

A Qubit, a Coin, and an Advice String Walk Into a Relational Problem

Scott Aaronson* Harry Buhrman† William Kretschmer§

Abstract

Relational problems (those with many possible valid outputs) are different from decision problems, but it is easy to forget just *how* different. This paper initiates the study of FBQP/qpoly , the class of relational problems solvable in quantum polynomial-time with the help of polynomial-sized quantum advice, along with its analogues for deterministic and randomized computation (FP , FBPP) and advice ($/\text{poly}$, $/\text{rpoly}$).

Our first result is that $\text{FBQP}/\text{qpoly} \neq \text{FBQP}/\text{poly}$, *unconditionally*, with no oracle—a striking contrast with what we know about the analogous decision classes. The proof repurposes the separation between quantum and classical one-way communication complexities due to Bar-Yossef, Jayram, and Kerenidis. We discuss how this separation raises the prospect of near-term experiments to demonstrate “quantum information supremacy,” a form of quantum supremacy that would not depend on unproved complexity assumptions.

Our second result is that $\text{FBPP} \not\subseteq \text{FP}/\text{poly}$ —that is, *Adleman’s Theorem fails for relational problems*—unless $\text{PSPACE} \subset \text{NP}/\text{poly}$. Our proof uses $\text{IP} = \text{PSPACE}$ and time-bounded Kolmogorov complexity. On the other hand, we show that proving $\text{FBPP} \not\subseteq \text{FP}/\text{poly}$ will be hard, as it implies a superpolynomial circuit lower bound for PromiseBPEXP .

We prove the following further results:

- *Unconditionally*, $\text{FP} \neq \text{FBPP}$ and $\text{FP}/\text{poly} \neq \text{FBPP}/\text{poly}$ (even when these classes are carefully defined).
- $\text{FBPP}/\text{poly} = \text{FBPP}/\text{rpoly}$ (and likewise for FBQP). For *sampling* problems, by contrast, $\text{SampBPP}/\text{poly} \neq \text{SampBPP}/\text{rpoly}$ (and likewise for SampBQP).

*University of Texas at Austin and OpenAI. Email: aaronson@cs.utexas.edu. Supported by a Vannevar Bush Fellowship from the US Department of Defense, the Berkeley NSF-QLCI CIQC Center, a Simons Investigator Award, and the Simons “It from Qubit” collaboration.

†QuSoft, CWI, and University of Amsterdam. Email: buhrman@cwi.nl. Supported in part by the Dutch Research Council (NWO/OCW), Gravitation Programmes Quantum Software Consortium (project number 024.003.037) and Networks (project number 024.002.003).

§Simons Institute for the Theory of Computing and University of California, Berkeley. Email: kretsch@berkeley.edu. Supported by an NDSEG Fellowship and a Simons Quantum Postdoctoral Fellowship. Most of this work done while a student at UT Austin.

1 Introduction

Here is a basic and underappreciated fact: there are computational problems—not distributed or cryptographic tasks, but just pure computational problems—that *provably* admit only randomized solutions. One simple example is: “output an n -bit Kolmogorov-random string.” Another example is: “given as input a halting Turing machine M , output any string *other than* what M outputs when run on its own description.”

Both of these are *relational problems*, defined by a relation $R \subseteq \{0, 1\}^* \times \{0, 1\}^*$. Given an input x , the goal in such a problem is to output any y such that $(x, y) \in R$. The class FP consists of all relations R for which there exists a deterministic polynomial-time algorithm to find a y such that $(x, y) \in R$ whenever one exists. (Ironically, the F stands for “functional,” even though the whole point with relational problems is that they need *not* be functions.)

It is trickier to define FBPP and FBQP, the relational analogues of BPP and BQP respectively. For unlike with decision problems, we can no longer amplify success probabilities by taking majorities, so different allowed error probabilities could lead to different complexity classes. For this reason, a wide variety of definitions of FBPP have appeared in the literature [19, 4, 25, 27, 22]. Having said that, there is one choice that seems more natural than others, which Aaronson [4] made more than a decade ago and which we follow here.¹

Call the relation $R \subseteq \{0, 1\}^* \times \{0, 1\}^*$ *polynomially-bounded* if there exists a polynomial p such that $|y| \leq p(|x|)$ for all $(x, y) \in R$. Then:

Definition 1. FBPP is the class of polynomially-bounded relations $R \subseteq \{0, 1\}^* \times \{0, 1\}^*$ for which there exists a polynomial-time randomized algorithm A such that for all x for which there exists a y with $(x, y) \in R$ and all $\varepsilon > 0$,

$$\Pr[(x, A(x, 0^{1/\varepsilon})) \in R] \geq 1 - \varepsilon,$$

where the probability is over A 's outputs. FBQP is exactly the same except that A can now be a quantum algorithm.

A few comments on this definition: we require A to succeed for any given $\varepsilon > 0$ in order to avoid problems being in FBPP or FBQP for “accidental” reasons, i.e. that the fraction of strings y such that $(x, y) \notin R$ happens to fall below some arbitrary threshold. We allow time polynomial in $1/\varepsilon$ because, as we’ll see, there are natural reductions that need such time. We demand that R be polynomially-bounded because otherwise, A might achieve smaller and smaller error probabilities ε by outputting longer and longer strings, rather than “doing better and better on the same strings,” which is not what we intuitively wanted when we allowed $\text{poly}(n, 1/\varepsilon)$ time. Finally, we do not require that membership in the relation be efficiently verifiable, in contrast to Goldreich’s definition [19, Definition 3.1]. This is for fairness to quantum algorithms: we want FBQP to contain the relational analogues of problems like BosonSampling [6] and Random Circuit Sampling [7] that have played a central role in recently claimed demonstrations of quantum computational supremacy [15, 36]. However, it seems unlikely that such problems can admit efficient verification of membership in the relation.²

¹We also found that GPT-4 [30], when prompted to give a definition for FBPP, settled on one similar to Definition 1. See Appendix A for the transcript.

²For example, if the relation $R \in \text{FBQP}$ defined in [6, Corollary 5.10] had efficient verification of membership, then the Gaussian Permanent Estimation problem $|\text{GPE}|_{\pm}^2$ [6, Problem 1.2] would be solvable in FP^{PH} , thus refuting either the Permanent-of-Gaussians Conjecture [6, Conjecture 1.5], the Permanent Anti-Concentration Conjecture [6, Conjecture 1.6], or $\text{P}^{\#\text{P}} \not\subseteq \text{PH}$.

Already with FBPP and FBQP, some interesting phenomena rear their heads: for example, we’ll observe in Section 4 that $\text{FP} \neq \text{FBPP}$, unconditionally. Note that, because of the requirement to succeed with probability $1 - \varepsilon$ for any $\varepsilon > 0$, this does *not* immediately follow from the examples with which we opened the paper, but it does follow from modifications of those examples, involving time-bounded Kolmogorov complexity or the time-bounded halting problem.

The message of this paper is that the story of FBPP and FBQP becomes wilder still—even more divergent from expectations formed from decision problems—once we bring classical and quantum *advice* into the picture.

1.1 Advice Classes

Karp and Lipton [23] introduced the nonuniform complexity class P/poly and proved the famous theorem that $\text{NP} \subset \text{P/poly}$ would imply the collapse of the polynomial hierarchy. Meanwhile, Adleman [10] proved that $\text{BPP} \subset \text{P/poly}$. Indeed, it is not hard to see that

$$\text{BPP/rpoly} = \text{BPP/poly} = \text{P/rpoly} = \text{P/poly},$$

where */rpoly* means “with polynomial-sized randomized advice,” and P/rpoly is the class of languages that admit nonuniform polynomial-time bounded-error randomized algorithms in which the only randomness comes from the advice. Note that the */rpoly* advice is at least as powerful as the */poly* advice as the */poly* advice can be seen as a distribution with probability concentrated on a single string.

When we come to BQP, it’s natural to ask what happens when the advice can be a quantum state on polynomially many qubits—perhaps a highly entangled state that’s intractable to prepare on one’s own. To capture this question, in 2003 Nishimura and Yamakami [29] defined the class BQP/qpoly , or Bounded-Error Quantum Polynomial-Time with polynomial-size quantum advice.

Definition 2. *BQP/qpoly is the class of all languages $L \subseteq \{0,1\}^*$ for which there exists a polynomial-time quantum algorithm A , a polynomial p , and an infinite list of advice states $\{|\psi_n\rangle\}_{n \geq 1}$, where $|\psi_n\rangle$ is on $p(n)$ qubits, such that for all n and all $x \in \{0,1\}^n$,*

$$\Pr[A(x, |\psi_n\rangle) = L(x)] \geq \frac{2}{3}.$$

Studying BQP/qpoly is one way to formalize the old question of “how much information is in an n -qubit state.” On the one hand, if we think of an n -qubit state $|\psi\rangle$ as a unit vector in \mathbb{C}^{2^n} , then it seems $|\psi\rangle$ could provide an exponential amount of information—say, about every possible input $x \in \{0,1\}^n$. On the other hand, Holevo’s Theorem [21] implies that we can encode at most n bits into n qubits, in such a way that they can be reliably retrieved later by measuring them.

So then, does BQP/qpoly collapse with BQP/poly —that is, BQP with polynomial-sized *classical* advice—or could it be vastly more powerful?

A priori, it’s not even obvious that $\text{BQP/qpoly} \neq \text{ALL}$, where ALL is the class of all languages. Underscoring this worry, it’s easy to show (for example) that $\text{PostBQP/qpoly} = \text{ALL}$, where PostBQP means quantum polynomial time with postselected measurements. To see this, given a language L and an input length n , we just need to consider the advice state

$$|\psi_n\rangle = \frac{1}{\sqrt{2^n}} \sum_{z \in \{0,1\}^n} |z\rangle |L(z)\rangle,$$

where $L(z) = 1$ if $z \in L$ and $L(z) = 0$ otherwise. Then given an input $x \in \{0, 1\}^n$, we first measure $|\psi_n\rangle$ in the standard basis, then postselect on getting the outcome $|x\rangle|L(x)\rangle$.

Despite the sometimes unsettling power of randomized and quantum advice, in 2004, Aaronson [1] proved that $\text{BQP}/\text{qpoly} \subseteq \text{PostBQP}/\text{poly}$. Since Aaronson [2] also showed that $\text{PostBQP} = \text{PP}$, and since adding deterministic advice “commutes” with standard complexity class inclusions, this can be stated equivalently as $\text{BQP}/\text{qpoly} \subseteq \text{PP}/\text{poly}$.

