreases* exponentially with system size, because quantum error correction is not being performed. As these sampling experiments become exponentially close to uniform asymptotically [ABOIN96, GD18], one cannot argue that these experiments are hard to reproduce in an approximate sense.

Consequently, if there is any hope of finding a hard quantum signal present in these near-term experiments, it must be present in the exponentially diminishing deviations from the uniform distribution. A natural question is whether or not these small deviations from uniformity are hard to reproduce classically. Our second main result (Theorem 2) provides evidence in this direction, by extending the prior hardness of exact computation results for random quantum circuits to the noisy setting. That is, we prove it is hard to exactly *compute* the output probability of a noisy random quantum circuit, to finer precision than the rate of convergence to uniform. This can be interpreted as proving the existence of a quantum signal that survives the noise, thereby formalizing an intuition that underlies the optimism that the recent Google experiment is solving a hard problem.

A natural hope is that one might be able to use these results to show that one cannot *exactly* reproduce the results of the Google experiment in a *sampling* sense. That is, can one argue that no classical algorithm could sample from the output distribution of noisy random quantum circuits nearly exactly? Unfortunately we know we cannot do this using current techniques [AA11]. The reason is that these techniques reduce hardness of exact (or approximate) sampling to the hardness of *approximately* computing output probabilities using Stockmeyer’s algorithm [Sto85], and the amount of imprecision required to be tolerated ( $O(2^{-n})$ ) is intrinsic in the argument, even in the exact sampling case<sup>20</sup>. As previously discussed, the output probabilities of noisy random circuits cannot be hard to compute to this degree of imprecision, due to the fact the noisy distributions are exponentially close to uniform. While [AC17, AG19] took steps to try to address the simulation of such noisy experiments via a different angle, their approach assumes the fidelity remains a small constant as the system size scales, and requires making non-standard assumptions. Therefore, we believe fundamentally new ideas will be required – beyond the Stockmeyer approach – if one wishes to address the hardness of exactly simulating highly noisy circuits with decreasing fidelities.

Our main results also offer an interesting new message about the low-noise setting: any potential means for proving the hardness of sampling in the *low noise* setting cannot be extended to hold in the *high noise* setting. Our barrier holds in the setting of random circuits with superconstant depth<sup>21</sup>. This offers a significant barrier for improving the status quo of quantum supremacy

---

<sup>20</sup>Technically speaking, the imprecision incurred by Stockmeyer’s algorithm is multiplicative, and so depends on the output probability being estimated, which for random circuits is  $O(2^{-n})$  with very high probability due to known “anti-concentration” results [HL09, BHH16, HM18, BCG20, DHJB20].

<sup>21</sup>More formally, the conjectured rate of convergence of noisy circuits to the uniform distribution is  $O(e^{-dn\gamma})$  where  $n$  is the number of qubits and  $d$  is the depth of the circuit, and  $\gamma$  is the depolarizing noise rate per qubit. The robustness required for proving hardness of sampling in the low noise setting is  $O(2^{-n})$ . Therefore, for any circuit

results: the only chance of resolving the hardness of sampling problem (using the standard and only approach we have – via a reduction to approximate computing using Stockmeyer’s theorem) is by inventing new techniques that are explicitly *not resilient* to uncorrected noise! This immediately rules out any “convex method” – i.e. any method where the underlying error correcting code structure is invariant under convex combinations – including polynomial interpolation, which we prove in Theorem 2 is noise resilient.

Our results are complementary, but distinct from past barrier results, starting with that of Aaronson and Arkhipov [AA11], who first realized that polynomial interpolation cannot close this gap if the ensemble of circuits anticoncentrates (See Appendix B for more details). While there are subtle technical differences between our barriers<sup>22</sup>, they arise for similar reasons – namely the set of low degree polynomials is convex. Later, Aaronson and Chen [AC17] proved an oracle separation relative to which approximate quantum sampling is classically easy, but the PH remains infinite. While this rules out any relativizing technique for resolving the hardness of approximate sampling, their result is very sensitive to the distinction between approximate and exact sampling – indeed, worst-case hardness of exact sampling results do exist and relativize [BJS11]! Therefore it is unclear how to extend their barrier result to our setting in which we are interested in the hardness of exact sampling from the output distribution of high-noise and/or ideal random quantum circuits.

Most recently, a barrier result for hardness of sampling was shown by Napp *et al.* [NPD<sup>+</sup>20] who give an approximation algorithm for computing the output probabilities of very shallow (i.e., fixed constant) depth 2D random quantum circuits. This is interesting, because one can use the methods of Bouland *et al.* and Movassagh to prove the average case hardness of *exactly* computing the output probabilities of the same ensemble of circuits. The implication is that any technique that establishes hardness of computing output probabilities of random quantum circuits to within additive error  $O(2^{-n})$  must not work for very shallow depth quantum circuits. This is a limitation, since current techniques that use polynomial interpolation are not sensitive to depth. This result is nicely complementary to our barrier, and in fact we can use it to conclude that any technique that extends the hardness of computation results to hardness of sampling results must be sensitive to both depth *and* noise.

Interestingly the noise barrier identified in this paper, and the depth barrier identified by Napp *et al.* [NPD<sup>+</sup>20], rely on specific properties of random quantum circuits. We leave it as an open problem if analogous barriers apply to the BosonSampling case<sup>23</sup>. Of course it is already known that polynomial interpolation will not prove hardness of sampling for BosonSampling [AA11] – in the case of  $m = n^2$  modes this barrier shows polynomial interpolation cannot prove robust hardness beyond  $O(e^{-3n \log n})$  additive error. This leaves open the possibility of improving the exponent in our robustness result for BosonSampling from 6 to 3 using current techniques, but not beyond. However, the fact that our results for BosonSampling are closer to achieving the desired robustness – due to structural differences in the algebraic properties of BosonSampling vs. RCS – indicates that BosonSampling could have certain advantages for proving “quantum supremacy” in the low-noise regime going forward.

---

with  $d > 1/\gamma_0$ , where  $\gamma_0$  is the error detection threshold, any noise-agnostic proof technique cannot prove sufficient robustness to prove hardness of sampling in the low noise setting. Note the depth  $d$  must be sufficiently large to perform error detection as well.

<sup>22</sup>In particular our barrier does not require anticoncentration, but instead requires a distribution over random circuits which converges to uniformity under depolarizing noise. Both properties are related to the scrambling nature of random circuits, but it is not *a priori* clear they are identical.

<sup>23</sup>We note this might be possible using the techniques of e.g. [GPRS19, QBQGP20, ONFJ21] for simulating noisy BosonSampling experiments, though there are many details to be worked out.

## 4 Average-case hardness

In this section we develop the average-case hardness of computing the output probability of random quantum circuits. We first establish improved results in the noiseless setting, and then show that our techniques can be generalized to the noisy case. An important ingredient of our proof is a robust polynomial interpolation technique, which is independently developed in Section 5.

### 4.1 Worst-to-average-case reduction

Our first result is to establish a worst-to-average-case reduction for computing the output probability of random quantum circuits in the noiseless setting. That is, if there exists an algorithm that can approximate the output probability of *most* random quantum circuits over an architecture, then there exists an algorithm that can approximate the output probability of *every* quantum circuit over the same architecture. Our average-case hardness result tolerates a constant failure probability (i.e. it is hard to compute the output probability for a large constant fraction of random circuits), and is robust up to a small additive error  $2^{-O(m \log m)}$  for random circuits with  $m$  gates. Our result is therefore an improvement over previous results [BFNV19, Mov20] in both aspects.

The main idea of the worst-to-average-case reduction works as follows. Consider a circuit architecture over which computing the output probability up to an exponentially small additive error is  $\#P$ -hard in the worst case. Then, to prove average-case hardness, it suffices to construct a reduction that belongs to some finite level of the Polynomial Hierarchy. That is, given an arbitrary circuit  $C_0$ , our goal is to compute  $\mathbf{p}_0(C_0)$  using a procedure in the Polynomial Hierarchy, with access to an oracle  $\mathcal{O}$  that can compute  $\mathbf{p}_0(C)$  for a large fraction of random circuits  $C$ . For a circuit architecture  $\mathcal{A}$ , let  $\mathcal{H}_{\mathcal{A}}$  be the distribution over circuits such that each gate is independently drawn from the Haar measure. Let  $\{G_i\}_{i=1,\dots,m}$  be the quantum gates in  $C_0$ . We create a new circuit  $C_1$  by applying a “one-time pad” to  $C_0$ , where each gate  $G_i$  is replaced by

$$G_i \rightarrow H_i G_i, \tag{3}$$

where  $\{H_i\}$  is independently drawn from the Haar measure over the unitary group and has the same dimension as  $G_i$ . By the invariance of Haar measure,  $C_1$  is distributed the same as  $\mathcal{H}_{\mathcal{A}}$ . Therefore  $\{H_i\}$  can be understood as the “random seed” used for the one-time pad.

By definition,  $\mathcal{O}$  can compute an accurate approximation of  $\mathbf{p}_0(C_1)$  with high probability. However, this number alone does not contain any information on our desired quantity  $\mathbf{p}_0(C_0)$ . The main insight that allows us to correlate average-case solutions to the worst-case quantity, which was originally developed to show the average-case hardness for permanents [Lip89, GLR<sup>+</sup>91], is to create many random instances  $\{C_i\}$  by the following procedure: first sample a “random seed”  $\{H_i\}$ , then apply small and different perturbations to the random seed, and then apply each perturbed random seed to  $C_0$ . As a result, while each random instance is marginally distributed approximately according to  $\mathcal{H}_{\mathcal{A}}$ , they are close to each other and are correlated in a way that reveals the worst-case quantity  $\mathbf{p}_0(C_0)$ .

More specifically, suppose there is a way of perturbing the circuit  $C_1$  into a new circuit  $C(\theta)$  ( $\theta \in [0, 1]$ ), such that  $C(\theta) \approx C_1$  when  $\theta \ll 1$ , and  $C(1) = C_0$ . Moreover, suppose  $\mathbf{p}_0(C(\theta))$  for different values of  $\theta$  are correlated in a way such that  $\mathbf{p}_0(C(1))$  can be inferred from  $\mathbf{p}_0(C(\theta))$  for small values of  $\theta$ . Then, to compute  $\mathbf{p}_0(C_0)$ , it suffices to query  $\mathcal{O}$  with  $C(\theta_i)$  for many small  $\theta_i$ s.

One way to develop such a procedure is by perturbing the random seeds  $\{H_i\}$  used in the construction of  $C_1$ , which we define as follows.

**Definition 1** ( $\theta$ -perturbed random circuit distribution). *Suppose there is a transform on unitary matrices  $H \mapsto H(\theta)$  parameterized by a real number  $\theta \in [0, 1]$ , that satisfies  $H(0) = H$  is unchanged,*

and  $H(1) = I$  is the identity matrix. For any circuit  $C_0$  with gates  $\{G_i\}_{i=1,\dots,m}$  and a random seed  $\{H_i\}_{i=1,\dots,m}$ , the circuit  $C(\theta)$  is defined by replacing  $G_i$  with

$$G_i \mapsto H_i(\theta)G_i. \quad (4)$$

Denote the distribution induced by  $C(\theta)$  as  $\mathcal{H}_{\mathcal{A},\theta}$ . Note that  $\mathcal{H}_{\mathcal{A},0} = \mathcal{H}_{\mathcal{A}}$  and  $\mathcal{H}_{\mathcal{A},1} = \mathbf{1}_{C_0}$ .

As discussed above, in order for such a perturbation to be useful for the worst-to-average-case reduction, we require that the following (informal) properties are satisfied.

**Properties of the perturbed circuit distribution (informal).**

1. When  $\theta$  is small,  $\mathcal{H}_{\mathcal{A},\theta}$  is close to  $\mathcal{H}_{\mathcal{A}}$  in total variation distance. Therefore, when  $\mathcal{O}$  is given an input  $C(\theta) \sim \mathcal{H}_{\mathcal{A},\theta}$  from the perturbed distribution, it is guaranteed to return a correct approximation of  $\mathbf{p}_0(C(\theta))$  with high success probability.
2. An approximation of  $\mathbf{p}_0(C(1))$  can be inferred from approximations of  $\mathbf{p}_0(C(\theta))$  for small values of  $\theta$  with a procedure in the Polynomial Hierarchy.

Finding such a good perturbation method is non-trivial. Two proposals were developed in previous work in the context of noiseless circuits. In [BFNV19], they considered the truncated Taylor series of  $He^{-\theta \log H}$ , which is given by

$$H(\theta) = H \cdot \left( \sum_{k=0}^K \frac{(-\theta \log H)^k}{k!} \right), \quad (5)$$

and showed that  $\mathbf{p}_0(C(\theta))$  is a degree  $O(mK)$  polynomial in  $\theta$ . Property 2 is then satisfied by using polynomial interpolation techniques. In this approach,  $H(\theta)$  as defined by Eq. (5) is not unitary, and the resulting average-case hardness is for a circuit family that has a small deviation from  $\mathcal{H}_{\mathcal{A}}$  given by the truncation error of the Taylor series. While this truncation error does not affect the applicability of this approach to address hardness of approximate sampling, it becomes evident when applying to noisy circuits, as the requirement for accuracy is much higher (see appendix A for a detailed discussion).

Later, Movassagh [Mov20] developed a new perturbation method, called Cayley transform, that is able to stay within the unitary path.

**Definition 2** (Cayley transform [Mov20]). *The Cayley transform of a unitary matrix  $H$  parameterized by  $\alpha \in [0, 1]$  is a unitary matrix defined as*

$$H(\alpha) := \frac{\alpha I + (2 - \alpha)H}{(2 - \alpha)I + \alpha H}, \quad (6)$$

where  $A/B$  means  $A \cdot B^{-1}$ . Consider the diagonalization of a unitary matrix  $H = \sum_j e^{i\varphi_j} |\psi_j\rangle\langle\psi_j|$ , the following is an equivalent form of the Cayley transform,

$$H(\alpha) = \sum_j \frac{1 + i(1 - \alpha) \tan \frac{\varphi_j}{2}}{1 - i(1 - \alpha) \tan \frac{\varphi_j}{2}} |\psi_j\rangle\langle\psi_j|. \quad (7)$$

Note that  $H(0) = H$  and  $H(1) = I$ .

Note that we are using slightly different notations from [Mov20]. A self-consistent presentation of the properties of the Cayley transform is given in Appendix E.

While not having an evident polynomial structure as in the construction of [BFNV19], in [Mov20] they applied Cayley transform to the gates  $\{H_i\}$  parameterized by  $\alpha = \theta$ , and showed that the resulting  $\mathbf{p}_0(C(\theta))$  is a degree  $(O(m), O(m))$  rational function in  $\theta$ . A similar polynomial interpolation technique can be applied to infer  $\mathbf{p}_0(C(1))$  from  $\mathbf{p}_0(C(\theta))$  with small values of  $\theta$ .

In particular [Mov20], building on [BFNV19], obtained the following average-case hardness results (also see Appendix A).

**Theorem 3** ([BFNV19, Mov20]). *Let  $\mathcal{A}$  be a circuit architecture so that computing  $\mathbf{p}_0(C)$  to within additive error  $2^{-O(m)}$  is  $\#\mathsf{P}$ -hard in the worst case. The following results hold:*

1. *It is  $\#\mathsf{P}$ -hard to compute  $\mathbf{p}_0(C)$  exactly on input  $C \sim \mathcal{H}_{\mathcal{A}}$ , with success probability at least  $1 - \eta$  over the choice of  $C$ , for any constant  $\eta < \frac{1}{4}$ .*
2. *It is  $\#\mathsf{P}$ -hard to compute  $\mathbf{p}_0(C)$  up to an additive imprecision  $2^{-O(m^3)}$  on input  $C \sim \mathcal{H}_{\mathcal{A}}$ , with success probability at least  $1 - O(1)/m$  over the choice of  $C$ .*

We improve these results by proving an average-case hardness that tolerates a larger additive imprecision  $2^{-O(m \log m)}$ , while only requiring a constant success probability over the choice of random circuits. These improvements follows from two new techniques: the first is a robust polynomial interpolation argument which tolerates a constant failure probability, and the second is an improved error bound for long distance polynomial extrapolation. Both results are developed in Section 5 and are used as black boxes here.

