

Quantum Black-Box Separations: Succinct Non-Interactive Arguments from Falsifiable Assumptions

Gorjan Alagic¹, Dana Dachman-Soled², Manasi Shingane³, and Patrick Struck⁴

¹ QuICS, University of Maryland, and NIST

`galagic@umd.edu`

² University of Maryland

`danadach@umd.edu`

³ QuICS, University of Maryland,

`mshingan@umd.edu`

⁴ University of Konstanz

`patrick.struck@uni-konstanz.de`

Abstract. In their seminal work, Gentry and Wichs [GW11] established an impossibility result for the task of constructing an adaptively-sound SNARG via black-box reduction from a falsifiable assumption. An exciting set of recent SNARG constructions demonstrated that, if one adopts a weaker but still quite meaningful notion of adaptive soundness, then impossibility no longer holds [WW24a, WW24b, WZ24, MPV24]. These fascinating new results raise an intriguing possibility: is there a way to remove this slight weakening of adaptive soundness, thereby completely circumventing the Gentry-Wichs impossibility?

A natural route to closing this gap would be to use a *quantum black-box reduction*, i.e., a reduction that can query the SNARG adversary on superpositions of inputs. This would take advantage of the fact that [GW11] only consider *classical reductions*. In this work, we show that this approach cannot succeed. Specifically, we extend the Gentry-Wichs impossibility result to quantum black-box reductions, and thereby establish an important limit on the power of such reductions.

1 Introduction

In this work, we explore the limits of *quantum, black-box reductions* in the context of succinct non-interactive arguments (SNARGs) for NP. SNARGs allow one to

* We thank Chris Majenz, Kai-Min Chung, and Xiaodi Wu for helpful conversations. Patrick Struck acknowledges funding by the Bavarian State Ministry of Science and the Arts in the framework of the bidt Graduate Center for Postdocs and the German Research Foundation (DFG) – SFB 1119 – 236615297 (while working at University of Regensburg) as well as the Hector Foundation II. Part of the work was done while Patrick Struck was visiting the University of Maryland. This work was also supported in part by the Air Force Office of Scientific Research (award number FA9550-20-1-0108) This work was supported by NSF award # CNS-2154705.

verify the validity of a statement x for some NP language L with a proof π that is of length much smaller than the witness w . Since their introduction, SNARGs have found myriad applications in both theory [BISW17] and practice [SCG⁺14].

A SNARG satisfies *soundness* if no efficient adversary can produce a fake proof (for a false statement) that is accepted by the verifier. If this holds against adversaries who are allowed to select the false statement after seeing the common reference string (CRS) of the SNARG, we say that the SNARG satisfies *adaptive soundness*. In their seminal work, Gentry and Wichs [GW11] showed that, under a natural assumption, it is impossible to prove adaptive soundness of a SNARG via a reduction from a *falsifiable assumption* [Nao03]—so long as the reduction treats the adversary as a black-box.⁵ As adaptive soundness is a highly desirable property, SNARGs used in practice are typically based on non-falsifiable assumptions, e.g., the Random Oracle Model [BR93].

Recently, a line of exciting breakthrough results almost completely circumvented the Gentry-Wichs barrier [WW24a, WW24b, WZ24, MPV24]. These results construct SNARGs via black-box reduction from falsifiable assumptions, achieving soundness that holds against adaptive adversaries—provided the adversary declares an upper bound on the false statement size prior to seeing the CRS. We refer to this form of soundness as *bounded-length adaptive soundness*. The constructions make use of indistinguishability obfuscation (iO) and the reductions make clever use of complexity leveraging.

These fascinating new results leave open the question of whether it is possible to fully circumvent the impossibility barrier and achieve adaptive soundness in the Gentry-Wichs sense, i.e., where even the length of the false statement can be chosen adaptively⁶. Specifically, one may ask whether there are proof techniques that are not ruled out by Gentry-Wichs that could be exploited to construct an adaptively sound SNARG from a falsifiable assumptions. While non-black-box use of the adversary is in principle a possibility, non-black-box techniques are few and far between [BLV03, Bar01, BP12, BKP19] and we are unaware of examples where such techniques have been used to prove *soundness* against poly-time adversaries. It thus seems that a black-box technique is a more plausible option. Here, quantum computation presents a natural possibility: using a *quantum* black-box reduction. In this case, the reduction is a quantum polynomial-time (QPT) algorithm that can make superposition queries to the adversary. For QPT algorithms, this query model is natural: given the code of the adversary, one can construct a corresponding reversible circuit and then run it on arbitrary inputs in superposition using a quantum computer.

Using a quantum reduction instead of a classical one can be helpful. Indeed, there are several cases in the literature in which the relationship between two primitives can be established with respect to quantum reductions, but is unknown with respect to classical reductions. For example, a quantum reduc-

⁵ Note that [GW11] *does* rule out reductions that use *the underlying assumption* in a non-black-box way (e.g., under an iO).

⁶ Consider, e.g., applications to verifiable delegated computation, where the CRS is fixed during setup, before the size of the delegated computation is known.

tion is known from the short integer solution (SIS) problem to the learning with error (LWE) problem [SSTX09], whereas only the other direction is known classically [CLZ22]. Additionally, a quantum reduction was initially given by Regev from the short independent vectors problem (SIVP) to the LWE problem [Reg09]. A classical reduction was subsequently given to LWE, but from a different lattice problem—Gap-SVP [BLP+13].

This motivates the following open question:

Do quantum black-box reductions allow one to circumvent the Gentry-Wichs impossibility result for SNARGs?

1.1 Our Results

Our main result, stated informally in the theorem below, is that the Gentry-Wichs impossibility result for SNARGs holds even for quantum reductions.

Theorem 1 (informal). *Let L be an NP language with a subexponentially-quantum-hard subset membership problem. Let Π be a SNARG for L satisfying completeness and succinctness. Then, for any falsifiable assumption (C, c) , there is no quantum black-box reduction showing the soundness of Π based on the assumption (C, c) .*

The formal version is given in Section 4 as Theorem 19. Similar to [GW11], our result is in the setting of non-uniform adversaries, and rules out non-security-parameter-preserving reductions. It also rules out reductions from δ -exponential hard assumptions (see Theorem 20). However, it does not rule out reductions that use quantum non-uniform advice.

As part of our proof of Theorem 19, we establish the following fact about quantum-query indistinguishability of distribution families. It shows that, if a set of distribution pairs is each indistinguishable from one sample, then the entire set is indistinguishable against quantum queries (in the appropriate sense).

Theorem 2 (informal). *Let $\mathbf{A} = \{A_i\}$ and $\mathbf{B} = \{B_i\}$ be two distribution families of size k . If A_i and B_i are (s^*, ε^*) -sample-indistinguishable for all i , then \mathbf{A} and \mathbf{B} are (s, ε) -oracle indistinguishable against adversaries making at most q queries, for $s^* = s + \text{poly}(kq)/\varepsilon^*$ and $\varepsilon = \text{poly}(\varepsilon^*/kq)$.*

The formal version is given in Section 3.2 as Theorem 17. It is a generalization of a result of Zhandry [Zha12b] that establishes an equivalence between sample indistinguishability and quantum-query indistinguishability for efficiently-sampleable distributions.

1.2 Technical Overview

The Gentry-Wichs Meta-Reduction. At a high level, the Gentry-Wichs impossibility proof proceeds as follows. First, one constructs an inefficient SNARG adversary $\bar{\mathcal{P}}$ that produces (convincing) bogus proofs $\bar{\pi}(\bar{x})$ for statements $\bar{x} \notin L$.

By soundness, it follows that the reduction Σ with query access to $\overline{\mathcal{P}}$ must win the falsifiable assumption experiment. Second, one constructs an efficient simulator \mathcal{S} for $\overline{\mathcal{P}}$ —where the simulation can fool any polynomial-time distinguisher that only gets query access to either $\overline{\mathcal{P}}$ or \mathcal{S} . As Σ is poly-time, it follows that $\Sigma^{\mathcal{S}}$ is an efficient algorithm that breaks the falsifiable assumption, yielding a contradiction.

To construct the (inefficient yet simulatable) SNARG adversary $\overline{\mathcal{P}}$, Gentry and Wichs make key use of a certain “leakage lemma”. This lemma states that, if \mathcal{L} and $\tilde{\mathcal{L}}$ are indistinguishable distributions and $\pi(x)$ is some short auxiliary information for samples $x \leftarrow \mathcal{L}$, then there exists a choice of short auxiliary information $\tilde{\pi}(\tilde{x})$ for samples $\tilde{x} \leftarrow \tilde{\mathcal{L}}$ so that $(x, \pi(x))$ is indistinguishable from $(\tilde{x}, \tilde{\pi}(\tilde{x}))$. In the context of Gentry-Wichs, \mathcal{L} is a distribution on yes-instances of L , $\tilde{\mathcal{L}}$ is a distribution on no-instances, and $\pi(x)$ is the SNARG proof. The lemma then yields, for every CRS crs , an augmented no-distribution $(\tilde{x}, \tilde{\pi}(\tilde{x})) \leftarrow \tilde{\mathcal{L}}_{\text{crs}}^*$ corresponding to the SNARG yes-distribution $(x, \pi(x)) \leftarrow \mathcal{L}_{\text{crs}}^*$.

Then $\overline{\mathcal{P}}(\text{crs})$ is defined to output samples from $\tilde{\mathcal{L}}_{\text{crs}}^*$ while $\mathcal{S}(\text{crs})$ outputs samples from $\mathcal{L}_{\text{crs}}^*$. The leakage lemma then immediately implies that the verifier will accept the outputs of $\overline{\mathcal{P}}$, and that $\Sigma^{\overline{\mathcal{P}}}$ will thus win the falsifiable assumption experiment. By a standard hybrid argument, since $\overline{\mathcal{P}}(\text{crs})$ and $\mathcal{S}(\text{crs})$ are sample-indistinguishable for every crs , $\overline{\mathcal{P}}$ and \mathcal{S} are indistinguishable as oracles. It then follows that $\Sigma^{\mathcal{S}}$ will also win the falsifiable assumption experiment.

The above discussion assumed that Σ only queries its oracle on one security parameter, i.e., the one relevant for the falsifiable assumption challenge. If instead it chose to query with very small security parameters, it could easily distinguish $\overline{\mathcal{P}}$ and \mathcal{S} . To address this, Gentry and Wichs show how to create a table of responses from $\overline{\mathcal{P}}$ for small values of the security parameter and encode this table into the non-uniform advice of \mathcal{S} .

The Challenge of Quantum Reductions. A natural approach for the setting of quantum reductions is to simply follow the same overall meta-reduction strategy: construct an inefficient SNARG adversary $\overline{\mathcal{P}}$ and show that it is query-indistinguishable from an efficient simulator \mathcal{S} . However, this approach faces two significant obstacles.

The first obstacle is in the construction of $\overline{\mathcal{P}}$. In the original proof, this was done via a “leakage lemma” that shows the existence of a short, fake proof for any given no-instance. In the case of quantum reductions, this fake proof would now need to fool quantum circuits. While a “quantum leakage lemma” (QLL) of this kind is known ([CCL⁺17, Lemma 6.3]), it only yields indistinguishability of a single classical sample, and it is not known to hold against circuits with *quantum advice*. The latter shortcoming is particularly problematic for the Gentry-Wichs hybrid-by-queries approach. In their approach, a query distinguisher \mathcal{A} that gains advantage (in distinguishing $\overline{\mathcal{P}}$ from \mathcal{S}) from some particular query is used to construct a sample distinguisher \mathcal{B} , yielding a contradiction. As $\overline{\mathcal{P}}$ is inefficient, \mathcal{B} needs to have a memory snapshot of \mathcal{A} prior to the relevant query—but in our setting this snapshot is now a quantum state. Unfortunately, the known techniques for proving the QLL do not appear to extend to the case of quantum

advice [Chu23]. We remark that this obstacle indicates that even the case of *quantum reductions making only classical queries* appears to be nontrivial.

The second obstacle is in establishing the quantum-query-indistinguishability of the adversary $\overline{\mathcal{P}}$ and the simulator \mathcal{S} . Note that a single classical query by Σ amounts to submitting a CRS as input and receiving an instance-proof pair as a response. As discussed above, Gentry and Wichs view such a query as one sample from a distribution indexed by the CRS. In the case of quantum reductions, even a single query by Σ can be a superposition over samples from a superpolynomial number of distributions, one for each possible value of the CRS. As a result, even *one-query indistinguishability* does not obviously follow from the indistinguishability of one classical sample (as provided by the QLL). Upgrading one-query-indistinguishability to multi-query-indistinguishability also seems nontrivial. One potentially promising tool here is a result of Zhandry [Zha12b], showing that samples and quantum queries grant the same power for distinguishing distributions. However, in our case the reduction can query in superposition across a superpolynomial number of distributions, corresponding to different choices of input CRS. Thus it would seem that Zhandry’s result would only allow to out quantum reductions that are forced to query the CRS classically.

Our Approach. In the end, our approach is in fact to follow the overall meta-reduction strategy of Gentry-Wichs. However, we make the following careful adaptations in order to avoid the obstacles discussed above.

1. For constructing $\overline{\mathcal{P}}$, we use the QLL. This yields, for each CRS crs , a distribution $\overline{\mathcal{L}}_{\text{crs}}^*$ consisting of no-instances and fake-proofs. Unfortunately, the guarantee is only that $\overline{\mathcal{L}}_{\text{crs}}^*$ is indistinguishable from the distribution $\mathcal{L}_{\text{crs}}^*$ (consisting of yes-instances and SNARG-proofs) against non-uniform quantum circuits that receive *a single classical sample* and *no quantum advice*. The efficient simulator \mathcal{S} is constructed as in Gentry-Wichs: on input crs , it outputs a sample from $\mathcal{L}_{\text{crs}}^*$, i.e., yes-instances and SNARG-proofs.
2. Recall that, as the QLL does not yield indistinguishability against quantum advice, we cannot hope to show indistinguishability of $\overline{\mathcal{P}}$ and \mathcal{S} via a hybrid-by-queries approach. Instead, we use a combination of complexity leveraging and Zhandry’s small-range-distribution techniques [Zha12b] in order to hybrid over the CRS, as follows.
 - (a) We first show that, for any fixed crs , the distributions $\overline{\mathcal{P}}(\text{crs})$ and $\mathcal{S}(\text{crs})$ are quantum-query-indistinguishable. To do this, we generalize a result of Zhandry on the equivalence of samples and quantum queries for distinguishing distributions [Zha12b] to the case of inefficient distributions. Our proof places a small number of samples in the (classical) non-uniform advice, and then applies Zhandry’s small-range-distribution technique.
 - (b) We then show that quantum-query indistinguishability of $\overline{\mathcal{P}}$ and \mathcal{S} extends to the case where the query algorithm can also query crs in superposition. This amounts to proving Theorem 2. As this is done via a sequence of hybrids over the CRS, this introduces a loss factor of $2^{|\text{crs}|}$; however, by careful complexity leveraging, we can absorb this loss into the quantum hardness of the subset membership problem for L .

3. Finally, we show how the small-range-distribution technique can also be used to build a table for \mathcal{S} that allows it to simulate quantum queries on small security parameters appropriately. This allows us to extend our impossibility result to also rule out non-security-parameter-preserving reductions.

1.3 Related Works

Meta-Reductions and Quantum Reductions. The meta-reduction technique was first introduced by Boneh and Venkatesan [BV98] to rule out restricted forms of reductions. Since then it has been used to rule out black-box security reductions for cryptographic primitives under various assumptions. As discussed above, it was used by Gentry and Wichs to rule out black-box reductions for adaptively-sound SNARGs [GW11] (see also [CLMP13, CGKS23]). Other works ruled out arguments satisfying certain special soundness guarantees [Pas11], certain statistical NIZKs [Pas13], and many more [PV05, GBL08, HH09, BDG⁺13, Wic13, Bro16, BJLS16, KLX20]. The meta-reduction technique was also used in the quantum setting by [DLS22], who generalizes previous impossibility results for the Fiat-Shamir heuristic to the case of shared quantum resources.

Impossibility of Black-Box Constructions. Meta-reduction results (like ours) do not require that the reduction only uses the *assumption* in a black-box way. A related line of works rule out reductions that make black-box use of both the adversary and the assumption. Impagliazzo and Rudich [IR89] introduced a framework for ruling out black-box constructions by showing an oracle separation between one-way permutations (OWP) and key agreement. Black-box separations have also been considered in the quantum setting. These include separations between: OWP and OWF [CX22], collision resistance and public-key quantum money [AHY23], quantum digital signatures and pseudorandom states [CM24], OWFs and key agreement over classical channels (conditionally) [ACC⁺22].

Subexponential SNARGs. As described earlier, several recent results have nearly circumvented the barrier of [GW11]. Waters and Wu [WW24a] first constructed adaptively secure SNARGs from sub-exponentially secure one-way functions, iO, and algebraic assumptions. Subsequently, Waters and Zhandry [WZ24] constructed adaptively secure SNARGs from sub-exponentially secure one-way functions, iO, and LWE. Then Waters and Wu [WW24b] removed the algebraic assumption and constructed adaptively secure SNARGs from sub-exponentially secure one-way functions and iO. Concurrently, Mathialagan et al. [MPV24] showed that the SNARG constructed by [SW14] is adaptively secure in the designated verifier model. We remark that all of these works require the adversary to declare a maximum length of their statement prior to SNARG setup.

2 Preliminaries

2.1 General Notation

For an integer k , we let $[k] = \{1, 2, \dots, k\}$. We say that a function $\varepsilon : \mathbb{N} \rightarrow \mathbb{N}$ is negligible if for all polynomials p , $\varepsilon(n) < 1/p(n)$ for all sufficiently large $n \in \mathbb{N}$. We will at times consider sets (denoted by, e.g., \mathcal{X} or \mathcal{Y}), and various distributions (denoted by, e.g., \mathcal{D}) on these sets. It is implicit that these sets and distributions are all parameterized by some positive integer parameter n , typically the security parameter. Given two sets \mathcal{X} and \mathcal{Y} , let $\mathcal{Y}^{\mathcal{X}}$ denote the set of functions $f : \mathcal{X} \rightarrow \mathcal{Y}$.

For a distribution \mathcal{D} on a set \mathcal{X} and an element $x \in \mathcal{X}$, we let $\Pr[\mathcal{D}(x)] := \Pr_{x' \leftarrow \mathcal{D}}[x' = x]$. We denote the total variation distance between two distributions \mathcal{D}_1 and \mathcal{D}_2 over a finite set \mathcal{X} by

$$\Delta(\mathcal{D}_1, \mathcal{D}_2) := \frac{1}{2} \sum_{x \in \mathcal{X}} |\Pr[\mathcal{D}_1(x)] - \Pr[\mathcal{D}_2(x)]|. \quad (1)$$

We define ε -closeness of two distributions \mathcal{D}_1 and \mathcal{D}_2 in terms of their total variation distance $\Delta(\mathcal{D}_1, \mathcal{D}_2)$. That is, we say that two distributions are ε -close if $\Delta(\mathcal{D}_1, \mathcal{D}_2) \leq \varepsilon$. Note that ε here is typically a function of the security parameter.

If \mathcal{D} is a distribution on \mathcal{Y} and \mathcal{X} is a set, then $\mathcal{D}^{\mathcal{X}}$ denotes the induced distribution on the set of functions $\mathcal{Y}^{\mathcal{X}}$. To draw a sample f from $\mathcal{D}^{\mathcal{X}}$, sample for each input $x \in \mathcal{X}$ an output $f(x) \in \mathcal{Y}$ according to \mathcal{D} . As discussed above, the sets \mathcal{X} and \mathcal{Y} and the distribution \mathcal{D} are all indexed by n . In this work, we will only consider inputs sets \mathcal{X} and the outputs sets \mathcal{Y} over bit strings of polynomial length. That is, all inputs sets \mathcal{X} and output sets \mathcal{Y} will be of size $2^{\text{poly}(n)}$ for a fixed polynomial in n .

