Notions of Quantum Reductions and Impossibility of Statistical NIZK

Chuhan Lu and Nikhil Pappu

Portland State University
{chuhan,nikpappu}@pdx.edu

Abstract. Non-Interactive Zero-Knowledge Arguments (NIZKs) are cryptographic protocols that enable a prover to demonstrate the validity of an NP statement to a verifier with a single message, without revealing any additional information. The soundness and zero-knowledge properties of a NIZK correspond to security against a malicious prover and a malicious verifier respectively. Statistical NIZKs (S-NIZKs) are a variant of NIZKs for which the zero-knowledge property is guaranteed to hold information-theoretically. Previous works have shown that S-NIZKs satisfying a weak version of soundness known as static soundness exist based on standard assumptions. However, the work of Pass (TCC 2013) showed that S-NIZKs with the stronger *adaptive* soundness property are inherently challenging to obtain. The work proved that standard (black-box) proof techniques are insufficient to prove the security of an S-NIZK based on any standard (falsifiable) assumption. We extend this result to the setting where parties can perform quantum computations and communicate using quantum information, while the quantum security reduction is restricted to query the adversary classically. To this end, we adapt the well-known meta-reduction paradigm for showing impossibility results to the quantum setting. Additionally, we reinterpret our result using a new framework for studying quantum reductions, which we believe to be of independent interest.

Keywords: Non-Interactive Zero-Knowledge · Quantum Reductions · Black-Box Impossibilities

1 Introduction

Quantum computing is reshaping cryptography remarkably. Fundamental cryptosystems can be broken by efficient quantum algorithms (e.g., [Sho99]), and post-quantum cryptography has risen as a major effort to secure classical cryptography against quantum attacks. On the other hand, quantum computing can also be harnessed by honest users, and sometimes outperform what is possible classically. Notable examples in quantum cryptography range from unconditional key exchange to quantum money and advanced copy-protection primitives [BB14, Aar09, BI20, BL19], all

of which are provably impossibly in classical cryptography. Using quantum constructions, one can also sometimes bypass classical difficulties and obtain improved efficiency [BKS23, ABKK23, GJMZ23] and weaker assumptions [GLSV20, BCKM21, AQY22, CMS23]. Meanwhile, limitations of quantum advantages have been discovered, where quantum does not fare significantly better on various tasks [Lo97, CHS20, HY20, ABDS21, CLM23].

Motivated by this situation, we study the plausibility of quantum computing in bypassing a classical impossibility result related to an important cryptographic primitive: Non-Interactive Zero-Knowledge (NIZK) arguments [BFM88, DSDCO$^+$01]. NIZKs form critical building blocks for various primitives, such as in signature schemes [BG89], CCA secure encryption [Sah99], and cryptocurrencies [SCG$^+$14]. While many NIZKs have been constructed over the years [BDSMP91, GO94, FLS99, GOS06, SW14, CCH$^+$19, PS19], the case of an important variant known as statistical NIZK (S-NIZK) with adaptive (computational) soundness remains unsatisfactory. The only known constructions rely on non-standard assumptions [AF07], and the work of Pass [Pas13] proved that in fact they cannot be constructed from the natural and desirable assumptions formalized as *falsifiable assumptions* [Nao03, GW11]. We hence pose the following question:

Does the quantum setting provide an advantage for the construction of adaptively sound S-NIZKs for NP-complete languages?

To effectively study the plausibility of the primitive, we need a finer look at how the primitive utilizes the given assumption. For instance, in constructing a signature from a NIZK system, we can differentiate if the signature scheme runs the NIZK as a black-box, and whether a forger is run as a black-box to break the NIZK system in the security reduction. In this regard, the result of Pass rules out reductions with black-box access to the adversary from adaptively-sound S-NIZK to any falsifiable assumption. Moreover, the underlying assumption may be used in a non-black-box way, both by the reduction and the construction. Although classical frameworks for classifying reductions [RTV04, BBF13] have been studied systematically, this is largely missing in the quantum setting. We hence aim to develop a framework suited for the quantum setting, which is general enough to effectively capture reductions in both post-quantum and quantum cryptography.