To establish our reduction, we first consider the  $\theta$ -perturbed random circuit distribution instantiated by the Cayley transform.

**Definition 3** ( $\theta$ -Cayley perturbed random circuit distribution). *For any circuit  $C_0$  with gates  $\{G_i\}_{i=1,\dots,m}$  and a random seed  $\{H_i\}_{i=1,\dots,m}$ , the circuit  $C(\theta)$  is defined by replacing  $G_i$  with*

$$G_i \mapsto H_i(\theta)G_i, \tag{8}$$

where  $H_i(\theta)$  denotes the Cayley transform of  $H_i$  parameterized by  $\theta$ . Denote the distribution induced by  $C(\theta)$  as  $\mathcal{H}_{\mathcal{A},\theta}$ . Note that  $\mathcal{H}_{\mathcal{A},0} = \mathcal{H}_{\mathcal{A}}$  and  $\mathcal{H}_{\mathcal{A},1} = \mathbf{1}_{C_0}$ .

As previously shown by [Mov20], the output probability  $\mathbf{p}_0(C(\theta))$  is a low-degree rational function in  $\theta$ .

**Lemma 1** ([Mov20]). *For any circuit  $C_0$ , let  $C(\theta)$  be a circuit from the  $\theta$ -Cayley perturbed random circuit distribution as in Definition 3. Then  $\mathbf{p}_0(C(\theta))$  is a degree  $(O(m), O(m))$  rational function in  $\theta$ .*

*Proof.* Here we give a self-consistent proof as the details are useful for our later developments.

First, notice that the output amplitude can be written as the Feynman path integral,

$$\langle 0^n | C(\theta) | 0^n \rangle = \sum_{y_1, \dots, y_{m-1} \in \{0,1\}^n} \prod_{j=1}^m \langle y_j | (H_j(\theta)G_j \otimes I_{\text{else}}) | y_{j-1} \rangle, \tag{9}$$

where we have  $y_0 = y_m = 0^n$ ,  $H_j(\theta)G_j$  is a local gate and  $I_{\text{else}}$  denotes identity on all the other qubits. Let the diagonalization of  $H_j$  be  $H_j = \sum_l e^{i\varphi_{jl}} |\psi_{jl}\rangle\langle\psi_{jl}|$ . Then, consider an individual term

in the above sum,

$$\begin{aligned}
\prod_{j=1}^m \langle y_j | H_j(\theta) G_j | y_{j-1} \rangle &= \prod_{j=1}^m \langle y_j | \sum_l \frac{(1 + i(1-\theta) \tan \frac{\varphi_{jl}}{2}) |\psi_{jl}\rangle \langle \psi_{jl}|}{1 - i(1-\theta) \tan \frac{\varphi_{jl}}{2}} G_j | y_{j-1} \rangle \\
&= \prod_{j=1}^m \langle y_j | \frac{\sum_l (1 + i(1-\theta) \tan \frac{\varphi_{jl}}{2}) |\psi_{jl}\rangle \langle \psi_{jl}| \prod_{t \neq l} (1 - i(1-\theta) \tan \frac{\varphi_{jt}}{2})}{\prod_l (1 - i(1-\theta) \tan \frac{\varphi_{jl}}{2})} G_j | y_{j-1} \rangle.
\end{aligned} \tag{10}$$

Let

$$\begin{aligned}
Q_0(\theta) &:= \prod_{j=1}^m \prod_l \left( 1 - i(1-\theta) \tan \frac{\varphi_{jl}}{2} \right), \\
Q(\theta) &:= |Q_0(\theta)|^2.
\end{aligned} \tag{11}$$

Then Eq. (10) is a degree  $(O(m), O(m))$  rational function in  $\theta$  with  $Q_0(\theta)$  being the denominator. Furthermore, notice that  $Q_0(\theta)$  does not depend on the Feynman path  $\{y_j\}$ , and each term in the sum in Eq. (9) shares the same denominator  $Q_0(\theta)$ . Therefore,  $\langle 0^n | C(\theta) | 0^n \rangle$  is a degree  $(O(m), O(m))$  rational function in  $\theta$ , and so is  $\mathbf{p}_0(C(\theta)) = |\langle 0^n | C(\theta) | 0^n \rangle|^2$ .  $Q(\theta)$  is then the denominator of  $\mathbf{p}_0(C(\theta))$ . Finally, note that if  $C_0$  only consists of 2-qubit gates, then  $\mathbf{p}_0(C(\theta))$  is a degree  $(8m, 8m)$  rational function in  $\theta$ .  $\square$

Lemma 1 formally supports Property 2 of the  $\theta$ -Cayley perturbed random circuit distribution. Intuitively, this low-degree rational function structure allows us to apply polynomial interpolation techniques to obtain an approximation to the worst-case quantity, even though the worst-case point ( $\theta = 1$ ) is far from the average-case data points ( $0 < \theta \ll 1$ ). Details of this interpolation technique are presented in the proof of our main result (Theorem 4) and in Section 5. Before presenting the main result, it remains to establish Property 1 of the  $\theta$ -Cayley perturbed random circuit distribution, which is given in the following lemma.

**Lemma 2** ([Mov20]). *Let  $\mathcal{H}_{\mathcal{A},\theta}$  be the  $\theta$ -Cayley perturbed random circuit distribution as in Definition 3, and  $\mathcal{H}_{\mathcal{A}}$  be the distribution of Haar random circuits over  $\mathcal{A}$ . Then we have*

$$D_{\text{TV}}(\mathcal{H}_{\mathcal{A},\theta}, \mathcal{H}_{\mathcal{A}}) = O(m\theta), \tag{12}$$

where  $D_{\text{TV}}(\cdot, \cdot)$  denotes the total variation distance between probability distributions and  $m$  is the number of gates in  $\mathcal{A}$ .

*Proof.* This total variation distance can be bounded by considering the distribution of each individual gates. In [Mov20], it was shown that the total variation distance between the Cayley transformed random unitary  $H(\theta)$  and Haar random unitary is  $O(\theta)$  (also see Lemma 12 in Appendix E). Therefore by additivity of the total variation distance we have  $D_{\text{TV}}(\mathcal{H}_{\mathcal{A},\theta}, \mathcal{H}_{\mathcal{A}}) = O(m\theta)$ .  $\square$

Having established Property 1 and 2, we are now ready to state and prove our first result on the average-case hardness of random circuits.

**Theorem 4.** *Let  $\mathcal{A}$  be a circuit architecture so that computing  $\mathbf{p}_0(C)$  to within additive error  $2^{-O(m)}$  is  $\#P$ -hard in the worst case. Then the following problem is  $\#P$ -hard under a  $\text{BPP}^{\text{NP}}$  reduction: for any constant  $\eta < \frac{1}{4}$ , on input a random circuit  $C \sim \mathcal{H}_{\mathcal{A}}$  with  $m$  gates, compute the output probability  $\mathbf{p}_0(C)$  up to additive error  $\delta = \exp(-O(m \log m))$ , with probability at least  $1 - \eta$  over the choice of  $C$ .*

**Remark 1.** Another way of stating Theorem 4 is the following: suppose there exists an algorithm that belongs to some finite level of PH that, on input a random circuit  $C \sim \mathcal{H}_A$  with  $m$  gates, computes the output probability  $p_0(C)$  up to additive error  $\delta = \exp(-O(m \log m))$ , with probability at least  $1 - \eta$  ( $\eta < \frac{1}{4}$ ) over the choice of  $C$  as well as the randomness of the algorithm. Then PH collapses to a finite level.

*Proof.* Let  $\mathcal{O}$  be an algorithm that correctly approximates  $p_0(C)$  of a random circuit  $C \sim \mathcal{H}_A$  up to additive error  $\delta$ , with success probability at least  $1 - \eta$  over the choice of  $C$ . In the following, we show that there exists a  $\text{BPP}^{\text{NP}^\mathcal{O}}$  procedure that on input *any* circuit  $C_0$ , computes  $p_0(C_0)$  up to additive error  $\delta' = \delta \exp(O(m \log m))$ , with success probability at least  $\frac{2}{3}$ . The theorem statement then follows from the worst-case hardness of computing  $p_0(C_0)$  over  $\mathcal{A}$ .

Consider any circuit  $C_0$  with  $m$  gates over the architecture  $\mathcal{A}$ . Create a new circuit  $C_1$  as follows: for each gate  $G_i$  ( $i = 1, \dots, m$ ) in  $C_0$ , we replace  $G_i$  with  $H_i G_i$ , where  $\{H_i\}$  is independently drawn from the Haar measure over the unitary group and has the same dimension as  $G_i$ . By the invariance of Haar measure,  $C_1$  is distributed the same as  $\mathcal{H}_A$ .

Fix the random unitary gates  $\{H_i\}$ . Next, we apply the  $\theta$ -Cayley perturbation on  $C_0$  using the random seed  $\{H_i\}$  as in Definition 3 to get the perturbed circuit  $C(\theta)$ . By definition, we have  $C(0) = C_1 \sim \mathcal{H}_A$  and  $C(1) = C_0$ .

By Lemma 1,  $p_0(C(\theta))$  is a degree  $(O(m), O(m))$  rational function in  $\theta$ . Let  $P(\theta)$ ,  $Q(\theta)$  be the numerator and denominator of  $p_0(C(\theta))$ , respectively, then  $p_0(C(\theta)) = \frac{P(\theta)}{Q(\theta)}$ . Note that from the proof of Lemma 1, we have that

$$Q(\theta) = \prod_{j=1}^m \prod_l \left| 1 - i(1 - \theta) \tan \frac{\varphi_{jl}}{2} \right|^2. \quad (13)$$

The goal is to recover  $p_0(C(1)) = \frac{P(1)}{Q(1)}$  from values of  $p_0(C(\theta))$  for small  $\theta$ . To do this, first note that  $Q(\theta)$  is a known polynomial whose value can be computed in time  $O(m)$  for any  $\theta$ . Therefore, the problem can be reduced to a polynomial interpolation for  $P(\theta)$ .

To analyze the error for the polynomial interpolation, it is useful to establish bounds for  $Q(\theta)$ . We can write  $Q(\theta)$  as

$$Q(\theta) = \prod_{j=1}^m \prod_l \left| 1 - i(1 - \theta) \tan \frac{\varphi_{jl}}{2} \right|^2 = \prod_{j=1}^m \prod_l \left( 1 + (1 - \theta)^2 \tan^2 \frac{\varphi_{jl}}{2} \right), \quad (14)$$

where it is easy to see that  $Q(\theta) \geq 1$ . Next, call the set of Haar random gates  $\{H_j\}_{j=1, \dots, m}$   $\beta$ -good, if all eigenvalues  $\varphi_{jl}$  of all gates lie in the range  $[-\pi + \beta, \pi - \beta]$ . By Lemma 13, this happens with probability at least  $1 - O(m\beta)$ . Suppose we choose  $\beta = O(m^{-1})$  such that  $\{H_j\}$  is  $\beta$ -good with high constant probability. Then conditioned on  $\{H_j\}$  being  $\beta$ -good, we have

$$\begin{aligned} Q(\theta) &= \prod_{j=1}^m \prod_l \left( 1 + (1 - \theta)^2 \tan^2 \frac{\varphi_{jl}}{2} \right) \\ &\leq \prod_{j=1}^m \prod_l \left( 1 + \tan^2 \frac{\pi - \beta}{2} \right) \\ &\leq \prod_{j=1}^m \prod_l \left( 1 + \frac{4}{\beta^2} \right) \\ &= (1 + O(m^2))^{O(m)} \\ &= \exp(O(m \log m)). \end{aligned} \quad (15)$$

We also define the above upper bound of  $Q(\theta)$  as  $K$ , where  $K = \exp(O(m \log m))$ .

Next we apply the algorithm  $\mathcal{O}$  on input  $C(\theta)$ . Notice that by definition,  $\mathcal{O}$  works on inputs from the distribution  $\mathcal{H}_{\mathcal{A}}$ , while  $C(\theta)$  is distributed according to  $\mathcal{H}_{\mathcal{A},\theta}$  as in Definition 3. Therefore, the success probability of  $\mathcal{O}$  depends on the distance between the two distributions,

$$\Pr_{C(\theta) \sim \mathcal{H}_{\mathcal{A},\theta}} [|\mathcal{O}(C(\theta)) - \mathbf{p}_0(C(\theta))| \geq \delta] \leq \eta + D_{\text{TV}}(\mathcal{H}_{\mathcal{A},\theta}, \mathcal{H}_{\mathcal{A}}), \quad (16)$$

where  $D_{\text{TV}}$  denotes total variation distance. By Lemma 2, we have  $D_{\text{TV}}(\mathcal{H}_{\mathcal{A},\theta}, \mathcal{H}_{\mathcal{A}}) = O(m\theta)$ . Let  $\Delta = O(m^{-1})$  and restrict  $\theta$  in the interval  $[0, \Delta]$ , such that  $D_{\text{TV}}(\mathcal{H}_{\mathcal{A},\theta}, \mathcal{H}_{\mathcal{A}}) = O(m\Delta)$  is upper bounded by a small constant.

Conditioned on  $\mathcal{O}$  being successful and  $\{H_j\}$  being  $\beta$ -good, we have  $|\mathcal{O}(C(\theta)) - \frac{P(\theta)}{Q(\theta)}| \leq \delta$ . After seeing the output  $\mathcal{O}(C(\theta))$  of  $\mathcal{O}$ , we multiply it with  $Q(\theta)$  to get an estimate of  $P(\theta)$ , which satisfies

$$|\mathcal{O}(C(\theta))Q(\theta) - P(\theta)| \leq \delta Q(\theta) \leq \delta K. \quad (17)$$

By a simple union bound, the above equation holds with failure probability at most

$$\begin{aligned} \Pr [|\mathcal{O}(C(\theta))Q(\theta) - P(\theta)| \geq \delta K] &\leq \Pr [|\mathcal{O}(C(\theta)) - \mathbf{p}_0(C(\theta))| \geq \delta \vee \{H_j\} \text{ not } \delta\text{-good}] \\ &\leq \eta + O(m\Delta) + O(m\beta) \leq \eta' < \frac{1}{4}, \end{aligned} \quad (18)$$

for a suitable choice of constants in the  $O$ -notation for  $\Delta = O(m^{-1})$  and  $\beta = O(m^{-1})$ , such that  $\eta' < \frac{1}{4}$ .

