From estimation of quantum probabilities to simulation of quantum circuits

Hakop Pashayan¹, Stephen D. Bartlett¹, and David Gross²

¹Centre for Engineered Quantum Systems, School of Physics, The University of Sydney, Sydney, NSW 2006, Australia

²Institute for Theoretical Physics, University of Cologne, Germany

24 December 2019

Investigating the classical simulability of quantum circuits provides a promising avenue towards understanding the computational power of quantum systems. Whether a class of quantum circuits can be efficiently simulated with a probabilistic classical computer, or is provably hard to simulate, depends quite critically on the precise notion of “classical simulation” and in particular on the required accuracy. We argue that a notion of classical simulation, which we call EPSILON-simulation (or ϵ -simulation for short), captures the essence of possessing “equivalent computational power” as the quantum system it simulates: It is statistically impossible to distinguish an agent with access to an ϵ -simulator from one possessing the simulated quantum system. We relate ϵ -simulation to various alternative notions of simulation predominantly focusing on a simulator we call a *poly-box*. A poly-box outputs $1/\text{poly}$ precision additive estimates of Born probabilities and marginals. This notion of simulation has gained prominence through a number of recent simulability results. Accepting some plausible computational theoretic assumptions, we show that ϵ -simulation is strictly stronger than a poly-box by showing that IQP circuits and unconditioned magic-state injected Clifford circuits are both hard to ϵ -simulate and yet admit a poly-box. In contrast, we also show that these two notions are equivalent under an additional assumption on the sparsity of the output distribution (*poly-sparsity*).

1 Introduction and summary of main results

Which quantum processes can be efficiently simulated using classical resources is a fundamental and longstanding problem [1, 2, 3, 4, 5, 6]. Research in this area can be split into two broad classes: results showing the hardness of efficient classical simulation for certain quantum processes, and the development of efficient classical algorithms for simulating other quantum processes. Recently, there has been substantial activity on both sides of this subject. Works on boson sampling [7], instantaneous quantum polynomial (IQP) circuits [8, 9], various translationally invariant spin models [10, 11], quantum Fourier sampling [12], one clean qubit (also known as DQC1) circuits [13, 14], chaotic quantum circuits [15] and conjugated Clifford circuits [16] have focused on showing the difficulty of classically simulating these quantum circuits. On the other hand, there has been substantial recent progress in classically simulating various elements of quantum systems including matchgate circuits with generalized inputs and measurements [17] (see also [3, 4, 18] for earlier works in this direction), circuits with positive quasi-probabilistic representations [19, 20, 21], stabilizer circuits supplemented with a small number of T gates [22], stabilizer circuits with small coherent local errors [23], noisy IQP circuits [24], noisy boson sampling circuits [25], low negativity magic state injection in the fault tolerant circuit model [26], quantum circuits with polynomial bounded negativity [27], Abelian-group normalizer circuits [28, 29] and certain circuits with computationally tractable states and sparse output distributions [30]. In addition, there has been some

Stephen D. Bartlett: stephen.bartlett@sydney.edu.au

work on using small quantum systems to simulate larger quantum systems [31] as well as using noisy quantum systems to simulate ideal ones [32].

An important motivation for showing efficient classical simulability or hardness thereof for a given (possibly non-universal) quantum computer is understanding what properties of a quantum computer give rise to super-classical computational power. In this context, we desire classical simulability to imply that the computational power of the target quantum computer is “contained in classical”, and the hardness of classical simulability to imply that the target computational device can achieve at least some computational task beyond classical. Achieving these desiderata hinges crucially on the strength of the notion of simulation that is employed. As an extreme example, if one uses a notion of simulation that is too weak, then efficient classical “simulation” of universal quantum circuits may be possible (even if $\text{BQP} \not\subseteq \text{BPP}$). In such a case, the existence of a “simulator” does not imply that the computational power of the simulated system is contained within classical. As an opposite extreme, if one uses a notion of simulation that is too strong, then efficient classical “simulation” of even classical circuits may be impossible [33]. In this case, the non-existence of such a simulator does not imply that the computational power of the “un-simulable” system is outside of classical. Once we establish the notion of simulation that is neither “too strong” nor “too weak”, it will become evident that both too strong and too weak notions of simulations have been commonly used in the literature. To this end, we require a clear mathematical statement about which notion of simulation minimally preserves the computational power of the system it simulates.

From a computer science perspective, the computational power of a device can be characterized by the set of problems such a device can solve. However, when it comes to quantum devices that produce probabilistic output from an exponentially growing space, even the question of what problems these devices solve or what constitutes a solution is subtle. Given an efficient description of a quantum circuit, the exact task performed by a quantum computer is to output a sample from the probability distribution associated with the measurement outcomes of that quantum circuit. This suggests that for ideal quantum computers, sampling from the *exact* quantum distribution is what constitutes a solution. On the other hand, it is unclear what well justified necessary requirement fail to be met by an arbitrarily small departure from exact sampling. Perhaps due to these subtleties, the choice of notion of “classical simulation” for sampling problems lacks consensus and, under the umbrella term of *weak simulation*, a number of different definitions have been used in the literature. We will argue that some of these notions are too strong to be minimal and others are too weak to capture computational power. The cornerstone of this argument will be the concept of *efficient indistinguishability*; the ability of one agent to remain indistinguishable from another agent under the scrutiny of any interactive test performed by a computationally powerful referee whilst simultaneously employing resources that are polynomially equivalent.

Examples of definitions that we argue are too strong include simulators required to sample from exactly the target distribution or sample from a distribution that is exponentially close (in L_1 -norm) to the target distribution [34]. These also include a notion of simulation based on approximate sampling where the accuracy requirement is the very strong condition that every outcome probability is within a small *relative* error of the target probability [35, 8, 14]. From the perspective of efficient indistinguishability, these notions of simulation are not minimal since they rule out weaker notions of simulation that are nonetheless efficiently indistinguishable from the target quantum system.

An example of a notion of approximate weak simulation that we argue is too weak requires that the classical algorithm sample from a distribution that is within some small fixed constant L_1 -norm of the target distribution [24, 25, 9, 10, 11, 36, 16]. We argue that such a notion does not capture the full computational power of the target, since it cannot perform a task that can be performed by the target device, namely of passing some sufficiently powerful distinguishability test.

The focus of this paper will be on a notion of approximate weak simulation we call *efficient polynomially small in L_1 -norm* (EPSILON) simulation (or ϵ -simulation for short). This has been used in prior works including Refs. [7, 37, 12, 22]. We will advocate for this notion of simulation (over other definitions of weak simulation) by showing that an ϵ -simulator of a quantum computer achieves efficient indistinguishability and any simulator that achieves efficient indistinguishability satisfies the definition of an ϵ -simulator. Thus ϵ -simulation minimally captures computational

power. The notion of ϵ -simulation is also closely related to the definition of a sampling problem from Ref. [37] (where the definition includes an exact statement of what constitutes a solution to the sampling problem). In this language, an ϵ -simulator of a family of quantum circuits can be exactly defined as an efficient classical algorithm which can solve all sampling problems defined by the family of quantum circuits in the natural way. Thus, our result shows that a device can solve all sampling problems defined by a quantum computer if and only if the device is efficiently indistinguishable from the quantum computer.

The conceptual significance of ϵ -simulation as a notion that minimally captures computational power motivates the study of its relation to other notions of simulation. This is particularly important for translating the existing results on simulability and hardness into statements about computational power relative to classical. Such a comparison to the above-mentioned approximate weak simulators is clear but a comparison to simulators defined in terms of Born probability estimation can be significantly more involved. Simulators which output sufficiently accurate Born probability estimates can be called as subroutines in an efficient classical procedure in order to output samples from a desired target distribution. Such a procedure can be used to “lift” these simulators to an ϵ -simulator implying that the computational power of all families of quantum circuits simulable in this way is contained within classical.

Some commonly used notions of simulation such as *strong simulation* and multiplicative precision simulation require the ability to estimate Born probabilities extremely accurately. These simulators can be lifted to ϵ -simulators [3, 4, 35]. We focus on another notions of simulation that has been prominent in recent literature [27, 26, 23] which we call a *poly-box*. Compared to strong or multiplicative precision simulators, a poly-box has a much less stringent requirement on the accuracy of Born probability estimates that it produces. We discuss the significant conceptual importance of poly-boxes owing to the fact that they capture the computational power with respect to decision problems while simultaneously being weak enough to be admitted by IQP circuits, unconditioned magic-state injected Clifford circuits and possibly other intermediate models for quantum computation.

Assuming some complexity theoretic conjectures, we show that a poly-box is a strictly weaker notion of simulation than ϵ -simulation. However, if we impose a particular sparsity restriction on the target family of quantum circuits, then we show that a poly-box can be lifted to an ϵ -simulator, implying that the two notions are, up to efficient classical computation, equivalent under this sparsity restriction.

1.1 Outline of our main results

1.1.1 Indistinguishability and ϵ -simulation.

In Sec. 2, we motivate the use of a particular notion of efficient simulation, which we call ϵ -simulation. Essentially, we say that an algorithm can ϵ -simulate a family of quantum circuits, if for any $\epsilon > 0$, it can sample from a distribution that is ϵ -close in L_1 -norm to the true output distribution of the circuit, and if the algorithm runs in time polynomial in $1/\epsilon$ and in the number of qubits. We provide an operational meaning for this notion by showing that “possessing an ϵ -simulator” for a family of circuits is equivalent to demanding that even a computationally omnipotent referee cannot distinguish the simulator’s outputs from that of the target circuit family. Further, any simulator that satisfies efficient distinguishability also satisfies the definition of an ϵ -simulator. This is captured by the following theorem presented in Sec. 2.

Theorem 1. *Bob has an ϵ -simulator of Alice’s quantum computer if and only if given the hypothesis testing scenario considered in Sec. 2.3 there exists a strategy for Bob which jointly achieves indistinguishability and efficiency.*

1.1.2 Efficient outcome estimation: the poly-box.

A family of binary outcome quantum circuits, where each circuit is indexed by a bit-string, defines a decision problem as follows: Given a bit-string indexing a quantum circuit, decide which of the

circuit’s two possible outcomes is more likely¹.

Here, the only quantity relevant to the computation is the probability associated with the binary measurement outcome (decision). Hence, in this setting, simulation can be defined in terms of the accuracy to which these probabilities can be estimated. A commonly used notion of simulation known as *strong simulation* requires the ability to estimate Born probabilities extremely accurately. In Sec. 3, we will define a much weaker notion of simulation (*poly-box*²) which is a device that computes an additive polynomial precision estimate of the quantum probability (or marginal probability) associated with a specific outcome of a quantum circuit.

We show that families of quantum circuits must admit a poly-box in order to be ϵ -simulable.

Theorem 2. *If $\mathcal{C}$ is a family of quantum circuits that does not admit a poly-box algorithm, then $\mathcal{C}$ is not ϵ -simulable.*

We advocate the importance of this notion on the grounds that whether or not some given family of quantum circuits admits a poly-box informs our knowledge of the computational power of that family relative to classical. In particular:

- if a (possibly non-universal) quantum computer can be efficiently classically simulated in the sense of a poly-box, then such a quantum computer cannot solve decision problems outside of classical
- if a (possibly non-universal) quantum computer cannot be efficiently classically simulated in the sense of a poly-box, then such a quantum computer can solve a sampling problem outside of classical (Thm. 2)

We give three examples of poly-boxes. The first one is an estimator based on Monte Carlo sampling techniques applied to a quasiprobability representation. This follows the work of Ref. [27], where it was found that the efficiency of this estimator depends on the amount of “negativity” in the quasiprobability description of the quantum circuit. As a second example, we consider the family of circuits $\mathcal{C}_{\text{PROD}}$, for which the n -qubit input state ρ is an arbitrary product state (with potentially exponential negativity), transformations consist of Clifford unitary gates, and measurements are of $k \leq n$ qubits in the computational basis. We present an explicit poly-box for $\mathcal{C}_{\text{PROD}}$ in Sec. 3. As a third example, we also outline a construction of a poly-box for Instantaneous Quantum Polynomial-time (IQP) circuits $\mathcal{C}_{\text{IQP}}$ based on the work of Ref. [39].

1.1.3 From estimation to simulation.

For the case of very high precision probability estimation algorithms, prior work has addressed the question of how to efficiently lift these to algorithms for high precision approximate weak simulators. In particular, Refs. [3, 4, 35] (see also Appendix B) lift estimation algorithms with small *relative* error. In Appendix B, we also present a potentially useful algorithm for lifting small additive error estimators. In Sec. 4 we focus on the task of lifting an algorithm for a poly-box, to an ϵ -simulator. Since a poly-box is a much less precise probability estimation algorithm (in comparison to strong simulation), achieving this task in the general case is implausible (see Sec. 5). In Sec. 4, we will show that a poly-box can efficiently be lifted to an ϵ -simulator if we restrict the family of quantum distributions to those possessing a property we call *poly-sparsity*. This sparsity property measures “peakedness versus uniformness” of distributions and is related to the scaling of the smooth max-entropy of the output distributions of quantum circuits. Loosely, a *poly-sparse* quantum circuit can have its outcome probability distribution well approximated by specifying the probabilities associated with polynomially many of the most likely outcomes. We formalize this notion in Sec. 4.

Theorem 3. *Let $\mathcal{C}$ be a family of quantum circuits with a corresponding family of probability distributions $\mathbb{P}$. Suppose there exists a poly-box over $\mathcal{C}$, and that $\mathbb{P}$ is poly-sparse. Then, there exists an ϵ -simulator of $\mathcal{C}$.*

¹Technically, one is also promised that the given bit-string will only ever index a circuit where the probability of the two outcomes are bounded away from 50%.

²This notion is similar to a notion introduced by Ref. [38] where it was (using a terminology inconsistent with the present paper) referred to as weak simulation.

We emphasize that the proof of this theorem is constructive, and allows for new simulation results for families of quantum circuits for which it was not previously known if they were efficiently simulable. As an example, our results can be straightforwardly used to show that Clifford circuits with sparse outcome distributions and with small amounts of local unitary (non-Clifford) noise, as described in Ref. [23], are ϵ -simulable.

1.1.4 Hardness results.

Finally, in Sec. 5, we prove that the poly-box requirements of Theorem 2 is on its own not sufficient for ϵ -simulability. The challenge to proving such a result is identifying a natural family of non-poly-sparse quantum circuits for which a poly-box exists but for which ϵ -simulation is impossible.

We prove that the family $\mathcal{C}_{\text{PROD}}$ described above, which violates the poly-sparsity requirement, admits a poly-box. Then, by assuming a now commonly used “average case hardness” conjecture [7, 9, 10, 12, 36, 16], we show that the ability to perform ϵ -simulation of $\mathcal{C}_{\text{PROD}}$ implies the unlikely result that the polynomial hierarchy collapses to the third level. Loosely, this result suggests that there exist quantum circuits where the probability of any individual outcome (and marginals) can be efficiently estimated, but the system cannot be ϵ -simulated. Our hardness result closely follows the structure of several similar results, and in particular that of the IQP circuits result of Ref. [9].

Our proof relies on a conjecture regarding the hardness of estimating Born rule probabilities to within a small multiplicative factor for a substantial fraction of randomly chosen circuits from $\mathcal{C}_{\text{PROD}}$. This average case hardness conjecture (which we formulate explicitly as Conjecture 1) is a strengthening of the worst case hardness of multiplicative precision estimation of probabilities associated with circuits from $\mathcal{C}_{\text{PROD}}$. Worst case hardness can be shown by applying the result of Refs. [40, 41, 42, 43, 44] and follows from an analogous argument to Theorem 5.1 of Ref. [16].

Theorem 4. *If there exists an ϵ -simulator of $\mathcal{C}_{\text{PROD}}$ and Conjecture 1 holds, then the polynomial hierarchy collapses to the third level.*

We note that our hardness result is implied by the hardness results presented in Refs. [10, 11], however; our proof is able to use a more plausible average case hardness conjecture than these references due to the fact that we are proving hardness of ϵ -simulation rather than proving the hardness of the yet weaker notion of approximate weak simulation employed by these references.

In Appendix D we also present Theorem 7. This theorem shows that the properties of poly-sparsity and anti-concentration are mutually exclusive.

The flow chart in Fig. 1 summarizes the main results in this paper by categorizing any given family of quantum circuits in terms of its computation power based on whether or not the circuit family admits certain properties related to simulability.

2 Defining simulation of a quantum computer

While there has been a breadth of recent results in the theory of simulation of quantum systems, this breadth has been accompanied with a plethora of different notions of simulation. This variety brings with it challenges for comparing results. Consider the following results, which are all based on (often slightly) different notions of simulation. As a first example, the ability to perform *strong simulation* of certain classes of quantum circuits would imply a collapse of the polynomial hierarchy, while under a weaker (but arguably more useful) notion of simulation this collapse is only implied if additional mathematical conjectures hold true [7, 9]. As another example, Ref. [14] shows that the quantum complexity class BQP is contained in the second level of the polynomial hierarchy if there exist efficient classical probabilistic algorithms for sampling a particular outcome (from the quantum circuits considered) with a probability that is exponentially close to the true quantum probability in terms of additive error (or polynomially close in terms of multiplicative error). As additional examples, Refs. [27, 23] present efficient classical algorithms for additive polynomial precision estimates of Born rule probabilities. While many such technical results are crucially sensitive to these distinctions in the meaning of simulation, there is a growing need to connect the choice of simulation definition used in a proof against (or for) efficient classical simulability to a statement about proofs of quantum advantage (or ability to practically classically solve a quantumly solvable problem). In particular, to the non-expert it can be unclear what the

Figure 1: An overview of the main results. An arbitrary family of quantum circuits $\mathcal{C}$ is partially classified by its computational power relative to universal classical computers. The unclassified category (admits a poly-box and is not poly-sparse) is known to contain circuit families that are hard to ϵ -simulate assuming some plausible complexity theoretic conjectures. We give examples of circuits families in these categories. Here, $\mathcal{C}_{\text{UNIV}}^*$, $\mathcal{C}_{\text{STAB}}^*$, $\mathcal{C}_{\text{polyN}}^*$ and $\mathcal{C}_{\text{IQP}}^*$ refer to the following families of circuits: universal circuits, stabilizer circuits, circuits with polynomially bounded negativity and IQP circuits respectively. The circuit families $\mathcal{C}_e$ and $\mathcal{C}_{\text{PROD}}$ are discussed in some detail in Sec. 4.1 and 3.4.2 respectively. The presence of a superscript represents an upper bound on the number of qubits to be measured.

complexity of classical simulation (in each of the above mentioned notions of simulation) of a given quantum device says about the hardness of building a classical device that can efficiently solve the computational problems that are solvable by the quantum device.

In this section, we will discuss a meaningful notion of approximate weak simulation, which we call ϵ -simulation. This notion of simulation is a natural mathematical relaxation of exact weak simulation and has been used in prior works, e.g., in Refs. [7, 12, 22]. Further, this notion of simulation is closely related to the class of problems in complexity theory known as sampling problems [37]. Here, we define ϵ -simulation and prove that up to polynomial equivalence, an ϵ -simulator of a quantum computer is effectively a perfect substitute for any task that can be performed by the quantum computer itself. In particular, we will show that ϵ -simulators satisfy *efficient indistinguishability* meaning that they can remain statistically indistinguishable from (according to a computationally unbounded referee) and have a polynomially equivalent run-time to the quantum computer that they simulate. We argue that efficient indistinguishability is a natural choice of a rigorously defined global condition which minimally captures the concept of computational power. The accuracy requirements of ϵ -simulation are rigorously defined at the local level of each circuit and correspond to solving a sampling problem (as defined in [37]) based on the outcome distribution of the circuit. Thus our result shows that the ability to solve all sampling problems solvable by a quantum computer $\mathcal{C}$ is a necessary and sufficient condition to being efficiently indistinguishable from $\mathcal{C}$ or “computationally as powerful as $\mathcal{C}$ ”.

