

Founding Quantum Cryptography on Quantum Advantage

or, Towards Cryptography from #P-Hardness

Dakshita Khurana*

Kabir Tomer[†]

Abstract

Recent oracle separations [Kretschmer, TQC'21, Kretschmer et. al., STOC'23] have raised the tantalizing possibility of building quantum cryptography from sources of hardness that persist even if the polynomial hierarchy collapses. We realize this possibility by building quantum bit commitments and secure computation from *unrelativized*, well-studied mathematical problems that are conjectured to be hard for $P^{\#P}$ – such as approximating the permanents of complex Gaussian matrices, or approximating the output probabilities of random quantum circuits. Indeed, we show that as long as *any one of the conjectures* underlying sampling-based quantum advantage (e.g., BosonSampling [Aaronson-Arkhipov, STOC'11], Random Circuit Sampling [Boixo et. al., Nature Physics 2018], IQP [Bremner, Jozsa and Shepherd, Proc. Royal Society of London 2010]) is true, quantum cryptography can be based on the extremely mild assumption that $P^{\#P} \not\subseteq (io)BQP/qpoly$.

Our techniques uncover strong connections between the hardness of approximating the probabilities of outcomes of quantum processes, the existence of “one-way” state synthesis problems, and the existence of useful cryptographic primitives such as one-way puzzles and quantum bit commitments. Specifically, we prove that the following hardness assumptions are equivalent under BQP reductions.

- The **hardness of approximating the probabilities** of outcomes of certain efficiently sampleable distributions. That is, there exist quantumly efficiently sampleable distributions for which it is hard to approximate the probability assigned to a randomly chosen string in the support of the distribution (upto inverse polynomial relative error).
- The existence of **one-way puzzles**, where a quantum sampler outputs a pair of classical strings – a puzzle and its key – and where the hardness lies in finding the key corresponding to a random puzzle. These are known to imply quantum bit commitments [Khurana and Tomer, STOC'24].
- The existence of **state puzzles**, or one-way state synthesis, where it is hard to synthesize a secret quantum state given a public classical identifier. These capture the hardness of search problems with quantum secrets and classical challenges.

These are the first constructions of quantum cryptographic primitives (one-way puzzles, quantum bit commitments, state puzzles) from well-studied mathematical assumptions that do not imply the existence of classical cryptography.

Along the way, we also show that distributions that admit efficient quantum samplers but cannot be pseudo-deterministically efficiently sampled imply quantum commitments.

*UIUC and NTT Research. Email: dakshita@illinois.edu

[†]UIUC. Email: ktomer2@illinois.edu

Contents

1	Introduction	1
1.1	Our Results	2
1.1.1	One-Way Puzzles and Sampling-Based Quantum Advantage	3
1.1.2	One-way Puzzles and the Hardness of (Pseudo)-Deterministic Sampling	6
1.1.3	One-way Puzzles and the Hardness of State Synthesis	6
1.2	Perspective	7
1.3	Open Problems	9
2	Related Work	9
3	Technical Overview	11
3.1	One-way Puzzles from the Hardness of Approximating Probabilities	12
3.2	One-way Puzzles from the Hardness of Pseudo-Deterministic Sampling	14
3.3	One-Way Puzzles from Hard State Synthesis Problems	15
4	Preliminaries	19
4.1	Notation and Conventions	19
4.2	Quantum Cryptographic Primitives	20
5	Hardness of Approximating Probabilities implies One-way Puzzles	21
5.1	Definitions	21
5.2	Uniform Approximation Hardness Implies Native Approximation Hardness	23
5.3	Native Approximation Hardness Implies One-Way Puzzles	27
5.4	One-Way Puzzles Imply Native Approximation Hardness	32
6	The Hardness of Pseudo-Deterministic Sampling implies One-Way Puzzles	38
7	State Puzzles are Equivalent to One-Way Puzzles	42
	References	67
A	Instantiating Uniform Approximation Hardness (Definition 5.1)	72
A.1	Random Circuit Sampling	72
A.2	Boson Sampling	73
A.3	IQP Circuit Sampling	74
B	Flatness of Unitary 2-Design Output Distributions	74
C	Proof of Theorem 7.3	76
D	Quantum Advantage from One-Way Puzzles in Microcrypt	77

1 Introduction

Nearly all of modern classical cryptography relies on unproven computational hardness. Decades of studying cryptosystems based on various concrete mathematical problems led to the emergence of complexity-based cryptography. Modern cryptography allows us to categorize the hardness offered by mathematical problems into *generic* cryptographic primitives, studying how abstract primitives are reducible to one another, and setting aside which concrete implementation of the generic primitive is used. For example, a one-way function is a classically efficiently computable function that is hard to invert. Concrete candidates for one-way functions are known based on a variety of algebraic problems such as the hardness of discrete logarithms, quadratic residuosity or learning with errors. The existence of one-way functions is a fundamental hardness assumption, necessary for the existence of modern classical cryptography [LR86, IL89, ILL89].

Hardness in Quantum Cryptography. At the same time, the breakthrough ideas of Weisner [Wie83], Bennett and Brassard [BB84] demonstrated the possibility of quantum cryptography – specifically key agreement – without the need for unproven assumptions, and based solely on the nature of quantum information. Unfortunately, it was soon discovered [LC97, May97] that other fundamental quantum cryptographic primitives, including bit commitments and secure computation, *require* some form of computational hardness. It is now known [GLSV21, BCKM21] that (post-quantum) *one-way* functions suffice to build bit commitments and quantum secure computation. However, it is plausible that one-way functions are not *necessary* for quantum cryptography, and that sources of hardness even milder than the existence of one-way functions could suffice. Can we precisely quantify the amount of hardness that is necessary for quantum cryptography?

Recent works have attempted to address this problem by introducing quantum relaxations of generic classical primitives, and building quantum cryptography from these relaxations. Some examples are pseudorandom quantum states [JLS18] and one-way state generators/one-way puzzles [MY22b, KT24], that have been introduced as quantum analogues of pseudorandom generators and one-way functions respectively. All of these primitives are cryptographically “useful” in that they imply quantum bit commitments and secure computation [AQY22, MY22b, KT24, BJ24]. Furthermore, *relative to appropriately chosen oracles*, these primitives are weaker than their classical counterparts – more precisely, there exist oracles relative to which secure instantiations of one-way states/pseudorandom states exist, even when $BQP = QMA$ (respectively, $P = NP$) [Kre21, KQST23]. So while classical cryptography would completely break down if $P = NP$, there is the exciting possibility that quantum cryptography would continue to exist!

Beyond Oracle Worlds. At this point it is natural to ask if there are *any* concrete candidates for quantum cryptosystems in the *real* (unrelativized) world, based on mathematical assumptions that do not imply classical cryptography. Unfortunately, the answer so far has been a resounding no.

Let us explain what we’re looking for in more detail. The gold standard in modern cryptography is to base security on mathematical conjectures whose statement is scientifically interesting independently of the cryptographic application itself¹. We want to avoid assuming that a suggested scheme itself is secure; since such assumptions are construction-dependent and the proclaimed guarantee of provable security essentially loses its meaning (we refer the reader to a survey by Goldwasser and Kalai [GK16] for a more thought-provoking analysis). Underlying modern clas-

¹For instance, random circuits have been conjectured to output pseudorandom states [AQY22]— but this essentially assumes that a given construction is secure, and to our knowledge, has not been cryptanalyzed or theoretically studied. It is also known to not be true for certain architectures (e.g., BosonSampling outputs are not pseudorandom [AA14]), and it is unclear why this conjecture would be separated from the existence of one-way functions.

sical cryptography is a bedrock of concrete mathematical problems that have been introduced and cryptanalyzed extensively and often independently of their cryptographic application.

Going back to the state of affairs in quantum cryptography: so far, all provably secure constructions of quantum cryptosystems, from well-studied assumptions, rely on the existence of one-way functions. At the same time, building on all the excitement about the possibility of cryptography without one-way functions, there is a large body of work conceptualizing quantum generalizations of classical cryptographic primitives, and reducing them to one another. Besides quantum commitments, secure computation, pseudorandom and one-way states discussed above, other examples include quantum public-key and private-key encryption, signature schemes, etc.² In fact, there is even a name for a world in which one-way functions do not exist and yet secure quantum cryptography is possible: *Microcrypt*.

However, in the absence of real instantiations, a sensible objection to this body of work is that we may just be building fictional castles in the air – with no hope of securely realizing these cryptosystems without assuming the existence of one-way functions. Our work fundamentally refutes this objection.

We provide the first *unrelativized* constructions of quantum bit commitments and secure computation from well-studied mathematical assumptions that do not imply classical cryptography. This addresses a major unresolved problem in the area, namely offering unrelativized evidence that the hardness required for quantum cryptography is weaker than that required for classical cryptography.

Hardness beyond the Polynomial Hierarchy. Recall that relative to certain oracles, quantum bit commitments exist even if $P = NP$, i.e., even if all languages in the polynomial hierarchy (PH) can be efficiently decided. This indicates that we may be able to build quantum cryptography from problems that lie outside the polynomial hierarchy, and plausibly continue to be hard even if PH collapses. One natural complexity class that is believed to not be contained in PH is $\#P$, known to contain problems such as finding the permanent of a given real or complex valued matrix. Besides being $\#P$ -hard to compute, permanents also have worst-case to average-case reductions, opening up the splendid possibility of basing quantum cryptography directly on worst-case hardness.

But actually building cryptosystems from the hardness of computing permanents turns out to be tricky due to the following conceptual barrier. In most natural constructions of (quantum) cryptosystems, honest parties need to be able to sample hard problems together with their solutions – whereas it appears unlikely that random matrices can be quantumly efficiently sampled together with their permanents. By exploiting connections with sampling-based quantum advantage and using some indirection, we overcome this barrier. This allows us to obtain quantum cryptography from the average-case hardness of *approximating* permanents of complex-valued matrices, or the hardness of approximating probabilities of outcomes of random circuits.

1.1 Our Results

Our first conceptual contribution is a connection between *sampling-based quantum advantage* and a quantum cryptographic primitive called a *one-way puzzle*. One-way puzzles are notable for implying quantum commitments [KT24], which in turn imply quantum secure computation [BCKM21, GLSV21, AQY22].

A one-way puzzle [KT24] is a simple, cryptographically useful primitive with classical outputs – analogous to a randomized variant of a one-way function. It consists of a quantum polynomial

²We refer the reader to the graph at <https://sattath.github.io/qcrypto-graph/> for several additional examples.

time algorithm Samp and a Boolean function Ver . Samp outputs a pair of classical strings – a puzzle and key (s, k) – satisfying $\text{Ver}(s, k) = 1$. The security guarantee is that given a “puzzle” s , it is (quantum) computationally infeasible to find a key k such that $\text{Ver}(s, k) = 1$, except with negligible probability.

Through a connection with quantum advantage, our first result yields multiple unrelativized instantiations of one-way puzzles and quantum commitments (and therefore also secure computation) without one-way functions, based on well-studied mathematical problems that are conjectured to be $\text{P}^{\#\text{P}}$ -hard. We discuss these in the following subsection, where we begin by briefly describing what we mean by quantum advantage.

1.1.1 One-Way Puzzles and Sampling-Based Quantum Advantage

More than two decades ago, it was observed [TD02] that quantum computers can sample from distributions that likely cannot be reproduced on any classical device. Subsequent works have solidified complexity-theoretic evidence that the output distributions of a variety of quantum computations – including several types of non-universal computations – may be computationally intractable to simulate on a classical device (see, e.g., [SB09, BJS11, AA11, BMS16, FM17, BIS⁺18, BFNV19, KMM21, BFLL21, Kro22, Mov23, ZVBL23] and a recent survey [HE23]). We show that any of these quantum computations also yield quantum cryptography without one-way functions, under the same complexity conjectures that have been studied in context of quantum advantage and the mild additional assumption that $\text{P}^{\#\text{P}} \not\subseteq (\text{io})\text{BQP}/\text{qpoly}$.

At first, building quantum cryptography only from quantum advantage may appear surprising or even unlikely – quantum advantage is all about tasks that are hard for *classical* computers, whereas quantum cryptography asks for hardness against *quantum* computers. Perhaps one may be able to obtain (only) *classically secure* one-way puzzles from advantage, but how would one possibly obtain hardness against quantum machines from sampling tasks that are hard only for classical machines?

To understand why, we will peel back the layers a little bit. A key idea in sampling-based quantum advantage (originating in [AA11]) is to relate the hardness of classical sampling to the hardness of computing the probabilities of outcomes of quantumly efficiently sampled distributions. Specifically, major existing proposals for sampling-based quantum advantage assume that the following average-case problem is $\#\text{P}$ -hard:

Given the description of a quantum sampler, approximate the probability assigned by the sampler’s output distribution to a uniformly chosen string in its support.

Here, the approximation is required to have inverse polynomially small relative error. In addition, outputs of the sampler are assumed to satisfy a natural *anti-concentration* property – requiring that not all of the hardness of approximation should lie on points that have extremely tiny (e.g., doubly exponentially small) probability mass. We capture this with the following assumption that admits instantiations from many concrete conjectures corresponding to different frameworks for quantum advantage such as BosonSampling, universal random circuits, IQP circuit sampling, etc.

Assumption 1 ($\#\text{P}$ -Hardness of Approximating Probabilities). *There is a family of (uniform) efficiently sampleable distributions $\mathcal{C} = \{\mathcal{C}_n\}_{n \in \mathbb{N}}$ over quantum circuits C where each C has n -qubit outputs (and potentially additional junk qubits), such that there exist polynomials $p(\cdot)$ and $\gamma(\cdot)$ satisfying:*

1. **Anticoncentration.** *For all large enough $n \in \mathbb{N}$*

$$\Pr_{\substack{C \leftarrow \mathcal{C}_n \\ x \leftarrow \{0,1\}^n}} \left[\Pr_C[x] \geq \frac{1}{p(n) \cdot 2^n} \right] \geq \frac{1}{\gamma(n)}$$

2. **Hardness of Approximating Probabilities.** For any oracle \mathcal{O} satisfying that for all large enough $n \in \mathbb{N}$,

$$\Pr_{\substack{C \leftarrow \mathcal{C}_n \\ x \leftarrow \{0,1\}^n}} \left[|\mathcal{O}(C, x) - \Pr_C[x]| \leq \frac{\Pr_C[x]}{p(n)} \right] \geq \frac{1}{\gamma(n)} - \frac{1}{p(n)}$$

we have that $\text{P}^{\#\text{P}} \subseteq \text{BPP}^{\mathcal{O}}$.³ Here, $\Pr_C[x]$ denotes the probability of obtaining outcome x when the output register of $C|0\rangle$ is measured in the standard basis.

As described above, different models of quantum advantage are based on the conjectured $\#\text{P}$ -hardness of computing output probabilities for different circuit families; all of these imply Assumption 1. Some concrete, well-studied candidates include:

- **Random Circuit Sampling.** One of the leading candidates for quantum advantage today assumes the hardness of approximating output probabilities of random (universal) quantum circuits (see e.g. [BIS⁺18, BFN19]). Here, the distribution \mathcal{C} corresponds to a circuit with gates drawn from some universal set and according to some specified architecture, and quantum advantage follows as long as Assumption 1 holds for \mathcal{C} . Understanding the underlying architectures and attempting to prove/disprove the corresponding conjecture is now the focus of a large, and quickly growing, body of work (see e.g. [BIS⁺18, BFN19, KMM21, BFLL21, Kro22, Mov23, ZVBL23]).
- **BosonSampling.** Aaronson and Arkhipov [AA11] related the task of finding probabilities of outcomes of a BosonSampling experiment to computing the permanents of appropriate random matrices. They formulated the following two conjectures that together, form the complexity-theoretic basis for advantage based on BosonSampling. The *Permanent of Gaussians Conjecture (PGC)* posits that it is $\#\text{P}$ -hard to approximate the permanent of a matrix of independent random $\mathcal{N}(0, 1)$ Gaussian entries, and the *Permanent Anti-Concentration Conjecture (PACC)* says that with high probability over a matrix A sampled randomly as above, $|\text{Per}(A)| \geq \sqrt{n!}/\text{poly}(n)$. The PGC and PACC imply Assumption 1.
- **IQP, and beyond.** Other non-universal models of quantum computation, such as Instantaneous Quantum Polynomial (IQP) and Deterministic Quantum Computation with one quantum bit (DQC1) are also candidates for advantage due to their potential ease of implementation on near-term quantum devices [SB09, KL98, MFF14]. Again, quantum advantage assumes the hardness of approximating probabilities when \mathcal{C} corresponds to these types of circuits, along with anti-concentration of circuit outputs, which implies Assumption 1.

One way to state our main theorem is:

Theorem 1.1 (Informal). *Suppose Assumption 1 is true. Then, one-way puzzles (which imply quantum commitments) exist if and only if $\text{P}^{\#\text{P}} \not\subseteq (\text{io})\text{BQP}/\text{qpoly}$.*

We also provide a slightly different statement of our main theorem. Assumption 1, together with the (extremely mild) assumption that $\text{P}^{\#\text{P}} \not\subseteq \text{ioBQP}/\text{qpoly}$ ⁴ implies the following assumption, which is a reformulation of Assumption 1 to (1) require the adversary to succeed only on

³Our theorem statements remain unchanged if the hardness reduction is a $\text{BQP}^{\mathcal{O}}$ (instead of a $\text{BPP}^{\mathcal{O}}$) machine. In fact, we could even allow the hardness reduction to be a $\text{BPP}^{\text{NP}^{\mathcal{O}}}$ machine – all this would change for our results is that the worst-case assumption that we make, which currently says $\text{P}^{\#\text{P}} \not\subseteq \text{BQP}$ would change to $\text{P}^{\#\text{P}} \not\subseteq \text{BPP}^{\text{NP}^{\text{BQP}}}$.

⁴We note that $\text{P}^{\#\text{P}} = \text{P}^{\text{PP}}$ and $\text{PP} \subseteq \text{BQP}/\text{qpoly}$ implies a collapse of the counting hierarchy to QMA [Aar06, Yir24].

infinitely many n instead of all large enough n , (2) only ask that the probabilities be hard to approximate for non-uniform QPT machines (instead of requiring the approximation to be #P-hard), and (3) remove the anti-concentration requirement (i.e., $\Pr_{C_n}[x] \geq \frac{1}{p(n)2^n}$) while instead, sampling challenge instances x according to the output distribution of C_n . We state the assumption below.

Assumption 2 (Native Approximation Hardness). *There exists a family of (uniform) efficiently samplable distributions $\mathcal{D} = \{\mathcal{D}_n\}_{n \in \mathbb{N}}$ over classical strings such that there exists a polynomial $p(\cdot)$ such that for all QPT $\mathcal{A} = \{\mathcal{A}_n\}_{n \in \mathbb{N}}$, every (non-uniform, quantum) advice ensemble $|\tau\rangle = \{|\tau_n\rangle\}_{n \in \mathbb{N}}$, and large enough $n \in \mathbb{N}$,*

$$\Pr_{x \leftarrow \mathcal{D}_n} \left[\left(|\mathcal{A}(|\tau\rangle, x) - \Pr_{\mathcal{D}_n}[x]| > \frac{\Pr_{\mathcal{D}_n}[x]}{p(n)} \right) \right] \geq \frac{1}{p(n)}$$

Our main theorem can be restated as:

Theorem 1.2. (Informal) *Assumption 2 is equivalent to the existence of one-way puzzles, and implies the existence of quantum bit commitments.*

One may wonder whether Assumptions 1 and 2 really are *weaker* than the existence of (post-quantum) one-way functions. First, our equivalence in Theorem 1.2 also shows that Assumption 2 is implied by the existence of one-way functions (which trivially imply puzzles).

Next, the the hardness of approximating probabilities in Assumption 1 is based on conjectured hardness beyond the polynomial hierarchy. The resulting one-way puzzles and commitments therefore lie squarely in Microcrypt.

Additionally, suppose that Assumption 1 is stated as a problem: “for a particular quantumly samplable distribution (say RCS/BosonSampling), approximate the probabilities of outputs upto low relative error in the average case”. Do we expect the hardness of this problem to imply the existence of one-way functions? Suppose there existed a one-way function f and a BPP reduction that used an inverter for f as a black-box to approximate probabilities for the distribution in Assumption 1. This would imply that these probabilities can be approximated in BPP^{NP} . However, all existing strategies to *prove* that these distributions exhibit sampling based advantage (e.g., [AA11]) rely on these probabilities being *hard to approximate, on average* by a BPP^{NP} machine. Thus, conjectures made in context of quantum advantage already imply that these assumptions will not yield one-way functions (under BPP reductions)⁵.

Furthermore, our results are tight – quantum advantage is *necessary* for one-way puzzles in Microcrypt. Namely, for any one-way puzzle that does not also imply a one-way function, the sampler Q outputting puzzle and solution pairs cannot be (uniformly) classically simulated. Otherwise, a classical machine C that samples from a distribution that is close in statistical distance to the output of Q (for all large enough input lengths) will yield a one-way function as follows: on input x , return the puzzle output by $C(x)$. We formally prove that this yields a one-way function in Appendix D. This also means that any improvements to our results, e.g., basing one-way puzzles *only* on #P-hardness (without the need for additional conjectures) would yield sampling-based quantum advantage from only minimal worst-case complexity assumptions such as the non-collapse of PH, and without the need for unproven conjectures.

Finally, a powerful consequence of these results is that they serve as a tool to reduce the existence of other cryptosystems that become insecure in the presence of #P oracles, to the existence of quantum commitments, as we will demonstrate with state synthesis hardness in Section 1.1.3.

⁵Note that the typical way we build classical cryptography from hard mathematical problems *is* via BPP reductions. But as already discussed, the conjectured #P hardness of the problem implies that we do not expect BQP or even PH reductions to be able to use an inverter for a one-way function to approximate these probabilities.

1.1.2 One-way Puzzles and the Hardness of (Pseudo)-Deterministic Sampling

Given the connection developed above between quantum advantage conjectures and quantum cryptography, it is natural to wonder how far we can push this connection. For example, how generically can one claim that the existence of quantum advantage implies the existence of quantum commitments?

While we do not know how to build quantum cryptography in Microcrypt generically from the assumption that $(\text{Samp})\text{BQP} \neq (\text{Samp})\text{BPP}$, we do in fact obtain one-way puzzles if we assume the existence of distributions that cannot be *pseudo-deterministically* sampled⁶.

We show that distributions that admit a quantum sampler but do not admit a *pseudo-deterministic* quantum sampler imply the existence of quantum commitments.

In more detail, assume that there exists a distribution \mathcal{D} that can be efficiently quantumly sampled, such that every *pseudo-deterministic* efficient quantum sampler fails to sample from any distribution \mathcal{D}' with $\text{SD}(\mathcal{D}, \mathcal{D}') \leq \epsilon$ (where SD denotes statistical distance), and $\epsilon = \frac{1}{p(n)}$ for some fixed polynomial $p(\cdot)$. We show that the existence of any such distribution implies the existence of one-way puzzles.

Theorem 1.3. *[Informal] If there exists a quantumly sampleable distribution that does not admit a pseudo-deterministic sampler, then one-way puzzles and quantum commitments exist.*

Thus, if one-way puzzles do not exist, then every quantumly sampleable distribution can essentially be treated as being deterministically sampleable. We believe that this outlook may in the future yield simple ways to build cryptography from hardness assumptions associated with quantum sampleable distributions.

Our next set of results relates the hardness of state synthesis with one-way puzzles and commitments. This also builds on the conceptual connection between the hardness of approximating probabilities and quantum cryptography.

1.1.3 One-way Puzzles and the Hardness of State Synthesis

Another fundamental “quantum” search problem is what we will call a *state puzzle* – this is like a one-way puzzle except that the hard-to-find ‘key’ is not classical, but an arbitrary quantum state. A state puzzle consists of a quantum polynomial time algorithm \mathcal{G} that samples a (secret) quantum state $|\psi_s\rangle$ along with a (public) string s – here w.l.o.g. we may assume that for every s , $|\psi_s\rangle$ is pure. The security guarantee is that given the “puzzle” s , it is infeasible for non-uniform QPT machines to synthesize any state that noticeably overlaps with $|\psi_s\rangle$ in expectation. An extremely natural question, that we address in this work, is whether quantum search problems such as state puzzles also imply one-way puzzles and quantum commitments.

The complexity of synthesizing known quantum states has been studied in several works [Aar16, INN⁺22, Ros24], and it is known that a state puzzle can be inverted by a BQP machine with access to a $\#P$ oracle [Aar16, INN⁺22]. Thus, the existence of state puzzles implies that $\text{BQP}^{\#P} \not\subseteq (\text{io})\text{BQP}/\text{qpoly}$.

By our previous result (Informal Theorem 1.1), we know that Assumption 1 implies the existence of quantum commitments, as long as $\text{BQP}^{\#P} \not\subseteq (\text{io})\text{BQP}/\text{qpoly}$ ⁷. From the discussion in the previous paragraph, we can replace $\text{BQP}^{\#P} \not\subseteq (\text{io})\text{BQP}/\text{qpoly}$ with the existence of state puzzles,

⁶By pseudo-deterministic, we mean that the sampler when run multiple times on the same input, with high probability outputs the same result on all executions.

⁷While we state Theorem 1.1 as assuming that $\text{P}^{\#P} \not\subseteq (\text{io})\text{BQP}/\text{qpoly}$, we note that this is in fact equivalent to $\text{BQP}^{\#P} \not\subseteq (\text{io})\text{BQP}/\text{qpoly}$.

in the previous sentence. Thus, if Assumption 1 holds, then state puzzles imply one-way puzzles. However, Assumption 1 is about average-case hardness of approximating probabilities, and state puzzles capture the average-case hardness of synthesizing states. So can state puzzles imply one-way puzzles and quantum commitments unconditionally? We show that this is indeed the case, and that Assumption 1 is not needed for this implication.

Theorem 1.4 (Informal). *State puzzles imply quantum bit commitments.*

As an intermediate step, we again build one-way puzzles unconditionally from state puzzles. To obtain one-way puzzles, we introduce novel techniques that reduce phase estimation to calculating probabilities of outcomes of a distribution generated using the state itself. This in fact proves the following equivalence:

Theorem 1.5 (Informal). *The existence of state puzzles is equivalent to the existence of one-way puzzles.*

Public-Key Quantum Money. Quantum money aims to use states as banknotes, leveraging the no-cloning principle to prevent counterfeiting. In a simplified model (often called a “public-key quantum money mini-scheme” [AC12]), a sampler outputs a banknote $|\psi_s\rangle$ together with a classical serial number s that can be efficiently obtained from $|\psi_s\rangle$ without disturbing it. Furthermore, it is computationally intractable to generate “clones” of $|\psi_s\rangle$. It is easy to see that a mini-scheme is also a state puzzle, since if synthesizing $|\psi_s\rangle$ were easy given s , then one could clone a banknote $|\psi_s\rangle$ efficiently by first computing s and synthesizing $|\psi_s\rangle$ from s . This observation combined with Theorems 1.4 and 1.5 yields the following corollary.

Corollary 1.1. *Quantum money mini-schemes imply one-way puzzles and quantum bit commitments.*

We stress that in general state puzzles appear to be a weaker primitive than quantum money – unlike money, they (1) do not require unclonability, only the hardness of generating $|\psi_s\rangle$ given s , and (2) do not require $|\psi_s\rangle$ to be efficiently verifiable with respect to s .

Amplifying State Puzzles. It is also natural to consider a ‘weak’ variant of a state puzzle (analogous to weak one-way functions), where it is computationally intractable to synthesize any state that overwhelmingly overlaps with $|\psi_s\rangle$ in expectation. Our implication from state puzzles to one-way puzzles holds even when starting with a weak state puzzle, which combined with Theorem 1.5 yields the following amplification theorem for state puzzles.

Theorem 1.6. *Weak state puzzles are equivalent to (strong) one-way puzzles and state puzzles.*

1.2 Perspective

Finally, we reflect on the implications of these results in the broader context of understanding hardness in quantum cryptography.

Microcrypt is “Real”. As already discussed above, this work provides the strongest evidence so far for the existence of Microcrypt: there are real, unrelativized constructions of quantum cryptosystems based on well-studied mathematical assumptions that do not imply the existence of one-way functions (as otherwise, quantum advantage claims break down). Under mild complexity assumptions ($P^{\#P} \not\subseteq \text{ioBQP}^{\text{NP}}$), the resulting constructions remain secure even against an adversary that can access an NP oracle. This also rules out *non-black-box constructions* of one-way functions from one-way puzzles or quantum commitments (with a black-box reduction), assuming any of the probability approximation conjectures. For such a reduction would be able to

use an NP oracle – that inverts any one-way function – to also solve a $\#P$ -hard problem, which cannot happen unless $P^{\#P} \subseteq \text{ioBQP}^{\text{NP}}$. Previous oracle results inherently only ruled out *black-box/relativizing constructions* (with black-box reductions) of one-way functions from quantum commitments and puzzles.

We also note that the mathematical assumptions/conjectures underlying our primitives are not new and have previously been extensively studied in the completely different context of quantum advantage. It is only their application to the realm of cryptography that is new.

A Sharper Understanding of Microcrypt. Since disproving the probability approximation conjectures will have far-ranging consequences in quantum advantage, let us assume that at least one of these conjectures (originally made in the context of quantum advantage) is true. Then the existence of one-way puzzles is equivalent to $P^{\#P} \not\subseteq (\text{io})\text{BQP}/\text{poly}$.

This gives us a sharper perspective on the hardness of primitives in Microcrypt. For one, any primitives that become insecure in the presence of a $\#P$ oracle (e.g. state puzzles, but also any new primitives we come up with in the coming years) will imply one-way puzzles and therefore also, commitments.

This equivalence also enables new insights into separations. For instance, one-way puzzles (and commitments) should not imply quantum cryptographic primitives that break in the presence of a QMA oracle, assuming $P^{\#P} \not\subseteq (\text{io})\text{BQP}^{\text{QMA}}$ (see e.g. [Vya03] for some evidence in support of this assumption). Concretely, a public-key quantum money (PKQM) mini-scheme implies one-way puzzles/commitments (by Corollary 1.1) but becomes insecure given a search-QMA oracle, and therefore likely should not be implied by one-way puzzles/commitments. Besides PKQM, other natural primitives that break in the presence of a search-QMA oracle include public key encryption (PKE) with classical public keys (and quantum secret keys/ciphertexts), digital signatures (DS) with classical verification keys (and quantum secret keys/signatures). Under any of the hardness of approximation conjectures, $P^{\#P} \not\subseteq (\text{io})\text{BQP}^{\text{search-QMA}}$ will imply a separation between these forms of PKQM/PKE/DS and one-way puzzles/commitments.

On Minimal Assumptions for Quantum Cryptography. How hard is it to break quantum cryptography? This work strengthens evidence that it is at least $P^{\#P}$ -hard to break one-way puzzles (and therefore also, quantum commitments).

While the bound of $P^{\#P}$ is tight for one-way puzzles⁸, it is not known to be tight for quantum commitments. In fact, recent work [LMW24] demonstrated (relative to a random oracle) the existence of quantum commitments that remain secure against all efficient adversaries that make only a single query to an arbitrary Boolean oracle. They also conjecture that their single query lower bound extends to polynomially many queries, and therefore that there is *no* classical oracle \mathcal{O} such that $\text{BQP}^{\mathcal{O}}$ will break a commitment. Proving or disproving this conjecture remains an open problem – in the small chance that $\text{BQP}^{\#P}$ does successfully break commitments, then our results would say that commitments also imply one-way puzzles (assuming the probability approximation conjectures, although we suspect that the conjectures may not be necessary just like the case of state puzzles). If their conjecture is true, then there is also the fascinating possibility that (computationally secure) quantum commitments can be constructed *unconditionally*, i.e. without the need for unproven assumptions. However, so far, unconditionally secure commitments are only known in models where participants have access to (inefficiently prepared) quantum auxiliary input [Qia24, MNY24].

⁸One-way puzzles can be broken by a $P^{\#P}$ machine [CGG⁺23].

Beyond Cryptography. This work shows that BosonSampling and Random Circuit Sampling are not as “useless” as they are often made out to be. Our constructions use the ability to efficiently sample from these distributions in a crucial way, although we also need some additional quantum processing to finally obtain useful cryptography. Our work also uncovers a connection between sampling-based quantum advantage and the complexity of decoding Hawking radiation. In more detail, Assumption 1 together with the mild assumption that $P^{\#P} \not\subseteq BQP/qpoly$ implies the existence of hard-to-decode Hawking radiation. This follows by combining this work with a theorem of [Bra23] demonstrating an equivalence between the existence of quantum commitments and the hardness of black-hole radiation decoding.

1.3 Open Problems

We will now examine some remaining technical obstacles to gaining a complete understanding of hardness in quantum cryptography.

1. **Concrete Instantiations of Other Microcrypt Primitives.** An important open challenge is to base other quantum primitives like pseudorandom states and unitaries, digital signatures, quantum money, etc. on concrete, ideally well-studied, mathematical assumptions weaker than one-way functions. For example: can we prove, under quantum advantage conjectures, that the output of random circuits are pseudorandom states? Note that such a claim can only hold for specific circuit architectures: for instance, BosonSampling outputs are trivially distinguishable from random [AA14]. Other primitives such as quantum money mini-schemes will likely require different mathematical assumptions that may not be as hard as $P^{\#P}$, but also do not necessarily imply the existence of one-way functions.
2. **Connections between Quantum Advantage and Cryptography.** Ours is not the first work to use tools developed in the context of quantum advantage to obtain evidence for quantum cryptography in Microcrypt. Previously, [KOST23] built on the oracle separation of BQP and PH [RT22] to demonstrate, *relative to an oracle*, that quantum commitments exist even when $P = NP$. This work develops similar connections without relying on oracles. How far does this relationship between quantum advantage and cryptography in Microcrypt extend? For example, it may not be outrageous to imagine that the classical hardness of factoring implies a one-way puzzle, although we do not (yet) know how to prove this.
3. **Fully Quantum Search Problems.** We now know how to build commitments from various classical-quantum search problems: one-way state generators [MY22b, MY22a] – where the challenge is a quantum state and the solution is a classical key – imply commitments [KT24, BJ24]; and so do state puzzles, where the challenge is classical and the solution is a quantum state. These types of search problems are meaningful because they are often easily implied by other cryptographic primitives and protocols, and are much easier to cryptanalyze than decision problems. For example as demonstrated by this work, one-way puzzles form a useful link between conjectures in quantum advantage and the existence of quantum commitments. A useful next step towards understanding the relationships between cryptographic primitives is to study computational search problems where both the challenge and the solution are quantum states.

2 Related Work

We discuss related independent work below, in decreasing order of technical overlap.

Qian, Raizes and Zhandry [QRZ24] independently obtain commitments from state puzzles via completely different techniques⁹. Their techniques generalize to also imply commitments from a variant of state puzzles where the classical challenge is replaced by a clonable basis state¹⁰. On the other hand, our techniques yield a stronger theorem statement about state puzzles: we show that even weakly hard state puzzles imply one-way puzzles, strongly hard state puzzles and commitments. Using measurements to obtain a “classical handle” on quantum states is a technique that is unique to our work, and this allows us to obtain an implication to one-way puzzles as well as an amplification theorem for state puzzles. We also present candidates for state puzzles assuming #P-hardness and conjectures from quantum advantage. Other parts of our paper, including building cryptography from #P hardness and well-studied conjectures, have no technical overlap with their work.

Two other concurrent and independent works connect quantum cryptography with meta-complexity [CGGH24, HM24]. Their techniques also establish an equivalence (between the hardness of approximating probabilities with the existence of one-way puzzles ([HM24] prove only one direction of the equivalence). The proof of our first result also establishes such an equivalence, and this is the only part of our paper that overlaps with [CGGH24, HM24]. We note that in order to build puzzles from quantum advantage conjectures, we require (and obtain) a stronger implication: we build puzzles from the hardness of approximating probabilities upto $(1 \pm 1/\text{poly})$ multiplicative approximation factors for arbitrary polynomials poly , whereas techniques in [CGGH24, HM24] appear to be limited to constant factor approximations only. Additionally, unlike the assumptions we use, there is no independent evidence that the meta-complexity assumptions studied in [CGGH24, HM24] are separated from the existence of one-way functions.

Two additional works [MSY24, BHHP24] are thematically related, but do not have technical overlap with our work. The first, [MSY24], proves an equivalence between interactive (inefficiently verifiable/sampling-based) quantum advantage and (only) *classically-secure* one-way puzzles. Classically secure puzzles are a new primitive introduced in [MSY24] and are *not* known to imply classical or quantum cryptographic primitives (such as commitments). Therefore, [MSY24] do not build useful cryptography from the existence of quantum advantage, or from conjectures in quantum advantage. Additionally in this work, we observe that if one-way puzzles exist but one-way functions do not, this immediately implies advantage (see Appendix D). Combining this with prior work [MY24] obtaining interactive advantage from one-way functions, makes it easy to see why puzzles imply advantage. The other work [BHHP24] introduces *new* hardness assumptions that have not been previously scrutinized, and uses these to build pseudoentangled states, quantum trapdoor functions and pseudorandom unitaries (on the other hand our goal is to rely on well-founded assumptions). Unlike our work, it is not clear whether their assumptions are separated from the existence of one-way functions.

In summary, our primary conceptual contribution – building quantum cryptography from well studied assumptions weaker than one-way functions, and our primary technical contribution – converting state puzzles to a classical search problem by carefully replacing a #P oracle, do not overlap with any other work.

⁹QRZ verbally announced their results, to certain private audiences, before us. However, we only learned that they proved an implication from state puzzles to commitments towards the completion of our manuscript. Techniques in both works were developed independently and the manuscripts were publicly posted concurrently.

¹⁰We suspect that our techniques can also be extended (e.g., using efficient shadow tomography [HKP20]) to yield one-way puzzles from such primitives, but we leave a formal exploration to future work.

3 Technical Overview

We now provide an overview of our techniques.

Recall that our dream goal is to build quantum cryptography from the hardness of #P, for which computing permanents is a complete problem. An immediate obstacle to building cryptography from the hardness of computing permanents is that it appears difficult to efficiently sample random matrices together with their permanents. If such sampling were (quantumly) efficiently possible, we would be done: we would set our one-way puzzle to be the matrix, and the corresponding solution to its permanent.