To compute  $P(1)$ , we apply the algorithm  $\mathcal{O}$  to a set of circuits  $\{C(\theta_i)\}$ , where  $\theta_i$  ( $i = 1, \dots, O(m^2)$ ) is a set of equally spaced points in the interval  $[0, \Delta]$ , and different perturbed circuits  $C(\theta_i)$  shares the same random seed  $\{H_j\}$ . By Eq. (18), we obtain a set of points  $\{(\theta_i, y_i)\}$  such that

$$\Pr [|y_i - P(\theta_i)| \geq \delta K] \leq \eta' < \frac{1}{4}. \quad (19)$$

The problem is then reduced to a polynomial interpolation for  $P(\theta)$ , a degree  $O(m)$  polynomial, with the noisy data  $\{(\theta_i, y_i)\}$ . Using our robust Berlekamp-Welch theorem which we develop in the next section (see Theorem 7), on input  $\{(\theta_i, y_i)\}$  we can compute a number  $p \approx P(1)$  with access to an NP oracle, where the error can be bounded by

$$\begin{aligned} |p - P(1)| &\leq \delta K \exp(O(m \log \Delta^{-1}) + O(m)) \\ &= \delta \cdot \exp(O(m \log m)). \end{aligned} \quad (20)$$

Finally, notice that

$$\mathbf{p}_0(C_0) = \mathbf{p}_0(C(1)) = \frac{P(1)}{Q(1)} = P(1) \quad (21)$$

as  $Q(1) = 1$ , therefore by choosing  $\delta = \exp(-O(m \log m))$  with a sufficiently large constant, we can compute the worst-case output probability  $\mathbf{p}_0(C_0)$  up to additive error  $\exp(-O(m \log m))$ . The overall procedure is in  $\text{BPP}^{\text{NP}^{\mathcal{O}}}$ . If  $\mathcal{O}$  is an algorithm that belongs to some finite level of PH, then this procedure also belongs to a finite level of PH, and therefore by the worst-case  $\#P$  hardness the PH collapses to a finite level.  $\square$

In the above proof we used a result developed in Section 5 as a black box, which we refer to as robust Berlekamp-Welch algorithm. This algorithm allows us to do polynomial interpolation on noisy data, while tolerating a constant fraction of data points that can be arbitrarily wrong

(see Section 2 for a detailed discussion). Similar to the Learning with Errors problem for solving noisy linear equations, the polynomial interpolation problem with both noise and corruption seems unlikely to be solved in polynomial time. However, recall that given the  $\#P$  hardness of computing  $\text{Per}(C)$  in the worst case, in our worst-to-average-case reduction we are allowed to use any finite level of  $\text{PH}$ . We show how to do such a polynomial interpolation in Section 5 having access to a  $\text{NP}$  oracle. As a result, we are able to tolerate a constant failure probability in our average-case hardness.

Furthermore, our proof techniques can also be naturally applied to BosonSampling, where a fundamental question is to prove robust average-case hardness results for computing permanents of matrices with i.i.d. Gaussian entries [AA11, AA14]. Consider a BosonSampling experiment with  $n$  photons and  $m = n^c$  ( $c > 2$ ) output modes. Then assuming no collision, the output probability corresponding to a output pattern  $S = (s_1, \dots, s_m)$  ( $s_i \in \{0, 1\}$ ,  $\sum_i s_i = n$ ) is

$$\Pr[S] = |\text{Per}(A_S)|^2, \quad (22)$$

where  $A$  is a  $m \times n$  submatrix of a  $m \times m$  Haar random unitary matrix, and  $A_S$  is the  $n \times n$  submatrix of  $A$  whose rows are selected according to  $S$ . When  $c$  is a large enough constant,  $A_S$  is distributed close in total variation distance to matrices with i.i.d. complex Gaussian entries of mean 0 and variance  $\frac{1}{m}$ . We then obtain the following result using similar proof techniques:

**Corollary 1.** *For any constant  $\eta < \frac{1}{4}$ , it is  $\#P$ -hard under a  $\text{BPP}^{\text{NP}}$  reduction to compute the output probability of a  $n$ -photon  $m = n^c$ -mode BosonSampling experiment up to additive error  $\delta = e^{-(c+4)n \log n - O(n)}$ , with success probability  $1 - \eta$ .*

**Remark 2.** *As shown in [AA11], to prove the hardness of approximate sampling for BosonSampling experiments, it suffices to improve the constant in our result from  $c + 4$  to  $c - 1$ . This is because on average the output probability is roughly*

$$\frac{1}{\binom{m}{n}} \approx \frac{n!}{m^n} \approx e^{-(c-1)n \log n}. \quad (23)$$

*However, the barrier result shown by [AA11] implies that this improvement cannot be obtained using similar techniques based on polynomial interpolation. See Appendix B for a detailed discussion, where we also show that this barrier result does not rule out improving the constant to  $c + 1$ .*

*Proof.* Let  $\mathcal{G}^{n \times n}$  denote the distribution over  $n \times n$  matrices where each entry is independently distributed according to  $\mathcal{CN}(0, 1)$ , the standard complex Gaussian distribution. Suppose  $c$  is a large enough constant<sup>24</sup>, then an algorithm for approximating the output probability of a BosonSampling experiment can be used to approximate

$$p_X := \frac{|\text{Per}(X)|^2}{m^n}, \quad X \sim \mathcal{G}^{n \times n} \quad (24)$$

with a small loss in the success probability.

Let  $\mathcal{O}$  be an algorithm that, on input a random matrix  $X \sim \mathcal{G}^{n \times n}$ , correctly approximates  $p_X$  up to additive error  $\delta$ , with success probability at least  $1 - \eta$  over the choice of  $X$ . In the following, we show that there exists a  $\text{BPP}^{\text{NP}^{\mathcal{O}}}$  procedure that on input *any* 0/1 matrix  $X_0$ , computes  $|\text{Per}(X_0)|^2$  up to additive error  $\delta' = \delta \exp((c+4)n \log n + O(n))$ , with success probability at least  $\frac{2}{3}$ . The theorem statement then follows from the worst-case  $\#P$  hardness of computing  $\text{Per}(X_0)$ .

<sup>24</sup>The proof in [AA11] requires  $c$  to be at least 5, and it was conjectured that  $c > 2$  suffices.

Let  $X_1 \sim \mathcal{G}^{n \times n}$  and define

$$X(\theta) := (1 - \theta)X_1 + \theta X_0. \quad (25)$$

Then  $|\text{Per}(X(\theta))|^2$  is a degree  $d = 2n$  polynomial in  $\theta$ . By a result of [AA11], the total variation distance between the distribution of  $X(\theta)$  and  $\mathcal{G}^{n \times n}$  is  $O(n^2\theta)$ . This can be shown by calculating the distance for one entry, which is  $O(\theta)$ , and then multiply by  $n^2$  as the entries are independently distributed.

Consider  $O(n^2)$  uniformly spaced points  $\{\theta_i\}$  in  $[0, \Delta]$  with  $\Delta = O(n^{-2})$ . For a suitable choice of constants, we can guarantee that for each data point  $\theta_i$ ,

$$\Pr \left[ \left| \mathcal{O}(X(\theta_i)) - \frac{|\text{Per}(X(\theta_i))|^2}{m^n} \right| \geq \delta \right] \leq \eta + O(n^2\Delta) \leq \eta' \quad (26)$$

for some constant  $\eta' < \frac{1}{4}$ . The procedure then follows by sending the points  $\{(\theta_i, \mathcal{O}(X(\theta_i)))\}$  to the robust Berlekamp-Welch algorithm (Theorem 7) to obtain the desired approximation to the worst-case quantity  $|\text{Per}(X_0)|^2 = |\text{Per}(X(1))|^2$ . Overall we have a  $\text{BPP}^{\text{NP}^\mathcal{O}}$  procedure for computing  $|\text{Per}(X_0)|^2$  up to additive error

$$\delta \cdot m^n \cdot \exp(d \log \Delta^{-1} + O(d)) = \delta \cdot \exp((c + 4)n \log n + O(n)), \quad (27)$$

which concludes the proof.  $\square$

## 4.2 Noisy circuits and error detection

Next, we generalize our average-case hardness result to noisy random quantum circuits. We first define the computational problem of approximating the output probability of noisy quantum circuits, and show that worst-case hardness can be established via error detection. Then in Section 4.3 we show that our previous reduction can be directly applied to noisy circuits, therefore obtaining average-case hardness.

For simplicity, here we consider noise models that are *local* and *stochastic*. Namely, for a circuit  $C$  with  $m$  gates, there are  $r = O(m)$  independent noise channels, where each noise channel  $\mathcal{N}_k$  ( $k = 1, \dots, r$ ) acts on a constant number of neighboring qubits and has the following form,

$$\mathcal{N}_k = (1 - \gamma)\mathcal{I} + \gamma\mathcal{E}_k, \quad (28)$$

where  $\mathcal{I}$  is the identity channel,  $\mathcal{E}_k$  is an arbitrary CPTP map, and  $\gamma$  is the noise rate. We assume that the noise model  $\mathcal{N} = \{\mathcal{N}_k\}_{k=1, \dots, r}$  is fixed *a priori* and does not depend on the choice of gates. In principle, our hardness proof can be applied to more general noise models, as long as they are gate-independent and feasible for error correction/detection. In particular, the choice of a fixed noise rate for each noise channel is not essential and is for simplicity, as otherwise we can define  $\gamma$  as the maximum of all noise rates.

Our noise model (28) can be understood as follows: for each  $k = 1, \dots, r$ , with probability  $1 - \gamma$ , do nothing; with probability  $\gamma$ , apply  $\mathcal{E}_k$ . The overall probability of having no error is at least  $(1 - \gamma)^r$ . A simulation of the noisy circuit needs to take expectation over these stochastic paths. Equivalently, we can think of  $\mathcal{N}_k$  as quantum channels that deterministically acts on input density matrices. Simulating the output probability of a noisy circuit is therefore equivalent to computing the diagonal elements of the output density matrix of the noisy circuit.

**Definition 4.** Fix a noise model  $\mathcal{N}$ . An algorithm simulates a circuit  $C$  of  $n$  qubits under the noise model  $\mathcal{N}$  up to additive error  $\delta$ , if on input  $C$ , it computes a number  $p$  that satisfies

$$|p - \mathbf{p}_0(C, \mathcal{N})| \leq \delta, \quad (29)$$

where  $\mathbf{p}_0(C, \mathcal{N}) := \text{Tr} [|0^n\rangle\langle 0^n| \cdot \mathcal{C}_{\mathcal{N}}(|0^n\rangle\langle 0^n|)]$ , and  $\mathcal{C}_{\mathcal{N}}$  denotes the quantum channel induced by applying  $\mathcal{N}$  to  $C$ . We use  $\mathbf{p}_0(C)$  to denote the same output probability in the noiseless case.

To show the computational hardness of simulating noisy circuits, one straightforward idea is to apply fault-tolerant quantum error correction. Suppose  $\mathcal{A}$  is an architecture over which fault-tolerant error correction is possible. For example,  $\mathcal{A}$  can consist of nearest neighbor two qubit gates over a 1D or 2D grid. Let  $C_0$  be a quantum circuit whose output probability is  $\#P$ -hard to approximate within an exponentially small additive error. Then by fault-tolerance threshold theorems [Kit97, ABO08], when the noise rate of  $\mathcal{N}$  is below a constant threshold, we can construct a quantum circuit  $C$  over  $\mathcal{A}$  with a polynomial overhead in the size of  $C_0$ , such that even with the presence of noise, the output distribution of  $C$  is exponentially close to the output distribution of  $C_0$  in total variation distance. The computational hardness of simulating  $C$  with the noise model  $\mathcal{N}$  then follows from the hardness of simulating  $C_0$ .

An improvement of the above argument developed by Fujii [Fuj16] is to identify the fact that fault-tolerant *error detection* suffices. This has the benefit of having a smaller overhead in implementing fault-tolerance, and more importantly, a higher noise threshold [Rei06, AGP07]. Fujii's result, which establishes the hardness of sampling from the output distribution of noisy circuits in the worst case, follows from a similar post-selection argument by Bremner, Jozsa and Shepherd [BJS11]: following Aaronson's PostBQP = PP result [Aar05], to prove that a family of circuits are hard to sample from, it suffices to show that they are universal for quantum computation under post-selection. To apply this argument, note that applying standard threshold theorems for an error detecting code, we have that when post-selecting the “no error detected” event, the output distribution of a noisy circuit that implements error detection is exponentially close to the underlying logical circuit in total variation distance, which means that the error-detected noisy circuits are universal for BQP under post-selection. Therefore, when noise is below the error detection threshold, the output distribution of noisy circuits are hard to sample from in the worst case.

Different from Fujii's result, here we consider the worst-case hardness of computing the output probability of noisy circuits. We show that here error detection also suffices: when the noise rate is below the error detection threshold, computing the output probability up to an exponentially small additive error is hard for the Polynomial Hierarchy.

**Theorem 5.** *Fix a noise model  $\mathcal{N}$ , let  $\gamma_{\text{th}}$  be the corresponding fault-tolerant error detection threshold, and suppose the noise rate of  $\mathcal{N}$  satisfies  $\gamma < \gamma_{\text{th}}$ . Suppose there exists a PH procedure (that belongs to some finite level of PH) that on input  $C$  with  $m$  gates, computes  $\mathbf{p}_0(C, \mathcal{N})$  up to additive error  $\delta = \exp(-O(m \log m))$ . Then PH collapses to a finite level.*

The proof follows from a similar application of the threshold theorem, and is presented in Appendix C.

### 4.3 Average-case hardness of noisy circuits

We prove average-case hardness of noisy circuits by extending our worst-to-average-case reduction to the noisy case. We start by re-examining the two properties of the  $\theta$ -Cayley perturbed random circuit distribution in the noisy setting. First, note that Property 1 is unaffected, as we have a fixed noise model that is independent of the choice of gates. Second, in the following we show that Property 2 still holds, which follows from the linearity of quantum channels and the convex structure of the Feynman path integral.

**Lemma 3.** *For any circuit  $C_0$  and any fixed noise model  $\mathcal{N}$ , let  $C(\theta)$  be a circuit from the  $\theta$ -Cayley perturbed random circuit distribution as in Definition 3. Then  $\mathbf{p}_0(C(\theta), \mathcal{N})$  is a degree  $(O(m), O(m))$  rational function in  $\theta$ .*

*Proof.* The proof follows a similar argument as in Lemma 1. Consider a noise model  $\mathcal{N} = \{\mathcal{N}_j\}$ . We prove the lemma statement in a special case where  $r = m$  and each gate  $H_j(\theta)G_j$  is followed by a noise channel  $\mathcal{N}_j$ , which can be directly generalized to the general case where the noise channels are arbitrarily placed.

As each noise channel  $\mathcal{N}_j$  only acts on a constant number of qubits, it can be written as  $\mathcal{N}_j(\rho) = \sum_{l_j} E_{jl_j} \rho E_{jl_j}^\dagger$  where the summation is over constant number of terms, and  $\{E_{jl_j}\}_{l_j}$  are the noise operators of  $\mathcal{N}_j$  satisfying  $\sum_{l_j} E_{jl_j}^\dagger E_{jl_j} = I$ . Let  $\mathcal{U}_j$  be the unitary channel of the gate  $H_j(\theta)G_j$ . Then, we have

$$\begin{aligned} \mathbf{p}_0(C(\theta), \mathcal{N}) &= \text{Tr} [ |0^n\rangle\langle 0^n| \cdot \mathcal{N}_m \circ \mathcal{U}_m \circ \cdots \circ \mathcal{N}_1 \circ \mathcal{U}_1 (|0^n\rangle\langle 0^n|) ] \\ &= \sum_{l_1, \dots, l_m} \left| \langle 0^n | \prod_{j=1}^m E_{jl_j} H_j(\theta) G_j | 0^n \rangle \right|^2. \end{aligned} \quad (30)$$

Here for simplicity we omit the identity operators acting on qubits outside the support of the gates. The summation index  $l_1, \dots, l_m$  can be understood as a *noise trajectory*, where the noise operators are stochastically applied. Note that for each term in the above summation, the same analysis for  $\mathbf{p}_0(C(\theta))$  can be applied (cf. Lemma 1), which shows that  $\left| \langle 0^n | \prod_{j=1}^m E_{jl_j} H_j(\theta) G_j | 0^n \rangle \right|^2$  is a degree  $(O(m), O(m))$  rational function in  $\theta$  with denominator  $Q(\theta) = \prod_{j=1}^m \prod_l |1 - i(1-\theta) \tan \frac{\varphi_{jl}}{2}|^2$ . Note that this denominator does not depend on the noise trajectory. Therefore, each term in the summation shares the same denominator, and the sum  $\mathbf{p}_0(C(\theta), \mathcal{N})$  is also a degree  $(O(m), O(m))$  rational function in  $\theta$  with denominator  $Q(\theta)$ .  $\square$

Having established worst-case hardness of noisy circuits (Theorem 5) and the low-degree rational function property (Lemma 3) for the reduction, the proof of average-case hardness then follows from a similar argument as in Theorem 4.

