# The Power of Quantum Circuits in Sampling

Guy Blanc *Stanford*
Caleb Koch *Stanford*
Jane Lange *MIT*
Carmen Strassle *Stanford*
Li-Yang Tan *Stanford*

February 2, 2026

## 1 Introduction

A central goal of quantum complexity theory is to understand the relative power of quantum and classical circuits. For circuits computing functions $f : \{0,1\}^n \to \{0,1\}$ this is the question of whether $\mathsf{BQP/poly} = \mathsf{P/poly}$. In this paper we are interested in the complexity of *distributions*, rather than functions, over $\{0,1\}^n$. A standard way of measuring the complexity of a distribution is by the complexity of sampling from it: the size of the smallest circuit that generates samples from it. The most basic question in this setting, as in others, is whether quantumness buys us any appreciable power at all:

**Question 1** (Does every quantum sampler have a classical counterpart?)**.** Does every polynomial-size quantum circuit $C$ have a corresponding polynomial-size classical circuit whose output distribution well-approximates that of $C$'s?

As our main result, we answer Question 1 relative to a random oracle:

**Theorem 1.** *The following holds relative to a random oracle. For every sufficiently large $n$, there is an explicit distribution $\mathcal{D}^{(n)}$ over $\{0,1\}^n$ that can be sampled exactly by a polynomial-size quantum circuit but no subexponential-size classical circuit can approximate $\mathcal{D}^{(n)}$ even to TV distance $1 - 2^{-\tilde{\Omega}(\sqrt{n})}$.*

Theorem 1 provides evidence that quantum circuits are far more powerful than classical ones: In sampling, quantumness can buy us both a dramatic reduction in circuit size as well as a dramatic improvement in accuracy. Furthermore, these advantages hold not just for specific, structured instances (that may have been tailored for a separation), but even for generic, unstructured ones.

The analogue of Theorem 1 for functions remains open: It is not known if $\mathsf{BQP/poly} \neq \mathsf{P/poly}$ relative to a random oracle, and there are formal barriers to such a separation [AA14].

**Exact vs. approximate sampling.** Previously Aaronson and Arkhipov showed [AA11, Theorem 34], under the assumption that the polynomial hierarchy ($\mathsf{PH}$) is infinite, that there is a distribution—the "complete" boson distribution $\mathcal{D}^{(n)}$—that can be sampled exactly by a polynomial-size quantum circuit but cannot be sampled exactly by polynomial-size classical circuits.

Our results are incomparable. On one hand, [AA11]'s holds under the assumption that $\mathsf{PH}$ is infinite whereas Theorem 1 is a relativized statement. On the other hand, [AA11]'s classical hardness only holds for exact sampling whereas Theorem 1's holds for approximate sampling, and even with near-maximal TV distance. Aaronson and Arkhipov pointed to this aspect of their result

as its central drawback: "Why are we so worried about this issue? One obvious reason is that noise, decoherence, photon losses, etc. will be unavoidable features in any real implementation of a boson computer. As a result, not even the boson computer *itself* can sample exactly from the distribution $\mathcal{D}$! So it seems arbitrary and unfair to require this of a classical simulation algorithm."

There are several obstacles to such an approximate version of [AA11]'s result; we recap and discuss them in the full version. We take a different approach and work with a different hard distribution. Our work therefore does not carry any immediate implications for the complexity of boson sampling.

**Separations for constant-depth circuits.** Bravyi, Gosset, and König [BGK18] raised the analogue of Question 1 for constant-depth circuits (specifically, $\mathsf{QNC}^0$ and $\mathsf{NC}^0$), calling it a challenging open problem. Recent works [BP23, KOW24] answer this analogue, and in contrast to our result and [AA11]'s, these separations are unrelativized and unconditional. This mirrors the state of affairs for the circuit complexity of functions, where we have unconditional lower bounds against $\mathsf{NC}^0$ and $\mathsf{AC}^0$ [FSS81, Ajt83, Yao85, Hås86] but remain limited to relativized or conditional lower bounds against $\mathsf{P/poly}$.

**Representational vs. algorithmic quantum advantage.** Question 1 is representational in nature: it asks about the *existence* of classical counterparts of quantum samplers. This stands in contrast to the well-studied algorithmic question of *constructing* a classical counterpart. We now elaborate on the distinction and discuss why separations in the algorithmic setting do not answer Question 1.