2.1 Strong and weak simulation

We note that every quantum circuit has an associated probability distribution that describes the statistics of the measurement outcomes. We will refer to this as the *circuit’s quantum probability distribution*. As an example, Fig. 2 below depicts a quantum circuit. The output of running this circuit is a classical random variable $X = (X_1, \dots, X_k)$ that is distributed according to the quantum probability distribution.

Two commonly used notions of simulation are *strong simulation* and *weak simulation*. A weak simulator of a quantum circuit generates samples from the circuit’s quantum probability distribution. In the strict sense of the term, a weak simulator generates samples from the *exact* quantum probability distribution. Loosely, having a weak simulator for a quantum system is an equivalent resource to using the quantum system itself.

Figure 2: An example of a quantum circuit. This circuit acts on n qubits (or, in general, qudits). The initial state is a product state. The unitary operation $\mathcal{U}$ must be constructed out of a sequence of local unitary gates. The first k qubits in this example are each measured in a fixed basis, yielding outcome $(X_1, X_2, \dots, X_k)$. Qubits $i > k$, shown without a measurement, are traced over (marginalized).

The term *weak simulation* has also been used in reference to classical algorithms which sample from distributions which *approximate* the target probability distribution. There exist at least four distinct notions of approximate weak simulation appearing in the quantum computation literature. As background, we give a brief description of these here although the focus of this paper will be on only one of these and will be discussed in some detail later in this sections.

1. The first notion of approximate weak simulation requires that the classical algorithm sample from a distribution that is exponentially close (in L_1 -norm) to the target distribution. This notion was used in Ref. [34, 33].
2. Another notion of approximate weak simulation requires that the sampled distribution be sufficiently close to the target distribution so as to ensure that for every outcome x , the sampled distribution satisfies $|P_{\text{sampled}}(x) - P_{\text{target}}(x)| \leq \epsilon P_{\text{target}}(x)$ for some fixed $\epsilon > 0$. See Ref. [35] and also [8, 14] for related variants.
3. The third notion of approximate weak simulation requires that the classical algorithm sample from a distribution that is inverse polynomially close (in L_1 -norm) to the target distribution. This notion of simulation has been used in prior works, e.g., in Refs. [7, 37, 12, 22] both in the context of hardness of classical simulation and existence of classical simulators. We call this ϵ -simulation.
4. The final prominent example of approximate weak simulation, requires that the classical algorithm sample from a distribution that is within some small fixed constant L_1 -norm of the target distribution. This definition has predominantly featured in hardness proofs [9, 10, 11, 36, 16]. It has also feature in proofs of efficient classical simulability of noisy boson sampling circuits [25] and noisy IQP circuits [24].

A strong simulator, in contrast, outputs probabilities or marginal probabilities associated with the quantum distributions. More specifically, a strong simulator of a circuit is a device that outputs the quantum probability of observing any particular outcome or the quantum probability of an outcome marginalized³ over one or more of the measurements. Note that a strong simulator requires an input specifying the event for which the probability of occurrence is required. Taking Fig. 2 as an example, a strong simulator could be asked to return the probability of observing the event $(X_1, X_2) = (1, 0)$, marginalized over the measurements 3 to k . The requirement that a strong simulator can also output estimates of marginals is weaker than requiring them to estimate the quantum probability associated with any event (subset of the outcome space).

While the names ‘strong’ and ‘weak’ simulation suggest that they are in some sense different magnitudes of the same type of thing, we note that these two types of simulation produce different

³A distribution $P(x)$ over bit-strings $x \in \{0, 1\}^n$ is said to have a marginal distribution $P_{\{i_1, \dots, i_m\}}(\tilde{x}) = \sum_{x_{i_1}} \dots \sum_{x_{i_m}} P(x)$ (marginalized over the bits $\{i_1, \dots, i_m\}$ where $\tilde{x} \in \{0, 1\}^{n-m}$ is given by modifying the vector x by removing the entries $\{i_1, \dots, i_m\}$).

types of output. In particular, a strong simulator outputs probabilities. (More specifically, it outputs exponential additive precision estimates of Born rule probabilities and their marginals.) In contrast a weak simulator outputs samples (from the exact target probability distribution).

Ref. [33] provides a compelling argument advocating for the use of weak simulation in place of strong simulation by showing that there exist classically efficiently weak simulable probability distributions that are $\#P$ -hard to strong simulate, thus showing that aiming to classically strong simulate is an unnecessarily challenging goal. In a similar vein, here we will advocate for the notion of ϵ -simulation over other notions of simulation including the alternative notions of approximate weak simulation.

2.2 ϵ -simulation

A weak simulator, which generates samples from the exact quantum probability distribution, is a very strict notion. Often, it would be sufficient to consider a simulator that generates samples from a distribution that is only sufficiently close to the quantum distribution, for some suitable measure of closeness. Such a relaxation of the requirement of weak simulation has been used by several authors, e.g., in Refs. [34, 33, 7, 37, 12, 22, 9, 10, 36, 16, 25, 24]. Here, we define the notion of ϵ -simulation, which is a particular relaxation of the notion of weak simulation, and motivate its use.

We first define a notion of sampling from a distribution that is only *close* to a given distribution. Consider a discrete probability distribution $\mathcal{P}$. Let $B(\mathcal{P}, \epsilon)$ denote the ϵ ball around the target $\mathcal{P}$ according to the L_1 distance (or equivalently, up to an irrelevant constant, the total variation distance). We define ϵ -sampling of a probability distribution $\mathcal{P}$ as follows:

Definition 1. *Let $\mathcal{P}$ be a discrete probability distribution. We say that a classical device or algorithm can ϵ -sample $\mathcal{P}$ iff for any $\epsilon > 0$, it can sample from a probability distribution $\mathcal{P}^\epsilon \in B(\mathcal{P}, \epsilon)$. In addition, its run-time should scale at most polynomially in $1/\epsilon$.*

We note that the use of the L_1 -norm in the above is motivated by the fact that the L_1 -distance upper bounds on the one-shot success probability of distinguishing between two distributions. More details can be found in the proof of Theorem 1 in Appendix A.

The definition above does not require the device to sample from precisely the quantum probability distribution $\mathcal{P}$, but rather allows it to sample from any probability distribution $\mathcal{P}^\epsilon$ which is in the ϵ ball around the target probability distribution, $\mathcal{P}$. We note that the device or algorithm will in general take time (or other resources) that depends on the desired precision ϵ in order to output a sample, hence the efficiency requirement ensures that these resources scale at most polynomially in the precision $1/\epsilon$.

Definition 2. *We say that a classical device or algorithm can ϵ -simulate a quantum circuit if it can ϵ -sample from the circuit's associated output probability distribution $\mathcal{P}$.*

We note that each of the above mentioned notions of simulation refers to the simulation of a single quantum circuit. More generally, we may be interested in (strong, weak, or ϵ) simulators of uniform families of quantum circuits. In this setting we can discuss the efficiency of a simulator with respect to n , the number of qubits⁴. As an example, consider a family of circuits described by a mapping from $\mathcal{A}^*$ (finite strings over some finite alphabet $\mathcal{A}$) to some set of quantum circuits $\mathcal{C} = \{c_a \mid a \in \mathcal{A}^*\}$ where for each $a \in \mathcal{A}^*$, c_a is a quantum circuit with some efficient description⁵ given by the index a . In the case of strong (weak) simulation, we say that a device can efficiently strong (weak) simulate the family of quantum circuits $\mathcal{C}$ if the resources required by the device to strong (weak) simulate $c_a \in \mathcal{C}$ are upper-bounded by a polynomial in n . In the case of ϵ -simulation, we require that the simulator be able to sample a distribution within ϵ distance of the quantum distribution efficiently in both n and $1/\epsilon$.

⁴As a technical condition, we require the circuit size, run-time (or any other resources) as well as the length of the circuit's description to be upper-bounded by $\text{poly}(n)$.

⁵Such a description must satisfy the uniformity condition. This can be done by fixing a finite gate set, input state and measurement basis and explicitly defining an efficiently computable mapping between $\mathcal{A}^*$ and the gate sequence.

Definition 3. We say that a classical device or algorithm can ϵ -simulate a uniform family of quantum circuit $\mathcal{C}$ if for all $\epsilon > 0$ and for any $c \in \mathcal{C}$ (with number of qubits n and quantum distribution $\mathcal{P}$) it can sample from a probability distribution $\mathcal{P}^\epsilon \in B(\mathcal{P}, \epsilon)$ in run-time $O(\text{poly}(n, \frac{1}{\epsilon}))$.

2.3 ϵ -simulation and efficient indistinguishability

As noted earlier, this definition ensures that ϵ -simulation is a weaker form of simulation than exact weak simulation. However, we point out that the notion of exact sampling may be weakened in a number of ways, with the ϵ -simulation approach being well suited to many applications related to quantum simulators. As an example, if the definition of simulation allowed for a fixed but small amount of deviation in L_1 distance (as opposed to one that can be made arbitrarily small) then computational power of a simulator will immediately be detectably compromised. The above notion of ϵ -simulation requires a polynomial scaling between the precision ($1/\epsilon$) of the approximate sampling and the time taken to produce a sample. Below (Theorem 1), we will use a statistical indistinguishability argument to show that a polynomial scaling is precisely what should be demanded from a simulator. In particular, we will show that a run-time which scales sub-polynomially in $1/\epsilon$ puts unnecessarily strong demands on a simulator while a super-polynomial run-time would allow the simulator's output to be statistically distinguishable from the output of the device it simulates.

We now introduce the hypothesis testing scenario we consider.

Hypothesis testing scenario. Suppose Alice possesses a quantum computer capable of running a (possibly non-universal) family of quantum circuits $\mathcal{C}$, and Bob has some simulation scheme for $\mathcal{C}$ (whether it's an ϵ -simulator is to be decided). Further, suppose that a referee with unbounded computational power and with full knowledge of the specifications of $\mathcal{C}$, will request data from either Alice or Bob and run a test that aims to decide between the hypotheses:

- H_a : The requested data came from Alice's quantum computer or
- H_b : The requested data came from Bob's simulator.

The setup will be as follows: At the start of the test, one of Alice or Bob will be randomly appointed as "the candidate". Without knowing their identity, the referee will then enter into a finite length interactive protocol with the candidate (see Fig 3). Each round of the protocol will involve the referee sending a circuit description to the candidate requesting the candidate to run the circuit and return the outcome. The choice of requests by the referee may depend on all prior requests and data returned by the candidate. The rules by which the referee:

1. chooses the circuit requested in each round,
2. chooses to stop making further circuit requests and
3. decides on H_a versus H_b given the collected data

define the hypothesis test. The goal of the referee is as follows. For any given $\delta > 0$ decide H_a versus H_b such that $P_{\text{correct}} \geq \frac{1}{2} + \delta$ where P_{correct} is the probability of deciding correctly. Bob's goal is to come up with a (δ -dependent) strategy for responding to the referee's requests such that it jointly achieves:

- indistinguishability: for any $\delta > 0$ and for any test that the referee applies, $P_{\text{correct}} < \frac{1}{2} + \delta$ and
- efficiency: for every choice of circuit request sequence α , Bob must be able to execute his strategy using resources which are $O(\text{poly}(N(\alpha), \frac{1}{\delta}))$ where $N(\alpha)$ is the resource cost incurred by Alice for the same circuit request sequence.

We note that the referee can always achieve a success probability $P_{\text{correct}} = \frac{1}{2}$ simply by randomly guessing H_a or H_b . Importantly, the referee has complete control over the number of rounds in the test and additionally does not have any upper bound imposed on the number of rounds. Hence, P_{correct} is the ultimate one shot probability of the referee correctly deciding between H_a or H_b and in no sense can this probability be amplified through more rounds of information requests. As such, we will say that the referee achieves distinguishability between Alice and Bob

if $\forall \delta > 0$, there exists a test that the referee can apply ensuring that $P_{\text{correct}} \geq 1 - \delta$ (independent of Bob's strategy). Alternatively, we will say that Bob achieves indistinguishability (from Alice) if $\forall \delta > 0$, there exists a response strategy for Bob such that $P_{\text{correct}} \leq \frac{1}{2} + \delta$ (independent of what test the referee can apply). We will show that if Bob has an ϵ -simulator then there exists a strategy for Bob such that he jointly achieves indistinguishability (i.e. the referee cannot improve on a random guess by any fixed probability $\delta > 0$) and efficiency. In this case, Bob can at the outset choose any $\delta > 0$ and ensure that $P_{\text{correct}} < \frac{1}{2} + \delta$ for all strategies the referee can employ.

The efficiency requirement imposed on Bob's strategy is with respect to the resource cost incurred by Alice. Here we will define what this means and justify the rationale behind this requirement. Let us first note that for any circuit $c_a \in \mathcal{C}$, there are resource costs $R(c_a)$ incurred by Alice in order to run this circuit. This may be defined by any quantity as long as this quantity is upper and lower-bounded by some polynomial in the number of qubits. For example, $R(c_a)$ may be defined by run-time, number of qubits, number of elementary gates, number of qubits plus gates plus measurement, length of circuit description etc. Since this quantity is polynomially equivalent to the number of qubits, without loss of generality, we can treat n_a (the number of qubits used in circuit c_a) as the measure of Alice's resource cost $R(c_a)$. We now note that for a given test, the referee may request outcome data from some string of circuits $c_1, \dots, c_m \in \mathcal{C}$. Thus we define the resource cost for Alice to meet this request by $N := n_1 + \dots + n_m$.

Bob's resource cost (run-time) with respect to each circuit $c_a \in \mathcal{C}$ is polynomially dependent on both n_a and the inverse of his choice of accuracy parameter ϵ . Thus, Bob's strategy is defined by the rules by which he chooses ϵ_j , the accuracy parameter for his response in the j^{th} round⁶. Thus, for a given sequence of circuit requests $a_1, \dots, a_m \in \mathcal{A}^*$, Bob will incur a resource cost $T = t_1 + \dots + t_m$ where $t_j \sim \text{poly}(n_{a_j}, 1/\epsilon_j)$ is Bob's resource in the j^{th} round. Thus the efficiency condition requires that there exists some polynomial $f(x, y)$ such that for all $\delta > 0$ and for all possible request sequences $\alpha = (a_1, \dots, a_m)$, $T(\alpha) \leq f(N(\alpha), \frac{1}{\delta})$. The efficiency requirement imposed on Bob's strategy thus naturally requires that the resource costs of Alice and Bob be polynomial equivalent for the family of tests that the referee can apply.

Theorem 1. *Bob has an ϵ -simulator of Alice's quantum computer if and only if given the hypothesis testing scenario considered above, there exists a strategy for Bob which jointly achieves indistinguishability and efficiency.*

The proof for this theorem can be found in Appendix A. The proof uses the fact that the L_1 distance between Alice and Bob's output distributions over the entire interactive protocol can be used to upper bound the probability of correctly deciding between H_a and H_b . Further, we show that the total L_1 distance between Alice and Bob's output distributions over the entire interactive protocol grows at most additively in the L_1 distance of each round of the protocol. We also note that an ϵ -simulator allows Bob to ensure that the L_1 distance of each round decays like an inverse quadratic ensuring that the sum of the L_1 distances converges to the desired upper bound. The convergence of the inverse quadratic series, which is an inverse polynomial, thus motivates the significance of ϵ -simulators i.e. simulators with run-time $O(\text{poly}(n, 1/\epsilon))$.

We note that the "if" component of the theorem says that meeting the definition of ϵ -simulator is necessary for achieving efficient indistinguishability, thus the notion of simulation cannot be weakened any further without compromising efficient indistinguishability.

Throughout this paper we view a quantum computer as a uniform family of quantum circuits $\mathcal{C} = \{c_a \mid a \in \mathcal{A}^*\}$. We note that by committing to the circuit model of quantum computation, our language including important definitions such as ϵ -simulation are not necessarily well suited to other models of computation unless these are first translated to the circuit model. For example, in a quantum computational model that makes use of intermediate measurements, such as the measurement based quantum computing (MBQC) model, consider a procedure where a part of the state is measured then conditioned on the outcome, a second measurement is conducted. This procedure (consisting of 2 rounds of measurement) can be described as a single circuit in the circuit model, but cannot be broken up into two rounds involving two separate circuits. This limitation

⁶Bob must possess some computational power in order to execute these rules. We will only require that Bob have some small amount of memory (to keep count of the rounds in the protocol) and compute simple arithmetic functions of this counter.

becomes apparent when we consider the hypothesis testing scenario. If the referee is performing a multi-round query, expecting the candidate to possess an MBQC-based quantum computer, then even Alice with a quantum computer may be unable to pass the test unless her computer operates in an architecture that can maintain quantum coherence between rounds. In the setting we consider, such a query by the referee is not allowed.

2.4 ϵ -simulation and computational power

In addition to the technical contribution of Theorem 1, we wish to make an argument for the conceptual connection between computational power and efficient indistinguishability. Intuitively, we wish to say that an agent A is at least as computationally powerful as agent B if A can “do” every task that B can do using an equivalent amount of resources. In our setting, we can restrict ourselves to polynomially equivalent resources and the most general task of sampling from a target probability distribution given an efficient description of it. However, defining what constitutes an acceptable solution to the sampling task is not only of central importance but also difficult to conceptually motivate. Given a description of a probability distribution, can anything short of sampling *exactly* from the specified distribution constitute success? An answer in the negative seems unsatisfactory because very small deviations⁷ from exact sampling are ruled out. However, an answer in the positive presents the subtlety of specifying the exact requirement for achieving the task. It is easy to offer mathematically reasonable requirements for what constitutes success at the local level of each task but significantly more difficult to conceptually justify these as precisely the right notion. In our view, this difficulty arises because a well formed conceptually motivated requirement at the local level of each task must be inherited from a global requirement imposed at the level of the agent across their performance on any possible task.

We advocate for efficient indistinguishability as the right choice of global requirement for defining computational power and implicitly defining what constitutes a solution to a sampling task. If an agent is efficiently indistinguishable from another then, for any choice of $\delta > 0$ chosen at the outset, the referee cannot assign any computational task to the candidate to observe a consequence that will improve (over randomly guessing) their ability to correctly decide between H_a and H_b by a probability δ . Thus, there is no observable consequence⁸ to substituting an agent with another efficiently indistinguishable agent. For these reasons, we argue that in the setting where the agents are being used as computational resources, an agent’s ability to (efficiently and indistinguishably) substitute another naturally defines containment of computational power. In light of this, the “only if” component of Theorem 1 says that, the computational power of Bob (given an ϵ -simulator of $\mathcal{C}$) contains that of Alice (given $\mathcal{C}$) and the “if” component says that an ϵ -simulator is the *minimal* simulator that achieves this since any simulator to achieve efficient indistinguishability is an ϵ -simulator.

The referee can be seen as a mathematical tool for bounding the adversarial ability of any natural process to distinguish an agent from an efficiently indistinguishable substitute. As such one may argue for further generalization of the concept of efficient indistinguishability from one which is defined with respect to (w.r.t.) a computationally unbounded referee to a notion dependent on the computational power of the referee. If we take the view that the computational power of all agents within this universe is bounded by universal quantum computation, then a particularly interest generalization is efficient indistinguishability w.r.t. a referee limited to universal quantum computation. We return to this generalization in the discussion, elsewhere focusing on efficient indistinguishability w.r.t. a computationally unbounded referee.