**Theorem 6.** *Fix a noise model  $\mathcal{N}$ . Let  $\mathcal{A}$  be a circuit architecture that can implement universal fault-tolerant error detection, and suppose the noise rate  $\gamma < \gamma_{\text{th}}$  is smaller than the error detection threshold. Suppose there exists a PH procedure (that belongs to some finite level of PH) that, on input a random circuit  $C \sim \mathcal{H}_{\mathcal{A}}$  with  $m$  gates, computes the output probability  $\mathbf{p}_0(C, \mathcal{N})$  up to additive error  $\delta = \exp(-O(m \log m))$ , with probability at least  $1 - \eta$  ( $\eta < \frac{1}{4}$ ) over the choice of  $C$  as well as the randomness of the algorithm. Then PH collapses to a finite level.*

*Proof.* Let  $\mathcal{O}$  be an algorithm that correctly approximates  $\mathbf{p}_0(C, \mathcal{N})$  of a random circuit  $C \sim \mathcal{H}_{\mathcal{A}}$  up to additive error  $\delta$ , with success probability at least  $1 - \eta$  over the choice of  $C$ . In the following, we show that there exists a  $\text{BPP}^{\text{NP}^{\mathcal{O}}}$  procedure that on input *any* circuit  $C_0$ , computes  $\mathbf{p}_0(C_0, \mathcal{N})$  up to additive error  $\delta' = \delta \exp(O(m \log m))$ , with success probability at least  $\frac{2}{3}$ . The theorem statement then follows from Theorem 5, the worst-case hardness result.

Following the same reduction technique as in Theorem 4, we sample a set of random unitary gates  $\{H_i\}$  and apply the  $\theta$ -Cayley perturbation on  $C_0$  using the random seed  $\{H_i\}$  as in Definition 3 to get the perturbed circuit  $C(\theta)$ . By definition, we have  $C(0) = C_1 \sim \mathcal{H}_{\mathcal{A}}$  and  $C(1) = C_0$ .

By Lemma 3,  $\mathbf{p}_0(C(\theta), \mathcal{N})$  is a degree  $(O(m), O(m))$  rational function in  $\theta$ . Let  $P(\theta)$ ,  $Q(\theta)$  be the numerator and denominator of  $\mathbf{p}_0(C(\theta), \mathcal{N})$ , respectively, then  $\mathbf{p}_0(C(\theta), \mathcal{N}) = \frac{P(\theta)}{Q(\theta)}$ . Note that from the proof of Lemma 3, we have that  $Q(\theta) = \prod_{j=1}^m \prod_l |1 - i(1-\theta) \tan \frac{\varphi_{jl}}{2}|^2$  with  $Q(1) = 1$ .

The goal is to recover  $\mathbf{p}_0(C_0, \mathcal{N})$  from values of  $\mathbf{p}_0(C(\theta), \mathcal{N})$  for small  $\theta$ . To do this, first note that  $Q(\theta)$  is a known polynomial whose value can be computed in time  $O(m)$  for any  $\theta$ . Therefore,

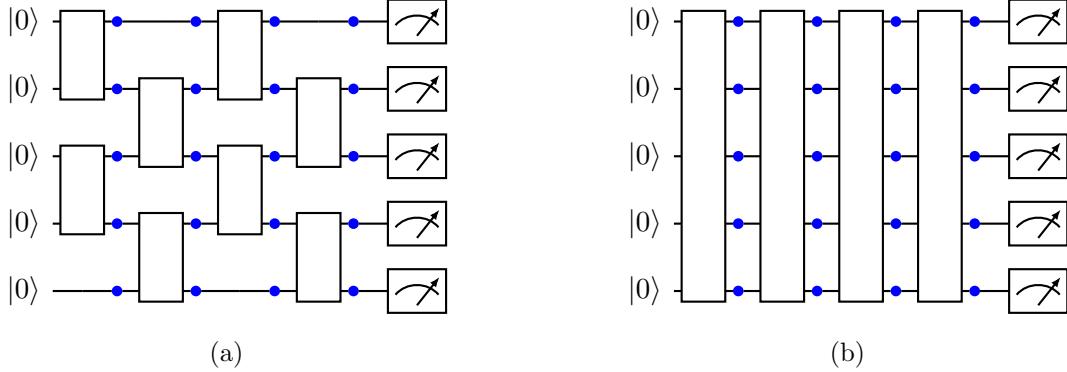


Figure 1: Noisy random circuit sampling, where each white box is a unitary gate independently drawn from the Haar measure, and blue dots represent noise. (a) Experimental implementation of random circuit sampling, where the circuit consists of nearest-neighbor two-qubit gates in an alternating architecture. (b) A toy model for studying the convergence to uniform distribution, where each layer is one global random unitary gate.

the problem can be reduced to a polynomial interpolation for  $P(\theta)$ , as

$$\mathsf{p}_0(C_0, \mathcal{N}) = \mathsf{p}_0(C(1), \mathcal{N}) = \frac{P(1)}{Q(1)} = P(1). \quad (31)$$

The rest of the proof is therefore the same as in Theorem 4.  $\square$

An important distinction between the noisy random circuit model and the ideal model brings interesting implications to our hardness result. In noisy random circuits, it was shown [ABOIN96, GD18] that the output distribution converges to the uniform distribution very quickly as circuit depth increases. Therefore, trivially outputs  $\frac{1}{2^n}$  is already a good approximation algorithm for  $\mathsf{p}_0(C, \mathcal{N})$ , while this is not the case for ideal random circuits, whose output probability has the Porter-Thomas distribution which is far from uniform in total variation distance. In other words, the amount of additive imprecision that we can tolerate in the hardness proof has an *upper bound*: it cannot be larger than the distance between the noisy output distribution and the uniform distribution.

This behavior can be understood as follows: the decoherence caused by noise leaks quantum information to the environment (here we consider incoherent noise models), and the amount of information that remains in the system decreases as circuit depth increases, which eventually goes to 0 and corresponds to the maximally mixed state. It was first shown by Aharonov *et al.* [ABOIN96] that the output distribution of *any* circuit with i.i.d. depolarizing noise converges to the uniform distribution with rate  $e^{-\gamma d}$ , where  $\gamma$  is the noise rate on each qubit at a time step, and  $d$  is circuit depth. Gao and Duan [GD18] proved a similar convergence speed for random circuits with general i.i.d. Pauli noise. However, it was conjectured [BIS<sup>+</sup>18] that the actual convergence speed for noisy random circuits is much faster, with rate  $e^{-\gamma nd}$  (or  $e^{-\gamma O(m)}$ ), which was indeed observed in experiments [AAB<sup>+</sup>19]. While the proof of this behavior for the local random circuit model (Fig. 1a) is out of the scope for this paper, we conjecture that it is true at least for some parameter region of  $(\gamma, n, d)$ . Furthermore, we provide a rigorous analysis for the toy model with global random gates (Fig. 1b), which we present in Appendix D.

**Lemma 4.** *The output distribution of the global noisy random circuit model (Fig. 1b) converges to the uniform distribution with speed  $e^{-\Omega(\gamma nd)}$ . More specifically, let  $p_e$  denote the output distribution*

of the noisy circuit with i.i.d. depolarizing noise with noise rate  $\gamma$ . Then

$$\mathbb{E} \left[ 2^n \sum_x p_e^2(x) - 1 \right] \leq \exp(-c\gamma n(d - d_0)), \quad (32)$$

where  $c, d_0 > 0$  are universal constants, and the expectation is over the global random circuit ensemble.

To apply this result, we can use a simple Markov's inequality which suggests that the output distribution of *most* noisy quantum circuits satisfy  $2^n \sum_x p_e^2(x) - 1 = e^{-\Omega(\gamma n d)}$ . Then, for each of these circuits, the total variation distance between the noisy distribution and the uniform distribution can be bounded by

$$D_{\text{TV}}(p_e, p_{\text{uniform}}) = \frac{1}{2} \sum_x \left| p_e(x) - \frac{1}{2^n} \right| \leq \frac{1}{2} \sqrt{2^n \sum_x p_e^2(x) - 1} \quad (33)$$

via Cauchy–Schwarz inequality.

An immediate observation is that this convergence to uniformity behavior implies that our hardness result in the noisy case is essentially optimal. Indeed, for a family of noisy random circuits where the distance to uniform is bounded by  $e^{-O(m)}$ , our hardness result says that it is still hard to compute the output probability to within  $e^{-O(m \log m)}$  additive imprecision, while it is trivial to achieve  $e^{-O(m)}$  additive imprecision. The key point here is that our proof is noise and architecture agnostic, as long as they enables error detection. Therefore, it is impossible to improve our hardness result to have better robustness, as long as the agnostic property remains.

The same limitation also applies to the noiseless case. In our proof, the noisy hardness result is a direct generalization of the noiseless case due to the linearity and convex structure, which means that our proof for the noiseless case cannot be improved either, as long as it can still be generalized to the noisy case. Therefore, any attempt to address the hardness of approximate sampling via a reduction to computing output probabilities (which requires a robustness of  $2^{-n}/\text{poly}(n)$ ) in the ideal case cannot be naturally adapted to noisy circuits. In this sense our result can be viewed as a complement of [NPD<sup>+</sup>20], which shows that the proof technique needs to be sensitive to depth because of an efficient simulation algorithm for shallow random circuits.

## 5 Robust Berlekamp–Welch

We prove the following theorem for robust polynomial interpolation. All polynomials considered here have real coefficients, and our results can be generalized to the complex case, as for polynomials with complex coefficients we can deal with the real and imaginary parts separately. Below our theorem is stated with success probability  $\frac{2}{3}$ , which can be amplified to  $1 - 1/\exp(d)$  by taking the median of  $\text{poly}(d)$  independent experiments.

**Theorem 7** (Robust Berlekamp–Welch). *For a degree  $d$  polynomial  $P(x)$ , suppose there is a set of data points  $D = \{(x_i, y_i)\}$  such that  $|D| = 100d^2$  and  $\{x_i\}$  is equally spaced in the interval  $[0, \Delta]$  ( $\Delta < 1$ ). Furthermore, assume that each point  $(x_i, y_i)$  satisfies*

$$\Pr [|y_i - P(x_i)| \geq \delta] \leq \eta, \quad (34)$$

where  $\eta < \frac{1}{4}$  is a constant, then there exists a  $\mathsf{P}^{\mathsf{NP}}$  algorithm which, takes  $D$  as input, returns a number  $p$  such that  $|p - P(1)| \leq \delta e^{d \log \Delta^{-1} + O(d)}$ , with success probability at least  $\frac{2}{3}$ .

**Remark 3.** Note that our result can be improved to using only  $O(d)$  data points, by applying a result of Rakhmanov [Rak07]. However, the number of data points we use does not affect our main results, as long as it is polynomial in  $d$ . Here we give a simpler proof using  $O(d^2)$  data points.

*Proof.* For simplicity, our proof is presented by specifying  $\eta = \frac{1}{300}$ , which can be naturally generalized to any  $\eta < \frac{1}{4}$ . Say a point  $(x_i, y_i)$  is correct if  $|y_i - P(x_i)| \leq \delta$ . By Eq. (34), the expected number of wrong points is less than  $\frac{1}{300}$  fraction. By Markov's inequality, with probability at least  $\frac{2}{3}$ , 0.99 fraction of points in  $D$  are correct. In the following we show that conditioned on 0.99 fraction of points being correct (denote the set of correct points as  $F$ ), there exists a deterministic  $\mathsf{P}^{\mathsf{NP}}$  algorithm that computes a number  $p$  that satisfies  $|p - P(1)| \leq \delta e^{d \log \Delta^{-1} + O(d)}$ . This implies the statement of the theorem.

Consider the following computational problem:

**Problem 1.** Given  $100d^2$  points  $D$ , decide if there exists at least 0.99 fraction of them (denoted as  $F'$ ) and a degree  $d$  polynomial  $Q(x)$  that satisfy  $|y_i - Q(x_i)| \leq \delta, \forall (x_i, y_i) \in F'$ .

Problem 1 is in NP because a certificate  $(F', Q)$  can be efficiently verified by checking each point in  $F'$ .

When giving  $D$  in the statement of the theorem as input, this problem has a satisfying certificate  $(F, P)$ . Having access to a NP oracle, as this problem is guaranteed to have a solution, we can find a certificate  $(F', Q)$  by performing binary search in the certificate space, which may be different from  $(F, P)$ . However, in the following we show that any solution will satisfy the requirement of the theorem, and the algorithm simply outputs  $Q(1)$ .

Let  $R(x) = P(x) - Q(x)$ . As  $|F \cap F'| \geq 0.98|D|$ , there are at least 0.98 fraction of points in  $D$  that satisfies

$$|R(x_i)| \leq |P(x_i) - y_i| + |Q(x_i) - y_i| = 2\delta. \quad (35)$$

Recall that  $D$  is a set of equally spaced points in  $[0, \Delta]$ . Lemma 6 says that in such condition  $R(x)$  can be uniformly bounded in the interval  $[0, \Delta]$ :

$$|R(x)| \leq \delta e^{O(d)}, \quad \forall x \in [0, \Delta]. \quad (36)$$

Finally, as  $R$  is uniformly bounded in  $[0, \Delta]$ , we can use Lemma 8 to bound the error at  $x = 1$ ,

$$|R(1)| = |P(1) - Q(1)| \leq \delta e^{O(d)} e^{d \log \Delta^{-1} + d \log 8} = \delta e^{d \log \Delta^{-1} + O(d)}, \quad (37)$$

which concludes the proof.  $\square$

The above proof needs several additional results which we develop in the following. First, recall the following result of Paturi, which gives a bound of how fast a polynomial can grow when it is uniformly bounded in a small interval.

**Lemma 5** (Paturi [Pat92]). *Let  $P$  be a degree  $d$  polynomial which satisfies  $|P(x)| \leq \varepsilon$  for  $x \in [0, \Delta]$  ( $\Delta < 1$ ). Then we have  $|P(1)| \leq \varepsilon e^{4d\Delta^{-1}}$ .*

*Proof.* This is implied by Fact 2 and Corollary 2 of [Pat92].  $\square$

The first missing step in the proof of Theorem 7, which is in Eq. (36), is to show the following: if a degree  $d$  polynomial is bounded by  $\delta$  for a constant fraction of a set of uniformly spaced points in  $[0, \Delta]$ , then the maximum value of the polynomial in the interval  $[0, \Delta]$  is at most  $\delta e^{O(d)}$ . We show this in the following lemma, which is the main technical part of the full proof.

**Lemma 6.** Let  $P$  be a degree  $d$  polynomial which satisfies  $|P(x)| \leq 1$  for any  $x \in A$ , where  $A \subseteq B$  and  $B$  is a set of equally spaced points in  $[0, 1]$  such that  $|B| = 100d^2$  and  $|A| \geq 0.98|B|$ . Then  $P$  satisfies  $|P(x)| \leq 2^{O(d)}$ ,  $\forall x \in [0, 1]$ .