A sampling problem is defined by a map from instances $C \in \{0,1\}^*$ to distributions $\mathcal{D}_C$. A classical (resp. quantum) algorithm that solves the sampling problem receives $C$ as input and constructs a classical (resp. quantum) sampler for $\mathcal{D}_C$.[1] The canonical sampling problem in this context, Circuit Sampling, has $C$ being the description of a quantum circuit and $\mathcal{D}_C$ being its output distribution. Circuit Sampling trivially admits an efficient quantum algorithm, and the question is whether it admits an efficient classical algorithm. An especially well-studied variant is Random Circuit Sampling [BIS+18, BFNV19], the average-case version where $\boldsymbol{C}$ is drawn according to a suitable distribution over quantum circuits.

The complexity of a sampling problem is determined jointly by the circuit complexity of sampling from $\mathcal{D}_C$ and the complexity of *constructing* a sampler for $\mathcal{D}_C$ from $C$. An efficient algorithm for the sampling problem implies the existence of small samplers for $\mathcal{D}_C$ for every instance $C$, but the converse does not necessarily hold. Consequently, quantum advantage for sampling problems conflate the following possibilities:

> **Possibility 1:** There is a polynomial-size quantum sampler with no equivalent polynomial-size classical counterpart.
>
> **Possibility 2:** Every polynomial-size quantum sampler has an equivalent polynomial-size classical counterpart, but it is hard to construct such a counterpart efficiently.

Both are interesting statements, but ideally we would isolate which is true. Separations such as Theorem 1 and [AA11]'s result isolate the former, thereby addressing Question 1.

[1] Alternatively, we can also define the algorithm to be one that receives $C$ as input and generates a random sample from $\mathcal{D}_C$. These definitions are equivalent by Cook–Levin, for the same reason that $\mathsf{P} = \mathsf{P}$-uniform $\mathsf{P/poly}$.

Another well-studied problem is Fourier Sampling [Aar10, AA15]. This problem is specified by an oracle ensemble $\{\mathcal{O}^{(n)} : \{0,1\}^n \to \{0,1\}\}_{n\in\mathbb{N}}$ and $\mathcal{D}^{(n)}$ is the Fourier distribution of $\mathcal{O}^{(n)}$. This problem is quantumly easy with a single query to $\mathcal{O}^{(n)}$. Existing random-oracle lower bounds apply to approximate sampling [AA15, AC17], but only rule out classical uniform algorithms (i.e. Turing machines), not non-uniform ones (i.e. circuits). These lower bounds therefore again leave open the possibility that the $\mathcal{D}^{(n)}$'s of Fourier Sampling all have small classical circuits for all oracle ensembles, and such circuits are just hard to construct. See the full version for an in-depth discussion, including why salting [CDGS18] does not apply in this context.

## 2 Technical Overview

**Yamakawa–Zhandry.** Our starting point is the work of Yamakawa and Zhandry [YZ22] who introduced an NP *search* problem, NullCodeword, and a quantum vs. classical query separation for it. In NullCodeword, $\mathcal{C} \subseteq \Sigma^n$ is a code (satisfying certain list-recoverable and list-decodable properties) and the algorithm is given query access to an oracle $\mathcal{O} : \Sigma \to \{0,1\}$. Its goal is to find a codeword whose coordinates all map to 0 under $\mathcal{O}$, i.e. an element of the set: $\mathcal{C}_{\mathcal{O}} \coloneqq \{c \in \mathcal{C} \colon \mathcal{O}(c_1) = \cdots = \mathcal{O}(c_n) = 0\}$. [YZ22] gave a $\mathrm{poly}(n)$-query quantum algorithm that solves NullCodeword w.h.p. relative to a random oracle $\boldsymbol{\mathcal{O}}$. In fact, they gave a $\mathrm{poly}(n)$-query quantum algorithm that samples $\mathrm{Unif}(\mathcal{C}_{\boldsymbol{\mathcal{O}}})$ w.h.p. relative to $\boldsymbol{\mathcal{O}}$. On the other hand, they proved that every $2^{n^{\Omega(1)}}$-query classical algorithm for NullCodeword fails w.h.p. relative to $\boldsymbol{\mathcal{O}}$.

**Overview of our approach.** Setting aside the distinction between search and sampling for now, recall that lower bounds against $q$-query algorithms straightforwardly give oracle lower bounds against all time-$q$ uniform algorithms, but not all size-$q$ circuits. For the former, it suffices to identify, for every time-$q$ uniform algorithm, a sufficiently large input length $n$ on which it fails. For the latter we need a stronger statement with the order of quantifiers switched: There is a sufficiently large $n$ on which *all* size-$q$ circuits fail.

One way of achieving the latter is to prove that $q$-query algorithms succeed with *tiny* probability ($\leq q^{-\Omega(q)}$), so small that we can afford a union bound over all size-$q$ circuits. NullCodeword cannot satisfy this: The trivial algorithm that simply outputs a uniform random element of $\Sigma^n$ succeeds with probability $\geq 2^{-\Theta(n)}$. We want to be able to handle $q$ being any $\mathrm{poly}(n)$, ideally even $2^{n^{\Omega(1)}}$.