⁷For example consider the scenario that whenever an agent is asked to sample from some distribution $\mathcal{P}$, they output samples from exactly $\mathcal{P}$ every time, possibly with one exception. In particular, a memory bit stores if the exception has ever taken place. If it has occurred, then forever more, when the agent is asked to sample from $\mathcal{P}$, an exact sample is produced. If the exception has not yet taken place then with some very small probability, the agent will output the most likely outcome instead of an exact sample from $\mathcal{P}$.

⁸Since the observer is the computationally unbounded referee, then any event is an observable consequence i.e. if we let S be the set of all possible responses across all rounds from both Alice and Bob, then any element of the power set of S is an observable consequence.

3 Probability Estimation

As described in the previous section, an exact (or approximate in the sense of Ref. [34]) weak simulator produces outcomes sampled from the exact (or exponentially close to the exact) Born rule probability distribution associated with a quantum circuit. The notion of ϵ -simulation is a weaker notion of simulation, a fact we aim to exploit by constructing algorithms for ϵ -simulation that would not satisfy the above-mentioned stronger notions of simulation. In this paper, we describe an approach to ϵ -simulation of quantum circuits based on two components: first, estimating Born rule probabilities for specific outcomes of a quantum circuit to a specified precision, and then using such estimates to construct a simulator. In this section, we describe this first component, coined a *poly-box*. We motivate and define poly-boxes, discuss their conceptual importance and give a number of important examples. In the next section, we employ such an estimator to construct an ϵ -simulator under certain conditions.

3.1 Born rule probabilities and estimators

Consider the description $c = \{\rho, U, \mathcal{M}\}$ of some ideal quantum circuit, with ρ an initial state, $U = U_L U_{L-1} \cdots U_1$ a sequence of unitary gates, and $\mathcal{M}$ a set of measurement operators (e.g., projectors). Associated with the measurement $\mathcal{M} = \{E_x \mid x \in \{0, 1\}^k\}$ is a set of possible measurement outcomes $x \in \{0, 1\}^k$. The Born rule gives us the exact quantum predictions associated with observing any particular outcome x :

$$\mathcal{P}(x) := \text{tr}(U\rho U^\dagger E_x). \quad (1)$$

Further, probabilities associated with events $S \subseteq \{0, 1\}^k$ are given by:

$$\mathcal{P}(S) := \sum_{x \in S} \text{tr}(U\rho U^\dagger E_x) \quad (2)$$

The task of efficiently classically estimating these probabilities with respect to general quantum circuits is of great practical interest, but is known to be hard even for rather inaccurate levels of estimation. For example, given a circuit c_a from a family of universal quantum circuits with a Pauli Z measurement of the first qubit only, deciding if $\mathcal{P}_a(0) > \frac{2}{3}$ or $< \frac{1}{3}$ is BQP-complete.

Monte Carlo methods are a common approach to estimating Born rule probabilities that are difficult to calculate directly [27, 26, 23]. Let p be an unknown parameter we wish to estimate, e.g., a Born rule probability. In a Monte Carlo approach, p is estimated by observing a number of random variables $X_1, \dots, X_s$ and computing some function of the outcomes $\hat{p}_s(X_1, \dots, X_s)$, chosen so that $\hat{p}_s$ is close to p in expectation. In this case, $\hat{p}_s$ is an estimator of p .

We first fix some terminology regarding the precision of an estimator, and how this precision scales with resources. We say that an estimator $\hat{p}_s$ of p is *additive* (ϵ, δ) -precision if:

$$\Pr(|p - \hat{p}_s| \geq \epsilon) \leq \delta, \quad \text{additive } (\epsilon, \delta)\text{-precision.} \quad (3)$$

We say that $\hat{p}_s$ is *multiplicative* (ϵ, δ) -precision if:

$$\Pr(|p - \hat{p}_s| \geq \epsilon p) \leq \delta, \quad \text{multiplicative } (\epsilon, \delta)\text{-precision.} \quad (4)$$

In the case where $p \leq 1$ is a probability, a multiplicative precision estimator is more accurate than an additive precision estimator.

For any estimator based on the Monte Carlo type of approach described above, there is a polynomial (typically linear) resource cost associated with the number of samples s . For example, the time taken to compute $\hat{p}_s$ will scale polynomially in s . More generally, s may represent some resource invested in computing the estimator $\hat{p}_s$ such as the computation run-time. For this reason, we may wish to classify additive/multiplicative (ϵ, δ) -precision estimators by how s scales with $1/\epsilon$ and $1/\delta$. We say that $\hat{p}_s$ is an additive *polynomial precision estimator* of p if there exists a polynomial $f(x, y)$ such that for all $\epsilon, \delta > 0$, $\hat{p}_s$ is an additive (ϵ, δ) -precision estimator for all $s \geq f(\epsilon^{-1}, \log \delta^{-1})$. We say that $\hat{p}_s$ is a multiplicative *polynomial precision estimator* of p if there

exists a polynomial $f(x, y)$ such that for all $\epsilon, \delta > 0$, $\hat{p}_s$ is a multiplicative (ϵ, δ) -precision estimator for all $s \geq f(\epsilon^{-1}, \delta^{-1})$ ⁹.

A useful class of polynomial additive precision estimators is given by application of the Hoeffding inequality. Suppose $\hat{p}_1$ resides in some interval $[a, b]$ and is an unbiased estimator of p (i.e. $\mathbb{E}(\hat{p}_1) = p$). Let $\hat{p}_s$ be defined as the average of s independent observations of $\hat{p}_1$. Then, by the Hoeffding inequality, we have:

$$\Pr(|p - \hat{p}_s| \geq \epsilon) \leq 2 \exp\left(\frac{-2s\epsilon^2}{(b-a)^2}\right), \quad (5)$$

for all $\epsilon > 0$. We note that for $s(\epsilon^{-1}, \log \delta^{-1}) \geq \frac{(b-a)^2}{2\epsilon^2} \log(2\delta^{-1})$, $\hat{p}_s$ is an additive (ϵ, δ) -precision estimator of p . With this observation, we see that additive polynomial precision estimators can always be constructed from unbiased estimators residing in a bounded interval.

As an important example let us consider one way an agent can generate Born probability estimates when given access to some classical processing power and a family of quantum circuits $\mathcal{C}$. Given a description of an event S and a description of a quantum circuit $c_a \in \mathcal{C}$, the agent can efficiently estimate $p = \mathcal{P}_a(S)$. In this example, the agent can construct the estimator $\hat{p}_s$ by independently running the circuit s times. On each of the runs $i = 1, \dots, s$, she observes if the outcome x is in the event S (in this case, $X_i = 1$) or not in S (in this case, $X_i = 0$). We then define $\hat{p}_s = \frac{1}{s} \sum_{i=1}^s X_i$. Using the Hoeffding inequality, it is easy to show that the Born rule probability estimator $\hat{p}_s$ is an additive polynomial precision estimator of p . Thus, for all $a \in \mathcal{A}^*$, $\epsilon, \delta > 0$, there is a choice of $s \in \mathbb{N}$ such that this procedure can be used to compute an estimate $\hat{p}$ of $p := \mathcal{P}_a(S)$ such that $\hat{p}$ satisfies the accuracy requirement:

$$\Pr(|p - \hat{p}| \geq \epsilon) \leq \delta \quad (6)$$

and the run-time required to compute the estimate $\hat{p}$ is $O(\text{poly}(n, \epsilon^{-1}, \log \delta^{-1}))$.

Let us now discuss an important aspect that we have been ignoring: namely the restrictions that need to be placed on the events S . We first note that since each event S is an element of the power set of $\{0, 1\}^k$, the total number of events grows doubly exponentially implying that any polynomial length description of events can only index a tiny fraction of the set of all events. Even once we make a particular choice as to how (and hence which) events are indexed by polynomial length descriptions, deciding if a bit-string x is in the event S is not computationally trivial (with the complexity depending on the indexing). Since the estimation procedure requires a computational step where the agent checks whether x is in S , there will be restrictions place on the allowed events depending on the computational limitations of the agent and the complexity of the indexing of events.

When discussing poly-boxes, we will be interested in the restricted set of events $S \in \{0, 1, \bullet\}^k$. We use this notation to indicate the set of all specific outcomes and marginals. Specifically, $S \in \{0, 1, \bullet\}^k$ is a subset of $\{0, 1\}^k$ where $\bullet$ is a “wild card” single qubit measurement outcome and hence is consistent with both a 0 and a 1 element. For example, $S = (0, \bullet, 1) := \{(0, 0, 1), (0, 1, 1)\}$. If S is a vector with bits $x_1, \dots, x_{k-m}$ in positions $i_1, \dots, i_{k-m}$ and a $\bullet$ in each of the positions $j_1, \dots, j_m$; then S represents the set of 2^m outcomes where the qubits numbered $i_1, \dots, i_{k-m}$ produced single qubit measurement outcomes corresponding to $x_1, \dots, x_{k-m}$ while the remaining qubits (those numbered $j_1, \dots, j_m$) produce either a 0 or a 1 measurement outcome. The probability corresponding to such an event S is the marginal probability associated with observing the outcome bit-string $x_1, \dots, x_{k-m}$ on the qubits numbered $i_1, \dots, i_{k-m}$ marginalized over the qubits $j_1, \dots, j_m$.

⁹The observant reader will notice that additive and multiplicative precision estimators have different scalings in δ . Of course one can define an alternative notion of additive estimation where $s \geq f(\epsilon^{-1}, \delta^{-1})$ or an alternative notion of multiplicative estimation where $s \geq f(\epsilon^{-1}, \log \delta^{-1})$. Here, we have chosen to define the notions that are most useful as motivated by the existence of techniques and associated inequalities bounding their performance. In particular, Hoeffding’s inequality allows the construction of additive (ϵ, δ) -precision estimators while Chebyshev’s inequality motivates the multiplicative (ϵ, δ) -precision estimator definition.

3.2 The poly-box: generating an additive polynomial precision estimate

Given a family of quantum circuits $\mathcal{C}$, we will be interested in constructing an ϵ -simulator of $\mathcal{C}$ using estimates of Born rule probabilities associated with circuits in $\mathcal{C}$. For this purpose we define a poly-box over $\mathcal{C}$.

Definition 4. (poly-box). A poly-box over a family of quantum circuits $\mathcal{C} = \{c_a \mid a \in \mathcal{A}^*\}$ with associated family of probability distributions $\mathbb{P} = \{\mathcal{P}_a \mid a \in \mathcal{A}^*\}$ is a classical algorithm that, for all $a \in \mathcal{A}^*$, $\epsilon, \delta > 0$ and $S \in \{0, 1, \bullet\}^{k_n}$, can be used to compute an estimate $\hat{p}$ of $\mathcal{P}_a(S)$ such that $\hat{p}$ satisfies the accuracy requirement:

$$\Pr(|p - \hat{p}| \geq \epsilon) \leq \delta \quad (7)$$

and, the run-time required to compute the estimate $\hat{p}$ is $O(\text{poly}(n, \epsilon^{-1}, \log \delta^{-1}))$.

Eq. (7), gives an upper bound on the probability that the computed estimate, $\hat{p}$, is far from the target quantity. This probability is over the potential randomness in the process used to generate the estimate $\hat{p}$. In addition we implicitly assume that the output of this process is independent of prior output. In particular, let $\alpha = (a, \epsilon, \delta, S)$ be an input into a poly-box and $\hat{p}_\alpha$ the observed output. Then, we implicitly assume that the probability distribution of $\hat{p}_\alpha$ only depends on the choice of input α and in particular is independent of prior output.

Note that a poly-box over a family of quantum circuits $\mathcal{C} = \{c_a \mid a \in \mathcal{A}^*\}$ with associated family of probability distributions $\mathbb{P} = \{\mathcal{P}_a \mid a \in \mathcal{A}^*\}$ is a classical algorithm that can be used to compute additive polynomial precision estimators $\hat{p}_s$ of $\mathcal{P}_a(S)$ for all $a \in \mathcal{A}^*$, $s \in \mathbb{N}$, $S \in \{0, 1, \bullet\}^{k_n}$ efficiently in s and n .

3.3 Conceptual significance of a poly-box

Whether or not a family of quantum circuits $\mathcal{C}$ admits a poly-box has bearing on both the complexity of sampling problems and decision problems solvable by $\mathcal{C}$, and so we will find that the notion of a poly-box is a useful concept. We first note that the existence of a poly-box is a necessary condition for ϵ -simulation.

Theorem 2. If $\mathcal{C}$ is a family of quantum circuits that does not admit a poly-box algorithm, then $\mathcal{C}$ is not ϵ -simulable.

Proof. We note that given an ϵ -simulator of $\mathcal{C}$, a poly-box over $\mathcal{C}$ can be constructed in the obvious way simply by observing the frequency with which the ϵ -simulator outputs outcomes in S and using this observed frequency as the estimator for $\mathcal{P}(S)$. $\square$

A poly-box over $\mathcal{C}$ is not only necessary for the existence of an ϵ -simulator over $\mathcal{C}$, but as we will show in Theorem 3, combined with an additional requirement, it is also sufficient. In addition, we note that if $\mathcal{C}$ admits a poly-box then all “generalized decision problems” solvable by $\mathcal{C}$ are solvable within BPP. As an illustrative but unlikely example, suppose there exists a classical poly-box over a universal quantum circuit family $\mathcal{C}_{\text{UNIV}}$. Then, for any instance x of a decision problem L in BQP, there is a quantum circuit $c_a \in \mathcal{C}_{\text{UNIV}}$ that decides if $x \in L$ (correctly on at least $2/3$ of the runs), simply by outputting the decision “ $x \in L$ ” when the first qubit measurement outcome is 1 on a single run of c_a and conversely, outputting the decision “ $x \notin L$ ” when the first qubit measurement outcome is 0. We note that, in order to decide if $x \in L$ one does not need the full power of an ϵ -simulator over $\mathcal{C}_{\text{UNIV}}$. In fact it is sufficient to only have access to the poly-box over $\mathcal{C}_{\text{UNIV}}$. Given a poly-box over $\mathcal{C}_{\text{UNIV}}$, one can request an (ϵ, δ) -precision estimate $\hat{p}$ for the probability p that the sampled outcome from c_a is in $S = (1, \bullet, \dots, \bullet)$. For $\epsilon < 1/6$ and $\delta < 1/3$, one may decide “ $x \in L$ ” if $\hat{p} \geq 1/2$ and “ $x \notin L$ ” otherwise. This will result in the correct decision with probability $\geq 2/3$ as required. A poly-box over $\mathcal{C}$ offers the freedom to choose any $S \in \{0, 1, \bullet\}^n$ which can in general be used to define a broader class of decision problems. Of course in the case of $\mathcal{C}_{\text{UNIV}}$, this freedom cannot be exploited because for every choice of a and $S \neq (1, \bullet, \dots, \bullet)$, there is a alternative easily computable choice of a' such that the probability that a run of $c_{a'} \in \mathcal{C}_{\text{UNIV}}$ results in an outcome in $(1, \bullet, \dots, \bullet)$ is identical to the probability that a run of $c_a \in \mathcal{C}_{\text{UNIV}}$ results in an outcome in S . However, since we are considering the general case of not necessarily universal families of quantum

circuits, it is feasible that a poly-box over $\mathcal{C}$ will be computationally more powerful than a poly-box over $\mathcal{C}$ restricted to only estimating probabilities of events of the form $S = (1, \bullet, \dots, \bullet)$. On the other hand, we do not wish to make poly-boxes exceedingly powerful. If we view a poly-box over $\mathcal{C}$ as a black box containing an agent with access to $\mathcal{C}$ and processing an estimation algorithm as per the aforementioned example, then by restricting the allowable events as above and choosing such a simple method of indexing these, we are able to limit the additional computational power given to the agent and/or poly-box.

3.4 Examples of poly-boxes

3.4.1 Poly-boxes from quasiprobability representations

There are a number of known algorithms for constructing poly-boxes over certain non-universal families of quantum circuits [2, 45, 27, 22, 23]. In particular, we focus on the algorithm presented in Ref. [27], which can be used to construct a poly-box over any family of quantum circuits $\mathcal{C}$ where the *negativity* of quantum circuits grows at most polynomially in the circuit size. We refer the interested reader to Ref. [27] for a definition of the negativity of a quantum circuit, but note that this quantity depends on the initial state, sequence of unitaries and the final POVM measurement that defines the circuit. For general quantum circuits, the negativity can grow exponentially in both the number of qudits and the depth of the circuit.

A key application of this approach is to Clifford circuits. In odd dimensions, stabilizer states, Clifford gates, and measurements in the computational basis do not contribute to the negativity¹⁰ of a Monte Carlo based estimator. Including product state preparations or measurements that are not stabilizer states, or non-Clifford gates such as the T gate, may contribute to the negativity of the circuit. Nonetheless, these non-Clifford operations can be accommodated within the poly-box provided that the total negativity is bounded polynomially. In addition, a poly-box exists for such circuits even in the case where the negativity of the initial state, *or* of the measurement, is exponential [27].

3.4.2 A poly-box over $\mathcal{C}_{\text{PROD}}$

As a nontrivial example of a class of Clifford circuits for which there exists a poly-box, consider the family of circuits $\mathcal{C}_{\text{PROD}}$. This family consists of quantum circuits with an n -qubit input state ρ that is an arbitrary product state¹¹ (with potentially exponential Wigner function negativity [27] in the input state). The allowed transformations are non-adaptive Clifford unitary gates, and $k \leq n$ qubits are measured at the end of the circuit, in the computational basis. Such a circuit family has been considered by Jozsa and Van den Nest [41], where it was referred to as INPROD, OUTMANY, NON-ADAPT. This circuit family will be discussed again in Sec. 5 where we will show the classical hardness of simulating this family according to another notion of simulation. Aaronson and Gottesman [2] provide the essential details of a poly-box for this family of circuits; for completeness, we present an explicit poly-box for $\mathcal{C}_{\text{PROD}}$ in the following lemma.

Lemma 1. *A classical poly-box exists for the Clifford circuit family $\mathcal{C}_{\text{PROD}}$.*

Proof. Give an arbitrary circuit $c = \{\rho, U, \mathcal{M}\} \in \mathcal{C}_{\text{PROD}}$ and an event $S \in \{0, 1, \bullet\}^n$ we construct an estimator $\hat{p}_s$ of the probability $\mathcal{P}(S)$ as follows:

1. Let $\Pi = \otimes_{i=1}^n \Pi_i$ be the projector corresponding to S . Here, we set:

$$\Pi_i = \begin{cases} \frac{I+Z}{2} & \text{if the } i^{\text{th}} \text{ entry of } S \text{ is } 0 \\ \frac{I-Z}{2} & \text{if the } i^{\text{th}} \text{ entry of } S \text{ is } 1 \\ I & \text{if the } i^{\text{th}} \text{ entry of } S \text{ is } \bullet \end{cases} \quad (8)$$

¹⁰with respect to either the phase point operator or stabilizer states choice of frame

¹¹As an additional technical requirement, we impose that the input product state is generated from $|0\rangle^{\otimes n}$ by the application of polynomially many gates from a universal single qubit gate set with algebraic entries.