*Proof.* Denote  $A$  as good points at which  $P$  is bounded, and denote  $B - A$  as bad points. We first prove the bound under a special case, where the good points are exactly the first 0.98 fraction. More specifically, for the set of uniformly spaced points

$$B = \{0, \delta, 2\delta, \dots, (100d^2 - 1)\delta\} \subseteq [0, 1] \quad (38)$$

where  $\delta = \frac{1}{100d^2}$ , suppose the polynomial is bounded at the first 0.98 fraction of points,

$$|P(x)| \leq 1, \quad \forall x \in A = \{0, \delta, \dots, (98d^2 - 1)\delta\}. \quad (39)$$

When a polynomial is bounded at consecutive uniformly spaced points in an interval, it can also be uniformly bounded at all points in that interval. Using Lemma 7 we get

$$|P(x)| \leq O(1), \quad \forall x \in [0, (98d^2 - 1)\delta]. \quad (40)$$

As  $P(x)$  is uniformly bounded on a constant subinterval of  $[0, 1]$ , we can then use Lemma 5 to obtain the desired bound  $|P(x)| \leq 2^{O(d)}$ ,  $\forall x \in [0, 1]$ . Also note that both inequalities are saturated by Chebyshev polynomial, so the bound is tight.

We proceed by proving the bound for the general case where  $A$  can be arbitrarily distributed in  $B$ . Let  $M = \max_{x \in [0, 1]} |P(x)|$  and suppose the maximum is achieved at  $x^* \in [\frac{1}{2}, 1]$ . Consider the roots of  $P$

$$P(x) = L \prod_{k=1}^d (x - r_k), \quad L > 0, \quad (41)$$

where  $r_k$  can have real and imaginary parts. We construct a new polynomial  $Q(x) = L \prod_{k=1}^d (x - r'_k)$  according to the following rule:

$$r'_k = \begin{cases} r_k, & \operatorname{Re}(r_k) \leq x^*, \\ 2x^* - \operatorname{Re}(r_k) + i \operatorname{Im}(r_k), & \operatorname{Re}(r_k) > x^*. \end{cases} \quad (42)$$

Let  $M' = \max_{x \in [0, 1]} |Q(x)|$ . We show that this construction is useful because  $Q(x)$  has the following properties:

1.  $\deg(Q) = \deg(P) = d$ .

*Proof.* By construction.

2.  $M' \geq M$ .

*Proof.*  $M' \geq |Q(x^*)| = |P(x^*)| = M$ .

3.  $|Q(x)| \leq |P(x)|$ ,  $\forall x \in [0, x^*]$ .

*Proof.* This follows because  $|x - r'_k| \leq |x - r_k|$ ,  $\forall k$ ,  $\forall x \in [0, x^*]$ .

4.  $M' = |Q(1)|$ .

*Proof.*  $|Q(x)|$  is monotonic increasing at  $[x^*, 1]$  because  $\forall x^* \leq x < y \leq 1$ , we have  $|x - r'_k| \leq |y - r'_k|$ . Also,  $\forall x \in [0, x^*]$ ,  $|Q(x)| \leq |P(x)| \leq M = |Q(x^*)|$ . This suggests that the maximum occurs at  $x = 1$ .

According to property 2, we proceed by proving an upper bound for  $M'$ . Let  $B' = \{0, \delta, 2\delta, \dots, (50d^2 - 1)\delta\} \subseteq [0, \frac{1}{2}]$ . Among these points, at least 0.96 fraction of them are good for  $Q(x)$ , denoted by  $A'$ . More specifically,

$$|Q(x)| \leq |P(x)| \leq 1, \forall x \in A' \subseteq B'. \quad (43)$$

If  $A'$  is a consecutive set starting from 0 (i.e.  $A' = \{0, \delta, 2\delta, \dots, (|A'| - 1)\delta\}$ ), then we can obtain the desired bound by repeating the argument at the beginning of this proof. When this is not the case, there exists  $y \in B'$  such that  $y$  is bad and  $y + \delta$  is good. We proceed by converting  $Q(x)$  to a new polynomial to eliminate this kind of points, while maintaining the good properties of  $Q$ .

Let  $y$  be the maximum element in  $B'$  such that  $y$  is bad and  $y + \delta$  is good (assuming such  $y$  exists). By definition, all good points larger than  $y$  are consecutive, denoted as

$$A' \cap [y, \frac{1}{2}] = \{y + \delta, y + 2\delta, \dots, y + T\delta\} \quad (44)$$

for some  $T \geq 1$ . We perform the following operation on  $Q(x) = L \prod_{k=1}^d (x - r_k)$ , resulting in a new polynomial  $R(x) = L \prod_{k=1}^d (x - r'_k)$ :

$$r'_k = \begin{cases} r_k, & \operatorname{Re}(r_k) \leq y, \\ \operatorname{Re}(r_k) - \delta + i \operatorname{Im}(r_k), & \operatorname{Re}(r_k) > y. \end{cases} \quad (45)$$

We prove the following properties for  $R$ .

5.  $|R(x - \delta)| \leq 1, \forall x \in A' \cap [y, \frac{1}{2}]$ .

*Proof.*  $\forall x \in A' \cap [y, \frac{1}{2}]$ , notice that  $|x - \delta - r'_k| = |x - r_k|$  when  $\operatorname{Re}(r_k) > y$ , and  $|x - \delta - r'_k| \leq |x - r_k|$  when  $\operatorname{Re}(r_k) \leq y$ . Therefore  $|R(x - \delta)| \leq |Q(x)| \leq 1$ .

6.  $|R(x)| \leq 1, \forall x \in A' \cap [0, y]$ .

*Proof.*  $\forall x \in A' \cap [0, y]$ , notice that  $|x - r'_k| = |x - r_k|$  when  $\operatorname{Re}(r_k) \leq y$ , and  $|x - r'_k| \leq |x - r_k|$  when  $\operatorname{Re}(r_k) > y$ . Therefore  $|R(x)| \leq |Q(x)| \leq 1$ .

From property 5-6 we conclude that  $R(x)$  has the same number of good points as  $Q(x)$  (if some point outside of  $(A' \cap [0, y]) \cup \{y, y + \delta, \dots, y + (T - 1)\delta\}$  is good for  $R$ , we still label it as bad). As  $y$  is good for  $R$ , we also have  $\max\{z \in B' : z + \delta \in A'\} < y$ . Finally we show that  $R$  has a larger maximum.

7. Let  $M'' = \max_{x \in [0, 1]} |R(x)|$ . Then  $M'' \geq M'$ .

*Proof.* As  $|1 - r'_k| \geq |1 - r_k|$ , we have  $M'' \geq |R(1)| \geq |Q(1)| = M'$ .

Next, we repeat the above process until  $\{z \in B' : z + \delta \in A'\}$  is empty, and denote the resulting polynomial as  $\tilde{R}$ , for which property 5-7 still holds. As  $\tilde{R}$  is bounded by 1 at  $\{0, \delta, 2\delta, \dots, (48d^2 - 1)\delta\}$ , we obtain

$$M' \leq \max_{x \in [0, 1]} |\tilde{R}(x)| \leq 2^{O(d)} \quad (46)$$

by repeating the argument at the beginning of the proof, which concludes the full proof.  $\square$

A missing ingredient in the above proof is to show that when a polynomial is bounded on a set of  $O(d^2)$  equally spaced points in an interval, then it can be uniformly bounded on that interval. As discussed in Remark 3, the number of points needed can be improved to  $O(d)$  by using a powerful result of Rakhmanov [Rak07].

**Lemma 7.** Let  $P$  be a degree  $d$  polynomial which satisfies  $|P(x)| \leq c$  for equally spaced points  $x \in \{0, \frac{a}{N}, \frac{2a}{N}, \dots, a\}$  in the interval  $[0, a]$ , where  $\frac{d^2}{N} \leq 1 - \varepsilon$ . Then  $P$  is uniformly bounded as  $|P(x)| \leq \frac{c}{\varepsilon}, \forall x \in [0, a]$ .

*Proof.* To prove a uniform bound for  $P(x)$ , we use a Markov's inequality [Lor05] to bound the maximum derivative,

$$\max_{x \in [0, a]} |P'(x)| \leq \frac{2d^2}{a} \max_{x \in [0, a]} |P(x)|. \quad (47)$$

Let  $M = \max_{x \in [0, a]} |P(x)|$  which is achieved at  $t \frac{a}{N} \leq x \leq (t+1) \frac{a}{N}$ . Then

$$\frac{2d^2}{a} M \geq \max_{x \in [0, a]} |P'(x)| \geq (M - c) \frac{2N}{a}, \quad (48)$$

which gives  $M \leq \frac{c}{\varepsilon}$ .  $\square$

We are now ready to fill in the final missing part in the proof of Theorem 7 which is in Eq. (37). A first idea is to directly apply Paturi's Lemma 5 which gives the final error bound  $\delta e^{O(d\Delta^{-1})}$ . However, this is clearly over pessimistic: for degree  $d$  polynomials, the error growth should not be much bigger than  $(1/\Delta)^d$  when  $\Delta$  is very small. Based on this intuition, our following result gives an improvement for the error bound of long distance polynomial extrapolation.

**Lemma 8.** Let  $P$  be a degree  $d$  polynomial that satisfies  $|P(x)| \leq \delta, \forall x \in [0, \Delta]$  ( $\Delta < 1$ ). Then we have

$$|P(1)| \leq \delta e^{d \log \Delta^{-1} + d \log 8}. \quad (49)$$

*Proof.* Let  $Q(x) = P(\frac{x+1}{2}\Delta)$ . Then  $Q(x)$  is a degree  $d$  polynomial that satisfies  $|Q(x)| \leq \delta, \forall x \in [-1, 1]$ . Next we use a well-known fact about polynomials, namely that the coefficients of bounded polynomials are at most exponential in the degree. Let  $Q(x) = \sum_{i=0}^d a_i x^i$ . Lemma 9 implies that

$$\sum_{i=0}^d |a_i| \leq 4^d \delta. \quad (50)$$

Then we have

$$\begin{aligned} |P(1)| &= \left| Q\left(\frac{2}{\Delta} - 1\right) \right| \\ &\leq \sum_{i=0}^d |a_i| \left(\frac{2}{\Delta} - 1\right)^i \\ &\leq \sum_{i=0}^d |a_i| \left(\frac{2}{\Delta}\right)^d \\ &\leq \delta \frac{8^d}{\Delta^d} = \delta e^{d \log \Delta^{-1} + d \log 8}. \end{aligned} \quad (51)$$

$\square$

**Lemma 9** (Lemma 4.1 in [She13]). Let  $P(x) = \sum_{i=0}^d a_i x^i$  be a polynomial. Then

$$\sum_{i=0}^d |a_i| \leq 4^d \max_{x \in [-1, 1]} |P(x)|. \quad (52)$$

## 5.1 Alternative proof via variable rescaling

In Section 2.2.3 we first introduced a proof technique via variable rescaling, which gives the same asymptotic bounds as the results presented above, but yields worse constants. In the following we explicitly calculate the constants when applying this variable rescaling technique to BosonSampling, where the goal is to establish robust average case hardness for random permanents.

We consider the following general framework for polynomial interpolation. Consider a degree  $d$  polynomial  $P(x)$  that is bounded on a large fraction of points in the interval  $[0, \Delta]$  ( $\Delta < 1$ ). That is, suppose there is a set of data points  $D = \{x_i\}$  which is equally spaced in the interval  $[0, \Delta]$ , such that

$$|P(x_i)| \leq \delta \quad (53)$$

for at least  $1 - \eta$  fraction of points in  $D$ . Here we assume that  $|D|$  is a large enough polynomial in  $d$  (for example,  $|D| = 1000d^2$ ), and  $\eta$  is a small enough constant (for example,  $\eta = 0.001$ ). Our results in this subsection suggest that  $|P(1)| \leq \delta e^{O(d\Delta^{-1})}$ , and we will also calculate the constant in the  $O$  notation.

The first step is to bound the polynomial  $P(x)$  in the interval  $[0, \Delta]$ . Consider the set of points  $\{0, t, 2t, \dots, (|D|-1)t\}$  where  $t = \frac{\Delta}{|D|}$ , and we know that  $P(x)$  is bounded at  $1 - \eta$  fraction of these points. From the proof of Lemma 6, it was shown that without loss of generality, we can assume that the polynomial is bounded at the first  $\frac{1}{2} - \eta$  fraction of points in  $D$ . Combining Lemma 7 and Lemma 5 which introduces a constant 4 in the exponent, we have

$$|P(x)| \leq 1.01\delta e^{4d(1/2-\eta)^{-1}} \leq \delta e^{4 \times d \times 2.01}, \quad \forall x \in [0, \Delta]. \quad (54)$$

This allows us to directly invoke Lemma 5 and obtain

$$|P(1)| \leq \delta e^{4 \times d \times 2.01} \cdot e^{4d\Delta^{-1}} = \delta e^{4d(\Delta^{-1} + 2.01)}. \quad (55)$$

Next we plug in the parameters for BosonSampling. The natural polynomial  $\text{Per}(X(\theta))$  is a degree  $n$  polynomial in  $\theta$ . Therefore, the polynomial  $|\text{Per}(X(\theta^k))|^2$  is a degree  $2kn$  polynomial in  $\theta$ . When  $\theta \in [0, \Delta]$ , the total variation distance between  $X(\theta^k)$  and Gaussian permanents is bounded by  $2.01n^2\Delta^k$  [AA11], which we require to be a small constant (say 0.001). This then gives  $\Delta^{-1} = c^{1/k} \cdot n^{2/k}$ , for some large constant  $c \sim 2000$ , and the exponent is

$$4 \cdot 2kn \cdot (2000^{1/k} \cdot n^{2/k} + 2.01). \quad (56)$$

Let  $k = t \log n$  for some constant  $t$ , then the exponent becomes

$$8t \cdot (2000^{1/t \log n} \cdot e^{2/t} + 2.01) \cdot n \log n. \quad (57)$$

Consider large  $n$  limit, and optimize with respect to  $t$ , the constant prefactor is

$$\min_{t>0} 8t \cdot (1.01 \cdot e^{2/t} + 2.01) = 69.7 < 70. \quad (58)$$

Consider BosonSampling experiment with  $n$  photons and  $m = n^c$  ( $c > 2$ ) output modes. The goal is to prove average case hardness for the output probability up to additive imprecision

$$\frac{n!}{m^n} \approx e^{-(c-1)n \log n}. \quad (59)$$

In the worst case, the output probability is hard within  $1/m^n = e^{-cn \log n}$  additive error. Therefore our proof shows average case hardness with robustness

$$\delta = e^{-(c+70)n \log n}. \quad (60)$$

## Acknowledgements

We especially thank Dorit Aharonov as well as Juani Bermejo-Vega, Jonas Haferkamp, Dominik Hangleiter, and Jens Eisert, for early discussions regarding hardness in the presence of noise that were influential in formulating the questions posed in this manuscript. We also thank Scott Aaronson, Roozbeh Bassirianjahromi, Abhinav Deshpande, Raul Garcia-Patron, Sam Gunn, Ramis Movassagh, Chinmay Nirkhe, Bryan O’Gorman and Umesh Vazirani for helpful comments and discussions. This work was supported by Vannevar Bush faculty fellowship N00014-17-1-3025, DOE QSA grant #FP00010905 and NSF QLCI Grant No. 2016245. B.F. acknowledges support from AFOSR (YIP number FA9550-18-1-0148 and FA9550-21-1-0008) and the National Science Foundation under Grant CCF-2044923 (CAREER). We would like to thank the Simons Institute for the Theory of Computing where part of this work was done.