This upper bound on the power of BQP/qpoly has a few implications. First, it immediately implies that $\text{BQP}/\text{qpoly} \neq \text{ALL}$, since $\text{PP}/\text{poly} \neq \text{ALL}$ is easy to show by a counting argument. Second, it means that there is no hope, in the present state of complexity theory, of proving that $\text{BQP}/\text{poly} \neq \text{BQP}/\text{qpoly}$. For any such proof would imply $\text{BQP}/\text{poly} \neq \text{PP}/\text{poly}$, and hence (for example) that PSPACE does not have polynomial-size circuits. At best, one could hope to show that $\text{BQP}^A/\text{poly} \neq \text{BQP}^A/\text{qpoly}$ for some oracle A . As it happens, even *this* is still open, although Aaronson and Kuperberg [9] showed the existence of a unitary oracle U such that $\text{BQP}^U/\text{poly} \neq \text{BQP}^U/\text{qpoly}$, and there’s been recent progress toward replacing this with an ordinary classical oracle [16, 28].

1.2 Relational Complexity Classes with Advice

In quantum computing, it has repeatedly been found that it’s easier to see the advantages of quantum algorithms over classical ones once we switch attention from decision problems to relational and sampling problems. This is what happened, for example, with BosonSampling [6], Random Circuit Sampling [7], and other sampling-based approaches to demonstrating quantum supremacy. It is also what happened with the recent breakthrough of Yamakawa and Zhandry [34], which achieved an exponential quantum speedup relative to a random oracle—but only by switching from decision problems (where the *Aaronson-Ambainis Conjecture* [5] asserts that no such separation is possible) to NP search problems, a particular kind of relational problem.³

In this paper, then, we do something that could’ve been done at any point in the past 20 years, but apparently wasn’t: namely, we ask about the advantages of quantum over classical *advice* on relational problems.

Definition 3. *FBQP/qpoly is the class of polynomially-bounded relations $R \subseteq \{0, 1\}^* \times \{0, 1\}^*$ for which there exists a polynomial-time quantum algorithm Q , a polynomial $p(n, m)$, and an infinite list of advice states $\{|\psi_{n,m}\rangle\}_{n,m \geq 1}$, where $|\psi_{n,m}\rangle$ is on $p(n, m)$ qubits, such that for all x for which there exists a y such that $(x, y) \in R$ and all m ,*

$$\Pr[(x, Q(x, 0^m, |\psi_{n,m}\rangle)) \in R] \geq 1 - \frac{1}{m}.$$

The one subtlety in this definition is that the advice state $|\psi_{n,m}\rangle$ is allowed to depend, not only on the input length n , but on the desired error probability $\varepsilon = 1/m$. We claim that this is simply the “right” choice: efficiency in this setting means time polynomial in n and $1/\varepsilon$, so the advice ought to be allowed to depend on both parameters as well.

Of course, it doesn’t make much sense to feed quantum advice to a classical complexity class (e.g., FP/qpoly). On the other hand, it’s sensible to consider FP/rpoly : this corresponds to classical algorithms that *only* get to use random bits if they come from the advice.

³If one just wants a superpolynomial quantum speedup relative to a random oracle for *some* relational problem—not necessarily an NP search problem—then Aaronson [3] showed that the problem of outputting large Fourier coefficients of a random Boolean function does the job.

Definition 4. FP/rpoly is the class of polynomially-bounded relations $R \subseteq \{0, 1\}^* \times \{0, 1\}^*$ for which there exists a polynomial-time deterministic classical algorithm A , a polynomial $p(n, m)$, and an infinite list of advice distributions $\{\mathcal{D}_{n,m}\}_{n,m \geq 1}$, where $\mathcal{D}_{n,m}$ is supported on $\{0, 1\}^{p(n,m)}$, such that for all x for which there exists a y such that $(x, y) \in R$ and all m ,

$$\Pr_{r \sim \mathcal{D}_{n,m}} [(x, A(x, 0^m, r)) \in R] \geq 1 - \frac{1}{m}.$$

We can similarly define FBPP/rpoly , FBQP/poly , and other possible combinations; we omit the details.

1.3 Our Results

We show that switching attention to relational problems dramatically changes the picture of randomized and quantum computation in the presence of advice.

Our first result is that quantum advice *unconditionally* provides more power than classical advice to solve relational problems:

Theorem 5. $\text{FBQP}/\text{qpoly} \neq \text{FBQP}/\text{rpoly}$.

So in particular, $\text{FBQP}/\text{qpoly} \neq \text{FBQP}/\text{poly}$. Indeed, we shall see that FBQP/qpoly is not contained in FC/poly for arbitrarily powerful uniform complexity classes \mathcal{C} : for example, the class of all computable problems. This is despite the fact that FBQP/qpoly does *not* equal ALL —as can be seen, for example, by considering its restriction to Boolean-valued problems, where it coincides with $\text{PromiseBQP}/\text{qpoly} \subseteq \text{PromisePP}/\text{poly}$.

As we discuss in Section 3, this complexity class separation suggests the possibility of a near-term experiment, which would run an FBQP/qpoly protocol in order to check explicitly whether an entangled state of n qubits (where, say, $n \approx 20$) encodes $\gg n$ bits of classical information. We hope further work will clarify whether such an experiment is feasible with current devices.

Theorem 5 is nonconstructive, and does not give an explicit example of a relation in FBQP/qpoly but not FBQP/rpoly (not counting, e.g., the use of brute force to find the lexicographically first relation that works). We leave the “explicitization” of this separation as one of our central challenges.

Our second result shows that Adleman’s Theorem [10], that $\text{BPP} \subset \text{P}/\text{poly}$, almost certainly does *not* extend to relational problems:⁴

Theorem 6. *If $\text{FBPP} \subset \text{FP}/\text{poly}$, then $\text{PSPACE} \subset \text{NP}/\text{poly}$ (and hence PH collapses).*

We complement Theorem 6 with a result showing that an unconditional proof of $\text{FBPP} \not\subseteq \text{FP}/\text{poly}$ is unlikely in the current state of complexity theory, as it would imply breakthrough circuit lower bounds:

⁴After this manuscript first appeared, Ilango, Li, and Williams [22] implicitly established a conceptually similar result that $\text{FBPP} \not\subseteq \text{FP}/\text{poly}$ under plausible assumptions. They show a conditional lower bound for the range avoidance problem AVOID , which lies in FBPP (whenever the stretch is at least linear). Roughly, [22, Theorem 28] shows that if subexponentially-secure indistinguishability obfuscation exists and coNP is not in NP/poly infinitely often, then $\text{AVOID} \not\subseteq \text{FP}/\text{poly}$. Comparatively, our result seems to weaken the assumption required to separate FBPP from FP/poly , but only in the sense that non-collapse of PH is a better-tested assumption than the existence of indistinguishability obfuscation.

Theorem 7. $\text{FBPP} \subseteq \text{FP}^{\text{PromiseBPEXP}}$. Hence, if $\text{PromiseBPEXP} \subset \text{PromiseP/poly}$, then $\text{FBPP} \subset \text{FP/poly}$.

We also show that, when FP and FBPP are either both given advice or both *not* given advice, the separation between them becomes unconditional:

Theorem 8. $\text{FP} \neq \text{FBPP}$.

Theorem 9. $\text{FP/poly} \neq \text{FBPP/poly}$.

This underscores yet another difference between decision and relational problems: if \mathcal{C}_1 and \mathcal{C}_2 are two uniform classes of promise problems, then the question of whether $\mathcal{C}_1 \subset \mathcal{C}_2/\text{poly}$ is equivalent to the question of whether $\mathcal{C}_1/\text{poly} \subseteq \mathcal{C}_2/\text{poly}$, since an advice string can just be appended to the input. With relational complexity classes such as FBPP, however, this equivalence is no longer immediate, since it doesn't account for how the length of the advice can depend on the error bound ε .

A last question is whether our separation between classical and quantum advice, in the relational setting, extends to a separation between deterministic and randomized advice. We show that the answer is no:

Theorem 10. $\text{FBPP/rpoly} = \text{FBPP/poly} = \text{FP/rpoly}$ and $\text{FBQP/rpoly} = \text{FBQP/poly}$.

See Figure 1 for the complexity class inclusion diagram that emerges from our results about relational classes.

We remark that several of our results are sensitive to the choices we made in defining FBPP and its variants, especially in regards to error reduction. In Section 5, we explore the consequences of choosing some alternative error bounds in Definition 1. There, we find that Theorems 6, 8, and 9 no longer hold unconditionally if we demand exponential error reduction, meaning that the algorithm outputs a sample consistent with the relation with probability $1 - \varepsilon$ in time *polylogarithmic* in $1/\varepsilon$.⁵ So, our results could be interpreted in two different ways: either as showing a striking contrast between relational and decisional classes, or as showing the remarkable power of FBPP when we don't demand exponential error reduction. We leave it to the reader to decide, and hope that this work inspires more discussion about subtleties in the definitions.

1.4 Quantum Communication Complexity

As it turns out, essentially everything we need to prove Theorem 5 was proved 20 years ago, by Bar-Yossef, Jayram, and Kerenidis [12]—though the fact that this is so is buried in their paper. These authors considered separations between randomized and quantum *one-way communication complexities*. That is, they considered the setting where Alice has an input x , Bob has an input y , and Alice can send a message m_x to Bob, which should then allow Bob to compute some joint property of x and y .

Let T be a task, which might be the evaluation of a Boolean function $f(x, y)$, but might also be a sampling or relational problem. We define $D^1(T)$, $R^1(T)$, and $Q^1(T)$ to be the minimum number of bits sent from Alice to Bob in any deterministic, bounded-error randomized, or bounded-error

⁵By contrast, Theorem 5 is unaffected by such a change in definition, because the FBQP/qpoly algorithm used in our proof will turn out to be errorless.

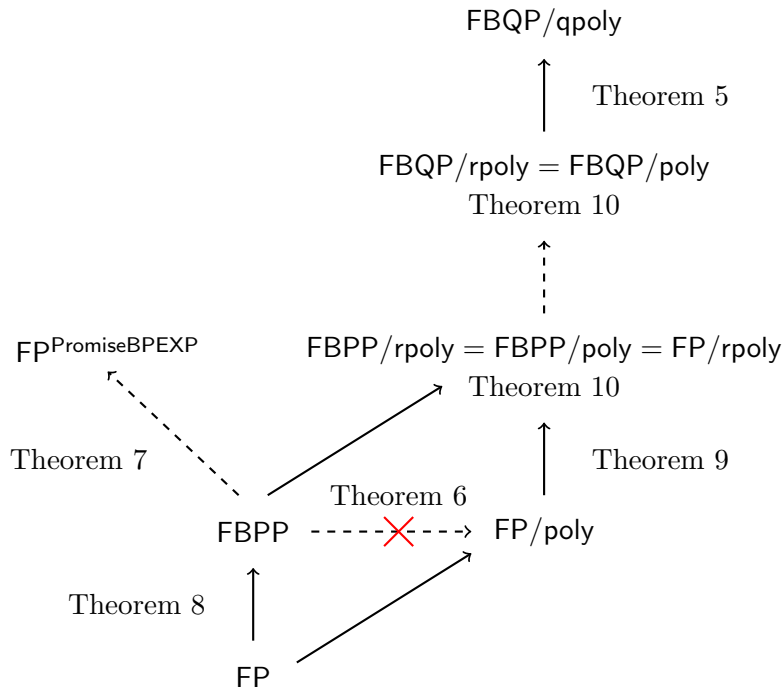


Figure 1: Relationships among classes of relational problems considered in this paper. A solid arrow from \mathcal{C}_1 to \mathcal{C}_2 indicates strict containment ($\mathcal{C}_1 \subsetneq \mathcal{C}_2$). A dashed arrow indicates a containment $\mathcal{C}_1 \subseteq \mathcal{C}_2$ that we conjecture to be strict, but a proof of strictness would require a breakthrough in complexity theory. A crossed dashed arrow indicates non-containment ($\mathcal{C}_1 \not\subseteq \mathcal{C}_2$) under the assumption that PH does not collapse.

quantum one-way communication protocol respectively that lets Bob perform the task for all valid input pairs (x, y) (with the number of bits maximized over all such input pairs). We assume no shared randomness or entanglement.