1.1 Our Contributions

Impossibility of S-NIZK in the Quantum Setting. We give a negative answer to the aforementioned question by demonstrating that for any S-NIZK protocol for an NP-complete language, its adaptive soundness cannot be reduced to a falsifiable assumption using a quantum black-box reduction. The result holds even if the protocol utilizes quantum computation and communication. We stress that there exist non-trivial languages in NP for which S-NIZKs can be constructed (even classically) from such assumptions [BR90, BFM88]. Moreover, it is possible that quantum S-NIZK protocols exist for certain languages, for which classical S-NIZK protocols do not exist. However, our result says that quantum protocols do not help in regards to the hardest languages in NP. To show our result, we require the existence of certain distributions. Specifically, the existence of an efficiently sampleable distribution $X_{\mathcal{L}}$ over statements in $\mathcal{L}$ and a distribution $\widetilde{X}_{\mathcal{L}}$ over statements in $\{0,1\}^{\star} \setminus \mathcal{L}$, that are indistinguishable by QPT algorithms with quantum advice. Assuming the existence of post-quantum one-way functions (w.r.t. quantum-advice), every NP-complete language satisfies this property.

Our result builds upon the classical impossibility of Pass [Pas13] and employs the meta-reduction paradigm [BV98], which has been used to establish cryptographic impossibility results [Cor02, DOP05, GW11, Wic13, Pas13, BDSG+13, MP18, DLS22]. Along the way, we extend this paradigm to the quantum setting. This extension proves to be subtle because quantum algorithms derive randomness from entanglement, unlike classical algorithms that use a random tape as input. Our result is stated as follows:

Theorem 1 (Informal). *Let Π be a non-interactive quantum protocol for an NP-complete language $\mathcal{L}$, satisfying the statistical zero-knowledge property. Let R be a quantum black-box reduction that has classical access to $\mathcal{A}$, such that for every attacker $\mathcal{A}$ that breaks the adaptive soundness of Π, $\mathsf{R}^{\mathcal{A}}$ breaks some falsifiable assumption C. Then, assumption C is false, assuming the existence of post-quantum one-way functions.*

In the context of our impossibility result, we restrict our attention to quantum reductions that access the attacker via classical queries. In Section 4.3, we discuss some challenges in extending the impossibility to reductions making superposition queries. Although classical queries to a quantum attacker of a quantum protocol might seem overly restrictive, we believe this restriction is meaningful in this context. This is because the attacker only takes a Common Reference String (CRS) as input, in contrast to several (usually multi-round) quantum protocols where it expects

a quantum input as part of the protocol [GLSV21, CMS23]. Consequently, quantum NIZK protocols with a classical setup in the literature only invoke the prover classically [Shm21, BCKM21] in their soundness reductions. On a different note, examples where superposition access provides an advantage seem to be contrived [Zha12, BZ13, AMRS20], unlike advantages arising from the use of quantum communication [GLSV21, CMS23, BCKM21].

While our work focuses on the CRS model, one could also consider the designated verifier model, where the verifier possesses a secret verification key in addition to the CRS. The result of Pass extends to this model, as noted in the work of Campanelli et al. [CGKS23]. We note that it is also relatively straightforward to extend our impossibility to this setting, even for quantum secret keys.

Framework for Quantum Reductions. In cryptographic proofs, a reduction is typically employed to demonstrate that the security of some *construction* of a primitive $\mathcal{P}$ can be achieved using an *implementation* of another primitive $\mathcal{Q}$. This reduction, represented as $\mathcal{P} \to \mathcal{Q}$ (i.e. $\mathcal{P}$ reduces to $\mathcal{Q}$), consists of: (1) a construction G of $\mathcal{P}$ using an implementation f of $\mathcal{Q}$; and (2) a security reduction R that transforms any successful attack on G into one on f. Thus, if f is secure, then so is G.