## References

- [AA11] Scott Aaronson and Alex Arkhipov. The computational complexity of linear optics. In *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing*, STOC ’11, page 333–342, New York, NY, USA, 2011. Association for Computing Machinery.
- [AA14] Scott Aaronson and Alex Arkhipov. Bosonsampling is far from uniform. *Quantum Info. Comput.*, 14(15–16):1383–1423, November 2014.
- [AAB<sup>+</sup>19] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C Bardin, Rami Barends, Rupak Biswas, Sergio Boixo, Fernando GSL Brandao, David A Buell, et al. Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779):505–510, 2019.
- [Aar05] Scott Aaronson. Quantum computing, postselection, and probabilistic polynomial-time. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 461(2063):3473–3482, 2005.
- [ABO08] Dorit Aharonov and Michael Ben-Or. Fault-tolerant quantum computation with constant error rate. *SIAM Journal on Computing*, 38(4):1207–1282, 2008.
- [ABOIN96] Dorit Aharonov, Michael Ben-Or, Russell Impagliazzo, and Noam Nisan. Limitations of noisy reversible computation. *arXiv:quant-ph/9611028*, 1996.
- [AC17] Scott Aaronson and Lijie Chen. Complexity-theoretic foundations of quantum supremacy experiments. In *Proceedings of the 32nd Computational Complexity Conference*, CCC ’17, Dagstuhl, DEU, 2017. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.
- [AG19] Scott Aaronson and Sam Gunn. On the classical hardness of spoofing linear cross-entropy benchmarking. *arXiv:1910.12085*, 2019.
- [AGP07] Panos Aliferis, Daniel Gottesman, and John Preskill. Accuracy threshold for postselected quantum computation. *arXiv:quant-ph/0703264*, 2007.
- [AH14] Scott Aaronson and Travis Hance. Generalizing and derandomizing Gurvits’s approximation algorithm for the permanent. *Quantum Information & Computation*, 14(7&8):541–559, 2014.

[BCG20] Boaz Barak, Chi-Ning Chou, and Xun Gao. Spoofing linear cross-entropy benchmarking in shallow quantum circuits. *arXiv:2005.02421*, 2020.

[BFNV19] Adam Bouland, Bill Fefferman, Chinmay Nirkhe, and Umesh Vazirani. On the complexity and verification of quantum random circuit sampling. *Nature Physics*, 15(2):159–163, 2019.

[BHH16] Fernando GSL Brandao, Aram W Harrow, and Michał Horodecki. Local random quantum circuits are approximate polynomial-designs. *Communications in Mathematical Physics*, 346(2):397–434, 2016.

[BIS<sup>+</sup>18] 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, 2018.

[BJS11] Michael J. Bremner, Richard Jozsa, and Dan J. Shepherd. Classical simulation of commuting quantum computations implies collapse of the polynomial hierarchy. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 467(2126):459–472, 2011.

[BMS16] Michael J. Bremner, Ashley Montanaro, and Dan J. Shepherd. Average-case complexity versus approximate simulation of commuting quantum computations. *Phys. Rev. Lett.*, 117:080501, Aug 2016.

[BMS17] Michael J Bremner, Ashley Montanaro, and Dan J Shepherd. Achieving quantum supremacy with sparse and noisy commuting quantum computations. *Quantum*, 1:8, 2017.

[CPS99] Jin-Yi Cai, Aduri Pavan, and D Sivakumar. On the hardness of permanent. In *Annual Symposium on Theoretical Aspects of Computer Science*, pages 90–99. Springer, 1999.

[DHJB20] Alexander M Dalzell, Nicholas Hunter-Jones, and Fernando GSL Brandão. Random quantum circuits anti-concentrate in log depth. *arXiv:2011.12277*, 2020.

[FGHP99] Stephen Fenner, Frederic Green, Steven Homer, and Randall Pruim. Determining acceptance possibility for a quantum computation is hard for the polynomial hierarchy. *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, 455(1991):3953–3966, 1999.

[FR99] Lance Fortnow and John Rogers. Complexity limitations on quantum computation. *Journal of Computer and System Sciences*, 59(2):240 – 252, 1999.

[Fuj16] Keisuke Fujii. Noise threshold of quantum supremacy. *arXiv:1610.03632*, 2016.

[GD18] Xun Gao and Luming Duan. Efficient classical simulation of noisy quantum computation. *arXiv:1810.03176*, 2018.

[GLR<sup>+</sup>91] Peter Gemmell, Richard Lipton, Ronitt Rubinfeld, Madhu Sudan, and Avi Wigderson. Self-testing/correcting for polynomials and for approximate functions. In *Proceedings of the Twenty-Third Annual ACM Symposium on Theory of Computing*, STOC ’91, page 33–42, New York, NY, USA, 1991. Association for Computing Machinery.

[GPRS19] Raúl García-Patrón, Jelmer J Renema, and Valery Shchesnovich. Simulating boson sampling in lossy architectures. *Quantum*, 3:169, 2019.

[GS92] Peter Gemmell and Madhu Sudan. Highly resilient correctors for polynomials. *Information processing letters*, 43(4):169–174, 1992.

[Gur05] Leonid Gurvits. On the complexity of mixed discriminants and related problems. In *Proceedings of the 30th International Conference on Mathematical Foundations of Computer Science*, MFCS’05, page 447–458, Berlin, Heidelberg, 2005. Springer-Verlag.

[Gur06] Venkatesan Guruswami. List decoding in average-case complexity and pseudorandomness. In *2006 IEEE Information Theory Workshop-ITW’06 Punta del Este*, pages 32–36. IEEE, 2006.

[HL09] Aram W. Harrow and Richard A. Low. Random quantum circuits are approximate 2-designs. *Communications in Mathematical Physics*, 291(1):257–302, Oct 2009.

[HM18] Aram Harrow and Saeed Mehraban. Approximate unitary  $t$ -designs by short random quantum circuits using nearest-neighbor and long-range gates. *arXiv:1809.06957*, 2018.

[HP06] Fumio Hiai and Denes Petz. *The Semicircle Law, Free Random Variables and Entropy (Mathematical Surveys & Monographs)*. American Mathematical Society, USA, 2006.

[Kit97] A Yu Kitaev. Quantum computations: algorithms and error correction. *Russian Mathematical Surveys*, 52(6):1191–1249, dec 1997.

[KMM21] Yasuhiro Kondo, Ryuhei Mori, and Ramis Movassagh. Fine-grained analysis and improved robustness of quantum supremacy for random circuit sampling. *arXiv:2102.01960 [quant-ph]*, 2021.

[Lip89] Richard J. Lipton. New directions in testing. In Joan Feigenbaum and Michael Merritt, editors, *Distributed Computing And Cryptography, Proceedings of a DIMACS Workshop, Princeton, New Jersey, USA, October 4-6, 1989*, volume 2 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 191–202. DIMACS/AMS, 1989.

[Lor05] G.G. Lorentz. *Approximation of Functions*. AMS Chelsea Publishing Series. Holt, Rinehart and Winston, 2005.

[Mov20] Ramis Movassagh. Quantum supremacy and random circuits. *arXiv:1909.06210v4 [quant-ph]*, 2020.

[NPD<sup>+</sup>20] John Napp, Rolando L. La Placa, Alexander M. Dalzell, Fernando G. S. L. Brandao, and Aram W. Harrow. Efficient classical simulation of random shallow 2D quantum circuits. *arXiv:2001.00021*, 2020.

[ODMZ20] Michał Oszmaniec, Ninnat Dangniam, Mauro ES Morales, and Zoltán Zimborás. Fermion sampling: a robust quantum computational advantage scheme using fermionic linear optics and magic input states. *arXiv:2012.15825*, 2020.

[ONFJ21] Changhun Oh, Kyungjoo Noh, Bill Fefferman, and Liang Jiang. Classical simulation of lossy boson sampling using matrix product operators. *arXiv:2101.11234*, 2021.

[Pat92] Ramamohan Paturi. On the degree of polynomials that approximate symmetric boolean functions (preliminary version). In *Proceedings of the Twenty-fourth Annual ACM Symposium on Theory of Computing*, STOC '92, pages 468–474, New York, NY, USA, 1992. ACM.

[QBQGP20] Haoyu Qi, Daniel J Brod, Nicolás Quesada, and Raúl García-Patrón. Regimes of classical simulability for noisy Gaussian boson sampling. *Physical review letters*, 124(10):100502, 2020.

[Rak07] E. A. Rakhmanov. Bounds for polynomials with a unit discrete norm. *Annals of Mathematics*, 165(1):55–88, 2007.

[Rei06] Ben W Reichardt. Error-detection-based quantum fault tolerance against discrete Pauli noise. *arXiv:quant-ph/0612004*, 2006.

[She13] Alexander A. Sherstov. Making polynomials robust to noise. *Theory of Computing*, 9(18):593–615, 2013.

[Sto85] Larry Stockmeyer. On approximation algorithms for  $\#P$ . *SIAM Journal on Computing*, 14(4):849–861, 1985.

[TD02] Barbara M Terhal and David P DiVincenzo. Adaptive quantum computation, constant depth quantum circuits and Arthur-Merlin games. *arXiv:quant-ph/0205133*, 2002.

[TO92] Seinosuke Toda and Mitsunori Ogiwara. Counting classes are at least as hard as the polynomial-time hierarchy. *SIAM Journal on Computing*, 21(2):316–328, 1992.

[WB86] Lloyd R Welch and Elwyn R Berlekamp. Error correction for algebraic block codes, December 30 1986. US Patent 4,633,470.

[ZSW20] Yiqing Zhou, E Miles Stoudenmire, and Xavier Waintal. What limits the simulation of quantum computers? *Physical Review X*, 10(4):041038, 2020.

[ZWD<sup>+</sup>20] 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, 2020.

## A Relationship between truncated Taylor series and Cayley path

A line of research initiated by Aaronson and Arkhipov [AA11] aims to establish complexity-theoretic evidence for the classical hardness of quantum supremacy experiments by providing evidence for the so-called “approximate sampling” conjecture, which in the context of random circuit sampling can be summarized as follows:

**Conjecture 1** (Approximate sampling). *No polynomial-time algorithm can do the following: on input a random circuit  $C \sim \mathcal{H}_A$ , obtain a sample from a distribution  $D'_C$  that satisfies  $D_{\text{TV}}(D_C, D'_C) \leq \frac{1}{\text{poly}(n)}$ , with probability at least  $1 - \frac{1}{\text{poly}(n)}$  over the choice of  $C$  as well as the internal randomness of the algorithm. Here  $D_C$  denotes the (ideal) output distribution of  $C$ .*

One approach to proving the approximate sampling conjecture is by showing that the output probability of random circuits are  $\#P$ -hard to approximate on average. Here “approximate” corresponds to a  $\frac{2^{-n}}{\text{poly}(n)}$  additive imprecision, which follows naturally from Stockmeyer’s approximate counting theorem [Sto85].

**Conjecture 2.** *It is  $\#P$ -hard to compute  $\mathbf{p}_0(C)$  up to additive imprecision  $\frac{2^{-n}}{\text{poly}(n)}$  on input a random circuit  $C \sim \mathcal{H}_A$ , with probability at least  $1 - \frac{1}{\text{poly}(n)}$  over the choice of  $C$  as well as the internal randomness of the algorithm.*

**Theorem 8** ([BFNV19]). *Conjecture 2 implies Conjecture 1.*

Although Conjecture 2 has not been proven so far, necessary conditions of Conjecture 2 were established in the form of average-case hardness of  $\mathbf{p}_0(C)$  that is robust up to a smaller additive imprecision. An open direction is therefore to improve these results to match the  $\frac{2^{-n}}{\text{poly}(n)}$  additive imprecision required by Conjecture 2.

As introduced in Section 4, [BFNV19] developed an approach based on truncated Taylor series. Consider the circuit distribution  $\mathcal{H}_{A,\theta,K}$  from Definition 1 which is instantiated by Eq. (5). They proved the exact average-case hardness for circuits drawn from  $\mathcal{H}_{A,\theta,K}$ .

**Theorem 9** ([BFNV19]). *It is  $\#P$ -hard to exactly compute  $\frac{3}{4} + \frac{1}{\text{poly}(n)}$  of the probabilities  $\mathbf{p}_0(C')$  over the choice of  $C'$  from the distribution  $\mathcal{H}_{A,\theta,K}$  where  $\theta = \frac{1}{\text{poly}(n)}$ ,  $K = \text{poly}(n)$ .*

They also showed that the average-case hardness still holds up to an additive imprecision of  $2^{-\text{poly}(m)}$ , but did not quantify the degree of the polynomial involved<sup>25</sup>. Note that here circuits drawn from the distribution  $\mathcal{H}_{A,\theta,K}$  are not unitary and has a small truncation error from the truncation of the Taylor series in Eq. (5). Therefore, Theorem 9 (as well as the  $2^{-O(m^3)}$  robust version) does not imply the exact version of Conjecture 2 (that is, the exact average-case hardness of circuits drawn from  $\mathcal{H}_A$ ). However, as the truncation error is much smaller than the additive imprecision requirement of Conjecture 2, it is shown that the *approximate* version of Theorem 9 is actually equivalent to Conjecture 2.

**Theorem 10** ([BFNV19]). *Conjecture 2 is equivalent to the following: it is  $\#P$ -hard to compute  $\mathbf{p}_0(C')$  up to additive imprecision  $\frac{2^{-n}}{\text{poly}(n)}$  on input a random circuit  $C' \sim \mathcal{H}_{A,\theta,K}$  where  $\theta = \frac{1}{\text{poly}(n)}$ ,  $K = \text{poly}(n)$ , with probability at least  $1 - \frac{1}{\text{poly}(n)}$  over the choice of  $C'$  as well as the internal randomness of the algorithm.*

As discussed in Section 4, Movassagh’s subsequent approach based on the Cayley transform [Mov20] does not have the truncation problem, and therefore they were able to prove the exact version of Conjecture 2. Movassagh also showed a similar robust version of his theorem with quantified robustness of  $2^{-O(m^3)}$  additive imprecision. Therefore these two approaches are equivalent from the viewpoint of providing evidence for the approximate sampling conjecture: to prove Conjecture 2, it is necessary and sufficient to improve the robustness of either one of these approaches to  $\frac{2^{-n}}{\text{poly}(n)}$ .

Our first result can be viewed as making progress in this direction. Note that in the main text our proof is based on a generalization of the Cayley path approach, but our techniques can also be applied to truncated Taylor series and obtain similar results. On the other hand, for our second result on noisy circuits, the Cayley path approach is more preferable, as in this case we are no longer in the regime of providing evidence for the approximate sampling conjecture, but are interested in tiny deviations from uniformity in the output distribution of noisy circuits.

<sup>25</sup>In retrospect, setting  $\theta = \frac{1}{m^2 \text{polylog}(m)}$  and  $K = \text{polylog}(m)$  yields robustness  $2^{-\tilde{O}(m^3)}$  to computing this quantity, but these parameter settings or the arguments justifying them were not given in the original paper.

## B Limitations of polynomial interpolation techniques

Here we give a self-contained description of the barrier to proving hardness of approximate sampling using polynomial interpolation techniques. This argument was first developed by Aaronson and Arkhipov for BosonSampling [AA11, Section 9.2], which can also be naturally generalized to random circuit sampling. They show that, assuming anti-concentration (which is proven for random circuit sampling), the average-case hardness proof based on polynomial interpolation cannot tolerate the amount of error required for showing hardness of approximate sampling.