Throughout the paper, we will consider tuples (i.e., finite indexed sequences) of distributions. A tuple of distributions on a set \mathcal{Y} is defined by an indexed set $\mathbf{D} = \{\mathcal{D}_i\}_{i \in [k]}$ where each \mathcal{D}_i is a distribution on \mathcal{Y} . Given a set \mathcal{X} and a distribution tuple \mathbf{D} on a set \mathcal{Y} , we define a distribution $\mathbf{D}^{\mathcal{X}}$ on the set $\mathcal{Y}^{[k] \times \mathcal{X}}$, as follows. For a function f sampled from $\mathbf{D}^{\mathcal{X}}$, each input $(i, x) \in [k] \times \mathcal{X}$ is mapped to a sample $f(i, x) \in \mathcal{Y}$ drawn from \mathcal{D}_i .

2.2 Quantum Circuit Families and Oracle Algorithms

Throughout this paper we consider the quantum circuit model of computation. A *quantum algorithm* is an infinite family $\mathcal{C} = \{\mathcal{C}_m\}_{m \in \mathbb{N}}$ of quantum circuits over some universal gate set, indexed by the input size m . In general, the gate set can include unitary gates, measurement gates, and “gates” for allocating ancilla qubits and discarding qubits. In this work, we will consider algorithms that can be augmented in two different ways: with *advice* (which is instance-independent), and with *auxiliary information* (which can be instance-dependent).

Non-Uniform Algorithms. Given a function $s: \mathbb{N} \rightarrow \mathbb{N}$ such that the size (i.e., number of gates) of \mathbf{C}_m satisfies $|\mathbf{C}_m| \leq s(m)$ for all m , we say that \mathbf{C} is a *size- s quantum circuit family*. We remark that this notion takes into account quantum algorithms that receive classical (non-uniform) advice. Indeed, such advice can simply be written into the circuit itself, e.g., as X gates to be applied to ancillas in the $|0\rangle$ state. Of course, the size of the advice is then counted towards the total circuit-size budget of $s(n)$. On the other hand, *quantum* non-uniform advice is *not* captured by this definition, since a quantum advice state of s (or fewer) qubits will in general not have a description of s (or fewer) bits. In this work, we will primarily deal with non-uniform algorithms that only use classical advice.

Polynomial-Time Algorithms. If there exists a polynomial p such that \mathbf{C} is a size- p quantum circuit family, then we say that \mathbf{C} is an *efficient* quantum algorithm (or a quantum polynomial-time algorithm, or a QPT). We further define a *uniform* QPT to be a QPT such that there exists a classical polynomial-time Turing machine that, on input 1^n , outputs a classical description of the gates of C_n . In this work, we will primarily be concerned with non-uniform QPTs, i.e., QPTs that are not necessarily uniform.

Oracle Algorithms. A *quantum oracle algorithm* (or quantum oracle circuit family) is a quantum algorithm \mathbf{C} whose circuits can make use of an additional “oracle” gate that implements some classical function $f: \mathcal{X} \rightarrow \mathcal{Y}$ (on bits) via the unitary operator

$$U_f : |x\rangle_{\mathcal{X}} |y\rangle_{\mathcal{Y}} \mapsto |x\rangle_{\mathcal{X}} |y \oplus f(x)\rangle_{\mathcal{Y}} . \quad (2)$$

Here we have slightly abused notation by letting \mathcal{X} and \mathcal{Y} denote both the input and output set (respectively) of f as well as the quantum registers containing the input and output (respectively). Prior to invoking U_f , a circuit of \mathbf{C} can prepare the input qubit registers \mathcal{X} and \mathcal{Y} in any state (by constructing it using an appropriate gate sequence), including states that are entangled with other registers. After the execution of U_f , the qubit registers \mathcal{X} and \mathcal{Y} are both returned to the circuit. The function f above is not necessarily fixed, so that we can consider running \mathbf{C} with oracles for various functions. A quantum algorithm \mathbf{C} that makes use of a quantum oracle U_f for a classical function f will be denoted by $\mathbf{C}^{(f)}$.

We remark that access to U_f implies access to its controlled version CU_f at the cost of one additional classical global query. The details of this can be found in Appendix A. Relying on this fact, we will assume that “quantum query access” means access to CU_f , and we will refer to access to U_f and CU_f interchangeably.

We highlight a particular case of interest where f describes a probabilistic classical function. In that case, we can write $f: \mathcal{I} \times \mathcal{X} \rightarrow \mathcal{Y}$ where \mathcal{I} is an input set and \mathcal{X} is a coins register, and the unitary U_f from (2) has the form of Fig. 1.

2.3 Quantum Indistinguishability of Distributions

We will consider several forms of quantum-computational indistinguishability of distributions. Without loss of generality, we assume that a quantum algorithm




Fig. 1. General quantum query interface $|f\rangle$ for $f: \mathcal{I} \times \mathcal{X} \rightarrow \mathcal{Y}$.

whose task is to distinguish between two settings (e.g., the input is a sample from one of two distributions) will always end by applying a two-outcome measurement and outputting the resulting measurement label, e.g., 0 or 1. We call such an algorithm a *quantum distinguisher*.

Sample-Indistinguishability. In the simplest scenario, the task of the quantum distinguisher is to distinguish between two distributions based on a single sample. We now define this notion of sample-indistinguishability, and then record an important fact about the power of distinguishers equipped with (short) auxiliary information.

Definition 3 ((s, ε)-Sample-Indistinguishability). *Let $s: \mathbb{N} \rightarrow \mathbb{N}$ and $\varepsilon: \mathbb{N} \rightarrow [0, 1]$. We say that two distributions \mathcal{D}_0 and \mathcal{D}_1 are (s, ε)-sample-indistinguishable if for every size- s quantum circuit family \mathbf{A} ,*

$$\left| \Pr_{x \leftarrow \mathcal{D}_0} [\mathbf{A}(1^n, x) = 1] - \Pr_{x \leftarrow \mathcal{D}_1} [\mathbf{A}(1^n, x) = 1] \right| \leq \varepsilon(n). \quad (3)$$

More generally, we say that two mixed quantum states ρ_0 and ρ_1 are (s, ε)-sample-indistinguishable if for every size- s quantum circuit family \mathbf{A} ,

$$|\Pr[\mathbf{A}(1^n, \rho_0) = 1] - \Pr[\mathbf{A}(1^n, \rho_1) = 1]| \leq \varepsilon(n). \quad (4)$$

Recalling Section 2.2, we emphasize that the above definition takes into account distinguishers that receive *classical* non-uniform advice. Such advice is hard-coded into the circuit description of \mathbf{A} and thus counts towards the total circuit-size budget s .

In Definition 3, we also considered the case where the task is to distinguish between two quantum states ρ_0 and ρ_1 . In our applications, we will only care about the special case where the quantum state ρ_b consists of two parts: a classical sample from some distribution, and some (quantum) auxiliary information. In that sense, we will remain in the setting of distinguishing distributions. The following lemma shows how sample indistinguishability of distributions can be transformed into sample indistinguishability even in the presence of (sufficiently short) quantum auxiliary information. This notion of “auxiliary information” (which is instance-dependent, and can be quantum) should not be confused with “advice” (which is instance-independent, and is classical in our setting).

Lemma 4 ([CCL⁺17, Lemma 6.3]). *For any $n, l, s' \in \mathbb{N}$ and $\varepsilon > 0$, the following holds for $s = \text{poly}(s', n, 2^l, 1/\varepsilon)$ and $\varepsilon' = 2\varepsilon$. Let \mathcal{D}_0 be a distribution*

on $\{0,1\}^n$, and let

$$\rho_{XB} = \sum_{x \in \{0,1\}^n} \Pr[\mathcal{D}_0(x)] |x\rangle\langle x|_X \otimes \rho_B^x \quad (5)$$

be a state with n classical bits (in register X) and ℓ qubits (in register B). For every distribution \mathcal{D}_1 that is (s, ε) -sample-indistinguishable from \mathcal{D}_0 , there exists a (possibly unbounded) quantum circuit family \mathcal{C} such that

$$\sigma_{XB} = \sum_{x \in \{0,1\}^n} \Pr[\mathcal{D}_1(x)] |x\rangle\langle x|_X \otimes \mathcal{C}(x)_B \quad (6)$$

and ρ_{XB} are (s', ε') -sample-indistinguishable.

Next, we observe that Lemma 4 also holds in the setting where (i.) the auxiliary information ρ_B^x is a classical string sampled from some distribution, and (ii.) \mathcal{C} is required to produce a classical output. Given a distribution \mathcal{D} on a set \mathcal{X} and an arbitrary probabilistic algorithm aux from \mathcal{X} to \mathcal{Y} , we define the augmented distribution $\mathcal{D}^{+\text{aux}}$ on the set $\mathcal{X} \times \mathcal{Y}$ as follows. To draw a sample from $\mathcal{D}^{+\text{aux}}$, first draw a sample x from \mathcal{D} , then compute $z \leftarrow \text{aux}(x)$, and finally output the result (x, z) .

Corollary 5 (Quantum Leakage Lemma (QLL)). *For any $n, l, s' \in \mathbb{N}$ and $\varepsilon > 0$, the following holds for $s = \text{poly}(s', n, 2^l, 1/\varepsilon)$ and $\varepsilon' = 2\varepsilon$. Let \mathcal{D}_0 be a distribution over $\{0,1\}^n$, and let aux_0 be an arbitrary algorithm with n input bits and ℓ output bits. For every distribution \mathcal{D}_1 that is (s, ε) -sample-indistinguishable from \mathcal{D}_0 , there exists an algorithm aux_1 such that the distributions $\mathcal{D}_0^{+\text{aux}_0}$ and $\mathcal{D}_1^{+\text{aux}_1}$ are (s', ε') -sample-indistinguishable.*

Proof. The distribution $\mathcal{D}_0^{+\text{aux}_0}$ is a special case of the state ρ_{XB} in Lemma 4, where ρ_B^x is the distribution produced by aux_0 on input x . Applying Lemma 4 to this particular ρ_{XB} , we get that there exists a quantum circuit \mathcal{C} such that $\mathcal{D}_0^{+\text{aux}_0}$ is (s', ε') -indistinguishable from the (quantum) state generated by first drawing a sample x from \mathcal{D}_1 and then outputting x along with $\mathcal{C}(x)$. Since (s', ε') -indistinguishability holds against all adversaries (of size at most s'), it in particular holds against all adversaries (of size at most s') who begin their computation by completely measuring the register B . This implies that, if we set aux_1 to be the algorithm that first applies \mathcal{C} and then measures the output, then $\mathcal{D}_0^{+\text{aux}_0}$ and $\mathcal{D}_1^{+\text{aux}_1}$ are (s', ε') -sample-indistinguishable distributions. \square

Oracle-Indistinguishability. We now consider the case of distinguishing between two distributions using an oracle that contains a table of samples. Recall from Section 2.1 that, given a set \mathcal{X} and a distribution \mathcal{D} on \mathcal{Y} , $\mathcal{D}^{\mathcal{X}}$ denotes the induced distribution on functions $f : \mathcal{X} \rightarrow \mathcal{Y}$.

Definition 6 ((s, q, ε) -Oracle-Indistinguishability). *Let $s, q : \mathbb{N} \rightarrow \mathbb{N}$ and $\varepsilon : \mathbb{N} \rightarrow [0, 1]$, and let \mathcal{D}_0 and \mathcal{D}_1 be two distributions over a set \mathcal{Y} . We say*

that \mathcal{D}_0 and \mathcal{D}_1 are (s, q, ε) -oracle-indistinguishable if, for every set \mathcal{X} and every size- s quantum (oracle) circuit family \mathbf{A} that makes at most q quantum queries,

$$\left| \Pr_{\mathcal{O} \leftarrow \mathcal{D}_0^{\mathcal{X}}}[\mathbf{A}^{|\mathcal{O}\rangle}(1^n) = 1] - \Pr_{\mathcal{O} \leftarrow \mathcal{D}_1^{\mathcal{X}}}[\mathbf{A}^{|\mathcal{O}\rangle}(1^n) = 1] \right| \leq \varepsilon(n) \quad (7)$$

for all sufficiently large n .

We emphasize that the above definition only considers distinguishers with *classical* non-uniform advice. Similarly to the case of sample indistinguishability, such advice is hard-coded into the circuit description of \mathbf{A} and thus counts towards the total circuit-size budget s . No quantum advice is allowed.

Note that Definition 6 allows for independently selecting the circuit-size upper bound s and the query upper bound q . While one might normally only control the circuit size s (with s also providing an upper bound for the total number of queries), in our applications it will be important to choose significantly different upper bounds for s and q . Specifically, we will need to consider cases where s is superpolynomial (or even unbounded) while q is polynomial. We say that two distributions are (q, ε) -oracle indistinguishable if they are indistinguishable (as in Definition 6 above) against all unbounded circuits that make at most q queries and have distinguishing advantage at least ε .

Indistinguishability from Small-Range Distributions. Let \mathcal{D} be a distribution on \mathcal{Y} , and let \mathcal{X} be a set. Recall that a sample f from the distribution $\mathcal{D}^{\mathcal{X}}$ is a function $f \in \mathcal{Y}^{\mathcal{X}}$ such that, for every $x \in \mathcal{X}$, $f(x)$ is sampled according to \mathcal{D} . Given a positive integer $r \leq |\mathcal{X}|$, a corresponding “small-range distribution” $\text{SR}_r^{\mathcal{D}}(\mathcal{X})$ on $\mathcal{Y}^{\mathcal{X}}$ is defined as follows. To sample a function g according to $\text{SR}_r^{\mathcal{D}}(\mathcal{X})$, do:

1. Sample a function $H \leftarrow \mathcal{D}^{[r]}$.
2. Sample a uniformly random function $K: \mathcal{X} \rightarrow [r]$.
3. For every $x \in \mathcal{X}$, define $g(x) = H(K(x))$.

The following result shows that distributions are oracle-indistinguishable from their associated small-range distributions, even against unbounded quantum distinguishers.

Lemma 7 ([Zha12a, Corollary 7.5]). *The output distributions of a quantum algorithm making q quantum queries to an oracle drawn from either $\mathcal{D}^{\mathcal{X}}$ or $\text{SR}_r^{\mathcal{D}}(\mathcal{X})$ are $\ell(q)/r$ -close where $\ell(q) = \pi^2(2q)^3/3$.*

Note that the above lemma holds against unbounded distinguishers.

2.4 Quantum Black-Box Reductions for SNARGs

Our work is concerned with provable security for SNARGs, defined below. Note that we require *adaptive* soundness, i.e., soundness even against adversaries that are allowed to select their instance after seeing the common reference string.

Definition 8 (Succinct Non-Interactive Arguments for NP). A SNARG system Π consists of three polynomial-time classical algorithms $\Pi = (\mathbf{G}, \mathbf{P}, \mathbf{V})$:

- $(\text{crs}, \text{priv}) \leftarrow \mathbf{G}(1^n)$: The generation algorithm takes as input the security parameter n and outputs a common reference string crs as well as private verification information priv .
- $\pi \leftarrow \mathbf{P}(\text{crs}, x, w)$: The prove algorithm takes in the crs , a statement x , and a witness w , and outputs a proof π .
- $b \leftarrow \mathbf{V}(\text{priv}, x, \pi)$: the verification algorithm takes as input priv , a statement x , and a proof π , and outputs a bit b (1 denotes acceptance, 0 denotes rejection).

We say that Π is a succinct non-interactive argument (SNARG) for a language L in \mathcal{NP} with corresponding relation \mathcal{R} if it satisfies the following three properties:

Completeness: For all $(x, w) \in \mathcal{R}$,

$$\Pr \left[\mathbf{V}(\text{priv}, x, \pi) = 0 \mid \begin{array}{l} (\text{crs}, \text{priv}) \leftarrow \mathbf{G}(1^n) \\ \pi \leftarrow \mathbf{P}(\text{crs}, x, w) \end{array} \right] = \text{negl}(n). \quad (8)$$

Soundness: For all PPTs $\bar{\mathbf{P}}$,

$$\Pr \left[\mathbf{V}(\text{priv}, x, \pi) = 1 \wedge x \notin L \mid \begin{array}{l} (\text{crs}, \text{priv}) \leftarrow \mathbf{G}(1^n) \\ (x, \pi) \leftarrow \bar{\mathbf{P}}(\text{crs}) \end{array} \right] = \text{negl}(n). \quad (9)$$

Succinctness: All proofs π produced by \mathbf{P} satisfy $|\pi| = \text{poly}(n)(|x| + |w|)^{o(1)}$.

We say that a PPT $\bar{\mathbf{P}}$ breaks a SNARG system $\Pi = (\mathbf{G}, \mathbf{P}, \mathbf{V})$ if the probability that \mathbf{V} outputs 1 in the soundness experiment with $\bar{\mathbf{P}}$ exceeds $1/p(n)$ for some polynomial p and infinitely many n . We will also sometimes consider inefficient $\bar{\mathbf{P}}$ that can win the soundness experiment, but this will of course not amount to a break of Π .

Remark. Note that we define soundness to hold against PPT adversaries (as opposed to QPT). We consider PPT adversaries because ruling out quantum reductions with oracle access to classical adversaries is a strictly weaker setting than when the reduction is given oracle access to quantum adversaries. Since our goal is to rule out a certain class of reduction, it is a stronger result to rule out reductions in the weaker setting.

Next, we define the notion of a falsifiable assumption. We say that an algorithm is *interactive* if it can produce outputs and receive inputs during multiple points in its execution. Here we are only concerned with classical interactions. The inputs received and outputs produced by an interactive quantum algorithm will thus always be classical. An “interaction” between two interactive algorithms \mathbf{C} and \mathbf{A} can itself be viewed as an algorithm, denoted by $\langle \mathbf{C}, \mathbf{A} \rangle$. When we discuss such an interaction, we will assume that \mathbf{C} and \mathbf{A} are designed to receive inputs and send outputs in a way that allows them to interact in the appropriate order.

Definition 9 (Post-Quantum Falsifiable Assumption). A post-quantum falsifiable assumption is a pair (\mathbf{C}, c) where \mathbf{C} is an interactive uniform poly-time classical algorithm (“the challenger”) and c is a constant (“the guessing probability”). The output of \mathbf{C} is a single bit; we let the output 1 mean **win**.

Such an assumption (\mathbf{C}, c) is said to be true if, for every QPT \mathbf{A} ,

$$\Pr[\langle \mathbf{C}, \mathbf{A} \rangle(1^n) = \mathbf{win}] \leq c + \text{negl}(n), \quad (10)$$

where the probability is taken over the coins of \mathbf{C} and the measurements of \mathbf{A} .

For $\delta > 0$, we say that a post-quantum falsifiable assumption (\mathbf{C}, c) is δ -exponentially hard if, for every adversary \mathbf{A} that runs in time $2^{O(n^\delta)}$,

$$\Pr[\langle \mathbf{C}, \mathbf{A} \rangle = \mathbf{win}] \leq c + \frac{1}{2^{O(n^\delta)}}, \quad (11)$$

where the probability is taken over the coins of \mathbf{C} and the measurements of \mathbf{A} .

We say that a QPT \mathbf{A} breaks the assumption (\mathbf{C}, c) if the probability that it wins the security experiment $\langle \mathbf{C}, \mathbf{A} \rangle$ exceeds $c + 1/p(n)$ for some polynomial p and infinitely many n . We say that a QPT \mathbf{A} breaks the δ -exponential version of the assumption (\mathbf{C}, c) if the probability that it wins the security experiment $\langle \mathbf{C}, \mathbf{A} \rangle$ exceeds $c + 1/2^{O(n^\delta)}$ for infinitely many n .