2. For each i where the i^{th} entry of S is not $\bullet$, $\Pi_i = \frac{I \pm Z}{2}$. In these cases, define a local Pauli operator P_i by sampling either I or $\pm Z$ with equal probability. For each i where the i^{th} entry of S is a $\bullet$, we deterministically set $P_i = I$.
3. We construct the n -qubit Pauli operator $P := \otimes_{i=1}^n P_i$, (including its sign $\pm$).
4. Using the Gottesman-Knill theorem [2], we compute the Pauli operator $P' = \otimes_{i=1}^n P'_i := U^\dagger P U$.
5. We compute the single sample estimate $\hat{p}_1$ using the equation:

$$\hat{p}_1 := \text{tr}(\rho P') = \prod_{i=1}^n \text{tr}(\rho_i P'_i). \quad (9)$$

6. We compute the estimator $\hat{p}_s$ by computing s independent single sample estimates and taking their average.

It is straightforward to show that the expectation value of $\hat{p}_s$ is the target quantum probability $p := \mathcal{P}(S)$. Further, the single sample estimates are bounded in the interval $[-1, 1]$. Hence, by the Hoeffding inequality,

$$\Pr(|\hat{p}_s - p| \geq \epsilon) \leq 2e^{-\frac{s\epsilon^2}{2}}. \quad (10)$$

This algorithm can be executed efficiently in s and in n and produces additive polynomial precision estimates of $\mathcal{P}(S)$ for any circuit $c \in \mathcal{C}_{\text{PROD}}$ and any $S \in \{0, 1, \bullet\}^n$ and is thus a poly-box. $\square$

3.4.3 A poly-box over $\mathcal{C}_{\text{IQP}}$

As an additional example, we note that $\mathcal{C}_{\text{IQP}}$, the family of Instantaneous Quantum Polynomial-time (IQP) quantum circuits [46, 47, 8, 9] that consist of computational basis preparation and measurements with all gates diagonal in the X basis admits a poly-box. One can construct such a poly-box over $\mathcal{C}_{\text{IQP}}$ by noting that Proposition 5 from Ref. [39] gives a closed form expression for all Born rule probabilities and marginals of these circuits. This expression:

$$\mathcal{P}_P(S) = \mathbb{E}_{r \in \text{span}\{\vec{e}_i \mid i \in \{i_1, \dots, i_k\}\}} \left[(-1)^{r \cdot s} \alpha \left(P_r, \frac{\pi}{4} \right) \right] \quad (11)$$

is an expectation value over 2^k vectors in $\mathbb{Z}_2^n$ where:

- $\{i_1, \dots, i_k\}$ are the set of indices where the entries of S are in $\{0, 1\}$;
- $s \in \mathbb{Z}_2^n$ is defined by $s_i = S_i$ when $i \in \{i_1, \dots, i_k\}$ and $s_i = 0$ otherwise;
- P_r is the *affinification* of the $m \times n$ binary matrix P which defines a Hamiltonian of the IQP circuit constructed from Pauli X operators according to $H_P := \sum_{i=1}^m \otimes_{j=1}^n X^{P_{ij}}$;
- $\alpha(P, \theta)$ is the normalized version of the weight enumerator polynomial (evaluated at $e^{-2i\theta}$) of the code generated by the columns of P .

We note that this is an expectation over exponentially many terms which have their real part bounded in the interval $[-1, 1]$. Further, for each r , the quantity $\alpha \left(P_r, \frac{\pi}{4} \right)$ can be evaluated efficiently using Vertigan's algorithm [48] and Ref. [39]. As such, one can construct an additive polynomial precision estimator for all Born rule probabilities and marginals simply by evaluating the expression:

$$\hat{p}_1 = \text{Re} \left[(-1)^{r \cdot s} \alpha \left(P_r, \frac{\pi}{4} \right) \right] \quad (12)$$

for polynomially many independent uniformly randomly chosen $r \in \text{span}\{\vec{e}_i \mid i \in \{i_1, \dots, i_k\}\}$ and computing the average over all choices. This can be shown to produce a poly-box by application of the Hoeffding inequality.

4 From estimation to simulation

Given the significance of ϵ -simulation as the notion that minimally preserves computational power, here we turn our attention to the construction of an efficient algorithms for lifting a poly-box to an ϵ -simulators. We give strong evidence that in the general case, such a construction is not possible. This suggests that a poly-box is statistically distinguishable from an ϵ -simulator and hence computationally less powerful. However, by restricting to a special family of quantum circuits, we show an explicit algorithm for lifting a poly-box to an ϵ -simulator. Combined with Theorem 2 this shows that within this restricted family a poly-box is computationally equivalent to an ϵ -simulator.

The significance of ϵ -simulation also motivates the need to understand the relationship to other simulators defined in terms of Born probability estimation. At the end of this section and in Appendices B and C we present two algorithms which lift an estimator of probabilities and marginals to a sampler.

4.1 A poly-box is not sufficient for ϵ -simulation

This section focuses on the relation between poly-boxes and ϵ -simulation. With a poly-box, one can efficiently estimate Born rule probabilities of outcomes of a quantum circuit with additive precision. However, assuming $\text{BQP} \neq \text{BPP}$, a poly-box alone is not a sufficient computational resource for ϵ -simulation. We illustrate this using a simple but somewhat contrived example, wherein an encoding into a large number of qubits is used to obscure (from the poly-box) the computational power of sampling.

Define a family of quantum circuits $\mathcal{C}_e$ using a universal quantum computer as an oracle as follows:

1. take as input a quantum circuit description $a \in \mathcal{A}^*$ (this is a description of some quantum circuit with n qubits);
2. call the oracle to output a sample outcome from this quantum circuit. Label the first bit of the outcome by X ;
3. sample an n -bit string $Y \in \{0, 1\}^n$ uniformly at random;
4. output $Z = (X \oplus \text{Par}(Y), Y) \in \{0, 1\}^{n+1}$, where $\text{Par}(Y)$ is the parity function on the input bit-string Y .

We note that $\mathcal{C}_e$ cannot admit an ϵ -simulator unless $\text{BQP} \subseteq \text{BPP}$, since simple classical post processing reduces the ϵ -simulator over $\mathcal{C}_e$ to an ϵ -simulator over universal quantum circuits restricted to a single qubit measurement.

We now show that $\mathcal{C}_e$ admits a poly-box:

1. take as input $a \in \mathcal{A}^*$, $\epsilon, \delta > 0$ and $S \in \{0, 1, \bullet\}^{n+1}$. Our poly-box will output probability estimates that are deterministically within ϵ of the target probabilities and hence we can set $\delta = 0$;
2. if S specifies a marginal probability i.e. $k < n + 1$, then the poly-box outputs the estimate 2^{-k} (where k is the number of non-marginalized bits in S); otherwise,
 - (a) small ϵ case: if $\epsilon < 1/2^n$, explicitly compute the quantum probability $p := \Pr(X = 1)$;
 - (b) large ϵ case: if $\epsilon \geq 1/2^n$, output the probability $2^{-(n+1)}$ as a guess.

This algorithm is not only a poly-box over $\mathcal{C}_e$ but it in fact outputs probability estimates that have exponentially small precision.

Lemma 2. *For all $a \in \mathcal{A}^*$, $\epsilon > 0$ and $S \in \{0, 1, \bullet\}^n$, the above poly-box can output estimates within ϵ additive error of the target probability using $O(\text{poly}(n, 1/\epsilon))$ resources. Further, the absolute difference between estimate and target probabilities will be $\leq \min \{2^{-(n+1)}, \epsilon\}$.*

Proof. We note that the resource cost of this algorithm is $O(\text{poly}(n, 1/\epsilon))$. Since in the case of small ϵ it is $O(\text{poly}(2^n)) \subseteq O(\text{poly}(1/\epsilon))$ and in the case of large ϵ it is $O(n)$.

We now consider the machine's precision by considering the case with no marginalization and the case with marginalization separately. We restrict the below discussion to the large ϵ case as the estimates are exact in the alternate case.

Let $z = (z_0, \dots, z_n) \in \{0, 1\}^{n+1}$ be fixed and define $z' := (z_1, \dots, z_n)$. Then,

$$\Pr(Z = z) = \Pr(Z_0 = z_0 \mid Y = z') \Pr(Y = z') = \Pr(X = z_0 \oplus \text{Par}(z')) 2^{-n}. \quad (13)$$

So for $S = z$ (i.e. no marginalization), we have an error given by $\max_{r \in \{p, 1-p\}} |2^{-(n+1)} - \frac{r}{2^n}| \leq 2^{-(n+1)}$.

For the case where $S_i = \bullet$ (i.e. there is marginalization over the i^{th} bit only and $k = n$), we note that the quantum marginal probability $p(S)$ is given exactly by:

$$p(S) = \sum_{z_i=0}^1 \Pr(Z = z) = \sum_{z_i=0}^1 \Pr(X = z_0 \oplus \text{Par}(z')) 2^{-n} = p 2^{-n} + (1-p) 2^{-n} = 2^{-k}, \quad (14)$$

where $z_j := S_j$ for $j \neq i$. This implies that for all $k < n+1$, the quantum probability is exactly 2^{-k} . Thus, in the worst case (no marginalization and $\epsilon \geq 2^{-n}$), the error is $\leq 2^{-(n+1)}$. $\square$

This example clearly demonstrates that the existence of a poly-box for a class of quantum circuits is not sufficient for ϵ -simulation. In the following, we highlight the role of the *sparsity* of the output distribution in providing, together with a poly-box, a sufficient condition for ϵ -simulation.

4.2 Sparsity and sampling

Despite the fact that in general the existence of a poly-box for some family $\mathcal{C}$ does not imply the existence of an ϵ -simulator for $\mathcal{C}$, for some quantum circuit families, a poly-box suffices. Here, we show that one can construct an ϵ -simulator for a family of quantum circuits $\mathcal{C}$ provided that there exists a poly-box over $\mathcal{C}$ and that the family of probability distributions corresponding to $\mathcal{C}$ satisfy an additional constraint on the *sparsity* of possible outcomes. We begin by reviewing several results from Schwarz and Van den Nest [30] regarding sparse distributions. In Ref. [30], they define the following property of discrete probability distributions:

Definition 5. (*ϵ -approximately t -sparse*). A discrete probability distribution is t -sparse if at most t outcomes have a non-zero probability of occurring. A discrete probability distribution is ϵ -approximately t -sparse if it has a L_1 distance less than or equal to ϵ from some t -sparse distribution.

The lemma below is a (slightly weakened) restatement of Theorem 11 from Ref. [30].

Lemma 3. (Theorem 11 of Ref. [30]). Let $\mathcal{P}$ be a distribution on $\{0, 1\}^k$ that satisfies the following conditions:

1. $\mathcal{P}$ is promised to be ϵ -approximately t -sparse, where $\epsilon \leq 1/6$;
2. For all $S \in \{0, 1, \bullet\}^k$, there exists an (s, k) -efficient randomized classical algorithm for sampling from $\hat{p}_s$, an additive polynomial estimator of $\mathcal{P}(S)$.

Then it is possible to classically sample from a probability distribution $\mathcal{P}' \in B(\mathcal{P}, 12\epsilon + \delta)$ efficiently in k , t , ϵ^{-1} and $\log \delta^{-1}$.

We note that for every discrete probability distribution $\mathcal{P}$, there is some unique minimal function $t(\epsilon)$ such that for all $\epsilon \geq 0$, $\mathcal{P}$ is ϵ -approximately t -sparse. We note that if this function is upper-bounded by a polynomial in ϵ^{-1} , then a randomized classical algorithm for sampling from estimators of $\mathcal{P}(S)$ can be extended to a randomized classical algorithm for sampling from some probability distribution $\mathcal{P}' \in B(\mathcal{P}, \epsilon)$ efficiently in ϵ^{-1} . This fact motivates the following definition:

Definition 6. (*poly-sparse*) Let $\mathcal{P}$ be a discrete probability distribution. We say that $\mathcal{P}$ is poly-sparse if there exists a polynomial $P(x)$ such that for all $\epsilon > 0$, $\mathcal{P}$ is ϵ -approximately t -sparse whenever $t \geq P(\frac{1}{\epsilon})$.

Let $\mathbb{P}$ be a family of probability distributions with $\mathcal{P}_a \in \mathbb{P}$ a distribution over $\{0, 1\}^{k_a}$. We say that $\mathbb{P}$ is poly-sparse if there exists a polynomial $P(x)$ such that for all $\epsilon > 0$ and $a \in \mathcal{A}^*$, $\mathcal{P}_a$ is ϵ -approximately t -sparse whenever $t \geq P(k_a/\epsilon)$.

The notion of poly-sparse is related to the notion of smooth max entropy $H_{\max}^\epsilon$. In particular, $\mathbb{P}$ is poly-sparse iff there exists a polynomial $P(x)$ such that for every $\mathcal{P} \in \mathbb{P}$ with domain cardinality 2^n , we have:

$$2^{H_{\max}^\epsilon(\mathcal{P})} \leq P\left(\frac{n}{\epsilon}\right) \quad (15)$$

where $H_{\max}^\epsilon(\mathcal{P}) := \inf_{\mathcal{P}'} \log_2 |\text{Supp}(\mathcal{P}')|$, $|\text{Supp}(\mathcal{P}')|$ is the cardinality of the support of the distribution $\mathcal{P}'$ and the infimum is taken over all distributions $\mathcal{P}'$ subject to $\frac{1}{2} \|\mathcal{P}' - \mathcal{P}\|_1 \leq \epsilon$. This notion was first defined in Ref. [49] where it corresponds to the ϵ -smooth Rényi entropy of order $\alpha = 0$.

4.3 Conditions for ϵ -simulation

With this notion of output distributions that are poly-sparse, we are in a position to state our main theorem of this section:

Theorem 3. Let $\mathcal{C}$ be a family of quantum circuits with a corresponding family of probability distributions $\mathbb{P}$. Suppose there exists a poly-box over $\mathcal{C}$, and that $\mathbb{P}$ is poly-sparse. Then, there exists an ϵ -simulator of $\mathcal{C}$.

Proof. Let $a \in \mathcal{A}^*$ and $\epsilon > 0$ be arbitrary. Then there exist $t = t(a, \epsilon)$ such that $\mathcal{P}_a$ is ϵ -approximately t -sparse. Further, due to the existence of the efficient classical poly-box over $\mathcal{C}$, for all $S \in \{0, 1, \bullet\}^{k_a}$, there exists an (s, k_a) -efficient randomized classical algorithm for sampling from an additive polynomial estimator of $\mathcal{P}_a(S)$. Thus by Lemma 3, it is possible to classically sample from a probability distribution $\mathcal{P}_a^\epsilon \in B(\mathcal{P}_a, \epsilon)$ efficiently in ϵ^{-1}, t and k_a . We note that here we have removed the dependence on δ since we can make $\delta \leq \epsilon$ whilst remaining efficient in ϵ^{-1}, t and k_a . Finally, since poly-sparsity guarantees the existence of a $t(a, \epsilon)$ that can be upper-bounded by a polynomial in $\frac{k_a}{\epsilon}$, we arrive at the desired result. $\square$

As an example, consider families of quantum circuits $\mathcal{C}$ where each circuit of size n can only produce outcomes from some set of size at most $\text{poly}(n)$. Then $\mathcal{C}$ is poly-sparse (even if the output distributions are uniform over the $\text{poly}(n)$ sized support). Hence, if $\mathcal{C}$ also admits a poly-box, then by Thm. 3 one can with high probability repeatedly sample from this space of $\text{poly}(n)$ outcomes hidden within an exponentially large space of bit-strings.

We have shown that having a poly-box and a poly-sparsity guarantee for a family of quantum circuits gives us an ϵ -simulator. We emphasize that the proof of this Theorem is constructive, and allows for new simulation results for families of quantum circuits for which it was not previously known if they were efficiently simulable. As an example, our results can be straightforwardly used to show that Clifford circuits with sparse outcome distributions and with small amounts of local unitary (non-Clifford) noise, as described in Ref. [23], are ϵ -simulable.

Theorem 3 requires a promise of poly-sparsity. Since this is a property of infinite families of probability distributions, one cannot hope to algorithmically verify (or even falsify) it through sampling from member distributions. Nevertheless, for distributions generated by some particular family of quantum circuits, a proof that this property holds may be possible.

In summary, the results of Thms. 3 and 2 imply that in order to construct an ϵ -simulator of any particular family of quantum circuits, it is necessary to construct a poly-box and further, if the family is poly-sparse, this is also sufficient. In Sec. 4.1, we also showed that there exists a somewhat artificial family of quantum circuits $\mathcal{C}_e$ with respect to which a poly-box is insufficient for ϵ -simulation. In the next section, we show that this phenomenon also occurs with much more natural families of quantum circuits.

4.4 On lifting stronger estimators to approximate samplers

In contrast to poly-boxes, certain stronger notions of simulation based on Born rule probability estimation can be lifted to ϵ -simulators (or even stronger approximate weak simulators). In Appendices B and C we present two such efficient classical algorithms.

The algorithm presented in Appendix C uses an estimator with multiplicative precision to construct an ϵ -simulator (it can in fact construct an approximate weak simulator based on the stronger notion from Ref. [35]). This algorithm exploits the fact that ratios of multiplicative precision estimators are multiplicative precision in order to sequentially, one qubit’s measurement outcome at a time, sample from the marginal probability of the next qubit’s measurement conditioned on the sampled measurement outcomes of the prior measurements. This algorithm and its variants have been presented in Refs. [3, 4, 35] and are well known within the simulation-of-quantum-circuits community.

The algorithm presented in Appendix B uses an estimator with exponentially small additive precision to construct an ϵ -simulator (it can in fact construct an approximate weak simulator based on the stronger notion from Ref. [34]). This algorithm aims to map a bit-string r (approximately representing a uniformly sampled point from the unit interval) to a bit-string representing the outcome of running the circuit. Such a mapping is defined for every ordering of the measurement outcomes. This algorithm makes intuitive use of marginal probability estimates to do a binary search for the measurement outcome corresponding to r . This technique avoids computing ratios of probability estimates making it useful in regimes where additive errors are small but larger than some of the probabilities in the target distribution. Hence, this algorithm has some advantages compared to that of Appendix C. In particular, it can be used to lift an additive ε precision estimator to a sampler from within L_1 distance $O(2^n \varepsilon)$. This can be used to construct a ϵ -simulator in certain cases where the algorithm in Appendix C would fail. An example is when one has access to an estimator with additive precision $\varepsilon = 2^{-n} \kappa$ where $\kappa > 0$ can be made arbitrarily small in run-time $O(\text{poly}(n, 1/\kappa))$.

5 Hardness results

In the previous section, we have shown that one can construct an ϵ -simulator for a family of quantum circuits $\mathcal{C}$ given a poly-box for this family together with a promise of poly-sparsity of the corresponding probability distribution. We also discussed a contrived construction of a family of quantum circuits that admits a poly-box but is not ϵ -simulable (unless $\text{BQP}=\text{BPP}$). In this section, we provide strong evidence (dependent only on standard complexity assumptions and a variant of the now somewhat commonly used [7, 9, 10, 11, 12, 36, 16] “average case hardness” conjecture) that a condition such as poly-sparsity is necessary even for natural families of quantum circuits. One such family has already been identified by noting that $\mathcal{C}_{\text{IQP}}$ admits a poly-box and is likely hard to ϵ -simulate [9]. Here, we also show the likely hardness of ϵ -simulating the non-poly-sparse Clifford circuit family $\mathcal{C}_{\text{PROD}}$ (defined in Sec. 3). These results mean that at least two (and possibly more) of the intermediate models of quantum computing have the property that the probability of individual outcomes and marginals can be estimated to $1/\text{poly}(n)$ additive error but due to non-sparsity, their ϵ -simulability is implausible.