Consider the output probability of a  $n$ -photon  $m = n^c$  mode BosonSampling experiment. Following the notation in the proof of Corollary 1, the goal is to prove average-case hardness of computing  $p_X = |\text{Per}(X)|^2/m^n$  for  $X \sim \mathcal{G}^{n \times n}$ . To do so, we consider

$$X(\theta) := (1 - \theta)X_1 + \theta X_0 \quad (61)$$

where  $X_0$  is an arbitrary 0/1 matrix and  $X_1 \sim \mathcal{G}^{n \times n}$ . The goal is to recover  $|\text{Per}(X_0)|^2$  using polynomial interpolation, given approximate values of  $|\text{Per}(X(\theta))|^2$  for small values of  $\theta$ . The main idea of the barrier result is to show the following:

*There exists a polynomial that agrees well with  $|\text{Per}(X(\theta))|^2$  when  $\theta$  is small, but is far from  $|\text{Per}(X_0)|^2$  when  $\theta = 1$ .*

This statement implies that it is impossible to recover the worst-case value from noisy average-case observations. To construct such a polynomial  $P(\theta)$ , consider

$$\begin{aligned} P(\theta) &:= |\text{Per}(X(\theta))|^2 + tQ(\theta), \\ Q(\theta) &:= |\text{Per}((1 - \theta)X_1 + \theta \mathbf{1}_{n \times n})|^2 \end{aligned} \quad (62)$$

where  $\mathbf{1}$  denotes the all-one matrix and  $t > 0$ . Then the error of the polynomial  $P(\theta)$  is given by

$$\left| P(\theta) - |\text{Per}(X(\theta))|^2 \right| = t |Q(\theta)|. \quad (63)$$

Using the anti-concentration property of  $\text{Per}(X_1)$ , with high probability the error is given by

- average-case error ( $\theta \ll 1$ ):  $t |Q(\theta)| \leq t \cdot \text{poly}(n) \cdot n!$ ,
- worst-case error ( $\theta = 1$ ):  $t |Q(1)| = t (n!)^2$ .

To prove hardness of approximate sampling, we require  $t = 1/q(n)$  for a sufficiently large polynomial  $q$ . However, in this case the worst-case error becomes too large and can be achieved by existing classical algorithms [Gur05, AH14].

On the other hand, choosing  $t = 1/(n!)^2$  is sufficient for worst-case hardness. In this case we can tolerate an additive error  $\text{poly}(n)/n!$  for average-case permanents. Consider the output probability and normalize the errors by  $m^n$ , we conclude that the barrier result of Aaronson and Arkhipov does not rule out the possibility of showing hardness of computing the output probability  $p_X$  up to additive error

$$e^{-(c+1)n \log n}. \quad (64)$$

The same argument can also be applied to random circuit sampling, which shows that proving hardness of computing  $\mathsf{p}_0(C)$  up to additive error  $2^{-n}/\text{poly}(n)$ , which is required for proving hardness of sampling, is impossible using polynomial interpolation techniques. On the other hand, this barrier result does not rule out the possibility of showing hardness of computing  $\mathsf{p}_0(C)$  up to additive error  $2^{-O(n)}$ .

Interestingly, using essentially the same techniques, our result is a log factor far from the barrier for random circuit sampling, while for BosonSampling our result is only a constant factor away.

## C Worst-case hardness for noisy circuits from error detection

It is a well-known result that computing the output probability  $\mathsf{p}_0(C)$  for quantum circuits  $C$  is  $\#\mathsf{P}$ -hard in the worst case [FR99, TD02, BJS11, BMS16], where the hardness is robust under constant multiplicative approximation or exponentially small additive approximation. Here, we consider a slightly different approach based on [FGHP99]. There, it was shown that determining if  $\mathsf{p}_0(C) \neq 0$  is hard for the complexity class  $\mathsf{coC}=\mathsf{P}$ , which corresponds to deciding whether an efficiently computable Boolean function  $f : \{0, 1\}^n \rightarrow \{-1, 1\}$  satisfies

$$\sum_{x \in \{0,1\}^n} f(x) \neq 0. \quad (65)$$

Importantly, it was shown [TO92] that  $\mathsf{coC}=\mathsf{P}$  is hard for  $\mathsf{PH}$  under randomized reductions. Therefore, determining if  $\mathsf{p}_0(C) \neq 0$  is hard for  $\mathsf{PH}$  unless  $\mathsf{PH}$  collapses.

In the following, we show that we can assume an exponentially small gap in the above decision problem without loss of generality. Consider the following promise decision problem.

**Problem 2.** *On input a  $n$ -qubit quantum circuit  $C$ , decide if  $\mathsf{p}_0(C) = 0$ , or  $\mathsf{p}_0(C) \geq \frac{1}{2^{2n}}$ .*

We show that this gapped variant is still hard for  $\mathsf{PH}$  following a similar argument as in [FGHP99].

**Lemma 10.** *Problem 2 is hard for  $\mathsf{coC}=\mathsf{P}$  under a polynomial time reduction.*

*Proof.* Consider an efficiently computable Boolean function  $f : \{0, 1\}^{n_0} \rightarrow \{-1, 1\}$ . Using the standard technique for classical reversible computation, we can create a  $n$ -qubit ( $n \geq n_0$ ) unitary circuit  $U_f$  in polynomial time, such that  $U_f |x\rangle |0\rangle = f(x) |x\rangle |0\rangle$  for all  $x \in \{0, 1\}^{n_0}$ . Let  $C$  be the Fourier sampling circuit,  $C = H^{\otimes n_0} U_f H^{\otimes n_0}$ . Then we have

$$\mathsf{p}_0(C) = \frac{1}{2^{2n_0}} \left( \sum_{x \in \{0,1\}^{n_0}} f(x) \right)^2. \quad (66)$$

As  $\sum_{x \in \{0,1\}^{n_0}} f(x)$  takes integer value, we know that either  $\mathsf{p}_0(C) = 0$ , or  $\mathsf{p}_0(C) \geq \frac{1}{2^{2n_0}} \geq \frac{1}{2^{2n}}$ . Therefore, deciding Problem 2 on input  $C$  is equivalent to deciding if  $\sum_{x \in \{0,1\}^{n_0}} f(x) \neq 0$ .  $\square$

Next, we show that simulating a noisy circuit up to additive error  $\exp(-O(m \log m))$  is as hard as Problem 2, and therefore is hard for  $\mathsf{PH}$ . The proof requires the use of fault-tolerant quantum error detection. Given an arbitrary noise model  $\mathcal{N}$  in the form of Eq. (28), its local and stochastic property implies that there exists a constant  $\gamma_{\text{th}} > 0$  such that fault-tolerant quantum error detection is possible when  $\gamma < \gamma_{\text{th}}$ . Different from quantum error correction, the fault-tolerance by error detection is *conditional*, or via *post selection*. Given a logical circuit  $C_0$  of  $n_0$  qubits and  $m_0$  gates, suppose  $C$  implements  $C_0$  with an error detection code using  $n$  qubits and  $m$  gates. Suppose that  $C$  has two output registers  $x, z$ , where  $x$  denotes the logical output, and  $z$  denotes the error detection syndrome, where  $z = 0$  implies that no error is detected throughout the circuit. The fault-tolerance property guarantees that conditioned on no error being detected, the output distribution of  $C$  is close to the logical circuit  $C_0$ , that is,

$$\Pr_{C_0}[x] \approx \Pr_{C, \mathcal{N}}[x|z = 0] = \frac{\Pr_{C, \mathcal{N}}[x, z = 0]}{\Pr_{C, \mathcal{N}}[z = 0]}, \quad x \in \{0, 1\}^{n_0}. \quad (67)$$

Here  $\Pr_{C_0}$  denotes the output distribution of the logical circuit  $C_0$  and  $\Pr_{C, \mathcal{N}}$  denotes the output distribution of  $C$  with noise model  $\mathcal{N}$ .

As we eventually want to decide Problem 2 which has an exponentially small gap, we would like  $C$  to be a very accurate simulation of  $C_0$ . This can be done by choosing a code distance  $d = \text{poly}(n_0)$ : while introducing a polynomial overhead  $n = \text{poly}(n_0)$  and  $m = \text{poly}(m_0)$ , this guarantees that the total variation distance between the output distribution of  $C_0$  and the output distribution of  $C$  (conditioned on  $z = 0$ ) is bounded by  $\exp(-\text{poly}(m_0))$ . We thus have the following property, which follows from standard threshold theorems [Kit97, Rei06, AGP07, ABO08].

**Lemma 11.** *Fix a noise model  $\mathcal{N}$ , let  $\gamma_{\text{th}}$  be the corresponding fault-tolerant error detection threshold, and suppose the noise rate of  $\mathcal{N}$  satisfies  $\gamma < \gamma_{\text{th}}$ . For any logical circuit  $C_0$  of  $n_0$  qubits and  $m_0$  gates and a desired polynomial  $p$ , there exists a circuit  $C$  that can be constructed in polynomial time, with  $n = \text{poly}(n_0)$  qubits and  $m = \text{poly}(m_0)$  gates, such that*

$$\left| \mathbf{p}_0(C_0) - \frac{\mathbf{p}_0(C, \mathcal{N})}{\Pr_{C, \mathcal{N}}[z = 0]} \right| \leq \exp(-p(m_0)). \quad (68)$$

Now we are ready to prove the worst-case hardness result.

**Theorem 11.** *Fix a noise model  $\mathcal{N}$ , let  $\gamma_{\text{th}}$  be the corresponding fault-tolerant error detection threshold, and suppose the noise rate of  $\mathcal{N}$  satisfies  $\gamma < \gamma_{\text{th}}$ . It is  $\text{coC}_=\text{P}$ -hard to compute  $\mathbf{p}_0(C, \mathcal{N})$  on input  $C$  with  $m$  gates, up to additive error  $\exp(-O(m \log m))$ .*

*Proof.* We prove a polynomial time reduction from Problem 2. On input a logical circuit  $C_0$  of  $n_0$  qubits and  $m_0$  gates, we would like to decide if  $\mathbf{p}_0(C_0) = 0$ , or  $\mathbf{p}_0(C_0) \geq \frac{1}{2^{2n_0}}$ . Choose  $p$  such that  $\exp(-p(m_0)) \ll \frac{1}{2^{2n_0}}$ , and construct a circuit  $C$  with  $n = \text{poly}(n_0)$  qubits and  $m = \text{poly}(m_0)$  gates according to Lemma 11. Consider the probability  $\Pr_{C, \mathcal{N}}[z = 0]$  of detecting no errors. We bound it in two directions. First,

$$\Pr_{C, \mathcal{N}}[z = 0] \geq \Pr[\text{no error occurs}] = \exp(-O(\gamma m)). \quad (69)$$

On the other hand,

$$\begin{aligned} \Pr_{C, \mathcal{N}}[z = 0] &= \Pr[\text{no error occurs}] + \Pr[\text{some error occurs but undetected}] \\ &\leq \exp(-O(\gamma m)) (1 + \exp(-\text{poly}(m_0))), \end{aligned} \quad (70)$$

where the second term takes into account both the probability of the error pattern as well as the probability of beating the code distance (also see [Fuj16]). Next, consider the two cases of Problem 2,

1. When  $\mathbf{p}_0(C_0) = 0$ , we have

$$\begin{aligned} \mathbf{p}_0(C, \mathcal{N}) &\leq \exp(-p(m_0)) \Pr_{C, \mathcal{N}}[z = 0] \\ &\leq \exp(-p(m_0)) \exp(-O(\gamma m)). \end{aligned} \quad (71)$$

2. When  $\mathbf{p}_0(C_0) \geq \frac{1}{2^{2n_0}}$ , we have

$$\begin{aligned} \mathbf{p}_0(C, \mathcal{N}) &\geq \left( \frac{1}{2^{2n_0}} - \exp(-p(m_0)) \right) \Pr_{C, \mathcal{N}}[z = 0] \\ &\geq \left( \frac{1}{2^{2n_0}} - \exp(-p(m_0)) \right) \exp(-O(\gamma m)). \end{aligned} \quad (72)$$

The above bounds can be explicitly computed. Moreover, they have a large gap, which is at least

$$O\left(\frac{1}{2^{2n_0}}\right) \exp(-O(\gamma m)) \gg \exp(-O(m \log m)). \quad (73)$$

Therefore, computing  $\mathsf{p}_0(C, \mathcal{N})$  up to additive error  $\exp(-O(m \log m))$  is sufficient for deciding Problem 2.  $\square$

Using the fact that  $\text{coC-P}$  is hard for  $\text{PH}$  under randomized reductions [TO92], Theorem 5 is a direct corollary of Theorem 11.

## D Proof of Lemma 4

The proof uses techniques that are developed by Harrow and Low [HL09], which were developed to show that random quantum circuits converge to unitary 2-design. Below we show that noisy circuits correspond to entirely different convergence behavior. These techniques can also be applied to analyze local random circuits, which we leave as important future work.

Consider the quantity  $CP$  (collision probability/classical purity) defined as follows,

$$CP = 2^n \mathbb{E}_{U \sim \text{GRC}(n, d)} \left[ \sum_x p_e^2(x) \right] - 1. \quad (74)$$

Here  $\text{GRC}(n, d)$  denotes the global random circuit distribution on  $n$  qubits with depth  $d$ , as shown in Fig. 1b, and  $p_e$  denotes the output distribution of the noisy circuit. We will drop the subscript  $U \sim \text{GRC}(n, d)$  unless necessary. For simplicity, here we consider i.i.d. depolarizing noise, where each blue dot represents a single qubit noise channel

$$\mathcal{E}(\rho) = (1 - \gamma)\rho + \gamma \frac{I}{2} \text{Tr}[\rho], \quad (75)$$

where  $\gamma$  is the noise rate and  $I$  is the identity matrix. Using Cauchy–Schwarz inequality, it is easy to see that for each specific circuit, the classical purity is an upper bound of the total variation distance, as

$$D_{\text{TV}}(p_e, p_{\text{uniform}}) = \frac{1}{2} \sum_x \left| p_e(x) - \frac{1}{2^n} \right| \leq \frac{1}{2} \sqrt{2^n \sum_x p_e^2(x) - 1}. \quad (76)$$

To see why we use  $CP$  to represent the distance, notice that  $CP$  is a linear function of the two-copy output density matrix  $\rho_e \otimes \rho_e$ . Therefore, by linearity of expectation, we can consider an initial state  $|0^n\rangle\langle 0^n|^{\otimes 2}$  going through an average quantum channel; that is, to understand the convergence behavior of  $CP$ , it suffices to track the evolution of the second moment of the density matrix  $\mathbb{E}[\rho_e \otimes \rho_e]$  as the circuit depth increases. More specifically, we can write any Hermitian operator in the Fourier basis as

$$\mathbb{E}[\rho_e \otimes \rho_e] = \frac{1}{2^n} \sum_{p,q} \delta(p, q) \sigma_p \otimes \sigma_q, \quad (77)$$

where  $p, q \in \{0, 1, 2, 3\}^n$  are  $2n$ -bit strings,  $\sigma_p, \sigma_q \in \{I, X, Y, Z\}^{\otimes n}$  are the corresponding Pauli operators indexed by  $0 \rightarrow I, 1 \rightarrow X, 2 \rightarrow Y, 3 \rightarrow Z$ , and

$$\delta(p, q) = \frac{1}{2^n} \text{Tr}[\mathbb{E}[\rho_e \otimes \rho_e] \sigma_p \otimes \sigma_q] = \frac{1}{2^n} \mathbb{E} \text{Tr}[(\rho_e \otimes \rho_e) \sigma_p \otimes \sigma_q] \quad (78)$$

are the Fourier coefficients, where the second equality follows from linearity.

As we take expectation over random quantum circuits  $U \sim \text{GRC}(n, d)$ , the Fourier coefficients  $\delta(p, q)$  only depend on  $(\mathcal{E}, n, d)$ , i.e. noise channel, system size, and circuit depth, and they contain all the information needed to evaluate  $CP$ . More specifically we have

$$CP = 2^n \mathbb{E} \left[ \sum_x p_e^2(x) \right] - 1 = 2^n \sum_{p \in \{0,3\}^n \setminus \{0^n\}} \delta(p, p) = 2^n \sum_{p \in \{0,3\}^n \setminus \{0^n\}} \alpha(p), \quad (79)$$

where we have defined  $\alpha(p) := \delta(p, p)$ .