We obtain our result in two steps, the first of which is a hardness amplification lemma for NullCodeword. We consider a variant that we call *$k$-fold* NullCodeword (and denote as $\text{NullCodeword}^{\otimes k}$), where the goal is to output $k$ distinct elements of $\mathcal{C}_{\boldsymbol{\mathcal{O}}}$. We show that this is indeed a much harder problem than NullCodeword—by choosing $k$ to be sufficiently large, we can drive the success probability of any $q$-query algorithm down to $\leq q^{-\Omega(q)}$. We then show how this extreme hardness of $\text{NullCodeword}^{\otimes k}$ translates into similar hardness of sampling from the solutions space of NullCodeword, i.e. sampling $\mathrm{Unif}(\mathcal{C}_{\boldsymbol{\mathcal{O}}})$. In preserving the tiny $\leq q^{-\Omega(q)}$ success probability of $q$-query algorithms, we are able to union bound over all size-$q$ circuits and show that they must all fail at sampling $\mathrm{Unif}(\mathcal{C}_{\boldsymbol{\mathcal{O}}})$. This along with the quantum easiness of sampling $\mathrm{Unif}(\mathcal{C}_{\boldsymbol{\mathcal{O}}})$ shown in [YZ22] yields Theorem 1.

## References


[AA11] Scott Aaronson and Alex Arkhipov. The computational complexity of linear optics. In *Proceedings of the 43rd Annual ACM Symposium on Theory of computing (STOC)*, pages 333–342, 2011. 1, 1, 1

[AA14] Scott Aaronson and Andris Ambainis. The need for structure in quantum speedups. *Theory of Computing*, 10(6):133–166, 2014. 1

[AA15] Scott Aaronson and Andris Ambainis. Forrelation: A problem that optimally separates quantum from classical computing. In *Proceedings of the 47th Annual ACM Symposium on Theory of Computing (STOC)*, pages 307–316, 2015. 1

[Aar10] Scott Aaronson. BQP and the polynomial hierarchy. In *Proceedings of the 42nd Annual ACM Symposium on Theory of Computing (STOC)*, pages 141–150, 2010. 1

[AC17] Scott Aaronson and Lijie Chen. Complexity-theoretic foundations of quantum supremacy experiments. In *Proceedings of the 32nd Computational Complexity Conference (CCC)*, pages 22–1, 2017. 1

[Ajt83] Miklós Ajtai. $\Sigma^1_1$-formulae on finite structures. *Annals of Pure and Applied Logic*, 24(1):1–48, 1983. 1

[BFNV19] Adam Bouland, Bill Fefferman, Chinmay Nirkhe, and Umesh Vazirani. Quantum supremacy and the complexity of random circuit sampling. *Proceedings of the 10th Innovations in Theoretical Computer Science conference (ITCS)*, 2019. 1

[BGK18] Sergey Bravyi, David Gosset, and Robert König. Quantum advantage with shallow circuits. *Science*, 362(6412):308–311, 2018. 1

[BIS+18] Sergio Boixo, Sergei Isakov, Vadim Smelyanskiy, Ryan Babbush, Nan Ding, Zhang Jiang, Michael Bremner, John Martinis, and Hartmut Neven. Characterizing quantum supremacy in near-term devices. *Nature Physics*, 14(6):595–600, 2018. 1

[BP23] Adam Bene Watts and Natalie Parham. Unconditional quantum advantage for sampling with shallow circuits. *arXiv preprint arXiv:2301.00995*, 2023. 1

[CDGS18] Sandro Coretti, Yevgeniy Dodis, Siyao Guo, and John Steinberger. Random oracles and non-uniformity. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT)*, pages 227–258, 2018. 1

[FSS81] Merrick Furst, James Saxe, and Michael Sipser. Parity, circuits, and the polynomial-time hierarchy. In *Proceedings of the 22nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 260–270, 1981. 1

[Hås86] Johan Håstad. *Computational Limitations for Small Depth Circuits*. MIT Press, Cambridge, MA, 1986. 1

[KOW24] Daniel Kane, Anthony Ostuni, and Kewen Wu. Locality bounds for sampling hamming slices. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing (STOC)*, pages 1279–1286, 2024. 1

[Yao85] Andrew Yao. Separating the polynomial-time hierarchy by oracles. In *Proceedings of the 26th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1–10, 1985. 1

[YZ22] Takashi Yamakawa and Mark Zhandry. Verifiable quantum advantage without structure. In *Proceedings of the 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 69–74, 2022. 2, 2