In the absence of such samplers, we turn to the rich literature on sampling based quantum advantage (e.g., [SB09, BJS11, AA11, BMS16, FM17, BIS⁺18, BFNV19, KMM21, BFL21, Kro22, Mov23, ZVBL23]). This line of work obtains quantum advantage from #P hardness by building on the following observations:

1. There exist quantumly sampleable distributions \mathcal{X} such that $\Pr_{\mathcal{X}}[x]$ – i.e., the probability assigned by \mathcal{X} to any string x – equals the square of the permanent of a corresponding unitary matrix. Moreover, permanents are known to be #P-hard to compute on average, and even #P-hard to *approximate*, upto inverse polynomial relative error, in the worst case.
2. If there existed a *classical sampler* that was able to sample *exactly* from \mathcal{X} , then $P^{\#P} = BPP^{\text{NP}}$ which is highly implausible because it would collapse the polynomial hierarchy (by Toda’s theorem).

Why would the existence of an exact classical sampler imply that $P^{\#P} = BPP^{\text{NP}}$? This is because given any deterministic sampler \mathcal{O} for \mathcal{X} , universal hashing makes it possible to approximate $\Pr_{\mathcal{X}}[x]$ for every x to within a multiplicative constant in $BPP^{\text{NP}^{\mathcal{O}}}$ [Sto83]. Since there is at least one x for which approximating $\Pr_{\mathcal{X}}[x]$ is #P-hard, this puts $P^{\#P} \subseteq BPP^{\text{NP}^{\mathcal{O}}}$.

Ruling out the possibility of classically *approximately* sampling from \mathcal{X} is not as straightforward. Suppose there exists a classical sampler that samples from a distribution \mathcal{Y} such that $\text{SD}(\mathcal{X}, \mathcal{Y}) \leq \epsilon$ for some small constant ϵ . The arguments above break down because this sampler may lead to large errors in approximating $\Pr_{\mathcal{X}}[x]$ for certain x , and it is no longer possible to rely only on the *worst case* hardness of approximating $\Pr_{\mathcal{X}}[x]$. What is done instead is that permanents are conjectured to be #P-hard to approximate even in the average case (see, e.g., [AA11]). This conjecture, combined with a type of rerandomization or *hiding* property of the sampler implies that probabilities $\Pr_{\mathcal{X}}[x]$ are #P-hard to *approximate on average for uniform choice of x* .

Furthermore, $\Pr_{\mathcal{X}}[x]$ are assumed to *anti-concentrate*, i.e. they must not be too small too often – suppose that an overwhelming fraction of strings x had $\Pr_{\mathcal{X}}[x] < \frac{1}{2^{2n}}$, then a classical sampler could sample from a statistically close distribution while still assigning incorrect probabilities (say $\frac{1}{2^{2n}}$) to an overwhelming fraction of the x values.

In particular, the following assumption cleanly captures the hardness implied by a variety of known sampling-based quantum advantage conjectures for RCS, BosonSampling, IQP, DQC, etc.

Assumption 1 (#P-Hardness of Approximating Probabilities). *There is a family of (uniform) efficiently sampleable distributions $\mathcal{C} = \{\mathcal{C}_n\}_{n \in \mathbb{N}}$ over quantum circuits C where each C has n -qubit outputs (and potentially additional junk qubits), such that there exist polynomials $p(\cdot)$ and $\gamma(\cdot)$ satisfying:*

1. **Anticoncentration.** *For all large enough $n \in \mathbb{N}$*

$$\Pr_{\substack{C \leftarrow \mathcal{C}_n \\ x \leftarrow \{0,1\}^n}} \left[\Pr_C[x] \geq \frac{1}{(p(n) \cdot 2^n)} \right] \geq \frac{1}{\gamma(n)}$$

2. **Hardness of Approximating Probabilities.** For any oracle \mathcal{O} satisfying that for all large enough $n \in \mathbb{N}$,

$$\Pr_{\substack{C \leftarrow \mathcal{C}_n \\ x \leftarrow \{0,1\}^n}} \left[|\mathcal{O}(C, x) - \Pr_C[x]| \leq \frac{\Pr_C[x]}{p(n)} \right] \geq \frac{1}{\gamma(n)} - \frac{1}{p(n)}$$

we have that $\text{P}^{\#\text{P}} \subseteq \text{BPP}^{\mathcal{O}}$. Here, $\Pr_C[x]$ denotes the probability of obtaining outcome x when the output register of $C|0\rangle$ is measured in the standard basis.

Assumption 1 yields a new route to obtaining one-way puzzles, as we describe next. For conceptual reasons, we will begin by reformulating the one-way puzzle primitive in terms of the hardness of post-selected sampling.

The Hardness of Post-Selected Sampling implies One-way Puzzles. For efficiently sampleable distribution \mathcal{X} , consider the task of *post-selected* sampling from \mathcal{X} . Namely, a challenger samples $x^* \leftarrow \mathcal{X}$, $i \in [n]$, and outputs x^* truncated to its first $i - 1$ bits: $(x^*_{[1\dots i-1]})$. The post-selected sampling task is to sample from $x \leftarrow \mathcal{X}$ conditioned on their first $(i - 1)$ bits of x being $x^*_{[1\dots i-1]}$.

Observe that the hardness of post-selected sampling immediately implies the existence of a (distributional) one-way puzzle: the puzzle sampler samples $x = x_{[1\dots n]} \leftarrow \mathcal{X}$, $i \leftarrow [n]$ and outputs $\text{puz} = i, x_{[1\dots i-1]}$ and $\text{sol} = (x_{[i\dots n]})$. Sampling from the distribution over sol corresponding to a puzzle puz is exactly the task of post-selected sampling. We will make use of this implication (together with the fact that distributional one-way puzzles imply one-way puzzles [CGG24]) in the upcoming subsections.

Finally, we note that the hardness of post-selected sampling is related to the hardness of *universal extrapolation* [IL90], and the hardness of approximating probabilities is related to the hardness of *universal approximation* [OW93, IL90]. The existence of universal extrapolators and universal approximators for classically sampleable distributions is known to be equivalent to the existence of one-way functions [IL90]. This work will implicitly show and exploit analogous equivalences between one-way puzzles and the hardness of universal approximation/extrapolation for quantumly sampled distributions.

3.1 One-way Puzzles from the Hardness of Approximating Probabilities

Our first key insight is that an *exact* post selected sampler makes it easy to compute probabilities of strings, upto inverse polynomial multiplicative error. In more detail, given an *exact* post-selected sampler $S_{\mathcal{X}}$ for \mathcal{X} , there is a polynomial-time machine R parameterized by a polynomial $p(\cdot)$ – that with oracle access to $S_{\mathcal{X}}$ and on input a string $v \in \{0, 1\}^n$ – outputs an approximation of $\Pr_{\mathcal{X}}[v]$ that is accurate upto inverse polynomial relative error.

$R^{S_{\mathcal{X}}}(v) :$

- Parse v as a sequence of bits $v[1]v[2] \dots v[n]$.
- Set $\text{prefix} = \perp$, and set $\text{pr} = 1$.
- For $i \in [n]$, do the following:
 - Run S on input prefix $p(n)$ times, and let pr_i denote the fraction of times that the first bit of the output is $v[i]$.
 - Set $\text{pr} = \text{pr} \cdot \text{pr}_i$, and $\text{prefix} = \text{prefix}||v[i]$.

- Output pr.

By Chernoff bounds, as long as for each prefix (denoted by prefix_v) of v , $\Pr_{\mathcal{X}}[\text{prefix}_v] > \frac{1}{q(n)} \cdot \frac{1}{2^{|\text{prefix}_v|}}$ for some fixed polynomial $q(\cdot) < p(\cdot)$, the reduction R above outputs an approximation to $\Pr_{\mathcal{X}}[v]$ with small inverse polynomial relative error. On the other hand, if $\Pr_{\mathcal{X}}[v] \ll \frac{1}{q(n)} \cdot \frac{1}{2^n}$ the reduction can fail, so we do not get a good approximation of probabilities in the worst case, and are only able to contradict *average-case hardness* of approximating $\Pr_{\mathcal{X}}[x]$ (Assumption 2). In the actual analysis, we crucially use the fact that challenges v that are sent to R , are sampled according to the distribution \mathcal{X} , and therefore strings v for which $\Pr_{\mathcal{X}}[v] \ll \frac{1}{q(n)} \cdot \frac{1}{2^n}$ are sampled with relatively low probability.

Note that the analysis described so far applies if we had a *perfect* post-selected sampler. But in the definition of a (distributional) one-way puzzle, not only do we want the hardness of sampling *exactly* from the target distribution, we also need it to be hard for adversaries to sample from a distribution \mathcal{D}' that is inverse-polynomially close (in statistical distance) to the target distribution. This gives an adversary/post-selected sampler the flexibility to introduce arbitrary errors in sampling, while maintaining overall low statistical distance from the target distribution. However, note that this latter requirement enforces that the adversary can only introduce high relative errors on values v (and their prefixes) for which $\Pr_{\mathcal{X}}[v]$ is low. We use this observation together with a more sophisticated analysis to show that the reduction R described above will still output a low relative error approximation to $\Pr_{\mathcal{X}}[x]$ on average, as long as $\Pr_{\mathcal{X}}[x]$ is not too small (i.e., $\Pr_{\mathcal{X}}[v] \geq \frac{1}{q(n)} \cdot \frac{1}{2^n}$).

This, combined with Assumption 1 yields a quantum polynomial time machine that solves $\#P$ -hard problems, contradicting the assumption that $P^{\#P} \not\subseteq (\text{io})\text{BQP}/\text{qpoly}$. This completes a sketch of our proof that Assumption 1 implies one-way puzzles. Our actual proof first further abstracts out the properties we need from Assumption 1 along with $P^{\#P} \not\subseteq (\text{io})\text{BQP}/\text{qpoly}$ into a different assumption (described below). The analysis above is then applied to prove that Assumption 2 implies one-way puzzles.

Assumption 2 (Native Approximation Hardness). *There exists a family of (uniform) efficiently samplable distributions $\mathcal{D} = \{\mathcal{D}_n\}_{n \in \mathbb{N}}$ over classical strings such that there exists a polynomial $p(\cdot)$ such that for all QPT $\mathcal{A} = \{\mathcal{A}_n\}_{n \in \mathbb{N}}$, every (non-uniform, quantum) advice ensemble $|\tau\rangle = \{|\tau_n\rangle\}_{n \in \mathbb{N}}$, and large enough $n \in \mathbb{N}$,*

$$\Pr_{x \leftarrow \mathcal{D}_n} \left[\left(\left| \mathcal{A}(|\tau\rangle, x) - \Pr_{\mathcal{D}_n}[x] \right| > \frac{\Pr_{\mathcal{D}_n}[x]}{p(n)} \right) \right] \geq \frac{1}{p(n)}$$

We point out one important technical issue that arises from the mismatch between the complexity-theoretic style of Assumption 1 and the cryptographic style of Assumption 2. We would like to use a BQP/qpoly adversary that contradicts Assumption 2 to show that $P^{\#P} \subseteq (\text{io})\text{BQP}/\text{qpoly}$, as long as Assumption 1 holds. To contradict Assumption 2, the BQP/qpoly adversary only needs to succeed in approximating probabilities on infinitely many $n \in \mathbb{N}$. On the other hand, Assumption 1 converts any adversary that approximates probabilities *for every large enough* $n \in \mathbb{N}$ into an oracle that can solve $\#P$ -hard problems for large enough n . It is at first unclear why these two statements can be put together to obtain the implication we want, i.e., $P^{\#P} \subseteq (\text{io})\text{BQP}/\text{qpoly}$. But on opening things up, we observe that Assumption 1 guarantees a black-box reduction that on input length n , queries an approximator adversary on polynomially many input lengths, each polynomially related to n . By modifying our distribution for Assumption 2 to output samples for each of these input lengths, we ensure that the approximator adversary answers all of the reductions queries correctly, infinitely often. This allows us to conclude that $P^{\#P} \subseteq (\text{io})\text{BQP}/\text{qpoly}$, as desired.

An Equivalence between Puzzles and the Hardness of Approximating Probabilities. We also prove a reverse implication, i.e., the existence of one-way puzzles implies that Assumption 2 holds. To prove this, we would like to use one-way puzzles to define a distribution \mathcal{D} such that we can invert the puzzle given a probability estimator for strings in the support of \mathcal{D} .

Defining the distribution to be the same as the output of the one-way puzzle sampler does not work. This is because even given an estimator that perfectly computes probabilities 100% of the time, it is unclear how to find a key k corresponding to a puzzle s by using an oracle that on input any $(s||k)$ outputs $\Pr_{\text{Samp}}[(s||k)]$.

Instead, following [CGG⁺23], we will aim to perform a binary search for k , given s . Indeed as a first attempt, our distribution \mathcal{D} will be defined as follows: run the puzzle sampler Samp to obtain puzzle and key (s, k) , sample $i \leftarrow [n]$ and then output $(i, s, k_{1\dots i})$ where $k_{1\dots i}$ denotes a truncation of k to the first i bits.

Now given s , it may at first appear easy to search for consecutive bits of k using a probability estimator for \mathcal{D} : first, run the estimator on $(1, s, 0)$ to obtain pr_0 and $(1, s, 1)$ to obtain pr_1 . Pick bit b for which $\text{pr}_b \geq \text{pr}_{1-b}$, and set the first bit of k , i.e. k_1 to b . Next, run the estimator on $(1, s, k_1 0)$ to obtain $\text{pr}_{k_1,0}$ and $(1, s, k_1 1)$ to obtain $\text{pr}_{k_1,1}$. Pick bit b' for which $\text{pr}_{k_1,b'} \geq \text{pr}_{k_1,1-b'}$, set the second bit of k to b' , and continue the process to iteratively find a key k . Indeed, this works if the estimator always returns correct probabilities, even on strings that are not in support of the distribution.

However, our assumption only requires the adversarial probability approximator to return (approximately) correct probabilities on strings that have non-zero probability mass. The construction and analysis above breaks down given such an approximator: in particular, we began by running the estimator on $(1, s, 0)$ and $(1, s, 1)$ to see if keys for s had a higher probability of beginning with 0 or with 1. Note that if every actual preimage key of s had first bit 0, the point $(1, s, 1)$ would be assigned zero probability mass in the distribution, meaning the adversary may return arbitrary values on $(1, s, 1)$ to confuse our search algorithm, without any penalty. The same problem can arise even if we have extremely low, but non-zero probability mass on certain points. We resolve this by modifying the distribution, as we describe next.

We will run the puzzle sampler Samp to obtain (s, k) , sample $i \leftarrow [n]$, and sample bit $c \leftarrow \{0, 1\}$. If $c = 0$, set $\tilde{k} = k_{1\dots i}$ and if $c = 1$ set $\tilde{k} = k_{1\dots i-1}||\beta$ for $\beta \leftarrow \{0, 1\}$. Output (i, s, \tilde{k}) . Intuitively adding some probability mass to both 0 and 1 on the last bit, we force an adversary to pay a penalty in statistical distance whenever they send the binary search algorithm described above down an incorrect path. The analysis requires some additional care to account for various types of errors, but we are able to conclude that any probability approximator for the above distribution implies an inverter for the one-way puzzle.

We refer the reader to Section 5 for additional details and a formal proof of the equivalence.

3.2 One-way Puzzles from the Hardness of Pseudo-Deterministic Sampling

Next, we discuss why the existence of a distribution \mathcal{X} that is quantumly efficiently sampleable, but is not efficiently *pseudo-deterministically* sampleable implies the existence of one-way puzzles.

An ϵ -pseudo-deterministic sampler is a QPT machine \mathcal{Q} along with (non-uniform, quantum) advice ensemble $|\tau\rangle = \{|\tau_n\rangle\}_{n \in \mathbb{N}}$, that satisfies the following property for all large enough $n \in \mathbb{N}$: for at least $1 - \epsilon(n)$ fraction of random strings $r \in \{0, 1\}^n$,

$$\exists y \text{ s.t. } \Pr[\mathcal{Q}(|\tau\rangle; r) = y] = 1 - \text{negl}(n)$$

A distribution \mathcal{X} is ϵ -pseudo-deterministically sampleable with $\epsilon(\cdot)$ error if there exists an efficient quantum ϵ -pseudo-deterministic sampler that outputs distribution \mathcal{D} such that $\text{SD}(\mathcal{X}, \mathcal{D}) \leq \epsilon(n)$.

We prove that if one-way puzzles do not exist, then for every polynomial $q(\cdot)$, every quantumly sampleable distribution is also $\frac{1}{q(n)}$ -pseudo-deterministically sampleable.

Our key insight here is that if post-selected sampling is easy, then every distribution can be pseudo-deterministically sampled by using the post-selected sampler to approximate probabilities. We describe a reduction, parameterized by a polynomial $p(\cdot)$ that with access to post-selected sampler $\mathcal{S}_{\mathcal{X}}$ for \mathcal{X} , samples pseudo-deterministically from \mathcal{X} .

$\mathcal{R}^{\mathcal{S}_{\mathcal{X}}}(r)$:

- Parse r as a sequence of blocks of randomness $r[1]r[2] \dots r[n]$, each block of size n bits.
- Set prefix = \perp .
- For $i \in [1, n]$, do the following:
 - Run \mathcal{S} on input prefix $p(n)$ times, and let pr_i denote the fraction of times that the first bit of the output is 0.
 - If $r[i] \leq \text{pr}_i \cdot 2^n$, set $x[i] = 0$. Otherwise set $x[i] = 1$.
 - Set prefix = prefix|| $x[i]$.
- Output prefix.

This insight can be turned into a formal proof that for every polynomial $q(\cdot)$, there is a polynomial $p(\cdot)$ such that \mathcal{R} parameterized with $p(\cdot)$ samples $\frac{1}{q(n)}$ -pseudo-deterministically from \mathcal{X} . Using analysis that is similar to the previous section, we prove that the next-bit probabilities computed by \mathcal{R} are approximately correct (on average). This lets us show that except for some bad choices of randomness (which are close to actual probability thresholds), the output of \mathcal{R} is (almost) deterministic. Moreover, the distribution output by \mathcal{R} has inverse polynomial statistical distance from \mathcal{X} as long as \mathcal{S} is a perfect post-selected sampler.

As before, the non-existence of one-way puzzles only guarantees that the adversary can sample from a distribution that is (arbitrarily) inverse-polynomially close to the post-selected distribution. With some more care, we are able to show that the reduction \mathcal{R} described above, even given access to such an adversary, will pseudodeterministically sample from \mathcal{X} . This completes an overview of our technique, and we encourage the reader to see Section 6 for a complete proof.

3.3 One-Way Puzzles from Hard State Synthesis Problems

Finally, we use the conceptual equivalence between the existence of one-way puzzles and the hardness of approximating probabilities to demonstrate an equivalence between one-way puzzles and a natural notion of hard state synthesis problems, which we call state puzzles.

A state puzzle consists of a QPT sampler \mathcal{G} that outputs pairs $(s, |\psi_s\rangle)$ such that given s , it is quantum computationally infeasible to output a state that overlaps noticeably with $|\psi_s\rangle$.

State Puzzles with Real, Positive Amplitudes. Consider any state $|\psi\rangle = \sum_x \alpha_x |x\rangle$ where all amplitudes α_x are real and positive. Measuring this state results in a distribution over computational basis terms, i.e. a distribution D_ψ where $\Pr_{D_\psi}[x] = |\alpha_x|^2$. Intuitively, the hardness of computing probabilities in D_ψ should be related to the hardness of synthesizing $|\psi\rangle$. We use this conceptual connection to obtain an equivalence between state puzzles and one-way puzzles, as follows.

Let us begin by recalling a procedure due Aaronson [Aar16] that calls a classical (PP) oracle to synthesize a state. Let $m(\cdot)$ be a large enough polynomial, and \mathcal{O} be an oracle that on input

(x, i) outputs the value $|\alpha_{x_{1\dots i}|1}|^2 / (|\alpha_{x_{1\dots i}|0}|^2 + |\alpha_{x_{1\dots i}|1}|^2)$, where $x_{1\dots i}|b$ denotes x truncated to its first i bits then concatenated with b , and for any $\ell \leq [n], t \in \{0, 1\}^\ell, |\alpha_t|^2 = \sum_{z \in \{0, 1\}^{n-\ell}} |\alpha_{t|z}|^2$. Intuitively, the oracle outputs the probability that a string sampled from D_ψ will have 1 as its $(i+1)^{\text{th}}$ bit given that the first i bits are $x_{1\dots i}$. Call this probability $p_{1|x_{1\dots i}}$.

- Set $i = 0$. Initialize register A to $|0^n, i\rangle$ and initialize an auxiliary register B to $|0^{m(n)}\rangle$.
- While $i \leq n$,
 - Query the oracle \mathcal{O} on the contents of the A register and CNOT the result on the B register. Denote the the resulting state by

$$\sum_{x \in \{0, 1\}^{i-1}} \beta_x |x, 0^{n-i+1}, i\rangle_A |p_{1|x_{1\dots i}}\rangle_B.$$

- Apply an efficient unitary to the previous state to obtain state

$$\sum_{x \in \{0, 1\}^i} \beta_x |x\rangle (\beta_{x_0}|0\rangle + \beta_{x_1}|1\rangle) |0^{n-i}, i\rangle_A |p_{1|x_{1\dots i}}\rangle_B.$$

where $\beta_{x_0} = \sqrt{1 - p_{1|x_{1\dots i}}}$ and $\beta_{x_1} = \sqrt{p_{1|x_{1\dots i}}}$.

- Use a related call to uncompute the auxiliary information and obtain

$$\sum_{x \in \{0, 1\}^{i+1}} \beta_x |x, 0^{n-i}, i\rangle_A |0^{m(n)}\rangle_B$$

- Set $i = i + 1$, also update the last part of A to $|i + 1\rangle$.

It is straightforward to observe that this process results in a state close to the desired state, upto precision errors in the probabilities. Our goal will be to simulate this procedure with access to an adversary breaking an appropriately defined distributional one-way puzzle.

The one-way puzzle sampler runs the state puzzle sampler \mathcal{G} to obtain $(s, |\psi_s\rangle)$. It then measures $|\psi_s\rangle$ in the standard basis to obtain string x , and samples $i \leftarrow [0, n - 1]$. Finally, it outputs $(s, i, x_{1\dots i})$ as the puzzle, and x_{i+1} as the key.

Assume there exists a perfect distributional inverter for this one-way puzzle. On input s , the reduction queries the one-way puzzle inverter iteratively for $i \in [n]$, starting with a state initialized to $|0^n, 0\rangle \otimes |0^{m(n)}\rangle$. At each step, the reduction queries the one-way puzzle inverter to obtain (coherently) for each basis string $x_{1\dots i}$, multiple samples of the next bit x_{i+1} distributed according to the target distribution. These samples are then used to obtain an estimate p_{x_i} of the probability $|\alpha_{x_i||0}|^2 / (|\alpha_{x_i||0}|^2 + |\alpha_{x_i||1}|^2)$, which is then used to build the state iteratively for each i by applying the same unitary as the one in Aaronson's algorithm described above.

Here, the one-way puzzle inverter may entangle its output on every query with arbitrary junk states on an auxiliary register; and we need to be able to uncompute these junk states. We cannot apply the standard trick of CNOT-ing our probability estimates on a fresh register and uncomputing, since the CNOT will end up disturbing the adversary's state and uncomputing may not remove junk. However, since the probability estimate $p_{1|x_{1\dots i}}$ is guaranteed to be close to the actual probability $|\alpha_{x_{1\dots i}|0}|^2 / (|\alpha_{x_{1\dots i}|0}|^2 + |\alpha_{x_{1\dots i}|1}|^2)$, we are able to use this estimate to compute each step of the synthesis algorithm (i.e., use the probability estimate to insert appropriate amplitudes on $|x_{1\dots i}0\rangle$ and $|x_{1\dots i}1\rangle$ coherently for each $x_{1\dots i}$) and remove junk at the end of each step.

This process applied iteratively for $i \in [0, n - 1]$ yields a state whose trace distance from $|\psi\rangle$ is at most $1/q(n)$ for arbitrary polynomial $q(\cdot)$, as long as the one-way puzzle inverter is $1/p(n)$ -close to the target distribution for some polynomial $p(\cdot)$ related to $q(\cdot)$.

Recovering Phase Information. The technique described above is limited to states with real, positive amplitudes. We need to work harder when the states to be synthesized carry non-trivial phase information. In particular, the one-way puzzle cannot be based only on measuring the state $|\psi_s\rangle$ in the standard basis. Instead, our one-way puzzle will be obtained by randomly choosing to either measure the state in the standard basis as before, or measuring phase information.

Given any state puzzle sampler \mathcal{G} , the one-way puzzle sampler does the following.

- Obtain $(s, |\zeta_s\rangle) \leftarrow \mathcal{G}(1^n)$.
- Sample a random 2-design C and let $|\psi_s\rangle = C(|\zeta_s\rangle)$.
- Sample $c \leftarrow \{0, 1\}$.
- If $c = 0$, then as before, measure $|\psi_s\rangle$ in the standard basis to obtain string x . Sample $i \leftarrow [0, n - 1]$. Output $(s, C, c, i, x_{1\dots i})$ as the puzzle and x_{i+1} as the key, where x_i denotes the first i bits of x and $x[i + 1]$ denotes the $(i + 1)^{\text{th}}$ bit of x .
- If $c = 1$, choose a two-to-one function f defined by a random shift Δ , i.e. $f(x) = f(x \oplus \Delta)$, then apply f to the register containing $|\psi_s\rangle$ and measure the output to obtain a residual state of the form

$$\left(\cos(\theta/2)|x_0\rangle + \sin(\theta/2)e^{-i\phi}|x_1\rangle \right)_A \otimes |f, f(x_0)\rangle_B,$$

for some $x_0, x_1 = x_0 \oplus \Delta$, and some $\theta \in [0, \pi), \phi \in [0, 2\pi)$. Next sample bit $d \leftarrow \{0, 1\}$, and

- If $d = 0$, measure the A register in basis $(|x_0\rangle + |x_1\rangle, |x_0\rangle - |x_1\rangle)$ to obtain bit z .
- If $d = 1$, measure the A register in basis $(|x_0\rangle + i|x_1\rangle, |x_0\rangle - i|x_1\rangle)$ to obtain bit z .

Output (s, C, c, x_0, x_1, d) as the puzzle and z as the key.

Denote by $|\tilde{\psi}_s\rangle$ the state that corresponds to $|\psi_s\rangle$ with its phase information removed. That is,

$$|\tilde{\psi}_s\rangle = \sum_x a_x |x\rangle$$

and

$$|\psi_s\rangle = \sum_x \alpha_x |x\rangle$$

where every a_x is real and positive and $\alpha_x = a_x e^{i\phi_x}$ for $\phi_x \in [0, 2\pi)$.

As already described above, a perfect distributional puzzle inverter for the corresponding input can be queried on $c = 0$ to synthesize a state close to $|\tilde{\psi}_s\rangle$. We discuss how the puzzle inverter queried on $c = 1$ can be used to find and insert phases $e^{i\phi_x}$ on every basis term $|x\rangle$ in $|\tilde{\psi}_s\rangle$.

Recall that when $c = 1$, the puzzle is generated by applying a two-to-one function and measuring its output, which results in state

$$|\psi\rangle_{f, x_0} := (\cos(\theta/2)|x_0\rangle + \sin(\theta/2)e^{-i\phi}|x_1\rangle)_A \otimes |f, f(x_0)\rangle_B.$$

Measuring $|\psi\rangle_{f, x_0}$ in basis

$$|x_0\rangle + |x_1\rangle, |x_0\rangle - |x_1\rangle$$

results in outcome $|x_0\rangle + |x_1\rangle$ with probability $(1 + \sin \theta \cos \phi)/2$; and in basis

$$|x_0\rangle + i|x_1\rangle, |x_0\rangle - i|x_1\rangle$$

results in $|x_0\rangle + i|x_1\rangle$ with probability $(1 + \sin \theta \sin \phi)/2$.

The puzzle itself is (s, C, c, x_0, x_1, d) where $d \leftarrow \{0, 1\}$. If $d = 0$, the key is the outcome of measuring $|\psi\rangle_{f, x_0}$ in the first basis, which is 0 with probability $(1 + \sin \theta \cos \phi)/2$ and 1 otherwise. If $d = 1$, the key is the outcome of measuring $|\psi\rangle_{f, x_0}$ in the second basis, which is 0 with probability $(1 + \sin \theta \sin \phi)/2$ and 1 otherwise.

The unifying technique in this work is to use a one-way puzzle inverter to approximate probabilities, and that is exactly what we will do at this point. For fixed f and x_0 , we will use the one-way puzzle inverter to estimate the probability values $(1 + \sin \theta \cos \phi)/2$ and $(1 + \sin \theta \sin \phi)/2$ corresponding to the state $|\psi\rangle_{f, x_0}$; which will then help us approximate the relative phase $e^{i\phi}$ between the basis terms $|x_0\rangle$ and $|x_1\rangle$ in $|\psi\rangle$. Here, $x_1 = x_0 \oplus \Delta$ for Δ indicating the shift chosen by function f . This gives us a way to use the puzzle inverter to compute the relative phase between pairs of terms in $|\psi_s\rangle$. Next, we would like to “insert” the resulting phases in the state $|\tilde{\psi}_s\rangle$ to recover our state $|\psi_s\rangle$.

For this, let us first imagine sampling and fixing a uniformly random anchor $y \in \{0, 1\}^n$. We will then compute the phase on every standard basis term $|x\rangle$ in $|\psi_s\rangle$, relative to $|y\rangle$. Namely, we coherently estimate the relative phases $\phi_{x_0, y}$ for all $|\psi\rangle_{x_0, f}$ – where f applies the shift $\Delta = x_0 \oplus y$. This yields the state

$$|\tilde{\psi}_s\rangle = \sum_x \alpha_x |x\rangle |\tilde{\phi}_{xy}\rangle$$

which can be transformed into

$$|\tilde{\psi}_s\rangle = \sum_x \alpha_x e^{i\tilde{\phi}_{xy}} |x\rangle$$

by applying an efficient unitary that applies the phase and then uncomputes $\tilde{\phi}_{xy}$. As long as the estimates $\tilde{\phi}_{xy}$ were approximately correct, this state is close (upto global phase) to the state $|\tilde{\psi}_s\rangle$.

On Errors in Phase Estimation. Note that the choice of anchor y in the process above affects errors: for instance, if $\langle y | \psi_s \rangle = 0$, the puzzle inverter is allowed to return arbitrary values that may be completely uncorrelated with actual phases in $|\psi_s\rangle$. So if we picked an anchor y for which $\langle y | \psi_s \rangle = 0$, we could end up synthesizing a state that is orthogonal to $|\psi_s\rangle$. The same problem persists even for anchors y for which $\langle y | \psi_s \rangle$ is extremely small, since the puzzle inverter is not perfect, and is allowed to sample from a distribution that has inverse polynomial statistical distance from the target distribution. To synthesize a state close to $|\psi_s\rangle$ in the presence of these errors, we sample our anchor y by measuring the state $|\tilde{\psi}_s\rangle$ in the standard basis, which outputs y proportionally to the probability mass of $|y\rangle$ in $|\psi_s\rangle$.

The accuracy of the phase estimate $\tilde{\phi}_{xy}$ depends on the relative probability mass on $|x\rangle$ and $|y\rangle$ in the state. In particular, to meaningfully recover an estimate of the phase $\tilde{\phi}_{xy}$ with small relative error, we require the ratio α_x/α_y to be at most $p(n)$ for some polynomial $p(\cdot)$. This is where the Clifford operator helps: since $|\psi_s\rangle$ was obtained by applying a random Clifford operator to $|\zeta_s\rangle$, we can use properties of 2-designs along with a careful step-wise Chebyshev bound to argue that the total probability mass on basis terms $|y\rangle$ for which $|\alpha_y|^2 > p(n) \cdot 2^n$ for some polynomial $p(\cdot)$, is less than $1/q(n)$ for a related polynomial $q(\cdot)$. This allows us to condition our analysis on obtaining anchors y for which $|\alpha_y|^2 \leq p(n) \cdot 2^n$. With some additional care in the analysis, we are able to bound the error in reconstructing the state $|\psi_s\rangle$ as a function of the error allowed to the one-way puzzle inverter. In particular, we can reconstruct $|\psi_s\rangle$ to arbitrary small inverse polynomial

error $1/p_1(n)$ as long as the one-way puzzle inverter samples from a distribution that is at most $1/p_2(n)$ -far from the target distribution, for a related polynomial $p_2(\cdot)$.

This gives us an equivalence between weak state puzzles and distributional one-way puzzles. Noting that distributional one-way puzzles are equivalent to (strong) one-way puzzles, this implies an equivalence between weak state puzzles and one-way puzzles. In the following subsection, we describe why one-way puzzles imply state puzzles (with pure state secrets), thereby yielding an amplification theorem for state puzzles.

One-way Puzzles Imply (Strong) State Puzzles. By definition, one-way puzzles imply a form of state puzzles where the state $|\psi_s\rangle$ is replaced with a mixture over classical keys of a one-way puzzle. By purifying this mixture, we obtain a strong state puzzle with pure states. Combined with the results described above, this shows that weak state puzzles are equivalent to strong state puzzles. We refer the reader to Section 7 for a formal statement of the equivalence, and proofs of results related to state puzzles.

4 Preliminaries

In this section, we discuss some notation and preliminary information, including definitions, that will be useful in the rest of the exposition.

4.1 Notation and Conventions

We write $\text{negl}(\cdot)$ to denote any *negligible* function, which is a function f such that for every constant $c \in \mathbb{N}$ there exists $N \in \mathbb{N}$ such that for all $n > N$, $f(n) < n^{-c}$. We will use $\text{SD}(A, B)$ to denote the statistical distance between (classical) distributions A and B .

Quantum conventions. A register X is a named Hilbert space \mathbb{C}^{2^n} . A pure state on register X is a unit vector $|\psi\rangle \in \mathbb{C}^{2^n}$, and we say that $|\psi\rangle$ consists of n qubits. A mixed state on register X is described by a density matrix $\rho \in \mathbb{C}^{2^n \times 2^n}$, which is a positive semi-definite Hermitian operator with trace 1.

A *quantum operation* F is a completely-positive trace-preserving (CPTP) map from a register X to a register Y , which in general may have different dimensions. That is, on input a density matrix ρ , the operation F produces $F(\rho) = \tau$ a mixed state on register Y . A *unitary* $U : X \rightarrow X$ is a special case of a quantum operation that satisfies $U^\dagger U = U U^\dagger = \mathbb{I}^X$, where \mathbb{I}^X is the identity matrix on register X . A *projector* Π is a Hermitian operator such that $\Pi^2 = \Pi$, and a *projective measurement* is a collection of projectors $\{\Pi_i\}_i$ such that $\sum_i \Pi_i = \mathbb{I}$.

We say a quantum circuit C outputs strings in $\{0, 1\}^n$ if C acts on $|0\rangle$ to produce an n -qubit output register (potentially along with a junk register). The output of the circuit is the outcome of measuring the output register of $C|0\rangle$ in the computational basis.

Theorem 4.1 (Additive Chernoff Bound). *For every $i \in [n]$, let X_i be an independent Bernoulli random variable that takes value 1 with probability p . Let $X := \sum_i X_i/n$. Then for $\delta > 0$:*

$$\Pr[|X - p| \geq \delta/\sqrt{n}] \leq 2e^{-\delta^2}$$

Definition 4.1 (Unitary 2-design (from [Mel24])). *Let D be a distribution over n -qubit unitaries. D is a unitary 2-design if and only if for all $O \in \mathcal{L}((\mathbb{C}^2)^{\otimes 2})$*

$$\mathbb{E}_{U \leftarrow D} [U^{\otimes 2} O U^{\dagger \otimes 2}] = \mathbb{E}_{U \leftarrow \mu_H} [U^{\otimes 2} O U^{\dagger \otimes 2}]$$

where μ_H is the Haar measure.

We also use the following theorem showing that the trace distance of pure states is upper bounded by their Euclidean distance.

Theorem 4.2. *Let $|\psi\rangle$ and $|\phi\rangle$ be two pure states such that $\|\psi\rangle - |\phi\rangle\| \leq \epsilon$. Then*

$$\text{TD}(|\psi\rangle\langle\psi|, |\phi\rangle\langle\phi|) \leq \epsilon$$

Proof. We have that

$$\begin{aligned} \epsilon^2 &\geq \|\psi\rangle - |\phi\rangle\|^2 \\ &= (\langle\psi| - \langle\phi|)(|\psi\rangle - |\phi\rangle) \\ &= 2 - (\langle\phi|\psi\rangle + \langle\psi|\phi\rangle) \\ &= 2 - 2\text{Re}(\langle\phi|\psi\rangle) \\ &\geq 2 - 2|\langle\phi|\psi\rangle|, \end{aligned}$$

which can be rearranged to

$$|\langle\phi|\psi\rangle| \geq 1 - \frac{\epsilon^2}{2}.$$

By the identity for trace distance of pure states,

$$\begin{aligned} \text{TD}(|\psi\rangle\langle\psi|, |\phi\rangle\langle\phi|) &= \sqrt{1 - |\langle\psi|\phi\rangle|^2} \\ &= \sqrt{1 - \left(1 - \frac{\epsilon^2}{2}\right)^2} \\ &= \sqrt{\epsilon^2 - \frac{\epsilon^4}{4}} \\ &\leq \epsilon. \end{aligned}$$

This completes the proof. □

4.2 Quantum Cryptographic Primitives

Definition 4.2 (One-way Puzzles). *A one-way puzzle is a pair of sampling and verification algorithms $(\text{Samp}, \text{Ver})$ with the following syntax.*

- $\text{Samp}(1^n) \rightarrow (s, k)$, is a QPT algorithm that outputs a pair of classical strings (s, k) . We refer to s as the puzzle and k as its key. Without loss of generality we may assume that $k \in \{0, 1\}^n$.
- $\text{Ver}(s, k) \rightarrow \top$ or \perp , is a Boolean function that maps every pair of classical strings (k, s) to either \top or \perp .

These satisfy the following properties.