Clearly $D^1(T) \geq R^1(T) \geq Q^1(T)$ for all tasks T . A natural question is how large the separations between the measures can be. It's well-known that D^1 and R^1 can be exponentially separated: for example, for the N -bit EQUALITY function EQ, we have $D^1(\text{EQ}) = N$ while $R^1(\text{EQ}) = O(\log N)$. But what about R^1 versus Q^1 ?

To study this, Bar-Yossef, Jayram, and Kerenidis [12] defined a relation problem called *Hidden Matching* or HM. Here Alice is given a string $x \in \{0, 1\}^N$ (with N even), while Bob is given a perfect matching y on the set $[N]$, consisting of $N/2$ edges. Bob's goal is to output $(i, j, x_i \oplus x_j)$ for *some* edge $(i, j) \in y$. The key result is then the following:

Theorem 11 ([12]). $Q^1(\text{HM}) = O(\log N)$, whereas $R^1(\text{HM}) = \Omega(\sqrt{N})$.

Crucially for us, Bar-Yossef, Jayram, and Kerenidis actually proved the following stronger statement:

Theorem 12 ([12, Proof of Theorem 4.1, page 373]). *Let \mathcal{M} be any set of perfect matchings on $[N]$ that is pairwise edge-disjoint and satisfies $|\mathcal{M}| = \Omega(N)$. Let μ be the distribution over inputs to HM in which Alice's input is uniform in $\{0, 1\}^N$ and Bob's input is uniform in \mathcal{M} . Then, any deterministic one-way protocol for HM that errs with probability at most $1/8$ with respect to μ requires $\Omega(\sqrt{N})$ bits of communication.*

To prove Theorem 5, in Section 2, we adapt Theorem 11 to the setting of FBQP/qpoly, treating the advice as one-way communication from an advisor to the FBQP algorithm.

To understand the situation more deeply, recall the result of Aaronson [1] from before, that $\text{BQP/qpoly} \subseteq \text{PP/poly}$. A direct analogue of that result for one-way communication complexity [1] says that $D^1(f)$ and $Q^1(f)$ are close whenever Bob's input is small:

Theorem 13 ([1]). *For all Boolean functions $f : \{0, 1\}^n \times \{0, 1\}^m \rightarrow \{0, 1\}$ (partial or total),*

$$D^1(f) = O(mQ^1(f) \log Q^1(f)).$$

This paper is pointing out that Theorem 13, and $\text{BQP/qpoly} \subseteq \text{PP/poly}$, both *fail catastrophically* for sampling and relational problems. This seems not to have been known even to experts who we asked. One reason, perhaps, is that the original separation of Bar-Yossef, Jayram, and Kerenidis [12] was partly overshadowed by the later work of Gavinsky et al. [18]. The latter modified the Hidden Matching relational problem to obtain a partial Boolean function, called *Boolean Hidden Matching* or BHM. They then showed that $Q^1(\text{BHM}) = O(\log N)$ whereas $R^1(\text{BHM}) = \Omega(\sqrt{N})$.

We are calling attention to a surprising difference between the original Hidden Matching separation and the later Boolean Hidden Matching one. Namely: we can make Bob's input "small" (say, $O(\log n)$ bits) in the HM separation, *even though we cannot do the same in the BHM separation*. For Boolean f , Theorem 13 shows that an exponential gap between $D^1(f)$ and $Q^1(f)$ is possible only when Bob's input is "large."

1.5 Other Proofs

Let us make a few remarks about our other results, proved in Sections 2 and 4. To show that $\text{FBPP/rpoly} = \text{FBPP/poly}$, we just take deterministic advice that consists of $O(n/\epsilon^2)$ independent samples from the randomized advice distribution, and then appeal to a Hoeffding and union

bound. To show that $\text{FP} \neq \text{FBPP}$, we consider the problem of outputting an n -bit string with large time-bounded Kolmogorov complexity.⁶ To show that $\text{FBPP} \subseteq \text{FP}^{\text{PromiseBPEXP}}$, we give a simple polynomial-time algorithm that builds a string in the relation one bit at a time, using the PromiseBPEXP oracle.

Finally, and most interestingly, to show that a “relational Adleman’s Theorem” ($\text{FBPP} \subseteq \text{FP/poly}$) is unlikely to hold, we build on an old idea due to Buhrman and Torenvliet [14]. We show that, if the problem of generating strings of high conditional time-bounded Kolmogorov complexity were in FP/poly , then in the $\text{IP} = \text{PSPACE}$ protocol [32], we could replace the randomized verifier by a deterministic polynomial-size circuit. Roughly speaking, the verifier replaces each random challenge with a string of high time-bounded Kolmogorov complexity conditioned on the prior transcript of the protocol. To argue that this derandomization is sound, we just have to show that the “bad” choices of randomness (i.e. those that cause the verifier to accept when it should reject) all have low conditional time-bounded Kolmogorov complexity. We complete the proof by observing that this derandomization would put PSPACE into NP/poly .

1.6 Sampling Problems

We conclude with some results about sampling problems, which are closely related to relation problems. A sampling problem is defined by a collection of probability distributions \mathcal{D}_x . Given an input x , the goal is to output a sample from \mathcal{D}_x , either exactly or approximately. Like for relational problems, we call a sampling problem $S = \{\mathcal{D}_x\}_{x \in \{0,1\}^*}$ *polynomially-bounded* if there exists a polynomial p such that for every x , \mathcal{D}_x is a distribution over strings of length at most $p(|x|)$. Again following Aaronson [4], we define the basic complexity class like so:

Definition 14. *SampBQP is the class of polynomially-bounded sampling problems $S = \{\mathcal{D}_x\}_{x \in \{0,1\}^*}$ for which there exists a polynomial-time quantum algorithm Q such that for all x and all $\varepsilon > 0$,*

$$\|\mathcal{D}_Q(x, 0^{1/\varepsilon}) - \mathcal{D}_x\| \leq \varepsilon,$$

where $\mathcal{D}_Q(x, 0^{1/\varepsilon})$ represents Q ’s output distribution on input $(x, 0^{1/\varepsilon})$ and $\|\cdot\|$ represents total variation distance.

Again, we can consider the classical analogue SampBPP (the deterministic version, SampP , doesn’t make much sense). We can also combine with deterministic, randomized, and quantum advice like in Definition 3, to get SampBPP/poly , SampBQP/qpoly , and so on. For example:

Definition 15. *SampBPP/rpoly is the class of polynomially-bounded sampling problems $S = \{\mathcal{D}_x\}_{x \in \{0,1\}^*}$ for which there exists a polynomial-time randomized algorithm A , a polynomial $p(n, m)$, and an infinite list of advice distributions $\{\mathcal{D}_{n,m}\}_{n,m \geq 1}$, where $\mathcal{D}_{n,m}$ is supported on $\{0,1\}^{p(n,m)}$, such that for all x and all m ,*

$$\|\mathcal{D}_A(x, 0^m, \mathcal{D}_{n,m}) - \mathcal{D}_x\| \leq \frac{1}{m},$$

where $\mathcal{D}_A(x, 0^m, \mathcal{D}_{n,m})$ represents A ’s output distribution on input $(x, 0^m, y)$ averaged over $y \sim \mathcal{D}_{n,m}$ and $\|\cdot\|$ represents total variation distance.

⁶An alternative approach (not shown here) is to prove $\text{FP} \neq \text{FBPP}$ using a direct diagonalization. The core idea of this argument is captured in [19, Section 3.1].

Definition 16. *SampBQP/qpoly* is the class of polynomially-bounded sampling problems $S = \{\mathcal{D}_x\}_{x \in \{0,1\}^*}$ for which there exists a polynomial-time quantum algorithm Q , a polynomial $p(n, m)$, and an infinite list of advice states $\{|\psi_{n,m}\rangle\}_{n,m \geq 1}$, where $|\psi_{n,m}\rangle$ is on $p(n, m)$ qubits, such that for all x and all m ,

$$\|\mathcal{D}_Q(x, 0^m, |\psi_{n,m}\rangle) - \mathcal{D}_x\| \leq \frac{1}{m},$$

where $\mathcal{D}_Q(x, 0^m, |\psi_{n,m}\rangle)$ represents Q 's output distribution on input $(x, 0^m, |\psi_{n,m}\rangle)$ and $\|\cdot\|$ represents total variation distance.

Note that our separations will also hold for the *exact* versions of these sampling classes, but the ε -approximate versions are more robust and seem of greater interest.

Our basic results, proved in Section 6, are as follows. First, we show that sampling classes are more powerful with randomized advice than with deterministic advice:

Theorem 17. $\text{SampBPP/poly} \neq \text{SampBPP/rpoly}$ and $\text{SampBQP/poly} \neq \text{SampBQP/rpoly}$.

To prove Theorem 17, we simply choose a probability distribution over $\{0,1\}^n$ randomly for each n , then appeal to a counting argument.

Second, as a straightforward corollary of Theorem 5, we show that quantum advice provides more power than classical advice for sampling problems:

Theorem 18. $\text{SampBQP/rpoly} \neq \text{SampBQP/qpoly}$.

Theorem 17 contrasts with the situation for relational problems, where $\text{FBPP/poly} = \text{FBPP/rpoly}$ by Theorem 10. This is noteworthy because Aaronson [4] used Kolmogorov complexity to prove a general connection between sampling problems and relational problems. This connection had the following implication, among others:

Theorem 19 ([4]). $\text{FBPP} = \text{FBQP}$ if and only if $\text{SampBPP} = \text{SampBQP}$.

Yet as we now see, the “equivalence” does not force the question of the power of randomized advice to have the same answer for sampling problems that it has for relational problems.

See Figure 2 for a complexity class inclusion diagram that summarizes our results about sampling classes.

2 Deterministic, Randomized, and Quantum Advice

We start this section by observing that for relational problems, randomized advice gives no more power than deterministic advice.

Theorem 10. $\text{FBPP/rpoly} = \text{FBPP/poly} = \text{FP/rpoly}$ and $\text{FBQP/rpoly} = \text{FBQP/poly}$.