Definition 10 (Quantum Black-Box Reduction for SNARGs). Let Π be a SNARG and (\mathbf{C}, c) a post-quantum falsifiable assumption. A black-box reduction (establishing the soundness of Π , based on the assumption (\mathbf{C}, c)) is an interactive QPT oracle algorithm Σ such that for every PPT adversary \mathcal{E} that breaks the SNARG system Π , $\Sigma^{|\mathcal{E}|}$ breaks the assumption (\mathbf{C}, c) .

In the above, Σ is granted full quantum oracle access to the PPT \mathcal{E} . Such a quantum oracle is of the form $U_{\mathcal{E}}$ (recall (2)) where we view \mathcal{E} as a deterministic function with input set $\mathcal{I} \times \mathcal{X}$ and output set \mathcal{Y} . Here \mathcal{I} is the set of valid common reference strings, \mathcal{X} is the coin space of \mathcal{E} and \mathcal{Y} is the set of instances and (potentially fake) proofs. Hence

$$U_{\mathcal{E}} : |\text{crs}\rangle_{\mathcal{I}} |\text{coins}\rangle_{\mathcal{X}} \mapsto |y \oplus \mathcal{E}(\text{crs}, \text{coins})\rangle_{\mathcal{Y}}. \quad (12)$$

Each set and each corresponding qubit register above are indexed by the security parameter. Given a Boolean circuit of \mathcal{E} , one can construct a quantum circuit for $U_{\mathcal{E}}$ using generic methods for reversible circuit synthesis. This is a reasonably general notion of black-box reduction for quantum adversaries.

Note that the oracle (12) allows the reduction Σ to select the coins of \mathcal{E} in any manner at all. In particular, it need not sample them uniformly at random. As part of our proof technique, we will also consider reductions Σ that must sample the coins of \mathcal{E} uniformly. This amounts to only granting Σ quantum oracle access to the distribution tuple defined by \mathcal{E} . Specifically, for each $\text{crs} \in \mathcal{I}$, we let $\mathcal{D}_{\mathcal{E}(\text{crs})}$ be the output distribution that \mathcal{E} produces on input crs when using uniformly random coins. This defines a tuple $\mathbf{D}_{\mathcal{E}} = \{\mathcal{D}_{\mathcal{E}(\text{crs})}\}_{\text{crs} \in \mathcal{I}}$ of distributions

(see Section 2.1). Recall that a sample from $\mathbf{D}_{\mathcal{E}}^{\mathcal{X}}$ is then a function $f : \mathcal{I} \times \mathcal{X} \rightarrow \mathcal{Y}$ such that each $f(\text{crs}, x)$ is a sample from $\mathcal{D}_{\mathcal{E}(\text{crs})}$. We let $\Sigma^{|\mathcal{E}|s}$ denote the process of first sampling f according to $\mathbf{D}_{\mathcal{E}}^{\mathcal{X}}$, and then running Σ with oracle

$$U_{\mathcal{E}}^s : |\text{crs}\rangle_{\mathcal{I}} |\text{coins}\rangle_{\mathcal{X}} \mapsto |y \oplus f(\text{crs}, \text{coins})\rangle_{\mathcal{Y}} . \quad (13)$$

In this case, we say that Σ has quantum *distributional* oracle access to \mathcal{E} .

2.5 Quantum-hard Membership Problems

Definition 11 (Subset Membership Problem). *A subset membership problem for the class NP consists of an NP language L with instance-witness relation R , and*

- *A distribution family $\mathcal{L} = \{\mathcal{L}_n\}_{n \in \mathbb{N}}$ over the language L and $\bar{\mathcal{L}} = \{\bar{\mathcal{L}}_n\}_{n \in \mathbb{N}}$ over $\bar{L} = \{0, 1\}^* \setminus L$. $\bar{\mathcal{L}}$ need not be efficiently samplable.*
- *An efficient sampling algorithm $(x, w) \leftarrow \text{Sam}(1^n)$ whose support lies in the relation R and whose projection to the first coordinate yields the distribution family $\mathcal{L} = \{\mathcal{L}_n\}_{n \in \mathbb{N}}$*

Definition 12 (Quantum-hard Subset Membership Problem). *Let $(\mathcal{L}, \bar{\mathcal{L}}, \text{Sam})$ be a subset membership problem over an NP language L , and let $s : \mathbb{N} \rightarrow \mathbb{N}$ and $\varepsilon : \mathbb{N} \rightarrow [0, 1]$. We say $(\mathcal{L}, \bar{\mathcal{L}}, \text{Sam})$ is (s, ε) -quantum-hard if distributions \mathcal{L} and $\bar{\mathcal{L}}$ are (s, ε) -sample-indistinguishable. We say the problem is sub-exponentially quantum hard if there exists a constant $\delta > 0$ such that the problem is (s, ε) -quantum-hard for $s(n) = 2^{\Omega(n^\delta)}$ and $\varepsilon(n) = 1/2^{\Omega(n^\delta)}$.*

As discussed in [GW11], the existence of any (sub-)exponentially hard subset membership problem $(\mathcal{L}_n, \bar{\mathcal{L}}_n, \text{Sam})$ implies the existence of a $(2^{n^\delta}, 2^{-n^\delta})$ -hard membership problem for an arbitrarily large constant δ by defining the problem as $(\mathcal{L}_{p(n)}, \bar{\mathcal{L}}_{p(n)}, \text{Sam})$ for a sufficiently large polynomial $p(\cdot)$. The same reasoning applies to (sub-)exponentially quantum-hard subset membership problems.

3 Quantum Oracle-Indistinguishability

In this Section, we prove some results about quantum oracle indistinguishability of distributions and tuples of distributions.

3.1 Oracle-Indistinguishability of Distributions

Zhandry [Zha12a, Theorem 1.1] shows that, for both poly-time distinguishers and unbounded distinguishers, the task of distinguishing distributions using quantum queries is not easier than distinguishing using individual samples. Below we give a generalized version of this result. Our generalization is parameterized by circuit size, query count, and distinguishing advantage, handles non-uniform distinguishers, and allows for distributions that are not efficiently samplable. The proof is similar to that of Zhandry; the main difference is that the non-uniform advice is used to store samples from the two distributions.

Theorem 13. *There exists a polynomial p such that the following holds. Let \mathcal{Y} be a set, and let \mathcal{D}_1 and \mathcal{D}_2 be distributions over \mathcal{Y} . If \mathcal{D}_1 and \mathcal{D}_2 are (s^*, ε^*) -sample-indistinguishable, then they are also (s, q, ε) -oracle-indistinguishable, for all s, q , and ε that satisfy*

$$s^*(n) \geq p(s(n)) + \frac{216q(n)^3 \log |\mathcal{Y}|}{\varepsilon(n)} \quad (14)$$

$$\varepsilon^*(n) \leq \frac{\varepsilon(n)^2}{216q(n)^3} \quad (15)$$

for all sufficiently large n .

Recalling Definition 3 and Definition 6, we note that the above theorem only considers non-uniform quantum adversaries with classical advice, and that this advice is accounted for in the circuit-size bound. In particular, quantum advice⁷ is not allowed.

Proof. Let \mathcal{D}_1 and \mathcal{D}_2 be (s^*, ε^*) -sample-indistinguishable distributions on \mathcal{Y} . Assume, for sake of contradiction, that there exists a set \mathcal{X} and a size- s quantum circuit family \mathcal{B} that makes at most q queries and satisfies

$$\left| \Pr_{\mathcal{O} \leftarrow \mathcal{D}_1^{\mathcal{X}}} [\mathcal{B}^{|\mathcal{O}\rangle}(1^n) = 1] - \Pr_{\mathcal{O} \leftarrow \mathcal{D}_2^{\mathcal{X}}} [\mathcal{B}^{|\mathcal{O}\rangle}(1^n) = 1] \right| > \varepsilon(n),$$

where s, q , and ε satisfy Equations (14) and (15) (for a certain polynomial p that we will describe later in the proof). By Lemma 7, it holds that

$$\left| \Pr_{\mathcal{O} \leftarrow \text{SR}_r^{\mathcal{D}_1}(\mathcal{X})} [\mathcal{B}^{|\mathcal{O}\rangle}(1^n) = 1] - \Pr_{\mathcal{O} \leftarrow \text{SR}_r^{\mathcal{D}_2}(\mathcal{X})} [\mathcal{B}^{|\mathcal{O}\rangle}(1^n) = 1] \right| > \frac{\varepsilon(n)}{2}. \quad (16)$$

for $r = 4 \cdot 27(q(n))^3 / \varepsilon(n) = 108(q(n))^3 / \varepsilon(n)$.

We now define $r + 1$ hybrid oracle distributions $\mathsf{H}_0, \dots, \mathsf{H}_r$. Each hybrid will first draw r -many samples y_1, y_2, \dots, y_r , and then produce an oracle; the oracle will, for each input x , return y_i for uniformly random $i \in [r]$. Hybrid H_i will sample y_1, y_2, \dots, y_i according to \mathcal{D}_2 and $y_{i+1}, y_{i+2}, \dots, y_r$ according to \mathcal{D}_1 . Note that in H_0 the oracle is drawn from $\text{SR}_r^{\mathcal{D}_1}(\mathcal{X})$, while in H_r the oracle is drawn from $\text{SR}_r^{\mathcal{D}_2}(\mathcal{X})$. Defining

$$\varepsilon_i := \left| \Pr_{\mathcal{O} \leftarrow \mathsf{H}_{i+1}} [\mathcal{B}^{|\mathcal{O}\rangle}(1^n) = 1] - \Pr_{\mathcal{O} \leftarrow \mathsf{H}_i} [\mathcal{B}^{|\mathcal{O}\rangle}(1^n) = 1] \right|,$$

it follows by Equation (16) that $\sum_{i=1}^r \varepsilon_i > \varepsilon(n)/2$.

Next, we construct a sample distinguisher \mathcal{A} as follows. As its non-uniform advice, \mathcal{A} receives $2r$ samples: r samples from \mathcal{D}_1 and r samples from \mathcal{D}_2 . As

⁷ One can state and prove an appropriate version of Theorem 13 that allows for quantum advice. We choose to present the simpler classical-advice version because the particular application we consider in this work does not allow for quantum advice.

its challenge, \mathcal{A} obtains a sample y , drawn either from \mathcal{D}_1 or \mathcal{D}_2 . It samples a uniformly random index $i \leftarrow [r]$ and then uses its samples to construct an oracle $\mathcal{O}_i \leftarrow \mathcal{H}_i$ (Note that only r out of the $2r$ samples are used). Conditioned on a particular index choice i' for i , \mathcal{A} simulates \mathcal{B} in either hybrid $\mathcal{H}_{i'-1}$ (if y is drawn from \mathcal{D}_1) or hybrid $\mathcal{H}_{i'}$ (if y is drawn from \mathcal{D}_2). Hence, for all i' ,

$$\left| \Pr_{y \leftarrow \mathcal{D}_1} [\mathcal{A}(1^n, y) = 1 | i = i'] - \Pr_{y \leftarrow \mathcal{D}_2} [\mathcal{A}(1^n, y) = 1 | i = i'] \right| = \varepsilon_{i'}. \quad (17)$$

Since i is sampled independently of all other events, it follows that

$$\left| \Pr_{y \leftarrow \mathcal{D}_1} [\mathcal{A}(1^n, y) = 1] - \Pr_{y \leftarrow \mathcal{D}_2} [\mathcal{A}(1^n, y) = 1] \right| \quad (18)$$

$$= \frac{1}{r} \sum_{i'=1}^r \varepsilon_{i'} = \frac{\varepsilon^*(n)}{108q(n)^3} \sum_{i'=1}^r \varepsilon_{i'} \quad (19)$$

$$> \frac{\varepsilon(n)}{108q(n)^3} \cdot \frac{\varepsilon(n)}{2} \quad (20)$$

$$= \frac{(\varepsilon(n))^2}{216q(n)^3} \geq \varepsilon^*(n), \quad (21)$$

where the last inequality follows from the assumption that Equation (15) holds. Finally, we compute the circuit size $|\mathcal{A}(1^n)|$ of \mathcal{A} . The circuit of \mathcal{A} consists of the circuit for \mathcal{B} , the $2r$ samples from \mathcal{D}_1 and \mathcal{D}_2 , and the gates needed to implement the circuit of the oracle \mathcal{O}_i (once for each query made by \mathcal{B}). To construct \mathcal{O}_i , we use a $2q$ -wise independent function [Zha12a] to simulate the random selection of y_i . Overall, the construction is straightforward to implement with a number of gates that is polynomial in q and the number of input and output bits of the oracle function. Thus there exists a polynomial p' such that

$$|\mathcal{A}(1^n)| < s(n) + 2r \log |\mathcal{Y}| + p'(q(n), \log |\mathcal{X}|, \log |\mathcal{Y}|) \quad (22)$$

$$\leq s(n) + \frac{216(q(n))^3 \log |\mathcal{Y}|}{\varepsilon(n)} + p'(q(n), \log |\mathcal{X}|, \log |\mathcal{Y}|) \quad (23)$$

$$\leq s(n) + \frac{216(q(n))^3 \log |\mathcal{Y}|}{\varepsilon(n)} + p'(s(n)) \quad (24)$$

where the last inequality follows from the assumption that Equation (14) holds. We also assumed that the number of gates of \mathcal{B} is larger than both the number of its queries and the number of input qubits for a single oracle query.

Defining p by $p(m) := p'(m) + m$, we have that $|\mathcal{A}(1^n)| < s^*(n)$. Together with Equation (21), this implies that \mathcal{A} violates the (s^*, ε^*) -sample-indistinguishability of \mathcal{D}_1 and \mathcal{D}_2 . This yields a contradiction and completes the proof. \square

3.2 Oracle-Indistinguishability for Distribution Tuples.

We now consider the task of distinguishing between two tuples of distributions $\mathbf{D}_0 = \{\mathcal{D}_{0,i}\}_{i \in [k]}$ and $\mathbf{D}_1 = \{\mathcal{D}_{1,i}\}_{i \in [k]}$ using an oracle that returns samples.

Similarly to Definition 6, the distinguisher's oracle will be a function f sampled from the set $\mathbf{D}_b^{\mathcal{X}}$, for some bit b and set \mathcal{X} . Recall that f has two inputs: a distribution index i and an input $x \in \mathcal{X}$; the output is then a sample $f(i, x)$ from the distribution $\mathcal{D}_{b,i}$. In our setting, the distinguisher is allowed to query both the distribution index i and the input x in superposition, as in Fig. 1.

Definition 14 ((s, q, ε) -oracle-indistinguishability for tuples). *Let $s, q : \mathbb{N} \rightarrow \mathbb{N}$ and $\varepsilon : \mathbb{N} \rightarrow [0, 1]$, and let $\mathbf{D}_0 = \{\mathcal{D}_{0,i}\}_{i \in [k]}$ and $\mathcal{D}_1 = \{\mathcal{D}_{1,i}\}_{i \in [k]}$ be two distribution tuples over an output set \mathcal{Y} . We say that \mathbf{D}_0 and \mathbf{D}_1 are (s, q, ε) -oracle-indistinguishable if, for every input set \mathcal{X} and every size- s quantum (oracle) circuit family \mathbf{A} that makes at most q queries,*

$$\left| \Pr_{\mathbf{O} \leftarrow \mathbf{D}_0^{\mathcal{X}}} [\mathbf{A}^{|\mathbf{O}\rangle}(1^n) = 1] - \Pr_{\mathbf{O} \leftarrow \mathbf{D}_1^{\mathcal{X}}} [\mathbf{A}^{|\mathbf{O}\rangle}(1^n) = 1] \right| \leq \varepsilon(n) \quad (25)$$

for all sufficiently large n .

We begin with a generalization of Lemma 7, showing that tuples are oracle-indistinguishable from the associated small-range distributions, even against unbounded quantum distinguishers. Given a tuple $\mathbf{D} = \{\mathcal{D}_i^{\mathcal{X}}\}_{i \in [k]}$ of distributions, we define the associated tuple $\text{SR}_r^{\mathbf{D}} = \{\text{SR}_r^{\mathcal{D}_i}(\mathcal{X})\}_{i \in [k]}$ of small-range distributions.

Lemma 15. *Let \mathcal{Y} be a set, and let $\mathbf{D} = \{\mathcal{D}_i^{\mathcal{X}}\}_{i \in [k]}$ be a tuple of distributions over \mathcal{Y} . Then, for any \mathcal{X} , the output distributions of a quantum algorithm making q quantum queries to either $\mathbf{D}^{\mathcal{X}}$ or $\text{SR}_r^{\mathbf{D}}(\mathcal{X})$ are $\frac{k\ell(q)}{r}$ -close where $\ell(q) = \frac{\pi^2(2q)^3}{3}$.*

Proof. Assume towards a contradiction that there exists a distinguisher \mathcal{B} making q quantum queries whose output distributions are not $\frac{k\ell(q)}{r}$ -close. Without loss of generality, we can assume that

$$\left| \Pr_{\mathbf{O} \leftarrow \mathbf{D}^{\mathcal{X}}} [\mathcal{B}^{|\mathbf{O}\rangle}(1^n) = 1] - \Pr_{\mathbf{O} \leftarrow \text{SR}_r^{\mathbf{D}}(\mathcal{X})} [\mathcal{B}^{|\mathbf{O}\rangle}(1^n) = 1] \right| > \frac{k\ell(q)}{r}. \quad (26)$$

We will show that there exists an index i^* and a quantum algorithm \mathcal{A} such that the output distributions of \mathcal{A} making q queries to $\mathcal{D}_{i^*}^{\mathcal{X}}$ and $\text{SR}_r^{\mathcal{D}_{i^*}}(\mathcal{X})$ are not $\frac{\ell(q)}{r}$ -close. This will violate Lemma 7. Define a hybrid sequence of distribution tuples $\mathbf{R}(i)$ for $i \in \{0, 1, \dots, k\}$:

$$\mathbf{R}(i) = (\underbrace{\mathcal{D}_{0,1}, \dots, \mathcal{D}_{0,i-1}}_{i-1}, \mathcal{D}_{0,i}, \underbrace{\text{SR}_r^{\mathcal{D}_{1,i+1}}, \dots, \text{SR}_r^{\mathcal{D}_{1,k}}}_{k-i}). \quad (27)$$

Next, define

$$\varepsilon_i := \left| \Pr_{\mathbf{O} \leftarrow \mathbf{R}(i-1)^{\mathcal{X}}} [\mathcal{B}^{|\mathbf{O}\rangle}(1^n) = 1] - \Pr_{\mathbf{O} \leftarrow \mathbf{R}(i)^{\mathcal{X}}} [\mathcal{B}^{|\mathbf{O}\rangle}(1^n) = 1] \right|. \quad (28)$$

Since $\mathbf{D}_0 = \mathbf{R}(0)$ and $\mathbf{D}_1 = \mathbf{R}(k)$, our initial assumption implies

$$\frac{k\ell(q)}{r} < \left| \Pr_{\mathbf{O} \leftarrow \mathbf{R}(0)^{\mathcal{X}}}[\mathcal{B}^{|\mathbf{O}\rangle}(1^n) = 1] - \Pr_{\mathbf{O} \leftarrow \mathbf{R}(k)^{\mathcal{X}}}[\mathcal{B}^{|\mathbf{O}\rangle}(1^n) = 1] \right| \leq \sum_{i=1}^k \varepsilon_i. \quad (29)$$

Choosing $i^* := \arg \max_i \{\varepsilon_i\}$, we have $\frac{k\ell(q)}{r} < k\varepsilon_{i^*}$ and hence $\varepsilon_{i^*} > \frac{\ell(q)}{r}$.