- **Correctness.** *Outputs of the sampler pass verification with overwhelming probability, i.e.,*

$$\Pr_{(s,k) \leftarrow \text{Samp}(1^n)} [\text{Ver}(s, k) = \top] = 1 - \text{negl}(n)$$

- **Security.** Given s , it is (quantum) computationally infeasible to find k satisfying $\text{Ver}(s, k) = \top$, i.e., for every quantum polynomial-sized adversary \mathcal{A} and every quantum advice state $|\tau\rangle = \{|\tau_n\rangle\}_{n \in \mathbb{N}}$,

$$\Pr_{(s,k) \leftarrow \text{Samp}(1^n)} [\text{Ver}(s, \mathcal{A}(|\tau\rangle, s)) = \top] = \text{negl}(n)$$

Definition 4.3 (ε -Distributional One-way Puzzles). For $\varepsilon : \mathbb{N} \rightarrow \mathbb{R}$, a ε -distributional one-way puzzle is defined by a quantum polynomial-time generator $\text{Samp}(1^n)$ that outputs a pair of classical strings (s, k) such that for every quantum polynomial-time adversary \mathcal{A} , every (non-uniform, quantum) advice ensemble $|\tau\rangle = \{|\tau_n\rangle\}_{n \in \mathbb{N}}$, for large enough $n \in \mathbb{N}$,

$$\text{SD}(\{s, k\} \{s, \mathcal{A}(|\tau\rangle, s)\}) \geq \varepsilon(n)$$

where $(s, k) \leftarrow \text{Samp}(1^n)$.

We will sometimes simply refer to distributional one-way puzzles. This is taken to mean $1/p(n)$ -distributional one-way puzzles for some non-zero polynomial p .

The following theorem shows that distributional one-way puzzles can be amplified to (standard) one-way puzzles.

Theorem 4.3 (Theorem 33 from [CGG24], rephrased). *If there exists a polynomial $p(\cdot)$ for which $1/p(n)$ -distributional one-way puzzles exist, then one-way puzzles exist.*

Definition 4.4 (ε -Pseudo-deterministic Hard Distributions). An algorithm \mathcal{B} is ε -pseudo-deterministic for $\varepsilon : \mathbb{N} \rightarrow \mathbb{R}$ if

$$\Pr_r [\exists y \text{ s.t. } \Pr [\mathcal{B}(r) \neq y] \leq \text{negl}(n)] > 1 - \varepsilon(n)$$

where r is a uniformly random string.

A family of efficiently sampleable distributions $D = \{D_n\}$ is ε -pseudo-deterministic hard if for all quantum polynomial-time ε -pseudodeterministic adversaries \mathcal{A} that take (non-uniform, quantum) advice ensemble $|\tau\rangle = \{|\tau_n\rangle\}_{n \in \mathbb{N}}$, for all large enough $n \in \mathbb{N}$

$$\text{SD}(\mathcal{A}(|\tau\rangle; r), D_n) > \varepsilon(n)$$

where r is a uniformly random string.

5 Hardness of Approximating Probabilities implies One-way Puzzles

In this section we define notions of distributions with hard to approximate probabilities and prove their equivalence with one-way puzzles.

5.1 Definitions

We begin by presenting a hardness assumption that has been extensively studied in the literature on sampling based quantum advantage [AA11, BMS16, FM17, BIS⁺18, BFN⁺19, KMM21, Kro22, Mov23, HE23]). Here, the adversary is given uniformly sampled $x \leftarrow \{0, 1\}^n$ together with a quantum circuit C on n qubits (with possibly $m = m(n)$ ancillas). The adversary's task is to estimate $\Pr_C[x]$, i.e. the probability that the output state of $C|0^{n+m}\rangle$ results in outcome x when measured in the standard basis. The adversary is required to output a low *relative error* approximation to

$\Pr_C[x]$, i.e. for some polynomial $p(\cdot)$ and for sufficiently many x , the adversary must output a value y such that

$$\left(1 + \frac{1}{p(n)}\right) \Pr_C[x] \leq y \leq \left(1 - \frac{1}{p(n)}\right) \Pr_C[x]$$

The outputs of circuit C are required to satisfy an additional property called anticoncentration, which says that a noticeable fraction of strings x have probability mass in C above a particular threshold of $1/(p(n) \cdot 2^n)$.

Definition 5.1 (Uniform Approximation Hardness). *A family of (uniform) efficiently sampleable distributions $\mathcal{C} = \{\mathcal{C}_n\}_{n \in \mathbb{N}}$ over quantum circuits C that output classical strings in $\{0, 1\}^n$ has uniform approximation hardness if it has the following properties. There exist polynomials $p(\cdot)$ and $\gamma(\cdot)$ such that:*

1. **Anticoncentration.** For all large enough $n \in \mathbb{N}$

$$\Pr_{\substack{C \leftarrow \mathcal{C}_n \\ x \leftarrow \{0,1\}^n}} [\Pr_C[x] \geq 1/(p(n) \cdot 2^n)] \geq 1/\gamma(n)$$

2. **Approximate Hardness of Sampling.** For any oracle \mathcal{O} satisfying that for all large enough $n \in \mathbb{N}$,

$$\Pr_{\substack{C \leftarrow \mathcal{C}_n \\ x \leftarrow \{0,1\}^n}} \left[|\mathcal{O}(C, x) - \Pr_C[x]| \leq \frac{\Pr_C[x]}{p(n)} \right] \geq 1/\gamma(n) - \frac{1}{p(n)}$$

we have that $\text{P}^{\#\text{P}} \subseteq \text{BPP}^{\mathcal{O}}$. Here, $\Pr_C[x]$ denotes that probability the C outputs x .

While Definition 5.1 captures the hardness conjectures proposed in the quantum advantage literature, it is slightly inconvenient for cryptographic applications. There are a few reasons for this: the biggest one is that only adversaries that succeed for *all* large enough n will help solve $\#\text{P}$. The standard adversarial model for cryptography allows for adversaries that succeed on infinitely many n . Second, hardness is defined for strings sampled uniformly, whereas for our applications, it will be more suitable to have hardness defined for strings sampled from the output of the circuit itself. Finally, we only require hardness against all polynomial-sized quantum machines, whereas the definition above requires $\#\text{P}$ hardness of the approximation task.

We will therefore define and build from Definition 5.1 (along with the assumption that $\text{P}^{\#\text{P}} \not\subseteq \text{iQBQP}/\text{qpoly}$) a related but more “crypto-friendly” primitive that will simplify the implication to puzzles (Theorem 5.1).

Definition 5.2 (Native Approximation Hardness). *A family of (uniform) efficiently sampleable distributions $\mathcal{D} = \{\mathcal{D}_n\}_{n \in \mathbb{N}}$ over classical strings has native approximation hardness if there exists a polynomial p such that for all QPT $\mathcal{A} = \{\mathcal{A}_n\}_{n \in \mathbb{N}}$, every (non-uniform, quantum) advice ensemble $|\tau\rangle = \{|\tau_n\rangle\}_{n \in \mathbb{N}}$, and large enough $n \in \mathbb{N}$,*

$$\Pr_{x \leftarrow \mathcal{D}_n} \left[|\mathcal{A}(|\tau\rangle, x) - \Pr_{\mathcal{D}_n}[x]| \leq \frac{\Pr_{\mathcal{D}_n}[x]}{p(n)} \right] \leq 1 - \frac{1}{p(n)}$$

Next we will show (Theorem 5.2) that the existence of distribution families satisfying Definition 5.2 implies the existence of one-way puzzles (Definition 4.2). Finally we will show the reverse implication (Theorem 5.3), namely that the existence of one-way puzzles implies the existence of distribution families satisfying Definition 5.2.

5.2 Uniform Approximation Hardness Implies Native Approximation Hardness

Theorem 5.1. *If $\text{P}^{\#\text{P}}/\text{qpoly} \not\subseteq \text{ioBQP}/\text{qpoly}$ and there exists a family of distributions $\mathcal{C} = \{\mathcal{C}_n\}_{n \in \mathbb{N}}$ that satisfies Definition 5.1, then there exists a family of distributions $\mathcal{D} = \{\mathcal{D}_n\}_{n \in \mathbb{N}}$ that satisfies Definition 5.2.*

Proof. Let L be a PP-complete language. \mathcal{C} satisfies Definition 5.1, so there must exist exist polynomials q and γ along with oracle PPT $R_L^{(\cdot)}$ such that:

1. For all large enough $n \in \mathbb{N}$

$$\Pr_{\substack{C \leftarrow \mathcal{C}_n \\ x \leftarrow \{0,1\}^n}} [\Pr_C[x] \geq 1/(q(n) \cdot 2^n)] \geq 1/\gamma(n)$$

2. Let \mathcal{O} be any oracle such that for all large enough $n \in \mathbb{N}$

$$\Pr_{\substack{C \leftarrow \mathcal{C}_n \\ x \leftarrow \{0,1\}^n}} \left[|\mathcal{O}(C, x) - \Pr_C[x]| \leq \frac{\Pr_C[x]}{q(n)} \right] \geq \frac{1}{\gamma(n)} - \frac{1}{q(n)}$$

then for all $n \in \mathbb{N}$ and for all $x \in \{0, 1\}^n$, $\Pr[R_L^{\mathcal{O}}(x) = L(x)] \geq 1 - \text{negl}(n)$. We call the set of all such oracles \mathbb{O} .

Let t_R be a polynomial such that $R_L^{(\cdot)}$ runs in time $t_R(n)$.

Since L is PP-complete, $\text{BQP}^L/\text{qpoly} = \text{BQP}^{\text{PP}}/\text{qpoly} = \text{BQP}^{\#\text{P}}/\text{qpoly}$. Let $L' \in \text{BQP}^L/\text{qpoly}$ such that $L' \notin \text{ioBQP}/\text{qpoly}$. Then there exists an oracle QPT $\mathcal{M}^{(\cdot)} = \{\mathcal{M}_n^{(\cdot)}\}_{n \in \mathbb{N}}$ and (non-uniform, quantum) advice ensemble $|\sigma\rangle = \{|\sigma_n\rangle\}_{n \in \mathbb{N}}$ such that for all $n \in \mathbb{N}$, for all $x \in \{0, 1\}^n$,

$$\Pr[\mathcal{M}^L(|\sigma\rangle, x) = L'(x)] \geq 1 - \text{negl}(n)$$

We will drop the advice from the notation since it is always implicitly provided to \mathcal{M} . Additionally, let $t_{\mathcal{M}}$ be a polynomial such that $|\mathcal{M}_n| \leq t_{\mathcal{M}}(n)$.

Our overall proof strategy will be to show that (given an adversary that estimates probabilities for a carefully constructed distribution \mathcal{D}) we can replace the oracle to L in \mathcal{M}^L with an efficient quantum algorithm (with advice). This is accomplished by using $R_L^{(\cdot)}$ with an appropriately constructed oracle to decide L instead. We wish to define \mathcal{D}_n in such a way that any adversary that estimates the probabilities of $y \leftarrow \mathcal{D}_n$ will also allow us to estimate $\Pr_C[x]$ for $C \leftarrow \mathcal{C}_n$ and $x \leftarrow \{0, 1\}^n$. We may initially try setting \mathcal{D}_n to be the induced distribution on (C, x) . The adversary will in this case provide an estimate of $\Pr_{\mathcal{C}_n}[C] \cdot \Pr_C[x]$. To compute $\Pr_C[x]$ we therefore also need to compute $\Pr_{\mathcal{C}_n}[C]$. This is accomplished by having \mathcal{D}_n consist of two modes, one in which the output is (C, x) and one in which the output is C alone. This allows us to compute estimates for both $\Pr_{\mathcal{C}_n}[C] \cdot \Pr_C[x]$ and $\Pr_{\mathcal{C}_n}[C]$ and therefore estimate $\Pr_C[x]$.

The above strategy is still insufficient for instantiating an oracle for $R_L^{(\cdot)}$. This is because an adversary that breaks the security of \mathcal{D}_n may only do so for infinitely many values of n . The reduction $R_L^{(\cdot)}$, on the other hand, requires an oracle that breaks security for every large enough value of n . While it is tempting to think that instantiating $R_L^{(\cdot)}$ with an infinitely-often oracle would lead to a reduction that also succeeds infinitely often, this is not necessarily the case. $R_L^{(\cdot)}$ may query its oracle on a variety of input sizes, and only succeed if the oracle performs well on all of them. We must therefore be able to answer queries for all input sizes up to $t_R(m)$ when running $R_L^{(\cdot)}$ on inputs of size m . Since $R_L^{(\cdot)}$ will ultimately be queried by $\mathcal{M}_n(x)$ for $x \in \{0, 1\}^n$, m can be

as large as $t_{\mathcal{M}}(n)$. We therefore modify \mathcal{D}_n in such a way that we can use an adversary breaking its security to answer queries of all sizes upto $t_R(t_{\mathcal{M}}(n))$.

Let t be any polynomial such that $t(n) > t_R(t_{\mathcal{M}}(n))$ and let p be any polynomial such that $p(n) > nq^4(t(n)) \cdot (t(n))^3$. We define \mathcal{D}_n as follows:

- Sample $\ell \leftarrow [t(n)]$
- Sample $C \leftarrow \mathcal{C}_\ell$
- Sample $x \leftarrow C$
- Sample $b_{\text{mode}} \leftarrow \{0, 1\}$
- If $b_{\text{mode}} = 0$:
 - Output $(b_{\text{mode}}, C, 0^\ell)$
- Else:
 - Output (b_{mode}, C, x)

We will prove that \mathcal{D} satisfies Definition 5.2. Suppose for the sake of contradiction that this is not the case. Then there exists a QPT $\mathcal{A} = \{\mathcal{A}_n\}_{n \in \mathbb{N}}$ and (non-uniform, quantum) advice ensemble $|\tau\rangle = \{|\tau_n\rangle\}_{n \in \mathbb{N}}$ such that for infinitely many $n \in \mathbb{N}$,

$$\Pr_{y \leftarrow \mathcal{D}_n} \left[|\mathcal{A}(|\tau\rangle, y) - \Pr_{\mathcal{D}_n}[y]| \leq \frac{\Pr_{\mathcal{D}_n}[y]}{p(n)} \right] > 1 - \frac{1}{p(n)}$$

Fix any such adversary \mathcal{A} , any such advice ensemble $|\tau\rangle$, and any such large enough $n \in \mathbb{N}$. We will drop the advice from the notation since it is always implicitly provided to the adversary. We will now use \mathcal{A} to build a QPT \mathcal{B} that calls \mathcal{A} internally such that for all $x \in \{0, 1\}^n$, with high probability, $R^{\mathcal{B}}(x) = L(x)$. Specifically, we define \mathcal{B} as the algorithm that takes input (C, x) and returns $\mathcal{A}(1, C, x) / \mathcal{A}(0, C, 0^{|x|})$.

First we show that we can replace the oracle \mathcal{O} with \mathcal{B} for queries of size less than $t(n)$. We may view the above inequality as bounding the probability that $\mathcal{A}(y)$ is too far from $\Pr_{\mathcal{D}_n}[y]$, i.e.

$$\Pr_{y \leftarrow \mathcal{D}_n} \left[|\mathcal{A}(y) - \Pr_{\mathcal{D}_n}[y]| > \frac{\Pr_{\mathcal{D}_n}[y]}{p(n)} \right] < \frac{1}{p(n)}$$

We can split this error probability into cases indexed by the values of b_{mode} and ℓ . First consider when $b_{\text{mode}} = 0$. For all $\ell \in [t(n)]$

$$\Pr[b_{\text{mode}} = 0] \cdot \Pr[\ell] \cdot \Pr_{C \leftarrow \mathcal{C}_\ell} \left[\left| \mathcal{A}(0, C, 0^\ell) - \Pr_{\mathcal{D}_n}[0, C, 0^\ell] \right| > \frac{\Pr_{\mathcal{D}_n}[0, C, 0^\ell]}{p(n)} \right] < \frac{1}{p(n)}$$

which can be simplified to

$$\forall \ell \in [t(n)], \Pr_{C \leftarrow \mathcal{C}_\ell} \left[\left| 2t(n) \cdot \mathcal{A}(0, C, 0^\ell) - \Pr_{\mathcal{C}_\ell}[C] \right| > \frac{\Pr_{\mathcal{C}_\ell}[C]}{p(n)} \right] < \frac{2t(n)}{p(n)} \quad (1)$$

Similarly, consider when $b_{\text{mode}} = 1$. For all $\ell \in [t(n)]$

$$\Pr[b_{\text{mode}} = 1] \cdot \Pr[\ell] \cdot \Pr_{\substack{C \leftarrow \mathcal{C}_\ell \\ x \leftarrow C}} \left[\left| \mathcal{A}(1, C, x) - \Pr_{\mathcal{D}_n}[1, C, x] \right| > \frac{\Pr_{\mathcal{D}_n}[1, C, x]}{p(n)} \right] < \frac{1}{p(n)}$$

which can be simplified to

$$\forall \ell \in [t(n)], \Pr_{\substack{C \leftarrow \mathcal{C}_\ell \\ x \leftarrow C}} \left[\left| 2t(n) \cdot \frac{\mathcal{A}(1, C, x)}{\Pr_{\mathcal{C}_\ell}[C]} - \Pr_C[x] \right| > \frac{\Pr_C[x]}{p(n)} \right] < \frac{2t(n)}{p(n)} \quad (2)$$

Let $\delta := \sqrt{\frac{2t(n)}{p(n)}}$. We will now show for every large enough ℓ the existence of a “good” set of (C, x) that is sampled with high enough probability and for which $\mathcal{B}(C, x)$ gives a good estimate of $\Pr_C[x]$ with high probability. For all $\ell \in [t(n)]$ let \mathbb{C}_ℓ be defined as follows.

$$\mathbb{C}_\ell := \left\{ C \text{ s.t. } \Pr \left[\left| 2t(n) \cdot \mathcal{A}(0, C, 0^\ell) - \Pr_{\mathcal{C}_\ell}[C] \right| > \frac{\Pr_{\mathcal{C}_\ell}[C]}{p(n)} \right] > \delta \right\}$$

Intuitively, \mathbb{C}_ℓ is the set of C output by \mathcal{C}_ℓ such that with all but δ probability, $2t(n) \cdot \mathcal{A}(0, C, 0^\ell)$ is a good estimate for $\Pr_{\mathcal{C}_\ell}[C]$. By a Markov argument on (1), for all $\ell \in [t(n)]$

$$\Pr_{C \leftarrow \mathcal{C}_\ell} [C \in \mathbb{C}_\ell] \leq \delta$$

Similarly, for all $\ell \in [t(n)]$ let \mathbb{G}_ℓ be defined as follows.

$$\mathbb{G}_\ell := \left\{ (C, x) \text{ s.t. } \Pr \left[\left| 2t(n) \cdot \frac{\mathcal{A}(1, C, x)}{\Pr_{\mathcal{C}_\ell}[C]} - \Pr_C[x] \right| > \frac{\Pr_C[x]}{p(n)} \right] > \delta \right\}$$

Intuitively, \mathbb{G}_ℓ is the set of (C, x) where C is output by \mathcal{C}_ℓ and x is an ℓ -bit string such that with all but δ probability, $2t(n) \cdot \frac{\mathcal{A}(1, C, x)}{\Pr_{\mathcal{C}_\ell}[C]}$ is a good estimate for $\Pr_C[x]$. By a Markov argument on (2), for all $\ell \in [t(n)]$

$$\Pr_{\substack{C \leftarrow \mathcal{C}_\ell \\ x \leftarrow C}} [(C, x) \in \mathbb{G}_\ell] \leq \delta$$

Claim 5.1. For all large enough $\ell \in [t(n)]$, let $\mathbb{B}_\ell := \{(C, x) \text{ s.t. } C \notin \mathbb{C}_\ell \wedge (C, x) \notin \mathbb{G}_\ell\}$

- For all $(C, x) \in \mathbb{B}_\ell$, $\Pr \left[|\mathcal{B}(C, x) - \Pr_C(x)| \leq \frac{3\Pr_C(x)}{p(n)} \right] \geq 1 - 2\delta$
- $\Pr_{\substack{C \leftarrow \mathcal{C}_\ell \\ x \leftarrow \{0,1\}^\ell}} [(C, x) \in \mathbb{B}_\ell] \geq 1/\gamma(\ell) - 2\delta \cdot q(\ell)$

Proof. For the first part of the claim, note that by the definitions of \mathbb{C}_ℓ , for all $(C, x) \in \mathbb{B}_\ell$, with probability atleast $(1 - \delta)$

$$\left| 2t(n) \cdot \mathcal{A}(0, C, 0^\ell) - \Pr_{\mathcal{C}_\ell}[C] \right| > \frac{\Pr_{\mathcal{C}_\ell}[C]}{p(n)}$$

which can be rearranged as

$$\frac{\mathcal{A}(0, C, 0^\ell)}{(1 - 1/p(n))} < \frac{\Pr_{\mathcal{C}_\ell}[C]}{2t(n)} < \frac{\mathcal{A}(0, C, 0^\ell)}{(1 + 1/p(n))}$$

Additionally, by the definitions of \mathbb{G}_ℓ , for all $(C, x) \in \mathbb{B}_\ell$, with probability atleast $(1 - \delta)$

$$\left| 2t(n) \cdot \frac{\mathcal{A}(1, C, x)}{\Pr_{\mathcal{C}_\ell}[C]} - \Pr_C[x] \right| > \frac{\Pr_C[x]}{p(n)}$$

which can similarly be rewritten as

$$\Pr_C[x] \cdot (1 - 1/p(n)) < 2t(n) \cdot \frac{\mathcal{A}(1, C, x)}{\Pr_{\mathcal{C}_\ell}[C]} < \Pr_C[x] \cdot (1 + 1/p(n))$$

Both events occur simultaneously with probability at least $1 - 2\delta$, in which case we may apply the bounds for $\frac{\Pr_{\mathcal{C}_\ell}[C]}{2t(n)}$ to the previous inequality.

$$\Pr_C[x] \cdot (1 - 1/p(n))^2 < \frac{\mathcal{A}(1, C, x)}{\mathcal{A}(0, C, 0^\ell)} < \Pr_C[x] \cdot (1 + 1/p(n))^2$$

which for large enough n gives

$$\left| \Pr_C[x] - \frac{\mathcal{A}(1, C, x)}{\mathcal{A}(0, C, 0^\ell)} \right| < \frac{3\Pr_C[x]}{p(n)}$$

which concludes the proof of part 1.

For the second part of the claim, first note that since $\Pr_{C \leftarrow \mathcal{C}_\ell}[(C, x) \in \mathbb{G}_\ell] \leq \delta$ and $\Pr_{C \leftarrow \mathcal{C}_\ell}[C \in \mathbb{C}_\ell] \leq \delta$

$$\Pr_{C \leftarrow \mathcal{C}_\ell}[C \in \overline{\mathbb{B}}_\ell] \leq 2\delta$$

where $\overline{\mathbb{B}}_\ell$ represents the complement of \mathbb{B}_ℓ . Let $\mathbb{A}_\ell := \{(C, x) : \Pr_C[x] \geq 1/(q(\ell) \cdot 2^\ell)\}$. By the anticoncentration property we know that for large enough $\ell \in [t(n)]$

$$\Pr_{\substack{C \leftarrow \mathcal{C}_\ell \\ x \leftarrow \{0,1\}^\ell}}[(C, x) \in \mathbb{A}] \geq 1/\gamma(\ell)$$

We aim to bound $\Pr_{\substack{C \leftarrow \mathcal{C}_\ell \\ x \leftarrow \{0,1\}^\ell}}[(C, x) \in \mathbb{A}_\ell \cap \overline{\mathbb{B}}_\ell]$ as follows

$$\begin{aligned} \Pr_{\substack{C \leftarrow \mathcal{C}_\ell \\ x \leftarrow \{0,1\}^\ell}}[(C, x) \in \mathbb{A}_\ell \cap \overline{\mathbb{B}}_\ell] &= \sum_{(C,x) \in \mathbb{A}_\ell \cap \overline{\mathbb{B}}_\ell} \Pr_{\mathcal{C}_\ell}[C] \cdot 1/2^\ell \\ &\leq \sum_{(C,x) \in \mathbb{A}_\ell \cap \overline{\mathbb{B}}_\ell} \Pr_{\mathcal{C}_\ell}[C] \cdot q(\ell) \cdot \Pr_C[x] \\ &= q(\ell) \cdot \Pr_{\substack{C \leftarrow \mathcal{C}_\ell \\ x \leftarrow C}}[(C, x) \in \mathbb{A}_\ell \cap \overline{\mathbb{B}}_\ell] \\ &\leq q(\ell) \cdot \Pr_{\substack{C \leftarrow \mathcal{C}_\ell \\ x \leftarrow C}}[(C, x) \in \overline{\mathbb{B}}_\ell] \\ &\leq 2\delta \cdot q(\ell) \end{aligned}$$

where the first step follows from the definition of \mathbb{A} and the last step follows from Claim 5.1. We can now bound $\Pr_{\substack{C \leftarrow \mathcal{C}_\ell \\ x \leftarrow \{0,1\}^\ell}}[(C, x) \in \mathbb{B}_\ell]$ as follows

$$\begin{aligned} \Pr_{\substack{C \leftarrow \mathcal{C}_\ell \\ x \leftarrow \{0,1\}^\ell}}[(C, x) \in \mathbb{B}_\ell] &\geq \Pr_{\substack{C \leftarrow \mathcal{C}_\ell \\ x \leftarrow \{0,1\}^\ell}}[(C, x) \in \mathbb{A}_\ell \cap \mathbb{B}_\ell] \\ &= \Pr_{\substack{C \leftarrow \mathcal{C}_\ell \\ x \leftarrow \{0,1\}^\ell}}[(C, x) \in \mathbb{A}_\ell] - \Pr_{\substack{C \leftarrow \mathcal{C}_\ell \\ x \leftarrow \{0,1\}^\ell}}[(C, x) \in \mathbb{A}_\ell \cap \overline{\mathbb{B}}_\ell] \\ &\geq 1/\gamma(\ell) - 2\delta \cdot q(\ell) \end{aligned}$$

which concludes the proof. \square

Claim 5.2. For all m such that $t_R(m) \leq t(n)$, for all $x \in \{0, 1\}^m$

$$\Pr[R_L^{\mathcal{B}}(x) = L(x)] \geq 1 - \text{negl}(m) - 2\delta \cdot t(n)$$

Proof. Since $R_L^{(\cdot)}(x)$ runs in time atmost $t_R(|x|)$, the size of the longest query to made is atmost $t_R(m) \leq t(n)$. Additionally, the number of queries made is also atmost $t_R(m) \leq t(n)$. By Claim 5.1 and by noting that $p(n) > nq^4(t(n)) \cdot (t(n))^3$ and $\delta = \sqrt{\frac{2t(n)}{p(n)}}$, for every $\ell \leq t(n)$, there exists a set \mathbb{B}_ℓ

- For all $(C, x) \in \mathbb{B}_\ell$, $\Pr \left[|\mathcal{B}(C, x) - \Pr_C(x)| \leq \frac{\Pr_C(x)}{q(t(n))} \right] \geq 1 - 2\delta$
- $\Pr_{\substack{C \leftarrow \mathcal{C}_\ell \\ x \leftarrow \{0,1\}^\ell}} [(C, x) \in \mathbb{B}_\ell] \geq 1/\gamma(\ell) - 1/q(t(n))$

Let \mathcal{O}' be the set of oracles such that for every $\ell \leq t(n)$, for all $(C, x) \in \mathbb{B}_\ell$,

$$|\mathcal{O}(C, x) - \Pr_C(x)| \leq \frac{\Pr_C(x)}{q(\ell)}$$

Since $\ell \leq t(n)$ it is easy to see that $\mathcal{O}' \subseteq \mathcal{O}$. $R_L^{(\cdot)}(x)$ succeeds with probability $1 - \text{negl}(m)$ for every $\mathcal{O} \in \mathcal{O}'$, so it succeeds with probability atleast $1 - \text{negl}(m)$ given any random distribution over oracles in \mathcal{O}' . We will now show that with high probability, \mathcal{B} is such a distribution over oracles.

Assume WLOG that $R_L^{\mathcal{B}}(x)$ queries any string atmost once during its execution. We only need to consider queries of length atmost $t(n)$ since no query is longer than $t(n)$. If for every query $(C, x) \in \mathbb{B}_\ell$ made to \mathcal{B} where $x \in \{0, 1\}^\ell$, it was the case that $|\mathcal{B}(C, x) - \Pr_C(x)| \leq \frac{\Pr_C(x)}{q(n)}$ then the outputs of \mathcal{B} are distributed according to some distribution over \mathcal{O}' . For each query this occurs independently with probability atleast $1 - 2\delta$ and there are at most $t(n)$ queries, so this event occurs with probability atleast $1 - 2\delta \cdot t(n)$. Therefore, $\Pr[R_L^{\mathcal{B}}(x) = L(x)] \geq 1 - \text{negl}(m) - 2\delta \cdot t(n)$.¹¹ \square

Noting that $p(n) > nq^4(t(n)) \cdot (t(n))^3$ and n is large enough, we get $2\delta \cdot t(n) \leq 1/3$. Therefore, by repeating in parallel and taking the majority outcome we obtain an oracle PPT $S^{(\cdot)}$ where for all m such that $t_R(m) \leq t(n)$, for all $y \in \{0, 1\}^m$, $\Pr[S^{\mathcal{B}}(y) = L(y)] \geq 1 - \text{negl}(m)$. Finally, we will use $\mathcal{M}^{S^{\mathcal{B}}}$ to decide L' on strings of length n . For $x \in \{0, 1\}^n$, $\mathcal{M}_n^{S^{\mathcal{B}}}(x)$ can only query strings of length at most $t_{\mathcal{M}}(n)$. Since $t_R(t_{\mathcal{M}}(n)) \leq t(n)$, this means that $S^{\mathcal{B}}$ will correctly answer the queries with all but negligible probability. As a result, for all $x \in \{0, 1\}^n$, $\Pr[\mathcal{M}_n^{S^{\mathcal{B}}}(x) = L'(x)] \geq 1 - \text{negl}(n)$.

Finally, since $S^{\mathcal{B}}$ is a polynomial size quantum circuit with quantum advice, $\mathcal{M}^{S^{\mathcal{B}}}$ may also be expressed as a polynomial size quantum circuit with quantum advice that decides L' for infinitely many input lengths. This contradicts the assumption that $L' \notin \text{ioBQP}/\text{qpoly}$ and concludes the proof of the theorem. \square

5.3 Native Approximation Hardness Implies One-Way Puzzles

Theorem 5.2. The existence of families of distributions that satisfy Definition 5.2 implies the existence of one-way puzzles (Definition 4.2)

¹¹A slightly different version of Definition 5.1 requires the oracle to approximate probabilities upto arbitrary relative error ϵ (given $1^{1/\epsilon}$ as input) and with probability $1/\gamma(\ell) - \delta$ (given $1^{1/\delta}$ as input) over the randomness of the input. In this case, note that since the size of the largest query is atmost $t(n)$, ϵ and δ are atleast $1/t(n)$. \mathcal{B} can therefore still be used to answer such oracle queries by setting $p(n)$ to be large enough to achieve relative error $1/t(n)$ with probability $1/\gamma(\ell) - 1/t(n)$ over the randomness of the input, and Theorem 5.1 will be unaffected by the change in the definition.

Proof. Let $\mathcal{D} = \{\mathcal{D}_n\}_{n \in \mathbb{N}}$ be a family of distributions that satisfies Definition 5.2. Therefore there exists a polynomial q such that for all QPT $\mathcal{A} = \{\mathcal{A}_n\}_{n \in \mathbb{N}}$, every (non-uniform, quantum) advice ensemble $|\tau\rangle = \{|\tau_n\rangle\}_{n \in \mathbb{N}}$, and large enough $n \in \mathbb{N}$,

$$\Pr_{x \leftarrow \mathcal{D}_n} \left[|\mathcal{A}(|\tau\rangle, x) - \Pr_{\mathcal{D}_n}[x]| \leq \frac{\Pr_{\mathcal{D}_n}[x]}{q(n)} \right] \leq 1 - \frac{1}{q(n)}$$

We may also assume w.l.o.g. that the outputs of \mathcal{D}_n are of n bits. Define $\text{Samp}(1^n)$ as follows:

- Sample $i \leftarrow [0, n-1]$
- Sample $x \leftarrow \mathcal{D}_n$
- Output puzzle $x_{1\dots i}$ and key x_{i+1}

Let $p(n)$ be any polynomial such that $p(n) > n^6 q(n)^3$. We will prove that $\text{Samp}(1^n)$ is a $1/p(n)$ -distributional one-way puzzle. Since Theorem 4.3 shows that distributional one-way puzzles can be amplified to obtain (strong) one-way puzzles, this suffices to prove the theorem.

Suppose for the sake of contradiction that $\text{Samp}(1^n)$ is *not* a $1/p(n)$ -distributional one-way puzzle. Then there exists a QPT $\mathcal{A} = \{\mathcal{A}_n\}_{n \in \mathbb{N}}$ and (non-uniform, quantum) advice ensemble $|\tau\rangle = \{|\tau_n\rangle\}_{n \in \mathbb{N}}$ such that for infinitely many $n \in \mathbb{N}$,

$$\text{SD}(\{x_{1\dots i}, x_{i+1}\}, \{x_{1\dots i}, \mathcal{A}(|\tau\rangle, x_{1\dots i})\}) \leq 1/p(n)$$

where $x_{1\dots i}, x_{i+1} \leftarrow \text{Samp}(1^n)$. Fix any such adversary \mathcal{A} , any such advice ensemble $|\tau\rangle$, and any such large enough $n \in \mathbb{N}$. We will drop the advice from the notation since it is always implicitly provided to the adversary.

First we define some useful terms. For $i \in [0, n-1]$, $z \in \{0, 1\}^i$, $z' \in \{0, 1\}^{n-i}$ and $b \in \{0, 1\}$:

- Define $p_z := \Pr_{x \leftarrow \mathcal{D}_n}[x_{1\dots i} = z]$, i.e. the probability that the first i bits of x sampled from \mathcal{D}_n match z .
- Define $p_{z'|z} := \Pr_{x \leftarrow \mathcal{D}_n}[x_{i+1\dots n} = z' \mid x_{1\dots i} = z]$ i.e. the probability that the last $n-i$ bits of x sampled from \mathcal{D}_n match z' conditioned on the first i bits of x matching z .
- Similarly define $p_{b|z} := \Pr_{x \leftarrow \mathcal{D}_n}[x_{i+1} = b \mid x_{1\dots i} = z]$ i.e. the probability that the $i+1$ -th bit of x sampled from \mathcal{D}_n is b conditioned on the first i bits of x matching z .
- Define $\tilde{p}_{b|z} := \Pr[\mathcal{A}(z) = b]$

We will use the adversary to build \mathcal{A}' that contradicts the security of \mathcal{C} . We will first build an estimator \mathcal{E}_b that takes input z and estimates $\tilde{p}_{b|z}$.

For $i \in [0, n-1]$, $z \in \{0, 1\}^i$, $b \in \{0, 1\}$ define the algorithm $\mathcal{E}_b(z)$ as:

- For $j = 1$ to $16n^7 q(n)^4$:
 - $X_j \leftarrow \mathbb{I}\{\mathcal{A}(z) = b\}$
- Return $\sum_j X_j / 16n^7 q(n)^4$

where \mathbb{I} is an indicator function. We now define the algorithm \mathcal{A}' that takes input x and computes an approximation to p_x . Define $\mathcal{A}'(x)$ as

- Return $\prod_{j \in [0, n-1]} \mathcal{E}_{x_{j+1}}(x_{1\dots j})$

Observe that $p_x = \prod_{j=0}^{n-1} p_{x_{j+1}|x_{1\dots j}}$. On a high level, to approximate p_x it therefore suffices to approximate $p_{x_{j+1}|x_{1\dots j}}$ for every j and multiply the approximations. We cannot directly estimate $p_{x_{j+1}|x_{1\dots j}}$ however \mathcal{A}' uses $\mathcal{E}_{x_{j+1}}(x_{1\dots j})$ to estimate $\tilde{p}_{x_{j+1}|x_{1\dots j}}$ instead. We will use the fact that \mathcal{A} distributionally inverts Samp to argue that this is sufficient to (on average) obtain an approximation to p_x .

First, we show that $\mathcal{E}_b(z)$ is a good approximation of $\tilde{p}_{b|z}$ with high probability.

Claim 5.3. For all $i \in [0, n-1]$, $z \in \{0, 1\}^i$, $b \in \{0, 1\}$

$$\Pr \left[|\mathcal{E}_b(z) - \tilde{p}_{b|z}| \leq \frac{1}{4n^3q(n)^2} \right] \geq 1 - 2e^{-n}$$

Proof. Follows from setting $\delta = \sqrt{n}$ in the additive Chernoff bound (Theorem 4.1). \square

The estimate we obtain has an inverse polynomial additive error. If the value being estimated is too small, then this leads to a large relative error in the estimate. We define the set \mathbb{B} as containing all x such that $p_{x_{j+1}|x_{1\dots j}}$ is at least $1/n^2q(n)$ for every index j . Formally

$$\mathbb{B} := \left\{ x \in \{0, 1\}^n \text{ s.t. } \forall j \in [0, n-1], p_{x_{j+1}|x_{1\dots j}} \geq 1/n^2q(n) \right\}$$

Intuitively, \mathbb{B} contains strings for which the additive error induced by Claim 5.3 only leads to a small relative error. Next we show that with high probability x sampled from \mathcal{D}_n is in \mathbb{B} .