Proof. We first prove that $\text{FBPP/rpoly} = \text{FBPP/poly}$. The proof of $\text{FBQP/rpoly} = \text{FBQP/poly}$ is identical but with quantum algorithms in place of randomized algorithms, so we omit it. Let R be a relational problem in FBPP/rpoly , decided by an algorithm A . Fix an input length n and an $\varepsilon > 0$. Let $\mathcal{D}_{n,\varepsilon}$ be the distribution over advice strings. Then for all $x \in \{0,1\}^n$, we must have

$$\Pr_{w \sim \mathcal{D}_{n,\varepsilon}} [(x, A(x, 0^{1/\varepsilon}, w)) \in R] \geq 1 - \varepsilon.$$




Figure 2: Relationships among classes of sampling problems considered in this paper. A solid arrow from \mathcal{C}_1 to \mathcal{C}_2 indicates strict containment ($\mathcal{C}_1 \subsetneq \mathcal{C}_2$). A dashed arrow indicates a containment $\mathcal{C}_1 \subseteq \mathcal{C}_2$ that we conjecture to be strict, but a proof of strictness would require a breakthrough in complexity theory.

In our FBPP/poly simulation, we'll take (say) $k = 100n/\varepsilon^2$ independent samples w_1, \dots, w_k from $\mathcal{D}_{n, \varepsilon/2}$ as the advice. Given an input x , we'll then just pick $i \in \{1, \dots, k\}$ uniformly at random and output $A(x, 0^{2/\varepsilon}, w_i)$. By Hoeffding's inequality, we have that for any fixed $x \in \{0, 1\}^n$,

$$\Pr_{w_1, \dots, w_k} \left[\Pr_i \left[A(x, 0^{2/\varepsilon}, w_i) \in R \right] < 1 - \varepsilon \right] \leq \exp(-k\varepsilon^2/2).$$

Hence, by a union bound over all $x \in \{0, 1\}^n$, there exists some choice of w_1, \dots, w_k that allows the FBPP/poly simulation to succeed with probability at least $1 - \varepsilon$ on every x .

Lastly, we also have $\text{FBPP/rpoly} = \text{FP/rpoly}$, since the randomized advice to an FP machine can include as many uniformly random bits as are needed to simulate any desired FBPP machine. \square

We now prove the unconditional separation between FBQP with quantum advice and FBQP with classical advice.

Theorem 5. $\text{FBQP/qpoly} \neq \text{FBQP/rpoly}$.

Proof. From Theorem 10, it suffices to show that $\text{FBQP/qpoly} \neq \text{FBQP/poly}$. Let $F = \{f_n\}_{n \geq 1}$ be an infinite family of Boolean functions, with

$$f_n : \{0, 1\}^n \rightarrow \{0, 1\}.$$

Then we define the following relation problem:

$$R_F = \{(x, (y, b)) : x, y \in \{0, 1\}^n, b \in \{0, 1\}, f_n(y) \oplus f_n(y \oplus x) = b\}.$$

In other words, given an input $x \in \{0, 1\}^n$, the problem is to output another string $y \in \{0, 1\}^n$, along with a bit b , such that $f_n(y)$ and $f_n(y \oplus x)$ XOR to b .

We first show that, for all F , this problem is in FBQP/qpoly . The quantum advice state is simply

$$|\psi_n\rangle := \frac{1}{\sqrt{2^n}} \sum_{y \in \{0, 1\}^n} (-1)^{f_n(y)} |y\rangle.$$

Given an input $x \in \{0, 1\}^n$, along with $|\psi_n\rangle$, the algorithm is now as follows. If $x = 0^n$, then just output $(y, 0)$ for any $y \in \{0, 1\}^n$. Otherwise, first find a matrix $A \in \mathbb{F}_2^{(n-1) \times n}$ whose nullspace is $\{0, x\}$. Then map $|\psi_n\rangle$ to

$$\frac{1}{\sqrt{2^n}} \sum_{y \in \{0, 1\}^n} (-1)^{f_n(y)} |y\rangle |Ay\rangle$$

and measure the $|Ay\rangle$ register in the computational basis, to reduce the $|y\rangle$ register to the form

$$\frac{(-1)^{f_n(y)} |y\rangle + (-1)^{f_n(y \oplus x)} |y \oplus x\rangle}{\sqrt{2}}$$

for some y . Then measure the above state in the $\{|y\rangle \pm |y \oplus x\rangle\}$ basis, to learn the relative phase $b := f(y) \oplus f(y \oplus x)$. Finally, output y, b . This algorithm succeeds with certainty for every x .

By contrast, Theorem 12 implies that, with probability 1 over the choice of F , the problem is not in FBQP/poly, or indeed in FBQP/rpoly. For each possible input $x \neq 0^n$ gives rise to a matching $\mathcal{M}_x := \{(y, y \oplus x) \mid y \in \{0, 1\}^n\}$ on $\{0, 1\}^n$, and these matchings are pairwise edge-disjoint. So, if we imagine that Alice holds the truth table of a random Boolean function f_n , consisting of $N = 2^n$ bits, while Bob holds a random index x of the matching, we find that Alice must send $\Omega(\sqrt{N}) = \Omega(2^{n/2})$ classical bits to Bob to allow him to satisfy the relation R_F with a success probability of at least $7/8$.

In the actual problem, of course, the function f_n is fixed for each n , rather than chosen by an Alice, and the FBQP/poly algorithm Q 's behavior depends on the f_n 's via the classical advice, rather than a message from Alice. Given a choice of F , let $a_{F,n,m} \in \{0, 1\}^{\text{poly}(n,m)}$ be the advice string for inputs of length n with error $1/m$. Then in order for Q to be correct on $x \in \{0, 1\}^n$, we require that for all m ,

$$\Pr[(x, Q(x, 0^m, a_{F,n,m})) \in R_F] \geq 1 - 1/m.$$

If we imagine that $F = \{f_n\}_{n \geq 1}$ is chosen uniformly at random, then we can bound the probability that Q satisfies this condition on all inputs of length n , i.e.

$$\begin{aligned} & \Pr_F [\forall x \in \{0, 1\}^n : \Pr[(x, Q(x, 0^m, a_{F,n,m})) \in R_F] \geq 1 - 1/m] \\ & \leq \Pr_{F, x \sim \{0, 1\}^n} [\Pr[(x, Q(x, 0^m, a_{F,n,m})) \in R_F] \geq 1 - 1/m] \\ & \leq \frac{m}{m-1} \Pr_{F, x \sim \{0, 1\}^n} [(x, Q(x, 0^m, a_{F,n,m})) \in R_F], \end{aligned}$$

where the last line uses Markov's inequality. Choose $m = 16$, so that $a_{F,n,m}$ is a string of length $\text{poly}(n) = o(2^{n/2})$. Then combining the above bound with Theorem 12 implies that

$$\Pr_F [\forall x \in \{0, 1\}^n : \Pr[(x, Q(x, 0^m, a_{F,n,m})) \in R_F] \geq 1 - 1/m] \leq \frac{16}{15} \cdot \frac{7}{8} = \frac{14}{15}$$

for all sufficiently large n . Moreover, this probability is independent for each $n \in \mathbb{N}$, because each f_n is chosen independently, so the overall probability that any choice of advice allows Q to compute R_F is at most $\prod_{n=1}^{\infty} 14/15 = 0$. This is to say that a uniformly random F satisfies $R_F \notin \text{FBQP/poly}$ with probability 1. \square

Note that, in the proof of $R_F \notin \text{FBQP/poly}$, we nowhere needed the fact that the algorithm was an efficient quantum algorithm (i.e., FBQP), but only that the algorithm succeeds with bounded error. Hence we can conclude more generally that $R_F \notin \text{FC/poly}$ for uniform complexity classes \mathcal{C} with arbitrarily large computational power, such as PSPACE, EXP, BPEXP, R, and so on. We additionally get $R_F \notin \text{FBQP/rpoly}$, because $\text{FBQP/rpoly} = \text{FBQP/poly}$.⁷ On the other hand, we cannot say that $R_F \notin \text{FC/rpoly}$ for any \mathcal{C} , because of the way the success conditions of certain complexity classes interact with randomized advice: as an example, $\text{PostBPP/rpoly} = \text{ALL}$, and so a reasonably defined relational analogue FPostBPP/rpoly certainly *would* contain R_F .

It is interesting to ask just how efficient we can make the quantum algorithm of Theorem 5. We describe how to implement the measurement on $|\psi_n\rangle$ via a simpler circuit, without the need to compute the matrix-vector multiplication Ay . We claim the following: first, the quantum circuit for measuring $|\psi_n\rangle$ and learning the output string y, b can be taken to be a stabilizer circuit. Second, this stabilizer circuit has $O(n)$ size and can be constructed in $O(n)$ time.

To see why, suppose for example that the input x is 001111. Suppose we measure $|\psi_n\rangle$ according to the circuit in Figure 3, and get the result $z = z_1z_2z_3z_4z_5z_6$. We claim that this measurement result corresponds to collapsing the input state to

$$\frac{|y\rangle + (-1)^b|x \oplus y\rangle}{\sqrt{2}},$$

where $y = z_1z_2z_3z_4z_50$ and $b = z_6$. The easiest way to see why is to consider the resulting state when we apply the inverse circuit to $|z\rangle$.

For a general x of Hamming weight $k \geq 1$, we choose an arbitrary i for which $x_i = 1$, and let qubit i play the role of measuring b . The circuit will consist of $k - 1$ CNOT gates between qubit i and the other qubits j for which $x_j = 1$, followed by a single Hadamard gate on qubit i to measure b .

One more comment: our proof of Theorem 5 was nonconstructive, in the sense that we did not exhibit any particular F such that $R_F \notin \text{FBQP/poly}$, but merely used counting to show that a random F works with probability 1. Thus, it is natural to wonder whether we could find an “explicit” F , say $F \in \text{FPSPACE}$, such that $R_F \notin \text{FBQP/poly}$ under a plausible hardness assumption. We do not know, but we would like to observe that for the *promise problem* versions of these classes—namely PromiseBQP/poly and PromiseBQP/qpoly —general principles imply (perhaps surprisingly) that, if there is any separation at all, then the separation can be witnessed “explicitly”:

Proposition 20. *Suppose $\text{PromiseBQP/poly} \neq \text{PromiseBQP/qpoly}$. Then there is a PromisePP problem in PromiseBQP/qpoly but not in PromiseBQP/poly .*

Proof. We prove the contrapositive. Suppose

$$\text{PromisePP} \cap \text{PromiseBQP/qpoly} \subseteq \text{PromiseBQP/poly}.$$

Let $\Pi = (\Pi_Y, \Pi_N)$ be a promise problem in PromiseBQP/qpoly . Aaronson [1] showed that $\text{PromiseBQP/qpoly} \subseteq \text{PromisePP/poly}$. Let $\{w_n\}_{n \geq 1}$ be the advice strings for the PromisePP/poly machine. Then we define a new promise problem Π' , whose yes-instances have the form (x, w_n)

⁷Alternatively, $R_F \notin \text{FBQP/rpoly}$ can be shown directly by a small modification of the above proof: simply replace the advice string $a_{F,n,m}$ with a sample from an advice distribution. This works because Theorem 12 lower-bounds *randomized* one-way communication complexity, not just deterministic, by Yao’s principle [35].