Our hardness result for classical ϵ -simulation of $\mathcal{C}_{\text{PROD}}$ closely follows the structure of several similar results, and in particular that of the IQP circuits result of Ref. [9]. We note that this hardness result is implied by the hardness results presented in Refs. [10, 11], however; our proof is able to use a more plausible average case hardness conjecture than these references due to the fact that we are proving hardness of ϵ -simulation rather than proving the hardness of the yet weaker notion of approximate weak simulation employed by these references.

Despite the existence of a poly-box over $\mathcal{C}_{\text{PROD}}$, we show that there cannot exist a classical ϵ -simulator of this family unless the average case hardness conjecture fails or the polynomial hierarchy collapses to the third level. We note that the hardness of *exact* weak simulation of $\mathcal{C}_{\text{PROD}}$ was shown in Ref. [41]. In contrast here we show the hardness of ϵ -simulation for this family. Our proof relies on a conjecture regarding the hardness of estimating Born rule probabilities to within a small multiplicative factor for a substantial fraction of randomly chosen circuits from $\mathcal{C}_{\text{PROD}}$.

This average case hardness conjecture is a strengthening of the worst case hardness of multiplicative precision estimation of probabilities associated with circuits from $\mathcal{C}_{\text{PROD}}$.

The hardness of ϵ -simulating $\mathcal{C}_{\text{PROD}}$ circuits is shown by first noting that the existence of a classical ϵ -simulator implies, via the application of the Stockmeyer approximate counting algorithm [50], the existence of an algorithm (in the third level of the PH) for estimating the probabilities associated with the output distribution of the ϵ -simulator to within a multiplicative factor. These estimates can then be related to estimates of the exact quantum probabilities by noting two points:

1. that the deviation between the ϵ -simulator's probability of outputting a particular outcome and that of the exact quantum probability will be exponentially small for the vast majority of outcomes. We show this fact using Markov's inequality.
2. that a significant portion of outcomes associated with randomly chosen circuit in $\mathcal{C}_{\text{PROD}}$ must have outcome probabilities larger than a constant fraction of 2^{-n} . We show this property using our proof that these circuits anti-concentrate.

These observations are combined to show that if there exists an ϵ -simulator of $\mathcal{C}_{\text{PROD}}$, then there exists a classical algorithm (in the third level of the PH) that can estimate Born rule outcome probabilities to within a multiplicative factor for almost 50% of circuits sampled from $\mathcal{C}_{\text{PROD}}$. This is in contradiction with Conjecture 1 thus implying that an ϵ -simulator does not exist.

5.1 Conjecture regarding average case hardness

We begin by stating our conjecture that multiplicative precision estimation of $\mathcal{C}_{\text{PROD}}$ is $\#P$ -hard in the average case.

Conjecture 1. *There exist an input product state ρ over n qubits such that given a uniformly random Clifford unitary U acting on n qubits, estimating $p := \text{tr}(U\rho U^\dagger|0\rangle\langle 0|)$ to within a multiplicative error of $1/\text{poly}(n)$ for 49% or more of the sampled Clifford unitaries is $\#P$ -hard.*

We note that this average case hardness conjecture has an analogous worst case hardness version¹². The worst case hardness can be proven by applying the result of Refs. [40, 41, 42, 43, 44] and by an argument essentially identical to the proof of Theorem 5.1 in Ref. [16]. We omit the proof here but note that this proof relies on three key facts:

1. that estimating Born rule probabilities for universal (indeed even IQP) circuits that use a gate set with algebraic entries, to within any multiplicative factor in the open interval $(1, \sqrt{2})$ is $\#P$ -hard [44, 42];
2. for gate sets with algebraic entries, all non-zero output probabilities are lower bounded by some inverse exponential [43];
3. that $\mathcal{C}_{\text{PROD}}$ circuits with post-selection (or adaptivity) are universal for quantum computation [40, 41].

We emphasize that similar conjectures are commonly used in related hardness proofs, such as Refs. [7, 9, 10, 12, 36, 16].

5.2 Anti-concentration of outcomes for $\mathcal{C}_{\text{PROD}}$

Next, we prove that Clifford circuits chosen uniformly at random from the family $\mathcal{C}_{\text{PROD}}$ satisfy an anti-concentration property.

Lemma 4. *Let d be a prime. For each $n \in \mathbb{N}$, let c_n be an n -qudit Clifford circuit chosen by fixing an arbitrary n qudit input state ρ , applying a uniformly random Clifford unitary U acting on n qudits and doing a computational basis measurement on all qudits. Then for all $\alpha \in (0, 1)$ and for any fixed choice of $x \in \{0, \dots, d-1\}^n$:*

$$\Pr_U \left(p_x \geq \frac{\alpha}{d^n} \right) > \frac{(1-\alpha)^2}{2}, \quad (16)$$

¹²This is the same statement as per Conjecture 1 but with “49% or more” replaced by “100%”.

where $p_x := \text{tr}(U\rho U^\dagger |x\rangle\langle x|)$ is the Born rule probability for the outcome x .

Proof. We use the unitary 2-design property of the Clifford group.

$$\mathbb{E}(p_x) = \text{tr}(\mathbb{E}(U\rho U^\dagger)|x\rangle\langle x|) = \text{tr}(\mathbb{1}/d^n |x\rangle\langle x|) = \frac{1}{d^n} \quad (17)$$

$$\begin{aligned} \mathbb{E}(p_x^2) &= \text{tr}(\mathbb{E}(U \otimes U(\rho \otimes \rho)U^\dagger \otimes U^\dagger)|x\rangle\langle x| \otimes |x\rangle\langle x|) \\ &= \frac{\text{tr}(P_{\text{Sym}}(\rho \otimes \rho))}{\text{tr} P_{\text{Sym}}} \text{tr}(P_{\text{Sym}}|x\rangle\langle x| \otimes |x\rangle\langle x|) \\ &= \frac{2\text{tr}(P_{\text{Sym}}(\rho \otimes \rho))}{d^n(d^n + 1)} \\ &= \frac{(\text{tr}(\rho^2) + (\text{tr}\rho)^2)}{d^n(d^n + 1)} \\ &\leq \frac{2}{d^n(d^n + 1)}, \end{aligned} \quad (18)$$

where $P_{\text{Sym}} = \frac{1}{2}(\mathbb{1} + \text{SWAP})$ is the projection onto the symmetric subspace of $\mathbb{C}^{d^n} \otimes \mathbb{C}^{d^n}$. We use the Paley-Zygmund inequality, which states that for a non-negative random variable R with finite variance, and for any $\alpha \in (0, 1)$:

$$\Pr(R \geq \alpha \mathbb{E}[R]) \geq (1 - \alpha)^2 \frac{\mathbb{E}^2[R]}{\mathbb{E}[R^2]}, \quad (\text{Paley-Zygmund inequality}) \quad (19)$$

Application of this inequality with Eqs. (17-18) then gives the desired result. $\square$

We point out that the property of anti-concentration is inconsistent with poly-sparsity. This result is shown in Theorem 7 of Appendix D.

5.3 Hardness theorem

We are now in a position to prove our main theorem:

Theorem 4. *If there exists an ϵ -simulator of $\mathcal{C}_{\text{PROD}}$ and Conjecture 1 holds, then the polynomial hierarchy collapses to the third level.*

Proof. Assuming there exists an ϵ -simulator of $\mathcal{C}_{\text{PROD}}$, we can treat the ϵ -simulator as a deterministic Turing machine with a random input. Let $\mathcal{T}$ be the Turing machine that takes as an input $\epsilon > 0$ (representing the L_1 error required), $r \in \{0, 1\}^{\text{poly}(n/\epsilon)}$ (representing the random bit-string) and $d_c \in \mathcal{A}^{\text{poly}(n)}$ (representing an efficient description of an n qubit circuit $c \in \mathcal{C}_{\text{PROD}}$) and outputs an outcome $X^\epsilon \in \{0, 1\}^k$ with the correct statistics (over uniformly random r inputs) up to ϵ in L_1 distance in time $\text{poly}(n, 1/\epsilon)$. That is, the output satisfies:

$$\|p - p^\epsilon\|_1 := \sum_{x \in \{0, 1\}^k} |p_x - p_x^\epsilon| \leq \epsilon \quad (20)$$

where $p_x := \Pr(X = x)$ is the probability of observing outcome x on a single run of the quantum circuit c and $p_x^\epsilon := \Pr_{r \sim \text{unif}}(X^\epsilon = x)$ is the probability of observing outcome x on a single run of the Turing machine $\mathcal{T}$ for a uniformly distributed random r and fixed ϵ, d_c inputs.

We now note that the problem of computing the proportion p_x^ϵ of bit-strings r that result in $\mathcal{T}(\epsilon, r, d_c) = x$ is a problem in $\#P$. Thus, the Stockmeyer algorithm gives us a means of estimating p_x^ϵ to within a multiplicative error in the complexity class FBPP^{NP} .

More precisely, there exists an algorithm in FBPP^{NP} which will output an estimate $\tilde{p}_x^\epsilon$ such that:

$$|p_x^\epsilon - \tilde{p}_x^\epsilon| \leq \frac{p_x^\epsilon}{\text{poly}(n)} \quad (21)$$

Thus we have that for all c and for all x :

$$\begin{aligned}
|p_x - \tilde{p}_x^\epsilon| &\leq |p_x - p_x^\epsilon| + |p_x^\epsilon - \tilde{p}_x^\epsilon| \\
&\leq |p_x - p_x^\epsilon| + \frac{p_x^\epsilon}{\text{poly}(n)} \\
&\leq |p_x - p_x^\epsilon| + \frac{p_x + |p_x - p_x^\epsilon|}{\text{poly}(n)} \\
&= |p_x - p_x^\epsilon| \left(1 + \frac{1}{\text{poly}(n)}\right) + \frac{p_x}{\text{poly}(n)}
\end{aligned} \tag{22}$$

We note that the expectation value of $|p_x - p_x^\epsilon|$ over random choice of $x \sim \text{unif}(\{0, 1\}^k)$ is upper-bounded by $2^{-n}\epsilon$. That is:

$$\mathbb{E}_x[|p_x - p_x^\epsilon|] = \frac{1}{2^k} \sum_x |p_x - p_x^\epsilon| = \frac{1}{2^k} \|p - p^\epsilon\|_1 \leq \frac{\epsilon}{2^k}$$

Restricting our attention to circuits in $\mathcal{C}_{\text{PROD}}$ where all of the qubits are measured i.e. $k = n$, we have:

$$\mathbb{E}_x[|p_x - p_x^\epsilon|] \leq \frac{\epsilon}{2^n} \tag{23}$$

We apply Markov's inequality, which states that for R a non-negative random variable and $\gamma > 0$:

$$\Pr\left(R \geq \frac{\mathbb{E}[R]}{\gamma}\right) \leq \gamma, \quad (\text{Markov's inequality}) \tag{24}$$

we have that for all $\beta > 0$:

$$\Pr_x\left(|p_x - p_x^\epsilon| \geq \frac{\mathbb{E}_x[|p_x - p_x^\epsilon|]}{\beta}\right) \leq \beta \tag{25}$$

That is:

$$\Pr_x\left(|p_x - p_x^\epsilon| < \frac{\epsilon}{\beta 2^n}\right) > (1 - \beta) \tag{26}$$

Applying this to the upper bound in Eq. (22), we find that for all $\beta > 0$:

$$\Pr_x\left(|p_x - \tilde{p}_x^\epsilon| < \frac{\epsilon}{\beta 2^n} \left(1 + \frac{1}{\text{poly}(n)}\right) + \frac{p_x}{\text{poly}(n)}\right) > (1 - \beta) \tag{27}$$

For any fixed choices of $\alpha \in (0, 1)$, $\beta, \epsilon > 0$, let us define the following events:

- Event A: $\frac{p_x}{\alpha} \geq \frac{1}{2^n}$
- Event B: $|p_x - \tilde{p}_x^\epsilon| < \frac{\epsilon}{\beta 2^n} \left(1 + \frac{1}{\text{poly}(n)}\right) + \frac{p_x}{\text{poly}(n)}$.

By Eq. (16), we have $\Pr_U(A) > \frac{(1-\alpha)^2}{2}$ and by Eq. (27), we have $\Pr_x(B) > (1 - \beta)$. Recall that the intersection bound tells us that $\Pr(A \cap B) \geq \max\{0, \Pr(A) + \Pr(B) - 1\}$ for events A and B . Thus, we have $\Pr(A \cap B) \geq \frac{(1-\alpha)^2 - 2\beta}{2}$. This immediately implies the following:

$$\Pr_{U,x}\left(|p_x - \tilde{p}_x^\epsilon| < \frac{\epsilon p_x}{\alpha \beta} \left(1 + \frac{1}{\text{poly}(n)}\right) + \frac{p_x}{\text{poly}(n)}\right) > \frac{(1 - \alpha)^2 - 2\beta}{2} \tag{28}$$

This can be further simplified by incorporating the randomness over x into the uniform randomness over the Clifford unitaries. Specifically, let $y \in \{0, 1\}^n$ be arbitrarily fixed. Further, let $U_x := \otimes_{i=1}^n X^{x_i}$. Then, noting that for all n qubit Cliffords V :

$$\Pr_{U, U_x} (U_x U = V) = \Pr_{U, U_x} (U = U_x V) \quad (29)$$

$$= \Pr_U (U = V) \quad (30)$$

$$= \Pr_U (U_y U = V) \quad (31)$$

where probabilities over U_x are chosen uniformly over all $x \in \{0, 1\}^n$. Applying this to Eq. (28) we find that for all $y \in \{0, 1\}^n$ and for all n qubit product states ρ :

$$\Pr_U \left(|p_y - \tilde{p}_y^\epsilon| < \frac{\epsilon p_y}{\alpha \beta} \left(1 + \frac{1}{\text{poly}(n)} \right) + \frac{p_y}{\text{poly}(n)} \right) > \frac{(1 - \alpha)^2 - 2\beta}{2} \quad (32)$$

We recall that for an ϵ -simulator, $\epsilon > 0$ can be made polynomially small efficiently in run-time and n . Thus, as an example, we may assign the following scaling to α, β, ϵ :

$$\alpha = \frac{1}{n}, \quad \beta = \frac{1}{2n^2}, \quad \epsilon = \frac{\alpha \beta}{n}. \quad (33)$$

This argument shows that the existence of an ϵ -simulator of $\mathcal{C}_{\text{PROD}}$ implies that there exists an algorithm in Δ_3^P that can for any fixed product states ρ and measurement outcomes $x \in \{0, 1\}^n$, output an $O(1/n)$ multiplicative precision estimate of $p_x := \text{tr}(U \rho U^\dagger |x\rangle \langle x|)$ for almost 50% of randomly uniformly chosen Clifford unitaries U acting on n qubits. That is:

$$\Pr_U \left[|p_x - \tilde{p}_x^\epsilon| < p_x O(1/n) \right] > \frac{1}{2} - \frac{1}{n} \quad (34)$$

By conjecture 1, this is $\#P$ -hard. This implies that a $\#P$ -hard problem is solved in FBPP^{NP} . By Toda's theorem [51], this collapses the polynomial hierarchy to its third level. $\square$

6 Discussion

There is a substantial and growing body of results showing the classical “simulability” of some quantum computers and the hardness of “simulability” of others. We hope that the results presented here will significantly inform the interpretation of this literature in relation to the comparison of the computational power of the relevant quantum computer to the computational power of a universal classical computer. For some family of quantum circuits $\mathcal{C}$, these results typically make statements of the form either:

- Simulability: $\mathcal{C}$ can be classically “simulated” or
- Hardness: $\mathcal{C}$ can be classically “simulated” implies some implausible outcome

In the case of simulability proofs, our results show that whenever the notion of simulation used is stronger or equivalent to ϵ -simulation, the useful computational power of $\mathcal{C}$ is contained within classical. Further, if the notion of simulation is a poly-box (a weaker notion than ϵ -simulation), this still applies provided that $\mathcal{C}$ is poly-sparse. If $\mathcal{C}$ is not known to be poly-sparse but admits a poly-box then, we can still conclude that without non-trivial classical post-processing, $\mathcal{C}$ is incapable of solving decision problems outside of the complexity class BPP.

In the case of hardness proofs, our results show that whenever the notion of simulation used is weaker or equivalent to ϵ -simulation, it is plausible that the useful computational power of $\mathcal{C}$ is beyond classical. However, for proofs of hardness based on weaker notions of simulation, it may be possible to alter the proof such that it shows the hardness of ϵ -simulation (rather than a yet weaker notion) with the added benefit that now the hardness is more plausible.

Some hardness results show the implausibility of classically simulating $\mathcal{C}$ with respect to a notion of simulation much stronger than ϵ -simulation. Even if quantum computers can reliably

achieve such a notion of simulation, these results cannot be seen as showing the implausibility of the existence of efficient classical devices that can be used as a perfectly good computational substitute to $\mathcal{C}$.

The perspective of efficient indistinguishability gives us a natural avenue to defining the set of all problems solvable by a quantum device. We have seen that the minimal notion of simulation to achieve this is ϵ -simulation; a significantly weaker notion of simulation than many of the notions used in literature [34, 35, 8, 14]. Thus, the gap between classical and quantum computational power can be closed not only by the development of more powerful classical simulation algorithms but also by significantly reducing the computational hurdle classical devices must overcome in order to act as efficient substitutes to quantum computers. Our results exploit this feature in order to show that any family of quantum circuits that both admits a poly-box and satisfies the poly-sparsity condition can be ϵ -simulated. The existence of multiple known constructions of poly-boxes (see Refs. [2, 45, 27, 22, 23]) over restricted families of quantum circuits, and in particular Ref. [27], demonstrates the significant advantages offered by weakening the minimal requirements on classical simulators from the stronger notions of weak simulation to that of ϵ -simulation.

For any given family of quantum circuits, poly-sparsity can be trivially guaranteed by upper bounding the number of measured qubits by $\log n$. However, the condition of poly-sparsity permits significantly more complex probability distribution families (including families with exponentially growing support). Future exploration of how to non-trivially guarantee poly-sparsity offers yet more potential for identifying interesting families of quantum circuits that are ϵ -simulable using the techniques outlined here.

In this paper, we have argued that ϵ -simulation minimally captures computational power. However, the term “minimally” is with respect to the computational power of the referee, which is unbounded in the setting we considered. This raises the importance of future work aimed at defining the notion of simulation which minimally captures the computational power of a quantum computer with respect to a referee that is computationally bounded to universal quantum computation (or equivalent). In light of this observation, our work suggests that even requiring a simulator to be capable of solving all sampling problems (as defined in Ref. [37]) solvable by the quantum device is too strong to be minimal (w.r.t. a universal quantum bounded referee). Future results in this direction would inform us on precisely how to further weaken the notion of sampling problems and to define a yet weaker complexity class than SampBQP (or more generally $\text{Samp}\mathcal{C}$) that (w.r.t. a universal quantum bounded referee) minimally captures computational power.

In an experimental setting where there is a constant lower bound to the noise present in the quantum device, the minimal requirements for efficient indistinguishability become yet weaker. In this setting, it is plausible that for IQP circuits and boson sampling circuits, classical computation can achieve the minimal requirements for efficient indistinguishability w.r.t. a universal quantum bounded referee. This possibility is supported by the existence of classical algorithms for simulating noisy IQP circuits [24] and noisy boson sampling circuits [25]. In the constant lower bounded noise setting, these algorithms fail to achieve efficient indistinguishability w.r.t. a computationally unbounded referee. However, whether or not they achieve efficient indistinguishability w.r.t. a universal quantum bounded referee remains a question to be resolved.