Claim 5.4.

$$\Pr_{x \leftarrow \mathcal{D}_n} [x \in \mathbb{B}] \geq 1 - 2/nq(n)$$

Proof. For each index j we define the set \mathbb{S}_j as the set of strings x such that $p_{x_{j+1}|x_{1\dots j}}$ is less than $1/n^2q(n)$. Intuitively, if a string x is not in \mathbb{B} , it must be in \mathbb{S}_j for some j . Therefore, we can prove the claim by bounding the probability of sampling a string in \mathbb{S}_j for every index j . Formally, for all $j \in [0, n-1]$

$$\mathbb{S}_j := \left\{ x \text{ s.t. } p_{x_{j+1}|x_{1\dots j}} < 1/n^2q(n) \right\}$$

Now, since $p_x = \prod_{j=0}^{n-1} p_{x_{j+1}|x_{1\dots j}}$, the probability of sampling x depends on $p_{x_{j+1}|x_{1\dots j}}$ for every index j . If $x \in \mathbb{S}_j$ then $p_{x_{j+1}|x_{1\dots j}}$ is small, which allows us to bound the probability of sampling such an x , i.e.

$$\begin{aligned} \Pr_{x \leftarrow \mathcal{D}_n} [x \in \mathbb{S}_j] &= \sum_{x \in \mathbb{S}_j} \Pr_{\mathcal{D}_n} [x] \\ &= \sum_{x \in \mathbb{S}_j} p_{x_{1\dots j}} \cdot p_{x_{j+1}|x_{1\dots j+1}} \cdot p_{x_{j+2\dots n}|x_{1\dots j+1}} \\ &< \sum_{x \in \mathbb{S}_j} p_{x_{1\dots j}} \cdot 1/n^2q(n) \cdot p_{x_{j+2\dots n}|x_{1\dots j+1}} \\ &\leq \sum_{x \in \{0,1\}^n} p_{x_{1\dots j}} \cdot 1/n^2q(n) \cdot p_{x_{j+2\dots n}|x_{1\dots j+1}} \\ &= 2/n^2q(n) \end{aligned}$$

By a union bound, this implies

$$\Pr_{x \leftarrow \mathcal{D}_n} [\exists j \text{ s.t. } x \in \mathbb{S}_j] \leq 2/nq(n)$$

Since for all $x \notin \mathbb{B}, \exists j$ s.t. $x \in \mathbb{S}_j$

$$\Pr_{x \leftarrow \mathcal{D}_n} [x \in \mathbb{B}] \geq 1 - 2/nq(n)$$

which concludes the proof of the claim. \square

Next we define the set of strings x such that for all j a good estimate of $\tilde{p}_{x_{j+1}|x_{1\dots j}}$ is also a good estimate of $p_{x_{j+1}|x_{1\dots j}}$. Define the set \mathbb{D} as follows.

$$\mathbb{D} := \left\{ x \in \{0, 1\}^n \text{ s.t. } \forall j \in [0, n-1], \left| \tilde{p}_{1|x_{1\dots j}} - p_{1|x_{1\dots j}} \right| \leq 1/4n^3q(n)^2 \right\}$$

We now use the fact that \mathcal{A} is a distributional inverter to show that with high probability x sampled from \mathcal{D}_n is in \mathbb{D} .

Claim 5.5.

$$\Pr_{x \leftarrow \mathcal{D}_n} [x \in \mathbb{D}] \geq 1 - 4n^5q(n)^2/p(n)$$

Proof. Recall that

$$\text{SD}(\{x_{1\dots i}, x_{i+1}\}, \{x_{1\dots i}, \mathcal{A}(x_{1\dots i})\}) \leq 1/p(n)$$

where $(x_{1\dots i}, x_{i+1}) \leftarrow \text{Samp}(1^n)$. Since i is chosen uniformly, we may split the statistical distance into terms for each value of i , i.e. for $x \leftarrow \mathcal{D}_n$

$$\begin{aligned} \sum_{j \in [0, n-1]} \Pr_{i \leftarrow [0, n-1]} [i = j] \cdot \text{SD}(\{x_{1\dots j}, x_{j+1}\}, \{x_{1\dots j}, \mathcal{A}(x_{1\dots j})\}) &\leq 1/p(n) \\ \implies \forall j \in [0, n-1], \text{SD}(\{x_{1\dots j}, x_{j+1}\}, \{x_{1\dots j}, \mathcal{A}(x_{1\dots j})\}) &\leq n/p(n) \end{aligned}$$

Expanding the statistical distance term, we may rewrite the expression as

$$\forall j, \mathbb{E}_{x \leftarrow \mathcal{D}_n} \left[\left| \tilde{p}_{1|x_{1\dots j}} - p_{1|x_{1\dots j}} \right| \right] \leq n/p(n)$$

By a Markov argument,

$$\forall j, \Pr_{x \leftarrow \mathcal{D}_n} \left[\left| \tilde{p}_{1|x_{1\dots j}} - p_{1|x_{1\dots j}} \right| \geq \frac{1}{4n^3q(n)^2} \right] \leq 4n^4q(n)^2/p(n)$$

By a union bound

$$\Pr_{x \leftarrow \mathcal{D}_n} \left[\forall j, \left| \tilde{p}_{1|x_{1\dots j}} - p_{1|x_{1\dots j}} \right| \geq \frac{1}{4n^3q(n)^2} \right] \leq 4n^5q(n)^2/p(n)$$

which by the definition of \mathbb{D} implies

$$\Pr_{x \leftarrow \mathcal{D}_n} [x \in \mathbb{D}] \geq 1 - 4n^5q(n)^2/p(n)$$

concluding the proof of the claim. \square

To complete the proof we will first show that for all $x \in \mathbb{B} \cap \mathbb{D}$, $\mathcal{A}'(x)$ is a good approximation to p_x with high probability. Then we show that with high probability $x \in \mathbb{B} \cap \mathbb{D}$ when x is sampled from \mathcal{D}_n . To show the former, we will need the following lemma. The lemma shows that for some real values $\{a_i, b_i\}_{i \in [0, n-1]}$, if b_i is a good (i.e. low relative error) estimate of a_i for all i then $\prod_i b_i$ is a good estimate of $\prod_i a_i$.

Lemma 5.1. For all $i \in [0, n - 1]$, let a_i, b_i, δ be values such that

- $0 < a_i \leq 1$
- $0 \leq \delta < 1/n$
- $\frac{|a_i - b_i|}{a_i} \leq \delta$

Then

$$\frac{|\prod_i a_i - \prod_i b_i|}{\prod_i a_i} \leq 2n\delta$$

Proof. Let $a := \prod_i a_i$ and $b := \prod_i b_i$. Then

$$b = \prod_i b_i = \prod_i a_i \cdot \frac{b_i}{a_i} = a \cdot \prod_i \left(1 + \frac{a_i - b_i}{a_i}\right)$$

Since $\left|\frac{a_i - b_i}{a_i}\right| \leq \delta \leq 1/n$ and $a > 0$

$$\begin{aligned} a(1 - \delta)^n &\leq b \leq a(1 + \delta)^n \\ \implies a(1 - n\delta) &\leq b \leq a/(1 - n\delta) \end{aligned}$$

Therefore,

$$\begin{aligned} \left|1 - \frac{b}{a}\right| &\leq \max\left(n\delta, \frac{n\delta}{1 - n\delta}\right) \\ &\leq 2n\delta \end{aligned}$$

which concludes the proof of the lemma. \square

Now we can show that for all $x \in \mathbb{B} \cap \mathbb{D}$, with high probability $\mathcal{A}'(x)$ is a good approximation of p_x .

Claim 5.6. $\forall x \in \mathbb{B} \cap \mathbb{D}$

$$\Pr \left[\left| \mathcal{A}'(x) - \Pr_{\mathcal{D}_n}[x] \right| \leq \frac{\Pr_{\mathcal{D}_n}[x]}{q(n)} \right] \geq 1 - 2ne^{-n}$$

Proof. We start by noting that for all j , $\mathcal{E}_{x_{j+1}}(x_{1..j})$ is close $\tilde{p}_{x_{j+1}|x_{1..j}}$ with high probability. Formally, by Claim 5.3, for all $j \in [0, n - 1]$

$$\Pr \left[\left| \mathcal{E}_{x_{j+1}}(x_{1..j}) - \tilde{p}_{x_{j+1}|x_{1..j}} \right| \leq \frac{1}{4n^3 q(n)^2} \right] \geq 1 - 2e^{-n}$$

By a union bound

$$\Pr \left[\forall j \in [0, n - 1], \left| \mathcal{E}_{x_{j+1}}(x_{1..j}) - \tilde{p}_{x_{j+1}|x_{1..j}} \right| \leq \frac{1}{4n^3 q(n)^2} \right] \geq 1 - 2ne^{-n}$$

Since $x \in \mathbb{D}$, $\tilde{p}_{x_{j+1}|x_{1..j}}$ is close to $p_{x_{j+1}|x_{1..j}}$. Formally, $\forall j \in [0, n - 1]$

$$\left| p_{x_{j+1}|x_{1..j}} - \tilde{p}_{x_{j+1}|x_{1..j}} \right| \leq \frac{1}{4n^3 q(n)^2}$$

By the triangle inequality, this shows that $\mathcal{E}_{x_{j+1}}(x_{1\dots j})$ is close $p_{x_{j+1}|x_{1\dots j}}$ with high probability.

$$\Pr \left[\forall j \in [0, n-1], \left| \mathcal{E}_{x_{j+1}}(x_{1\dots j}) - p_{x_{j+1}|x_{1\dots j}} \right| \leq \frac{1}{2n^3q(n)^2} \right] \geq 1 - 2ne^{-n}$$

This gives us a bound on the additive error. To obtain a bound on relative error, we note that for all $x \in \mathbb{B}$, $\forall j \in [0, n-1]$, $p_{x_{j+1}|x_{1\dots j}} \geq 1/n^2q(n)$. Therefore we can divide by $p_{x_{j+1}|x_{1\dots j}}$

$$\Pr \left[\forall j \in [0, n-1], \frac{\left| \mathcal{E}_{x_{j+1}}(x_{1\dots j}) - p_{x_{j+1}|x_{1\dots j}} \right|}{p_{x_{j+1}|x_{1\dots j}}} \leq \frac{1}{2nq(n)} \right] \geq 1 - 2ne^{-n}$$

Finally we use Lemma 5.1 to show that the product of good estimates for $p_{x_{j+1}|x_{1\dots j}}$ is a good estimate for p_x . For all $j \in [0, n-1]$, let $a_j := p_{x_{j+1}|x_{1\dots j}}$ and $b_j := \mathcal{E}_{x_{j+1}}(x_{1\dots j})$. Let $\delta := 1/2nq(n)$. Note that $\prod_j a_j = \prod_j p_{x_{j+1}|x_{1\dots j}} = p_x$ and $\prod_j b_j = \mathcal{A}'(x)$. Then applying Lemma 5.1 to the above

$$\Pr \left[\frac{|\mathcal{A}'(x) - p_x|}{p_x} \leq 1/q(n) \right] \geq 1 - 2ne^{-n}$$

which after rearranging gives

$$\Pr \left[|\mathcal{A}'(x) - p_x| \leq p_x/q(n) \right] \geq 1 - 2ne^{-n}$$

concluding the proof of the claim. \square

Finally, combining Claim 5.4 and Claim 5.5 we can show that

$$\Pr_{x \leftarrow \mathcal{D}_n} [x \in \mathbb{B} \cap \mathbb{D}] \geq 1 - 4n^5q(n)^2/p(n) - 2/nq(n)$$

By Claim 5.6

$$\begin{aligned} \Pr_{x \leftarrow \mathcal{D}_n} \left[\left| \mathcal{A}'(x) - \Pr_{\mathcal{D}_n}[x] \right| \leq \frac{\Pr_{\mathcal{D}_n}[x]}{q(n)} \right] &\geq \Pr_{x \leftarrow \mathcal{D}_n} [x \in \mathbb{B} \cap \mathbb{D}] \cdot (1 - 2e^{-n}) \\ &\geq (1 - 4n^5q(n)^2/p(n) - 2/nq(n)) (1 - 2e^{-n}) \end{aligned}$$

Since $p(n) > n^6q(n)^3$ and n is large enough,

$$\begin{aligned} \Pr_{x \leftarrow \mathcal{D}_n} \left[\left| \mathcal{A}'(x) - \Pr_{\mathcal{D}_n}[x] \right| \leq \frac{\Pr_{\mathcal{D}_n}[x]}{p(n)} \right] &\geq (1 - 4n^5q(n)^2/p(n) - 2/nq(n)) (1 - 2e^{-n}) \\ &> (1 - 4/nq(n) - 2/nq(n)) (1 - 2e^{-n}) \\ &> (1 - 1/q(n)) \end{aligned}$$

which contradicts Definition 5.2, concluding the proof of the theorem. \square

5.4 One-Way Puzzles Imply Native Approximation Hardness

Theorem 5.3. *The existence of distributional one-way puzzles (Definition 4.3) implies the existence of families of distributions that satisfy Definition 5.2.*

Let $\text{Samp}(1^n)$ be a $1/q(n)$ -distributional one-way puzzle for some polynomial q that samples n bit puzzles and n bit keys. Let $\mathcal{D} = \{\mathcal{D}_n\}_{n \in \mathbb{N}}$ be a family of distributions where \mathcal{D}_n is defined as follows:

1. $(s, k) \leftarrow \text{Samp}(1^n)$
2. $i \leftarrow [0, n - 1]$
3. $x \leftarrow s \| k_{1\dots i}$
4. $b_{\text{mode}} \leftarrow \{0, 1\}$
5. If $b_{\text{mode}} = 0$, then $\beta \leftarrow \{0, 1\}$. If $b_{\text{mode}} = 1$, then $\beta \leftarrow k_{i+1}$
6. Output (x, β)

Let p be a polynomial such that $p(n) > 5nq(n)$. We will prove that \mathcal{D} satisfies Definition 5.2, i.e. for all QPT $\mathcal{A} = \{\mathcal{A}_n\}_{n \in \mathbb{N}}$, every (non-uniform, quantum) advice ensemble $|\tau\rangle = \{|\tau_n\rangle\}_{n \in \mathbb{N}}$, and large enough $n \in \mathbb{N}$,

$$\Pr_{x \leftarrow \mathcal{D}_n} \left[|\mathcal{A}(|\tau\rangle, x) - \Pr_{\mathcal{D}_n}[x]| \leq \frac{\Pr_{\mathcal{D}_n}[x]}{p(n)} \right] \leq 1 - \frac{1}{p(n)}$$

Note that this suffices to prove the theorem.

Suppose for the sake of contradiction that \mathcal{D} does not satisfy Definition 5.2. Therefore there exists a QPT $\mathcal{A} = \{\mathcal{A}_n\}_{n \in \mathbb{N}}$ and (non-uniform, quantum) advice ensemble $|\tau\rangle = \{|\tau_n\rangle\}_{n \in \mathbb{N}}$ such that for infinitely many $n \in \mathbb{N}$,

$$\Pr_{x \leftarrow \mathcal{D}_n} \left[|\mathcal{A}(|\tau\rangle, x) - \Pr_{\mathcal{D}_n}[x]| \leq \frac{\Pr_{\mathcal{D}_n}[x]}{p(n)} \right] > 1 - 1/p(n).$$

Fix any such adversary \mathcal{A} , any such advice ensemble $|\tau\rangle$, and any such large enough $n \in \mathbb{N}$. We will drop the advice from the notation since it is always implicitly provided to the adversary. We will show that we can use \mathcal{A} to contradict one-wayness of Samp.

First we define some useful terms. For all $j \in [0, n - 1]$, $s \in \{0, 1\}^n$, $z \in \{0, 1\}^j$, and $b \in \{0, 1\}$ we define:

- Define $p_s := \Pr_{s', k \leftarrow \text{Samp}(1^n)}[s = s']$ i.e. the probability that $\text{Samp}(1^n)$ samples puzzle s .
- Define $p_{b|sz} := \Pr_{s', k \leftarrow \text{Samp}(1^n)}[k_{i+1} = b \mid s' = s \wedge k_{1\dots i} = z]$ i.e. the probability that the $i+1$ -th bit of key k sampled by $\text{Samp}(1^n)$ equals b conditioned on $\text{Samp}(1^n)$ sampling puzzle s and the first i bits of k matching z .

For all $s \in \{0, 1\}^n$, $j \in [0, n - 1]$, $k \in \{0, 1\}^n$, define $\mathcal{S}(s, k_{1\dots j}) :=$

1. $\tilde{a}_1 \leftarrow \mathcal{A}(s \| k_{1\dots j} \| 1)$
2. $\tilde{a}_0 \leftarrow \mathcal{A}(s \| k_{1\dots j} \| 0)$
3. $\pi := \frac{3\tilde{a}_1 - \tilde{a}_0}{2(\tilde{a}_1 + \tilde{a}_0)}$
If $\pi > 1$, set $\pi \leftarrow 1$
If $\pi < 0$, set $\pi \leftarrow 0$
4. Sample $k_{j+1} \leftarrow \text{Bern}(\pi)$, return k_{j+1} .

Intuitively, $\mathcal{S}(s, k_{1\dots j})$ aims to sample from the distribution on k_{j+1} induced by sampling from $\text{Samp}(1^n)$ conditioned on s and $k_{1\dots j}$, i.e. to sample k_{j+1} with probability $p_{k_{j+1}|sk_{1\dots j}}$. Define $\tilde{p}_{b|sz} := \Pr[\mathcal{S}(s, z) = b]$.

For all $s \in \{0, 1\}^n$, define $\mathcal{A}'(s) :=$

1. For $j = 0$ to $n - 1$:
 - $k_{j+1} \leftarrow \mathcal{S}(s, k_{1..j})$ (i.e. if $j = 0$, $k_1 \leftarrow \mathcal{S}(s)$)
2. Return k .

Intuitively, $\mathcal{A}'(s)$ aims to use the sampler \mathcal{S} to sample bit-by-bit from the distribution induced on keys k sampled by $\text{Samp}(1^n)$ conditioned on sampling puzzle s . More formally, we will prove that

$$\text{SD}(\{s, k\}, \{s, \mathcal{A}'(s)\}) \leq \frac{1}{q(n)},$$

where $(s, k) \leftarrow \text{Samp}(1^n)$.

Note that $\Pr_{\text{Samp}(1^n)}[(s, k)]$ may be expressed in terms of the probability of sampling s and the probability of sampling k_{j+1} conditioned on having sampled s and $k_{1..j}$.

$$\Pr_{\text{Samp}(1^n)}[(s, k)] = p_s \cdot \prod_{i=1}^{n-1} p_{k_{i+1}|sk_{1..i}}$$

The probability distribution $\{s, \mathcal{A}'(s)\}_{s, k \leftarrow \text{Samp}(1^n)}$ may also be expressed similarly

$$\Pr_{(s', k') \leftarrow \text{Samp}(1^n)}[s' = s] \cdot \Pr_{\mathcal{A}'(s)}[k] = p_s \cdot \prod_{i=1}^{n-1} \tilde{p}_{k_{i+1}|sk_{1..i}}$$

To argue that the two distributions are close, we will define a series of hybrid distributions interpolating between them. For all $j \in [0, n - 1]$, define distribution \tilde{D}_j on $\{0, 1\}^n \times \{0, 1\}^n$ as follows. For all $s \in \{0, 1\}^n$ and $k \in \{0, 1\}^n$

$$\begin{aligned} \Pr_{\tilde{D}_j}[s, k] &:= p_s \cdot p_{k_1|s} \cdot p_{k_2|sk_1} \cdot p_{k_3|sk_{1,2}} \cdots p_{k_j|sk_{1..j-1}} \cdot \tilde{p}_{k_{j+1}|sk_{1..j}} \cdots \tilde{p}_{k_n|sk_{1..n-1}} \\ &= p_s \cdot \prod_{i=0}^{j-1} p_{k_{i+1}|sk_{1..i}} \prod_{i=j}^{n-1} \tilde{p}_{k_{i+1}|sk_{1..j}} \end{aligned}$$

Note that $\tilde{D}_0 = \{s, \mathcal{A}'(s)\}_{s, k \leftarrow \text{Samp}(1^n)}$ and $\tilde{D}_n = \{s, k\}_{s, k \leftarrow \text{Samp}(1^n)}$. Therefore, by the triangle inequality for statistical distance

$$\begin{aligned} \text{SD}(\{s, k\}, \{s, \mathcal{A}'(s)\}) &= \text{SD}(\tilde{D}_0, \tilde{D}_n) \\ &\leq \sum_{j=0}^{n-1} \text{SD}(\tilde{D}_{j+1}, \tilde{D}_j) \end{aligned} \tag{3}$$

where $(s, k) \leftarrow \text{Samp}(1^n)$. To upper bound $\text{SD}(\{s, k\}, \{s, \mathcal{A}'(s)\})$, it therefore suffices to upper

bound $\text{SD}(\tilde{D}_{j+1}, \tilde{D}_j)$. For any $j \in [0, n-1]$

$$\begin{aligned}
\text{SD}(\tilde{D}_{j+1}, \tilde{D}_j) &= \frac{1}{2} \cdot \sum_{s,k} \left| \Pr_{\tilde{D}_{j+1}}[s, k] - \Pr_{\tilde{D}_j}[s, k] \right| \\
&= \frac{1}{2} \cdot \sum_{s,k} p_s \cdot \prod_{i=0}^{j-1} p_{k_{i+1}|sk_{1\dots i}} \left| p_{k_{j+1}|sk_{1\dots j}} - \tilde{p}_{k_{j+1}|sk_{1\dots j}} \right| \prod_{i=j}^{n-1} \tilde{p}_{k_{i+1}|sk_{1\dots i}} \\
&= \frac{1}{2} \cdot \sum_{s, k_{1\dots j+1}} p_s \cdot \prod_{i=0}^{j-1} p_{k_{i+1}|sk_{1\dots i}} \left| p_{k_{j+1}|sk_{1\dots j}} - \tilde{p}_{k_{j+1}|sk_{1\dots j}} \right| \cdot \sum_{k_{j+2\dots n}} \prod_{i=j}^{n-1} \tilde{p}_{k_{i+1}|sk_{1\dots i}} \\
&= \frac{1}{2} \cdot \sum_{s, k_{1\dots j+1}} p_s \cdot \prod_{i=0}^{j-1} p_{k_{i+1}|sk_{1\dots i}} \left| p_{k_{j+1}|sk_{1\dots j}} - \tilde{p}_{k_{j+1}|sk_{1\dots j}} \right| \tag{4}
\end{aligned}$$

The value of $\left| p_{k_{j+1}|sk_{1\dots j}} - \tilde{p}_{k_{j+1}|sk_{1\dots j}} \right|$ expresses how far the output distribution of $\mathcal{S}(s, k_{1\dots j})$ is from the distribution of k_{j+1} conditioned on $s, k_{1\dots j}$. In the next subclaim we show that this term is small if for all $b \in \{0, 1\}$, $\mathcal{A}(s, k_{1\dots j}, b)$ is close to $\Pr_{\mathcal{D}_n}[s|k_{1\dots j}|b]$ with high probability.

Claim 5.7. For all $s \in \{0, 1\}^n, k \in \{0, 1\}^n, j \in [0, n-1]$, define $\tau_{s, k_{1\dots j}}$ as follows.

$$\tau_{s, k_{1\dots j}} := \sum_{b \in \{0, 1\}} \Pr \left[\left| \mathcal{A}(s, k_{1\dots j}, b) - \Pr_{\mathcal{D}_n}[s|k_{1\dots j}|b] \right| > \frac{\Pr_{\mathcal{D}_n}[s|k_{1\dots j}|b]}{p(n)} \right]$$

then

$$\left| p_{k_{j+1}|sk_{1\dots j}} - \tilde{p}_{k_{j+1}|sk_{1\dots j}} \right| \leq 7/p(n) + \tau_{s, k_{1\dots j}}$$

Proof. First we note that since by definition $p_{1|sk_{1\dots j}} + p_{0|sk_{1\dots j}} = 1$ and $\tilde{p}_{1|sk_{1\dots j}} + \tilde{p}_{0|sk_{1\dots j}} = 1$

$$\left| p_{1|sk_{1\dots j}} - \tilde{p}_{1|sk_{1\dots j}} \right| = \left| p_{0|sk_{1\dots j}} - \tilde{p}_{0|sk_{1\dots j}} \right|$$

For $b \in \{0, 1\}$, let $a_b := \Pr_{\mathcal{D}_n}[s|k_{1\dots j}|b]$. By the construction of \mathcal{D}_n , for $b \in \{0, 1\}$

$$a_b = \Pr[i = j] \cdot \Pr_{s', k' \leftarrow \text{Samp}(1^n)}[s' = s \wedge k'_{1\dots j} = k_{1\dots j}] \cdot \left(\frac{1}{4} + \frac{p_{b|sk_{1\dots j}}}{2} \right)$$

This allows us to express $p_{b|sk_{1\dots j}}$ in terms of a_b . More precisely we can say

$$p_{1|sk_{1\dots j}} = \frac{3a_1 - a_0}{2(a_0 + a_1)}$$

Recall that $\mathcal{S}(s, k_{1\dots j})$ computes \tilde{a}_1 and \tilde{a}_0 , and then calculates $\pi = \frac{3\tilde{a}_1 - \tilde{a}_0}{2\tilde{a}_1 + 2\tilde{a}_0}$. Finally it samples a bit according to $\text{Bern}(\pi)$. Therefore we can bound the distance between $p_{1|sk_{1\dots j}}$ and $\tilde{p}_{1|sk_{1\dots j}}$ in the case where for all $b \in \{0, 1\}$, \tilde{a}_b is close to a_b .

SubClaim 5.1. If for some sampled \tilde{a}_0 and \tilde{a}_1 the following two conditions hold:

- $|\tilde{a}_0 - a_0| \leq \frac{a_0}{p(n)}$
- $|\tilde{a}_1 - a_1| \leq \frac{a_1}{p(n)}$

Then $|\pi - p_{1|sk_{1\dots j}}| \leq \frac{6}{p(n)}$.

Proof. Let $t := a_1/a_0$ and $\tilde{t} := \tilde{a}_1/\tilde{a}_0$. Since $a_0 \geq 0$ and $a_1 \geq 0$ we can bound \tilde{t} using t as follows. $\tilde{a}_0 \geq (1 - 1/p(n)) \cdot a_0$ and $\tilde{a}_1 \leq (1 + 1/p(n)) \cdot a_1$ so for large enough n

$$\begin{aligned}\tilde{t} &\leq \frac{1 + 1/p(n)}{1 - 1/p(n)} \cdot t \\ &\leq (1 + 3/p(n)) \cdot t\end{aligned}$$

Similarly, $\tilde{a}_0 \leq (1 + 1/p(n)) \cdot a_0$ and $\tilde{a}_1 \geq (1 - 1/p(n)) \cdot a_1$ so for large enough n

$$\begin{aligned}\tilde{t} &\geq \frac{1 - 1/p(n)}{1 + 1/p(n)} \cdot t \\ &\geq (1 - 3/p(n)) \cdot t\end{aligned}$$

Now, $p_{1|sk_1\dots j} = \frac{3a_1 - a_0}{2(a_0 + a_1)} = \frac{3t - 1}{2(t + 1)}$ and $\pi = \frac{3\tilde{a}_1 - \tilde{a}_0}{2(\tilde{a}_0 + \tilde{a}_1)} = \frac{3\tilde{t} - 1}{2(\tilde{t} + 1)}$ so

$$\begin{aligned}|p_{1|sk_1\dots j} - \pi| &= \left| \frac{3t - 1}{2(t + 1)} - \frac{3\tilde{t} - 1}{2(\tilde{t} + 1)} \right| \\ &= \left| \frac{3}{2} - \frac{4}{2(t + 1)} - \frac{3}{2} + \frac{4}{2(\tilde{t} + 1)} \right| \\ &= \left| \frac{2}{(\tilde{t} + 1)} - \frac{2}{(t + 1)} \right| \\ &= \left| \frac{2(t - \tilde{t})}{(\tilde{t} + 1)(t + 1)} \right|\end{aligned}$$

We have shown above that $|t - \tilde{t}| \leq 3t/p(n)$

$$\begin{aligned}|p_{1|sk_1\dots j} - \pi| &\leq \left| \frac{4t/p(n)}{(\tilde{t} + 1)(t + 1)} \right| \\ &= \frac{6}{p(n)} \cdot \left| \frac{t}{(\tilde{t} + 1)(t + 1)} \right| \\ &= \frac{6}{p(n)} \cdot \left| \frac{1}{(\tilde{t} + 1)(1 + 1/t)} \right| \\ &\leq \frac{6}{p(n)}\end{aligned}$$

□

The subclaim shows that when $|\tilde{a}_0 - a_0| \leq \frac{a_0}{p(n)}$ and $|\tilde{a}_1 - a_1| \leq \frac{a_1}{p(n)}$ then $|\pi - p_{1|sz}| \leq \frac{6}{p(n)}$. Additionally note that $|\pi - p_{1|sz}|$ cannot exceed 1. We can therefore unconditionally bound $|\pi - p_{1|sz}|$ in terms of the probability of sampling such \tilde{a}_0, \tilde{a}_1 . Define τ' as

$$\tau' := \Pr \left[\left(|\tilde{a}_0 - a_0| > \frac{a_0}{p(n)} \right) \vee \left(|\tilde{a}_1 - a_1| > \frac{a_1}{p(n)} \right) \right]$$

We can therefore express $\left| p_{k_{j+1}|sk_1\dots j} - \tilde{p}_{k_{j+1}|sk_1\dots j} \right|$ as

$$\begin{aligned}\left| p_{k_{j+1}|sk_1\dots j} - \tilde{p}_{k_{j+1}|sk_1\dots j} \right| &\leq (1 - \tau') \cdot 6/p(n) + \tau' \cdot 1 \\ &= 6/p(n) + (1 - 6/p(n))\tau' \\ &\leq 6/p(n) + \tau'\end{aligned}$$

The statement of the claim follows from the observation that $\tau_{s,k_{1\dots j}} \geq \tau'$ □

We can use the claim to rewrite (4) as

$$\begin{aligned} \text{SD}(\tilde{D}_{j+1}, \tilde{D}_j) &\leq \frac{1}{2} \cdot \sum_{s,k_{1\dots j+1}} p_s \cdot \prod_{i=0}^{j-1} p_{k_{i+1}|sk_{1\dots i}} \cdot (6/p(n) + \tau_{s,k_{1\dots j}}) \\ &= 3/p(n) + \frac{1}{2} \cdot \sum_{s,k_{1\dots j+1}} p_s \cdot \prod_{i=0}^{j-1} p_{k_{i+1}|sk_{1\dots i}} \cdot \tau_{s,k_{1\dots j}} \end{aligned}$$

which can be plugged into (3) to get

$$\begin{aligned} \text{SD}(\{s, k\}, \{s, \mathcal{A}'(s)\}) &\leq 3n/p(n) + \frac{1}{2} \cdot \sum_{j,s,k_{1\dots j+1}} p_s \cdot \prod_{i=0}^{j-1} p_{k_{i+1}|sk_{1\dots i}} \cdot \tau_{s,k_{1\dots j}} \\ &= 3n/p(n) + \frac{1}{2} \sum_{j,s,k_{1\dots j+1}} \Pr_{s',k' \leftarrow \text{Samp}(1^n)} [s' = s \wedge k'_{1\dots j} = k_{1\dots j}] \cdot \tau_{s,k_{1\dots j}} \quad (5) \end{aligned}$$

Recall that by assumption, with high probability over $x \leftarrow \mathcal{D}_n$, $|\mathcal{A}(x) - \Pr_{\mathcal{D}_n}[x]|$ is bounded, i.e.:

$$\Pr_{x \leftarrow \mathcal{D}_n} \left[|\mathcal{A}(x) - \Pr_{\mathcal{D}_n}[x]| > \frac{\Pr_{\mathcal{D}_n}[x]}{p(n)} \right] < 1/p(n)$$

which can be expressed as the following sum

$$\sum_x \Pr_{\mathcal{D}_n}[x] \cdot \Pr \left[|\mathcal{A}(x) - \Pr_{\mathcal{D}_n}[x]| > \frac{\Pr_{\mathcal{D}_n}[x]}{p(n)} \right] < 1/p(n)$$

and by the construction of \mathcal{D}_n

$$\begin{aligned} \Pr_{\mathcal{D}_n}[s \| k_{1\dots j} \| b] &= \Pr[i = j] \cdot \Pr_{s',k' \leftarrow \text{Samp}(1^n)} [s' = s \wedge k'_{1\dots j} = k_{1\dots j}] \cdot \left(\frac{1}{4} + \frac{pb|sk_{1\dots j}}{2} \right) \\ &= \frac{1}{n} \cdot \Pr_{s',k' \leftarrow \text{Samp}(1^n)} [s' = s \wedge k'_{1\dots j} = k_{1\dots j}] \cdot \left(\frac{1}{4} + \frac{pb|sk_{1\dots j}}{2} \right) \\ &\geq \frac{\Pr_{s',k' \leftarrow \text{Samp}(1^n)} [s' = s \wedge k'_{1\dots j} = k_{1\dots j}]}{4n} \end{aligned}$$

so the above sum can be rewritten as

$$\begin{aligned} 1/p(n) &> \sum_{j,x=s \| k_{1\dots j+1} \| b} \frac{\Pr_{s',k' \leftarrow \text{Samp}(1^n)} [s' = s \wedge k'_{1\dots j} = k_{1\dots j}]}{4n} \cdot \Pr \left[|\mathcal{A}(x) - \Pr_{\mathcal{D}_n}[x]| > \frac{\Pr_{\mathcal{D}_n}[x]}{p(n)} \right] \\ &= \sum_{j,s,k_{1\dots j+1}} \frac{\Pr_{s',k' \leftarrow \text{Samp}(1^n)} [s' = s \wedge k'_{1\dots j} = k_{1\dots j}]}{4n} \cdot \tau_{s,k_{1\dots j}} \end{aligned}$$

Plugging this back into (5) and recalling that $p(n) > 5nq(n)$

$$\begin{aligned} \text{SD}(\{s, k\}, \{s, \mathcal{A}'(s)\}) &\leq 3n/p(n) + \frac{1}{2} \cdot 4n/p(n) \\ &= 5n/p(n) \\ &< 1/q(n) \end{aligned}$$

which contradicts the security of Samp.

6 The Hardness of Pseudo-Deterministic Sampling implies One-Way Puzzles

In this section we prove the following theorem.

Theorem 6.1. *If $1/q(n)$ -pseudo-deterministic hard distributions (Definition 4.4) exist for some non-zero polynomial q , then one-way puzzles exist.*

Proof. Let D_n be a $1/q(n)$ -pseudo-deterministic hard distribution on n bits. We define a candidate puzzle $\text{Samp}(1^n)$ as follows:

- Sample $i \leftarrow [0, n - 1]$
- Sample $x \leftarrow D_n$
- Output puzzle $x_{1\dots i}$ and key x_{i+1}

Let $p(n)$ be a polynomial greater than $2nq(n)$. We prove that $\text{Samp}(1^n)$ is a $1/p(n)$ -distributional one-way puzzle. Since Theorem 4.3 shows that distributional one-way puzzles can be amplified to obtain (strong) one-way puzzles, this suffices to prove the theorem.

Assume for the sake of contradiction that Samp is not a $1/p(n)$ -distributional one-way puzzle. Therefore, there exists a QPT $\mathcal{A} = \{\mathcal{A}_n\}_{n \in \mathbb{N}}$ and (non-uniform, quantum) advice ensemble $|\tau\rangle = \{|\tau_n\rangle\}_{n \in \mathbb{N}}$ such that for infinitely many $n \in \mathbb{N}$,

$$\text{SD}(\{x_{1\dots i}, x_{i+1}\}, \{x_{1\dots i}, \mathcal{A}(|\tau\rangle, x_{1\dots i})\}) < \frac{1}{p(n)} \quad (6)$$

where $x_{1\dots i}, x_{i+1} \leftarrow \text{Samp}(1^n)$. Fix any such adversary \mathcal{A} , any such advice ensemble $|\tau\rangle$, and any such large enough $n \in \mathbb{N}$. We will drop the advice from the notation since it is always implicitly provided to the adversary.

We first define an estimator E that takes input a string z and performs the following:

1. For j in $[25n^4q(n)]$:

$$X_j \leftarrow \mathcal{A}(z)$$

2. Return $\sum_j X_j / 25n^4q(n)$

Intuitively, E outputs an estimate of the probability that $\mathcal{A}(z)$ outputs 1. Define the algorithm $\mathcal{A}'(z; R)$ that takes input z and randomness $R \in [2^n]$ and performs the following:

1. $e \leftarrow E(z)$
2. If $2^n \cdot e > R$ then return 1 else return 0

Define the algorithm $\mathcal{B}(1^n; R_1, R_2, \dots, R_n)$ that takes randomness $\{R_i\}_{i \in [n]}$ where $R_i \in [2^n]$ and performs the following:

1. For $j = 0$ to $n - 1$:

$$x_{j+1} \leftarrow \mathcal{A}'(x_{1\dots j}; R_{j+1}) \quad (\text{i.e. if } j = 0, x_1 \leftarrow \mathcal{A}'(\epsilon))$$

2. Return x .

We will now show that \mathcal{B} contradicts the security of \mathcal{D} .

Claim 6.1 (Correctness).

$$\text{SD}(\mathcal{B}(1^n), D_n) < 1/q(n)$$

Proof. We first show that the the output distribution of \mathcal{A}' is negligibly close to that of \mathcal{A} .

SubClaim 6.1. For any $z \in \{0, 1\}^{<n}$,

$$|\Pr[\mathcal{A}'(z) = 1] - \Pr[\mathcal{A}(z) = 1]| \leq 1/2^n$$

Proof. \mathcal{A}' uses E to obtain a probability estimate e , and then uses the random coins R to output 1 if $2^n \cdot e > R$, else output 0. Therefore, for any estimate e , the probability that \mathcal{A}' outputs 1 equals the probability that $2^n \cdot e > R$. Since R is uniformly sampled from $[2^n]$, this probability is $\frac{\lfloor 2^n \cdot e \rfloor}{2^n}$ which is at most $1/2^n$ far from e . For a fixed z , comparing the probability of $\mathcal{A}'(z)$ returning 1 and the expectation of $E(z)$ we therefore obtain

$$\begin{aligned} |\Pr[\mathcal{A}'(z) = 1] - \mathbb{E}[E(z)]| &= \left| \sum_e \Pr[E(z) = e] \cdot \frac{\lfloor 2^n \cdot e \rfloor}{2^n} - \sum_e \Pr[E(z) = e] \cdot e \right| \\ &= \left| \sum_e \Pr[E(z) = e] \cdot \left(\frac{\lfloor 2^n \cdot e \rfloor}{2^n} - e \right) \right| \\ &\leq \sum_e \Pr[E(z) = e] \cdot \left| \left(\frac{\lfloor 2^n \cdot e \rfloor}{2^n} - e \right) \right| \\ &\leq \sum_e \Pr[E(z) = e] \cdot \frac{1}{2^n} \\ &\leq \frac{1}{2^n} \end{aligned}$$

Additionally, since $E(z)$ simply returns the average of random variables X_i , each of which has expected value $\Pr_{\mathcal{A}}[\mathcal{A}(z) = 1]$, the expected value of $E(z)$ is also $\Pr_{\mathcal{A}}[\mathcal{A}(z) = 1]$. Therefore

$$|\Pr[\mathcal{A}'(z) = 1] - \Pr[\mathcal{A}(z) = 1]| \leq 1/2^n$$

□

A straightforward consequence of SubClaim 6.1 is that \mathcal{A}' is a valid adversary for Samp. Formally,

SubClaim 6.2.

$$\text{SD}(\{x_{1..i}, x_{i+1}\}, \{x_{1..i}, \mathcal{A}'(x_{1..i})\}) < \frac{1}{p(n)} + 1/2^n.$$

where $i \leftarrow [0, n-1]$ and $x \leftarrow D_n$.