Figure 3: A circuit for measuring the state $|\psi_n\rangle$ in Theorem 5. Note that the circuit will depend on the input x ; an example with $x = 001111$ is shown. First, all qubits i such that $x_i = 0$ are measured in the computational basis. Next, the qubits i such that $x_i = 1$ are measured to determine two things: (i) a computational basis state modulo a NOT gate being applied to each qubit; and (ii) the relative phase between those two basis states, one with the NOT gates applied and the other without (this is the purpose of the sole Hadamard gate).

for $x \in \Pi_Y \cap \{0, 1\}^n$, and whose no-instances have the form (x, w_n) for $x \in \Pi_N \cap \{0, 1\}^n$. Clearly $\Pi' \in \text{PromisePP}$. We also have $\Pi' \in \text{PromiseBQP/qpoly}$, since we can just ignore w_n . By assumption, then, $\Pi' \in \text{PromiseBQP/poly}$. But moving w_n back to the advice, this means that $\Pi \in \text{PromiseBQP/poly}$ as well. Therefore $\text{PromiseBQP/qpoly} = \text{PromiseBQP/poly}$. \square

So in particular, under a plausible complexity assumption, namely $\text{BQP/poly} \neq \text{BQP/qpoly}$, there *is* an explicit problem in FBQP/qpoly but not in FBQP/poly —for observe that promise problems are a special case of relational problems.

Of course, for relational problems we would like to do better than this observation, by constructing an explicit problem in $\text{FBQP/qpoly} \setminus \text{FBQP/poly}$ under a “standard” hardness assumption, one about complexity classes like EXP or PSPACE or P/poly .

3 Toward Quantum Information Supremacy

The extreme simplicity of the circuit in Figure 3 to measure the advice state $|\psi_n\rangle$ —namely, a linear number of 1- and 2-qubit Clifford gates—raises the question of whether an experiment “witnessing” the separation between FBQP/rpoly and FBQP/qpoly might be feasible with current technology, and on a large enough scale to be interesting.

As it happens, Kumar, Kerenidis, and Diamanti [24] reported an experimental demonstration of the Bar-Yossef-Jayram-Kerenidis Hidden Matching protocol [12] in 2018. However, their experiment used an optical coherent state with 2^n modes, and tiny average photon number per mode, in order to simulate n qubits. It therefore didn’t directly test the question of whether n qubits (encoded, say, using n entangled particles) require $\exp(n) \gg n$ bits to simulate classically. This is the question that we propose to test now.

We are finally in the era of small programmable quantum computers: devices that can run more-or-less arbitrary circuits (subject to locality constraints) on ~ 100 qubits and ~ 1000 gates,

and then extract a detectable signal on measurement. Within the past few years, these devices have been used to assert the milestone of “quantum computational supremacy”—that is, a clear advantage over currently-available classical algorithms and hardware—for contrived tasks such as Random Circuit Sampling and BosonSampling (see, e.g., [15, 36]).

Notably, these quantum supremacy tasks are validated using, e.g., Google’s *Linear Cross-Entropy Benchmark* [15], which lets us see them as literally *relational problems*. In Random Circuit Sampling, for example, we are given as input a classical description of an n -qubit quantum circuit C . We are then asked to output any distinct strings $s_1, \dots, s_k \in \{0, 1\}^n$ that satisfy an inequality such as

$$\sum_{i=1}^k |\langle s_i | C | 0^n \rangle|^2 \geq \frac{1.002k}{2^n}.$$

When, say, $k \approx n$, the above yields a relational problem in FBQP that is plausibly conjectured not to be in FBPP (see, e.g., Aaronson and Gunn [8]).

Of course, any conjecture of this sort rests on unproved computational hardness assumptions: if nothing else, then $P \neq PSPACE$! Alas, “standard” hardness assumptions have not sufficed here. The fundamental drawback of current quantum supremacy experiments is that, with each concession that we need to make to experimental reality—for example, depolarizing noise, photon losses, a limited number of qubits (to keep the classical verification feasible), limited circuit depth (to control the noise), limited qubit connectivity, etc.—the relevant hardness assumptions move to shakier and shakier ground. Furthermore, the worry about classical spoofing is far from hypothetical! Classical algorithms for simulating noisy random quantum circuits *have* improved, both in theory and in practice (see, e.g., [17, 31, 11])—if not enough to kill the current claims of quantum supremacy outright, then enough to call them into reasonable doubt.

Thus, wouldn’t it be great if a meaningful quantum supremacy experiment could be designed based on *no* unproved hardness assumptions? Such an experiment might, for example, try to falsify the hypothesis that every “realistic” entangled state of n qubits is secretly describable using $p(n)$ classical bits, for some small polynomial p , whether due to noise or experimental limitations or even a breakdown of quantum mechanics itself. Note that this marks a fundamental difference compared to existing experiments based on (say) Bell inequality violations—whereas the Bell/CHSH experiments test the nonlocal nature of quantum correlations, they in no way test the exponential dimensionality of Hilbert space. As far as we know, the experiments we propose here would be the first directly to test the latter without relying on any unproved computational assumptions.

In such an experiment, we might first prepare a random n -qubit entangled state $|\psi\rangle$, then measure $|\psi\rangle$ in a basis B chosen randomly and “on the fly”—just like Alice’s and Bob’s measurement bases in the Bell/CHSH experiment are ideally chosen when the entangled photons are already in flight. We would repeat this process many times, with a new random basis B each time, collect statistics on the measurement outcomes, and then argue that no $p(n)$ -bit classical digest of $|\psi\rangle$ could possibly have allowed those statistics to be reproduced.

This is exactly what we suggest to do, by leveraging the unconditional separation between FBQP/rpoly and FBQP/qpoly, which in turn is based on the unconditional separation between randomized and quantum one-way communication complexities. Note that such an experiment would almost certainly be impractical with current hardware, if the required measurements on $|\psi\rangle$ were complicated ones. In practice, then, it is essential that we have not merely an exponential separation in one-way communication complexities for a relational problem, but one wherein Bob’s measurements are “simple”—which is the content of FBQP/rpoly \neq FBQP/qpoly.

The detailed consideration of such an experiment is beyond this paper’s scope. Briefly, though, laying the groundwork for this experiment would involve refining and improving the FBQP/rpoly \neq FBQP/qpoly separation in several ways. Firstly, one would want a separation between randomized and quantum advice length that was as *quantitatively tight* as possible, and that also included concrete bounds for particular small numbers of qubits n , such as 20 or 30. Secondly, one would want to account for the complexity of *preparing* the advice state $|\psi\rangle$: for example, what if we chose $|\psi\rangle = C|0^n\rangle$, where C is a random quantum circuit with m gates? Thirdly, one would, if possible, want the measurements of $|\psi\rangle$ to be *even simpler* than the one from Figure 3, and lower-depth: for example, could one even measure each qubit separately? Fourthly, one would want an analysis that accounted for $|\psi\rangle$ being *extremely noisy*—as it will be, in any near-term implementation—and that carefully quantified the experimental resources needed to achieve a clear separation between randomized and quantum advice length even in the teeth of the noise. See Section 7 for further discussion of these challenges, especially the quantitative tightness one.

Once all the challenges are taken into account, the separation between randomized and quantum information achievable with current devices might be rather modest, as it was in the earlier work by Kumar, Kerenidis, and Diamanti [24]. For example, perhaps it will be possible to perform an experiment with $n \approx 20$ qubits, using 2-qubit gates of 99.8% fidelity or whatever, to verify that any secretly classical description of the qubits’ state (even a probabilistic description) would need at least ~ 100 bits, in order to explain the observed success at measuring the state to solve a relational problem. In our view, though, this would already be a historic result, sufficient to disturb the certainty of those who regard the vastness of Hilbert space as just a theoretical fiction. We hope to address some of the challenges in future work.

4 The Power of FBPP

In this section we show several senses in which FBPP behaves differently from its decision-problem counterpart. We start with an unconditional separation between FP and FBPP.

Theorem 8. FP \neq FBPP.

Proof. Recall the definition of Levin’s time-bounded Kolmogorov complexity: given a string y ,

$$\text{Kt}(y) := \min_{P:P()=y} (|P| + \log_2 t(P)) :$$

that is, we minimize the length of a program P (in some fixed programming language) plus the logarithm of P ’s runtime, over all programs P that output y given a blank input. Now consider the relation

$$R = \{(x, y) : |x| = |y|, \text{Kt}(y) \geq \frac{|y|}{2}\}.$$

We first show that $R \in \text{FBPP}$. Given an input of length n , the strategy depends on the allowed error probability ε . If $\varepsilon \geq \frac{1}{2^{n/2}}$, then we can simply output a uniformly random $y \in \{0, 1\}^n$; a counting argument will then imply that $\text{Kt}(y) \geq \frac{n}{2}$ with probability at least $1 - \varepsilon$. If, on the other hand, $\varepsilon < \frac{1}{2^{n/2}}$, then we can use brute force to find and output the lexicographically first string $y \in \{0, 1\}^n$ such that $\text{Kt}(y) \geq \frac{n}{2}$. This takes time exponential in n , so polynomial in $\frac{1}{\varepsilon}$.

Next we show that $R \notin \text{FP}$. Let A be a deterministic algorithm with polynomial running time $p(n)$; then for all n , we clearly have

$$\text{Kt}(A(0^n)) \leq |A| + \log_2 n + \log_2 p(n).$$

But the above is less than $\frac{n}{2}$ for all sufficiently large n , which proves that A cannot compute R . \square

Next we show that the unconditional separation still holds if FP and FBPP both have polynomial-sized advice.

Theorem 9. $\text{FP/poly} \neq \text{FBPP/poly}$.

Proof. For each n and $x \in \{0, 1\}^n$, we choose a subset $S_x \subset \{0, 1\}^n$ uniformly at random and independently subject to $|S_x| = 2^{n/2}$. We then define the following relational problem R :

$$R = \{(x, y) : |y| = |x|, y \notin S_x\}.$$

In other words: given as input an n -bit string x , the problem is to output an n -bit string y that is not in S_x .

We first claim that $R \notin \text{FP/poly}$, with probability 1 over the choice of S_x 's. This is because, for any fixed FP/poly algorithm C and any n , we have

$$\Pr_{\{S_x\}_{x \in \{0,1\}^n}} [(x, C(x)) \in R \forall x \in \{0, 1\}^n] \leq \left(1 - \frac{1}{2^{n/2}}\right)^{2^n} = \frac{1}{\exp(2^{n/2})},$$

which remains small even after we take a union bound over all possible C 's.