We now construct a distinguisher \mathcal{A} for the problem of distinguishing $\mathcal{D}_{i^*}^{\mathcal{X}}$ from $\text{SR}_r^{\mathcal{D}_{i^*}}(\mathcal{X})$. Given an oracle U_f (as defined in (2)) where either $f \leftarrow \mathcal{D}_{i^*}^{\mathcal{X}}$ or $f \leftarrow \text{SR}_r^{\mathcal{D}_{i^*}}(\mathcal{X})$, $\mathcal{A}^{(f)}$ proceeds as follows.

1. For each $j \in [k] \setminus i^*$, sample a function f_j so that $f_j \leftarrow \mathcal{D}_j^{\mathcal{X}}$ for $j < i^*$ and $f_j \leftarrow \text{SR}_r^{\mathcal{D}_j}(\mathcal{X})$ for $j > i^*$. Use f_j to construct, for each $j \in [k] \setminus i^*$, a circuit for the unitary U_{f_j} .
2. Use the circuits from the previous step together with the oracle gate for CU_f to construct a circuit for the unitary

$$V = \sum_{j \in [k] \setminus i^*} |j\rangle \langle j| \otimes U_{f_j} + |i^*\rangle \langle i^*| \otimes U_f. \quad (30)$$

(Recall from Section 2.2 that access to U_f implies access to CU_f .)

3. Run \mathcal{B} , answering all of its queries with V . Output whatever \mathcal{B} outputs.

Observe that if $f \leftarrow \mathcal{D}_{i^*}^{\mathcal{X}}$ then V simulates an oracle drawn from $\mathbf{R}(i^*)$, and if $f \leftarrow \text{SR}_r^{\mathcal{D}_{i^*}}(\mathcal{X})$ then V simulates an oracle drawn from $\mathbf{R}(i^* - 1)$. Therefore

$$\left| \Pr_{\mathbf{O} \leftarrow \mathcal{D}_{i^*}^{\mathcal{X}}}[\mathcal{A}^{|\mathbf{O}\rangle}(1^n) = 1] - \Pr_{\mathbf{O} \leftarrow \text{SR}_r^{\mathcal{D}_{i^*}}(\mathcal{X})}[\mathcal{A}^{|\mathbf{O}\rangle}(1^n) = 1] \right| > \frac{\ell(q)}{r}. \quad (31)$$

This contradicts Lemma 7 and the theorem statement follows. \square

Remark. While the indistinguishability bound in Lemma 15 scales with the size k of the distribution tuple, it is likely that a better bound is possible. However, we elect to use this bound since the proof is straightforward and the bound is sufficient for our application.

Corollary 16. *Let \mathcal{Y} be a set, and let $\mathbf{D} = \{\mathcal{D}_i^{\mathcal{X}}\}_{i \in [k]}$ be a tuple of distributions over \mathcal{Y} and let $r, r' > 0$. Then, for any \mathcal{X} , the output distributions of a quantum algorithm making q quantum queries to either $\text{SR}_{r_1}^{\mathbf{D}}(\mathcal{X})$ or $\text{SR}_{r_2}^{\mathbf{D}}(\mathcal{X})$ are $k\ell(q)\left(\frac{1}{r_1} - \frac{1}{r_2}\right)$ -close where $\ell(q) = \frac{\pi^2(2q)^3}{3}$.*

Proof. Directly follows from two applications of Lemma 15.

Theorem 17. *There exists a polynomial p such that the following holds. Let $\mathbf{D}_0 = \{\mathcal{D}_{0,i}\}_{i \in [k]}$ and $\mathbf{D}_1 = \{\mathcal{D}_{1,i}\}_{i \in [k]}$ be tuples of distributions over an output set \mathcal{Y} for a function $k(n)$. If $\mathcal{D}_{0,i}$ and $\mathcal{D}_{1,i}$ are (s^*, ε^*) -sample-indistinguishable*

for all $i \in [k]$, then \mathbf{D}_0 and \mathbf{D}_1 are (s, q, ε) -oracle-indistinguishable for all s, q , and ε satisfying the following conditions for all sufficiently large n :

$$s^*(n) \geq s(n) + \frac{q(n)^3 k(n)^3 p(n)}{\varepsilon(n)} \quad (32)$$

$$\varepsilon^*(n) \leq O\left(\frac{\varepsilon(n)^2}{k(n)^2 q(n)^3}\right) \quad (33)$$

Proof. Assume towards a contradiction that there exists a set \mathcal{X} and a size- $s(n)$ (non-uniform, oracle) distinguisher \mathcal{B} that makes at most q queries and satisfies

$$\left| \Pr_{\mathcal{O} \leftarrow \mathbf{D}_0^{\mathcal{X}}}[\mathcal{B}^{|\mathcal{O}|}(1^n) = 1] - \Pr_{\mathcal{O} \leftarrow \mathbf{D}_1^{\mathcal{X}}}[\mathcal{B}^{|\mathcal{O}|}(1^n) = 1] \right| > \varepsilon(n)$$

where s, q , and ε satisfy Equations (32) and (33) for the given s^* and ε^* . We will argue that there exists an index i^* and a distinguisher \mathcal{A}_{i^*} (constructed from \mathcal{B}) that distinguishes oracles drawn from $\mathcal{D}_{0,i^*}^{\mathcal{X}}$ and $\mathcal{D}_{1,i^*}^{\mathcal{X}}$ with a smaller circuit and greater probability than allowed by Theorem 13.

Define a hybrid sequence of distribution tuples $\mathbf{R}(i)$ for $i \in \{1, \dots, k\}$:

$$\mathbf{R}(i) = (\underbrace{\mathcal{D}_{0,1}, \dots, \mathcal{D}_{0,i-1}}_{i-1}, \mathcal{D}_{0,i}, \underbrace{\mathcal{D}_{1,i+1}, \dots, \mathcal{D}_{1,k}}_{k-i}). \quad (34)$$

Next, define

$$\varepsilon_i := \left| \Pr_{\mathcal{O} \leftarrow \mathbf{R}(i-1)^{\mathcal{X}}}[\mathcal{B}^{|\mathcal{O}|}(1^n) = 1] - \Pr_{\mathcal{O} \leftarrow \mathbf{R}(i)^{\mathcal{X}}}[\mathcal{B}^{|\mathcal{O}|}(1^n) = 1] \right|. \quad (35)$$

Since $\mathbf{D}_0 = \mathbf{R}(0)$ and $\mathbf{D}_1 = \mathbf{R}(k)$, our initial assumption implies

$$\varepsilon < \left| \Pr_{\mathcal{O} \leftarrow \mathbf{R}(0)^{\mathcal{X}}}[\mathcal{B}^{|\mathcal{O}|}(1^n) = 1] - \Pr_{\mathcal{O} \leftarrow \mathbf{R}(k)^{\mathcal{X}}}[\mathcal{B}^{|\mathcal{O}|}(1^n) = 1] \right| \leq \sum_{i=1}^k \varepsilon_i. \quad (36)$$

Choose $i^* = \arg \max_i \{\varepsilon_i\}$. Then we have $\varepsilon < k\varepsilon_{i^*}$ and hence $\varepsilon_{i^*} > \varepsilon/k$.

Next, we set $r = \frac{4(k-1)\ell(q)}{\varepsilon_{i^*}}$ with $\ell(q)$ being the polynomial from Lemma 7. We then define additional tuples $\mathbf{R}'(i^*)$ and $\mathbf{R}''(i^*)$ below. These tuples are sandwiched between the already-defined $\mathbf{R}(i^* - 1)$ and $\mathbf{R}(i^*)$ in a natural way.

$$\begin{aligned} \mathbf{R}(i^* - 1) &= (\mathcal{D}_{0,1}, \dots, \mathcal{D}_{0,i^*-1}, \mathcal{D}_{1,i^*}, \mathcal{D}_{1,i^*+1}, \dots, \mathcal{D}_{1,k}) \\ \mathbf{R}'(i^*) &= (\text{SR}_r^{\mathcal{D}_{0,1}}, \dots, \text{SR}_r^{\mathcal{D}_{0,i^*-1}}, \mathcal{D}_{1,i^*}, \text{SR}_r^{\mathcal{D}_{1,i^*+1}}, \dots, \text{SR}_r^{\mathcal{D}_{1,k}}) \\ \mathbf{R}''(i^*) &= (\text{SR}_r^{\mathcal{D}_{0,1}}, \dots, \text{SR}_r^{\mathcal{D}_{0,i^*-1}}, \mathcal{D}_{0,i^*}, \text{SR}_r^{\mathcal{D}_{1,i^*+1}}, \dots, \text{SR}_r^{\mathcal{D}_{1,k}}) \\ \mathbf{R}(i^*) &= (\underbrace{\mathcal{D}_{0,1}, \dots, \mathcal{D}_{0,i^*-1}}_{i^*-1}, \mathcal{D}_{0,i^*}, \underbrace{\mathcal{D}_{1,i^*+1}, \dots, \mathcal{D}_{1,k}}_{k-i^*}). \end{aligned}$$

By Lemma 15,

$$\left| \Pr_{\mathcal{O} \leftarrow \mathbf{R}(i^*-1)^{\mathcal{X}}}[\mathcal{B}^{|\mathcal{O}|}() = 1] - \Pr_{\mathcal{O} \leftarrow \mathbf{R}'(i^*)^{\mathcal{X}}}[\mathcal{B}^{|\mathcal{O}|}() = 1] \right| \leq \frac{(k-1)\ell(q)}{r} \quad (37)$$

$$\left| \Pr_{\mathbf{O} \leftarrow \mathbf{R}''(i^*)^{\mathcal{X}}}[\mathcal{B}^{|\mathbf{O}\rangle}() = 1] - \Pr_{\mathbf{O} \leftarrow \mathbf{R}(i^*)^{\mathcal{X}}}[\mathcal{B}^{|\mathbf{O}\rangle}() = 1] \right| \leq \frac{(k-1)\ell(q)}{r}. \quad (38)$$

Recalling our choice of r above, we obtain that

$$\left| \Pr_{\mathbf{O} \leftarrow \mathbf{R}'(i^*)^{\mathcal{X}}}[\mathcal{B}^{|\mathbf{O}\rangle}() = 1] - \Pr_{\mathbf{O} \leftarrow \mathbf{R}''(i^*)^{\mathcal{X}}}[\mathcal{B}^{|\mathbf{O}\rangle}() = 1] \right| \geq \frac{\varepsilon_{i^*}}{2}. \quad (39)$$

We now construct a non-uniform algorithm \mathcal{A}_{i^*} for the task of, given an oracle for some function g , distinguishing whether g was drawn from $\mathcal{D}_{0,i^*}^{\mathcal{X}}$ or $\mathcal{D}_{1,i^*}^{\mathcal{X}}$:

1. The non-uniform advice consists of the index i^* and $r(k-1)$ distribution samples. Specifically, for each $i < i^*$, there are r -many samples from $\mathcal{D}_{0,i}$, and for each $i > i^*$ there are r -many samples from $\mathcal{D}_{1,i}$.
2. \mathcal{A}_{i^*} uses the samples to construct, for each $i \neq i^*$, a function $f_i \leftarrow \text{SR}_r^{\mathcal{D}_{c,i}}(\mathcal{X})$ (where $c = 0$ iff $i < i^*$) and a reversible circuit for the unitary U_{f_i} (see Equation (2)).
3. Using the circuits for U_{f_i} and the oracle gate U_g , \mathcal{A}_{i^*} constructs a circuit for

$$V = |i^*\rangle\langle i^*| \otimes U_g + \sum_{i \in [k] \setminus i^*} |i\rangle\langle i| \otimes U_{f_i} \quad (40)$$

4. \mathcal{A}_{i^*} runs \mathcal{B} with V as the oracle. The details of the simulation can be found Appendix B. \mathcal{A} ends by outputting whatever \mathcal{B} outputs.

Observe that if $g \leftarrow \mathcal{D}_{1,i^*}^{\mathcal{X}}$, then \mathcal{A}_{i^*} simulates \mathcal{B} with access to the tuple $\mathbf{R}'(i^*)$, and if $g \leftarrow \mathcal{D}_{0,i^*}^{\mathcal{X}}$, then \mathcal{A}_{i^*} simulates \mathcal{B} with access to the tuple $\mathbf{R}''(i^*)$. Together with Equation (39), we get

$$\left| \Pr_{\mathbf{O} \leftarrow \mathcal{D}_{0,i^*}^{\mathcal{X}}}[\mathcal{A}_{i^*}^{|\mathbf{O}\rangle}() = 1] - \Pr_{\mathbf{O} \leftarrow \mathcal{D}_{1,i^*}^{\mathcal{X}}}[\mathcal{A}_{i^*}^{|\mathbf{O}\rangle}() = 1] \right| \geq \frac{\varepsilon_{i^*}}{2} > \frac{\varepsilon}{2k}. \quad (41)$$

By Theorem 13, \mathcal{D}_{0,i^*} and \mathcal{D}_{1,i^*} are $(\tilde{s}, q, \tilde{\varepsilon})$ -oracle-indistinguishable where

$$\varepsilon^* \leq \frac{\tilde{\varepsilon}^2}{216q^3}. \quad (42)$$

Setting $\tilde{\varepsilon} = \frac{\varepsilon}{2k}$, it then follows by assumption that

$$\left| \Pr_{\mathbf{O} \leftarrow \mathcal{D}_{0,i^*}^{\mathcal{X}}}[\mathcal{A}_{i^*}^{|\mathbf{O}\rangle}() = 1] - \Pr_{\mathbf{O} \leftarrow \mathcal{D}_{1,i^*}^{\mathcal{X}}}[\mathcal{A}_{i^*}^{|\mathbf{O}\rangle}() = 1] \right| > \frac{\varepsilon^2}{864k^2q^3}. \quad (43)$$

It remains to verify that the circuit size of \mathcal{A}_{i^*} violates Equation (32) (for a certain polynomial p). The circuit of \mathcal{A}_{i^*} consists of the circuit of \mathcal{B} , the $r(k-1)$ samples used to construct the small-range distributions, and the circuit used to construct the oracle V . The circuit to construct V constructs the small range distributions for every distribution index and will generate a sample controlled on a specific index. To do this, the oracle for V will utilize a $2q$ -wise independent

function for each index and query (which can be implemented using $O(q^2)$ many gates⁸). Overall, the construction of V can be implemented with a number of gates that is linear in k and polynomial in q , the size of the input, and the size of the output of the oracle. Recall that the size of the input and output are both polynomial in n . Therefore, there exists a certain polynomial p such that

$$|\mathcal{A}_{i^*}(1^n)| = s(n) + r(k-1) \log |\mathcal{Y}| + kq \cdot p(q, \log |\mathcal{X}|, \log |\mathcal{Y}|) \quad (44)$$

$$= s(n) + \frac{4k(k-1)^2 l(q(n)) \log |\mathcal{Y}|}{\varepsilon} + kq \cdot p(q, \log |\mathcal{X}|, \log |\mathcal{Y}|) \quad (45)$$

$$\leq s(n) + \frac{106q(n)^3 k^3 \log |\mathcal{Y}|}{\varepsilon} + kq^3 \cdot p(n) \quad (46)$$

$$< s^*(n). \quad (47)$$

This yields a contradiction to Theorem 13 since we have constructed a distinguisher \mathcal{A}_{i^*} that is of size $|\mathcal{A}_{i^*}(1^n)| < s^*(n)$ but distinguishes the distributions \mathcal{D}_{0,i^*} and \mathcal{D}_{1,i^*} with probability greater than the bound specified in Equation (15) which then concludes the proof. \square

4 SNARG Impossibility

In this section we describe the meta-reduction for post-quantum SNARGs. We show if there exists an efficient quantum black-box reduction from a SNARG to a falsifiable assumption, then there exists an efficient quantum attack on the assumption (cf. Section 4.2). Before that, we show that there exists a (possibly inefficient) classical SNARG adversary $\bar{\mathcal{P}}$ that can be efficiently simulated by a PPT simulator \mathcal{S} for a QPT distinguisher \mathcal{D} that is given access to only classical non-uniform advice (cf. Section 4.1). This implies an efficient attack for the sub-exponentially hard subset membership problem $(\text{Sam}, \mathcal{L}, \bar{\mathcal{L}})$. Note that we rule out reductions in the non-security parameter preserving setting, in the \mathcal{D} is permitted quantum queries to both the random coins used by $\bar{\mathcal{P}}$ (resp. \mathcal{S}), the crs, and the security parameter. In this setting, \mathcal{D} may query its oracle on security parameters $m \neq n$. We provide a proof for the security parameter preserving setting (where \mathcal{D} is restricted to querying its oracle on security parameters $m = n$) in Section C. Readers may find it helpful to read the security parameter preserving proof as a warm up.

4.1 A Poly-Time-Simulatable Adversary

Lemma 18. *Let L be a language in NP with a sub-exponentially quantum-hard subset membership problem. Let $\Pi = (\mathbf{G}, \mathbf{P}, \mathbf{V})$ be a SNARG for the language L that satisfies the completeness and succinctness properties as in Definition 8. Then, there is an algorithm $\bar{\mathcal{P}}$, satisfying the following:*

⁸ A 2q-wise independent hash function can be implemented by multiplying degree k polynomials over a finite field.

- $\overline{\mathcal{P}}$ is a stateful and computationally unbounded algorithm. On input $(1^n, \text{crs})$, it always outputs some (x, π) with $x \notin L$ of size $|x| = \text{poly}(n)$, such that:

$$\Pr \left[\mathbf{V}(\text{priv}, \bar{x}, \bar{\pi}) = 1 \mid \begin{array}{l} (\text{crs}, \text{priv}) \leftarrow \mathbf{G}(1^n) \\ (\bar{x}, \bar{\pi}) \leftarrow \overline{\mathcal{P}}(1^n, \text{crs}) \end{array} \right] \geq 1 - \text{negl}(n).$$

- $\overline{\mathcal{P}}$ is poly-time oracle simulatable. That is, for all QPT oracle distinguishers \mathcal{D} and all polynomials $\rho(n)$ there exists a PPT simulator $\tilde{\mathcal{S}}$ such that for sufficiently large n :

$$|\Pr[\mathcal{D}^{\overline{\mathcal{P}}}(1^n) = 1] - \Pr[\mathcal{D}^{\tilde{\mathcal{S}}(1^n)}(1^n) = 1]| < \frac{1}{\rho(n)}.$$

The oracle distinguisher \mathcal{D} can query its oracle using a security parameter m and need not set $m = n$. Recall that \mathcal{D} has quantum oracle access to its oracle and its coins as defined in Equation (12). The simulator $\tilde{\mathcal{S}}$ is given 1^n as input and runs in time polynomial in n on any query.

Proof (Sketch). At a high level, we use the same meta-reduction technique used by Gentry and Wichs. The proof is outlined in the following.

1. **Define a distribution over valid statements and proofs.** We first choose an (s^*, ε^*) -quantum-hard subset membership problem $(\mathcal{L}_n, \bar{\mathcal{L}}_n, \text{Sam})$ for L with $s^*(n) = 2^{n^{d'}}$ and $\varepsilon^*(n) = 2^{-n^{d'}}$ for a constant d' we will pick below. As discussed in Section 2.5, the existence of $(\mathcal{L}_n, \bar{\mathcal{L}}_n, \text{Sam})$ follows from the existence of any quantum-hard subset membership problem. We also define, for every n and every crs , a “(yes-instance, true proof)” distribution $\mathcal{L}_{n, \text{crs}}^{+\text{aux}}$ as follows. To draw a sample from $\mathcal{L}_{n, \text{crs}}^{+\text{aux}}$, generate $(x, w) \leftarrow \text{Sam}(1^n)$, compute $\pi \leftarrow \Pi.P(1^n, \text{crs}, x, w)$, and finally output (x, π) as the sample. Note that x is drawn according to \mathcal{L}_n .