Proof. The proof follows directly from SubClaim 6.1 and inequality (6). □

We now define some helpful terms. For $i \in [0, n-1]$, $z \in \{0, 1\}^i$

- Define $p_z := \Pr_{x \leftarrow D_n}[x_{1..i} = z]$, i.e. the probability that the first i bits of x sampled from D_n match z .
- Define $p_{b|z} := \Pr_{x \leftarrow D_n}[x_{i+1} = b | x_{1..i} = z]$ i.e. the probability that the $i+1$ -th bit of x sampled from D_n is b conditioned on the first i bits of x matching z .

- Define $\tilde{p}_{b|z} := \Pr[\mathcal{A}(z) = b]$.

Also, for all $j \in [0, n]$, define distribution \tilde{D}_j on $\{0, 1\}^n$ as follows. For all $x \in \{0, 1\}^n$:

$$\begin{aligned} \Pr_{\tilde{D}_j}[x] &:= p_{x_1} \cdot p_{x_2|x_1} \cdot p_{x_3|x_{1,2}} \cdots p_{x_j|x_{1\dots j-1}} \cdot \tilde{p}_{x_{j+1}|x_{1\dots j}} \cdots \tilde{p}_{x_n|x_{1\dots n-1}} \\ &= \prod_{i=0}^{j-1} p_{x_{i+1}|x_{1\dots i}} \prod_{i=j}^{n-1} \tilde{p}_{x_{i+1}|x_{1\dots i}} \end{aligned}$$

First, note that $\tilde{D}_n = D_n$.

$$\begin{aligned} \Pr_{\tilde{D}_n}[x] &= \prod_{i=0}^{n-1} p_{x_{i+1}|x_{1\dots i}} \\ &= p_x \\ &= \Pr_{D_n}[x] \end{aligned}$$

Also note that $\tilde{D}_0 = \mathcal{B}(1^n)$.

$$\begin{aligned} \Pr_{\tilde{D}_0}[x] &= \prod_{i=0}^{n-1} \tilde{p}_{x_{i+1}|x_{1\dots i}} \\ &= \Pr[B(1^n) = x] \end{aligned}$$

By the triangle inequality for statistical distance

$$\begin{aligned} \text{SD}(\mathcal{B}(1^n), D_n) &= \text{SD}(\tilde{D}_0, \tilde{D}_n) \\ &\leq \sum_{j=0}^{n-1} \text{SD}(\tilde{D}_{j+1}, \tilde{D}_j) \end{aligned}$$

The statistical distance between \tilde{D}_{j+1} and \tilde{D}_j can be expressed as

$$\begin{aligned} \text{SD}(\tilde{D}_{j+1}, \tilde{D}_j) &= \frac{1}{2} \cdot \sum_x \left| \Pr_{\tilde{D}_{j+1}}[x] - \Pr_{\tilde{D}_j}[x] \right| \\ &= \frac{1}{2} \cdot \sum_x \prod_{i=0}^{j-1} p_{x_{i+1}|x_{1\dots i}} \left| p_{x_{j+1}|x_{1\dots j}} - \tilde{p}_{x_{j+1}|x_{1\dots j}} \right| \prod_{i=j}^{n-1} \tilde{p}_{x_{i+1}|x_{1\dots i}} \\ &= \frac{1}{2} \cdot \sum_{x_{1\dots j+1}} \prod_{i=0}^{j-1} p_{x_{i+1}|x_{1\dots i}} \left| p_{x_{j+1}|x_{1\dots j}} - \tilde{p}_{x_{j+1}|x_{1\dots j}} \right| \cdot \sum_{x_{j+2\dots n}} \prod_{i=j}^{n-1} \tilde{p}_{x_{i+1}|x_{1\dots i}} \\ &= \frac{1}{2} \cdot \sum_{x_{1\dots j+1}} \prod_{i=0}^{j-1} p_{x_{i+1}|x_{1\dots i}} \left| p_{x_{j+1}|x_{1\dots j}} - \tilde{p}_{x_{j+1}|x_{1\dots j}} \right| \end{aligned}$$

For any x , expanding $\prod_{i=0}^{j-1} p_{x_{i+1}|x_{1\dots i}}$ shows that the expression equals $p_{x_{1\dots j}}$, i.e. the probability of obtaining first j bits $x_{1\dots j}$ when sampling from D_n . Therefore

$$\text{SD}(\tilde{D}_{j+1}, \tilde{D}_j) = \frac{1}{2} \cdot \sum_{x_{1\dots j+1}} p_{x_{1\dots j}} \left| p_{x_{j+1}|x_{1\dots j}} - \tilde{p}_{x_{j+1}|x_{1\dots j}} \right|$$

Now the right hand side of the equation equals $\text{SD}(\{x_{1\dots j}, x_{j+1}\}, \{x_{1\dots j}, \mathcal{A}'(x_{1\dots j})\})$ when $x \leftarrow D_n$, which gives the following upper bound for $\text{SD}(\mathcal{B}(1^n), D_n)$.

$$\text{SD}(\mathcal{B}(1^n), D_n) \leq \sum_j \text{SD}(\{x_{1\dots j}, x_{j+1}\}, \{x_{1\dots j}, \mathcal{A}'(x_{1\dots j})\})$$

where $x \leftarrow D_n$. If we also consider $j \leftarrow [0, n-1]$ then

$$\begin{aligned} \text{SD}(\mathcal{B}(1^n), D_n) &\leq n \cdot \text{SD}(\{x_{1\dots j}, x_{j+1}\}, \{x_{1\dots j}, \mathcal{A}'(x_{1\dots j})\}) \\ &< n \cdot (1/p(n) + 1/2^n) \\ &< 2n/p(n) \end{aligned}$$

where the second step follows from SubClaim 6.2 and the last step holds for large enough n . Since $p(n) \geq 2nq(n)$, this implies

$$\text{SD}(\mathcal{B}(1^n), D_n) < 1/q(n)$$

which concludes the proof of the claim. \square

Claim 6.2 (Pseudo-determinism).

$$\Pr_{R_1, \dots, R_n}[\exists y \text{ s.t. } \Pr[\mathcal{B}(1^n; R_1, \dots, R_n) \neq y] \leq 1/2^n] > 1 - 1/q(n)$$

Proof. $\mathcal{B}(1^n; R_1, \dots, R_n)$ consists of loop where in the i -th iteration \mathcal{A}' is run with randomness R_i . The input of \mathcal{A}' in the i -th iteration (apart from the random coins R_i) is completely determined by the output of previous iterations. It therefore suffices to show that for each iteration i , for most strings R_i , the output of \mathcal{A}' is pseudo-deterministic, i.e. for the i -th iteration there exists an output y such that \mathcal{A}' with random coins R_i outputs y with high probability.

SubClaim 6.3. For any $i \in [0, n-1]$, $z \in \{0, 1\}^i$, $R \in [2^n]$, we say that $\mathcal{A}'(z; R)$ has determinism error atmost ϵ if there exists y such that

$$\Pr[\mathcal{A}'(z; R) = y] \geq 1 - \epsilon$$

Then for all $i \in [0, n-1]$, for all $z \in \{0, 1\}^i$,

$$\Pr_{R \leftarrow [2^n]}[\mathcal{A}'(z; R) \text{ has determinism error atmost } (1/n2^n)] > 1 - 1/nq(n)$$

Proof. $\mathcal{A}'(z; R)$ uses $E(z)$ to obtain a probability estimate e , and then uses the random coins R to output 1 if $2^n \cdot e > R$, else output 0. Let π be the probability that $A(z) = 1$. $E(z)$ is the average of $p(n)$ independent random variables that are 1 with probability π and are 0 otherwise. Therefore by setting $\delta = n$ in the additive Chernoff bound (Theorem 4.1)

$$\Pr_E[|\pi - E(z)| \geq 1/5nq(n)] \leq 2e^{-n^2}$$

Suppose $|2^n \cdot \pi - R| \geq 2/5nq(n)$. Then with probability atleast $2e^{-n^2}$, $2^n \cdot E(z)$ and $2^n \cdot \pi$ are on the same side of R , i.e. if $R < 2^n \cdot \pi$ then $R < 2^n \cdot E(z)$ with probability atleast $2e^{-n^2}$, while if $R > 2^n \cdot \pi$ then $R > 2^n \cdot E(z)$ with probability atleast $2e^{-n^2}$. Since the output of $\mathcal{A}'(z; R)$ is entirely determined by whether or not $R < 2^n \cdot E(z)$, $\mathcal{A}'(z; R)$ therefore has determinism error atmost $2e^{-n^2}$ which is less than $(1/n2^n)$ for large enough n

All that remains to be shown is that $|2^n \cdot \pi - R| \geq 2/5nq(n)$ holds with high enough probability. Since R is uniformly sampled from $[2^n]$, the number of values of R such that $|2^n \cdot \pi - R| < 2/5nq(n)$ is atmost $1 + 2^n \cdot 4/5nq(n)$. The probability that R is not one of these values is therefore atleast $1 - 4/5nq(n) - 1/2^n$ which is greater than $1 - 1/nq(n)$ for large enough n . \square

For each iteration i , with probability atleast $1 - 1/nq(n)$, $\mathcal{A}'(z_i; R_i)$ has determinism error at most $1/(n2^n)$ (where z_i is the input in the i -th iteration). Then

$$\Pr_{R_1, \dots, R_n} [\forall i, \mathcal{A}'(z_i; R_i) \text{ has determinism error at most } 1/(n2^n)] > 1 - n/nq(n) = 1/q(n)$$

The determinism error of \mathcal{B} is at most the sum of the determinism error of its iterations, therefore

$$\Pr_{R_1, \dots, R_n} [\mathcal{B}(1^n; R_1, \dots, R_n) \text{ has determinism error at most } 1/2^n] > 1 - 1/q(n)$$

which concludes the proof of the claim. \square

Claim 6.1 and Claim 6.2 show that \mathcal{B} contradicts the security of \mathcal{D} which concludes the proof of the theorem. \square

7 State Puzzles are Equivalent to One-Way Puzzles

We define state puzzles, which capture the hardness of synthesizing a (secret) quantum state $|\psi_s\rangle$ corresponding to a (public) classical string s , and are implied by quantum money.

Definition 7.1 (State Puzzles). *A state puzzle is defined by a quantum polynomial-time generator $\mathcal{G}(1^n)$ that outputs a classical-quantum state $(s, |\psi_s\rangle)$ such that given s , it is (quantum) computationally infeasible to output ρ that overlaps noticeably with $|\psi_s\rangle$.*

Formally, for every quantum polynomial-time adversary \mathcal{A} , every (non-uniform, quantum) advice ensemble $|\tau\rangle = \{|\tau_n\rangle\}_{n \in \mathbb{N}}$, for large enough $n \in \mathbb{N}$,

$$\mathbb{E}_{\substack{(s, |\psi_s\rangle) \leftarrow \mathcal{G}(1^n) \\ \rho \leftarrow \mathcal{A}(|\tau\rangle, s)}}} [\text{Tr}(|\psi_s\rangle\langle\psi_s|\rho)] \leq \text{negl}(n)$$

We also define a weaker version of state puzzles, where we require that the state output by \mathcal{A} must fail to project onto $|\psi_s\rangle\langle\psi_s|$ with noticeable probability.

Definition 7.2 (ε -Weak State Puzzles). *For $\varepsilon : \mathbb{N} \rightarrow \mathbb{R}$, a ε -weak state puzzle is defined by a quantum polynomial-time generator $\mathcal{G}(1^n)$ that outputs a classical-quantum state $(s, |\psi_s\rangle)$ such that given s , it is (quantum) computationally infeasible to output ρ that almost completely overlaps with $|\psi_s\rangle$*

Formally, for every quantum polynomial-time adversary \mathcal{A} , every (non-uniform, quantum) advice ensemble $|\tau\rangle = \{|\tau_n\rangle\}_{n \in \mathbb{N}}$, for large enough $n \in \mathbb{N}$,

$$\mathbb{E}_{\substack{(s, |\psi_s\rangle) \leftarrow \mathcal{G}(1^n) \\ \rho \leftarrow \mathcal{A}(|\tau\rangle, s)}}} [\text{Tr}(|\psi_s\rangle\langle\psi_s|\rho)] \leq 1 - \varepsilon(n)$$

We will sometimes simply refer to weak state puzzles. This is taken to mean $1/p(n)$ -weak state puzzles for some non-zero polynomial p .

In this section we prove the equivalence of (weak and standard) state puzzles and (distributional and standard) one-way puzzles. We first prove that the existence of weak state puzzles implies the existence of one-way puzzles.

Theorem 7.1. *If $1/q(n)$ -weak state puzzles (Definition 7.2) exist for some non-zero polynomial $q(\cdot)$ then $1/p(n)$ -distributional one-way puzzles (Definition 4.3) exist for some non-zero polynomial $p(\cdot)$.*

Since Theorem 4.3 shows that distributional one-way puzzles can be amplified to obtain (strong) one-way puzzles, the following is a corollary of Theorem 7.1.

Corollary 7.1. *If $1/q(n)$ -weak state puzzles (Definition 7.2) exist for some non-zero polynomial $q(\cdot)$ then one-way puzzles (Definition 4.2) exist.*

Proof. (of Theorem 7.1) For some polynomial $q(\cdot)$, let \mathcal{G} be a $1/q(n)$ -weak state puzzle (Definition 7.2) that outputs $s, |\$s\rangle$, where $s \in \{0, 1\}^n$ and $|\$s\rangle \in \{\mathbb{C}^2\}^{\otimes n}$. Define the algorithm $\text{Samp}(1^n)$ as follows:

1. Sample $c \leftarrow \mathcal{C}$ where \mathcal{C} is the Clifford group for n qubits.
2. Sample $s, |\$s\rangle \leftarrow \mathcal{G}(1^n)$
3. Compute $|\$s, c\rangle := c|\$s\rangle$
4. Sample $b_{\text{mode}} \leftarrow \{0, 1\}$
5. If $b_{\text{mode}} = 0$:
 - (a) Sample $i \leftarrow [0, n - 1]$
 - (b) Measure the first i bits of $|\$s, c\rangle$ in the computational basis to obtain measurement output x and residual state $|\$x\rangle$. If $i = 0$ then x is the empty string and $|\$x\rangle := |\$s, c\rangle$
 - (c) Measure the $i + 1$ -th bit of $|\$s, c\rangle$ (i.e. the first bit of $|\$x\rangle$) to obtain measurement output β .
 - (d) Let $\pi = (s, c, b_{\text{mode}}, i, x)$. Output puzzle π and key β .
6. If $b_{\text{mode}} = 1$:
 - (a) Sample $r \leftarrow \{0, 1\}^n \setminus \{0^n\}$
 - (b) For $z \in \{0, 1\}^n$, define $f_r(z) := \min(z, z \oplus r)$.
 - (c) Apply $\sum_z |z\rangle\langle z| \otimes X^{f_r(z)}$ to $|\$s, c\rangle|0\rangle$
 - (d) Measure the second register in computational basis to obtain measurement outcome x_0 , and set $x_1 = x_0 \oplus r$. Let the residual state on the first register be $|\psi_{\text{post}}\rangle$.
 - (e) Sample $b_{\text{order}} \leftarrow \{0, 1\}$.
 - (f) $y_0 := x_{b_{\text{order}}}$ and $y_1 := x_{1-b_{\text{order}}}$
 - (g) Define $V_{y_0, y_1, b}$ ¹² as a unitary that maps
 - $|y_0\rangle \mapsto \frac{|y_0\rangle + i^b |y_1\rangle}{\sqrt{2}}$
 - $|y_1\rangle \mapsto \frac{|y_0\rangle - i^b |y_1\rangle}{\sqrt{2}}$
 - (h) Sample $b_{\text{rot}} \leftarrow \{0, 1\}$.
 - (i) Apply $V_{y_0, y_1, b_{\text{rot}}}$ to $|\psi_{\text{post}}\rangle$ and measure in computational basis to obtain outcome y .
 - (j) If $y = y_0$ then $\beta \leftarrow 0$, else $\beta \leftarrow 1$.

¹²We can implement $V_{y_0, y_1, 0}$ as follows. Let U_0 be a unitary that maps $|0\rangle$ to $|\varphi_0\rangle := \frac{|y_0\rangle + |y_1\rangle}{\sqrt{2}}$ and let U_1 be a unitary that maps $|0\rangle$ to $|\varphi_1\rangle := \frac{|y_0\rangle - |y_1\rangle}{\sqrt{2}}$. Given a state $|y_b\rangle|0\rangle$, apply U_b to the second register to get $|y_b\rangle|\varphi_b\rangle$. Then apply U_0^\dagger to the second register. If $b = 0$ this results in $|y_0\rangle|0\rangle$, else this results in $|y_1\rangle|\varphi\rangle$ where $|\varphi\rangle = U_0^\dagger|\varphi_1\rangle$ is some state orthogonal to $|0\rangle$. Then coherently perform the operation that applies X^{y_0} to the first register if the second register is $|0\rangle$ and applies X^{y_1} to the first register if the second register is any other computational basis state. This results in $|0\rangle|0\rangle$ if $b = 0$ and $|0\rangle|\varphi\rangle$ otherwise. Finally, apply U_0 to the second register to obtain $|0\rangle|\varphi_b\rangle$ and output the second register. $V_{y_0, y_1, 1}$ can be implemented similarly.

(k) Let $\pi = (s, c, b_{\text{mode}}, y_0, y_1, b_{\text{rot}})$. Output puzzle π and key β .

Let k be a constant greater than 6 such that for large enough n , $n^k \geq q(n)^3$ and let $p(\cdot)$ be a polynomial such that $p(n) \geq n^{64k}$. We will prove that Samp is a $1/p(n)$ -distributional one-way puzzle. Note that this suffices to prove Theorem 7.1.

Assume for the sake of contradiction that Samp is *not* a $1/p(n)$ -distributional one-way puzzle. By the definition of $1/p(n)$ -distributional one-way puzzle, there exists a QPT $\mathcal{A} = \{A_n\}_{n \in \mathbb{N}}$ and (non-uniform, quantum) advice ensemble $|\tau\rangle = \{|\tau_n\rangle\}_{n \in \mathbb{N}}$ such that for infinitely many $n \in \mathbb{N}$,

$$\{\pi, \beta\}_{(\pi, \beta) \leftarrow \text{Samp}(1^n)} \approx_{1/p(n)} \{\pi, \mathcal{A}(|\tau\rangle, \pi)\}_{(\pi, \beta) \leftarrow \text{Samp}(1^n)}$$

Fix any such adversary \mathcal{A} , any such advice ensemble $|\tau\rangle$, and any such large enough $n \in \mathbb{N}$. We will use this adversary to build a reduction that contradicts the security of \mathcal{G} . For all $s \in \{0, 1\}^n$ and $c \in \mathcal{C}_n$, define $\Delta_{s,c}$ as follows:

- Let D_0 be the distribution of (π, β) when π, β is sampled by Samp conditioned on $s = s$ and $c = c$.
- Let D_1 be the distribution of $(\pi, A(\pi, |\tau\rangle))$ when π, β is sampled by Samp conditioned on $s = s$ and $c = c$.
- $\Delta_{s,c} := SD(D_0, D_1)$

Intuitively, $\Delta_{s,c}$ is the adversary's error in sampling from the true distribution when $s = s$ and $c = c$. We will use the adversary to synthesize an approximation of $|\mathcal{S}_{s,c}\rangle$ for a random choice of s, c , following the pattern of Aaronson's synthesis algorithm. The algorithm queries a PP oracle to obtain the values of probabilities and phases. The reduction cannot query a PP oracle, so we will replace the query responses with estimates obtained by querying the adversary.

First, we perform the real amplitude step of Aaronson synthesis.

Claim 7.1 (Real Amplitude Synthesis). *Let $s \in \{0, 1\}^n$ and $c \in \mathcal{C}_n$. Let $|\tau_{\text{amp}}\rangle := |0\rangle|\tau\rangle^{\otimes np(n)}$. Then there exists an efficient unitary $\widetilde{M}_{s,c}$ such that*

$$\left| \widetilde{M}_{s,c}|0\rangle|\tau_{\text{amp}}\rangle - |\mathcal{S}_{s,c}^*\rangle|\tau_{\text{amp}}\rangle \right| \leq \sqrt{3n^3/p(n)^{1/4}} + 2\sqrt{n\Delta_{s,c}}$$

Additionally, there is a uniform circuit family that takes (s, c) as input and implements $\widetilde{M}_{s,c}$.

Proof. Fix any s and c . We first define some terms that will be useful for the proof.

- Interpret $|\mathcal{S}_{s,c}\rangle$ as $\sum_{z \in \{0,1\}^n} a_z e^{-i\phi_z} |z\rangle$ where $a_z \geq 0$ and $\phi_z \in [0, 2\pi)$.
- $|\mathcal{S}_{s,c}^*\rangle := \sum_{z \in \{0,1\}^n} a_z |z\rangle$. Intuitively, $|\mathcal{S}_{s,c}^*\rangle$ represents $|\mathcal{S}_{s,c}\rangle$ with the phase information removed, i.e., with real amplitudes.
- For all $i \in [1, n]$, for all $z \in \{0, 1\}^i$,
 - $p_z := \Pr[x = z | b_{\text{mode}} = 0 \wedge i = i']$ represents the probability that the first i bits equal z when measuring $|\mathcal{S}_{s,c}\rangle$ in the computational basis.
- For all $i \in [0, n-1]$, for all $z \in \{0, 1\}^i$, for all $b \in \{0, 1\}$,
 - $p_{b|z} := \Pr[\beta = b | x = z \wedge b_{\text{mode}} = 0 \wedge i = i']$ represents the probability that the $i+1$ th bit equals b conditioned on the first i bits equalling z when measuring $|\mathcal{S}_{s,c}\rangle$ in the computational basis.

$$- \tilde{p}_{b|z} := \Pr[b = \mathcal{A}(s, c, 0, i, z, |\tau\rangle)]$$

We first use A to obtain a good estimate for $p_{1|z}$ by defining an estimator $\mathcal{E}_i(z)$. For all $i \in [0, n-1]$, for all $z \in \{0, 1\}^i$ define $\mathcal{E}_i(z)$ that takes advice $|\tau\rangle^{\otimes p(n)}$ as follows:

- For $j = 1$ to $p(n)$:
 - $B_j \leftarrow A(s, c, 0, i, z, |\tau\rangle)$
- Return $\sum_j B_j / p(n)$

The rest of the construction of $\widetilde{M}_{s,c}$ proceeds as in the real amplitude step of Aaronson synthesis, except with the oracle queries replaced with estimates given by $\mathcal{E}_i(\cdot)$.

Superposition queries to $\mathcal{E}_i(\cdot)$ may produce entangled junk so we must later on uncompute to remove the junk. To do so we define E_i to be the purification of \mathcal{E}_i that acts on input register Z , output register V , and advice register V' , i.e.,

$$E_i|z\rangle_Z|0\rangle_V|\tau_1\rangle_{V'} = |z\rangle_Z \sum_v \sqrt{\Pr[v = \mathcal{E}_i(z)]} |v\rangle_V |\text{junk}_v\rangle_{V'}$$

where $|\tau_1\rangle := |\tau\rangle^{\otimes p(n)}|0\rangle$ and $|\text{junk}_v\rangle$ is some normalized state.

For all $v \geq 0$, define $|\psi_v\rangle := \sqrt{v}|1\rangle + \sqrt{1-v}|0\rangle$ and let P be a unitary that maps $|v\rangle_V|0\rangle_{Z'}$ to $|v\rangle_V|\psi_v\rangle_{Z'}$ ¹³. Let $X = \{X_i\}_{i \in [1, n]}$, $A = \{A_i\}_{i \in [1, n]}$, $A' = \{A'_i\}_{i \in [1, n]}$ be collections of registers where each X_i is a single qubit. Define \widetilde{M}_i that acts on $X_{1 \dots i+1}$, A_i , and A'_i as follows:

- Let $Z := X_{1 \dots i}$, $Z' := X_{i+1}$, $V := A_i$, $V' := A'_i$
- Apply $(E_i)^\dagger P E_i$ to $ZZ'VV$

Define $\widetilde{M}_{s,c} := \widetilde{M}_{n-1} \widetilde{M}_{n-2} \dots \widetilde{M}_0$. It is easy to see that $\widetilde{M}_{s,c}$ can be implemented efficiently given (s, c) . The rest of the proof of the claim is therefore dedicated to proving correctness of the construction.

We now show that the above construction satisfies Claim 7.1 by a series of subclaims. First we note that with high probability, $\mathcal{E}_i(z)$ is close to $\tilde{p}_{1|z}$.

SubClaim 7.1. For all $i \in [0, n-1]$, for all $z \in \{0, 1\}^i$

$$\Pr \left[\left| \mathcal{E}_i(z) - \tilde{p}_{1|z} \right| \geq \frac{n}{\sqrt{p(n)}} \right] \leq 2e^{-2n^2}$$

Proof. Follows from the definition of $\tilde{p}_{1|z}$ and setting δ to be n in the additive Chernoff bound (Theorem 4.1). \square

Next we note that since E_i , P , and $(E_i)^\dagger$ do not affect the computational basis state on the Z register, for all s, c , for all $i \in [0, n-1]$, for all $z \in \{0, 1\}^i$, there exists a state $|\sigma_z\rangle$ such that

$$(E_i)^\dagger P E_i |z\rangle_Z |0\rangle_{Z'} |\tau_1\rangle_{V'} = |z\rangle_Z |\sigma_z\rangle_{Z'} V V'$$

The next subclaim shows that we can synthesize a state close to $|\psi_{\tilde{p}_{1|z}}\rangle$.

¹³We may only be able to efficiently implement P upto some exponentially small error, however, this small error will not affect our result so we will elide it for the sake of clarity.

SubClaim 7.2. For all $i \in [0, n-1]$, for all $z \in \{0, 1\}^i$, if n is large enough

$$\left| |\sigma_z\rangle_{Z'VV'} - |\psi_{\tilde{p}_{1|z}}\rangle_{Z'}|0\rangle_V|\tau_1\rangle_{V'} \right| \leq \sqrt{3n}/p(n)^{1/4}$$

Proof. First we apply $(E_i)^\dagger PE_i$ to both terms and note this does not change the absolute value of their difference.

$$\begin{aligned} \left| |\sigma_z\rangle_{Z'VV'} - |\psi_{\tilde{p}_{1|z}}\rangle_{Z'}|0\rangle_V|\tau_1\rangle_{V'} \right|^2 &= \left| |z\rangle_Z|\sigma_z\rangle_{Z'VV'} - |z\rangle_Z|\psi_{\tilde{p}_{1|z}}\rangle_{Z'}|0\rangle_V|\tau_1\rangle_{V'} \right|^2 \\ &= \left| (E_i)^\dagger PE_i |z\rangle_Z|0\rangle_{Z'V}|\tau_1\rangle_{V'} - |z\rangle_Z|\psi_{\tilde{p}_{1|z}}\rangle_{Z'}|0\rangle_V|\tau_1\rangle_{V'} \right|^2 \\ &= \left| PE_i |z\rangle_Z|0\rangle_{Z'V}|\tau_1\rangle_{V'} - E_i |z\rangle_Z|\psi_{\tilde{p}_{1|z}}\rangle_{Z'}|0\rangle_V|\tau_1\rangle_{V'} \right|^2 \end{aligned} \quad (7)$$

Next, we use the definition of E_i and P to expand each term. Expanding $E_i |z\rangle_Z|\psi_{\tilde{p}_{1|z}}\rangle_{Z'}|0\rangle_V|\tau_1\rangle_{V'}$

$$E_i |z\rangle_Z|\psi_{\tilde{p}_{1|z}}\rangle_{Z'}|0\rangle_V|\tau_1\rangle_{V'} = |z\rangle_Z|\psi_{\tilde{p}_{1|z}}\rangle_{Z'} \sum_v \sqrt{\Pr[v = \mathcal{E}_i(z)]} |v\rangle_V |\text{junk}_v\rangle_{V'} \quad (8)$$

Expanding $PE_i |z\rangle_Z|0\rangle_{Z'V}|\tau_1\rangle_{V'}$

$$\begin{aligned} PE_i |z\rangle_Z|0\rangle_{Z'V}|\tau_1\rangle_{V'} &= P |z\rangle_Z \sum_v \sqrt{\Pr[v = \mathcal{E}_i(z)]} |0\rangle_{Z'} |v\rangle_V |\text{junk}_v\rangle_{V'} \\ &= |z\rangle_Z \sum_v \sqrt{\Pr[v = \mathcal{E}_i(z)]} |\psi_v\rangle_{Z'} |v\rangle_V |\text{junk}_v\rangle_{V'} \end{aligned} \quad (9)$$

Plugging (8) and (9) into (7) gives

$$\begin{aligned} &\left| |\sigma_z\rangle_{Z'VV'} - |\psi_{\tilde{p}_{1|z}}\rangle_{Z'}|0\rangle_V|\tau_1\rangle_{V'} \right|^2 \\ &= \left| |z\rangle_Z \sum_v \sqrt{\Pr[v = \mathcal{E}_i(z)]} \left(|\psi_v\rangle_{Z'} - |\psi_{\tilde{p}_{1|z}}\rangle_{Z'} \right) |v\rangle_V |\text{junk}_v\rangle_{V'} \right|^2 \\ &= \left| \sum_v \sqrt{\Pr[v = \mathcal{E}_i(z)]} \left(|\psi_v\rangle - |\psi_{\tilde{p}_{1|z}}\rangle \right) |v\rangle \right|^2 \\ &= \sum_v \Pr[v = \mathcal{E}_i(z)] \left| |\psi_v\rangle - |\psi_{\tilde{p}_{1|z}}\rangle \right|^2 \end{aligned} \quad (10)$$

SubClaim 7.1 shows that all but a negligible fraction of the probability mass in the output of \mathcal{E}_i is on v that are $n/\sqrt{p(n)}$ close to $\tilde{p}_{1|z}$. We can therefore bound the contribution v values that are not close, i.e. for $\mathbb{V} := \left\{ v : |v - \tilde{p}_{1|z}| \leq \frac{n}{\sqrt{p(n)}} \right\}$, we see that terms not in \mathbb{V} contribute negligible amounts to the sum.

$$\begin{aligned} &\sum_v \Pr[v = \mathcal{E}_i(z)] \left| |\psi_v\rangle - |\psi_{\tilde{p}_{1|z}}\rangle \right|^2 \\ &= \sum_{v \in \mathbb{V}} \Pr[v = \mathcal{E}_i(z)] \left| |\psi_v\rangle - |\psi_{\tilde{p}_{1|z}}\rangle \right|^2 + \sum_{v \notin \mathbb{V}} \Pr[v = \mathcal{E}_i(z)] \left| |\psi_v\rangle - |\psi_{\tilde{p}_{1|z}}\rangle \right|^2 \\ &\leq \sum_{v \in \mathbb{V}} \Pr[v = \mathcal{E}_i(z)] \left| |\psi_v\rangle - |\psi_{\tilde{p}_{1|z}}\rangle \right|^2 + 4 \sum_{v \notin \mathbb{V}} \Pr[v = \mathcal{E}_i(z)] \\ &\leq \sum_{v \in \mathbb{V}} \Pr[v = \mathcal{E}_i(z)] \left| |\psi_v\rangle - |\psi_{\tilde{p}_{1|z}}\rangle \right|^2 + 8e^{-2n^2} \end{aligned} \quad (11)$$

where the fourth step uses SubClaim 7.1 to show that $\sum_{v \notin \mathbb{V}} \Pr[v = \mathcal{E}_i(z)]$. Now, for any $v \in \mathbb{V}$, since v and $\tilde{p}_{1|z}$ are close, we can bound the first term. By definition, $|\psi_v\rangle = \sqrt{v}|1\rangle + \sqrt{1-v}|0\rangle$ and $|\psi_{\tilde{p}_{1|z}}\rangle = \sqrt{\tilde{p}_{1|z}}|1\rangle + \sqrt{1-\tilde{p}_{1|z}}|0\rangle$. Therefore, for $v \notin \mathbb{V}$

$$\begin{aligned}
& \left| |\psi_v\rangle - |\psi_{\tilde{p}_{1|z}}\rangle \right|^2 \\
&= \left| \sqrt{v}|1\rangle + \sqrt{1-v}|0\rangle - \sqrt{\tilde{p}_{1|z}}|1\rangle - \sqrt{1-\tilde{p}_{1|z}}|0\rangle \right|^2 \\
&\leq \left(\sqrt{v} - \sqrt{\tilde{p}_{1|z}} \right)^2 + \left(\sqrt{1-v} - \sqrt{1-\tilde{p}_{1|z}} \right)^2 \\
&\leq \left| \left(\sqrt{v} - \sqrt{\tilde{p}_{1|z}} \right) \left(\sqrt{v} + \sqrt{\tilde{p}_{1|z}} \right) \right| + \left| \left(\sqrt{1-v} - \sqrt{1-\tilde{p}_{1|z}} \right) \left(\sqrt{1-v} + \sqrt{1-\tilde{p}_{1|z}} \right) \right| \\
&\leq 2|v - \tilde{p}_{1|z}| \\
&\leq 2n/\sqrt{p(n)}
\end{aligned}$$

where the last step follows directly from the definition of \mathbb{V} . Substituting this bound in (11) and plugging the result into (10) gives

$$\begin{aligned}
\left| |\sigma_z\rangle_{Z \setminus \mathbb{V}'} - |\psi_{\tilde{p}_{1|z}}\rangle_{Z'} |0\rangle_{\mathbb{V}} |\tau_1\rangle_{\mathbb{V}'} \right|^2 &\leq \sum_{v \in \mathbb{V}} \Pr[v = \mathcal{E}_i(z)] \left| |\psi_v\rangle - |\psi_{\tilde{p}_{1|z}}\rangle \right|^2 + 8e^{-2n^2} \\
&\leq \sum_{v \in \mathbb{V}} \Pr[v = \mathcal{E}_i(z)] \cdot 2n/\sqrt{p(n)} + 8e^{-2n^2} \\
&\leq 2n/\sqrt{p(n)} + 8e^{-2n^2} \leq 3n/\sqrt{p(n)}
\end{aligned}$$

where the last step uses the fact that n is large enough. This concludes the proof of SubClaim 7.2. \square

For all $i \in [0, n-1]$, define the unitary M_i that for all $z \in \{0, 1\}^i$ maps $|z\rangle_{X_{1\dots i}} |0\rangle_{X_{i+1}} |0\rangle_{A_i} |\tau_1\rangle_{A'_i}$ to $|z\rangle_{X_{1\dots i}} |\psi_{p_{1|z}}\rangle_{X_{i+1}} |0\rangle_{A_i} |\tau_1\rangle_{A'_i}$. These represent steps in the amplitude step of Aaronson synthesis. Define $|\tau_2\rangle_{A'} := \bigotimes_j |\tau_1\rangle_{A'_j}$. Note that when z is the empty string, $p_{b|z} = p_b$.