Aiming to tighten the separation between simulability and hardness is an important goal toward a deeper understanding of the computational power of quantum versus classical circuits. Specifically, the aim is to move towards a full classification of simulability by gradually reducing the “unclassified” space (of parameters describing a quantum computer that are both outside the range to ensure simulability and outside the range to ensure hardness of simulability). By focusing on the tension between anti-concentration and poly-sparsity our work has made modest progress in this direction with potential for further consolidation and progress with respect to this aim.

We have shown the poly-sparsity and anti-concentration properties to be mutually exclusive. If we assume that the polynomial hierarchy does not collapse and restrict to quantum computers that admit a poly-box and average case hardness (some plausible candidates being $\mathcal{C}_{\text{IQP}}$, $\mathcal{C}_{\text{PROD}}$, and their poly-sparse restricted counterparts) we see that either poly-sparsity holds ensuring ϵ -simulability or anti-concentration holds ensuring hardness.

For general quantum circuit families, poly-sparsity and anti-concentration are not exhaustive. Future work directed towards finding interesting spaces of quantum computers where the two notions are exhaustive would help to classify more of the yet unclassified computers in Fig. 1

(admits a poly-box and not poly-sparse). Restricted to this setting, all quantum computers that admit the appropriate average case hardness property would admit a hardness proof. Further, such work can give a much needed new perspective on the peculiar nature of the transition from ϵ -simulable to hardness that IQP and magic state injected Clifford circuit families undergo as they transition from poly-sparse to non-poly-sparse. In particular, this may shed light on whether this behavior (shared by IQP, magic state injected Clifford circuit and possibly others) is common to intermediate models of quantum computing for a good reason or simply a coincidence.

Our work establishes the conceptual importance of a poly-box as a notion of simulation. Through the Hoeffding inequality and powerful sampling techniques such as Monte Carlo simulations, we inherit a number of important examples of poly-boxes including IQP circuits, magic state injected Clifford circuit and circuits with polynomially bounded negativity (see also Refs. [2, 45, 27, 22, 23]). This is of immediate practical interest as admitting a poly-box is sufficient for many useful problems such as finding certain expectation values or estimating the probability associated with certain events.

Whether or not a family of quantum circuits $\mathcal{C}$ admits a poly-box significantly informs our understanding of the computational power of $\mathcal{C}$ relative to classical. Simulability of a family of quantum circuits $\mathcal{C}$ according to the notion of a poly-box, implies that, without additional classical computational resources, $\mathcal{C}$ cannot solve decision problems outside of classical. If $\mathcal{C}$ is also poly-sparse (with binary outcome circuits being a very special case) then even an agent with universal classical computational power and access to the quantum computer $\mathcal{C}$ is confined to universal classical computational power. However, when supplemented with a universal classical computer, if $\mathcal{C}$ admits a poly-box but is not poly-sparse then it may be capable of solving decision problems beyond BPP. This possibility is not ruled out by our analysis and is consistent with the fact that $\mathcal{C}_{\text{PROD}}$ and $\mathcal{C}_{\text{IQP}}$ circuits both admit hardness proofs.

There is something conceptually unclear about circuit families that admit poly-boxes and a hardness proof of the type presented in Sec. 5. In particular, it is unclear if these admit a poly-box purely due to the restriction placed on the types of events that a poly-box can be queried about, or if hardness of ϵ -simulation could manifest even in circuit families which allow efficient classical polynomial precision estimation of probabilities associated with any family of events decidable in BPP. In the latter case, an agent with access to such circuits cannot solve any decision problem outside of BPP even given access to a universal classical computer. The former case leaves open the possibility that these families of circuits will behave like $\mathcal{C}_e$ (introduced in Sec. 4.1) where some appropriate classical post-processing of outcome samples will render them more powerful than BPP (assuming $\text{BQP} \not\subseteq \text{BPP}$). This question is closely related to an open question raised by Aaronson in Ref. [37].

It is surprising that examples of families of circuits that admit a poly-box and a plausible hardness proof are far from rare and in fact may be typical among intermediate models of quantum computing. In addition to the families we have shown to be in this category ($\mathcal{C}_{\text{PROD}}$ and $\mathcal{C}_{\text{IQP}}$), we note that linear optical networks $\mathcal{C}_{\text{LON}}$ and circuits with polynomially bounded negativity $\mathcal{C}_{\text{poly}\mathcal{N}}$ are also plausible candidates. We note that due to an algorithm by Gurvits [6] (see also Ref. [52]), the family of linear optical quantum circuits considered in the boson sampling setting of Ref. [7] admit additive polynomial precision estimators of individual outcome probabilities. However, there is no known poly-box over this family since it is currently unclear how to produce such estimators for *all* marginal probabilities. Alternatively, $\mathcal{C}_{\text{poly}\mathcal{N}}$ is known to admit a poly-box [27]. Also, for odd prime d it contains the qudit generalization of $\mathcal{C}_{\text{PROD}}$ which is both universal under post-selection and anti-concentrates. Hence an average case hardness conjecture is also plausible implying that $\mathcal{C}_{\text{poly}\mathcal{N}}$ admits a proof of hardness essentially identical to that of $\mathcal{C}_{\text{PROD}}$. In light of these considerations we are optimistic that useful and computationally interesting applications can be found for intermediate models of quantum computation.

Acknowledgments

The authors are grateful to Marco Tomamichel, Ryan Mann, Michael Bremner, Dax Koh, Daniel Brod and Joel Wallman for helpful discussions. This research is supported by the ARC via the Centre of Excellence in Engineered Quantum Systems (EQUS) project number CE170100009 and

by the U.S. Army Research Office through grant W911NF-14-1-0103. HP also acknowledges support from the Australian Institute for Nanoscale Science and Technology Postgraduate Scholarship (John Makepeace Bennett Gift). The work of DG is supported by the Excellence Initiative of the German Federal and State Governments (ZUK 81), the DFG within the CRC 183 (project B01), and the DAAD.

References

- [1] R. P. Feynman, “Simulating Physics with Computers,” *Int. J. Theor. Phys.*, **21**, 467–488 (1982).
- [2] S. Aaronson and D. Gottesman, “Improved simulation of stabilizer circuits,” *Phys. Rev. A*, **70**, 052328 (2004).
- [3] L. G. Valiant, “Quantum circuits that can be simulated classically in polynomial time,” *SIAM Journal on Computing*, **31**, 1229–1254 (2002).
- [4] B. M. Terhal and D. P. DiVincenzo, “Classical simulation of noninteracting-fermion quantum circuits,” *Phys. Rev. A* **65**, 032325 (2002).
- [5] S. D. Bartlett, B. C. Sanders, S. L. Braunstein, and K. Nemoto, “Efficient classical simulation of continuous variable quantum information processes,” *Phys. Rev. Lett.* **88**, 097904 (2002).
- [6] L. Gurvits, “On the complexity of mixed discriminants and related problems,” in: Jędrzejowicz J., Szepietowski A. (eds) *Mathematical Foundations of Computer Science 2005. MFCS 2005. Lecture Notes in Computer Science*, vol 3618. Springer, Berlin, Heidelberg (2005).
- [7] S. Aaronson and A. Arkhipov, “The computational complexity of linear optics,” in *Proceedings of the forty-third annual ACM symposium on Theory of computing*, 333–342, ACM (2011).
- [8] M. J. Bremner, R. Jozsa, and D. J. Shepherd, “Classical simulation of commuting quantum computations implies collapse of the polynomial hierarchy,” *Proc. R. Soc. A* **467**, (2010).
- [9] M. J. Bremner, A. Montanaro, and D. J. Shepherd, “Average-case complexity versus approximate simulation of commuting quantum computations,” *Phys. Rev. Lett.* **117**, 080501 (2016).
- [10] X. Gao, S.-T. Wang, and L.-M. Duan, “Quantum supremacy for simulating a translation-invariant Ising spin model,” *Phys. Rev. Lett.* **118**, 040502 (2017).
- [11] J. Bermejo-Vega, D. Hangleiter, M. Schwarz, R. Raussendorf, and J. Eisert, “Architectures for quantum simulation showing a quantum speedup,” *Phys. Rev. X* **8**, 021010 (2018).
- [12] B. Fefferman and C. Umans, “The power of quantum fourier sampling,” *arXiv preprint arXiv:1507.05592*, (2015).
- [13] T. Morimae, K. Fujii, and J. F. Fitzsimons, “Hardness of classically simulating the one-clean-qubit model,” *Phys. Rev. Lett.* **112**, 130502 (2014).
- [14] T. Morimae, K. Fujii, and H. Nishimura, “Power of one nonclean qubit,” *Phys. Rev. A* **95**, 042336 (2017).
- [15] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, R. Babbush, N. Ding, Z. Jiang, M. J. Bremner, J. M. Martinis, and H. Neven, “Characterizing quantum supremacy in near-term devices,” *Nature Phys.* **14**, 595-600 (2018).
- [16] A. Bouland, J. F. Fitzsimons, and D. E. Koh, “Complexity Classification of Conjugated Clifford Circuits,” in *33rd Computational Complexity Conference (CCC 2018)* (R. A. Servedio, ed.), vol. 102 of *Leibniz International Proceedings in Informatics (LIPIcs)*, (Dagstuhl, Germany), pp. 21:1–21:25, Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2018.
- [17] D. J. Brod, “Efficient classical simulation of matchgate circuits with generalized inputs and measurements,” *Phys. Rev. A* **93**, 062332 (2016).
- [18] R. Jozsa and A. Miyake, “Matchgates and classical simulation of quantum circuits,” *Proc. R. Soc. A* **464**, 3089–3106 (2008).

- [19] V. Veitch, C. Ferrie, D. Gross, and J. Emerson, “Negative quasi-probability as a resource for quantum computation,” *New J. Phys.* **14**, 113011 (2012).
- [20] V. Veitch, N. Wiebe, C. Ferrie, and J. Emerson, “Efficient simulation scheme for a class of quantum optics experiments with non-negative Wigner representation,” *New J. Phys.* **15**, 013037 (2013).
- [21] A. Mari and J. Eisert, “Positive Wigner functions render classical simulation of quantum computation efficient,” *Phys. Rev. Lett.* **109**, 230503 (2012).
- [22] S. Bravyi and D. Gosset, “Improved classical simulation of quantum circuits dominated by Clifford gates,” *Phys. Rev. Lett.* **116**, 250501 (2016).
- [23] R. S. Bennink, E. M. Ferragut, T. S. Humble, J. A. Laska, J. J. Nutaro, M. G. Pleszkoch, and R. C. Pooser, “Unbiased simulation of near-Clifford quantum circuits,” *Phys. Rev. A* **95**, 062337 (2017).
- [24] M. J. Bremner, A. Montanaro, and D. J. Shepherd, “Achieving quantum supremacy with sparse and noisy commuting quantum computations,” *Quantum* **1**, 8 (2017).
- [25] M. Oszmaniec and D. J. Brod, “Classical simulation of photonic linear optics with lost particles,” *New J. Phys.* **20**, 092002 (2018).
- [26] M. Howard and E. Campbell, “Application of a resource theory for magic states to fault-tolerant quantum computing,” *Phys. Rev. Lett.* **118**, 090501 (2017).
- [27] H. Pashayan, J. J. Wallman, and S. D. Bartlett, “Estimating outcome probabilities of quantum circuits using quasiprobabilities,” *Phys. Rev. Lett.* **115**, 070501 (2015).
- [28] M. Van Den Nest, “Efficient classical simulations of quantum Fourier transforms and normalizer circuits over Abelian groups,” *arXiv preprint arXiv:1201.4867*, 2012.
- [29] J. Bermejo-Vega and M. Van Den Nest, “Classical simulations of Abelian-group normalizer circuits with intermediate measurements,” *Quantum Info. Comput.* **14**, 181–216 (2014).
- [30] M. Schwarz and M. Van Den Nest, “Simulating quantum circuits with sparse output distributions,” *arXiv preprint arXiv:1310.6749*, 2013.
- [31] S. Bravyi, G. Smith, and J. A. Smolin, “Trading classical and quantum computational resources,” *Phys. Rev. X*, **6**, 021043 (2016).
- [32] K. Temme, S. Bravyi, and J. M. Gambetta, “Error mitigation for short-depth quantum circuits,” *Phys. Rev. Lett.* **119**, 180509 (2017).
- [33] M. Van Den Nest, “Classical simulation of quantum computation, the Gottesman-Knill theorem, and slightly beyond,” *Quantum Info. Comput.* **10**, 258–271 (2010).
- [34] R. Jozsa and N. Linden, “On the role of entanglement in quantum-computational speed-up,” *Proc. R. Soc. A* **459**, 2011–2032 (2003).
- [35] B. M. Terhal and D. P. DiVincenzo, “Adaptive quantum computation, constant depth quantum circuits and arthur-merlin games,” *Quantum Info. Comput.* **4**, 134–145 (2004).
- [36] T. Morimae, “Hardness of classically sampling the one-clean-qubit model with constant total variation distance error,” *Phys. Rev. A* **96**, 040302 (2017).
- [37] S. Aaronson, “The equivalence of sampling and searching,” *Theory of Computing Systems*, **55**, 281–298 (2014).
- [38] M. Van Den Nest, “Simulating quantum computers with probabilistic methods,” *Quantum Info. Comput.* **11**, 784–812 (2011).
- [39] D. Shepherd, “Binary matroids and quantum probability distributions,” *arXiv preprint arXiv:1005.1744*, (2010).
- [40] S. Bravyi and A. Kitaev, “Universal quantum computation with ideal clifford gates and noisy ancillas,” *Phys. Rev. A* **71**, 022316 (2005).
- [41] R. Jozsa and M. Van Den Nest, “Classical simulation complexity of extended Clifford circuits,” *Quantum Info. Comput.* **14**, 633–648 (2014).

- [42] L. A. Goldberg and H. Guo, “The complexity of approximating complex-valued Ising and Tutte partition functions,” *comput. complex.* **26**, 765–833 (2017).
- [43] G. Kuperberg, “How hard is it to approximate the Jones polynomial?,” *Theory of Computing*, **11**, 183–219 (2015).
- [44] K. Fujii and T. Morimae, “Commuting quantum circuits and complexity of Ising partition functions,” *New J. Phys.* **19**, 033003 (2017).
- [45] D. Stahlke, “Quantum interference as a resource for quantum speedup,” *Phys. Rev. A* **90**, 022302 (2014).
- [46] D. Shepherd and M. J. Bremner, “Temporally unstructured quantum computation,” *Proc. R. Soc. A* **465**, 1413–1439 (2009).
- [47] D. J. Shepherd, “Quantum complexity: restrictions on algorithms and architectures,” *arXiv preprint arXiv:1005.1425* (2010).
- [48] D. Vertigan, “Bicycle dimension and special points of the Tutte polynomial,” *J. Combin. Theory Ser. B* **74**, 378–396 (1998).
- [49] R. Renner and S. Wolf, “Smooth Rényi entropy and applications,” in *International Symposium on Information Theory, 2004. ISIT 2004. Proceedings.*, pp. 233 (2004).
- [50] L. Stockmeyer, “The complexity of approximate counting,” in *Proceedings of the fifteenth annual ACM symposium on Theory of computing*, pp. 118–126, ACM, (1983).
- [51] S. Toda, “PP is as hard as the polynomial-time hierarchy,” *SIAM Journal on Computing* **20**, 865–877 (1991).
- [52] S. Aaronson and T. Hance, “Generalizing and derandomizing Gurvits’s approximation algorithm for the permanent,” *Quantum Info. Comput.* **14**, 541–559 (2014).

A Statistical indistinguishability proof

We first show a well known connection between the optimal probability of choosing the correct hypothesis in a hypothesis test and the L_1 distance.

Suppose $\mathcal{P}_1$ and $\mathcal{P}_2$ are probability distribution over some finite set I , and suppose a sample X is observed from the distribution $\mathcal{Q}$ where either $\mathcal{Q} = \mathcal{P}_1$ (hypothesis H_1) or $\mathcal{Q} = \mathcal{P}_2$ (hypothesis H_2). Then, any hypothesis test must have some H_1 acceptance region $A_1 \subseteq I$ and some H_2 acceptance region $A_2 := A_1^c \subseteq I$. The probability of a type I error is $\alpha := \Pr(X \in A_2 \mid X \sim \mathcal{P}_1)$ and the probability of a type II error is $\beta := \Pr(X \in A_1 \mid X \sim \mathcal{P}_2)$. The L_1 distance between $\mathcal{P}_1$ and $\mathcal{P}_2$ can be written as:

$$\begin{aligned} \|\mathcal{P}_1 - \mathcal{P}_2\|_1 &:= \sum_{x \in I} |\mathcal{P}_1(x) - \mathcal{P}_2(x)| \\ &= 2 \sup_{A_1 \subseteq I} [\mathcal{P}_1(A_1) - \mathcal{P}_2(A_1)] \\ &= 2 \sup_{A_1 \subseteq I} [(1 - \mathcal{P}_1(A_1^c) - \mathcal{P}_2(A_1))] \\ &= 2(1 - \alpha^* - \beta^*) \end{aligned}$$

where, the second equality can be verified by noting that the supremum is achieved when $A_1 = \{x \in I \mid \mathcal{P}_1(x) \geq \mathcal{P}_2(x)\}$. Here, α^* and β^* are the type I and type II errors for the optimal choice of acceptance region / hypothesis test. We note that if a priori, H_1 and H_2 are equally likely, then the probability of choosing the correct hypothesis, based on a single sample, using the optimal test is thus given by:

$$\begin{aligned} P_{\text{correct}} &= 1 - \Pr(X \in A_2 \mid X \sim \mathcal{P}_1)\Pr(X \sim \mathcal{P}_1) - \Pr(X \in A_1 \mid X \sim \mathcal{P}_2)\Pr(X \sim \mathcal{P}_2) \\ &= 1 - \alpha^*\Pr(H_1) - \beta^*\Pr(H_2) \\ &= \frac{1}{2} + \frac{\|\mathcal{P}_1 - \mathcal{P}_2\|_1}{4}. \end{aligned} \tag{35}$$

The interactive protocol between the referee and the candidate will proceed as follows (see Figure 3):

1. Initially, the referee will fix a test by choosing a function $a(\cdot)$ that dictates how all gathered data in prior rounds determines the next circuit request. We note that while this can be further generalized by allowing stochastic maps (rather than functions), this has no bearing on our results and our proof can fairly easily be extended if required.
2. Initially the referee will make the circuit request $a_0 \in \mathcal{A}^*$
3. The response from the candidate is denoted by the random variable $\tilde{Y}_{a_0}$ and the string of random variables $a_0, \tilde{Y}_{a_0}$ will be represented by $\tilde{X}_1$
4. The referee may make another circuit request by applying the map a to $\tilde{X}_1$ thus defining the next circuit request $a(\tilde{X}_1)$.
5. On the $(j+1)^{\text{th}}$ round, the referee's circuit request will be represented by $a(\tilde{X}_j)$ and the response will be represented by $\tilde{Y}_{a(\tilde{X}_j)}$ where, $\tilde{X}_{j+1}$ represents the string of random variables $\tilde{X}_j, a(\tilde{X}_j), \tilde{Y}_{a(\tilde{X}_j)}$.
6. In addition, at the end of the j^{th} round for $j = 1, 2, \dots$, a fixed stochastic binary map h will be applied to $\tilde{X}_j$ with the outcome determining whether or non to halt the interactive procedure. We will assume that the test will eventually halt and represent the final round of any given test by $m \in \mathbb{N}$.
7. Finally, the referee will decide H_a vs H_b by applying a fixed binary map d to the full collected data set $\tilde{X}_m$.