In the classical setting, the work of Baecher, Brzuska, and Fischlin [BBF13] introduced a framework for categorizing various types of reductions based on their black-box nature. Central to their framework is their so called CAP notation. The notation employs three characters, each of which can either be 'B' (black-box) or 'N' (non-black-box), indicating access corresponding to the construction ($\underline{\mathbf{C}}$AP), adversary (C$\underline{\mathbf{A}}$P) and primitive (CA$\underline{\mathbf{P}}$) . For instance, a reduction from $\mathcal{P}$ to $\mathcal{Q}$ with black-box access for all components is denoted as $(\mathcal{P} \to \mathcal{Q}) - \mathrm{BBB}$, and described as follows:

- $\underline{\mathbf{B}}$BB: the construction of $\mathcal{P}$ makes *black-box use* of an implementation of $\mathcal{Q}$;
- B$\underline{\mathbf{B}}$B: the security reduction R makes *black-box use* of an attacker $\mathcal{A}$ that breaks the construction of $\mathcal{P}$;
- BB$\underline{\mathbf{B}}$: the security reduction R makes *black-box use* of implementation $\mathcal{Q}$.

In this study, we introduce a quantum counterpart of this notation, and refer to it as the Q-CAP system. This system aims to categorize various combinations that arise when different entities, such as the reduction, construction, or adversary, are implemented as quantum algorithms. Our

system uses the three letter approach of the CAP notation, and incorporates the $|\,\rangle$ notation to indicate the quantum components. For example, consider these variants of the classical BBB reduction:

- $|\mathrm{BBB}\rangle$: both the construction and reduction are quantum;
- $\mathrm{B}|\mathrm{BB}\rangle$: only the reduction is quantum;
- $|\mathrm{B}\rangle\mathrm{BB}$: only the construction is quantum.

Furthermore, we utilize superscripts to indicate query access power (c for classical queries and q for quantum queries); and subscripts to specify whether the queried party is classically (c) or quantumly (q) implemented. When the superscripts/subscripts represent "quantum", they can be omitted for brevity. We also consider reductions which obtain both oracle access to a unitary U along with oracle access to the reverse implementation $U^\dagger$ as a separate case. We use the letter 'S' standing for 'strong-black-box' to denote such access. Although such access is commonly assumed in quantum cryptography literature, there is some debate as to whether it should be considered black-box when compared with the classical black-box definition [DLS22]. We illustrate the notation's details in Fig. 1.

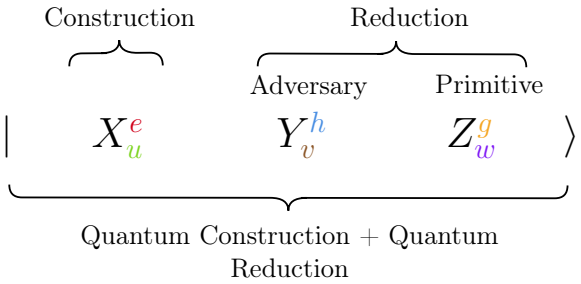

Fig. 1. Q-CAP Notation for $(\mathcal{P} \to \mathcal{Q})$: $X, Y, Z \in \{\mathrm{B}, \mathrm{S}, \mathrm{N}\}$ and $e, h, g, u, v, w \in \{c, q\}$

Notation	Meaning
B	Black-Box access
S	Non-Black-Box access
N	Strong-Black-Box access (i.e. black-box access to both U and $U^\dagger$)
c	Classical query/implementation
q	Quantum query/implementation
C-A-P	
X	Construction of $\mathcal{P}$ makes B/N/S use of an implementation of $\mathcal{Q}$
Y	Reduction makes B/N/S use of an adversary breaking $\mathcal{P}$
Z	Reduction makes B/N/S use of primitive $\mathcal{Q}$
Superscripts	
e	c/q-query access to implementation of $\mathcal{Q}$ by construction of $\mathcal{P}$
h	c/q-query access to an adversary breaking $\mathcal{P}$ by the reduction
g	c/q-query access to implementation of $\mathcal{Q}$ by the reduction
Subscripts	
u	Construction works for c/q-implementation of $\mathcal{Q}$
v	Reduction works for c/q-adversary $\mathcal{A}$
w	Reduction works for c/q-implementation of $\mathcal{Q}$

Our motivation in defining such a framework is to conveniently capture the quantum reductions demonstrated in the literature and to better understand the relationships between different types of reductions. The framework helps us identify the precise types of reductions that separation results rule out. For example, the work of Hosoyamada and Yamakawa [HY20] demonstrates that collision-resistant hash functions (CRHFs) cannot be reduced to one-way permutations (OWPs) using $|\mathrm{B}_c\mathrm{B}_c\mathrm{B}_c\rangle$ reductions. In other words, even quantum constructions and security reductions are ruled out, considering quantum superposition access to both the adversary and the OWP.