SubClaim 7.3. For all $i \in [0, n-1]$

$$M_i M_{i-1} \dots M_1 M_0 |0\rangle_{X_{1\dots i+1}A} |\tau_2\rangle_{A'} = \sum_{z \in \{0,1\}^{i+1}} \sqrt{p_z} |z\rangle_{X_{1\dots i+1}} |0\rangle_A |\tau_2\rangle_{A'}$$

Proof. We prove by induction on i . Consider the base case when $i = 0$. By definition, $|\psi_{\tilde{p}_{1|z}}\rangle = \sqrt{\tilde{p}_{1|z}}|1\rangle + \sqrt{1-\tilde{p}_{1|z}}|0\rangle$, so

$$\begin{aligned}
M_0 |0\rangle_{X_1} |0\rangle_A |\tau_2\rangle_{A'} &= |\psi_{p_{1|z}}\rangle_{X_1} |0\rangle_A |\tau_2\rangle_{A'} \\
&= (\sqrt{p_0}|0\rangle_{X_1} + \sqrt{p_1}|1\rangle_{X_1}) |0\rangle_A |\tau_2\rangle_{A'} \\
&= \sum_{z \in \{0,1\}} \sqrt{p_z} |z\rangle_{X_1} |0\rangle_A |\tau_2\rangle_{A'}
\end{aligned}$$

Suppose the SubClaim holds for all $i < i'$. Then by applying the induction hypothesis

$$\begin{aligned}
& M_{i'} M_{i'-1} \dots M_1 M_0 |0\rangle_{\mathbf{X}_{1\dots i'+1} \mathbf{A}} |\tau_2\rangle_{\mathbf{A}'} \\
&= M_{i'} \sum_{z \in \{0,1\}^{i'}} \sqrt{p_z} |z\rangle_{\mathbf{X}_{1\dots i}} |0\rangle_{\mathbf{X}_{i'+1}} |0\rangle_{\mathbf{A}} |\tau_2\rangle_{\mathbf{A}'} \\
&= \sum_{z \in \{0,1\}^{i'}} \sqrt{p_z} |z\rangle_{\mathbf{X}_{1\dots i'}} \left(\sqrt{p_{0|z}} |0\rangle_{\mathbf{X}_{i'+1}} + \sqrt{p_{1|z}} |1\rangle_{\mathbf{X}_{i'+1}} \right) |0\rangle_{\mathbf{A}} |\tau_2\rangle_{\mathbf{A}'} \\
&= \sum_{z \in \{0,1\}^{i'+1}} \sqrt{p_z} |z\rangle_{\mathbf{X}_{1\dots i'+1}} |z\rangle_{\mathbf{X}_{1\dots i'+1}} |0\rangle_{\mathbf{A}} |\tau_2\rangle_{\mathbf{A}'}
\end{aligned}$$

Which concludes the proof of SubClaim 7.3. \square

Next we show via hybrid argument that M_i can be replaced by \widetilde{M}_i one by one for each i . Let

$$\Delta_{s,c,i} := \frac{1}{2} \cdot \sum_{z \in \{0,1\}^i} p_z \left(|p_{0|z} - \widetilde{p}_{0|z}| + |p_{1|z} - \widetilde{p}_{1|z}| \right)$$

Note that by the definition of $\Delta_{s,c}$, the contribution to $\Delta_{s,c}$ in the case when $b_{\text{mode}} = 0$ may be written as the expectation over i of $\Delta_{s,c,i}$. Therefore

$$\begin{aligned}
\Delta_{s,c} &\geq \Pr[b_{\text{mode}} = 0] \cdot \sum_i \Pr[i] \cdot \Delta_{s,c,i} \\
&\geq \sum_i \Delta_{s,c,i} / 2n
\end{aligned}$$

SubClaim 7.4. For all $i \in [0, n]$, define

$$|\dot{c}_i\rangle := \widetilde{M}_{n-1} \widetilde{M}_{n-2} \dots \widetilde{M}_i M_{i-1} \dots M_0 |0\rangle_{\mathbf{X}_{\mathbf{A}}} |\tau_2\rangle_{\mathbf{A}'}$$

Then for all $i \in [0, n-1]$

$$\left| |\dot{c}_{i+1}\rangle - |\dot{c}_i\rangle \right| \leq \sqrt{3n/p(n)}^{1/4} + \sqrt{2\Delta_{s,c,i}}$$

Proof. By expanding the definitions,

$$\begin{aligned}
\left| |\dot{c}_{i+1}\rangle - |\dot{c}_i\rangle \right| &= \left| \widetilde{M}_{n-1} \widetilde{M}_{n-2} \dots \widetilde{M}_{i+1} \left(M_i - \widetilde{M}_i \right) M_{i-1} \dots M_0 |0\rangle_{\mathbf{X}_{\mathbf{A}}} |\tau_2\rangle_{\mathbf{A}'} \right| \\
&= \left| \left(M_i - \widetilde{M}_i \right) M_{i-1} \dots M_0 |0\rangle_{\mathbf{X}_{\mathbf{A}}} |\tau_2\rangle_{\mathbf{A}'} \right|
\end{aligned}$$

Applying SubClaim 7.3 to $M_{i-1} \dots M_0 |0\rangle_{\mathbf{X}_{\mathbf{A}}} |\tau_2\rangle_{\mathbf{A}'}$ we get

$$\left| |\dot{c}_{i+1}\rangle - |\dot{c}_i\rangle \right| = \left| \left(M_i - \widetilde{M}_i \right) \sum_{z \in \{0,1\}^i} \sqrt{p_z} |z\rangle_{\mathbf{X}_{1\dots i}} |0\rangle_{\mathbf{X}_{i+1\dots n} \mathbf{A}} |\tau_2\rangle_{\mathbf{A}'} \right|$$

Expanding $M_i |z\rangle_{\mathbf{X}_{1\dots i}} |0\rangle_{\mathbf{X}_{i+1\dots n} \mathbf{A}} |\tau_2\rangle_{\mathbf{A}'}$

$$M_i |z\rangle_{\mathbf{X}_{1\dots i}} |0\rangle_{\mathbf{X}_{i+1\dots n} \mathbf{A}} |\tau_2\rangle_{\mathbf{A}'} = |z\rangle_{\mathbf{X}_{1\dots i}} |\psi_{p_{1|z}}\rangle_{\mathbf{X}_{i+1}} |0\rangle_{\mathbf{X}_{i+2\dots n} \mathbf{A}} |\tau_2\rangle_{\mathbf{A}'}$$

Expanding $\widetilde{M}_i|z\rangle_{X_{1\dots i}}|0\rangle_{X_{i+1\dots n}A}|\tau_2\rangle_{A'}$

$$\begin{aligned} & \widetilde{M}_i|z\rangle_{X_{1\dots i}}|0\rangle_{X_{i+1\dots n}A}|\tau_2\rangle_{A'} \\ &= \widetilde{M}_i|z\rangle_{X_{1\dots i}}|0\rangle_{X_{i+1\dots n}A} \left(|\tau_1\rangle_{A'_1} \cdots |\tau_1\rangle_{A'_n} \right) \\ &= |z\rangle_{X_{1\dots i}}|\sigma_z\rangle_{X_{i+1}A_iA'_i}|0\rangle_{X_{i+2\dots n}A_{1\dots i-1,i+1\dots n}} \left(|\tau_1\rangle_{A'_1} \cdots |\tau_1\rangle_{A'_{i-1}}|\tau_1\rangle_{A'_{i+1}} \cdots |\tau_1\rangle_{A'_n} \right) \end{aligned}$$

When taking the norm of the difference, the registers $X_{i+2\dots n}, A_{1\dots i-1,i+1\dots n}, A'_{1\dots i-1,i+1\dots n}$ are identical in both states and may be ignored. Therefore by the triangle inequality

$$\begin{aligned} & \left| |\dot{C}_{i+1}\rangle - |\dot{C}_i\rangle \right| \\ &= \left| \sum_{z \in \{0,1\}^i} \sqrt{p_z}|z\rangle \left(|\psi_{p_{1|z}}\rangle|0\rangle|\tau_1\rangle - |\sigma_z\rangle \right) \right| \\ &= \left| \sum_{z \in \{0,1\}^i} \sqrt{p_z}|z\rangle \left(|\psi_{p_{1|z}}\rangle|0\rangle|\tau_1\rangle - |\psi_{\tilde{p}_{1|z}}\rangle|0\rangle|\tau_1\rangle + |\psi_{\tilde{p}_{1|z}}\rangle|0\rangle|\tau_1\rangle - |\sigma_z\rangle \right) \right| \\ &\leq \left| \sum_{z \in \{0,1\}^i} \sqrt{p_z}|z\rangle \left(|\psi_{p_{1|z}}\rangle|0\rangle|\tau_1\rangle - |\psi_{\tilde{p}_{1|z}}\rangle|0\rangle|\tau_1\rangle \right) \right| + \left| \sum_{z \in \{0,1\}^i} \sqrt{p_z}|z\rangle \left(|\psi_{\tilde{p}_{1|z}}\rangle|0\rangle|\tau_1\rangle - |\sigma_z\rangle \right) \right| \end{aligned}$$

Consider the first term. Intuitively, $\tilde{p}_{1|z}$ is on average $\Delta_{s,c,i}$ far from $p_{1|z}$ (over randomness of z). We can therefore bound the first term as follows.

$$\begin{aligned} & \left| \sum_{z \in \{0,1\}^i} \sqrt{p_z}|z\rangle \left(|\psi_{p_{1|z}}\rangle|0\rangle|\tau_1\rangle - |\psi_{\tilde{p}_{1|z}}\rangle|0\rangle|\tau_1\rangle \right) \right| \\ &= \sqrt{\sum_{z \in \{0,1\}^i} p_z \left| |\psi_{p_{1|z}}\rangle - |\psi_{\tilde{p}_{1|z}}\rangle \right|^2} \\ &= \sqrt{\sum_{z \in \{0,1\}^i} p_z \left| \left(\sqrt{p_{1|z}} - \sqrt{\tilde{p}_{1|z}} \right) |1\rangle + \left(\sqrt{p_{0|z}} - \sqrt{\tilde{p}_{0|z}} \right) |0\rangle \right|^2} \\ &= \sqrt{\sum_{z \in \{0,1\}^i} p_z \left(\left(\sqrt{p_{1|z}} - \sqrt{\tilde{p}_{1|z}} \right)^2 + \left(\sqrt{p_{0|z}} - \sqrt{\tilde{p}_{0|z}} \right)^2 \right)} \\ &\leq \sqrt{\sum_{z \in \{0,1\}^i} p_z \left(|p_{1|z} - \tilde{p}_{1|z}| + |p_{0|z} - \tilde{p}_{0|z}| \right)} \\ &\leq \sqrt{2\Delta_{s,c,i}} \end{aligned}$$

Now, consider the second term. We have shown in SubClaim 7.2 that we synthesize a state close

to $|\psi_{\tilde{p}_1|z}\rangle$, therefore

$$\begin{aligned}
& \left| \sum_{z \in \{0,1\}^i} \sqrt{p_z} |z\rangle \left(|\psi_{\tilde{p}_1|z}\rangle |0\rangle |\tau_1\rangle - |\sigma_z\rangle \right) \right| \\
&= \sqrt{\sum_{z \in \{0,1\}^i} p_z \left| |\psi_{\tilde{p}_1|z}\rangle |0\rangle |\tau_1\rangle - |\sigma_z\rangle \right|^2} \\
&\leq \sqrt{\sum_{z \in \{0,1\}^i} p_z \cdot 3n / \sqrt{p(n)}} \\
&= \sqrt{3n/p(n)}^{1/4}
\end{aligned}$$

Where the third step follows from SubClaim 7.2. Adding both error terms gives the final error and concludes the proof of SubClaim 7.4. \square

Now, summing up the errors from SubClaim 7.4

$$\|\dot{c}_n\rangle - |\dot{c}_0\rangle\| \leq \sum_{i \in [0, n-1]} \|\dot{c}_{i+1}\rangle - |\dot{c}_i\rangle\| \leq \left(n \cdot \sqrt{3n/p(n)}^{1/4} + \sum_i \sqrt{2\Delta_{s,c,i}} \right)$$

Applying Jensen's inequality to the final term, followed by the definition of $\Delta_{s,c,i}$

$$\|\dot{c}_n\rangle - |\dot{c}_0\rangle\| \leq \left(n \cdot \sqrt{3n/p(n)}^{1/4} + \sqrt{2n \sum_i \Delta_{s,c,i}} \right) \leq \sqrt{3n^3/p(n)}^{1/4} + 2\sqrt{n\Delta_{s,c}}$$

Finally, we note that by SubClaim 7.3,

$$|\dot{c}_n\rangle = \sum_{z \in \{0,1\}^n} \sqrt{p_z} |z\rangle |0\rangle |\tau_2\rangle = |\$\ast_{s,c}\rangle |0\rangle |\tau_2\rangle$$

and by definition of $\widetilde{M}_{s,c}$ and $|\dot{c}_0\rangle$

$$|\dot{c}_0\rangle = \widetilde{M}_{s,c} |0\rangle |\tau_2\rangle$$

Note that $|\tau_{\text{amp}}\rangle = |0\rangle |\tau_2\rangle$. Therefore

$$\left| \widetilde{M}_{s,c} |0\rangle |\tau_{\text{amp}}\rangle - |\$\ast_{s,c}\rangle |\tau_{\text{amp}}\rangle \right| \leq \sqrt{3n^3/p(n)}^{1/4} + 2\sqrt{n\Delta_{s,c}}$$

and therefore which concludes the proof of Claim 7.1. \square

Claim 7.1 shows how to construct $\widetilde{M}_{s,c}$ for each (s, c) . We now construct \widetilde{M} that takes (s, c) as input and applies $\widetilde{M}_{s,c}$.

Claim 7.2. Let $|\tau_{\text{amp}}\rangle := |0\rangle |\tau\rangle^{\otimes np(n)}$. Then there exists an efficient unitary \widetilde{M} such that for all $s \in \{0, 1\}^n$ and $c \in \mathcal{C}_n$

$$\left| \widetilde{M} |s, c\rangle |0\rangle |\tau_{\text{amp}}\rangle - |s, c\rangle |\$\ast_{s,c}\rangle |\tau_{\text{amp}}\rangle \right| \leq \sqrt{3n^3/p(n)}^{1/4} + 2\sqrt{n\Delta_{s,c}}$$

Proof. Let \widetilde{M} be defined as $\sum_{s,c} |s, c\rangle \langle s, c| \otimes \widetilde{M}_{s,c}$. The statement then follows directly from Claim 7.1. \square

Claim 7.2 shows how the adversary may be used to synthesize a state close to $|\mathbb{S}_{s,c}^*\rangle$. The next step in Aaronson's synthesis is to coherently apply a phase to each basis vector $|z\rangle$ in $|\mathbb{S}_{s,c}^*\rangle$ to obtain $|\mathbb{S}_{s,c}\rangle$. We will show how to use the adversary to perform a similar task. Finally we will apply c^\dagger to obtain an approximation to $|\mathbb{S}_s\rangle$.

Claim 7.3 (State Synthesis). *Let $|\tau_{\text{synth}}\rangle := |\tau\rangle^{\otimes(2n+1)p(n)}$. There exists an efficient algorithm \mathcal{B} such that*

$$\mathbb{E}_{s, |\mathbb{S}_s\rangle \leftarrow \mathcal{G}(1^n)} [\langle \mathbb{S}_s | \mathcal{B}(s, |\tau_{\text{synth}}\rangle) | \mathbb{S}_s \rangle] \geq 1 - 1/q(n)$$

Proof. We first construct an algorithm that takes input z' and estimates the value of $\phi_{z'/z} := \phi_{z'} - \phi_z$ for some fixed z , where $\phi_{z'}$ and ϕ_z are the arguments of the complex phases of $|z\rangle$ and $|z'\rangle$ respectively in $|\mathbb{S}_{s,c}\rangle$. For all s, c , for all $z, z' \in \{0, 1\}^n$ define $\mathcal{U}_z^{s,c}(z')$ that takes advice $|\tau\rangle^{\otimes p(n)}$ as follows:

- For $j = 1$ to $np(n)^{1/4}$:
 - $u_j \leftarrow A(s, c, 1, z, z', 0, |\tau\rangle)$
 - $v_j \leftarrow A(s, c, 1, z, z', 1, |\tau\rangle)$
- $u' \leftarrow \sum_j u_j / np(n)^{1/4}$ and $u \leftarrow 2u' - 1$
- $v' \leftarrow \sum_j v_j / np(n)^{1/4}$ and $v \leftarrow 2v' - 1$
- Return $\arctan2(v, u)$

The rest of the reduction proceeds as in the phase step of Aaronson's synthesis, except the oracle calls are replaced with estimates given by $\mathcal{U}_t^{s,c}(\cdot)$, where t' is a pivot chosen by measuring an approximation of $|\mathbb{S}_{s,c}^*\rangle$.

Superposition queries to $\mathcal{U}_z^{s,c}(\cdot)$ may produce entangled junk so we must later on uncompute to remove the junk. Let $U_z^{s,c}$ be the purification of $\mathcal{U}_z^{s,c}$ that acts on input register Z , output register V , and advice register V' , i.e.,

$$U_z^{s,c} |z'\rangle_Z |0\rangle_V |\tau_1\rangle_{V'} = |z'\rangle_Z \sum_v \sqrt{\Pr[v = \mathcal{U}_z^{s,c}(z')]} |v\rangle_V |\text{junk}_v\rangle_{V'}$$

where $|\tau_1\rangle := |\tau\rangle^{\otimes p(n)} |0\rangle$ and $|\text{junk}_v\rangle$ is some normalized state.

Let P' ¹⁴ be a unitary that maps $|v\rangle_V$ to $e^{-iv} |v\rangle_V$. Let $|\tau_{\text{amp}}\rangle$ and \widetilde{M} be as defined in Claim 7.2. Define the algorithm \mathcal{B}' that takes input (s, c) and advice $|\tau_{\text{amp}}\rangle^{\otimes 2}$ and $|\tau_1\rangle$ as follows:

1. Compute $\widetilde{M}|s, c\rangle |0^n\rangle |\tau_{\text{amp}}\rangle$ and measure the second register to get t' .
2. Compute $(U_{t'}^{s,c})^\dagger P' U_{t'}^{s,c} \left(\widetilde{M}|s, c\rangle_R |0^n\rangle_Z |\tau_{\text{amp}}\rangle_A \right) |0\rangle_V |\tau_1\rangle_{V'}$
3. Apply c^\dagger to register Z
4. Return Z

¹⁴We may only be able to efficiently implement P' upto some exponentially small error, however, this small error will not affect our result so we will elide it for the sake of clarity.

Finally, let \mathcal{B} be the algorithm that takes input s , samples $c \leftarrow \mathcal{C}_n$ and outputs $\mathcal{B}'(s, c)$ using advice $|\tau_{\text{amp}}\rangle^{\otimes 2}$ and $|\tau_1\rangle$. Since $|\tau_{\text{amp}}\rangle = |0\rangle|\tau\rangle^{\otimes np(n)}$ and $|\tau_1\rangle = |\tau\rangle^{\otimes p(n)}|0\rangle$, the algorithm can be executed using advice $|\tau_{\text{synth}}\rangle = |\tau\rangle^{\otimes (2n+1)p(n)}$.

Analysis. The reduction estimates $\phi_{z'z}$ by estimating the probabilities of certain binary outcome measurements. The candidate puzzle Samp is constructed so that any adversary that distributionally inverts the puzzle can be used to approximate the probabilities of the measurements with low error. However, small errors in the probability estimates can still result in large errors in the phase estimate. Essentially, $\mathcal{U}_z^{s,c}(z')$ gives a good estimate of $\phi_{z'z}$ if

- (a) the adversary has low error when $y_0 = z$ and $y_1 = z'$, and
- (b) the weights on z and z' in $|\mathbb{S}_{s,c}\rangle$ are not too far from each other.

Applying a unitary 2-design to the state flattens out the weights on the computational basis states, allowing us to argue that the weights on z' and z are close most of the time.

Before we can analyse the above algorithm, we need to define some helpful sets. Recall k is a constant such that $n^k \geq q(n)$ and $p(n) \geq n^{64k}$. For all s , we will define \mathbb{G}_s as the set of c where the probability that measuring $|\mathbb{S}_{s,c}\rangle$ in the computational basis results in a heavy z is less than $1/n^k$, where a string z is heavy if the probability mass of z in $|\mathbb{S}_{s,c}\rangle$ is greater than $n^{3k}/2^n$. Intuitively, \mathbb{G}_s is the set of c such that $|\mathbb{S}_{s,c}\rangle$ has weight roughly evenly spread over computational basis states z , i.e., the total weight on heavy z values is small. Formally, $\mathbb{G}_s := \left\{ c \text{ s.t. } \sum_{z: |\langle z|\mathbb{S}_{s,c}\rangle|^2 \leq n^{3k}/2^n} |\langle z|\mathbb{S}_{s,c}\rangle|^2 \leq 1/n^k \right\}$. We also define \mathbb{S} as the set of s, c such that the adversary has error less than $1/\sqrt{p}$ when $s = s$ and $c = c$. Formally, $\mathbb{S} := \{s, c \text{ s.t. } \Delta_{s,c} \leq 1/\sqrt{p}\}$.

SubClaim 7.5. For all s :

$$\Pr_c [c \in \mathbb{G}_s] \geq 1 - 1/n^k$$

Proof. The proof follows directly from the following theorem, proved in Appendix B.

Theorem 7.2 (Flatness of 2-designs). Let \mathcal{C} be a unitary 2-design on n qubits. Fix any n qubit state $|\psi\rangle$. For any $C \in \text{Supp}(\mathcal{C})$, let $p_C(x) := |\langle x|C|\psi\rangle|^2$ be the probability that measuring $C|\psi\rangle$ in the computational basis results in x . Then the following holds for all $k > 6$ and sufficiently large n . Define

$$\mathbb{G} := \left\{ C \in \text{Supp}(\mathcal{C}) : \sum_{x: p_C(x) \geq \frac{n^{3k}}{2^n}} p_C(x) \leq 1/n^k \right\}$$

Then

$$\Pr_{C \leftarrow \mathcal{C}} [C \in \mathbb{G}] \geq 1 - 1/n^k$$

□

SubClaim 7.6.

$$\Pr_{s,c} [s, c \in \mathbb{S}] \geq 1 - 1/p(n)^{1/2}$$

Proof. $\Delta_{s,c}$ is defined as the error when $s = s$ and $c = c$. We can write the total error of \mathcal{A} (which is upper bounded by $1/p(n)$) as the expectation of $\Delta_{s,c}$, i.e.

$$\begin{aligned} \frac{1}{p(n)} &\geq \sum_{s,c} \Pr[s = s] \cdot \Pr[c = c] \cdot \Delta_{s,c} \\ &\geq \sum_{s,c \notin \mathbb{S}} \Pr[s = s] \cdot \Pr[c = c] \cdot \Delta_{s,c} \\ &\geq \sum_{s,c \notin \mathbb{S}} \Pr[s = s] \cdot \Pr[c = c] \cdot \frac{1}{\sqrt{p(n)}}, \end{aligned}$$

where the third step follows from the definition of \mathbb{S} . Therefore, $\Pr_{s,c}[s, c \notin \mathbb{S}] \leq \frac{1}{\sqrt{p(n)}}$. \square

Define $\mathbb{S}' := \{s, c : s, c \in \mathbb{S} \wedge c \in \mathbb{G}_s\}$. We first note that with high probability $s, c \in \mathbb{S}'$.

SubClaim 7.7.

$$\Pr_{s,c}[s, c \in \mathbb{S}'] \geq 1 - 1/p(n)^{1/2} - 1/n^k$$

Proof. Follows from the definition of \mathbb{S}' and Subclaims 7.5 and 7.6. \square

The goal of the remainder of the proof is to show that when $(s, c) \in \mathbb{S}'$, $\mathcal{B}'(s, c)$ outputs a state close to $|\mathbb{S}_s\rangle\langle\mathbb{S}_s|$. Since w.h.p. $(s, c) \in \mathbb{S}'$, we can ignore case when $(s, c) \notin \mathbb{S}'$ at the cost of some small error probability which is incorporated in the final result. For the rest of the proof, we fix some $(s, c) \in \mathbb{S}'$, and drop the parameterization on (s, c) in the notation.

Define some terms that will be useful for the proof.

- Interpret $|\mathbb{S}_{s,c}\rangle$ as $\sum_{z \in \{0,1\}^n} a_z e^{-i\phi_z} |z\rangle$ where $a_z \geq 0$ and $\phi_z \in [0, 2\pi)$, and let $\alpha_z := a_z e^{-i\phi_z}$ and $\phi_{z'z} := \phi_{z'} - \phi_z$
- $|\mathbb{S}_{s,c}^*\rangle := \sum_{z \in \{0,1\}^n} a_z |z\rangle$. Intuitively, $|\mathbb{S}_{s,c}^*\rangle$ represents $|\mathbb{S}_{s,c}\rangle$ with the phase information removed, i.e., with real amplitudes.

We now describe how y_0 and y_1 are distributed during the execution of Samp. Define Y_0 as the distribution on y_0 induced by Samp. Let $R := \{0, 1\}^n \setminus \{0^n\}$.

SubClaim 7.8. For all $y_0 \in \{0, 1\}^n$

$$\Pr_{Y_0}[y_0] = \frac{1}{2} \left(|\alpha_{y_0}|^2 \left(1 - \frac{1}{|R|} \right) + \frac{1}{|R|} \right)$$

Proof. Recall that y_0 is generated by measuring the F register in the computational basis to obtain x_0 , and setting y_0 to be either x_0 or $x_0 \oplus r$ at random (where r is uniformly sampled from R).

$$\begin{aligned} \Pr_{Y_0}[y_0] &= \frac{1}{|R|} \sum_{r \neq 0} \frac{1}{2} \times \Pr[\text{measuring } F \text{ register gives } f_r(z) = \min(y_0, y_0 \oplus r)] \\ &= \frac{1}{2|R|} \sum_{r \neq 0} \left| \langle \mathbb{I} \otimes |f_r(y_0)\rangle\langle f_r(y_0)| \left(\sum_x \alpha_x |x\rangle |f_r(x)\rangle \right) \right|^2. \end{aligned}$$

Now, $f_r(y_0) = f_r(x)$ if and only if $x = y_0$ or $x = y_0 \oplus r$, so only the terms, $\alpha_{y_0}|y_0\rangle|f_r(y_0)\rangle$ and $\alpha_{y_0 \oplus r}|y_0 \oplus r\rangle|f_r(y_0)\rangle$, are in the support of the projector. Therefore

$$\begin{aligned} \Pr_{Y_0}[y_0] &= \frac{1}{2|R|} \sum_{r \neq 0} |\alpha_{y_0}|y_0\rangle + \alpha_{y_0 \oplus r}|y_0 \oplus r\rangle|^2 \\ &= \frac{1}{2|R|} \sum_{r \neq 0} (|\alpha_{y_0}|^2 + |\alpha_{y_0 \oplus r}|^2) \\ &= \frac{|\alpha_{y_0}|^2}{2} + \sum_{y \neq y_0} \frac{|\alpha_y|^2}{2|R|}. \end{aligned}$$

Since $\sum_y |\alpha_y|^2 = 1$,

$$\Pr_{Y_0}[y_0] = \frac{|\alpha_{y_0}|^2}{2} + \frac{1 - |\alpha_{y_0}|^2}{2R},$$

which after rearranging gives the statement in the claim. \square

Define $Y_1^{y_0}$ as the distribution on y_1 induced by Samp conditioned on sampling y_0 .

SubClaim 7.9. For all $y_1 \in \{0, 1\}^n$

$$\Pr_{Y_1^{y_0}}(y_1) = \frac{|\alpha_{y_1}|^2 + |\alpha_{y_0}|^2}{1 + |\alpha_{y_0}|^2(|R| - 1)}$$

Proof. By the definition of $Y_1^{y_0}$

$$\begin{aligned} \Pr_{Y_1^{y_0}}[y_1] &= \Pr[y_1|y_0] \\ &= \frac{\Pr[y_1 \wedge y_0]}{\Pr_{Y_0}[y_0]} \end{aligned}$$

Recall that y_0 and y_1 are a random permutation of x_0 and $x_0 \oplus r$ where x_0 is obtained by measuring the F register in the computational basis and r is uniformly sampled from R . Then, y_0 and y_1 are obtained with probability $\frac{1}{2}$ conditioned on $r = y_0 \oplus y_1$ and the measurement outcome of F is $f_r(y_0) = f_r(y_0 \oplus r) = f_r(y_1)$. Thus,

$$\begin{aligned} \Pr_{Y_1^{y_0}}[y_1] &= \frac{1}{2} \cdot \frac{\Pr[\text{measuring } F \text{ register gives } f_r(y_0) \wedge r = y_0 \oplus y_1]}{\Pr_{Y_0}[y_0]} \\ &= \frac{1}{|R|} \frac{|(\mathbb{I} \otimes |f_{y_0 \oplus y_1}(y_0)\rangle \langle f_{y_0 \oplus y_1}(y_0)|) (\sum_x \alpha_x |x\rangle |f_{y_0 \oplus y_1}(x)\rangle)|^2}{\Pr_{Y_0}[y_0]} \\ &= \frac{\frac{1}{|R|} (|\alpha_{y_0}|^2 + |\alpha_{y_1}|^2)}{|\alpha_{y_0}|^2 \left(1 - \frac{1}{|R|}\right) + \frac{1}{|R|}} \\ &= \frac{|\alpha_{y_0}|^2 + |\alpha_{y_1}|^2}{1 + |\alpha_{y_0}|^2(|R| - 1)}, \end{aligned}$$

where the third step follows from Subclaim 7.8. \square

Define Δ'_{y_0} as follows:

- Let D_0 be the distribution of (π, β) when π, β is sampled by Samp conditioned on Samp sampling y_0 and $s = s, c = c, b_{\text{mode}} = 1$.
- Let D_1 be the distribution of $(\pi, A(\pi, |\tau|))$ when π is sampled by Samp conditioned on Samp sampling y_0 and $s = s, c = c, b_{\text{mode}} = 1$.
- $\Delta'_{y_0} := SD(D_0, D_1)$

Define Δ'_{y_0, y_1} as follows:

- Let D_0 be the distribution of (π, β) when π, β is sampled by Samp conditioned on Samp sampling y_0 and y_1 , while $s = s, c = c, b_{\text{mode}} = 1$.
- Let D_1 be the distribution of $(\pi, A(\pi, |\tau|))$ when π is sampled by Samp conditioned on on Samp sampling y_0 and y_1 , while $s = s, c = c, b_{\text{mode}} = 1$.
- $\Delta'_{y_0, y_1} := SD(D_0, D_1)$

Intuitively, Δ'_{y_0} (and Δ'_{y_0, y_1}) represent the adversary's error conditioned on sampling y_0 (and y_1). Define the following sets

- $\mathbb{Z} := \{z \text{ s.t. } \Delta'_z \leq 1/p(n)^{1/4}\}$
- $\mathbb{Z}'_z := \{z' \neq z \text{ s.t. } \Delta'_{z, z'} \leq 1/2p(n)^{1/8}\}$
- $\mathbb{L} := \{z \text{ s.t. } |\alpha_z|^2 \geq \frac{1}{n^{3k}2^n}\}$
- $\mathbb{U} := \{z \text{ s.t. } |\alpha_z|^2 \leq \frac{n^{3k}}{2^n}\}$

Intuitively, \mathbb{Z} and \mathbb{Z}'_z are sets of strings on which the adversary's error is low, while $\mathbb{L} \cap \mathbb{U}$ is the set of strings whose probability mass in $|\$_{s,c}\rangle$ not too far from $1/2^n$. The next subclaim shows that

- (a) If z is sampled by measuring $|\$_{s,c}^*\rangle$, with high probability $z \in \mathbb{Z} \cap \mathbb{L} \cap \mathbb{U}$
- (b) If $z \in \mathbb{Z} \cap \mathbb{L} \cap \mathbb{U}$ then most of the probability mass of $|\$_{s,c}\rangle$ is on strings z' such that $z' \in \mathbb{Z}'_z \cap \mathbb{L} \cap \mathbb{U}$

SubClaim 7.10.

1. $\sum_{z \notin \mathbb{Z}} |\alpha_z|^2 \leq 3/p(n)^{1/4}$
2. $\sum_{z \notin \mathbb{L}} |\alpha_z|^2 \leq 1/n^{3k}$
3. $\sum_{z \notin \mathbb{U}} |\alpha_z|^2 \leq 1/n^k$
4. For all $z \in \mathbb{U} \cap \mathbb{Z}$, $\sum_{z' \notin \mathbb{Z}'_z} |\alpha_{z'}|^2 \leq 4n^{3k}/p(n)^{1/8}$

Proof. 1. Recall that $\Delta_{s,c}$ is the adversary's error in sampling when $s = s$ and $c = c$. $\Delta_{s,c}$ can therefore be expressed as the expectation over y_0 of Δ'_{y_0} , i.e.

$$\begin{aligned}\Delta_{s,c} &\geq \sum_{y_0} \Pr_{Y_0}[y_0] \Delta'_{y_0} \\ &\geq \sum_{y_0} \Delta'_{y_0} \left[\frac{|\alpha_{y_0}|^2}{2} \left(1 - \frac{1}{|R|}\right) + \frac{1}{2|R|} \right] \\ &\geq \sum_{y_0} \Delta'_{y_0} \left[\frac{|\alpha_{y_0}|^2}{2} \left(1 - \frac{1}{|R|}\right) \right] \\ &\geq \sum_{y_0} \Delta'_{y_0} \frac{|\alpha_{y_0}|^2}{3},\end{aligned}$$

where the second step uses Subclaim 7.8. By the definition of S' , $\Delta_{s,c} \leq \frac{1}{\sqrt{p(n)}}$, therefore

$$\begin{aligned}\frac{1}{\sqrt{p(n)}} &\geq \sum_{y_0} \frac{\Delta'_{y_0} |\alpha_{y_0}|^2}{3} \\ &\geq \sum_{y_0 \notin \mathbb{Z}} \frac{\Delta'_{y_0} |\alpha_{y_0}|^2}{3} \\ &\geq \sum_{y_0 \notin \mathbb{Z}} \frac{|\alpha_{y_0}|^2}{3p(n)^{\frac{1}{4}}},\end{aligned}$$

where the last step follows from the definition of \mathbb{Z} . After rearranging and relabeling,

$$\sum_{z \notin \mathbb{Z}} |\alpha_z|^2 \leq \frac{3}{p(n)^{\frac{1}{4}}}.$$

2. By the definition of \mathbb{L}

$$\begin{aligned}\sum_{z \notin \mathbb{L}} |\alpha_z|^2 &\leq \sum_{z \notin \mathbb{L}} \frac{1}{n^{3k} 2^n} \\ &\leq \sum_z \frac{1}{n^{3k} 2^n} \\ &= \frac{1}{n^{3k}}.\end{aligned}$$

3. Recall that $\mathbb{G}_s = \left\{ c \text{ s.t. } \sum_{z: |\langle z | \mathbb{S}_{s,c} \rangle|^2 \leq n^{3k}/2^n} |\langle z | \mathbb{S}_{s,c} \rangle|^2 \leq 1/n^k \right\}$. We fixed $s, c \in S'$, which by definition implies $c \in \mathbb{G}_s$. Therefore

$$\begin{aligned}\frac{1}{n^k} &\geq \sum_{z: |\langle z | \mathbb{S}_{s,c} \rangle|^2 \leq \frac{n^{3k}}{2^n}} |\langle z | \mathbb{S}_{s,c} \rangle|^2 \\ &= \sum_{z \notin \mathbb{U}} |\langle z | \mathbb{S}_{s,c} \rangle|^2 \\ &= \sum_{z \notin \mathbb{U}} |\alpha_z|^2\end{aligned}$$

4. By the definition of Δ'_{y_0} and Δ'_{y_0, y_1} , Δ'_{y_0} can be expressed as the expectation over y_1 of Δ'_{y_0, y_1} . Therefore

$$\begin{aligned}\Delta'_{y_0} &= \sum_{y_1} \Pr_{Y_1^{y_0}}[y_1] \cdot \Delta'_{y_0, y_1} \\ &\geq \sum_{y_1 \notin \mathbb{Z}'_{y_0}} \Pr_{Y_1^{y_0}}[y_1] \cdot \Delta'_{y_0, y_1} \\ &\geq \sum_{y_1 \notin \mathbb{Z}'_{y_0}} \Pr_{Y_1^{y_0}}[y_1] \cdot \frac{1}{2p(n)^{\frac{1}{8}}},\end{aligned}$$

where the last step follows from the definition of \mathbb{Z}'_{y_0} . Therefore,

$$\begin{aligned}\sum_{y_1 \notin \mathbb{Z}'_{y_0}} \Pr_{Y_1^{y_0}}[y_1] &\leq \Delta'_{y_0} \cdot 2p(n)^{\frac{1}{8}} \\ &\leq \frac{2}{p(n)^{\frac{1}{8}}},\end{aligned}$$

where the last step follows from $y_0 \in \mathbb{Z}$. By Subclaim 7.9,

$$\begin{aligned}\Pr_{Y_1^{y_0}}[y_0] &= \frac{|\alpha_{y_0}|^2 + |\alpha_{y_1}|^2}{1 + |\alpha_{y_0}|^2(|R| - 1)} \\ &\geq \frac{|\alpha_{y_1}|^2}{1 + n^{3k}(|R| - 1)2^{-n}} \\ &\geq \frac{|\alpha_{y_1}|^2}{2n^{3k}},\end{aligned}$$

where the second step uses $z \in \mathbb{U}$ and the last step holds for large enough n . Substituting in the previous equation, rearranging, and relabeling gives

$$\sum_{z' \in \mathbb{Z}'_z} |\alpha_{z'}|^2 \geq \frac{4n^{3k}}{p(n)^{\frac{1}{8}}}.$$

□

Next we show that $\mathcal{U}_{t'}(z')$ gives a good estimate for $\phi_{z't'}$ when $t' \in \mathbb{Z} \cap \mathbb{L} \cap \mathbb{U}$ and $z' \in \mathbb{Z}'_{t'} \cap \mathbb{L} \cap \mathbb{U}$.

SubClaim 7.11. For all $z \in \mathbb{L} \cap \mathbb{U} \cap \mathbb{Z}$ and for all $z' \in \mathbb{L} \cap \mathbb{U} \cap \mathbb{Z}'_z$ ¹⁵

$$\Pr \left[\left| e^{-i\mathcal{U}_z(z')} - e^{-i\phi_{z'z}} \right| > \frac{8\sqrt{2}n^{6k}}{p^{1/8}} \right] \leq \text{negl}(n)$$

Proof. Fix any $z \in \mathbb{L} \cap \mathbb{U} \cap \mathbb{Z}$ and any $z' \in \mathbb{L} \cap \mathbb{U} \cap \mathbb{Z}'_z$. Recall that for all x , $a_x = |\alpha_x|$ and $e^{-i\phi_x} = \alpha_x/|\alpha_x|$. Also recall that $\phi_{z'z} = \phi_{z'} - \phi_z$. Consider the state $|\psi_{\text{post}}\rangle$ conditioned on obtaining $y_0 = z$ and $y_1 = z'$ during the execution of Samp.

$$|\psi_{\text{post}}\rangle = \frac{\alpha_z|z\rangle + \alpha_{z'}|z'\rangle}{\sqrt{|\alpha_z|^2 + |\alpha_{z'}|^2}} = e^{-i\phi_z} \cdot \frac{a_z|z\rangle + a_{z'}e^{-i\phi_{z'z}}|z'\rangle}{\sqrt{|\alpha_z|^2 + |\alpha_{z'}|^2}}$$

¹⁵We show that the resulting complex phases are close instead of showing that the angles are close. This is because the estimate of the angle may have a 2π error, but this has no operational meaning when applying the phase.