Figure 3: The figure above shows the interactive protocol between the referee and the candidate. In each round of the protocol, the referee send a circuit description to the candidate. This circuit description is in general given by applying any fixed (possibly stochastic) map to all of the prior data collected by the referee. The candidate's responses ($\tilde{Y}$) may depend only on the circuit request from the current round and the round number. When the candidate is known to be Alice or Bob, we will represent the variables corresponding to $\tilde{Y}$ and $\tilde{X}$ by Y and X or Y' and X' respectively.

We will use the notation convention above but in the case when the candidate is fixed to be Alice, we will remove the tilde (i.e. $\tilde{X}, \tilde{Y} \rightarrow X, Y$) and alternatively when the candidate is fixed to be Bob, we will replace the tilde with a prime (i.e. $\tilde{X}, \tilde{Y} \rightarrow X', Y'$).

The set of all possible data collected by the referee (based on all probabilistic choices including the choice of the candidate) over the course of the entire test can be viewed as a tree where each branch corresponds to a distinct observed value of the random variable $\tilde{X}_m$. The two probability distribution $\mathcal{P}_1$ and $\mathcal{P}_2$ discussed above will each correspond to a distribution over all of the branches of this tree conditioned on the choice of candidate. We note that one can easily incorporate probabilistic choices by the referee into the formalism which will only result in an increase in the number of branches of the tree. Eq. (35) shows that Bob can ensure suppression of $P_{correct}$ by suppressing the L_1 distance $\|\mathcal{P}_1 - \mathcal{P}_2\|_1$. However, if Bob has an ϵ -simulator, he can only directly control the L_1 distance between his output and that of Alice for a given circuit request. The proof below culminating in Eq. (39) demonstrates that $\|\mathcal{P}_1 - \mathcal{P}_2\|_1$ is sub-additive in the L_1 distance of each circuit request thus Bob can upper bound $P_{correct}$ by bounding each round's L_1 distance in such a way as to ensure the sum converges to the desired bound for $\|\mathcal{P}_1 - \mathcal{P}_2\|_1$.

Proof. We now prove each direction of the “if and only if” statement of Thm. 1.

“ $\Rightarrow$ ” Here, we assume that Bob's simulation scheme is an ϵ -simulator over $\mathcal{C}$ and explicitly specify a strategy for Bob which simultaneously achieves indistinguishability and efficiency.

Bob's strategy will be as follows; if he becomes the candidate, then in the j^{th} round of the protocol, he will be asked to report the outcome of running some circuit indexed by $a_j \in \mathcal{A}^*$. In this case, Bob will ϵ -simulate the circuit c_{a_j} with the precision setting given by:

$$\epsilon_j = \frac{24\delta}{\pi^2 j^2}$$

We note that Bob's strategy as outlined above is fixed and independent of the referee's hypothesis test. Further, we note that for all $m \in \mathbb{N}$:

$$\begin{aligned} \sum_{j=1}^m \epsilon_j &\leq \sum_{j=1}^{\infty} \epsilon_j \\ &= 4\delta \end{aligned} \quad (36)$$

We define the map $\mathcal{E}[X, X']$ from any pair of random variables X with probability distribution $\mathcal{P}$ and X' with probability distribution $\mathcal{P}'$ to $\mathbb{R}$ as the L_1 distance between $\mathcal{P}$ and $\mathcal{P}'$.

We will show that for every test, the quantity on the LHS of Eq. (36) upper bounds $\mathcal{E}[X_m, X'_m]$. Hence:

$$\mathcal{E}[X_m, X'_m] \leq 4\delta \quad (37)$$

$$\begin{aligned} \mathcal{E}[X_{j+1}, X'_{j+1}] &= \sum_{\alpha, \beta} \left| \Pr(Y_{a(X_j)} = \beta | X_j = \alpha) \Pr(X_j = \alpha) - \Pr(Y'_{a(X'_j)} = \beta | X'_j = \alpha) \Pr(X'_j = \alpha) \right| \\ &= \sum_{\alpha, \beta} \left| \Pr(Y_{a(X_j)} = \beta | X_j = \alpha) \Pr(X_j = \alpha) - \Pr(Y_{a(X_j)} = \beta | X_j = \alpha) \Pr(X'_j = \alpha) \right. \\ &\quad \left. + \Pr(Y_{a(X_j)} = \beta | X_j = \alpha) \Pr(X'_j = \alpha) - \Pr(Y'_{a(X'_j)} = \beta | X'_j = \alpha) \Pr(X'_j = \alpha) \right| \\ &\leq \sum_{\alpha, \beta} \Pr(Y_{a(X_j)} = \beta | X_j = \alpha) |\Pr(X = \alpha) - \Pr(X' = \alpha)| \\ &\quad + \sum_{\alpha, \beta} \Pr(X'_j = \alpha) \left| \Pr(Y_{a(X_j)} = \beta | X_j = \alpha) - \Pr(Y'_{a(X'_j)} = \beta | X'_j = \alpha) \right| \\ &\leq \mathcal{E}[X_j, X'_j] + \sum_{\alpha} \mathcal{E}[Y_{a(\alpha)}, Y'_{a(\alpha)}] \Pr(X'_j = \alpha) \end{aligned} \quad (38)$$

where the sums are taken over α in the support of $\tilde{X}_j$ and β in $\bigcup_{a \in \mathcal{A}^*} \text{supp}(\tilde{Y}_a)$.

We note that the precision of Bob's response in any round only depends on the round number. Thus, Eq. (38) can be simplified by replacing $\mathcal{E}[Y_{a(\alpha)}, Y'_{a(\alpha)}]$ with the upper bound ϵ_{j+1} . Combined with the observation that $\mathcal{E}[X_1, X'_1] = \epsilon_1$, we have shown that:

$$\begin{aligned} \mathcal{E}[X_m, X'_m] &\leq \sum_{j=1}^m \epsilon_j \\ &\leq 4\delta \end{aligned} \quad (39)$$

This proves that Bob's strategy meets the indistinguishability property. We now consider the efficiency of the strategy. We recall that given a circuit request sequence α , Alice's and Bob's resource costs are represented by $N(\alpha)$ and $T(\alpha)$ respectively. Further, Alice's resource costs is lower bounded by m , the number of rounds of the Hypothesis test α .

By definition of ϵ -simulation, there exists $\kappa, c_1, c_2 \in \mathbb{N}$ such that for a given circuit index a , and precision ϵ , $T(a) \leq c_1 \left(\frac{N(a)}{\epsilon} \right)^\kappa + c_2$. For simplicity, we will set $c_1 = 1$ and $c_2 = 0$ as this is immaterial given sufficiently large $N(\alpha)$ and $\frac{1}{\epsilon}$. For $m = 1$, clearly the strategy is efficient. Hence,

given a string of inputs $\alpha = (a_1, \dots, a_m)$, with $m \geq 2$ we have:

$$T(\alpha) = \sum_{j=1}^m T(a_j) \quad (40)$$

$$\leq \sum_{j=1}^m \left(\frac{N(a_j)}{\epsilon_j} \right)^\kappa \quad (41)$$

$$= \sum_{j=1}^m \left(\frac{\pi^2 j^2 N(a_j)}{24\delta} \right)^\kappa \quad (42)$$

$$\leq \left(\frac{\pi^2}{24\delta} \right)^\kappa \left[\sum_{j=1}^{m-1} j^{2\kappa} + m^{2\kappa} [N(\alpha) - (m-1)]^\kappa \right] \quad (43)$$

$$\leq \left(\frac{\pi^2}{24\delta} \right)^\kappa \left[\left(\frac{m-0.5}{2\kappa+1} \right)^{2\kappa+1} + m^{2\kappa} N(\alpha)^\kappa - m^{2\kappa} (m-1)^\kappa \right] \quad (44)$$

$$\leq \left(\frac{\pi^2 m^2 N(\alpha)}{24\delta} \right)^\kappa \quad (45)$$

$$\leq \left(\frac{\pi^2 N(\alpha)^3}{24\delta} \right)^\kappa \quad (46)$$

$$\in O \left(\text{poly}(N(\alpha), \frac{1}{\delta}) \right) \quad (47)$$

where:

- in Eq. (43) we have used the fact that $N(a_j) \geq 1$ for all j hence the expression is maximized when α is chosen such that $N(a_j) = 1$ for $j = 1, \dots, m-1$ and $N(a_m) = N(a_1) + \dots + N(a_m) - (m-1)$;
- in Eq. (44) we have used integration to show the inequality for any $k \in \mathbb{N}$; $\sum_{j=1}^m j^k < \left(\frac{m+0.5}{k+1} \right)^{k+1}$; and
- in Eq. (44) we have also used the fact that for $\kappa > 1$ and $N \geq m > 0$, one can show that $(N-m)^\kappa \leq N^\kappa - m^\kappa$
- in Eq. (45) we have used the inequality $\left(\frac{m-0.5}{2\kappa+1} \right)^{2\kappa+1} - m^{2\kappa} (m-1)^\kappa \leq 0$ for $m \geq 2$ and $\kappa \geq 1$;

hence, there exists a polynomial $f(x, y)$ such that for all request strings α and $\delta > 0$, $T(\alpha) \leq f(N(\alpha), \frac{1}{\delta})$.

“ $\Leftarrow$ ”: We restrict ourselves to interactive protocols consisting of only one round. For each fixed circuit request, under the optimal choice of the decision map d , $\delta \propto \epsilon$ hence for all $c \in \mathcal{C}$ and for all $\epsilon > 0$, Bob must be able to sample from some distribution $\mathcal{P}^\epsilon \in B(\mathcal{P}, \epsilon)$. Further, since Bob’s strategy meets the efficiency condition, for every $a \in \mathcal{A}^*$, Bob must be able to output the sample using resources $\in O(\text{poly}(N(a), \frac{1}{\delta})) \subseteq O(\text{poly}(n, \frac{1}{\epsilon}))$. $\square$

B Strong simulation implies EPSILON-simulation

In this appendix, we will show that the existence of a classical strong simulator of a family of quantum circuits implies the existence of an ϵ -simulator (it can in fact construct an approximate weak simulator based on the stronger notion from Ref. [34]). This algorithm aims to map a bit-string (representing the outcome of running the circuit) to r , which is sampled uniformly from $[0, 1]$. While such a mapping is defined for every ordering of the measurement outcomes, it cannot be efficiently computed. This algorithm makes intuitive use of marginal probability estimates to

do a binary search for the bit-string corresponding to r . This technique avoids computing ratios of probability estimates making it useful in regimes where additive errors are small but larger than some of the probabilities in the target distribution.

We start by giving a more precise definition of a strong simulator (than was presented in Sec. 2.1).

Definition 7. (strong simulator). *A strong simulator of a uniform family of quantum circuits $\mathcal{C} = \{c_a \mid a \in \mathcal{A}^*\}$ with associated family of probability distributions $\mathbb{P} = \{\mathcal{P}_a \mid a \in \mathcal{A}^*\}$ is a classical algorithm that, for all $a \in \mathcal{A}^*$, $\epsilon, \delta > 0$ and $S \in \{0, 1, \bullet\}^{k_n}$, can be used to compute an estimate $\hat{p}$ of $\mathcal{P}_a(S)$ such that $\hat{p}$ satisfies the accuracy requirement:*

$$\Pr(|p - \hat{p}| \geq \epsilon) \leq \delta \quad (48)$$

and, the run-time required to compute the estimate $\hat{p}$ is $O(\text{poly}(n, \log \epsilon^{-1}, \log \delta^{-1}))$.

We point out that much like a poly-box, a strong simulator outputs estimates of Born probabilities. The key difference is that the precision of a strong simulator is exponential compared to the polynomial precision of a poly-box. In particular, for any polynomial f , a strong simulator can (efficiently in n) output estimates such that Eq. (48) is satisfied for $\epsilon \in \Omega(2^{-f(n)})$ (as opposed to a poly-box which generally requires $\epsilon \in \Omega(1/f(n))$). Hence, we note that the only difference between the definition of a strong simulator and that of a poly-box is the scaling of run-time in ϵ .

Theorem 5. *Let $\mathcal{C}$ be a uniform family of quantum circuits. If $\mathcal{C}$ admits a strong simulator, then $\mathcal{C}$ admits an ϵ -simulator.*

In fact we will prove an even stronger statement; that a strong simulator implies approximate weak simulation in the much stronger sense of approximate weak simulation used in Ref. [34] (exponentially small error in L_1 norm).

Before proving this theorem, we introduce an algorithm that uses output from a strong simulator to approximately sample from the output distribution of a circuit i.e. to produce output consistent with the definition of an ϵ -simulator. Without loss of generality, let $c \in \mathcal{C}$ be an arbitrary n qubit circuit with all n qubits measured. We will denote the quantum probabilities by p_S and the output of the strong simulator by $p_S^{\epsilon, \delta}$ suppressing the dependence on c .

To give a rough intuition, the algorithm will first sample a polynomial length bit-string $\tilde{r}$ which will be mapped to a probability $r \in [0, 1]$. This value will remain fixed and be used throughout the algorithm until a sample $\tilde{X}$ is generated from the approximate output distribution. This sample will be the output of the ϵ -simulator upon a single execution with the input (ϵ', c) . The sample $\tilde{X} = (\tilde{X}_1, \dots, \tilde{X}_n)$ will be generated by sampling one bit at a time starting with $\tilde{X}_1$. The choice of the j^{th} bit $\tilde{X}_j$ is based on the comparisons between the output of the strong simulator $p_S^{\epsilon, \delta}$ and the probability r . This n step process will require n calls to the strong simulator where in each call, the only variation in the inputs is the events S_j . Each event S_j will be chosen based on the previously sampled values $\tilde{X}_1, \dots, \tilde{X}_{j-1}$.

The algorithm will proceed as follows:

1. Fix $m \in \mathbb{N}$ and $\epsilon, \delta > 0$ based on $\mathcal{C}$ and the desired L_1 error upper bound, ϵ' (see later).
2. Sample $\tilde{r}$ uniformly from $\{0, 1\}^m$.
3. Compute $r = \sum_{i=1}^m \tilde{r}_i 2^{-i}$
4. Set $S := (s_1, \dots, s_n) = (\bullet, \dots, \bullet)$.
5. Set $j = 1$.
6. Set $s_j = 0$.
7. Set $S_j = S$.
8. Request $p_{S_j}^{\epsilon, \delta}$ from the strong simulator.
9. If $p_{S_j}^{\epsilon, \delta} \geq r$, then set $\tilde{X}_j = 0$ otherwise, set $\tilde{X}_j = 1$.

10. Set $s_j = \tilde{X}_j$.
11. If $j = n$, output the string $\tilde{X} = (\tilde{X}_1, \dots, \tilde{X}_n)$ and end.
12. Reset $j \rightarrow j + 1$ and go to step 6.

We now prove Theorem 5.

Proof. We wish to show that for all acceptable families of quantum circuits $\mathcal{C}$, choices of $c \in \mathcal{C}$ and $\epsilon' > 0$:

- there exist a polynomially bounded function $f(\epsilon', n)$ which determines m and
- there exist functions for determining ϵ, δ

such that given a strong simulator of $\mathcal{C}$, the above algorithm can be executed in run-time $O(\text{poly}(n, \epsilon'^{-1}))$ and produce output $\tilde{X}$ from a distribution $\tilde{\mathcal{P}}$ satisfying $\tilde{\mathcal{P}} \in B(\mathcal{P}, \epsilon')$.

We note that the probability distribution over $x \in \{0, 1\}^n$ defines a partitioning (up to sets of measure zero) of the unit interval into 2^n intervals¹³ V_x labeled by x such that the uniform measure on these intervals corresponds to the quantum probability of outcome x . That is, we fix the partitioning such that for all $x \in \{0, 1\}^n$:

$$\mu(V_x) = p_x. \quad (49)$$

To be specific, we can define $V_x = [v_x^-, v_x^+]$ where:

$$v_x^- = \sum_{x' < x} p_{x'} \quad (50)$$

$$v_x^+ = \sum_{x' \leq x} p_{x'} \quad (51)$$

$$(52)$$

where, the above order on bit strings x' and x is defined by lexicographical ordering.

We note that given a uniform sample p from the unit interval, p will, up to measure zero, be strictly identified with an outcome $x \in \{0, 1\}^n$ through the mapping $o : [0, 1] \setminus D \rightarrow \{0, 1\}^n$ implicitly defined by $p \in V_{o(p)}$ for all $p \in [0, 1] \setminus D$ where $D := \{p_S | S \in \{0, 1, \bullet\}^n\}$. Further, in the ideal case where the strong simulator produces output which is deterministically exact i.e. $p_S^{\epsilon, \delta} = p_S$ for all S , we note that the above algorithm would, for a given r , produce output $\tilde{X} = o(r)$. For r distributed uniformly on the unit interval, this ensures $\tilde{X}$ is sampled from exactly the quantum distribution. We thus note that two sources of error arise. The first is from the inaccuracies introduced by the strong simulator's output. The second is from having to approximate a uniform sample over $[0, 1]$ by a uniform sample over $\{0, 1\}^m$.

Let $\tilde{p}_x^{\epsilon, \delta}$ denote the probability $\Pr(\tilde{X} = x)$. Then, we have:

$$\tilde{p}_x^{\epsilon, \delta} = \sum_{\tilde{r} \in \{0, 1\}^m} 2^{-m} \Pr(\tilde{X} = x \mid r) \quad (53)$$

Given an interval $V = [v^-, v^+]$ and $\alpha \in \mathbb{R}$, we define:

$$V^\alpha = \begin{cases} [v^- - \alpha, v^+ + \alpha], & \text{if } \alpha \geq 0 \text{ or } v^+ - v^- \geq 2\alpha \\ [\frac{v^- + v^+}{2}, \frac{v^- + v^+}{2}], & \text{otherwise} \end{cases} \quad (54)$$

If $p_x \geq 2\epsilon$ and $r \in V_x^{-\epsilon}$ then:

$$\Pr(\tilde{X} = x \mid r) \geq (1 - \delta)^n. \quad (55)$$

This can be seen by noting that with probability $\geq 1 - \delta$, each requested probability estimate in step 8 will be within ϵ of the corresponding quantum probability resulting in $\tilde{X}_j = o(r)_j$.

¹³Here, we use a looser notion of interval by allowing points $p \in \mathbb{R}$ to constitute an intervals $[p, p]$.

Thus, we have:

$$\tilde{p}_x^{\epsilon, \delta} \geq \sum_{\substack{\tilde{r} \in \{0,1\}^m \\ r \in V_x^{-\epsilon}}} 2^{-m} (1 - \delta)^n \quad (56)$$

$$\geq l_x 2^{-m} (1 - \delta)^n \quad (57)$$

$$\geq [p_x - 2\epsilon - 2^{-m}] (1 - \delta)^n \quad (58)$$

where

$$l_x := \left\lfloor \frac{|V_x^{-\epsilon}|}{2^{-m}} \right\rfloor \quad (59)$$

$$= \left\lfloor \frac{p_x - 2\epsilon}{2^{-m}} \right\rfloor \quad (60)$$

is a lower bound on the number of bit strings $\tilde{r}$ which under the map in step 3 must be contained in the interval $V_x^{-\epsilon}$.

If $r \in \mathbb{R} \setminus V_x^{+\epsilon}$ then:

$$\Pr(\tilde{X} = x | r) \leq 1 - (1 - \delta)^n \quad (61)$$

since estimates within ϵ of the target probability at each of the n iterations of step 8 will result in $\tilde{X} \neq x$.