By contrast, we claim that $R \in \text{FBPP/poly}$. The algorithm is as follows: if $\varepsilon \geq 1/2^{n/2}$, then just output a uniformly random $y \in \{0, 1\}^n$. If, on the other hand, $\varepsilon < 1/2^{n/2}$, then the advice string can have size $2^n \varepsilon = \text{poly}(n, 1/\varepsilon)$, and can therefore just provide a giant list containing some $y \notin S_x$ for each possible $x \in \{0, 1\}^n$. \square

Finally, we give strong evidence that $\text{FBPP} \not\subseteq \text{FP/poly}$.

Theorem 6. *If $\text{FBPP} \subset \text{FP/poly}$, then $\text{PSPACE} \subset \text{NP/poly}$ (and hence PH collapses).*

Proof. We use an idea of Buhrman and Torenvliet [14], which is in turn based on the $\text{IP} = \text{PSPACE}$ theorem and conditional time-bounded Kolmogorov complexity.⁸

Given strings x and y , we define

$$\text{Kt}(y|x) := \min_{P:P(x)=y} (|P| + \log_2 t(P, x)),$$

or the *time-bounded Kolmogorov complexity of y conditioned on x* , to be the minimum, over all programs P (in some fixed programming language) such that $P(x) = y$, of the bit-length of P plus the log of its runtime on input x .

We now define the following relation:

$$R^* := \{(x, 0^n), y) : y \in \{0, 1\}^n, \text{Kt}(y|x) \geq n/2\}.$$

In other words, given as input x (which could have some arbitrary length $m = \text{poly}(n)$) and 0^n , the problem is to output an n -bit string y that cannot be computed too quickly by any short, deterministic program given x .

⁸Buhrman and Torenvliet [14] used these ideas to prove that $\text{PSPACE} \subseteq \text{NP}^{R_s^{CS}}$, where R_s^{CS} is an oracle to decide whether a given string has maximal space-bounded conditional Kolmogorov complexity. We, by contrast, are interested in the nonuniform complexity of *relational* problems. Outputting strings of large time-bounded conditional Kolmogorov complexity is a convenient relational problem for proving the implication we want.

Our first claim is that $R^* \in \text{FBPP}$. The argument is the same as in the proof of Theorem 8; the fact that we condition on x makes no difference. And so—we save this fact for later—if $\text{FBPP} \subset \text{FP/poly}$, then R^* is also decided by some polynomial-size family of circuits $\{C_{m,n}\}_{m,n \geq 1}$.

We now recall the relevant facts about the proof of $\text{IP} = \text{PSPACE}$ (see [32] for details). Let ϕ be an instance of $TQBF$, the canonical PSPACE -complete problem. Then to verify that $\phi \in TQBF$, the verifier engages the prover in an n -round conversation about a certain complicated (but polynomial-sized) arithmetic expression of the form

$$P = \sum_{x_1 \in \{0,1\}} R_{x_1} \prod_{x_2 \in \{0,1\}} R_{x_1} R_{x_2} \sum_{x_3 \in \{0,1\}} R_{x_1} R_{x_2} R_{x_3} \cdots \prod_{x_n \in \{0,1\}} R_{x_1} \cdots R_{x_n} \varphi(x_1, \dots, x_n)$$

over the finite field \mathbb{F}_q , where we take q to be a prime such that $q > 16^{n^2}$. The expression involves three types of quantifiers over variables: sums, products, and so-called *degree reduction operators* (these are the R_{x_i} 's).

The expression P is constructed by carefully arithmetizing ϕ to maintain the following properties:

1. $P = 1$ if $\phi \in TQBF$ while $P = 0$ if $\phi \notin TQBF$.
2. For any t , if we substitute field values $r_1, \dots, r_{t-1} \in \mathbb{F}_q$ at the first $t - 1$ quantifiers in an appropriate way, remove the t^{th} quantifier, and keep in place everything to the right of the t^{th} quantifier, then we are left with a univariate polynomial $h_t(x_i)$ of degree at most $\text{poly}(n)$ over \mathbb{F}_q , where x_i is the variable that appears in the t^{th} quantifier. (The whole purpose of the degree reduction operators is to ensure this.)

The conversation proceeds in $T \leq n^2 + 1$ rounds, one for each quantifier. At round t , the prover sends the verifier a univariate polynomial $g_t : \mathbb{F}_q \rightarrow \mathbb{F}_q$, and makes the crucial claim that $g_t = h_t$ as polynomials over \mathbb{F}_q , where h_t is the univariate polynomial discussed previously, which depends on the random finite field values $r_1, \dots, r_{t-1} \in \mathbb{F}_q$ chosen by the verifier in the previous rounds. After applying some preliminary checks, the verifier then tests the prover's claim by choosing a new $r_t \in \mathbb{F}_q$ uniformly at random and sending it to the prover, and the conversation continues. Finally, at the very last round, the verifier can check g_T directly against a polynomial obtained by arithmetizing ϕ .

The key fact is that, by the Fundamental Theorem of Algebra, if $g_t \neq h_t$ as polynomials, then $g_t(r)$ and $h_t(r)$ can coincide on at most $\max\{\deg(g_t), \deg(h_t)\} = n^{O(1)}$ values of r . And these are the only values of $r \in \mathbb{F}_q$ that can cause the protocol to fail at round t (in the sense that the verifier will now accept even though $\phi \notin TQBF$).

For technical reasons to be explained later, let $s \in \{0, 1\}^{n^{O(1)}}$ be a polynomial-sized string that is chosen uniformly at random and then fixed.

We now make the following observation: *in place of a uniformly random $r_t \in \mathbb{F}_q$, the verifier could send any $r_t \in \mathbb{F}_q$ such that $\text{Kt}(r_t | s, \phi, r_1, \dots, r_{t-1}, g_t) \geq 2n^2$.*

To see why, consider a “bad” r_t : that is, one such that $g_t(r_t) = h_t(r_t)$, even though $g_t \neq h_t$ as polynomials. As we said, there can be at most $n^{O(1)}$ such bad r_t 's. Furthermore, we claim that the complete list of bad r_t 's can be generated in $2^T n^{O(1)}$ time, given ϕ and s, r_1, \dots, r_{t-1} and g_t as input, with overwhelming probability over the choice of s .

To generate the list, we first compute h_t explicitly as a polynomial over \mathbb{F}_q , by simply “brute-forcing” every sum and product over a variable $x_i \in \{0, 1\}$ and every degree reduction operator

that appears to the right of the t^{th} quantifier. This takes time $2^{T-t}n^{O(1)}$, since we pick up a factor of 2 for every quantifier that needs to be brute-forced. We next factor the polynomial $g_t(r) - h_t(r)$ over the finite field \mathbb{F}_q —for example, by using the randomized algorithm due to Berlekamp [13], which runs in $\text{poly}(n, \log q) = \text{poly}(n)$ time, and which fails with probability at most $1/2^{p(n)}$, where p is a polynomial that we can make as large as needed by choosing a large enough randomness string s to feed to Berlekamp’s algorithm. Finally, from the degree-1 irreducible factors of $g_t - h_t$, we extract the solutions $r \in \mathbb{F}_q$ to $g_t(r) = h_t(r)$.

Thus, letting Π be a program that does the above, for any bad r_t , we have

$$\text{Kt}(r_t|s, \phi, r_1, \dots, r_{t-1}, g_t) \leq |\Pi| + \log_2 \deg(g_t - h_t) + \log_2(2^T n^{O(1)}),$$

which is less than $2T$ for all large enough n and $T \approx n^2$. And so, taking the contrapositive, if r_t has conditional time-bounded Kolmogorov complexity at least $2T$, then it cannot be bad.

But if we set $x := \langle s, \phi, r_1, \dots, r_{t-1}, g_t \rangle$, then the problem of finding an r_t such that

$$\text{Kt}(r_t|x, \phi, r_1, \dots, r_{t-1}, g_t) \geq 2T$$

can be solved by finding a $y \in \{0, 1\}^{4n^2}$ such that $((x, 0^{4n^2}), y)$ is in the relation R^* . And we said that, by the assumption $\text{FBPP} \subset \text{FP/poly}$, there is a polynomial-size circuit family $\{C_{m,n}\}_{m,n \geq 1}$ that does this.

Hence we can decide $TQBF$ in NP/poly , as follows. Given as input a $TQBF$ instance $\phi(x_1, \dots, x_n)$, the polynomial-sized advice provides a description of an appropriate circuit $C_{m,4n^2}$, along with a hardwired value for the randomness string s . Given ϕ and given this advice, the NP prover is asked to provide a complete transcript for the $\text{IP} = \text{PSPACE}$ protocol, *assuming that the verifier generates each of its messages using $C_{m,4n^2}$* . Finally, the NP verifier checks each step in this transcript, using C_n to make sure that the prover computed the IP verifier’s messages correctly.

By the reasoning above, this derandomization of $\text{IP} = \text{PSPACE}$ is sound: any failure would imply that $C_{m,4n^2}$ had generated a message of small time-bounded conditional Kolmogorov complexity. Or more precisely, this is true with overwhelming probability over the choice of s , which means that there must exist fixed s ’s that work when hardwired into the advice. Therefore $\text{PSPACE} \subset \text{NP/poly}$ as claimed. \square

We conclude this section by observing a barrier to any unconditional proof of $\text{FBPP} \not\subset \text{FP/poly}$, as it would lead to new circuit lower bounds.

Theorem 7. $\text{FBPP} \subseteq \text{FP}^{\text{PromiseBPEXP}}$. *Hence, if $\text{PromiseBPEXP} \subset \text{PromiseP/poly}$, then $\text{FBPP} \subset \text{FP/poly}$.*

Proof. We prove the first part of the theorem; the second part is an immediate consequence. Let $R \in \text{FBPP}$. For simplicity, suppose there exists a polynomial $p(|x|)$ such that for every $(x, y) \in R$, $|y| = p(|x|)$ (which can always be assumed under a suitable efficient encoding).

Let $A(x, 0^{1/\varepsilon})$ be the probabilistic algorithm for computing R with probability at least $1 - \varepsilon$ in time $\text{poly}(|x|, 1/\varepsilon)$. Fix $\varepsilon(|x|) = 4^{-p(|x|)}$.

Let Π be the following promise problem of, given an input (x, z) , to decide whether:

(YES) With probability at least $2/3 \cdot 3^{-|z|}$, the prefix of $A(x, 0^{1/\varepsilon(|x|)})$ is $z1$, or

(NO) With probability at most $1/3 \cdot 3^{-|z|}$, the prefix of $A(x, 0^{1/\varepsilon(|x|)})$ is $z1$,

promised that one of (YES) or (NO) is the case.

Observe that $\Pi \in \text{PromiseBPEXP}$: the algorithm runs $A(x, 0^{1/\varepsilon(|x|)})$ on (say) $100^{|z|}$ independent random strings, and outputs YES or NO depending on whether than more than a $1/2 \cdot 3^{-|z|}$ fraction of the strings begin with $z1$. A Chernoff bound guarantees that the algorithm is correct with high probability, so long as (x, z) is in the promise.