2. **Construct an inefficient Π -adversary $\overline{\mathcal{P}}$.** The adversary $\overline{\mathcal{P}}$ will receive queries of the form $(1^m, \text{crs})$. We choose a certain constant d' above and apply Corollary 5 to $\mathcal{D}_0 := \mathcal{L}_m$, $\mathcal{D}_1 := \bar{\mathcal{L}}_m$, and $\mathcal{D}_0^{\text{aux}_0} := \mathcal{L}_{m, \text{crs}}^{+\text{aux}}$, for every m and every crs . The result is an inefficient algorithm $\overline{\mathcal{P}}$ which, on input $(1^m, \text{crs})$, outputs a “(no-instance, fake proof)” pair $(\bar{x}, \bar{\pi})$ that is $(s'(m), \varepsilon'(m))$ -sample-indistinguishable from the “(yes-instance, true proof)” outputs (x, π) of $\mathcal{L}_{m, \text{crs}}^{+\text{aux}}$ — for certain superpolynomial $s'(m)$ and negligible $\varepsilon'(m)$. By setting $m = n$, it follows that $(\bar{x}, \bar{\pi}) \leftarrow \overline{\mathcal{P}}(1^n, \text{crs})$ is accepted by $\Pi.V(\text{crs}, \cdot)$, which establishes the first part of the lemma.

3. **Construct a Simulator $\tilde{\mathcal{S}}$.** Similar to [GW11] we define $\tilde{\mathcal{S}}(1^n)$ so that, on query $(1^m, \text{crs})$, it generates a response according to a certain threshold $m^*(n)$. If $m \geq m^*(n)$, the query response is a sample $(x, \pi) \leftarrow \mathcal{L}_{m, \text{crs}}^{+\text{aux}}$; if $m < m^*(n)$ the query response is drawn from a table \mathcal{T}_n that is provided to $\tilde{\mathcal{S}}(1^n)$ as non-uniform advice. The table is necessary because, for small m , the statements will be so short that the distinguisher can efficiently distinguish between a true statement x and a false statement \bar{x} without even considering the accompanying (true or fake) proof. It will turn out that a table \mathcal{T}_n containing $\text{poly}(n)$ -many $(\bar{x}, \bar{\pi})$ tuples generated from $\overline{\mathcal{P}}$ can be used appropriately by $\tilde{\mathcal{S}}(1^n)$ to fool even quantum queries.

4. **Show that $\bar{\mathcal{P}}$ and $\tilde{\mathcal{S}}(1^n)$ are Oracle-Indistinguishable.** Finally, we show that the $(s'(m), \varepsilon'(m))$ -sample-indistinguishability of $\bar{\mathcal{P}}$ and $\tilde{\mathcal{S}}(1^n)$ implies that $\bar{\mathcal{P}}$ and $\tilde{\mathcal{S}}(1^n)$ are in fact not distinguishable even by a QPT oracle algorithm making quantum queries.⁹ At first glance, it might seem that Theorem 17 could be used directly to upgrade sample indistinguishability (of $\mathcal{L}_{m, \text{crs}}^{+\text{aux}}$ and $\bar{\mathcal{L}}_{m, \text{crs}}^{+\text{aux}}$) to oracle indistinguishability (of $\bar{\mathcal{P}}$ and $\tilde{\mathcal{S}}(1^n)$). However, this is not the case: the distributions are indexed by both crs and m , and the distinguisher is permitted to query on any m and any crs ; as a result, Theorem 17 would only yield oracle indistinguishability with the parameters associated to the *smallest* m , which is not enough. Instead, we will combine a complexity leveraging argument with a sequence of hybrids over the possible security parameters, applying Theorem 17 repeatedly in each hybrid transition.

Proof (of Lemma 18).

1. Define a distribution over valid statements and proofs. We begin by choosing the hardness of the subset membership problem. By the succinctness of Π , we can choose a sufficiently large constant d such that the length of $\text{crs} \leftarrow \Pi.\mathcal{G}(1^n)$ is bounded by $O(n^d)$ and the length of any proof $\pi \leftarrow \Pi.\mathcal{P}(1^n, x, w)$ is bounded by $O(n^d)(|x| + |w|)^{o(1)}$. By the assumption on L , there exists an $(s^*(n), \varepsilon^*(n))$ -quantum-hard subset membership problem $(\mathcal{L}_n, \bar{\mathcal{L}}_n, \text{Sam})$ with

$$s^*(n) = 2^{cn^{d+3}} \quad \text{and} \quad \varepsilon^*(n) = 2^{-n^{d+3}}. \quad (48)$$

The explicit choice of c will be clear later in Equation (49). We define $\mathcal{L}_{n, \text{crs}}^{+\text{aux}}$ to be the distribution over “(yes-instance, true proof)” pairs (x, π) , generated honestly by running $(x, w) \leftarrow \text{Sam}(1^n)$, $\text{crs} \leftarrow \mathcal{G}(1^n)$ and $\pi \leftarrow \Pi.\mathcal{P}(\text{crs}, x, w)$. Note that $\ell_{\text{pf}}(n) = o(n^{d+3})$ bounds the length of such π .

2. Construct an Inefficient Π -adversary $\bar{\mathcal{P}}$. We now construct $\bar{\mathcal{P}}$. Recall that the queries to $\bar{\mathcal{P}}$ are of the form $(1^m, \text{crs})$, where m need not equal n .

By definition, \mathcal{L}_m and $\bar{\mathcal{L}}_m$ are $(s^*(m), \varepsilon^*(m))$ indistinguishable. Then by applying Corollary 5 to \mathcal{L}_m , $\bar{\mathcal{L}}_m$, and $\mathcal{L}_{m, \text{crs}}^{+\text{aux}}$, we obtain the existence of a distribution $\bar{\mathcal{L}}_{m, \text{crs}}^{+\text{aux}}$ that augments $\bar{\mathcal{L}}_m$ with bitstrings of size at most $\ell_{\text{pf}}(m)$ such that $\mathcal{L}_{m, \text{crs}}^{+\text{aux}}$ and $\bar{\mathcal{L}}_{m, \text{crs}}^{+\text{aux}}$ are $(s'(m), \varepsilon'(m))$ -sample indistinguishable.

Note that if a problem is $(s^*(n), \varepsilon^*(n))$ hard, then it is also $(s^*(n), \tilde{\varepsilon}(n))$ hard for $\tilde{\varepsilon} \geq \varepsilon^*$. We now carefully choose $(s^*(m), \tilde{\varepsilon}(m))$ to obtain the explicit bounds on $s'(m)$ and $\varepsilon'(m)$ given by Corollary 5. We obtain that, for $\tilde{\varepsilon} = 2^{-m^{d+2}}$, there exist some $(s'(m), \varepsilon'(m))$ such that $s^*(m) = \text{poly}(s'(m), \frac{1}{\tilde{\varepsilon}}, m, 2^{\ell_{\text{pf}}})$ and $\varepsilon'(m) = \Theta(\tilde{\varepsilon}(m))$. It follows that there exists a constant k such that

$$s'(m) = \Theta\left(\left(\frac{s^*}{\text{poly}(\frac{1}{\tilde{\varepsilon}}, 2^{\ell_{\text{pf}}}, m)}\right)^{1/c}\right) = \Theta\left(\frac{(2^{cm^{d+3}})^{1/c}}{\text{poly}(\frac{1}{\tilde{\varepsilon}}, 2^{\ell_{\text{pf}}}, m)^{1/c}}\right) \quad (49)$$

⁹ To achieve this, $\bar{\mathcal{P}}$ (resp., $\tilde{\mathcal{S}}$) will also need to apply a hard-wired random (resp., $2q$ -wise-independent) function to its coins.

$$= \frac{2^{\Omega(m^{d+3})}}{(\Theta(2^{k(m^{d+2})})\text{poly}(2^{o(m^{d+3})}, n))^{1/c}} \quad (50)$$

$$= 2^{\Omega(m^{d+3})} \quad (51)$$

and

$$\varepsilon'(m) = \Theta(\tilde{\varepsilon}(m)) = 2^{-\Theta(m^{d+2})}. \quad (52)$$

By the argument above, $\tilde{\mathcal{L}}_{m,\text{crs}}^{+\text{aux}}$ is $(s'(m), \varepsilon'(m))$ sample-indistinguishable from $\mathcal{L}_{m,\text{crs}}^{+\text{aux}}$. We then define $\bar{\mathcal{P}}$ to be the algorithm that, given a query $(1^m, \text{crs})$, samples and outputs $(\bar{x}, \bar{\pi}) \leftarrow \tilde{\mathcal{L}}_{m,\text{crs}}^{+\text{aux}}$.

Setting $m = n$ now yields the first part of the lemma. Specifically, since the size of \mathbf{V} is less than $s'(n)$, we have that

$$\begin{aligned} & \Pr \left[\mathbf{V}(\text{priv}, \bar{x}, \bar{\pi}) = 1 \mid \begin{array}{l} (\text{crs}, \text{priv}) \leftarrow \mathbf{G}(1^n) \\ (\bar{x}, \bar{\pi}) \leftarrow \bar{\mathcal{P}}(1^n, \text{crs}) \end{array} \right] \\ & \geq \Pr \left[\mathbf{V}(\text{priv}, x, \pi) = 1 \mid \begin{array}{l} (\text{crs}, \text{priv}) \leftarrow \mathbf{G}(1^n) \\ (x, \pi) \leftarrow \mathcal{L}_{n,\text{crs}}^{+\text{aux}} \end{array} \right] - \varepsilon'(n) \\ & \geq 1 - \text{negl}(n). \end{aligned} \quad (53)$$

Finally, we construct a ‘‘coin-randomized’’ version of $\bar{\mathcal{P}}$. Define $\bar{\mathcal{P}}_R$ to be the algorithm that first constructs a perfectly random function F (with appropriate domain and range) which for query $(1^m, \text{crs}, \text{coins})$, outputs $(1^m, \text{crs}, F(\text{coins}))$.

3. Construct a simulator $\tilde{\mathcal{S}}_R(1^n)$. To prove the second part of the lemma, our goal is to argue that $\bar{\mathcal{P}}_R$ can be replaced with an efficient simulator, $\tilde{\mathcal{S}}_R(1^n)$. Let \mathcal{D} be a QPT oracle distinguisher and let ρ be a polynomial. We will construct an $\tilde{\mathcal{S}}_R(1^n)$ such that

$$\Pr[\mathcal{D}^{\langle \bar{\mathcal{P}}_R \rangle}(1^n) = 1] - \Pr[\mathcal{D}^{\langle \tilde{\mathcal{S}}_R(1^n) \rangle}(1^n) = 1] < \frac{1}{\rho(n)} \quad (54)$$

for all sufficiently large n . We will first define the version of $\tilde{\mathcal{S}}(1^n)$ that does not randomize the coins provided on input. Later, we will define $\tilde{\mathcal{S}}_R(1^n)$ which acts exactly like $\tilde{\mathcal{S}}(1^n)$ except it first randomizes the provided coins with a $2q$ -wise independent function.

We define $\tilde{\mathcal{S}}(1^n)$ to respond to queries according to a threshold $m^*(n) = \lfloor \log^{1/d+1} n \rfloor$. On query $(1^m, \text{crs})$, $\tilde{\mathcal{S}}(1^n)$ will respond as follows. If $m \geq m^*(n)$, $\tilde{\mathcal{S}}(1^n)$ will respond with a sample $(x, \pi) \leftarrow \mathcal{L}_{m,\text{crs}}^{+\text{aux}}$. If $m < m^*(n)$, $\tilde{\mathcal{S}}$ will respond using a certain table \mathcal{T}_n .

The table \mathcal{T}_n contains values generated by running $\bar{\mathcal{P}}$; as $\bar{\mathcal{P}}$ is inefficient, \mathcal{T}_n will be provided to $\tilde{\mathcal{S}}(1^n)$ as non-uniform advice. Let $q(n) = \text{poly}(n)$ be an upper bound on the number of queries made by $\mathcal{D}(1^n)$. Define $r(n) := \ell(q(n)) \cdot \rho(n)$ where ℓ is the polynomial defined in Lemma 7. Let $\mathcal{X}_m = \{0, 1\}^{b(m)}$ be a set

over the random coins of length $b(m)$ for a polynomial $b(\cdot)$. The table \mathcal{T}_n consists of tuples of the form $(i, m, \text{crs}, \bar{x}, \bar{\pi})$, one for each triple (i, m, crs) where $i \in \{1, \dots, r(n)\}$, $m \in \{1, \dots, m^*(n)\}$, and $\text{crs} \in \{0, 1\}^{m^d}$. Let $\ell_{\text{st}} = |\bar{x}|$. Then, the final size of the table is

$$|\mathcal{T}_n| = \sum_{m=1}^{m^*(n)} r(n) \cdot 2^{m^d} (\ell_{\text{pf}}(m) + \ell_{\text{st}}(m)) \quad (55)$$

$$\leq 2^{m^*(n)^d} \cdot \ell(q(n)) \cdot \rho(n) \cdot m^*(n) (o(m^*(n)^{d+3}) + \text{poly}(m^*(n))) \quad (56)$$

$$= O\left(2^{(\log n)^{\frac{d}{d+1}}}\right) \quad (57)$$

$$= O(\text{poly}(n)) \quad (58)$$

In order to prepare for answering queries on security parameters $m \leq m^*$, $\tilde{\mathcal{S}}(1^n)$ will use the samples in \mathcal{T}_n to construct, for all $m \leq m^*$ and all $\text{crs} \in \{0, 1\}^{m^d}$, a function $f_{m, \text{crs}}$ drawn from the small-range distribution $\text{SR}_r^{\overline{\mathcal{P}}(1^m, \text{crs})}$ for domain \mathcal{X}_m . Then upon an input query $(1^m, \text{crs}, \text{coins})$, $\tilde{\mathcal{S}}(1^n)$ will respond with $f_{m, \text{crs}}(\text{coins})$. Note that the circuit size $|\tilde{\mathcal{S}}(1^n)| \leq \text{poly}(n) \cdot |\mathcal{T}_n|$.

Finally, we construct the coin-randomized version of $\tilde{\mathcal{S}}$. We define $\tilde{\mathcal{S}}_R(1^n)$ to first sample a $2q$ -wise-independent function f_{2q} , and then to respond to queries $(1^m, \text{crs}, \text{coins})$ by querying $\tilde{\mathcal{S}}(1^n)$ with $(1^m, \text{crs}, f_{2q}(\text{coins}))$ and returning the result. Note that ultimately this will allow the simulator to simulate the distribution over inefficient adversaries, $\overline{\mathcal{P}}_R$.

4. Show that $\overline{\mathcal{P}}$ and $\tilde{\mathcal{S}}$ are Oracle-Indistinguishable. We will now show that $\overline{\mathcal{P}}_R$ and $\tilde{\mathcal{S}}_R$ are oracle-indistinguishable against the given \mathcal{D} . We begin by proving the following claim.

Claim.

$$|\Pr[\mathcal{D}^{\overline{\mathcal{P}}_R}(1^n) = 1] - \Pr[\mathcal{D}^{\tilde{\mathcal{S}}_R(1^n)}(1^n) = 1]| \quad (59)$$

$$= |\Pr[\mathcal{D}^{\overline{\mathcal{P}}_s}(1^n) = 1] - \Pr[\mathcal{D}^{\tilde{\mathcal{S}}(1^n)_s}(1^n) = 1]|. \quad (60)$$

Proof. We first define quantum oracle access for \mathcal{D} . Let register 0 contain the security parameter, register 1 contain the crs , and let register 2 contain the random coins all respectively queried by \mathcal{D} . Let $\text{HW}(\cdot)$ denote the function that computes the hamming weight of a bitstring ν and let $m = \text{HW}(\nu)$ be the security parameter. We define the security parameter using $\text{HW}(\cdot)$ to ensure that the oracle is defined on all inputs.

Distributional Oracle Access to $\overline{\mathcal{P}}$ and $\tilde{\mathcal{S}}(1^n)$. We now define the oracles $U_{\overline{\mathcal{P}}}^{\mathcal{S}}$ and $U_{\tilde{\mathcal{S}}}^{\mathcal{S}}$ which define \mathcal{D} 's distributional oracle access to $\overline{\mathcal{P}}$ and $\tilde{\mathcal{S}}(1^n)$ (see Equation 13). $U_{\overline{\mathcal{P}}}^{\mathcal{S}}$ encodes a random oracle, RO , which it uses to first randomize the coins queried by \mathcal{D} . Similarly, the oracle $U_{\tilde{\mathcal{S}}}^{\mathcal{S}}$ encodes a $2q$ -wise independent function, f_{2q} , to randomize the coins.

$$U_{\overline{\mathcal{P}}}^{\mathcal{S}}: |\nu\rangle_0 |\text{crs}\rangle_1 |\text{coins}\rangle_2 |y\rangle_3 \quad (61)$$

$$\mapsto: |\nu\rangle_0 |\text{crs}\rangle_1 |\text{coins}\rangle_2 |y\rangle_3 |RO(\text{coins})\rangle_A \quad (62)$$

$$\mapsto |\nu\rangle_0 |\text{crs}\rangle_1 |\text{coins}\rangle_2 |y \oplus \overline{\mathcal{P}}(1^m, \text{crs}, RO(\text{coins}))\rangle_3 \quad (63)$$

The oracle for $\tilde{\mathcal{S}}(1^n)$ uses an ancilla register A to compute the security parameter relative to the threshold, $m^*(n) = \lfloor \log^{1/(d+1)} n \rfloor$. Since \mathcal{D} is a QPT algorithm, there exists a max security parameter $m_{\max}(n) \leq p(n)$. Therefore, $\nu \in \{0, 1\}^{m_{\max}(n)}$.

We first define two subroutines $U_{\mathcal{T}_n, m}$ and $U_{\mathcal{S}, m}$. Let \mathcal{S} be the algorithm run by the oracle when $m > m^*(n)$. That is, let \mathcal{S} be the PPT algorithm such that on input $(1^m, \text{crs}, \text{coins})$, it generates an honest statement and proof $(x, \pi) \leftarrow \mathcal{L}_{m, \text{crs}}^{+\text{aux}}$.