Let θ be the unique angle in $[0, \pi]$ such that $\cos \frac{\theta}{2} = \frac{a_z}{\sqrt{a_z^2 + a_{z'}^2}}$ and $\sin \frac{\theta}{2} = \frac{a_{z'}}{\sqrt{a_z^2 + a_{z'}^2}}$. Let $|\psi'\rangle := e^{i\phi_z} |\psi_{\text{post}}\rangle$. Then

$$|\psi'\rangle = \cos(\theta/2) |z\rangle + e^{-i\phi_{z'z}} \sin(\theta/2) |z'\rangle$$

Let A' be the probability that applying $V_{z,z',0}$ to $|\psi_{\text{post}}\rangle$ and measuring results in output z . We first obtain an expression for A' in terms of θ and $\phi_{z'z}$ by ignoring global phase and expanding

$$\begin{aligned} A' &= |\langle z | V_{z,z',0} | \psi_{\text{post}} \rangle|^2 \\ &= |\langle z | V_{z,z',0} | \psi' \rangle|^2 \\ &= \left| \langle z | \left(\frac{\cos(\theta/2) + e^{-i\phi_{z'z}} \sin(\theta/2)}{\sqrt{2}} |z\rangle + \frac{\cos(\theta/2) - e^{-i\phi_{z'z}} \sin(\theta/2)}{\sqrt{2}} |z'\rangle \right) \right|^2 \\ &= \left| \frac{\cos(\theta/2) + e^{-i\phi_{z'z}} \sin(\theta/2)}{\sqrt{2}} \right|^2 \\ &= \frac{\cos(\theta/2) + e^{-i\phi_{z'z}} \sin(\theta/2)}{\sqrt{2}} \cdot \frac{\cos(\theta/2) + e^{i\phi_{z'z}} \sin(\theta/2)}{\sqrt{2}} \\ &= \frac{\cos^2(\theta/2) + \sin^2(\theta/2) + 2 \cos(\theta/2) \sin(\theta/2) (e^{-i\phi_{z'z}} + e^{i\phi_{z'z}})}{2} \\ &= \frac{1 + \sin \theta (e^{-i\phi_{z'z}} + e^{i\phi_{z'z}})}{2} \\ &= \frac{1 + \sin \theta \cos \phi_{z'z}}{2} \end{aligned}$$

Similarly, let B' be the probability that applying $V_{z,z',1}$ to $|\psi_{\text{post}}\rangle$ and measuring results in output z . Then

$$\begin{aligned} B' &= |\langle z | V_{z,z',1} | \psi_{\text{post}} \rangle|^2 \\ &= |\langle z | V_{z,z',1} | \psi' \rangle|^2 \\ &= \left| \langle z | \left(\frac{\cos(\theta/2) + ie^{-i\phi_{z'z}} \sin(\theta/2)}{\sqrt{2}} |z\rangle + \frac{\cos(\theta/2) - ie^{-i\phi_{z'z}} \sin(\theta/2)}{\sqrt{2}} |z'\rangle \right) \right|^2 \\ &= \left| \frac{\cos(\theta/2) + ie^{-i\phi_{z'z}} \sin(\theta/2)}{\sqrt{2}} \right|^2 \\ &= \frac{\cos(\theta/2) + ie^{-i\phi_{z'z}} \sin(\theta/2)}{\sqrt{2}} \cdot \frac{\cos(\theta/2) - ie^{i\phi_{z'z}} \sin(\theta/2)}{\sqrt{2}} \\ &= \frac{\cos^2(\theta/2) + \sin^2(\theta/2) + 2i \cos(\theta/2) \sin(\theta/2) (e^{-i\phi_{z'z}} - e^{i\phi_{z'z}})}{2} \\ &= \frac{1 + \sin \theta (ie^{-i\phi_{z'z}} - ie^{i\phi_{z'z}})}{2} \\ &= \frac{1 + \sin \theta \sin \phi_{z'z}}{2} \end{aligned}$$

Let $A := 2A' - 1 = \sin \theta \cos \phi_{z'z}$ and $B := 2B' - 1 = \sin \theta \sin \phi_{z'z}$. We note that since $\theta \in [0, \pi]$, $\sin \theta \geq 0$ and therefore $\arctan 2(B, A) = \phi_{z'z}$.

Let $\tilde{A}' := \Pr[1 = \mathcal{A}(s, c, 1, z, z', 0)]$ and let $\tilde{B}' := \Pr[1 = \mathcal{A}(s, c, 1, z, z', 1)]$. Since $z' \in \mathbb{Z}'_z$

$$\begin{aligned} 1/2p(n)^{1/8} &\geq \Delta'_{z,z'} \\ &= \frac{1}{2} \left(\left| \tilde{A}' - A' \right| + \left| \tilde{B}' - B' \right| \right) \end{aligned}$$

where the second step follows directly from the definition of $\Delta'_{z,z'}$. As a result

$$\begin{aligned} |\tilde{A}' - A'| &\leq 1/p(n)^{1/8} \\ |\tilde{B}' - B'| &\leq 1/p(n)^{1/8} \end{aligned}$$

Let $\tilde{A} := 2\tilde{A}' - 1$ and $\tilde{B} := 2\tilde{B}' - 1$. Therefore

$$\begin{aligned} |\tilde{A} - A| &\leq 2/p(n)^{1/8} \\ |\tilde{B} - B| &\leq 2/p(n)^{1/8} \end{aligned}$$

i.e. \tilde{A} is close to A and \tilde{B} is close to B . Consider the case when \mathcal{U}_z is run on z' and internally samples u and v . By setting $\delta = \sqrt{n}$ in the additive Chernoff bound (Theorem 4.1), we see that u and v that are computed by \mathcal{U}_z are good approximations of \tilde{A} and \tilde{B} , and thus of A and B . Formally,

$$\begin{aligned} \Pr \left[|u - \tilde{A}| \geq \frac{2}{p(n)^{1/8}} \right] &\leq 2e^{-2n} \\ \Pr \left[|v - \tilde{B}| \geq \frac{2}{p(n)^{1/8}} \right] &\leq 2e^{-2n} \end{aligned}$$

Using the fact that \tilde{A} and \tilde{B} are close to A and B as shown above, we can bound the Euclidean distance between (u, v) and (A, B)

$$\begin{aligned} 1 - 4e^{-2n} &\leq \Pr \left[|v - \tilde{B}| \leq \frac{2}{p(n)^{1/8}} \text{ and } |u - \tilde{A}| \leq \frac{2}{p(n)^{1/8}} \right] \\ &\leq \Pr \left[|v - B| \leq \frac{4}{p(n)^{1/8}} \text{ and } |u - A| \leq \frac{4}{p(n)^{1/8}} \right] \\ &\leq \Pr \left[(v - B)^2 + (u - A)^2 \leq \left(\frac{4\sqrt{2}}{p(n)^{1/8}} \right)^2 \right] \end{aligned} \tag{12}$$

We will use the following theorem which we prove in Appendix C

Theorem 7.3. *Let $(x, y), (x^*, y^*) \in \mathbb{R}^2$ such that $\exists \gamma > 0, \gamma' > 0$ s.t.*

- $x^2 + y^2 \geq \gamma^2$
- $(x - x^*)^2 + (y - y^*)^2 \leq (\gamma')^2$
- $\gamma' < \gamma$

Then $|e^{-i \cdot \arctan 2(y,x)} - e^{-i \cdot \arctan 2(y^*,x^*)}| \leq 2\gamma'/\gamma$

Using the fact that both z and z' belong to $\mathbb{L} \cap \mathbb{U}$, we can show that (A, B) is atleast $1/n^{6k}$ far

from the origin.

$$\begin{aligned}
A^2 + B^2 &= \sin^2 \theta (\cos^2 \phi_{z'z} + \sin^2 \phi_{z'z}) \\
&= \sin^2 \theta \\
&= (2 \sin(\theta/2) \cos(\theta/2))^2 \\
&= \left(\frac{2a_z a_{z'}}{a_z^2 + a_{z'}^2} \right)^2 \\
&= \left(\frac{2}{a_z/a_{z'} + a_{z'}/a_z} \right)^2 \\
&\geq \left(\frac{2 \cdot \frac{1}{n^{3k} 2^n}}{2n^{3k}/2^n} \right)^2 \\
&\geq \left(1/n^{6k} \right)^2
\end{aligned}$$

where the fourth step follows from the definition of θ and the fifth step follows from bounds on a_z and $a_{z'}$ implied by $z, z' \in \mathbb{L} \cap \mathbb{U}$. Recall that $p(n) > n^{64k}$ and $k > 6$. Therefore, for large enough n , $p(n) > 4\sqrt{2/n^{6k}}$. Equation (12) thus implies that we can apply Theorem 7.3 with probability at least $1 - 4e^{-2n}$ when we set $(x, y) = (A, B)$, $(x^*, y^*) = (u, v)$, $\gamma = 1/n^{6k}$ and $\gamma' = 4\sqrt{2}/p(n)^{1/8}$. Formally,

$$\begin{aligned}
&\Pr \left[\left| e^{-i \cdot \arctan 2(v, u)} - e^{-i \cdot \arctan 2(B, A)} \right| \leq \left(\frac{8\sqrt{2}n^{6k}}{p(n)^{1/8}} \right) \right] \geq 1 - 4e^{-2n} \\
&\implies \Pr \left[\left| e^{-i \cdot \mathcal{U}_z(z')} - e^{-i \cdot \phi_{z'z}} \right| > \frac{8\sqrt{2}n^{6k}}{p(n)^{1/8}} \right] \leq \text{negl}(n)
\end{aligned}$$

which concludes the proof of the subclaim. \square

Next we show that if $z \in \mathbb{L} \cap \mathbb{U} \cap \mathbb{Z}$, then the oracle responses in the phase step of Aaronson's synthesis can be answered by the estimator's outputs.

SubClaim 7.12. For all $z \in \mathbb{L} \cap \mathbb{U} \cap \mathbb{Z}$, for sufficiently large n

$$\left| | \$_{s,c} \rangle_{\mathbb{Z}} | 0 \rangle_{\mathbb{V}} | \tau_1 \rangle_{\mathbb{V}'} - e^{-i\phi_z^{z,c}} (U_z)^\dagger P' U_z | \$_{s,c}^* \rangle_{\mathbb{Z}} | 0 \rangle_{\mathbb{V}} | \tau_1 \rangle_{\mathbb{V}'} \right| \leq \frac{8\sqrt{2}n^{6k}}{p(n)^{1/8}} + \frac{2}{\sqrt{n^k}} + \frac{2}{\sqrt{n^{3k}}} + \frac{4\sqrt{n^{3k}}}{p(n)^{1/16}}$$

Proof. Applying a global phase to the difference does not alter the magnitude, so we multiply by $e^{-i\phi_z}$

$$\begin{aligned}
\zeta &:= \left| | \$ \rangle_{\mathbb{Z}} | 0 \rangle_{\mathbb{V}} | \tau_1 \rangle_{\mathbb{V}'} - e^{-i\phi_z} U_z^\dagger P' U_z | \$^* \rangle_{\mathbb{Z}} | 0 \rangle_{\mathbb{V}} | \tau_1 \rangle_{\mathbb{V}'} \right| \\
&= \left| e^{i\phi_z} U_z | \$ \rangle_{\mathbb{Z}} | 0 \rangle_{\mathbb{V}} | \tau_1 \rangle_{\mathbb{V}'} - P' U_z | \$^* \rangle_{\mathbb{Z}} | 0 \rangle_{\mathbb{V}} | \tau_1 \rangle_{\mathbb{V}'} \right|
\end{aligned} \tag{13}$$

Next, we use the definition of U_z and P' to expand each term. Expanding $e^{i\phi_z} U_z | \$ \rangle_{\mathbb{Z}} | 0 \rangle_{\mathbb{V}} | \tau_1 \rangle_{\mathbb{V}'}|$

$$\begin{aligned}
e^{i\phi_z} U_z | \$ \rangle_{\mathbb{Z}} | 0 \rangle_{\mathbb{V}} | \tau_1 \rangle_{\mathbb{V}'} &= \sum_{z'} \alpha_{z'} e^{i\phi_z} | z' \rangle_{\mathbb{Z}} \sum_v \sqrt{\Pr[v = \mathcal{U}_z(z')]} | v \rangle_{\mathbb{V}} | \text{junk}_v \rangle_{\mathbb{V}'} \\
&= \sum_{z'} a_{z'} e^{-i\phi_{z'z}} | z' \rangle_{\mathbb{Z}} \sum_v \sqrt{\Pr[v = \mathcal{U}_z(z')]} | v \rangle_{\mathbb{V}} | \text{junk}_v \rangle_{\mathbb{V}'}
\end{aligned} \tag{14}$$

Expanding $P'U_z|s^*\rangle_Z|0\rangle_V|\tau_1\rangle_{V'}$

$$P'U_z|s^*\rangle_Z|0\rangle_V|\tau_1\rangle_{V'} = \sum_{z'} a_{z'}|z'\rangle_Z \sum_v e^{-iv} \sqrt{\Pr[v = \mathcal{U}_z(z')]}|v\rangle_V|\text{junk}_v\rangle_{V'} \quad (15)$$

Define $\mathbb{A} := \mathbb{L} \cap \mathbb{U} \cap \mathbb{Z}$ and $\mathbb{B} := \mathbb{L} \cap \mathbb{U} \cap \mathbb{Z}'_z$. Plugging (14) and (15) into (13) and squaring gives

$$\begin{aligned} \zeta^2 &= \left| \sum_{z'} a_{z'}|z'\rangle_Z \sum_v \left(e^{-i\phi_{z'z}} - e^{-iv} \right) \sqrt{\Pr[v = \mathcal{U}_z(z')]}|v\rangle_V|\text{junk}_v\rangle_{V'} \right|^2 \\ &= \sum_{z'} a_{z'}^2 \sum_v \Pr[v = \mathcal{U}_z(z')] \left| e^{-i\phi_{z'z}} - e^{-iv} \right|^2 \\ &\leq \sum_{z' \in \mathbb{B}} a_{z'}^2 \sum_v \Pr[v = \mathcal{U}_z(z')] \left| e^{-i\phi_{z'z}} - e^{-iv} \right|^2 + \sum_{z' \notin \mathbb{B}} 4a_{z'}^2 \end{aligned} \quad (16)$$

Let $\delta := \frac{8\sqrt{2}n^{6k}}{p(n)^{1/8}}$. For all $z' \in \mathbb{B}$, let $\mathbb{V}_{z'} := \{v \text{ s.t. } |e^{-iv} - e^{-i\phi_{z'z}}| \leq \delta\}$. Then SubClaim 7.11 shows that the probability that when $z' \in \mathbb{B}$, the probability that $\mathcal{U}_z(z')$ outputs $v \notin \mathbb{V}_{z'}$ is negligible. Therefore

$$\begin{aligned} &\sum_v \Pr[v = \mathcal{U}_z(z')] \left| e^{-i\phi_{z'z}} - e^{-iv} \right|^2 \\ &\leq \sum_{v \in \mathbb{V}_{z'}} \Pr[v = \mathcal{U}_z(z')] \left| e^{-i\phi_{z'z}} - e^{-iv} \right|^2 + \sum_{v \notin \mathbb{V}_{z'}} 4 \Pr[v = \mathcal{U}_z(z')] \\ &\leq \sum_{v \in \mathbb{V}_{z'}} \Pr[v = \mathcal{U}_z(z')] \left| e^{-i\phi_{z'z}} - e^{-iv} \right|^2 + \text{negl}(n) \\ &\leq \sum_{v \in \mathbb{V}_{z'}} \Pr[v = \mathcal{U}_z(z')] \delta^2 + \text{negl}(n) \\ &\leq \delta^2 + \text{negl}(n) \end{aligned}$$

Plugging back in (16) gives

$$\begin{aligned} \zeta^2 &\leq \sum_{z' \in \mathbb{B}} a_{z'}^2 \delta^2 + \sum_{z' \notin \mathbb{B}} 4a_{z'}^2 \\ &\leq \delta^2 + \sum_{z' \notin \mathbb{B}} 4a_{z'}^2 \\ &\leq \delta^2 + 4/n^k + 4/n^{3k} + 16n^{3k}/p^{1/8} \\ &\leq \frac{128n^{12k}}{p(n)^{1/4}} + 4/n^k + 4/n^{3k} + \frac{16n^{3k}}{p(n)^{1/8}} \end{aligned} \quad (17)$$

where the third step follows from parts 2, 3, and 4 of SubClaim 7.10 and the last step substitutes the value of δ . The final expression follows from taking the square root of both sides and noting that the square root function is subadditive. \square

We can now begin analyzing the algorithm $\mathcal{B}'(s, c)$, dropping the advice state from the input for notational convenience. For all t , define $|\sigma_t\rangle$ as follows:

$$|\sigma_t\rangle := (U_t)^\dagger P'U_t \left(\widetilde{M}_{\text{amp}}|s, c\rangle_{\mathbb{R}}|0^n\rangle_Z|\tau_{\text{amp}}\rangle_{\mathbb{A}} \right) |0\rangle_V|\tau_1\rangle_{V'}$$

and define ρ_t as the state on the Z register of $|\sigma_t\rangle\langle\sigma_t|$ after tracing out the remaining registers, i.e.

$$\rho_t = \text{Tr}_{\text{RAVV}'}(|\sigma_t\rangle\langle\sigma_t|)$$

Note that the output of $\mathcal{B}'(s, c)$ conditioned sampling t is $c^\dagger \rho_t c$. Next we show that $c^\dagger \rho_t c$ is close to $|\mathbb{S}_s\rangle\langle\mathbb{S}_s|$ when $t \in \mathbb{L} \cap \mathbb{U} \cap \mathbb{Z}$. By $(s, c) \in \mathcal{S}'$ and the definition of \mathbb{S} , $\Delta_{s,c} \leq 1/\sqrt{p(n)}$. Let $\delta_{\text{amp}} := (\sqrt{3n^3} + 2\sqrt{n})/p(n)^{1/4}$ (i.e. the error term in Claim 7.2 after substituting $\Delta_{s,c} \leq 1/\sqrt{p(n)}$) and let $\delta_{\text{phase}} := \frac{8\sqrt{2}n^{6k}}{p(n)^{1/8}} + \frac{2}{\sqrt{n^k}} + \frac{2}{\sqrt{n^{3k}}} + \frac{4\sqrt{n^{3k}}}{p(n)^{1/16}}$ (i.e. the error term in SubClaim 7.12).

SubClaim 7.13. *If $t \in \mathbb{L} \cap \mathbb{U} \cap \mathbb{Z}$ then*

$$\langle\mathbb{S}_s|c^\dagger \cdot \rho_t \cdot c|\mathbb{S}_s\rangle \geq 1 - \delta_{\text{amp}} - \delta_{\text{phase}}$$

Proof. Let

$$\begin{aligned} |\psi_1\rangle &:= |s, c\rangle|0^n\rangle|\tau_{\text{amp}}\rangle|0\rangle|\tau_1\rangle \\ |\psi_2\rangle &:= |s, c\rangle|\mathbb{S}_{s,c}^*\rangle|\tau_{\text{amp}}\rangle|0\rangle|\tau_1\rangle \\ |\psi_3\rangle &:= e^{i\phi_t}|s, c\rangle|\mathbb{S}_{s,c}\rangle_Z|\tau_{\text{amp}}\rangle|0\rangle|\tau_1\rangle \\ M_1 &:= \widetilde{M}_{\text{amp}} \\ M_2 &:= (U_t^{s,c})^\dagger P' U_t^{s,c}. \end{aligned}$$

By $(s, c) \in \mathcal{S}'$ and the definition of \mathbb{S} , $\Delta_{s,c} \leq 1/\sqrt{p(n)}$. Therefore, Claim 7.2 can be restated as

$$|M_1|\psi_1\rangle - |\psi_2\rangle| \leq \delta_{\text{amp}}.$$

Similarly, SubClaim 7.12 can be restated as

$$|M_2|\psi_2\rangle - |\psi_3\rangle| \leq \delta_{\text{phase}}.$$

By the triangle inequality and substituting the last two equations in the last step,

$$\begin{aligned} |M_2 M_1 |\psi_1\rangle - |\psi_3\rangle| &\leq |M_2 M_1 |\psi_1\rangle - M_2 |\psi_2\rangle| + |M_2 |\psi_2\rangle - |\psi_3\rangle| \\ &= |M_1 |\psi_1\rangle - |\psi_2\rangle| + |M_2 |\psi_2\rangle - |\psi_3\rangle| \\ &\leq \delta_{\text{amp}} + \delta_{\text{phase}}. \end{aligned}$$

By Theorem 4.2,

$$\text{TD}(M_2 M_1 |\psi_1\rangle\langle\psi_1| M_1^\dagger M_2^\dagger, |\psi_3\rangle\langle\psi_3|) \leq \delta_{\text{amp}} + \delta_{\text{phase}}.$$

Noting that $|\sigma_t\rangle$ is defined as $M_2 M_1 |\psi_1\rangle$

$$\text{TD}(|\sigma_t\rangle\langle\sigma_t|, |\psi_3\rangle\langle\psi_3|) \leq \delta_{\text{amp}} + \delta_{\text{phase}}.$$

We can trace out all but the Z register from both states without increasing the trace distance. Therefore,

$$\text{TD}(\rho_t, |\mathbb{S}_{s,c}\rangle\langle\mathbb{S}_{s,c}|) \leq \delta_{\text{amp}} + \delta_{\text{phase}}.$$

Consider the projector $|\mathbb{S}_{s,c}\rangle\langle\mathbb{S}_{s,c}|$. The projection succeeds on the state $|\mathbb{S}_{s,c}\rangle\langle\mathbb{S}_{s,c}|$ with probability 1, so it must succeed on ρ_t with probability atleast $1 - \text{TD}(\rho_t, |\mathbb{S}_{s,c}\rangle\langle\mathbb{S}_{s,c}|) \geq 1 - \delta_{\text{amp}} - \delta_{\text{phase}}$. Therefore,

$$\langle\mathbb{S}_s|c^\dagger \rho_t c|\mathbb{S}_s\rangle = \langle\mathbb{S}_{s,c}|\rho_t|\mathbb{S}_{s,c}\rangle \geq 1 - \delta_{\text{amp}} - \delta_{\text{phase}}.$$

□

Next we show that with high probability, $t \in \mathbb{L} \cap \mathbb{U} \cap \mathbb{Z}$. Let $\delta_{\text{samp}} := 3/p(n)^{1/4} - 1/n^k - 1/n^{3k}$ (i.e. the sum of error terms from parts 1, 2, and 3 of SubClaim 7.10).

SubClaim 7.14. *Let t be the outcome when measuring the Z register of $\widetilde{M}_{\text{amp}}|s, c\rangle_{\mathbb{R}}|0^n\rangle_{\mathbb{Z}}|\tau_{\text{amp}}\rangle_{\mathbb{A}}$ in the computational basis. Then*

$$\Pr[t \in \mathbb{L} \cap \mathbb{U} \cap \mathbb{Z}] \geq 1 - \delta_{\text{samp}} - \delta_{\text{amp}}$$

Proof. Consider the probability of obtaining a string t' upon measuring the second register of $|s, c\rangle|\$\ast_{s,c}\rangle|\tau_{\text{amp}}\rangle$.

$$\Pr[t'] = \left| \langle t' | \sum_z a_z |z\rangle \right|^2 = (a_z)^2$$

Therefore,

$$\begin{aligned} \Pr[t' \in \mathbb{L} \cap \mathbb{U} \cap \mathbb{Z}] &= \sum_{z \in \mathbb{L} \cap \mathbb{U} \cap \mathbb{Z}} (a_z)^2 \\ &\geq 1 - \sum_{z \notin \mathbb{L}} (a_z)^2 - \sum_{z \notin \mathbb{U}} (a_z)^2 - \sum_{z \notin \mathbb{Z}} (a_z)^2 \\ &\geq 1 - 3/p(n)^{1/4} - 1/n^k - 1/n^{3k} = 1 - \delta_{\text{samp}} \end{aligned}$$

where the last step follows from parts 1, 2, and 3 of SubClaim 7.10, noting that $a_z = |\alpha_z|$. By Claim 7.2 and Theorem 4.2 we know that $\widetilde{M}_{\text{amp}}|s, c\rangle|0\rangle|\tau_{\text{amp}}\rangle$ and $|s, c\rangle|\$\ast_{s,c}\rangle|\tau_{\text{amp}}\rangle$ are at most δ_{amp} apart in trace distance. Therefore, the output distributions upon measuring each state can be at most δ_{amp} far. Therefore,

$$\Pr[t \in \mathbb{L} \cap \mathbb{U} \cap \mathbb{Z}] \geq 1 - \delta_{\text{samp}} - \delta_{\text{amp}}$$

□

Putting together SubClaim 7.14 and SubClaim 7.13 shows that the expected overlap of $|\$\rangle\langle\$_s|$ and $\mathcal{B}'(s, c)$ is high.

SubClaim 7.15.

$$\mathbb{E}_{\mathcal{B}'} [\langle\$_s|\mathcal{B}'(s, c)|\$_s\rangle] \geq 1 - \delta_{\text{samp}} - 2\delta_{\text{amp}} - \delta_{\text{phase}}$$

Proof. SubClaim 7.13 shows that if t sampled by \mathcal{B}' is in $\mathbb{L} \cap \mathbb{U} \cap \mathbb{Z}$ then the overlap is at least $1 - \delta_{\text{amp}} - \delta_{\text{phase}}$, while SubClaim 7.14 shows that $t \in \mathbb{L} \cap \mathbb{U} \cap \mathbb{Z}$ occurs with probability at least $1 - \delta_{\text{samp}} - \delta_{\text{amp}}$. Putting these together

$$\begin{aligned} \mathbb{E}_{\mathcal{B}'} [\langle\$_s|\mathcal{B}'(s, c)|\$_s\rangle] &\geq \Pr[t \in \mathbb{L} \cap \mathbb{U} \cap \mathbb{Z}] \cdot (1 - \delta_{\text{amp}} - \delta_{\text{phase}}) \\ &\geq (1 - \delta_{\text{samp}} - \delta_{\text{amp}}) \cdot (1 - \delta_{\text{amp}} - \delta_{\text{phase}}) \\ &\geq 1 - \delta_{\text{samp}} - 2\delta_{\text{amp}} - \delta_{\text{phase}} \end{aligned}$$

where the first step follows from SubClaim 7.13 and the second step follows from SubClaim 7.14.

□

Finally, we show that \mathcal{B} achieves the claimed bound. We note that SubClaim 7.15 applies for arbitrary fixed $(s, c) \in \mathcal{S}'$ and SubClaim 7.7 shows that for s sampled by \mathcal{G} and $c \leftarrow \mathcal{C}_n$ the

probability that $(s, c) \in \mathbb{S}$ is atleast $1 - 1/p(n)^{1/2} - 1/n^k$. Putting them together

$$\begin{aligned} \mathbb{E}_{s, |\mathbb{S}_s\rangle \leftarrow \mathcal{G}(1^n)} [\langle \mathbb{S}_s | \mathcal{B}(s, |\tau_{\text{synth}}\rangle) | \mathbb{S}_s \rangle] &\geq \Pr[(s, c) \in \mathbb{S}'] \cdot (1 - \delta_{\text{samp}} - 2\delta_{\text{amp}} - \delta_{\text{phase}}) \\ &\geq (1 - 1/p(n)^{1/2} - 1/n^k) \cdot (1 - \delta_{\text{samp}} - 2\delta_{\text{amp}} - \delta_{\text{phase}}) \\ &\geq 1 - 1/p(n)^{1/2} - 1/n^k - \delta_{\text{samp}} - 2\delta_{\text{amp}} - \delta_{\text{phase}} \end{aligned}$$

All that remains it to simplify the expression. Substituting the definitions of δ_{samp} , δ_{amp} , and δ_{phase} , noting that $k > 6$ and $p(n) > n^{64k}$ and simplifying for large enough n

$$1 - 1/p(n)^{1/2} - 1/n^k - \delta_{\text{samp}} - 2\delta_{\text{amp}} - \delta_{\text{phase}} \geq 1 - 3/\sqrt{n^k}$$

Finally, since $n^k \geq q(n)^3$

$$\mathbb{E}_{s, |\mathbb{S}_s\rangle \leftarrow \mathcal{G}(1^n)} [\langle \mathbb{S}_s | \mathcal{B}(s, |\tau_{\text{synth}}\rangle) | \mathbb{S}_s \rangle] \geq 1 - 3/q(n)^{3/2} \geq 1 - 1/q(n)$$

which concludes the proof of the Claim. \square

Claim 7.3 shows the existence of an algorithm that contradicts the security of \mathcal{G} which concludes the proof of the theorem. \square

Next, we prove that the existence of distributional one-way puzzles implies the existence of (standard) state puzzles.

Theorem 7.4. *If one-way puzzles (Definition 4.2) exist then state puzzles (Definition 7.1) exist.*

Since Theorem 4.3 shows that distributional one-way puzzles can be amplified to obtain (strong) one-way puzzles, the following is a corollary of Theorem 7.4.

Corollary 7.2. *If $1/q(n)$ -distributional one-way puzzles (Definition 4.3) exist for some non-zero polynomial $q(\cdot)$ then state puzzles (Definition 7.1) exist.*

of Theorem 7.4. Let $(\text{Samp}(1^n), \text{Ver})$ be a one-way puzzle that samples n bit puzzles and n bit keys. Without loss of generality, we may assume that $\text{Samp}(1^n)$ is the algorithm that first applies a unitary U_n to $|0\rangle$ to obtain

$$U_n|0\rangle = \sum_{s,k} \sqrt{p_{s,k}} |\mu_{s,k}\rangle |s\rangle_S |k\rangle_K$$

where $\{\mu_{s,k}\}_{s,k}$ are a set of normalised states, s and k are an n -bit puzzle and an n -bit key respectively output by $\text{Samp}(1^n)$ with probability $p_{s,k}$. This is followed by a classical basis measurement of registers S and K to obtain puzzle s and key k .

Let $\mathcal{G}(1^n)$ be the following algorithm

1. Apply U_n to $|0\rangle$ to obtain $\sum_{s,k} \sqrt{p_{s,k}} |\mu_{s,k}\rangle |s\rangle_S |k\rangle_K$
2. Measure S in the classical basis to obtain string s and residual state $|\psi_s\rangle$.
3. Return $(s, |\psi_s\rangle)$.

We will prove that $\mathcal{G}(1^n)$ is a state puzzle, which suffices to prove the theorem. Suppose for the sake of contradiction that $\mathcal{G}(1^n)$ is not a state puzzle. Then there exists a polynomial $p(n)$. QPT $\mathcal{A} = \{\mathcal{A}_n\}_{n \in \mathbb{N}}$, and an advice ensemble $|\tau\rangle = \{|\tau_n\rangle\}_{n \in \mathbb{N}}$ such that for all large enough $n \in \mathbb{N}$,

$$\mathbb{E}_{(s, |\psi_s\rangle) \leftarrow \mathcal{G}(1^n)} [\langle \psi_s | \mathcal{A}(|\tau\rangle, s) | \psi_s \rangle] \geq \frac{1}{p(n)}$$

Fix any such adversary \mathcal{A} , any such advice ensemble $|\tau\rangle$, and any such large enough $n \in \mathbb{N}$. We will drop the advice from the notation since it is always implicitly provided to the adversary.

Define the algorithm \mathcal{A}' that takes input s , computes $\mathcal{A}(s)$, and outputs the outcome of measuring the K register of $\mathcal{A}(s)$ in the computational basis. We will show that \mathcal{A}' contradicts the security of the one-way puzzle ($\text{Samp}(1^n), \text{Ver}$).

Let p_s be the probability that $\text{Samp}(1^n)$ samples puzzle s . Note that this is identical to the probability that $\mathcal{G}(1^n)$ samples s since in both cases s is generated the same way. We first show that the expected trace distance between $|\psi_s\rangle\langle\psi_s|$ and $\mathcal{A}(s)$ is at most $\sqrt{1 - 1/p(n)}$.

Claim 7.4.

$$\mathbb{E}_{(s, |\psi_s\rangle) \leftarrow \mathcal{G}(1^n)} \left[\text{TD}\left(|\psi_s\rangle\langle\psi_s|, \mathcal{A}(s)\right) \right] \leq \sqrt{1 - 1/p(n)}$$

Proof. For any s , the fidelity $F(\mathcal{A}(s), |\psi_s\rangle)$ of $\mathcal{A}(s)$ and $|\psi_s\rangle$ is $\sqrt{\langle \psi_s | \mathcal{A}(s) | \psi_s \rangle}$. Therefore, by Uhlmann's Theorem, for any s

$$\begin{aligned} \text{TD}\left(|\psi_s\rangle\langle\psi_s|, \mathcal{A}(s)\right) &\leq \sqrt{1 - F(\mathcal{A}(s), |\psi_s\rangle)^2} \\ &= \sqrt{1 - \langle \psi_s | \mathcal{A}(s) | \psi_s \rangle} \end{aligned}$$

Expressing $\mathbb{E}_{(s, |\psi_s\rangle) \leftarrow \mathcal{G}(1^n)} \left[\text{TD}\left(|\psi_s\rangle\langle\psi_s|, \mathcal{A}(s)\right) \right]$ as a sum over s and applying the above

$$\begin{aligned} \mathbb{E}_{(s, |\psi_s\rangle) \leftarrow \mathcal{G}(1^n)} \left[\text{TD}\left(|\psi_s\rangle\langle\psi_s|, \mathcal{A}(s)\right) \right] &= \sum_s p_s \cdot \text{TD}\left(|\psi_s\rangle\langle\psi_s|, \mathcal{A}(s)\right) \\ &\leq \sum_s p_s \cdot \sqrt{1 - \langle \psi_s | \mathcal{A}(s) | \psi_s \rangle} \end{aligned}$$

Applying Jensen's inequality

$$\begin{aligned} \mathbb{E}_{(s, |\psi_s\rangle) \leftarrow \mathcal{G}(1^n)} \left[\text{TD}\left(|\psi_s\rangle\langle\psi_s|, \mathcal{A}(s)\right) \right] &\leq \sqrt{\sum_s p_s \cdot (1 - \langle \psi_s | \mathcal{A}(s) | \psi_s \rangle)} \\ &= \sqrt{1 - \sum_s p_s \cdot \langle \psi_s | \mathcal{A}(s) | \psi_s \rangle} \end{aligned}$$

Now, we can rewrite the fact that $\mathbb{E}_{(s, |\psi_s\rangle) \leftarrow \mathcal{G}(1^n)} [\langle \psi_s | \mathcal{A}(s) | \psi_s \rangle] \geq \frac{1}{p(n)}$ as a sum over s .

$$\sum_s p_s \cdot \langle \psi_s | \mathcal{A}(s) | \psi_s \rangle \geq 1/p(n)$$

which when plugged into the previous inequality implies

$$\mathbb{E}_{(s, |\psi_s\rangle) \leftarrow \mathcal{G}(1^n)} \left[\text{TD}\left(|\psi_s\rangle\langle\psi_s|, \mathcal{A}(s)\right) \right] \leq \sqrt{1 - 1/p(n)}$$

□

Claim 7.5.

$$\Pr_{(s,k) \leftarrow \text{Samp}(1^n)} [\text{Ver}(s, \mathcal{A}'(s)) = 1] \geq 1/3p(n)$$

Proof. By the correctness of the one-way puzzle

$$\Pr_{(s,k) \leftarrow \text{Samp}(1^n)} [\text{Ver}(s, k) = 1] \geq 1 - \text{negl}(n)$$

Let M be an algorithm that takes input $|\psi\rangle$ and returns the result k of measuring the K register in the computational basis. The distribution over s and k obtained by sampling $(s, |\psi_s\rangle)$ from $\mathcal{G}(1^n)$ and sampling k from $M(|\psi_s\rangle\langle\psi_s|)$ is therefore identical to the distribution obtained by sampling (s, k) from $\text{Samp}(1^n)$. We can express the probability that $\mathcal{A}'(s)$ successfully outputs a key that passes verification as a sum over s . Therefore

$$\Pr_{\substack{(s, |\psi_s\rangle) \leftarrow \mathcal{G}(1^n) \\ k \leftarrow M(|\psi_s\rangle\langle\psi_s|)}} [\text{Ver}(s, k) = 1] \geq 1 - \text{negl}(n)$$

which may be rewritten as

$$\Pr_{(s, |\psi_s\rangle) \leftarrow \mathcal{G}(1^n)} [\text{Ver}(s, M(|\psi_s\rangle\langle\psi_s|)) = 1] \geq 1 - \text{negl}(n)$$

and expressed as a sum over s as follows.

$$\sum_s p_s \cdot \Pr[\text{Ver}(s, M(|\psi_s\rangle\langle\psi_s|)) = 1] \geq 1 - \text{negl}(n)$$

For any s and any state ρ ,

$$\left| \Pr[\text{Ver}(s, M(|\psi_s\rangle\langle\psi_s|)) = 1] - \Pr[\text{Ver}(s, M(\rho)) = 1] \right| \leq \text{TD}(|\psi_s\rangle\langle\psi_s|, \rho)$$

which means that we can replace $|\psi_s\rangle\langle\psi_s|$ with $\mathcal{A}(s)$ at the cost of an error of $\text{TD}(|\psi_s\rangle\langle\psi_s|, \mathcal{A}(s))$, i.e.

$$\begin{aligned} 1 - \text{negl}(n) &\leq \sum_s p_s \cdot \left(\Pr[\text{Ver}(s, M(\mathcal{A}(s))) = 1] + \text{TD}(|\psi_s\rangle\langle\psi_s|, \mathcal{A}(s)) \right) \\ &\leq \sum_s p_s \cdot \left(\Pr[\text{Ver}(s, M(\mathcal{A}(s))) = 1] \right) + \mathbb{E}_{s, |\psi_s\rangle \leftarrow \mathcal{G}(1^n)} \left[\text{TD}(|\psi_s\rangle\langle\psi_s|, \mathcal{A}(s)) \right] \end{aligned}$$

The first term is exactly $\Pr_{s,k \leftarrow \text{Samp}(1^n)} [\text{Ver}(s, \mathcal{A}'(s)) = 1]$ and the second term is shown in Claim 7.4 upper bounded by $\sqrt{1 - p(n)}$. Therefore,

$$\begin{aligned} \Pr_{s,k \leftarrow \text{Samp}(1^n)} [\text{Ver}(s, \mathcal{A}'(s)) = 1] &\geq 1 - \text{negl}(n) - \sqrt{1 - p(n)} \\ &\geq 1 - \text{negl}(n) - (1 - p(n)/2) \\ &\geq p(n)/2 - \text{negl}(n) \\ &\geq p(n)/3 \end{aligned}$$

where the last step holds for large enough n . □

This shows that \mathcal{A}' contradicts the security of the one-way puzzle, concluding the proof of the theorem. □

We also have the following straightforward corollary, which follows from the implication of state puzzles from quantum money mini-schemes (and other unclonable primitives).

Corollary 7.3. *Quantum money mini-schemes imply one-way puzzles and quantum bit commitments.*

Acknowledgments

We thank Scott Aaronson, Lijie Chen and William Kretschmer for helpful conversations about the (im)possibility of quantumly efficiently sampling matrices jointly with their permanents. We thank Alexandra (Sasha) Levinshteyn for help with typesetting parts of this manuscript. Finally, we thank Daniel Apon, Tomoyuki Morimae, Barak Nehoran, Luowen Qian and Amit Sahai for useful comments on the writeup.

The authors were supported in part by AFOSR, NSF 2112890, NSF CNS-2247727 and a Google Research Scholar award. This material is based upon work supported by the Air Force Office of Scientific Research under award number FA9550-23-1-0543.