In the following we will study the evolution of these coefficients  $\delta(p, q)$  as gates and noise are applied. For the initial state  $|0^n\rangle\langle 0^n|^{ \otimes 2}$ , we have  $\delta(p, q) = \frac{1}{2^n}$  for any  $p, q \in \{0, 3\}^n$  and 0 otherwise. Next, consider the application of a global random unitary on the Pauli basis, it was shown [HL09] that

$$\begin{aligned} & \mathbb{E}_{U \sim \mathbb{U}(2^n)} \left[ U \sigma_1 U^\dagger \otimes U \sigma_2 U^\dagger \right], \quad \sigma_1, \sigma_2 \in \{I, X, Y, Z\}^{\otimes n} \\ &= \begin{cases} 0, & \sigma_1 \neq \sigma_2, \\ I^{\otimes 2n}, & \sigma_1 = \sigma_2 = I^{\otimes n}, \\ \frac{1}{2^{2n}-1} \left( \sum_{p \in \{0,1,2,3\}^n \setminus \{0^n\}} \sigma_p \otimes \sigma_p \right), & \sigma_1 = \sigma_2 \neq I^{\otimes n}. \end{cases} \end{aligned} \quad (80)$$

Therefore the off-diagonal coefficients  $\delta(p, q)$  ( $p \neq q$ ) becomes 0 for depth  $d \geq 1$ . From now on we only consider diagonal coefficients  $\alpha(p) = \delta(p, p)$ . The evolution of  $\alpha$  under a global random unitary is therefore given by

$$\alpha'(p) = \begin{cases} \alpha(p), & p = 0^n, \\ \frac{1}{2^{2n}-1} \sum_{q \in \{0,1,2,3\}^n \setminus \{0^n\}} \alpha(q), & p \neq 0^n. \end{cases} \quad (81)$$

It is easy to observe the following properties: (1)  $\alpha(0^n) = \frac{1}{2^n}$  is a constant; (2) for  $p \neq 0^n$ ,  $\alpha'(p)$  does not depend on  $p$ ; (3)  $\sum_{p \in \{0,1,2,3\}^n} \alpha(p)$  is preserved by unitary evolution.

Finally, consider the effect of depolarizing noise on a single qubit, which gives

$$\begin{aligned} I \otimes I &\rightarrow I \otimes I, \\ X \otimes X &\rightarrow (1 - \gamma)^2 X \otimes X, \\ Y \otimes Y &\rightarrow (1 - \gamma)^2 Y \otimes Y, \\ Z \otimes Z &\rightarrow (1 - \gamma)^2 Z \otimes Z. \end{aligned} \quad (82)$$

Therefore the effect of noise can be summarized as

$$\alpha(p) \rightarrow (1 - \gamma)^{2|p|} \alpha(p), \quad (83)$$

where  $|p|$  denotes the Hamming weight of  $p$ , i.e. the number of non-zero terms.

Combining everything together, we can write the evolution of  $\sum_{p \in \{0,1,2,3\}^n \setminus \{0^n\}} \alpha(p)$  in one cycle (i.e. one global gate followed by a layer of noise channels) as

$$\sum_{p \in \{0,1,2,3\}^n \setminus \{0^n\}} \alpha(p) \rightarrow \frac{\sum_{p \in \{0,1,2,3\}^n \setminus \{0^n\}} (1 - \gamma)^{2|p|}}{2^{2n} - 1} \sum_{p \in \{0,1,2,3\}^n \setminus \{0^n\}} \alpha(p). \quad (84)$$

Define the decay coefficient as

$$\beta = \frac{\sum_{p \in \{0,1,2,3\}^n \setminus \{0^n\}} (1 - \gamma)^{2|p|}}{2^{2n} - 1} = \frac{\sum_{k=1}^n \binom{n}{k} 3^k (1 - \gamma)^{2k}}{2^{2n} - 1} = \frac{(1 + 3(1 - \gamma)^2)^n - 1}{2^{2n} - 1}. \quad (85)$$

Therefore at depth  $d$  we have

$$\sum_{p \in \{0,1,2,3\}^n \setminus \{0^n\}} \alpha(p) = \frac{2^n - 1}{2^n} \beta^d. \quad (86)$$

To compute  $CP$  at depth  $d$ , note that we can first compute  $\sum_{p \in \{0,1,2,3\}^n \setminus \{0^n\}} \alpha(p)$  at depth  $d - 1$ . These coefficients are then uniformly redistributed by the global unitary, and then decay according to noise. Overall we have

$$\begin{aligned} CP &= 2^n \sum_{p \in \{0,3\}^n \setminus \{0^n\}} \alpha(p) \\ &= 2^n \sum_{p \in \{0,3\}^n \setminus \{0^n\}} (1 - \gamma)^{2|p|} \frac{\frac{2^n - 1}{2^n} \beta^{d-1}}{2^{2n} - 1} \\ &= \frac{(1 + (1 - \gamma)^2)^n - 1}{2^n + 1} \left( \frac{(1 + 3(1 - \gamma)^2)^n - 1}{2^{2n} - 1} \right)^{d-1}. \end{aligned} \quad (87)$$

To obtain an upper bound of  $CP$ , we have

$$\begin{aligned} CP &= \frac{(1 + (1 - \gamma)^2)^n - 1}{2^n + 1} \left( \frac{(1 + 3(1 - \gamma)^2)^n - 1}{2^{2n} - 1} \right)^{d-1} \\ &\leq \left( 1 - \left( \gamma - \frac{\gamma^2}{2} \right) \right)^n \cdot \left( 1 - \left( \frac{3\gamma}{2} - \frac{3\gamma^2}{4} \right) \right)^{nd-n} \\ &= \exp(-\Omega(\gamma n(d - d_0))). \end{aligned} \quad (88)$$

where  $d_0$  is a universal constant.

## E Random unitary and Cayley transform

In this section we introduce the Cayley transform [Mov20] and establish some useful properties. These properties are originally proved in [Mov20], and here we give a self-consistent presentation as we are using different notations.

We consider the Haar measure  $\mu_H$  over the unitary group  $\mathbb{U}(N)$  of  $N \times N$  unitary matrices. Any unitary matrix can be diagonalized as  $U = V\Lambda V^\dagger$ , where  $V$  is unitary and  $\Lambda = \text{diag}(e^{i\varphi_1}, \dots, e^{i\varphi_N})$ ,  $\varphi_j \in [-\pi, \pi]$ . From now on we use ‘‘eigenvalues’’ to refer to the angles  $\varphi_j$ . Consider the following transform for the eigenvalues,

$$\begin{aligned} f_\theta : [-\pi, \pi] &\rightarrow [-\pi, \pi], \\ f_\theta(\varphi) &= 2 \arctan \left( (1 - \theta) \tan \frac{\varphi}{2} \right). \end{aligned} \quad (89)$$

It is easy to verify that applying  $f_\theta$  to all eigenvalues of a unitary matrix is equivalent to applying the Cayley transform as given in Definition 2.

Define  $\mu_\theta$  ( $\theta \geq 0$ ) as the probability measure on  $\mathbb{U}(N)$  induced by applying  $f_\theta$  to all eigenvalues of the unitary matrix drawn from  $\mu_H$ . From the definition it is easy to see that  $\mu_0 = \mu_H$  and  $\mu_1$  is the indicator for the identity matrix.

In the following we establish the bound of the total variation distance between  $\mu_\theta$  and  $\mu_H$  when  $\theta$  is a small positive number.

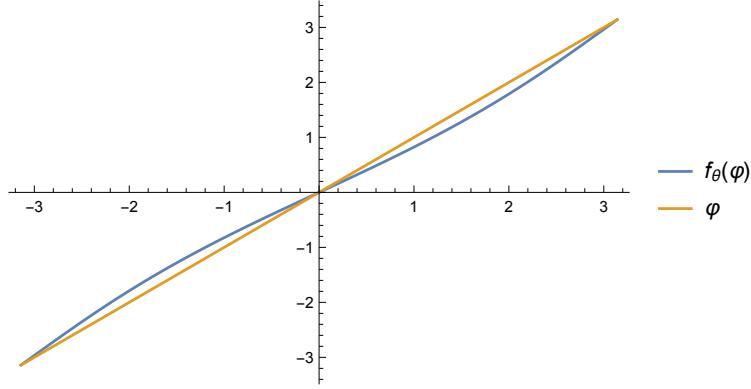


Figure 2: Plot of the eigenvalue transform  $f_\theta(\varphi)$  for  $\theta = 0.2$ .

**Lemma 12.** For  $N = O(1)$  and  $\theta \in (0, 0.1)$ , the total variation distance between  $\mu_\theta$  and  $\mu_H$  is bounded by

$$D_{\text{TV}}(\mu_\theta, \mu_H) = \int_{\mathbb{U}(N)} \frac{1}{2} \left| 1 - \frac{d\mu_\theta}{d\mu_H} \right| d\mu_H = O(\theta). \quad (90)$$

*Proof.* First, notice that  $\mu_\theta$  and  $\mu_H$  has the same eigenvector distribution. As the eigenvalues and eigenvectors are independently distributed in  $\mu_H$ , the integral only concerns the eigenvalue density.

To see this, we use the two-sided invariance of the Haar measure. Let  $U \sim_{\mu_H} \mathbb{U}(N)$  be expressed as  $U = V \Lambda V^\dagger$ , where  $V$  is the eigenvectors and  $\Lambda$  is the eigenvalues. For any unitary  $Q \in \mathbb{U}(N)$ ,  $Q U Q^\dagger$  is equal to  $U$  in distribution. Therefore,  $Q U Q^\dagger$  is also equal to  $U$  in distribution when  $Q$  is independently drawn from  $\mu_H$ . As  $Q U Q^\dagger = (QV)\Lambda(QV)^\dagger$ , the eigenvectors  $V$  is distributed the same as  $QV$  up to global phase, which is distributed the same as  $Q$ , which is independent from  $\Lambda$ . To summarize, the eigenvectors and eigenvalues of Haar random unitary are independently distributed, where the eigenvectors are distributed the same as  $\mu_H$  up to global phase, and we quote the eigenvalue density from [HP06]:

$$g_H(\varphi_1, \dots, \varphi_N) = \frac{1}{(2\pi)^N N!} \prod_{1 \leq k < l \leq N} |e^{i\varphi_k} - e^{i\varphi_l}|^2 \quad (91)$$

where  $\varphi_i \in [-\pi, \pi]$ . Let  $g_\theta$  be the eigenvalue density of  $\mu_\theta$ . The total variation distance is then reduced to the following integral over Lebesgue measure on  $[-\pi, \pi]^N$ ,

$$D_{\text{TV}}(\mu_\theta, \mu_H) = \int_{[-\pi, \pi]^N} \frac{1}{2} |g_\theta(\varphi_1, \dots, \varphi_N) - g_H(\varphi_1, \dots, \varphi_N)| d\varphi_1 \cdots d\varphi_N. \quad (92)$$

An example of the transform  $f_\theta$  is plotted in Fig. 2. It is easy to prove that  $f_\theta$  is monotonic and has range  $[-\pi, \pi]$ .

To prove the upper bound of the total variation distance, we need to calculate the transformed eigenvalue density  $g_\theta(\varphi_1, \dots, \varphi_N)$ , by applying  $f_\theta$  to  $\varphi_1, \dots, \varphi_N$  jointly distributed according to  $g_H$ . Then  $g_\theta$  can be written as

$$g_\theta(\varphi_1, \dots, \varphi_N) = g_H(f_\theta^{-1}(\varphi_1), \dots, f_\theta^{-1}(\varphi_N)) \prod_{i=1}^N \left| \frac{df_\theta^{-1}(\varphi_i)}{d\varphi_i} \right|, \quad \forall \varphi_i \in [-\pi, \pi]. \quad (93)$$

where the inverse transform of  $f_\theta$  is given by

$$f_\theta^{-1}(\varphi) = 2 \arctan \left( \frac{1}{1-\theta} \tan \frac{\varphi}{2} \right) \quad (94)$$

and its derivative given by

$$\frac{df_\theta^{-1}(\varphi)}{d\varphi} = \frac{1-\theta}{1-(2\theta-\theta^2)\cos^2 \frac{\varphi}{2}}. \quad (95)$$

It is easy to prove the following bounds by a simple calculation: assume  $\theta$  is small enough (say  $\theta < 0.1$ ), then for all  $\varphi \in [-\pi, \pi]$ ,

$$\begin{aligned} |f_\theta^{-1}(\varphi) - \varphi| &\leq O(\theta) \\ \left| \frac{df_\theta^{-1}(\varphi)}{d\varphi} - 1 \right| &\leq O(\theta), \end{aligned} \quad (96)$$

where the constants in the above two equations is independent of  $\varphi$ . The proof then follows by directly bounding the difference between the densities. Consider

$$\begin{aligned} &|g_\theta(\varphi_1, \dots, \varphi_N) - g_H(\varphi_1, \dots, \varphi_N)| \\ &= \left| g_H(f_\theta^{-1}(\varphi_1), \dots, f_\theta^{-1}(\varphi_N)) \prod_{i=1}^N \left| \frac{df_1^{-1}(\varphi_i)}{d\varphi_i} \right| - g_H(\varphi_1, \dots, \varphi_N) \right| \\ &\leq \left| \left( g_H(\varphi_1, \dots, \varphi_N) + \sum_i O(\theta) \right) \prod_{i=1}^N (1 + O(\theta)) - g_H(\varphi_1, \dots, \varphi_N) \right| \\ &\leq O(\theta), \end{aligned} \quad (97)$$

where the third line follows from Eq. (96) and the continuity of  $g_H$ . Finally the integral of the difference between the densities is still bounded by  $O(\theta)$ , which concludes the full proof.  $\square$

Note that in the matrix form of the Cayley transform

$$H(\theta) = \sum_j \frac{1+i(1-\theta)\tan \frac{\varphi_j}{2}}{1-i(1-\theta)\tan \frac{\varphi_j}{2}} |\psi_j\rangle\langle\psi_j|, \quad (98)$$

where  $H = \sum_j e^{i\varphi_j} |\psi_j\rangle\langle\psi_j|$ , the magnitude of the denominator in the above equation goes to infinity as  $\varphi_j$  is close to  $\pm\pi$ , which will affect the accuracy of the polynomial interpolation in our worst-to-average-case reduction. However, the probability of such an event is small when  $H$  is drawn from  $\mu_H$ , as shown in the following lemma.

**Lemma 13.** *The Haar distribution  $\mu_H$  over the unitary group  $\mathbb{U}(N)$  satisfies*

$$\Pr_{\mu_H} [\varphi_j \in [-\pi + \delta, \pi - \delta], \forall j] \geq 1 - \frac{N\delta}{\pi}, \quad (99)$$

where  $\varphi_j$  ( $j = 1 \dots N$ ) denotes the eigenvalues of the random unitary distributed according to  $\mu_H$ .

*Proof.* The proof is a simple union bound:

$$\begin{aligned} \Pr_{\mu_H} [\varphi_j \in [-\pi + \delta, \pi - \delta], \forall j] &= 1 - \Pr_{\mu_H} [\exists j : \varphi_j \notin [-\pi + \delta, \pi - \delta]] \\ &\geq 1 - \sum_{j=1}^N \Pr_{\mu_H} [\varphi_j \notin [-\pi + \delta, \pi - \delta]] \\ &= 1 - N \Pr_{\mu_H} [\varphi_1 \notin [-\pi + \delta, \pi - \delta]] \\ &= 1 - \frac{N\delta}{\pi}. \end{aligned} \quad (100)$$

$\square$