Next, we claim that $R \in \text{FP}^\Pi$. The algorithm for outputting $(x, y) \in R$ is as follows. Let $z_0 = \emptyset$. For each $i \in p(|x|)$, compute $z_i = z_{i-1}\Pi(x, z_{i-1})$. That is, we obtain z_i by appending 1 to z_{i-1} if the Π oracle answers YES on (x, z_{i-1}) , and by appending 0 otherwise. Finally, output $z_{p(|x|)}$.

The correctness of the algorithm follows by observing that after step i of the algorithm, $A(x, 0^{1/\varepsilon(|x|)})$ has probability at least 3^{-i} of outputting a string that starts with z_i . The proof is by induction on i : either (x, z_{i-1}) satisfies the promise, in which case $z_i = z_{i-1}\Pi(x, z_{i-1})$ appears as a prefix with probability at least $2 \cdot 3^{-i}$, or else both $z_{i-1}0$ and $z_{i-1}1$ appear as a prefix with probability at least 3^{-i} . Then, $y = z_{p(|x|)}$ must be a string that $A(x, 0^{1/\varepsilon(|x|)})$ outputs with probability at least $3^{-p(|x|)}$. But since A errs with probability at most $\varepsilon(|x|) = 4^{-p(|x|)}$, we conclude that $(x, y) \in R$. \square

5 Alternative Error Bounds

In this section, we consider some of the consequences of modifying the error bounds in the definition of FBPP.

Define the class FBPP_{\log} exactly the same way as FBPP, except that now the algorithm is required to take $\text{poly}(n, \log 1/\varepsilon)$ time rather than merely $\text{poly}(n, 1/\varepsilon)$. In other words, we mandate that the algorithm can reduce the error probability to an exponentially small quantity in polynomial time. Clearly $\text{FP} \subseteq \text{FBPP}_{\log} \subseteq \text{FBPP}$.

As we shall see here, the complexity situation for FBPP_{\log} differs dramatically from that for FBPP. First, we observe that FBPP_{\log} *cannot* be unconditionally separated from FP:

Proposition 21. *If $\text{P} = \text{NP}$, then $\text{FBPP}_{\log} = \text{FP}$.*

Proof. Let $R \in \text{FBPP}_{\log}$, and let p be a polynomial such that $|y| \leq p(n)$ for all $(x, y) \in R$ with $|x| \leq n$. Set $\varepsilon := 1/4^{p(n)}$. Then there exists a randomized algorithm A that, given $x \in \{0, 1\}^n$, outputs a $y \in \{0, 1\}^{\leq p(n)}$ such that $(x, y) \in R$, with success probability at least $1 - \varepsilon$, in $\text{poly}(n, \log 1/\varepsilon) = \text{poly}(n)$ time.

This means that, under the assumption $\text{P} = \text{NP}$ (and hence $\text{P} = \text{PH}$), in FP we can use Stockmeyer approximate counting [33] to find a $y \in \{0, 1\}^{\leq p(n)}$ such that (say) $\Pr[A(x) = y] \geq \frac{0.1}{2^{p(n)}}$, which must exist by an averaging argument. Such a y must then satisfy $(x, y) \in R$, by the assumption that A succeeds with probability at least $1 - \varepsilon$. \square

Second, we observe that the analogue of Adleman's Theorem [10] *does* hold for FBPP_{\log} :

Proposition 22. $\text{FBPP}_{\log} \subset \text{FP}/\text{poly}$.

Proof. Set $\varepsilon := 1/4^n$. Then the FBPP_{\log} machine takes $\text{poly}(n, \log 1/\varepsilon) = \text{poly}(n)$ time, and we can fix as the FP/poly advice a single randomness string that works for all inputs $x \in \{0, 1\}^n$, which must exist by the union bound. \square

Combining Proposition 22 with Theorem 6 implies that $\text{FBPP}_{\log} \neq \text{FBPP}$, unless the polynomial hierarchy collapses. To summarize, then, FBPP_{\log} behaves more like the decision class BPP than it does like FBPP.

A different choice would be to consider $\text{FBPP}_{\text{negl}}$, which we define as the class of all polynomially-bounded relations $R \subseteq \{0, 1\}^* \times \{0, 1\}^*$ for which there exists a polynomial-time randomized algorithm that, given x , outputs a y such that $(x, y) \in R$ (whenever one exists) with success probability at least $1 - \varepsilon(n)$, for some negligible function $\varepsilon(n) = \frac{1}{n^{\omega(1)}}$.

We have the following unconditional result, which supersedes the analogues of Theorems 6, 8, and 9:

Theorem 23. $\text{FBPP}_{\text{negl}} \not\subseteq \text{FP/poly}$.

Proof. The relation R that witnesses the separation is the same one from the proof of Theorem 9, involving a “bad output set” $S_x \subset \{0, 1\}^n$ chosen uniformly at random for each $x \in \{0, 1\}^n$ subject to $|S_x| = 2^{n/2}$. We already showed in Theorem 9 that $R \notin \text{FP/poly}$. For $R \in \text{FBPP}_{\text{negl}}$, the algorithm is just to output an n -bit string uniformly at random, independent of x . \square

What is the relationship between $\text{FBPP}_{\text{negl}}$ and FBPP? In one direction we have the following:

Proposition 24. $\text{FBPP}_{\text{negl}} \not\subseteq \text{FBPP}$.

Proof. Consider the relation

$$R = \{(x, y) : |x| = |y|, K(y) \geq \frac{|y|}{2}\},$$

where K is Kolmogorov complexity. We have $R \in \text{FBPP}_{\text{negl}}$ by simply outputting a random string of length $n = |x|$. On the other hand, if R had an FBPP algorithm, then by setting $\varepsilon := 1/4^n$ and then searching for the lexicographically first string that the algorithm output with probability at least (say) $\frac{0.1}{2^n}$, we could deterministically compute an n -bit string such that $K(x) \geq n/2$, which is impossible. \square

In the other direction, we leave open whether $\text{FBPP} \subset \text{FBPP}_{\text{negl}}$ or whether the two classes are incomparable. Clearly we do have $\text{FBPP}_{\log} \subset \text{FBPP}_{\text{negl}}$, by setting (say) $\varepsilon = 1/2^n$.

6 Sampling Problems

We now show that the analogue of Theorem 10 is *false* for sampling problems:

Theorem 17. $\text{SampBPP/poly} \neq \text{SampBPP/rpoly}$ and $\text{SampBQP/poly} \neq \text{SampBQP/rpoly}$.

Proof. Like in Theorem 10, we prove the separation involving FBPP, as the separation involving FBQP is completely analogous. It suffices to choose some family of nonempty sets $S_n \subset \{0, 1\}^n$, one for each n . Then consider the problem of outputting a uniformly random element of S_n on input $x \in \{0, 1\}^n$. This problem is clearly in SampBPP/rpoly , since we can take the randomized advice itself to be a uniformly random element of S_n . But if the S_n 's are chosen uniformly at random, then a counting argument shows that the problem has probability 0 of being in SampBPP/poly . \square

Lastly, we observe that Theorem 5 gives rise to a separation of sampling classes with randomized and quantum advice:

Theorem 18. $\text{SampBQP/rpoly} \neq \text{SampBQP/qpoly}$.

Proof. Consider the problem of sampling from the output distribution of the errorless algorithm that solves the relation problem R_F in the proof of Theorem 5. This sampling problem is in SampBQP/qpoly . On the other hand, if this sampling problem were in SampBQP/rpoly , this would imply $R_F \in \text{FBQP/rpoly}$, because sampling from the distribution within total variation distance ε would solve R_F with probability $1 - \varepsilon$. But this would violate Theorem 5. \square

To summarize, then, the classes SampBQP/poly , SampBQP/rpoly , and SampBQP/qpoly are *all* distinct.

7 Open Problems

We gave an example of a relational problem R_F in FBQP/qpoly but not in FBQP/rpoly : indeed, one that is easy to solve using n qubits of quantum advice, but requires $\Omega(2^{n/2})$ bits of classical randomized advice. While this separation is tight for R_F itself, is there a different relational problem, solvable with n qubits of advice, for which the lower bound on randomized advice reaches its maximum of $\Omega(2^n)$?

We note that Montanaro [26] showed in 2019 that $\Omega(2^n)$ bits of classical advice are needed to perform certain sampling tasks, for which n qubits of quantum advice suffice. In the other direction, Gosset and Smolin [20] have shown that, for decision and promise problems solvable with n qubits of quantum advice, $O(2^{n/2})$ bits of classical randomized advice suffice. The case of relational problems remains open. We conjecture that the answer is closer to 2^n than $2^{n/2}$.

As we discussed in Section 3, the gap between 2^n and $2^{n/2}$ would be *extremely* useful to close, as a prerequisite to any possible quantum information supremacy experiment based on the $\text{FBQP/rpoly} \neq \text{FBQP/qpoly}$ separation. Other complexity results that would bear directly on such an experiment include:

1. a lower bound on the amount of classical randomized advice needed to simulate *noisy* quantum advice for some relational problem—say, as a function of the fidelity δ between the actual advice state and a desired pure state;
2. a quantitative refinement of $\text{FBQP/rpoly} \neq \text{FBQP/qpoly}$ that took into account the circuit complexity of *preparing* the n -qubit quantum advice state (which, in practice, would likely have to be much less than 2^n); and
3. a reproof of $\text{FBQP/rpoly} \neq \text{FBQP/qpoly}$ in which the circuit to measure the quantum advice state was made *as simple as possible*—could it even measure each of the n qubits independently from the rest?

Moving on, is there any sense in which FBQP/qpoly contains “more” problems than FBQP/rpoly —i.e., can the classes be separated by counting the number of problems in each?

Can we separate FBQP/qpoly from FBQP/rpoly via an “explicit” problem, rather than relying on the probabilistic method? More concretely: can we show that, under some plausible hardness assumption, there is a relation in (say) $\text{FBQP/qpoly} \cap \text{FPSPACE}$ that is not in FBQP/rpoly ? Such a result would create more “symmetry” between that separation and our $\text{FP/poly} \neq \text{FP/rpoly}$ separation. For the latter, Theorem 6 gave an explicit relational problem that plausibly realizes

the nonconstructively proven separation: namely, the problem of outputting a string of high time-bounded Kolmogorov complexity, conditional on an input string x .

Would $\text{FBPP} \subset \text{FP/poly}$ have even stronger consequences than $\text{PSPACE} \subset \text{NP/poly}$, such as $\text{PSPACE} \subset \text{P/poly}$ or even $\text{EXP} \subset \text{P/poly}$? Also, is there a *relativizing* proof that $\text{FBPP} \subset \text{FP/poly}$ would have unlikely consequences?

Does $\text{FBPP}_{\log} = \text{FP}$ under some plausible derandomization assumption? Is $\text{FBPP} \subset \text{FBPP}_{\text{negl}}$? (Recall that FBPP_{\log} is the subclass of FBPP where we require $\text{poly}(n, \log 1/\varepsilon)$ time to achieve error ε , while $\text{FBPP}_{\text{negl}}$ is the variant where we require only that the error probability be negligible.)