References

- [AA11] Scott Aaronson and Alex Arkhipov. The computational complexity of linear optics. In *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing, STOC '11*, page 333–342, New York, NY, USA, 2011. Association for Computing Machinery.
- [AA14] Scott Aaronson and Alex Arkhipov. Bosonsampling is far from uniform. *Quantum Inf. Comput.*, 14(15-16):1383–1423, 2014.
- [Aar06] Scott Aaronson. Oracles are subtle but not malicious. In *21st Annual IEEE Conference on Computational Complexity (CCC 2006), 16-20 July 2006, Prague, Czech Republic*, pages 340–354. IEEE Computer Society, 2006.
- [Aar16] Scott Aaronson. The complexity of quantum states and transformations: From quantum money to black holes, 2016.
- [AC12] Scott Aaronson and Paul F. Christiano. Quantum money from hidden subspaces. In Howard J. Karloff and Toniann Pitassi, editors, *Proceedings of the 44th Symposium on Theory of Computing Conference, STOC 2012, New York, NY, USA, May 19 - 22, 2012*, pages 41–60. ACM, 2012.
- [AQY22] Prabhanjan Ananth, Luowen Qian, and Henry Yuen. Cryptography from pseudorandom quantum states. In Yevgeniy Dodis and Thomas Shrimpton, editors, *Advances in Cryptology - CRYPTO 2022 - 42nd Annual International Cryptology Conference, CRYPTO 2022, Santa Barbara, CA, USA, August 15-18, 2022, Proceedings, Part I*, volume 13507 of *Lecture Notes in Computer Science*, pages 208–236. Springer, 2022.
- [BB84] Charles H Bennett and Gilles Brassard. Quantum cryptography: Public key distribution and coin tossing. In *Proceedings of the IEEE International Conference on Computers, Systems, and Signal Processing*, pages 175–179, 1984.
- [BCKM21] James Bartusek, Andrea Coladangelo, Dakshita Khurana, and Fermi Ma. One-way functions imply secure computation in a quantum world. In Tal Malkin and Chris Peikert, editors, *Advances in Cryptology – CRYPTO 2021, Part I*, volume 12825 of *Lecture Notes in Computer Science*, pages 467–496, Virtual Event, August 16–20, 2021. Springer, Heidelberg, Germany.
- [BFLL21] Adam Bouland, Bill Fefferman, Zeph Landau, and Yunchao Liu. Noise and the frontier of quantum supremacy. In *62nd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2021, Denver, CO, USA, February 7-10, 2022*, pages 1308–1317. IEEE, 2021.

- [BFNV19] Adam Bouland, Bill Fefferman, Chinmay Nirkhe, and Umesh V. Vazirani. "quantum supremacy" and the complexity of random circuit sampling. In Avrim Blum, editor, *10th Innovations in Theoretical Computer Science Conference, ITCS 2019, January 10-12, 2019, San Diego, California, USA*, volume 124 of *LIPICs*, pages 15:1–15:2. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.
- [BHHP24] John Bostanci, Jonas Haferkamp, Dominik Hangleiter, and Alexander Poremba. Efficient quantum pseudorandomness from hamiltonian phase states. 2024.
- [BIS⁺18] Sergio Boixo, Sergei V. Isakov, Vadim N. Smelyanskiy, Ryan Babbush, Nan Ding, Zhang Jiang, Michael J. Bremner, John M. Martinis, and Hartmut Neven. Characterizing quantum supremacy in near-term devices. *Nature Physics*, 14(6):595–600, Jun 2018.
- [BJ24] Rishabh Batra and Rahul Jain. Commitments are equivalent to one-way state generators. *CoRR*, abs/2404.03220, 2024.
- [BJS11] Michael J. Bremner, Richard Jozsa, and Dan J. Shepherd. Classical simulation of commuting quantum computations implies collapse of the polynomial hierarchy. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 467(2126):459–472, 2011.
- [BMS16] Michael J. Bremner, Ashley Montanaro, and Dan J. Shepherd. Average-case complexity versus approximate simulation of commuting quantum computations. *Phys. Rev. Lett.*, 117:080501, Aug 2016.
- [Bra23] Zvika Brakerski. Black-hole radiation decoding is quantum cryptography. In *Advances in Cryptology – CRYPTO 2023: 43rd Annual International Cryptology Conference, CRYPTO 2023, Santa Barbara, CA, USA, August 20–24, 2023, Proceedings, Part V*, page 37–65, Berlin, Heidelberg, 2023. Springer-Verlag.
- [CGG⁺23] Bruno Cavalari, Eli Goldin, Matthew Gray, Peter Hall, Yanyi Liu, and Angelos Pelecanos. On the computational hardness of quantum one-wayness. *CoRR*, abs/2312.08363, 2023.
- [CGG24] Kai-Min Chung, Eli Goldin, and Matthew Gray. On central primitives for quantum cryptography with classical communication. In Leonid Reyzin and Douglas Stebila, editors, *Advances in Cryptology - CRYPTO 2024 - 44th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2024, Proceedings, Part VII*, volume 14926 of *Lecture Notes in Computer Science*, pages 215–248. Springer, 2024.
- [CGGH24] Bruno P. Cavalari, Eli Goldin, Matthew Gray, and Peter Hall. A meta-complexity characterization of quantum cryptography, 2024.
- [FM17] Keisuke Fujii and Tomoyuki Morimae. Commuting quantum circuits and complexity of ising partition functions. *New Journal of Physics*, 19(3):033003, mar 2017.
- [GBA⁺22] Daniel Grier, Daniel J. Brod, Juan Miguel Arrazola, Marcos Benicio de Andrade Alonso, and Nicolás Quesada. The Complexity of Bipartite Gaussian Boson Sampling. *Quantum*, 6:863, November 2022.

- [GK16] Shafi Goldwasser and Yael Tauman Kalai. Cryptographic assumptions: A position paper. In Eyal Kushilevitz and Tal Malkin, editors, *Theory of Cryptography - 13th International Conference, TCC 2016-A, Tel Aviv, Israel, January 10-13, 2016, Proceedings, Part I*, volume 9562 of *Lecture Notes in Computer Science*, pages 505–522. Springer, 2016.
- [GLSV21] Alex B. Grilo, Huijia Lin, Fang Song, and Vinod Vaikuntanathan. Oblivious transfer is in minicrypt. In Anne Canteaut and François-Xavier Standaert, editors, *Advances in Cryptology - EUROCRYPT 2021 - 40th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, October 17-21, 2021, Proceedings, Part II*, volume 12697 of *Lecture Notes in Computer Science*, pages 531–561. Springer, 2021.
- [HE23] Dominik Hangleiter and Jens Eisert. Computational advantage of quantum random sampling. *Rev. Mod. Phys.*, 95:035001, Jul 2023.
- [HKP20] Hsin-Yuan Huang, Richard Kueng, and John Preskill. Predicting many properties of a quantum system from very few measurements. *Nature Physics*, 16(10):1050–1057, Oct 2020.
- [HM24] Taiga Hiroka and Tomoyuki Morimae. Quantum cryptography from meta-complexity. Cryptology ePrint Archive, Paper 2024/1539, 2024.
- [IL89] Russell Impagliazzo and Michael Luby. One-way functions are essential for complexity based cryptography (extended abstract). In *30th Annual Symposium on Foundations of Computer Science*, pages 230–235, Research Triangle Park, NC, USA, October 30 – November 1, 1989. IEEE Computer Society Press.
- [IL90] Russell Impagliazzo and Leonid A. Levin. No better ways to generate hard NP instances than picking uniformly at random. In *31st Annual Symposium on Foundations of Computer Science*, pages 812–821, St. Louis, MO, USA, October 22–24, 1990. IEEE Computer Society Press.
- [ILL89] Russell Impagliazzo, Leonid A. Levin, and Michael Luby. Pseudo-random generation from one-way functions (extended abstracts). In *21st Annual ACM Symposium on Theory of Computing*, pages 12–24, Seattle, WA, USA, May 15–17, 1989. ACM Press.
- [INN⁺22] Sandy Irani, Anand Natarajan, Chinmay Nirkhe, Sujit Rao, and Henry Yuen. Quantum search-to-decision reductions and the state synthesis problem. In Shachar Lovett, editor, *37th Computational Complexity Conference, CCC 2022, July 20-23, 2022, Philadelphia, PA, USA*, volume 234 of *LIPICs*, pages 5:1–5:19. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022.
- [JLS18] Zhengfeng Ji, Yi-Kai Liu, and Fang Song. Pseudorandom quantum states. In Hovav Shacham and Alexandra Boldyreva, editors, *Advances in Cryptology – CRYPTO 2018, Part III*, volume 10993 of *Lecture Notes in Computer Science*, pages 126–152, Santa Barbara, CA, USA, August 19–23, 2018. Springer, Heidelberg, Germany.
- [KL98] E Knill and R Laflamme. Power of one bit of quantum information. *Physical Review Letters*, 81(25), 12 1998.

- [KMM21] Yasuhiro Kondo, Ryuhei Mori, and Ramis Movassagh. Quantum supremacy and hardness of estimating output probabilities of quantum circuits. In *62nd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2021, Denver, CO, USA, February 7-10, 2022*, pages 1296–1307. IEEE, 2021.
- [KQST23] William Kretschmer, Luowen Qian, Makrand Sinha, and Avishay Tal. Quantum cryptography in algorithmica. In Barna Saha and Rocco A. Servedio, editors, *Proceedings of the 55th Annual ACM Symposium on Theory of Computing, STOC 2023, Orlando, FL, USA, June 20-23, 2023*, pages 1589–1602. ACM, 2023.
- [Kre21] William Kretschmer. Quantum pseudorandomness and classical complexity. In Min-Hsiu Hsieh, editor, *16th Conference on the Theory of Quantum Computation, Communication and Cryptography, TQC 2021, July 5-8, 2021, Virtual Conference*, volume 197 of *LIPICs*, pages 2:1–2:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.
- [Kro22] Hari Krovi. Average-case hardness of estimating probabilities of random quantum circuits with a linear scaling in the error exponent, 2022.
- [KT24] Dakshita Khurana and Kabir Tomer. Commitments from quantum one-wayness. In Bojan Mohar, Igor Shinkar, and Ryan O’Donnell, editors, *Proceedings of the 56th Annual ACM Symposium on Theory of Computing, STOC 2024, Vancouver, BC, Canada, June 24-28, 2024*, pages 968–978. ACM, 2024.
- [LC97] Hoi-Kwong Lo and Hoi Fung Chau. Is quantum bit commitment really possible? *Physical Review Letters*, 78(17):3410, 1997.
- [LMW24] Alex Lombardi, Fermi Ma, and John Wright. A one-query lower bound for unitary synthesis and breaking quantum cryptography. In Bojan Mohar, Igor Shinkar, and Ryan O’Donnell, editors, *Proceedings of the 56th Annual ACM Symposium on Theory of Computing, STOC 2024, Vancouver, BC, Canada, June 24-28, 2024*, pages 979–990. ACM, 2024.
- [LR86] Michael Luby and Charles Rackoff. Pseudo-random permutation generators and cryptographic composition. In *18th Annual ACM Symposium on Theory of Computing*, pages 356–363, Berkeley, CA, USA, May 28–30, 1986. ACM Press.
- [May97] Dominic Mayers. Unconditionally secure quantum bit commitment is impossible. *Physical review letters*, 78(17):3414, 1997.
- [Mel24] Antonio Anna Mele. Introduction to haar measure tools in quantum information: A beginner’s tutorial. *Quantum*, 8:1340, 2024.
- [MFF14] Tomoyuki Morimae, Keisuke Fujii, and Joseph F. Fitzsimons. Hardness of classically simulating the one-clean-qubit model. *Phys. Rev. Lett.*, 112:130502, Apr 2014.
- [MNY24] Tomoyuki Morimae, Barak Nehoran, and Takashi Yamakawa. Unconditionally secure commitments with quantum auxiliary inputs. In Leonid Reyzin and Douglas Stebila, editors, *Advances in Cryptology - CRYPTO 2024 - 44th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2024, Proceedings, Part VII*, volume 14926 of *Lecture Notes in Computer Science*, pages 59–92. Springer, 2024.
- [Mov23] Ramis Movassagh. The hardness of random quantum circuits. *Nature Physics*, 19(11):1719–1724, Nov 2023.

- [MSY24] Tomoyuki Morimae, Yuki Shirakawa, and Takashi Yamakawa. Cryptographic characterization of quantum advantage. *Corr*, abs/2410.00499, 2024.
- [MY22a] Tomoyuki Morimae and Takashi Yamakawa. One-wayness in quantum cryptography. *CoRR*, abs/2210.03394, 2022.
- [MY22b] Tomoyuki Morimae and Takashi Yamakawa. Quantum commitments and signatures without one-way functions. In *Advances in Cryptology – CRYPTO 2022, Part I*, Lecture Notes in Computer Science, pages 269–295, Santa Barbara, CA, USA, August 2022. Springer, Heidelberg, Germany.
- [MY24] Tomoyuki Morimae and Takashi Yamakawa. Quantum advantage from one-way functions. In Leonid Reyzin and Douglas Stebila, editors, *Advances in Cryptology - CRYPTO 2024 - 44th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2024, Proceedings, Part V*, volume 14924 of *Lecture Notes in Computer Science*, pages 359–392. Springer, 2024.
- [OW93] Rafail Ostrovsky and Avi Wigderson. One-way functions are essential for non-trivial zero-knowledge. In *Second Israel Symposium on Theory of Computing Systems, ISTCS 1993, Natanya, Israel, June 7-9, 1993, Proceedings*, pages 3–17. IEEE Computer Society, 1993.
- [Qia24] Luowen Qian. Unconditionally secure quantum commitments with preprocessing. In Leonid Reyzin and Douglas Stebila, editors, *Advances in Cryptology - CRYPTO 2024 - 44th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2024, Proceedings, Part VII*, volume 14926 of *Lecture Notes in Computer Science*, pages 38–58. Springer, 2024.
- [QRZ24] Luowen Qian, Justin Raizes, and Mark Zhandry. Hard quantum extrapolations in quantum cryptography. *Corr*, abs/2409.16516, 2024.
- [Ros24] Gregory Rosenthal. Efficient quantum state synthesis with one query. In David P. Woodruff, editor, *Proceedings of the 2024 ACM-SIAM Symposium on Discrete Algorithms, SODA 2024, Alexandria, VA, USA, January 7-10, 2024*, pages 2508–2534. SIAM, 2024.
- [RT22] Ran Raz and Avishay Tal. Oracle separation of BQP and PH. *J. ACM*, 69(4):30:1–30:21, 2022.
- [SB09] Dan Shepherd and Michael J. Bremner. Temporally unstructured quantum computation. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 465(2105):1413–1439, 2009.
- [Sto83] Larry J. Stockmeyer. The complexity of approximate counting (preliminary version). In David S. Johnson, Ronald Fagin, Michael L. Fredman, David Harel, Richard M. Karp, Nancy A. Lynch, Christos H. Papadimitriou, Ronald L. Rivest, Walter L. Ruzzo, and Joel I. Seiferas, editors, *Proceedings of the 15th Annual ACM Symposium on Theory of Computing, 25-27 April, 1983, Boston, Massachusetts, USA*, pages 118–126. ACM, 1983.
- [TD02] Barbara M. Terhal and David P. DiVincenzo. Adaptive quantum computation, constant-depth circuits and arthur-merlin games. *Quantum Information and Computation*, pp 134-145, 2002.

- [Vya03] Mikhail N. Vyalyi. Qma=pp implies that PP contains PH. *Electron. Colloquium Comput. Complex.*, TR03-021, 2003.
- [Wie83] Stephen Wiesner. Conjugate coding. *SIGACT News*, 15:78–88, 1983.
- [Yir24] Justin Yirka. Even quantum advice is unlikely to solve PP. *CoRR*, abs/2403.09994, 2024.
- [ZVBL23] Alexander Zlokapa, Benjamin Villalonga, Sergio Boixo, and Daniel A. Lidar. Boundaries of quantum supremacy via random circuit sampling. *npj Quantum Information*, 9(1):36, Apr 2023.

A Instantiating Uniform Approximation Hardness (Definition 5.1)

In this section we will show how to instantiate uniform approximation hardness with well studied conjectures from the sampling-based quantum advantage literature. Specifically, we will import some conjectures from the literature on BosonSampling, Random Circuit Sampling, IQP and QQC1 sampling; and will discuss why they imply Definition 5.1.

A.1 Random Circuit Sampling

The exposition in this section is primarily taken from [BFNV19]. Define a circuit architecture $A := \{A_n\}_{n \in \mathbb{N}}$ as a family of graphs with $\text{poly}(n)$ vertices, where each vertex v has $\text{deg}_{\text{in}}(v) = \text{deg}_{\text{out}}(v) \in \{1, 2\}$. Intuitively, the vertices of the graph denote one or two qubit gates and the edges denote wires. A quantum circuit is instantiated by specifying the gate for each vertex. Define \mathcal{H}_A as the distribution over circuits formed by drawing a (one or two qubit) gate independently from the Haar measure for each vertex in A and assigning the gate to the vertex.

Definition A.1 (Anticoncentration). *For an architecture A , we say that RCS anticoncentrates on A if there exist constants $\kappa, \gamma > 0$ such that for all large enough n*

$$\Pr_{C \leftarrow \mathcal{H}_{A_n}} \left[\Pr_C[0^n] \geq \frac{1}{\kappa 2^n} \right] \geq \gamma$$

Definition A.2 (Hiding). *For an architecture A , we say that \mathcal{H}_A has the hiding property if for any $C \leftarrow \mathcal{H}_{A_n}$ and uniformly random $y \leftarrow \{0, 1\}^n$, C_y is distributed as \mathcal{H}_{A_n} where C_y is the circuit such that $\Pr_C[x] = \Pr_{C_y}[x \oplus y]$, i.e. the circuit C with X gates appended to every output wire where the value of the output bit in y is 1.*

Definition A.3 (Approximate Average-Case Hardness). *An architecture $A := \{A_n\}_{n \in \mathbb{N}}$ is said to be approximate average-case #P-hard to approximate if it has the following property. There exist functions $\epsilon(n) = 1/p(n)$ and $\delta(n) = 1/q(n)$ for some polynomials p and q such that for any oracle \mathcal{O} s.t. for all large enough n ¹⁶*

$$\Pr_{C \leftarrow \mathcal{H}_{A_n}} [|\mathcal{O}(C) - \Pr_C[0^n]| \leq \epsilon(n)/2^n] \geq 1 - \delta(n)$$

it holds that $\text{P}^{\#P} \subseteq \text{BPP}^{\mathcal{O}}$.

¹⁶In the literature a slightly different form is often used where \mathcal{O} takes $1^{1/\epsilon}$ and $1^{1/\delta}$ as input and must approximate to precision $\epsilon/2^n$ with error probability at most δ . The version we present is more convenient and cleaner for the purpose of building cryptography, but our results hold for both versions. See the proof of Theorem 5.1 for details.

RCS is based on the conjecture that there exists an architecture A that satisfies Definition A.1, Definition A.2, and Definition A.3 (see, e.g., [BFNV19]). We show that this implies the following conjecture, which directly leads to an instantiation of Definition 5.1.

Conjecture 1. *There exists an architecture A and polynomials p and q such that:*

1. **Anticoncentration.**

$$\Pr_{\substack{C \leftarrow \mathcal{H}_{A_n} \\ x \leftarrow \{0,1\}^n}} [\Pr_C[x] \geq 1/(p(n) \cdot 2^n)] \geq 1/\gamma(n)$$

2. **Hardness.** *For any oracle \mathcal{O} satisfying that for all large enough $n \in \mathbb{N}$,*

$$\Pr_{\substack{C \leftarrow D_n \\ x \leftarrow \{0,1\}^n}} \left[|\mathcal{O}(C, x) - \Pr_C[x]| \leq \frac{\Pr_C[x]}{p(n)} \right] \geq 1/\gamma(n) - \frac{1}{p(n)}$$

we have that $\mathsf{P}^{\#\mathsf{P}} \subseteq \mathsf{BPP}^{\mathcal{O}}$.

To see why the existence of an architecture satisfying Definition A.1, Definition A.2, and Definition A.3 implies Conjecture 1, it is first observed that anticoncentration holds directly from Definition A.1 and Definition A.2. Next, the hiding property implies that approximate average-case hardness holds even for an oracle that takes input $C \leftarrow \mathcal{H}_{A_n}$ and $x \leftarrow \{0, 1\}^n$ and estimates $\Pr_C(x)$. Finally, anticoncentration implies that estimating probabilities with small additive error on average implies the ability to estimate probabilities with small relative error on a large fraction of the set of anticoncentrated points.

A.2 Boson Sampling

This section imports conjectures that were made in [AA11] to obtain quantum advantage from Boson Sampling; and discusses why these conjectures imply Definition 5.1.

Conjecture 2. *[Permanent-of-Gaussians-Conjecture] There exist polynomials $p(\cdot), q(\cdot)$ such that for $\epsilon = 1/p(n), \delta = 1/q(n)$, if there exists an oracle \mathcal{O} that given as input a matrix $X \sim \mathcal{N}(0, 1)_{\mathbb{C}}^{n \times n}$ of i.i.d. Gaussians, can estimate $\text{Per}(X)$ to within error $\pm \epsilon(n) \cdot |\text{Per}(X)|$, with probability at least $1 - \delta(n)$ over X , then $\mathsf{P}^{\#\mathsf{P}} \subseteq \mathsf{BPP}^{\mathcal{O}}$.*

Conjecture 3. *[Permanent Anti-Concentration Conjecture] There exists a polynomial $p(\cdot)$ such that for all n and $\delta > 0$,*

$$\Pr_{X \sim \mathcal{N}(0,1)_{\mathbb{C}}^{n \times n}} \left[|\text{Per}(X)| < \frac{\sqrt{n!}}{p(n, 1/\delta)} \right] < \delta$$

Theorem A.1. *Conjectures 2 and 3 imply Definition 5.1.*

Proof. (Sketched, from [AA11, GBA⁺22]) It is shown in [AA11] how the probability that a randomly chosen linear optical network (i.e. circuit) outputs 0 is proportional to the square of the permanent of an appropriate submatrix of a Haar random matrix; where the submatrix itself is Gaussian. They also prove a hiding property, which argues that any given Gaussian matrix can be embedded into a Haar random matrix while keeping the location of the given submatrix hidden. The embedding procedure itself is known to be in $\mathsf{FBPP}^{\mathsf{NP}}$ and is conjectured in [AA11] to be in FBPP . Subsequent models of Boson Sampling such as (Bipartite) Gaussian Boson Sampling modify the experimental setup and work with Gaussian (as opposed to Haar random) matrices, which allow trivially embedding a Gaussian submatrix in FBPP .

Specifically, [GBA⁺22] efficiently (in BPP) reduce the task of computing permanents of Gaussian random matrices to the task of estimating the probabilities of (random) outputs of a bipartite Gaussian boson sampling setup. Under Conjectures 2 and 3, this proves that the resulting probabilities are #P hard to approximate, which implies the statement of Definition 5.1. \square

A.3 IQP Circuit Sampling

IQP refers to a class of randomly chosen commuting quantum circuits, which take as input the state $|0^n\rangle$, whose gates are diagonal in the Pauli-X basis, and whose n -qubit output is measured in the computational basis. Under *any one of* the two conjectures below, one coming from condensed matter physics and the other from computer science, the output distributions of IQP circuits have been proven classically hard to simulate (unless the polynomial hierarchy collapses) [SB09, BMS16].

Conjecture 4. Consider the partition function of the general Ising model,

$$Z(\omega) = \sum_{\mathbf{z} \in \{\pm 1\}^n} \omega \exp \left(\sum_{i < j} w_{ij} z_i z_j + \sum_{k=1}^n v_k z_k \right), \quad (18)$$

where the exponentiated sum is over the complete graph on n vertices, w_{ij} and v_k are real edge and vertex weights, and $\omega \in \mathbb{C}$. Let the edge and vertex weights be picked uniformly at random from the set $\{0, \dots, 7\}$.

Then it is #P-hard to approximate $|Z(e^{i\pi/8})|^2$ up to multiplicative error $1/4 + o(1)$ for a $1/24$ fraction of instances, over the random choice of weights.

Conjecture 5. Let $f : \{0, 1\}^n \rightarrow \{0, 1\}$ be a uniformly random degree-3 polynomial over \mathbb{F}_2 , and define $ngap(f) := (|\{x : f(x) = 0\}| - |\{x : f(x) = 1\}|)/2^n$. Then it is #P-hard to approximate $ngap(f)^2$ up to a multiplicative error of $1/4 + o(1)$ for a $1/24$ fraction of polynomials f .

Just like the case of random circuit sampling, IQP circuits satisfy the hiding property trivially due to the circuit architecture supporting the addition of random X gates. Furthermore, the structure of IQP circuits allows provable anti-concentration results for $ngap(f)$ and the partition function of the random Ising model. These properties combined help reduce the task of approximating the partition function of the general Ising model, or approximating $ngap(f)^2$ within the bounds in the conjecture, to the task of approximating output probabilities of randomly chosen IQP circuits. Thus *either one of* Conjectures 4 or 5 implies hardness according to Definition 5.1.

B Flatness of Unitary 2-Design Output Distributions

Theorem 7.2 (Flatness of 2-designs). Let \mathcal{C} be a unitary 2-design on n qubits. Fix any n qubit state $|\psi\rangle$. For any $C \in \text{Supp}(\mathcal{C})$, let $p_C(x) := |\langle x|C|\psi\rangle|^2$ be the probability that measuring $C|\psi\rangle$ in the computational basis results in x . Then the following holds for all $k > 6$ and sufficiently large n . Define

$$\mathbb{G} := \left\{ C \in \text{Supp}(\mathcal{C}) : \sum_{x: p_C(x) \geq \frac{n^{3k}}{2^n}} p_C(x) \leq 1/n^k \right\}$$

Then

$$\Pr_{C \leftarrow \mathcal{C}} [C \in \mathbb{G}] \geq 1 - 1/n^k$$

Proof. By the Chebyshev bound applied to random variable $p_C(x)$ for uniformly random x and $C \leftarrow \mathcal{C}$, for all $\alpha > 0$

$$\begin{aligned} \Pr_{\substack{x \leftarrow \{0,1\}^n \\ C \leftarrow \mathcal{C}}} \left[|p_C(x) - 1/2^n| \geq \frac{\alpha}{2^n} \right] &\leq \frac{2^{2n} (\mathbb{E}_{C,x}[p_C(x)^2] - \mathbb{E}_{C,x}[p_C(x)]^2)}{\alpha^2} \\ &\leq \frac{2^{2n} (\mathbb{E}_{C,x}[p_C(x)^2] - 1/2^{2n})}{\alpha^2} \end{aligned} \quad (19)$$

where the second step follows from the expectation over x of $p_C(x)$ is $1/2^n$ for all C . Since C is sampled from a unitary 2-design, we show a bound on the second moment. Fix any x . By the definition of $p_C(x)$,

$$\begin{aligned} \mathbb{E}_C[p_C(x)^2] &= \mathbb{E}_C \left[\left(\langle x|C^\dagger|\psi\rangle\langle\psi|C|x\rangle \right)^2 \right] \\ &= \mathbb{E}_C \left[\langle x|^{\otimes 2}C^{\otimes 2}|\psi\rangle\langle\psi|^{\otimes 2}C^{\dagger\otimes 2}|x\rangle^{\otimes 2} \right] \\ &= \langle x|^{\otimes 2}\mathbb{E}_C \left[C^{\otimes 2}|\psi\rangle\langle\psi|^{\otimes 2}C^{\dagger\otimes 2} \right] |x\rangle^{\otimes 2} \end{aligned}$$

By the definition of unitary 2-designs (Definition 4.1), $\mathbb{E}_C [C^{\otimes 2}|\psi\rangle\langle\psi|^{\otimes 2}C^{\dagger\otimes 2}]$ equals the second moment operator with respect to the Haar measure applied to $|\psi\rangle\langle\psi|^{\otimes 2}$, i.e. $\mathbb{E}_{U \leftarrow \mu_H} [U^{\otimes 2}|\psi\rangle\langle\psi|^{\otimes 2}U^{\dagger\otimes 2}]$ which is known to equal $\frac{\mathbb{I} + \mathbb{F}}{2^n(2^n + 1)}$ where \mathbb{I} is identity and \mathbb{F} is the flip operator (see Corollary 13 in [Mel24]). Therefore

$$\mathbb{E}_C \left[C^{\otimes 2}|\psi\rangle\langle\psi|^{\otimes 2}C^{\dagger\otimes 2} \right] = \frac{\mathbb{I} + \mathbb{F}}{2^n(2^n + 1)}$$

We can now compute $\mathbb{E}_C[p_C(x)^2]$ as follows

$$\begin{aligned} \mathbb{E}_C[p_C(x)^2] &= \langle x|^{\otimes 2} \frac{\mathbb{I} + \mathbb{F}}{2^n(2^n + 1)} |x\rangle^{\otimes 2} \\ &= \frac{2}{2^n(2^n + 1)} \end{aligned}$$

Since the above holds for any x , it also holds for uniformly sampled x , which implies

$$\mathbb{E}_{x,C}[p_C(x)^2] = \frac{2}{2^n(2^n + 1)}$$

Substituting in (19) gives

$$\Pr_{\substack{x \leftarrow \{0,1\}^n \\ C \leftarrow \mathcal{C}}} \left[|p_C(x) - 1/2^n| \geq \frac{\alpha}{2^n} \right] \leq \frac{1}{\alpha^2}$$

We therefore obtain a bound on the fraction of x with $p_C(x)$ above a threshold. However, we need to bound the total probability mass on such x . To do so we leverage the fact that the above bound holds for every α . We can therefore simultaneously bound the fraction of such x for every a sequence of α values. This turns out to be sufficient to bound the probability mass.

Define $f_\alpha(C) := \Pr_{x \leftarrow \{0,1\}^n} [p_C(x) \geq \frac{1+\alpha}{2^n}]$. Intuitively, this is the fraction of x with $p_C(x)$ above the required threshold. Therefore

$$\mathbb{E}_C[f_\alpha(C)] \leq 1/\alpha^2$$

By a Markov argument

$$\Pr_C[f_\alpha(C) \geq 1/\alpha^{1.5}] \leq 3/\alpha^{0.5}$$

For any $i \in \mathbb{N}$, setting α to n^{2i}

$$\Pr_C[f_{n^{2i}}(C) \geq 1/n^{3i}] \leq 3/n^i$$

Let $\mathbb{G}_i := \{C : f_{n^{2i}}(C) < 1/n^{3i}\}$. Then $\Pr_C[C \in \mathbb{G}_i] \geq 1 - 3/n^i$. Let $\mathbb{G}' := \bigcap_{i \geq k+3} \mathbb{G}_i$. Then

$$\begin{aligned} \Pr_C[C \in \mathbb{G}'] &\geq 1 - \left(3/n^{k+3} + 3/n^{k+4} + \dots\right) \\ &\geq 1 - \frac{3}{n^{k+3}}(1 + 1/n + 1/n^2 + \dots) \\ &\geq 1 - 6/n^{k+3} \end{aligned}$$

Fix any $C \in \mathbb{G}'$.

$$\begin{aligned} \sum_{x: p_C(x) \geq \frac{1+n^{2k+6}}{2^n}} p_C(x) &= \sum_{i \geq k+3} \sum_{x: \frac{1+n^{2i+2}}{2^n} \geq p_C(x) \geq \frac{1+n^{2i}}{2^n}} p_C(x) \\ &\leq \sum_{i \geq k+3} \sum_{x: \frac{1+n^{2i+2}}{2^n} \geq p_C(x) \geq \frac{1+n^{2i}}{2^n}} \frac{1+n^{2i+2}}{2^n} \\ &\leq \sum_{i \geq k+3} \frac{2^n}{n^{3i}} \cdot \frac{1+n^{2i+2}}{2^n} \\ &\leq \sum_{i \geq k+3} \frac{2^n}{n^{3i}} \cdot \frac{1+n^{2i+2}}{2^n} \\ &\leq \sum_{i \geq k+3} 2/n^{i-2} \\ &\leq 4/n^{k+1} \end{aligned}$$

where the third step uses the definitions of \mathbb{G}' , \mathbb{G}_i , and $f_{n^{2i}}$. For $C \in \mathbb{G}'$, all $k > 6$ and sufficiently large n

$$\sum_{x: p_C(x) \geq \frac{n^{3k}}{2^n}} p_C(x) \leq 1/n^k$$

and $\Pr_C[C \in \mathbb{G}'] \geq 1 - 1/n^k$, which concludes the proof. \square

C Proof of Theorem 7.3

Theorem 7.3. Let $(x, y), (x^*, y^*) \in \mathbb{R}^2$ such that $\exists \gamma > 0, \gamma' > 0$ s.t.

- $x^2 + y^2 \geq \gamma^2$
- $(x - x^*)^2 + (y - y^*)^2 \leq (\gamma')^2$

- $\gamma' < \gamma$

Then $|e^{-i \cdot \arctan2(y, x)} - e^{-i \cdot \arctan2(y^*, x^*)}| \leq 2\gamma'/\gamma$

Proof. We prove this theorem by a geometric argument. Let $X = (x, y)$ and $Y = (x^*, y^*)$ be points on the Cartesian plane, and let $O = (0, 0)$ be the origin. Let the (smaller) angle between rays OX and OY be ζ . Then $|\arctan2(y, x) - \arctan2(y^*, x^*)|$ is equal to ζ (upto a multiple of 2π offset). By expanding the expression we aim to bound

$$\begin{aligned} |e^{-i \cdot \arctan2(y, x)} - e^{-i \cdot \arctan2(y^*, x^*)}| &= \left| 2 \sin \left(\frac{\arctan2(y, x) - \arctan2(y^*, x^*)}{2} \right) \right| \\ &= \left| 2 \sin \left(\frac{\zeta}{2} \right) \right| \\ &\leq |\zeta| \end{aligned}$$

Therefore it suffices to show that ζ is upper bounded by $2\gamma'/\gamma$.

Let C be a circle of radius γ' centered at X . Since $|OX| = \sqrt{x^2 + y^2} \geq \gamma > \gamma'$, O is strictly external to the circle. Additionally, since $|XY| = \sqrt{(x - x^*)^2 + (y - y^*)^2} \leq \gamma'$, Y is on the circle or internal to it. Therefore, OY is either a tangent or secant line. The angle between OX and OY is maximized if OY is tangent to the circle, in which case $OY \perp XY$ and $\sin(\zeta) = |XY|/|OX| \leq \gamma'/\gamma$. Additionally note that ζ is acute, so $\zeta \leq 2 \sin(\zeta) \leq 2\gamma'/\gamma$, which concludes the proof. \square

D Quantum Advantage from One-Way Puzzles in Microcrypt

In this section we show that if one-way functions *do not* exist, the existence of one-way puzzles implies sampling based quantum advantage¹⁷. Let $(\text{Samp}, \text{Ver})$ be a one-way puzzle. Then we claim that the existence of an efficient classical algorithm that samples from a distribution less than $1/3$ statistically-far from the distribution of $\text{Samp}(1^n)$ on all large enough $n \in \mathbb{N}$ implies the existence of one-way functions.

Suppose there exists an efficient classical algorithm \mathcal{S} such that for all large enough n

$$\text{SD}(\text{Samp}(1^n), \mathcal{S}(1^n; r)) \leq 1/3$$

where r is a uniform string of n -bits.

Claim D.1. Define $f(x)$ for $x \in \{0, 1\}^n$ as follows:

- $(s', k') \leftarrow \mathcal{S}(1^n; x)$
- Output s'

Then f is a $1/n$ -distributional one-way function.

Proof. Suppose not, i.e. there exists a non-uniform PPT adversary \mathcal{A} such that for infinitely many $n \in \mathbb{N}$

$$\{s', x\} \approx_{1/n} \{s', \mathcal{A}(s')\} \tag{20}$$

where x is a uniform n -bit string and $(s', k') := \mathcal{S}(1^n; x)$.

Let \mathcal{S}' be the algorithm that on input x performs the following:

¹⁷The first version of this manuscript had the informal argument sketched out in the Results section, but we have added a formal proof to the current (updated) version.

- $s^*, k^* \leftarrow \mathcal{S}(1^n; x)$
- Return k^*

Consider the distribution $\{s', \mathcal{S}'(\mathcal{A}(s'))\}$. Since $\{s', \mathcal{A}(s')\}$ is at most $1/n$ far from $\{s', x\}$, $\{s', \mathcal{S}'(\mathcal{A}(s'))\}$ is also (at most) $1/n$ far from $\{s', \mathcal{S}'(x)\}$. Since (s', k') is obtained by running $\mathcal{S}(1^n; x)$ and $\mathcal{S}'(x)$ simply outputs k' , $\{s', \mathcal{S}'(x)\}$ is identical to $\{s', k'\}$. This gives

$$\{s', k'\} \approx_{1/n} \{s', \mathcal{S}'(\mathcal{A}(s'))\}$$

Since the distribution over (s', k') is $1/3$ close to the distribution of $\text{Samp}(1^n)$

$$\text{SD}(\{s, k\}, \{s', k'\}) \leq 1/3 \text{ and } \text{SD}(\{s, \mathcal{S}'(\mathcal{A}'(s))\}, \{s', \mathcal{S}'(\mathcal{A}(s'))\}) \leq 1/3$$

where $(s, k) \leftarrow \text{Samp}(1^n)$. Putting these together with the previous equation we get

$$\text{SD}(\{s, k\}, \{s, \mathcal{S}'(\mathcal{A}'(s))\}) \leq 2/3 + 1/n$$

Now, by the correctness of the one-way puzzle

$$\Pr_{s, k \leftarrow \text{Samp}(1^n)} [\text{Ver}(s, k) = 1] \geq 1 - \text{negl}(n)$$

which implies

$$\Pr_{s, k \leftarrow \text{Samp}(1^n)} [\text{Ver}(s, \mathcal{S}'(\mathcal{A}'(x))) = 1] \geq 1/3 - 1/n - \text{negl}(n)$$

which contradicts the security of the one-way puzzle. □

A Note on Quantum Advantage from One-way Puzzles. Since the existence of distributional one-way functions implies the existence of one-way functions [ILL89], if one-way puzzles exist but one-way functions do not, there exist distributions that can be efficiently sampled by quantum machines but cannot be efficiently sampled by classical machines. It was shown in [MY24] that the existence of one-way functions implies (interactive, inefficiently verifiable) proofs of quantumness. Combining the previous two facts yields a direct proof that the existence of one-way puzzles implies interactive quantum advantage.