Let $O \in \{\mathcal{T}_n, \mathcal{S}\}$. The bit contained in register A determines O . That is, if $m \leq m^*(n)$, then $O = \mathcal{T}_n$. If $m > m^*(n)$, then $O = \mathcal{S}$. Note that here we slightly abuse notation and use \mathcal{T}_n to both define the table of samples provided to $\tilde{\mathcal{S}}(1^n)$ as non-uniform advice and also the algorithm that samples from the table.

$$U_{O, m} : |\text{crs}\rangle_1 |\text{coins}\rangle_2 |y\rangle_3 |m \leq m^*(n)\rangle_A \quad (64)$$

$$\mapsto |\text{crs}\rangle_1 |\text{coins}\rangle_2 |y \oplus O(1^m, \text{crs}, \text{coins})\rangle_3 |m \leq m^*(n)\rangle_A \quad (65)$$

The oracle $U_{\tilde{\mathcal{S}}}^{\S}$ will output the response below and uncompute all ancilla registers.

$$U_{\tilde{\mathcal{S}}}^{\S} : |\nu\rangle_0 |\text{crs}\rangle_1 |\text{coins}\rangle_2 |y\rangle_3 \quad (66)$$

$$\mapsto |\nu\rangle_0 |\text{crs}\rangle_1 |\text{coins}\rangle_2 |y\rangle_3 |m \leq m^*(n)\rangle_A \quad (67)$$

$$\mapsto |\nu\rangle_0 U_{O, m} |\text{crs}\rangle_1 |\text{coins}\rangle_2 |y\rangle_3 \quad (68)$$

Recall that $q(n) \leq p(n)$ for a polynomial p . Therefore we can choose f_{2q} such that the output distribution generated by making q queries to f_{2q} is perfectly indistinguishable from that of making q queries to a random function [Zha12a, Fact 2]. Therefore it holds that

$$\Pr[\mathcal{D}^{|\overline{\mathcal{P}}\rangle_{\S}}(1^n) = 1] = \Pr[\mathcal{D}^{|\overline{\mathcal{P}}_R\rangle}(1^n) = 1] \quad (69)$$

$$\Pr[\mathcal{D}^{|\tilde{\mathcal{S}}\rangle_{\S}}(1^n) = 1] = \Pr[\mathcal{D}^{|\tilde{\mathcal{S}}_R\rangle}(1^n) = 1] \quad (70)$$

where the left hand side concerns *distributional* quantum oracle access (see Equation (13)) and the right hand side refers to normal quantum oracle access (see Equation (12)). This implies that

$$\Pr[\mathcal{D}^{|\overline{\mathcal{P}}\rangle_{\S}}(1^n) = 1] - \Pr[\mathcal{D}^{|\tilde{\mathcal{S}}\rangle_{\S}}(1^n) = 1] > \frac{1}{\rho(n)}. \quad (71)$$

□

By the above argument, it is sufficient to only consider \mathcal{D} to have distributional oracle access to its oracle. Assume towards contradiction that

$$|\Pr[\mathcal{D}^{|\overline{\mathcal{P}}\rangle_{\S}}(1^n) = 1] - \Pr[\mathcal{D}^{|\tilde{\mathcal{S}}(1^n)\rangle_{\S}}(1^n) = 1]| > \frac{1}{\rho(n)}. \quad (72)$$

We now define the first series of hybrids. Recall we set a threshold $m^*(n) = \lfloor \log^{1/(d+1)} n \rfloor$ and that since \mathcal{D} is a QPT algorithm, there exists a max security parameter $m_{\max}(n) \leq \text{poly}(n)$. We first define a function $r'(\cdot)$:

$$r'(j) = r \cdot 2^{j^d} \cdot m_{\max}(n).$$

where r will be a value dependent on the distinguishing advantage $\rho(n)$ and the number of queries. Looking ahead, $r'(j)$ will be used to determine the number of samples needed to sample from the output distribution of $\overline{\mathcal{P}}$ run with security parameter $j \in \{1, \dots, m_{\max}(n)\}$ to generate the analagous small-range distribution for $\overline{\mathcal{P}}, \text{SR}_{r'}^{\overline{\mathcal{P}}(1^m, \text{crs})}(\mathcal{X}_m)$.

- $\tilde{\mathcal{H}}_0$: In this hybrid, the distinguisher \mathcal{D} makes at most $q(n)$ queries to $U_{\overline{\mathcal{P}}}^{\mathbb{S}}$. Recall \mathcal{D} makes quantum queries on the security parameter, 1^m (recall $m = \text{HW}(\nu)$ for $\nu \in \{0, 1\}^{m_{\max}(n)}$), the common reference string, crs , and on the random coins, coins .
- $\tilde{\mathcal{H}}_1$: Let $r = 2\ell(q) \cdot \rho(n)$. We then define $\tilde{\mathcal{H}}_1$ to be the same as $\tilde{\mathcal{H}}_0$ except we first sample $r'(m_{\max}(n))$ many samples from the output distribution of $\overline{\mathcal{P}}$. We use these $r'(m_{\max}(n))$ many samples as non-uniform advice to first generate a small range distribution $\text{SR}_{r'}^{\overline{\mathcal{P}}(1^m, \text{crs})}(\mathcal{X}_m)$ for every security parameter $m \in \{1, \dots, m_{\max}(n)\}$ and every $\text{crs} \in \{0, 1\}^{m^d}$ for that particular security parameter m . Let U_f denote the oracle we construct using $\text{SR}_{r'}^{\overline{\mathcal{P}}(1^m, \text{crs})}(\mathcal{X}_m)$ for every m and crs and is defined via Algorithm 4.1.

Algorithm 1 Constructing small range oracle U_f

for Every security parameter $m \in \{1, \dots, m_{\max}(n)\}$ and all $\text{crs} \in \{0, 1\}^{m^d}$ **do**

1. Sample $\{(\bar{x}, \bar{\pi})_i^{m, \text{crs}}\}_{i \in [r']}$ $\leftarrow \overline{\mathcal{P}}(1^m, \text{crs})$

2. Sample $f_{m, \text{crs}} \leftarrow \text{SR}_{r'}^{\overline{\mathcal{P}}(1^m, \text{crs})}(\mathcal{X}_m)$

Construct U_f :

$$|1^m\rangle |\text{crs}\rangle |\text{coins}\rangle |y\rangle \mapsto |1^m\rangle |\text{crs}\rangle |\text{coins}\rangle |y \oplus f_{m, \text{crs}}(\text{coins})\rangle \quad (73)$$

end for

Then \mathcal{D} makes $q(n)$ queries U_f instead of $U_{\overline{\mathcal{P}}}^{\mathbb{S}}$.

To argue indistinguishability of hybrids $\tilde{\mathcal{H}}_0$ and $\tilde{\mathcal{H}}_1$ we apply Lemma 15. Recall that Lemma 15 holds against adversaries of unbounded size. We apply Lemma 15 with r samples for each security parameter $j \in \{1, \dots, m_{\max}(n)\}$ and associated $\text{crs} \in \{0, 1\}^{j^d}$ which tells us that the distinguishing advantage of the output distributions is

$$\sum_{j=1}^{m_{\max}(n)} \frac{2^{j^d} \cdot \ell(q)}{r'(j)} \leq \frac{2^{m_{\max}(n)^d} \cdot \ell(q) \cdot m_{\max}(n)}{r'(m_{\max}(n))} \quad (74)$$

$$= \frac{2^{m_{\max}(n)^d} \cdot \ell(q) \cdot m_{\max}(n)}{r \cdot 2^{m_{\max}(n)^d} \cdot m_{\max}(n)} \quad (75)$$

$$= \frac{1}{2\rho(n)}. \quad (76)$$

where we used the fact that since the security parameter must be queried in unary. It follows that $\tilde{\mathcal{H}}_0$ and $\tilde{\mathcal{H}}_1$ are $(\varepsilon_0(n), q(n))$ -oracle-indistinguishable where

$$\varepsilon_0(n) = \frac{1}{2\rho(n)} > n^{-O(1)}. \quad (77)$$

Then, it follows from the assumption that

$$\left| \Pr_{\mathcal{O} \leftarrow \tilde{\mathcal{H}}_1} [\mathcal{D}^{|\mathcal{O}|_s}(1^n) = 1] - \Pr_{\mathcal{O} \leftarrow \tilde{\mathcal{S}}} [\mathcal{D}^{|\mathcal{O}|_s}(1^n) = 1] \right| \geq \frac{1}{2\rho(n)}. \quad (78)$$

We now define the second sequence of hybrids over all security parameters. Our goal is to show that the overall distinguishing advantage of \mathcal{D} is at least as large as $\tilde{\varepsilon}(n) = \frac{1}{2\rho(n)}$. We define the hybrid $\mathcal{H}_{m_{\max}(n)}$ to be the experiment where \mathcal{D} queries an oracle generated from $\text{SR}_{r'(j)}^{\overline{\mathcal{P}}(1^m, \text{crs})}(\mathcal{X}_m)$ and \mathcal{H}_1 to be the experiment where \mathcal{D} queries $\tilde{\mathcal{S}}(1^n)$. More formally, we aim to show indistinguishability of the extreme hybrids:

$$\left| \Pr_{\mathcal{O} \leftarrow \tilde{\mathcal{H}}_1} [\mathcal{D}^{|\mathcal{O}|_s}(1^n) = 1] - \Pr_{\mathcal{O} \leftarrow \tilde{\mathcal{H}}_1} [\mathcal{D}^{|\mathcal{O}|_s}(1^n) = 1] \right| \geq \tilde{\varepsilon}(n). \quad (79)$$

We define \mathcal{H}_j to be the experiment that executes the following:

1. Recall $m = \text{HW}(s)$. Construct a distribution tuple $\mathbf{R} = \{R_{m, \text{crs}}\}$ for all $m \in \{1, \dots, m_{\max}(n)\}$, $\text{crs} \in \{0, 1\}^{m^d}$. Let $\varepsilon(j)$ be the distinguishing probability of \mathcal{H}_j and \mathcal{H}_{j-1} . Each distribution $R_{m, \text{crs}}$ is generated from either $\text{SR}_{r'(j)}^{\overline{\mathcal{P}}(1^m, \text{crs})}(\mathcal{X}_m)$ or $\mathcal{S}(1^n)$ where $r = \frac{2\ell(q)}{\varepsilon(j)}$. The threshold $m^*(n)$ and the j th security parameter determines the choice of oracle $R_{m, \text{crs}}$. The specific details of the construction of \mathbf{R} are described in Algorithm 2.
2. Run \mathcal{B} with an oracle $\mathcal{O}_j \leftarrow \mathbf{R}^{\mathcal{X}}$. The oracle is defined as:

$$\mathcal{O}_j = \sum_{\nu \in \{0, 1\}^{m_{\max}(n)}: m^*(n) < |\nu| \leq j} |\nu\rangle\langle\nu| \otimes U_f \otimes \mathbb{I} \quad (80)$$

$$+ \sum_{\nu \in \{0, 1\}^{m_{\max}(n)}: |\nu| > j} |\nu\rangle\langle\nu| \otimes U_{\mathcal{S}, m} \otimes \mathbb{I}. \quad (81)$$

We now argue indistinguishability of each hybrid \mathcal{H}_j . Formally,

$$\tilde{\varepsilon}(n) < \left| \Pr_{\mathcal{O} \leftarrow \tilde{\mathcal{H}}_{m_{\max}(n)}} [\mathcal{D}^{|\mathcal{O}|_s}(1^n) = 1] - \Pr_{\mathcal{O} \leftarrow \tilde{\mathcal{H}}_1} [\mathcal{D}^{|\mathcal{O}|_s}(1^n) = 1] \right| \quad (82)$$

$$\leq \sum_{j=2}^{m_{\max}(n)} \left| \Pr_{\mathcal{O} \leftarrow \tilde{\mathcal{H}}_j} [\mathcal{D}^{|\mathcal{O}|_s}(1^n) = 1] - \Pr_{\mathcal{O} \leftarrow \tilde{\mathcal{H}}_{j-1}} [\mathcal{D}^{|\mathcal{O}|_s}(1^n) = 1] \right| \quad (83)$$

Algorithm 2 Constructing \mathbf{R} in hybrid \mathcal{H}_j

for $j \in \{m_{\max}(n), \dots, m^*(n)\}, \text{crs} \in \{0, 1\}^{m^d}$ **do** Generate every $f_{m, \text{crs}} \leftarrow \text{SR}_{r(j)}^{\overline{\mathcal{P}}}(\mathcal{X}_j)$ according to Steps 1 and 2 of Algorithm 4.1 and hardcode every function as non-uniform advice.

if $m^*(n) \leq m < j$ **then**

 Set $R_{m, \text{crs}} = U_f$ (According to Step 3 of Algorithm 4.1).

else

 Set $R_{m, \text{crs}} = U_{\mathcal{S}, m}$ for every query made with security parameter m .

end if

end for

$$\leq \sum_{j=m^*(n)+2}^{m_{\max}(n)} \left| \Pr_{\mathcal{O} \leftarrow \tilde{\mathcal{H}}_j} [\mathcal{D}^{\mathcal{O}}_{\mathcal{S}}(1^n) = 1] - \Pr_{\mathcal{O} \leftarrow \tilde{\mathcal{H}}_{j-1}} [\mathcal{D}^{\mathcal{O}}_{\mathcal{S}}(1^n) = 1] \right| \quad (84)$$

$$\leq (m_{\max}(n) - m^*(n) + 1)(\varepsilon(j^*)). \quad (85)$$

where recall we set $\varepsilon(j)$ to be the distinguishing probability of \mathcal{H}_j and \mathcal{H}_{j-1} and $j^* = \arg \max_j \varepsilon(j)$. Note the third inequality follows from the fact that if $m \leq m^*(n)$ then the hybrids are exactly equivalent.

We now consider $\tilde{\mathcal{H}}_{j^*-1}$, which is a variant of the hybrid \mathcal{H}_{j^*-1} . We define $\tilde{\mathcal{H}}_{j^*-1}$ to be exactly the same as \mathcal{H}_{j^*-1} , except in $\tilde{\mathcal{H}}_{j^*-1}$ we set the number of samples used to be $r'(j^*)$ instead of $r'(j^* - 1)$. We now apply Corollary 16, with $r_1 = r'(j^*)$ and $r_2 = r'(j^* - 1)$ where $r \geq 2\ell(q) \cdot \frac{1}{\varepsilon(j^*)}$ for both functions. Note that the size of the distribution tuple is $k = (j^* - 1) \cdot 2^{j^*-1} \leq m_{\max}(n) \cdot 2^{j^*-1}$. We obtain that

$$\left| \Pr_{\mathcal{O} \leftarrow \tilde{\mathcal{H}}_{j^*-1}} [\mathcal{D}^{\mathcal{O}}_{\mathcal{S}}(1^n) = 1] - \Pr_{\mathcal{O} \leftarrow \mathcal{H}_{j^*-1}} [\mathcal{D}^{\mathcal{O}}_{\mathcal{S}}(1^n) = 1] \right| \quad (86)$$

$$\leq k\ell(q) \left(\frac{1}{r(j^* - 1)} - \frac{1}{r'(j^*)} \right) \quad (87)$$

$$\leq \left(\frac{m_{\max}(n) \cdot 2^{(j^*-1)^d} \cdot \ell(q) \cdot \varepsilon(j^*)}{2 \cdot \ell(q) \cdot m_{\max}(n)} \right) \left(\frac{1}{2^{(j^*-1)^d}} - \frac{1}{2^{(j^*)^d}} \right) \quad (88)$$

$$\leq \frac{\varepsilon(j^*)}{4}. \quad (89)$$

and

$$\left| \Pr_{\mathcal{O} \leftarrow \tilde{\mathcal{H}}_{j^*}} [\mathcal{D}^{\mathcal{O}}_{\mathcal{S}}(1^n) = 1] - \Pr_{\mathcal{O} \leftarrow \tilde{\mathcal{H}}_{j^*-1}} [\mathcal{D}^{\mathcal{O}}_{\mathcal{S}}(1^n) = 1] \right| \geq \tilde{\varepsilon}(n) - \frac{\varepsilon(j^*)}{4}. \quad (90)$$

We now argue that given \mathcal{D} we can construct a non-uniform distinguisher $\tilde{\mathcal{D}}$ of size $s^*(n)$ that oracle-distinguishes $\overline{\mathcal{P}}$ and $\tilde{\mathcal{S}}$ with probability greater than $\varepsilon^*(n)$ where $s^*(n)$ and $\varepsilon^*(n)$ are as specified by Theorem 17. Specifically, $\tilde{\mathcal{D}}$

uses \mathcal{D} and its own oracle $\tilde{\mathcal{O}}$ to distinguish distribution families $\text{SR}_{r'}^{\overline{\mathcal{P}}(1^m, \text{crs})}(\mathcal{X}_{j^*})$ and $\mathcal{S}(1^n)$ for security parameter j^* .

We first argue that $\tilde{\mathcal{D}}$ can simulate each of the above hybrids \mathcal{H}_j by simulating the associated oracle, \mathcal{O}_j . Recall that the optimal distinguishing probability of \mathcal{D} was denoted by security parameter j^* . We first generate the non-uniform advice, which contains at most $r'(j^*) = r \cdot m_{\max}(n) \cdot 2^{(j^*)^d}$ samples required to simulate U_f for every $m \in [j^*]$ and every $\text{crs} \in \{0, 1\}^{(j^*)^d}$. For security parameters $m > j$ in which $\tilde{\mathcal{D}}$ is required to simulate $U_{\mathcal{S}, m}$, $\tilde{\mathcal{D}}$ generates samples on the fly by running \mathcal{S} controlled on the registers which contain the security parameter m , crs , and coins respectively.

Note that there are two cases of interest (in the third case, i.e., $j^* < m^*(n)$, the hybrids are identical). When \mathcal{D} makes a query (in superposition) with security parameter $m^*(n) \leq m < j^*$ and is given oracle access to $\mathcal{O}_{j^*} = U_f$, then $\tilde{\mathcal{D}}$ simulates the hybrid \mathcal{H}_{j^*} . Alternatively, if \mathcal{D} makes a query (in superposition) with security parameter $m^*(n) \leq m < j^*$ and is given oracle access to $\mathcal{O}_{j^*} = U_{\mathcal{S}, m}$, then $\tilde{\mathcal{D}}$ simulates hybrid \mathcal{H}_{j^*-1} . Therefore it follows that

$$\left| \Pr_{\mathcal{O} \leftarrow \overline{\mathcal{P}}} [\tilde{\mathcal{D}}^{|\mathcal{O}\rangle_{\mathcal{S}}}(1^n) = 1] - \Pr_{\mathcal{O} \leftarrow \tilde{\mathcal{S}}(1^n)} [\tilde{\mathcal{D}}^{|\mathcal{O}\rangle_{\mathcal{S}}}(1^n) = 1] \right| \quad (91)$$

$$\geq \left| \Pr_{\mathcal{O} \leftarrow \mathcal{H}_{j^*}} [\mathcal{D}^{|\mathcal{O}\rangle_{\mathcal{S}}}(1^n) = 1] - \Pr_{\mathcal{O} \leftarrow \mathcal{H}_{j^*-1}} [\mathcal{D}^{|\mathcal{O}\rangle_{\mathcal{S}}}(1^n) = 1] \right| \quad (92)$$

$$\geq \tilde{\varepsilon}(n) - \frac{\varepsilon(j^*)}{4} \quad (93)$$

$$\geq \tilde{\varepsilon}(n) - \frac{\tilde{\varepsilon}(n)}{4(m_{\max}(n) - m^*(n) + 1)} \quad (94)$$

$$= \Omega(n^{-O(1)}). \quad (95)$$

in the second inequality we used that $m_{\max}(n) \leq p(n)$ (recall $p(n)$ is the polynomial defined by the maximum runtime of \mathcal{D}) and we used that $\tilde{\varepsilon}(n) \geq \rho(n)$.

We will now analyze the size of the circuit $\tilde{\mathcal{D}}$. We claim that we can simulate each hybrid \mathcal{H}_j for $j \in \{1, \dots, m_{\max}(n)\}$, given non-uniform advice, in size at most $s(j) = 2^{O(j^{d+2})}$. We can consider Algorithm 2 to be the circuit that is constructed using controlled oracle access to the appropriate unitary. That is, controlled on the register containing a security parameter m , the oracle generates an associated crs and then responds using either U_f or $U_{\mathcal{S}, m}$. As described earlier, we generated our non-uniform advice by sampling r -many samples for each $\text{crs} \in \{0, 1\}^{o(j^d)}$ and each $j \in [m_{\max}(n)]$. Note that the total number of hard-coded samples needed to simulate U_f for the j th security parameter is at most $r'(j) = 2^{j^d} \cdot r \cdot j \ll 2^{O(j^{d+2})}$ (where r is defined according to the small-range distribution).