Are there examples of problems in FBPP , FBQP , FP/rpoly , or the other relational classes studied in this paper, for which $\text{poly}(1/\varepsilon)$ (rather than, say, $\text{polylog}(1/\varepsilon)$) running time is actually needed to achieve error ε ? (Certainly we've given examples of *reductions* where such running time is needed.) For example, can we show that $\text{FBPP}_{\log} \neq \text{FBPP}$ *unconditionally*, without assuming noncollapse of the polynomial hierarchy?

8 Acknowledgments

We thank Daochen Wang and Alexandru Cojocaru for conversations that suggested the FBQP/poly versus FBQP/qpoly problem, and Lance Fortnow, Iordanis Kerenidis, Ashley Montanaro, Ronald de Wolf, and David Zuckerman for helpful discussions. We thank anonymous referees for constructive feedback.

References

- [1] S. Aaronson. Limitations of quantum advice and one-way communication. *Theory of Computing*, 1:1–28, 2005. Earlier version in CCC'2004. quant-ph/0402095. [pp. 4, 8, 13]
- [2] S. Aaronson. Quantum computing, postselection, and probabilistic polynomial-time. *Proc. Roy. Soc. London*, A461(2063):3473–3482, 2005. quant-ph/0412187. [p. 4]
- [3] S. Aaronson. BQP and the polynomial hierarchy. In *Proc. ACM STOC*, 2010. arXiv:0910.4698. [p. 4]
- [4] S. Aaronson. The equivalence of sampling and searching. In *Proc. Computer Science Symposium in Russia (CSR)*, 2011. arXiv:1009.5104, ECCC TR10-128. [pp. 2, 9, 10]
- [5] S. Aaronson and A. Ambainis. The need for structure in quantum speedups. In *Proc. Innovations in Theoretical Computer Science (ITCS)*, 2011. arXiv:0911.0996. [p. 4]
- [6] S. Aaronson and A. Arkhipov. The computational complexity of linear optics. *Theory of Computing*, 9(4):143–252, 2013. Earlier version in STOC'2011. ECCC TR10-170, arXiv:1011.3245. [pp. 2, 4]
- [7] S. Aaronson and L. Chen. Complexity-theoretic foundations of quantum supremacy experiments. In *Proc. Conference on Computational Complexity*, 2017. arXiv:1612.05903. [pp. 2, 4]
- [8] S. Aaronson and S. Gunn. On the classical hardness of spoofing linear cross-entropy benchmarking. *Theory of Computing*, 16:1–8, 2020. arXiv:1910.12085. [p. 15]

- [9] S. Aaronson and G. Kuperberg. Quantum versus classical proofs and advice. *Theory of Computing*, 3(7):129–157, 2007. Earlier version in CCC’2007. arXiv:quant-ph/0604056. [p. 4]
- [10] L. Adleman. Two theorems on random polynomial time. In *Proc. IEEE FOCS*, pages 75–83, 1978. [pp. 3, 5, 20]
- [11] D. Aharonov, X. Gao, Z. Landau, Y. Liu, and U. Vazirani. A polynomial-time classical algorithm for noisy random circuit sampling. arXiv:2211.03999, 2022. [p. 15]
- [12] Z. Bar-Yossef, T. S. Jayram, and I. Kerenidis. Exponential separation of quantum and classical one-way communication complexity. *SIAM J. Comput.*, 38(1):366–384, 2008. Earlier version in STOC’2004. ECCC TR04-036. [pp. 6, 8, 14]
- [13] E. Berlekamp. Factoring polynomials over large finite fields. *Math. Comp.*, 24:713–735, 1970. [p. 19]
- [14] H. Buhrman and L. Torenvliet. Randomness is hard. *SIAM J. Comput.*, 30(5):1485–1501, 2000. Earlier version in Proceedings of CCC’1998. [pp. 9, 17]
- [15] F. Arute et al. Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779):505–510, 2019. [pp. 2, 15]
- [16] B. Fefferman and S. Kimmel. Quantum vs. classical proofs and subset verification. In *Mathematical Foundations of Computer Science*, pages 1–23, 2018. arXiv:1510.06750. [p. 4]
- [17] X. Gao, M. Kalinowski, C. N. Chou, M. D. Lukin, B. Barak, and S. Choi. Limitations of linear cross-entropy as a measure for quantum advantage. arXiv:2112.01657, 2021. [p. 15]
- [18] D. Gavinsky, J. Kempe, I. Kerenidis, R. Raz, and R. de Wolf. Exponential separation for one-way quantum communication complexity, with applications to cryptography. *SIAM J. Comput.*, 38(5):1695–1708, 2008. Earlier version in STOC’2007. quant-ph/0611209. [p. 8]
- [19] O. Goldreich. In a world of $P = BPP$. *Studies in Complexity and Cryptography*, 6650:191–232, 2011. ECCC TR10-135. [pp. 2, 9]
- [20] D. Gosset and J. A. Smolin. A compressed classical description of quantum states. In *Proc. Theory of Quantum Computation, Communication, and Cryptography (TQC)*, pages 1–9, 2019. arXiv:1801.05721. [p. 22]
- [21] A. S. Holevo. Some estimates of the information transmitted by quantum communication channels. *Problems of Information Transmission*, 9:177–183, 1973. English translation. [p. 3]
- [22] R. Ilango, J. Li, and R. Williams. Indistinguishability obfuscation, range avoidance, and bounded arithmetic. In *Proc. ACM STOC*, pages 1076–1089, 2023. ECCC TR23-038. [pp. 2, 5]
- [23] R. M. Karp and R. J. Lipton. Turing machines that take advice. *Enseign. Math.*, 28:191–201, 1982. Earlier version in STOC’1980, pp. 302-309. [p. 3]
- [24] N. Kumar, I. Kerenidis, and E. Diamanti. Experimental demonstration of quantum advantage for one-way communication complexity surpassing best-known classical protocol. *Nature Communications*, 10(1):1–10, 2019. arXiv:1811.09154. [pp. 14, 16]

- [25] Ryan L. Mann. *Quantum computation and combinatorial structures*. PhD thesis, University of Technology Sydney, February 2019. [p. 2]
- [26] A. Montanaro. Quantum states cannot be transmitted efficiently classically. *Quantum*, 3(154), 2019. arXiv:1612.06546. [p. 22]
- [27] Tomoyuki Morimae and Takashi Yamakawa. Quantum advantage from one-way functions. arXiv:2302.04749, 2023. [p. 2]
- [28] A. Natarajan and C. Nirkhe. A classical oracle separation between QMA and QCMA. arXiv:2210.15380, 2022. [p. 4]
- [29] H. Nishimura and T. Yamakami. Polynomial time quantum computation with advice. *Inform. Proc. Lett.*, 90:195–204, 2003. ECCC TR03-059, quant-ph/0305100. [p. 3]
- [30] OpenAI. GPT-4 technical report. arXiv:2303.08774, 2023. [pp. 2, 25]
- [31] F. Pan, K. Chen, and P. Zhang. Solving the sampling problem of the Sycamore quantum circuits. *Phys. Rev. Lett.*, 129(9), 2022. arXiv:2111.03011. [p. 15]
- [32] A. Shamir. IP=PSPACE. *J. of the ACM*, 39(4):869–877, 1992. Earlier version in FOCS’1990, pp. 11-15. [pp. 9, 18]
- [33] L. J. Stockmeyer. The complexity of approximate counting. In *Proc. ACM STOC*, pages 118–126, 1983. [p. 20]
- [34] T. Yamakawa and M. Zhandry. Verifiable quantum advantage without structure. In *Proc. IEEE FOCS*, pages 69–74, 2022. arXiv:2204.02063. [p. 4]
- [35] A. C-C. Yao. Probabilistic computations: Toward a unified measure of complexity. In *Proc. IEEE FOCS*, pages 222–227, 1977. [p. 13]
- [36] H.-S. Zhong, H. Wang, Y.-H. Deng, M.-C. Chen, L.-C. Peng, Y.-H. Luo, J. Qin, D. Wu, X. Ding, Y. Hu, P. Hu, X.-Y. Yang, W.-J. Zhang, H. Li, Y. Li, X. Jiang, L. Gan, G. Yang, L. You, Z. Wang, L. Li, N.-L. Liu, C.-Y. Lu, and J.-W. Pan. Quantum computational advantage using photons. *Science*, 370(6523):1460–1463, 2020. arXiv:2012.01625. [pp. 2, 15]

A Conversation with GPT-4

Below is the transcript of a conversation, dated January 9, 2023, between one of the authors and a preliminary version of GPT-4 [30] about possible definitions for the class FBPP. For ease of reading, the conversation has been $\text{T}_{\text{E}}\text{X}$ ified (e.g. “FBPP” \rightarrow FBPP, “eps” \rightarrow ε). With mild prodding, GPT eventually settles on Definition 1.

SCOTT : What’s a good definition of FBPP – the generalization of BPP to relational problems?

GPT : FBPP could be defined as follows: FBPP (Functional BPP) is a complexity class for relational problems, defined analogously to BPP for decision problems. A relational problem is in FBPP if there is a probabilistic polynomial-time algorithm that, given an input, produces an output consistent with the relation with probability at least $2/3$ (or any other constant bounded away from $1/2$).

SCOTT : But unlike with BPP, amplification isn't possible, so the choice of constant (such as $2/3$) actually matters, right?

GPT : The choice of constant does matter, but amplification is still possible in many cases. Just like with BPP, the algorithm could be run multiple times to increase the confidence in the output. However, the amplification process is more complicated for FBPP, since it's not always clear how to combine the outputs from multiple runs of the algorithm to produce a single, more accurate output. This will depend on the specifics of the relational problem at hand. But in many cases, amplification should still be possible.

SCOTT : But in other cases not! Thus, can you suggest a definition of FBPP that would allow arbitrary error probabilities?

GPT : One possible alternative definition could be: FBPP (Functional BPP) is a complexity class for relational problems. A relational problem is in FBPP if there is a probabilistic polynomial-time algorithm that, given an input, produces an output consistent with the relation with probability strictly greater than $1/2$.

SCOTT : Or what about with probability $1 - \varepsilon$ for any specified $\varepsilon > 0$?

GPT : Yes, that would work as well. The definition could be: FBPP (Functional BPP) is a complexity class for relational problems. A relational problem is in FBPP if there is a probabilistic polynomial-time algorithm that, given an input, produces an output consistent with the relation with probability at least $1 - \varepsilon$ for any specified $\varepsilon > 0$.

SCOTT : And how is the running time of the algorithm allowed to depend on ε ?

GPT : The definition above doesn't specify, but one reasonable choice would be to require that the running time be polynomial in both the input size and $1/\varepsilon$. This would ensure that the algorithm is efficient even as the error probability is made arbitrarily small.