Thus, we also have:

$$\tilde{p}_x^{\epsilon, \delta} \leq \sum_{\substack{\tilde{r} \in \{0,1\}^m \\ r \in V_x^{+\epsilon}}} 2^{-m} \Pr(\tilde{X} = x | r) + \sum_{\substack{\tilde{r} \in \{0,1\}^m \\ r \in \mathbb{R} \setminus V_x^{+\epsilon}}} 2^{-m} 1 - (1 - \delta)^n \quad (62)$$

$$\leq u_x 2^{-m} + [1 - (1 - \delta)^n] \quad (63)$$

$$\leq [p_x + 2\epsilon + 2^{-m}] + [1 - (1 - \delta)^n] \quad (64)$$

where

$$u_x := \left\lfloor \frac{|V_x^{+\epsilon}|}{2^{-m}} \right\rfloor + 1 \quad (65)$$

$$= \left\lfloor \frac{p_x + 2\epsilon}{2^{-m}} \right\rfloor + 1 \quad (66)$$

is an upper bound on the number of bit strings $\tilde{r}$ which under the map in step 3 must be contained in the interval $V_x^{+\epsilon}$.

Thus $p_x \geq 2\epsilon$:

$$[-2\epsilon - 2^{-m}] (1 - \delta)^n - p_x [1 - (1 - \delta)^n] \leq \tilde{p}_x^{\epsilon, \delta} - p_x \leq [2\epsilon + 2^{-m}] + [1 - (1 - \delta)^n] \quad (67)$$

i.e.

$$|\tilde{p}_x^{\epsilon, \delta} - p_x| \leq [2\epsilon + 2^{-m}] + [1 - (1 - \delta)^n] \quad (68)$$

Also, if $p_x \leq 2\epsilon$, then:

$$-p_x \leq \tilde{p}_x^{\epsilon, \delta} - p_x \leq [2\epsilon + 2^{-m}] + [1 - (1 - \delta)^n] \quad (69)$$

thus, the bound from Eq. (68) also applies in this case.

This implies that:

$$\epsilon' \leq 2^n [2\epsilon + 2^{-m} + 1 - (1 - \delta)^n]. \quad (70)$$

Clearly there exist choices of polynomials f_1, f_2, f_3 such that for $\epsilon' = \frac{1}{\text{poly}(n)}$ or even $\epsilon' = 2^{-\text{poly}(n)}$, Eq. (70) can be satisfied by choosing $\epsilon \leq 2^{-f_1(n)}$, $\delta \leq 2^{-f_2(n)}$ and $m \geq f_3(n)$. We complete the proof by noting that these choices ensure that the run-time of the strong simulator and the above algorithm are efficient in n and $1/\epsilon'$. $\square$

C Multiplicative precision simulation implies EPSILON-simulation

In this section, we present an algorithm (very similar to Ref. [35]) which uses an estimator with multiplicative precision to construct an ϵ -simulator (it can in fact construct an approximate weak simulator based on the stronger notion from Ref. [35]). This algorithm exploits the fact that ratios of multiplicative precision estimators are multiplicative precision in order to sequentially, one qubit's measurement outcome at a time, sample from the marginal probability of the next qubit's measurement conditioned on the sampled outcomes of the prior measurements. This algorithm and its variants have also been presented in [3, 4] and are well known within the simulation-of-quantum-circuits community.

Here, we claim without proof that this algorithm lifts a classical multiplicative precision simulator of a family of quantum circuits to an approximate weak simulator based on the stronger notion from Ref. [35]. This result has been shown in Ref. [35], but we discuss it here for completeness.

We start by giving a definition of a multiplicative precision simulator.

Definition 8. (multiplicative precision simulator). *A multiplicative precision simulator of a uniform family of quantum circuits $\mathcal{C} = \{c_a \mid a \in \mathcal{A}^*\}$ with associated family of probability distributions $\mathbb{P} = \{\mathcal{P}_a \mid a \in \mathcal{A}^*\}$ is a classical algorithm that, for all $a \in \mathcal{A}^*$, $\epsilon, \delta > 0$ and $S \in \{0, 1, \bullet\}^{k_n}$, can be used to compute an estimate $\hat{p}$ of $\mathcal{P}_a(S)$ such that $\hat{p}$ satisfies the accuracy requirement:*

$$\Pr(|p - \hat{p}| \geq \epsilon p) \leq \delta \quad (71)$$

and, the run-time required to compute the estimate $\hat{p}$ is $O(\text{poly}(n, \epsilon^{-1}, \delta^{-1}))$.

We claim that a multiplicative precision simulator can be used to construct an ϵ -simulator.

Theorem 6. *Let $\mathcal{C}$ be a uniform family of quantum circuits. If $\mathcal{C}$ admits a multiplicative precision simulator, then $\mathcal{C}$ admits an ϵ -simulator.*

We omit a complete proof of this theorem as it makes straightforward use of standard techniques. However, we outline the algorithms which uses output from a multiplicative precision simulator to approximately sample from the output distribution of a circuit. Without loss of generality, let $c \in \mathcal{C}$ be an arbitrary n qubit circuit with all n qubits measured. We will denote the quantum probabilities by p_S and the output of the multiplicative precision simulator by $p_S^{\epsilon, \delta}$ suppressing the dependence on c .

The algorithm will proceed as follows:

1. Fix $m \in \mathbb{N}$ and $\epsilon, \delta > 0$ based on $\mathcal{C}$ and the desired L_1 error upper bound, ϵ' (see later).
2. Set $S := (s_1, \dots, s_n) = (\bullet, \dots, \bullet)$.
3. Set $p_{S_0}^{\epsilon, \delta} := 1$.
4. Set $j = 1$.
5. Set $s_j = 0$.
6. Set $S_j = S$.
7. Request $p_{S_j}^{\epsilon, \delta}$ from the multiplicative precision simulator.
8. Compute $c_j := p_{S_j}^{\epsilon, \delta} / p_{S_{j-1}}^{\epsilon, \delta}$.
9. Sample $\tilde{r}$ uniformly from $\{0, 1\}^m$.
10. Compute $r = \sum_{i=1}^m \tilde{r}_i 2^{-i}$.
11. If $c_j \geq r$, then set $\tilde{X}_j = 0$ otherwise, set $\tilde{X}_j = 1$.
12. Set $s_j = \tilde{X}_j$.
13. If $j = n$, output the string $\tilde{X} = (\tilde{X}_1, \dots, \tilde{X}_n)$ and end.

14. Reset $j \rightarrow j + 1$ and go to step 5.

We note that multiplicative precision estimate can divide each other and still produce a multiplicative precision estimate. Hence c_j computed in step 8 is a multiplicative precision estimate of the quantum conditional probability $p_{S_j}/p_{S_{j-1}} = \Pr(X_j = x_j \mid X_1 = x_1, \dots, X_{j-1} = x_{j-1})$. This ensures that for $\epsilon' = \frac{1}{\text{poly}(n)}$, there exist polynomials f_1, f_2, f_3 such that $\epsilon \leq 1/f_1(n)$, $\delta \leq 1/f_2(n)$ and $m \geq f_3(n)$ satisfy the desired accuracy.

D On poly-sparsity and anti-concentration

In this section we will prove that poly-sparsity and anti-concentration cannot simultaneously be satisfied by any family of quantum circuits. This result is proven in Theorem 7.

The condition of poly-sparsity forces the output distributions over exponentially many outcomes to concentrate on polynomially many outcomes. Alternatively, the property of anti-concentration forces the probabilities of observing any particular outcome, over random choices of circuits, to be low. Intuitively these properties do appear to oppose each other. However, since these properties are statements with respect to different probability spaces, we must first translate each property into a statement about a common probability space and with a common measure. This is done for anti-concentration and poly-sparsity in Lemma 5 and 6 respectively. We then state and prove our main claim in Theorem 7.

First let us restate the relevant definitions in some detail.

Definition 9. (*poly-sparse*) Let $\mathcal{P}$ be a discrete probability distribution. We say that $\mathcal{P}$ is *poly-sparse* if there exists a polynomial $P(x)$ such that for all $\epsilon > 0$, $\mathcal{P}$ is ϵ -approximately t -sparse whenever $t \geq P(\frac{1}{\epsilon})$.

Let $\mathbb{P}$ be a family of probability distributions with $\mathcal{P}_a \in \mathbb{P}$ a distribution over $\{0, 1\}^{k_a}$. We say that $\mathbb{P}$ is *poly-sparse* if there exists a polynomial $P(x)$ such that for all $\epsilon > 0$ and $a \in \mathcal{A}^*$, $\mathcal{P}_a$ is ϵ -approximately t -sparse whenever $t \geq P(k_a/\epsilon)$.

Definition 10. (*anti-concentration*) Let $\mathcal{C}$ be a family of quantum circuits with $\mathbb{P}$ its associated family of probability distributions. For all $n \in \mathbb{N}$ let σ_n be a probability measure over $\mathcal{A}^n$. We say that $\mathcal{C}$ *anti-concentrates* with respect to the set of measures $\Sigma := \{\sigma_n\}_{n \in \mathbb{N}}$ iff $\forall n \in \mathbb{N}, \forall x \in \{0, 1\}^n$ and $\forall \alpha \in (0, 1)$:

$$\Pr_{a \in \mathcal{A}^n} \left(\mathcal{P}_a(x) \geq \frac{\alpha}{2^n} \right) > \frac{(1 - \alpha)^2}{2} \quad (72)$$

where the probability is with respect to the measure σ_n .

Lemma 5. For each $n \in \mathbb{N}$, let σ_n be a probability measure over $\mathcal{A}^n$, let ν_n be any probability measure over $\{0, 1\}^n$ and let τ_n be a probability measure over $\mathcal{A}^n \times \{0, 1\}^n$ defined as the product measure $\sigma_n \times \nu_n$. Then $\mathcal{C}$ anti-concentrates with respect to $\{\sigma_n\}_{n \in \mathbb{N}}$ implies that $\forall n \in \mathbb{N}$ and $\forall \alpha \in (0, 1)$:

$$\Pr_{(a, x) \in \mathcal{A}^n \times \{0, 1\}^n} \left(\mathcal{P}_a(x) \geq \frac{\alpha}{2^n} \right) > \frac{(1 - \alpha)^2}{2} \quad (73)$$

where the probability is taken with respect to τ_n .

Proof. For each $n \in \mathbb{N}$, we first define the sets $S_x := \{(a, x) \in \mathcal{A}^n \times \{0, 1\}^n \mid \mathcal{P}_a(x) \geq \frac{\alpha}{2^n}\}$ and $S'_x := \{a \in \mathcal{A}^n \mid \mathcal{P}_a(x) \geq \frac{\alpha}{2^n}\}$. Let $S := \bigcup_{x \in \{0, 1\}^n} S_x$.

By the definition of anti-concentration, we have that $\forall n \in \mathbb{N}, \forall x \in \{0, 1\}^n$ and $\forall \alpha \in (0, 1)$:

$$\sigma_n(S_x) > \frac{(1 - \alpha)^2}{2}.$$

Let us fix ν_n and note that:

$$\sum_{x \in \{0, 1\}^n} \nu_n(\{x\}) \times \sigma_n(S'_x) > \sum_{x \in \{0, 1\}^n} \nu_n(\{x\}) \times \frac{(1 - \alpha)^2}{2} \quad (74)$$

The LHS of Eq. (74) simplifies as follows:

$$\begin{aligned} \sum_{x \in \{0,1\}^n} \nu_n(\{x\}) \times \sigma_n(S'_x) &= \sum_{x \in \{0,1\}^n} \tau_n(S_x) \\ &= \tau_n(S). \end{aligned}$$

By the fact that ν_n is a probability measure, the RHS simplifies to $\frac{(1-\alpha)^2}{2}$ thus proving the claim. $\square$

Lemma 6. For each $n \in \mathbb{N}$, let σ_n be any probability measure over $\mathcal{A}^n$, let u_n be the uniform probability measure over $\{0,1\}^n$ and let τ_n be a probability measure over $\mathcal{A}^n \times \{0,1\}^n$ defined as the product measure $\sigma_n \times \mu_n$. Then $\mathcal{C}$ is poly-sparse implies that $\forall \beta, \epsilon \in (0,1], \exists n_0 \in \mathbb{N}$ such that $\forall n \geq n_0$ and $\forall \gamma > 0$:

$$\Pr_{(a,x) \in \mathcal{A}^n \times \{0,1\}^n} \left(\mathcal{P}_a(x) \geq \frac{\epsilon}{2^n \gamma (1-\beta)} \right) \leq \gamma + \beta \quad (75)$$

where the probability is taken with respect to τ_n .

Proof. For any family of quantum circuits $\mathcal{C}$, $\forall n \in \mathbb{N}$ and for each $a \in \mathcal{A}^n$, a minimal function $t_a : [0,1] \rightarrow \mathbb{N}$ can be uniquely defined such that $\forall \epsilon \in [0,1]$, the probability of observing an outcome to be one of the $t_a(\epsilon)$ most likely outcomes (when circuit $c_a \in \mathcal{C}$ is run) is $\geq 1 - \epsilon$. Poly-sparsity implies that there exists a polynomial P such that $\forall \epsilon \in (0,1], \forall n \in \mathbb{N}, \forall a \in \mathcal{A}^n, t_a(\epsilon) \leq P(n/\epsilon)$.

We apply Markov's inequality, which states that for R a non-negative random variable and $\gamma > 0$:

$$\Pr \left(R \geq \frac{\mathbb{E}[R]}{\gamma} \right) \leq \gamma. \quad (\text{Markov's inequality})$$

For any fixed $n \in \mathbb{N}$, $a \in \mathcal{A}^n$, let us consider uniformly randomly sampling from one of the $2^n - t_a(\epsilon)$ least likely outcomes. In this case:

$$\begin{aligned} \mathbb{E}[\mathcal{P}_a(x)] &= \frac{\epsilon^+}{2^n - t_a(\epsilon)} \\ &\leq \frac{\epsilon}{2^n - t_a(\epsilon^+)} \\ &= \frac{\epsilon}{2^n - t_a(\epsilon) - 1} \end{aligned}$$

where $\epsilon^\pm$ are the two extremal points of the interval $V := [\epsilon^-, \epsilon^+)$ containing ϵ such that $\forall \kappa \in V, t_a(\kappa) = t_a(\epsilon)$.

This implies that $\forall \epsilon \in (0,1], \forall n \in \mathbb{N}, \forall a \in \mathcal{A}^n, \forall \gamma > 0$, if we uniformly randomly sample from one of the $2^n - t_a(\epsilon)$ least likely outcomes, then:

$$\begin{aligned} \Pr_{x \in \{0,1\}^n} \left(\mathcal{P}_a(x) \geq \frac{\mathbb{E}[\mathcal{P}_a(x)]}{\gamma} \right) &\leq \Pr_{x \in \{0,1\}^n} \left(\mathcal{P}_a(x) \geq \frac{\epsilon}{\gamma(2^n - t_a(\epsilon) - 1)} \right) \\ &\leq \gamma \end{aligned}$$

Hence, for x uniformly sampled over *all* bit-strings:

$$\Pr_{x \in \{0,1\}^n} \left(\mathcal{P}_a(x) \geq \frac{\epsilon}{\gamma(2^n - t_a(\epsilon) - 1)} \right) \leq \gamma + \frac{t_a(\epsilon)}{2^n}$$

i.e.

$$\Pr_{x \in \{0,1\}^n} (\mathcal{P}_a(x) \geq \eta) \leq c \quad (76)$$

where $\eta := \frac{\epsilon 2^{-n}}{\gamma(1 - t_a(\epsilon) 2^{-n} - 2^{-n})}$, $c := \gamma + \frac{t_a(\epsilon)}{2^n}$ and the probability is over the uniform measure u_n over $\{0,1\}^n$.

Let us note that by the fact that poly-sparsity requires that there is a polynomial P such that $t_a(\epsilon) \leq P(n/\epsilon)$, we have that $\forall \beta, \epsilon \in (0, 1]$, $\exists n_0 \in \mathbb{N}$ such that $\forall n \geq n_0$, $\forall a \in \mathcal{A}^n$ and $\forall \gamma > 0$:

$$\beta \geq \frac{t_a(\epsilon) + 1}{2^n}$$

implying that $\eta \leq \frac{\epsilon 2^{-n}}{\gamma(1-\beta)}$ and $c < \gamma + \beta$.

For all $n \in \mathbb{N}$, we define the sets

$$T_a := \{(a, x) \in \mathcal{A}^n \times \{0, 1\}^n \mid \mathcal{P}_a(x) \geq \eta\},$$

$T'_a := \{x \in \{0, 1\}^n \mid \mathcal{P}_a(x) \geq \eta\}$ and $T := \bigcup_{a \in \mathcal{A}^n} T_a$. We now note that Eq. (76) can be rewritten as:

$$u_n(T'_a) \leq c.$$

For any fixed sequence of measures σ_n over $\mathcal{A}^n$, we can define $\tau_n = \sigma_n \times u_n$ to get:

$$\sum_{a \in \mathcal{A}^n} \sigma_n(a) \times u_n(T'_a) \leq \sum_{a \in \mathcal{A}^n} \sigma_n(a) \times c. \quad (77)$$

The LHS of this equation can be simplified as follows:

$$\begin{aligned} \sum_{a \in \mathcal{A}^n} \sigma_n(a) \times u_n(T'_a) &= \sum_{a \in \mathcal{A}^n} \tau_n(T_a) \\ &= \tau_n(T). \end{aligned}$$

By the fact that σ_n is a probability measure, the RHS of Eq. (77) simplifies to c thus proving the claim. $\square$

Theorem 7. *For each $n \in \mathbb{N}$ let σ_n be a measures over $\mathcal{A}^n$ and let $\mathcal{C} = \bigcup_{n \in \mathbb{N}} \{c_a\}_{a \in \mathcal{A}^n}$ be a family of quantum circuits such that $\mathcal{C}$ anti-concentrates with respect to $\Sigma = \{\sigma_n\}_{n \in \mathbb{N}}$. Then $\mathcal{C}$ is not poly-sparse.*

Proof. We will show that the two conditions together give rise to a contradiction and hence are inconsistent. We apply Lemma 5 and set $\alpha = 1/8$ and for each $n \in \mathbb{N}$, $\nu_n = u_n$, the uniform measure over $\{0, 1\}^n$ giving:

$$\Pr_{(a,x) \in \mathcal{A}^n \times \{0,1\}^n} \left(\mathcal{P}_a(x) \geq \frac{1}{2^n \times 2} \right) > \frac{1}{8}. \quad (78)$$

We apply Lemma 6 and set $\beta = \gamma = 1/16$ and $\epsilon = \frac{1}{64} (1 - \frac{1}{16})$ giving:

$$\Pr_{(a,x) \in \mathcal{A}^n \times \{0,1\}^n} \left(\mathcal{P}_a(x) \geq \frac{1}{2^n \times 4} \right) \leq \frac{1}{8} \quad (79)$$

where both Eq. (78) and (79) are with respect to the same measure τ_n thus implying a contradiction. $\square$

Theorem 7 establishes that poly-sparsity and anti-concentration are mutually exclusive properties. However, these properties can nonetheless jointly fail to be satisfied. For example, we can take an infinite family of circuits (growing faster than 2^n) which is poly-sparse then for each n append a single circuit with output distribution that is uniform. This change ensures that the family now breaks poly-sparsity but is insufficient for anti-concentration to be instated. In addition, if we assume that the family admits a poly-box then we notice that by Theorem 3 this family was ϵ -simulable before the change but after the change no longer satisfies the requirements for the application of Theorem 3 despite the fact that we have only added a sequence of uniform distributions which are ϵ -simulable. A clean mathematical characterization of when poly-sparsity and anti-concentration can jointly fail may help inform how to best resolve this undesirable situation.