Let $\tilde{\mathcal{D}}_j(1^n)$ denote the distinguishing circuit used to simulate hybrid \mathcal{H}_j . Then the final circuit size of $\tilde{\mathcal{D}}$ consists of $|\mathcal{D}|$, the non-uniform advice for each hybrid,

and the circuit required to implement the $2q$ -wise independent hash function.

$$|\tilde{\mathcal{D}}_j(1^n)| \leq |\mathcal{D}(1^n)| + \frac{106 \cdot q^3 \cdot 2^{3j^d} \cdot j}{\varepsilon(j)} + \text{poly}(q(n), \log \mathcal{X}, \log \mathcal{Y}) \leq 2^{O(j^d)} \quad (96)$$

By using Theorem 17 we obtain that if $\bar{\mathcal{P}}$ and $\tilde{\mathcal{S}}$ are $(s'(m), \varepsilon'(m))$ -sample-indistinguishable, then they are also $(s(m), q(n), \varepsilon(m))$ -oracle-indistinguishable (up to the appropriate constraints). Then, by choosing $q(n) \leq p(n)$ and $\varepsilon^*(n) = 2^{-\Omega(n^{d+2})}$, it holds that for all $j \geq m^*(n)$, the probability $\varepsilon(j) \leq \varepsilon(m^*(n))$. Theorem 17 then yields

$$\varepsilon'(m^*(n)) \leq \frac{\varepsilon(m^*(n))^2}{864 \cdot q(n)^3 \cdot 2^{2m^*(n)^d}} \leq n^{-\omega(1)} \quad (97)$$

and, since the associated circuit size for any $j \geq m^*(n)$ is $s^*(j) \geq s(m^*(n))$,

$$s'(m^*(n)) \geq s(m^*(n)) + \frac{106 \cdot q(n)^3 \cdot 2^{3m^*(n)^d}}{\varepsilon^*(m^*(n))} \geq n^{\omega(1)}, \quad (98)$$

where we chose $s(n) = 2^{O((n)^{d+2})}$. However, in the final hybrid by Equation (96) we have constructed a distinguisher $\tilde{\mathcal{D}}_{m^*(n)}(1^n)$ of size

$$|\tilde{\mathcal{D}}_{m^*(n)}(1^n)| \leq 2^{O(m^*(n)^d)} = n^{O(1)} < s(n)$$

that has distinguishing advantage at least $n^{-O(1)} > \varepsilon(m^*(n))$. This contradicts Theorem 17 for all $j \geq m^*(n)$ and the lemma follows. \square

4.2 Main result

We are now ready to state and prove our main result about quantum black-box reductions for SNARGs. The following is a rigorous version of Theorem 1.

Theorem 19 (Impossibility of Reductions to SNARGs). *Let L be a language in NP with a subexponentially-quantum-hard subset membership problem. Let $\Pi = (\mathsf{G}, \mathsf{P}, \mathsf{V})$ be a SNARG for the language L that satisfies the completeness and succinctness properties as in Definition 8. Then, for any falsifiable assumption (C, c) , one of the following must hold:*

- A. *The assumption (C, c) is false.*
- B. *There is no quantum black-box reduction showing the soundness of Π based on the assumption (C, c) .*

The same conclusion holds if we assume that \mathcal{L} has an exponentially quantum-hard subset membership problem and we allow Π to only be slightly succinct.

Proof. As Π satisfies the conditions of Lemma 18, there exists an (inefficient) adversary $\bar{\mathcal{P}}$ that outputs no instances and “fake” proofs that are accepted by $\Pi.V$ with all but negligible probability.

Assume the negation of B, i.e., that there exists a quantum black-box reduction Σ that establishes the soundness of Π based on the falsifiable assumption (\mathbf{C}, c) . It follows that there exists some polynomial p such that

$$\Pr[\langle \mathbf{C}, \Sigma^{|\bar{\mathcal{P}}}\rangle(1^n) = \mathbf{win}] \geq c + 1/p(n).$$

By the second part of Lemma 18, there also exists an efficient algorithm $\tilde{\mathcal{S}}$ such that $\bar{\mathcal{P}}$ and $\tilde{\mathcal{S}}(1^n)$ are oracle-indistinguishable to QPT algorithms. As $(\mathbf{C}, \Sigma^{(\cdot)})$ is QPT, it follows that

$$\Pr[\langle \mathbf{C}, \Sigma^{|\tilde{\mathcal{S}}(1^n)}\rangle(1^n) = \mathbf{win}] \geq c + 1/p(n) - \text{negl}(n).$$

Thus, the QPT $\Sigma^{|\tilde{\mathcal{S}}(1^n)}$ breaks (\mathbf{C}, c) , meaning that A is true.

To prove the version of the theorem for exponentially quantum-hard subset membership problems and slightly succinct Π , we utilize the same argument as [GW11] using the parameters set for $s(n)$ and $\varepsilon(n)$ in Lemma 18. That is, recall that for proofs to satisfy for the slightly succinct property, the length of the proof for the exponentially quantum-hard subset membership problem is $O(n^b(|x| + |w|)^c) + o(|x| + |w|)$ where we set constants $c < 1$ and b . Then there is a choice of a constant d such that $d+3 > b/1-c$. This ensures that there exists an exponentially quantum-hard subset membership problem $(\mathcal{L}_n, \tilde{\mathcal{L}}_n, \text{Samp})$ where $s(n) = n^{d+3}$, $\varepsilon(n) = n^{-(d+3)}$, and $\ell_{\text{pf}} = o(n^{d+3})$. \square

4.3 Impossibility of Reductions from δ -exponentially Hard Falsifiable Assumptions

Similarly to [GW11], Theorem 19 also extends to reductions from δ -exponentially hard falsifiable assumptions. This is formalized in the following theorem.

Theorem 20. *Let L be a language in NP with a subexponentially-quantum-hard subset membership problem. Let $\Pi = (\mathbf{G}, \mathbf{P}, \mathbf{V})$ be a non-interactive proof system for the language L that satisfies the completeness and succinctness properties as in Definition 8. Then, for any δ -exponential version of a falsifiable assumption (\mathbf{C}, c) , one of the following must hold:*

- A. *The δ -exponential version of the assumption (\mathbf{C}, c) is false.*
- B. *There is no quantum black-box reduction showing the soundness of Π based on the δ -exponential version of assumption (\mathbf{C}, c) .*

Modifications to Lemma 18 The necessary changes to the proof of Lemma 18 follow similarly to those stated in [GW11, Appendix A]. The primary change is that for the case of δ -exponentially hard assumptions, the distinguisher \mathcal{D} is now allowed to run in time $2^{O(n^\delta)}$. Therefore, we now want to argue that for every \mathcal{D} that runs in time $2^{O(n^\delta)}$, there exists a SNARG adversary $\bar{\mathcal{P}}$ that can be simulated by a simulator \mathcal{S} of size $2^{O(n^\delta)}$ such that

$$\Pr[\mathcal{D}^{|\bar{\mathcal{P}}}\rangle(1^n) = 1] - \Pr[\mathcal{D}^{|\mathcal{S}}\rangle(1^n) = 1] < \frac{1}{2^{\omega(n^\delta)}}.$$

We will utilize the same proof, except with a modification to the choice of d which defines the hardness of the subset membership problem. Recall in the proof that by the succinctness of Π , it holds that there exists a sufficiently large constant d which bounds the length of the $\mathbf{crs} \leftarrow \Pi.\mathbf{G}(1^n)$. We define a constant $d' = \max(\delta, d)$ and then choose $s(n) = 2^{cn^{d'+3}}$ and $\varepsilon(n) = 2^{-n^{d'+3}}$. We also change the threshold of the non-security parameter preserving proof to be $m^*(n) = n^{\frac{\delta}{d'+1}}$. The resulting table \mathcal{T}_n is of size

$$|\mathcal{T}_n| = \sum_{m=1}^{m^*(n)} r(n) \cdot 2^{m^{d'}} (\ell_{pf}(m) + \ell_{st}(m)) \quad (99)$$

$$\leq \frac{2^{m^*(n)^{d'}} \ell(q(n))}{\rho(n)} \cdot m^*(n) (o(m^*(n)^{d'+3}) + \text{poly}(m^*(n))) \quad (100)$$

$$= O\left(2^{n^{\frac{\delta \cdot d'}{d'+1}}}\right) \quad (101)$$

$$= O(2^{n^\delta}) \quad (102)$$

We use the same hybrid argument over all security parameters $\{2^{n^\delta}, \dots, 1\}$ in decreasing order. Then finally, we obtain by Theorem 17

$$s'(m^*(n)) \geq s^*(m^*(n)) + \frac{106 \cdot q(n)^3 \cdot 2^{2n^{\frac{d' \cdot \delta}{d'+1}}}}{\varepsilon^*(m^*(n))} = 2^{O(n^{\delta \cdot \frac{d'+2}{d'+1}})} = 2^{\omega(n^\delta)} \quad (103)$$

and

$$\varepsilon'(m^*(n)) \leq \frac{\varepsilon^*(m^*(n))^2}{864 \cdot q(n)^3 \cdot 2^{2n^{\frac{d' \cdot \delta}{d'+1}}}} = 2^{O(-n^{\delta \cdot \frac{d'+2}{d'+1}})} \leq 2^{\omega(-n^\delta)} \quad (104)$$

Proof (of Theorem 20). The proof follows exactly like that of Theorem 19 except we consider quantum black-box reductions Σ of size $2^{O(n^\delta)}$. That is, we assume the existence of a quantum reduction Σ that establishes the soundness of the SNARG Π based on a δ -exponentially hard falsifiable assumption (\mathbf{C}, c) . It then follows that there exists a function $\nu(n) > \frac{1}{2^{-n^\delta}}$

$$\Pr[\langle \mathbf{C}, \Sigma^{\bar{\rho}} \rangle(1^n) = \mathbf{win}] \geq c + \frac{1}{\nu(n)}.$$

However, by Lemma 18, it holds that

$$\Pr[\langle \mathbf{C}, \Sigma^{\mathcal{S}} \rangle(1^n) = \mathbf{win}] \geq c + 1/\nu(n) - \frac{1}{2^{\omega(n^\delta)}}.$$

This contradicts the hardness of the δ -exponentially hard falsifiable assumption.

References

- ACC⁺22. Per Austrin, Hao Chung, Kai-Min Chung, Shiuan Fu, Yao-Ting Lin, and Mohammad Mahmoody. On the impossibility of key agreements from quantum random oracles. In *CRYPTO 2022, Part II*, 2022.

- AHY23. Prabhanjan Ananth, Zihan Hu, and Henry Yuen. On the (im)plausibility of public-key quantum money from collision-resistant hash functions. In *ASIACRYPT 2023, Part VIII*, 2023.
- Bar01. Boaz Barak. How to go beyond the black-box simulation barrier. In *42nd FOCS*, 2001.
- BDG⁺13. Nir Bitansky, Dana Dachman-Soled, Sanjam Garg, Abhishek Jain, Yael Tautman Kalai, Adriana López-Alt, and Daniel Wichs. Why “Fiat-Shamir for proofs” lacks a proof. In *TCC 2013*, 2013.
- BISW17. Dan Boneh, Yuval Ishai, Amit Sahai, and David J. Wu. Lattice-based SNARGs and their application to more efficient obfuscation. In *EUROCRYPT 2017, Part III*, 2017.
- BJLS16. Christoph Bader, Tibor Jäger, Yong Li, and Sven Schäge. On the impossibility of tight cryptographic reductions. In *Advances in Cryptology—EUROCRYPT 2016: 35th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Vienna, Austria, May 8-12, 2016, Proceedings, Part II 35*, pages 273–304. Springer, 2016.
- BKP19. Nir Bitansky, Dakshita Khurana, and Omer Paneth. Weak zero-knowledge beyond the black-box barrier. In *51st ACM STOC*, 2019.
- BLP⁺13. Zvika Brakerski, Adeline Langlois, Chris Peikert, Oded Regev, and Damien Stehlé. Classical hardness of learning with errors. In *45th ACM STOC*, 2013.
- BLV03. Boaz Barak, Yehuda Lindell, and Salil P. Vadhan. Lower bounds for non-black-box zero knowledge. In *44th FOCS*, 2003.
- BP12. Nir Bitansky and Omer Paneth. From the impossibility of obfuscation to a new non-black-box simulation technique. In *53rd FOCS*, 2012.
- BR93. Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In *ACM CCS 93*, 1993.
- Bro16. Daniel RL Brown. Breaking rsa may be as difficult as factoring. *Journal of Cryptology*, 29(1):220–241, 2016.
- BV98. Dan Boneh and Ramarathnam Venkatesan. Breaking RSA may not be equivalent to factoring. In *EUROCRYPT’98*, 1998.
- CCL⁺17. Yi-Hsiu Chen, Kai-Min Chung, Ching-Yi Lai, Salil P Vadhan, and Xiaodi Wu. Computational notions of quantum min-entropy. *arXiv preprint arXiv:1704.07309*, 2017.
- CGKS23. Matteo Campanelli, Chaya Ganesh, Hamidreza Khoshakhlagh, and Janno Siim. Impossibilities in succinct arguments: Black-box extraction and more. In *AFRICACRYPT 23*, 2023.
- Chu23. Kai-Min Chung. Private communication, 2023.
- CLMP13. Kai-Min Chung, Huijia Lin, Mohammad Mahmoody, and Rafael Pass. On the power of nonuniformity in proofs of security. In *ITCS ’13*, page 389–400, New York, NY, USA, 2013. Association for Computing Machinery.
- CLZ22. Yilei Chen, Qipeng Liu, and Mark Zhandry. Quantum algorithms for variants of average-case lattice problems via filtering. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 372–401. Springer, 2022.
- CM24. Andrea Coladangelo and Saachi Mutreja. On black-box separations of quantum digital signatures from pseudorandom states, 2024.
- CX22. Shujiao Cao and Rui Xue. The gap is sensitive to size of preimages: Collapsing property doesn’t go beyond quantum collision-resistance for preimages bounded hash functions. In *Annual International Cryptology Conference*, pages 564–595. Springer, 2022.

- DLS22. Frédéric Dupuis, Philippe Lamontagne, and Louis Salvail. Fiat-shamir for proofs lacks a proof even in the presence of shared entanglement. *arXiv preprint arXiv:2204.02265*, 2022.
- GBL08. Sanjam Garg, Raghav Bhaskar, and Satyanarayana V. Lokam. Improved bounds on security reductions for discrete log based signatures. In *Advances in Cryptology - CRYPTO 2008, 28th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 17-21, 2008*, volume 5157 of *Lecture Notes in Computer Science*, pages 93–107. Springer, 2008.
- GW11. Craig Gentry and Daniel Wichs. Separating succinct non-interactive arguments from all falsifiable assumptions. In *43rd ACM STOC*, 2011.
- HH09. Iftach Haitner and Thomas Holenstein. On the (im)possibility of key dependent encryption. In *TCC 2009*, 2009.
- IR89. Russell Impagliazzo and Steven Rudich. Limits on the provable consequences of one-way permutations. In *Proceedings of the twenty-first annual ACM symposium on Theory of computing*, pages 44–61, 1989.
- KLX20. Jonathan Katz, Julian Loss, and Jiayu Xu. On the security of time-lock puzzles and timed commitments. In *Theory of Cryptography: 18th International Conference, TCC 2020, Durham, NC, USA, November 16–19, 2020, Proceedings, Part III 18*, pages 390–413. Springer, 2020.
- MPV24. Surya Mathialagan, Spencer Peters, and Vinod Vaikuntanathan. Adaptively sound zero-knowledge SNARKs for UP. *Cryptology ePrint Archive*, Paper 2024/227, 2024. <https://eprint.iacr.org/2024/227>.
- Nao03. Moni Naor. On cryptographic assumptions and challenges. In *Annual International Cryptology Conference*, pages 96–109. Springer, 2003.
- Pas11. Rafael Pass. Limits of provable security from standard assumptions. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 109–118, 2011.
- Pas13. Rafael Pass. Unprovable security of perfect NIZK and non-interactive non-malleable commitments. In *TCC 2013*, 2013.
- PV05. Pascal Paillier and Damien Vergnaud. Discrete-log-based signatures may not be equivalent to discrete log. In *Advances in Cryptology - ASIACRYPT 2005, 11th International Conference on the Theory and Application of Cryptology and Information Security, Chennai, India, December 4-8, 2005, Proceedings*, volume 3788 of *Lecture Notes in Computer Science*, pages 1–20. Springer, 2005.
- Reg09. Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *Journal of the ACM (JACM)*, 56(6):1–40, 2009.
- SCG⁺14. Eli Ben Sasson, Alessandro Chiesa, Christina Garman, Matthew Green, Ian Miers, Eran Tromer, and Madars Virza. Zerocash: Decentralized anonymous payments from bitcoin. In *2014 IEEE symposium on security and privacy*, pages 459–474. IEEE, 2014.
- SSTX09. Damien Stehlé, Ron Steinfeld, Keisuke Tanaka, and Keita Xagawa. Efficient public key encryption based on ideal lattices. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 617–635. Springer, 2009.
- SW14. Amit Sahai and Brent Waters. How to use indistinguishability obfuscation: deniable encryption, and more. In *46th ACM STOC*, 2014.
- Wic13. Daniel Wichs. Barriers in cryptography with weak, correlated and leaky sources. In *ITCS 2013*, 2013.

- WW24a. Brent Waters and David J. Wu. Adaptively-sound succinct arguments for NP from indistinguishability obfuscation. Cryptology ePrint Archive, Paper 2024/165, 2024. <https://eprint.iacr.org/2024/165>.
- WW24b. Brent Waters and David J. Wu. A pure indistinguishability obfuscation approach to adaptively-sound SNARGs for NP. Cryptology ePrint Archive, Paper 2024/933, 2024. <https://eprint.iacr.org/2024/933>.
- WZ24. Brent Waters and Mark Zhandry. Adaptive security in SNARGs via iO and lossy functions. Cryptology ePrint Archive, Paper 2024/254, 2024. <https://eprint.iacr.org/2024/254>.
- Zha12a. Mark Zhandry. How to construct quantum random functions. In *53rd FOCS*, 2012.
- Zha12b. Mark Zhandry. Secure identity-based encryption in the quantum random oracle model. In *CRYPTO 2012*, 2012.

A Implementing Controlled Unitary Access

Claim. Let \mathcal{X} and \mathcal{Y} be sets, and let $f : \mathcal{X} \rightarrow \mathcal{Y}$. Let U_f be a unitary operator that implements quantum oracle access to f as defined in 2. Then q -quantum query access to U_f implies $q - 1$ -quantum query access to controlled U_f , CU_f , which acts as follows:

$$CU_f : |b\rangle_B |x\rangle_X |y\rangle_Y \rightarrow |b\rangle_B |x\rangle_X |y \oplus b \cdot f(x)\rangle_Y \quad (105)$$

where $b \in \{0, 1\}$ and B, X, Y represent the control, input and output registers respectively.

Proof. First, before constructing CU_f , make a single classical query to f to learn $f(0)$. Then to construct CU_f :

1. Prepare two ancilla registers $|0\rangle_{X'} |f(0)\rangle_{Y'}$.
2. Flip the bit b provided as input by computing $b \oplus 1$ in register B .
3. Apply $CSWAP_{B, X, X'}$ where B acts as the control register.
4. Query U_f on registers X and Y .
5. Apply a $CSWAP_{B, X, X'}$ where B acts as the control register.
6. Apply a $CNOT_{B, Y, Y'}$ where B acts as the control register.
7. Uncompute the bit flip $b \oplus 1$ in register B to return to b .

If $b = 0$, then the resulting state maps

$$|0\rangle_B |x\rangle_X |y\rangle_Y \rightarrow |0\rangle_B |x\rangle_X |y\rangle_Y \quad (106)$$

and if $b = 1$, then the resulting state maps

$$|1\rangle_B |x\rangle_X |y\rangle_Y \rightarrow |1\rangle_B |x\rangle_X |y \oplus f(x)\rangle_Y \quad (107)$$

which satisfies the equation in the claim. Note that the two instances in which queries were made to f were to learn $f(0)$ (via a single classical global query) and in step 4 of the algorithm to construct CU_f .

B Simulation of the Oracle in Theorem 17

Claim. \mathcal{A}_{i^*} simulates V defined in Equation (40).

- Proof.* 1. If register I contains the index $j = i^*$, we generate the responses from the single index oracle of \mathcal{A}_{i^*} , $U_{f_{i^*}}$. That is, controlled on the index register I , if $j = i^*$, then output a response via $U_{f_{i^*}}$ in output register Y .
2. \mathcal{A}_{i^*} simulates U_{f_j} for $j \neq i^*$:
- In this case, \mathcal{A}_{i^*} first constructs a small-range distribution $\text{SR}_r^{\mathcal{D}^{b,j}}(\mathcal{X})$ using the table of samples provided in the non-uniform advice for $b \in \{0, 1\}$ and $j \in [k] \setminus \{i^*\}$. If $j < i^*$, then \mathcal{A}_{i^*} constructs $\text{SR}_r^{\mathcal{D}^{0,j}}(\mathcal{X})$ and if $j > i^*$, then \mathcal{A}_{i^*} constructs $\text{SR}_r^{\mathcal{D}^{1,j}}(\mathcal{X})$. Note that the resulting families of small-range distributions are $\frac{k\ell(q)}{r}$ -close to the original family of distributions by Lemma 15.
 - Recall that in the setting where $j < i^*$, \mathcal{A}_{i^*} constructs the oracle U_{f_j} by first constructing a function $f_j \leftarrow \text{SR}_r^{\mathcal{D}^{0,j}}(\mathcal{X})$ and in the setting where $j > i^*$, \mathcal{A}_{i^*} constructs the oracle U_{f_j} by first constructing a function $f_j \leftarrow \text{SR}_r^{\mathcal{D}^{1,j}}(\mathcal{X})$ for all $j \in [k] \setminus \{i^*\}$. Then \mathcal{A}_{i^*} evaluates the function on the queried inputs in superposition. Then, it is clear that \mathcal{A}_{i^*} can simulate U_{f_j} for any $j \neq i^*$ since it is possible to reconstruct U_{f_j} given access to its full classical description. Explicitly, U_{f_j} is implemented via the following circuit

$$\sum_{j,x,y} |j\rangle_I |x\rangle_X |y\rangle_Y |0\rangle_A |0\rangle_B \quad (108)$$

$$\rightarrow \sum_{j,x,y} |j\rangle |x\rangle |y\rangle |b = (j > i^*)\rangle |0\rangle \quad (109)$$

$$+ \sum_{j,x,y} |j\rangle |x\rangle |y\rangle |b = (j = i^*)\rangle |0\rangle \quad (110)$$

$$\rightarrow \sum_{j,x,y} |j\rangle |x\rangle |y\rangle |b = (j > i^*)\rangle |\text{SR}_r^{\mathcal{D}^{b,j}}(\mathcal{X})\rangle \quad (111)$$

$$+ \sum_{j,x,y} |j\rangle |x\rangle |y\rangle |(j = i^*)\rangle |g\rangle \quad (112)$$

$$\rightarrow \sum_{j,x,y} |j\rangle |x\rangle |y \oplus f_{b,j}(x)\rangle |b = (j > i^*)\rangle |\text{SR}_r^{\mathcal{D}^{b,j}}(\mathcal{X})\rangle \quad (113)$$

$$+ \sum_{j,x,y} |j\rangle |x\rangle |y \oplus f_{i^*}(x) \cdot b\rangle |(j = i^*)\rangle |g\rangle \quad (114)$$

$$\rightarrow \sum_{j,x,y} |j\rangle |x\rangle |y \oplus f_{b,j}(x)\rangle |b = (j > i^*)\rangle |0\rangle \quad (115)$$

$$+ \sum_{j,x,y} |j\rangle |x\rangle |y\rangle |b = (j = i^*)\rangle |0\rangle \quad (116)$$

$$\rightarrow \sum_{j,x,y} |j\rangle |x\rangle |y \oplus f_{b,j}(x)\rangle |0\rangle |0\rangle \quad (117)$$

$$+ \sum_{j,x,y} |j\rangle |x\rangle |y \oplus g(x)\rangle |0\rangle |0\rangle \quad (118)$$

$$\rightarrow \sum_{j,x,y} |j\rangle |x\rangle |y \oplus f_{b,j}(x)\rangle |0\rangle |0\rangle \quad (119)$$

The first arrow is implemented by evaluating a bit b , indicating if $j < i^*$ or if $j > i^*$ in an ancilla register A . The second arrow is implemented applying a CNOT gate controlled on registers A and I and then selecting the distribution $\text{SR}_r^{\mathcal{D}_{b,j}}(\mathcal{X})$ which is constructed from the non-uniform advice. The third arrow follows from generating a sample from $\text{SR}_r^{\mathcal{D}_{b,j}}(\mathcal{X})$ controlled on the X register, which by definition is equivalent to evaluating $f_{b,j}(x)$ (see preliminaries for evaluation of a function construction from a small-range distributions). The fourth arrow follows from applying a CNOT controlled on register B and outputting the final sample in the output register, Y . In this step, we also uncompute the ancilla register B . The fifth arrow follows from uncomputing register A .

C Security Parameter Preserving Proof

In this section, we provide a proof for a special case of Lemma 18 in which the distinguisher \mathcal{D} is restricted to querying security parameters $m = n$. This is known as the security parameter preserving setting.

Proof (Sketch). Our proof uses the meta-reduction approach of Gentry-Wichs, proceeding as follows.

1. **Construct a simulator \mathcal{S} .** We begin by observing that, for any subset membership problem $(\mathcal{L}_n, \bar{\mathcal{L}}_n, \text{Samp})$ for L , the following holds. Define the algorithm $\mathcal{S}(1^n)$ that, on input $(1^n, \text{crs})$, samples $(x, w) \leftarrow \text{Sam}(1^n)$, then computes $\pi \leftarrow \text{II.P}(1^n, \text{crs}, x, w)$, and finally outputs (x, π) . By the completeness property of II , the output of $\mathcal{S}(1^n)$ is always accepted by II.V .
2. **Construct an inefficient II -adversary $\bar{\mathcal{P}}$ for every crs .** Next, we choose the subset membership problem to be sufficiently hard, so that the following holds for (certain) superpolynomial s and negligible ε : if Corollary 5 is applied to the distributions $\mathcal{L}_{n,\text{crs}}$, $\bar{\mathcal{L}}_{n,\text{crs}}$, and the distribution defined by \mathcal{S} the result is an inefficient algorithm $\bar{\mathcal{P}}$ whose outputs $(\bar{x}, \bar{\pi})$ are $(s'(n), \varepsilon'(n))$ -sample-indistinguishable from the outputs (x, π) of $\mathcal{S}(1^n)$, but with $\bar{x} \in \bar{L}$. It then follows that the no-instances and “fake proofs” outputted by $\bar{\mathcal{P}}$ are accepted by II.V with overwhelming probability for the specified crs . This establishes the first part of the lemma.
3. **Show that $\bar{\mathcal{P}}$ and \mathcal{S} are oracle-indistinguishable.** Finally, we show that our choice of s and ε above was such that the (s, ε) -sample-indistinguishability of $\bar{\mathcal{P}}$ and $\mathcal{S}(1^n)$ implies (via Theorem 17) that $\bar{\mathcal{P}}$ and $\mathcal{S}(1^n)$ are in fact not distinguishable even by a QPT oracle algorithm making quantum queries.¹⁰

¹⁰ Note that this statement is only true when \mathcal{S} and $\bar{\mathcal{P}}$ appropriately randomize their coins. The details of this process are explained later in the full proof.

That is, we observe that while $\overline{\mathcal{P}}$ (resp. $\mathcal{S}(1^n)$) is a randomized algorithm, we can rewrite them as deterministic functions that map a certain choice of security parameter, 1^n , common reference string, crs , and random coins coins to a certain sample $(\bar{x}, \bar{\pi})$ (resp. (x, π)). Therefore, by considering quantum queries to $\overline{\mathcal{P}}$ and \mathcal{S} as quantum queries to a distribution family over deterministic functions, we can apply Theorem 17, which establishes the second part of the lemma.

This finishes the proof sketch. \square

Proof. We begin by choosing the hardness of the subset membership problem. By the succinctness of Π , we can choose a sufficiently large constant d such that the length of $\text{crs} \leftarrow \Pi.\mathcal{G}(1^n)$ is bounded by $O(n^d)$ and the length of any proof $\pi \leftarrow \Pi.\mathcal{P}(x, w)$ is bounded by $O(n^d)(|x| + |w|)^{o(1)}$. By the assumption on L , there exists a $(s(n), \varepsilon(n))$ -hard subset membership problem $(\mathcal{L}_n, \bar{\mathcal{L}}_n, \text{Sam})$ with

$$s(n) = 2^{cn^{d+3}} \quad \varepsilon(n) = 2^{-n^{d+3}} \quad (120)$$

where c is defined by the polynomial in Corollary 5. The explicit choice of c will be clear later in Equation (49). Let $\ell_{pf}(n) = o(n^{d+3})$ be the length of the proof for $\pi \leftarrow \Pi.\mathcal{P}(1^n, x, w)$ when $\text{crs} \leftarrow \mathcal{V}(1^n)$ and $(x, w) \leftarrow \text{Sam}(1^n)$. Note that if a problem is $(s(n), \varepsilon(n))$ hard, then it is also $(s(n), \tilde{\varepsilon}(n))$ hard for $\tilde{\varepsilon} \geq \varepsilon$.

Let $\mathcal{L}_{n, \text{crs}}^{+\text{aux}}$ be the output distribution of \mathcal{S} . Recall that, by definition, \mathcal{L}_n and $\bar{\mathcal{L}}_n$ are (s, ε) indistinguishable. Then by applying Corollary 5 to $\mathcal{L}_n, \bar{\mathcal{L}}$, and $\mathcal{L}_{n, \text{crs}}^{+\text{aux}}$, we obtain the existence of a distribution $\bar{\mathcal{L}}_{n, \text{crs}}^{+\text{aux}}$ that augments $\bar{\mathcal{L}}_n$ with bitstrings of size at most $\ell_{pf}(n)$ such that $\mathcal{L}_{n, \text{crs}}^{+\text{aux}}$ and $\bar{\mathcal{L}}_{n, \text{crs}}^{+\text{aux}}$ are $(s'(n), \varepsilon'(n))$ -sample indistinguishable.

We now carefully choose $(s(n), \tilde{\varepsilon}(n))$ to obtain the explicit bounds on $s'(n)$ and $\varepsilon'(n)$ given by Corollary 5. We obtain that, for $\tilde{\varepsilon} = 2^{-n^{d+2}}$, there exist some $(s'(n), \varepsilon'(n))$ such that $s(n) = \text{poly}(s'(n), \frac{1}{\tilde{\varepsilon}}, n, 2^{\ell_{pf}})$ and $\varepsilon'(n) = \Theta(\tilde{\varepsilon}(n))$. It follows that

$$s'(n) = \Theta\left(\left(\frac{s}{\text{poly}(\frac{1}{\tilde{\varepsilon}}, 2^{\ell_{pf}}, n)}\right)^{1/c}\right) = \Theta\left(\frac{(2^{cn^{d+3}})^{1/c}}{\text{poly}(\frac{1}{\tilde{\varepsilon}}, 2^{\ell_{pf}}, n)^{1/c}}\right) \quad (121)$$

$$= \frac{2^{\Omega(n^{d+3})}}{(\Theta(2^{k(n^{d+2})})\text{poly}(2^{o(n^{d+3})}, n))^{1/c}} \quad (122)$$

$$= 2^{\Omega(n^{d+3})} \quad (123)$$

and

$$\varepsilon'(n) = \Theta(\tilde{\varepsilon}(n)) = 2^{-\Theta(n^{d+2})}. \quad (124)$$

By the argument above, $\bar{\mathcal{L}}_{n, \text{crs}}^{+\text{aux}}$ is $(s'(n), \varepsilon'(n))$ -sample-indistinguishable from $\mathcal{L}_{n, \text{crs}}^{+\text{aux}}$. We then define $\overline{\mathcal{P}}$ to be the algorithm that samples and outputs $(\bar{x}, \bar{\pi}) \leftarrow \bar{\mathcal{L}}_{n, \text{crs}}^{+\text{aux}}$.

The first part of the lemma follows from the fact that $\overline{\mathcal{P}}$ outputs statements and proofs $(\bar{x}, \bar{\pi})$ such that $\bar{x} \notin L$ that are still accepted by \mathbf{V} . Formally,

$$\begin{aligned} & \Pr \left[\mathbf{V}(\text{priv}, \bar{x}, \bar{\pi}) = 1 \left| \begin{array}{l} (\text{crs}, \text{priv}) \leftarrow \mathbf{G}(1^n) \\ (\bar{x}, \bar{\pi}) \leftarrow \overline{\mathcal{P}}(1^n, \text{crs}, \text{coins}) \end{array} \right. \right] \\ & \geq \Pr \left[\mathbf{V}(\text{priv}, x, \pi) = 1 \left| \begin{array}{l} (\text{crs}, \text{priv}) \leftarrow \mathbf{G}(1^n) \\ (x, \pi) \leftarrow \mathcal{S}(1^n, \text{crs}, \text{coins}) \end{array} \right. \right] - \varepsilon'(n) \\ & \geq 1 - \text{negl}(n). \end{aligned} \quad (125)$$

Where the first inequality follows from applying Corollary 5 to argue $(\bar{x}, \bar{\pi})$ is $(s'(n), \varepsilon'(n))$ -indistinguishable from (x, π) .

Security Parameter Preserving Reductions. As a warm-up, we start with the case of security parameter preserving reductions, i.e., when the distinguisher \mathcal{D} queries only the security parameter $m = n$. To prove the second part of the lemma, our goal is to argue that $\overline{\mathcal{P}}$ can be replaced with \mathcal{S} . Let \mathcal{D} be a QPT oracle distinguisher and let ρ be a polynomial such that

$$\Pr[\mathcal{D}^{\overline{\mathcal{P}}}(1^n) = 1] - \Pr[\mathcal{D}^{\mathcal{S}}(1^n) = 1] > \frac{1}{\rho(n)} \quad (126)$$

Consider the variant $\overline{\mathcal{P}}_{RO}$ of $\overline{\mathcal{P}}$, which encodes a random oracle, RO , and a variant of \mathcal{S} , $\mathcal{S}_{f_{2q}}$, which encodes a $2q$ -wise independent function, f_{2q} . Essentially, $\overline{\mathcal{P}}_{RO}$ and $\mathcal{S}_{f_{2q}}$ take as input the queried coins queries by \mathcal{D} and generate uniformly random coins via RO and f_{2q} , respectively. Let register 1 be the register that contains the crs , let register 2 be the register that contains the queried random coins, let register 3 be the output register, and let A be an ancilla register initialized to all zeros. Then we define oracles $U_{n, \overline{\mathcal{P}}}$ and $U_{n, \mathcal{S}}$ for $\overline{\mathcal{P}}_{RO}$ and $\mathcal{S}_{f_{2q}}$ respectively. Note that $\overline{\mathcal{P}}_{RO}$ and \mathcal{S} both are now deterministic functions of n, crs , and coins , and their oracles can be defined analogously to Equation (12). We provide the definitions of the oracles here for concreteness.

$$U_{n, \overline{\mathcal{P}}}: |\text{crs}\rangle_1 |\text{coins}\rangle_2 |y\rangle_3 |\bar{0}\rangle_A \quad (127)$$

$$\rightarrow |\text{crs}\rangle_1 |\text{coins}\rangle_2 |y\rangle_3 |RO(\text{coins})\rangle_A \quad (128)$$

$$\rightarrow |\text{crs}\rangle_1 |\text{coins}\rangle_2 |y \oplus \overline{\mathcal{P}}(1^n, \text{crs}, RO(\text{coins}))\rangle_3 |\bar{0}\rangle_A. \quad (129)$$

$$U_{n, \mathcal{S}}: |\text{crs}\rangle_1 |\text{coins}\rangle_2 |y\rangle_3 |\bar{0}\rangle_A \quad (130)$$

$$\rightarrow |\text{crs}\rangle_1 |\text{coins}\rangle_2 |y\rangle_3 |f_{2q}(\text{coins})\rangle_A \quad (131)$$

$$\rightarrow |\text{crs}\rangle_1 |\text{coins}\rangle_2 |y \oplus \mathcal{S}(1^n, \text{crs}, f_{2q}(\text{coins}))\rangle_3 |\bar{0}\rangle_A. \quad (132)$$

Note that we can a priori upper bound the number of queries q made by the \mathcal{D} since \mathcal{D} runs in time $\tilde{p}(n)$ for a fixed polynomial \tilde{p} . Therefore we can choose

f_{2q} such that the output distribution generated by making q queries to $f_{2q}()$ is perfectly indistinguishable from that of making q queries to a random function [Zha12a, Fact 2]. Therefore it holds that

$$\Pr[\mathcal{D}^{|\overline{\mathcal{P}}\rangle_{\mathcal{S}}}(1^n) = 1] = \Pr[\mathcal{D}^{|\overline{\mathcal{P}}_{RO}\rangle}(1^n) = 1] \quad (133)$$

$$\Pr[\mathcal{D}^{|\mathcal{S}\rangle_{\mathcal{S}}}(1^n) = 1] = \Pr[\mathcal{D}^{|\mathcal{S}_{f_{2q}}\rangle}(1^n) = 1] \quad (134)$$

where the left hand side concerns *distributional* quantum oracle access (see Equation (13)) and the right hand side refers to normal quantum oracle access (see Equation (12)). This implies that

$$\Pr[\mathcal{D}^{|\overline{\mathcal{P}}\rangle_{\mathcal{S}}}(1^n) = 1] - \Pr[\mathcal{D}^{|\mathcal{S}\rangle_{\mathcal{S}}}(1^n) = 1] > \frac{1}{\rho(n)}. \quad (135)$$

From this point, when we say the distinguisher is given oracle access to $\overline{\mathcal{P}}$ or \mathcal{S} , we will mean that the distinguisher is given oracle access to $\overline{\mathcal{P}}_{RO}$ and $\mathcal{S}_{f_{2q}}$ as defined above.

We begin by applying Theorem 17 to $\overline{\mathcal{P}}$ and \mathcal{S} . That is, by setting $\mathbf{D}_0 = \overline{\mathcal{P}}$ and by setting $\mathbf{D}_1 = \mathcal{S}$ we get that $\overline{\mathcal{P}}$ and \mathcal{S} are $(s^*(n), \varepsilon^*(n), q(n))$ -oracle-indistinguishable where $s^*(n)$ and $\varepsilon^*(n)$ satisfy the appropriate constraints. Note that while Theorem 17 holds for samples generated uniformly at random, we use Equation (69) to argue that the samples are perfectly indistinguishable to those sampled uniformly at random. For the specific parameters we choose $s^*(n) = 2^{n^{d+2}}$ and $\varepsilon^*(n) = 2^{-n^{d+2}} = 2^{-\Omega(n^{d+2})}$, since

$$|\mathcal{D}(1^n)| \leq s^*(n) + \frac{106(n)^3(|x| + |\pi|)2^{3n^d}}{\varepsilon^*(n)} \quad (136)$$

$$\leq s^*(n) + \frac{106q(n)^3(\text{poly}(n) + o(n^{d+3}))2^{3n^d}}{\varepsilon^*(n)} \quad (137)$$

$$\leq 2^{n^{d+2}} + \frac{106q(n)^3 \text{poly}(n)2^{3n^d}}{2^{-n^{d+2}}} \quad (138)$$

$$= 2^{O(n^{d+2})} \quad (139)$$

$$< s'(n) \quad (140)$$

and

$$|\Pr[\mathcal{D}^{|\overline{\mathcal{P}}\rangle_{\mathcal{S}}}(1^n) = 1] - \Pr[\mathcal{D}^{|\mathcal{S}\rangle_{\mathcal{S}}}(1^n) = 1]| \geq \frac{1}{\rho(n)} \quad (141)$$

$$> \frac{\varepsilon^*(n)}{864q(n)^3 \cdot 2^{2n^d}} \quad (142)$$

$$= \frac{2^{-n^{d+2}}}{864q(n)^3 \cdot 2^{2n^d}} \quad (143)$$

$$= 2^{-\Omega(n^{d+2})} \quad (144)$$

$$\geq \varepsilon'(n). \quad (145)$$

Then since $\varepsilon^*(n) \leq \rho(n)$, the lemma statement follows.