To the best of our knowledge, there is currently no consistent framework for studying quantum reductions. A unified framework would make it easier to compare different results and identify avenues for improvement. Compared to the categorizations of previous works on quantum separation results [HY20, CX21], we believe our notation presents more aspects of the reduction in a succinct way, making it easier to parse. In Section 3, we delve into the relations between the different reductions and illustrate several use cases of the framework.

1.2 Related Work

Pass' Classical Black-Box Impossibility. Our work is most related to the work of Pass [Pas13], since our main result is a generalization of the classical impossibility shown in that work. The work showed that there does not exist any S-NIZK protocol for any NP-complete language with a corresponding black-box security reduction to a falsifiable assumption, unless the assumption is false. Our impossibility differs in that it rules out quantum S-NIZK protocols (for NP-complete languages) that may also utilize quantum communication. Furthermore, we also consider quantum black-box reductions and quantum falsifiable assumptions. However, we restrict our attention to reductions that only make classical queries to the attacker, and describe some challenges in making the extension to quantum queries. Our proof follows the same template as that of Pass' work and utilizes the meta-reduction technique. However, we observe that even with classical queries, there are certain nuances posed by the fact that the internal states of the reduction, attacker, and zero-knowledge simulators are all quantum. These nuances arise due to the difficulty of quantum rewinding and the lack of a quantum analogue for programming a classical random tape.

S-NIZKs in the Quantum Setting. The work of Canetti et al. [CCH+19] constructs a static sound S-NIZK based on circular-secure LWE, which is a quantum secure assumption. The work of Peikert et al. [PS19] improved on this to obtain an S-NIZK based on plain LWE. Their construction was shown to be quantum-secure in the work of Coladangelo et al. [CVZ20]. The work of Morimae and Yamakawa [MY23] studied S-NIZKs for the complexity class QMA, which subsumes the class NP. However, this protocol assumes a stronger quantum setup.

Black-Box Separations. In cryptography, there have been several works demonstrating the impossibility of constructing a target cryptographic primitive by making black-box use of a simpler primitive [IR89, Sim98, GMR01, DOP05, GMMM18, RTV04, FS10]. Recently, several works have shown such impossibility results in the quantum setting [AHY23, CHS20, HY20, CX21, CLM23]. These impossibility results require that the construction of the target primitive makes black box use of the underlying primitive. For instance, the work of Ananth et al. [AHY23] rules out black-box constructions of public-key quantum money from CRHFs. On the other hand, starting from the work of Boneh et al. [BV98], the class of works employing the meta-reduction paradigm [Cor02, DOP05, GW11, Wic13, Pas13, BDSG+13, MP18, DLS22] only require that the security reduction makes black-box use of the attacker to break some crypto-

graphic assumption. Consequently, these impossibility results apply even to constructions that make non-black-box use of the underlying primitive. Moreover, this technique has been utilized to show separations between certain primitives and broad classes of assumptions, including concrete assumptions such as RSA [RSA78] and DDH [DH22].

The meta-reduction paradigm has also been employed to establish impossibilities in the quantum setting. The work of Jiang et al. [JZM21] demonstrated an impossibility concerning CCA-secure KEMs in the quantum setting by utilizing this technique. It is worth noting that their result assumes that the security reduction invokes the adversary only once, which is motivated by the difficulty of quantum rewinding. Hence, they do not run into our subtlety specific to the quantum setting. The work of Dupuis et al. [DLS22] also presented an impossibility in the quantum setting using the meta-reduction approach. Specifically, they extended the classical impossibility result of Bitansky et al. [BDSG$^+$13] concerning the soundness of instantiations of the Fiat-Shamir paradigm to the quantum setting, where the parties share quantum entanglement. The work of Dagdelen et al. [DFG13] also presents an impossibility about soundness of the Fiat-Shamir transform in the quantum random oracle model (QROM). These impossibility results do not encounter the previously mentioned subtlety due to relying on a stronger premise of statistical indistinguishability compared to computational indistinguishability in our case. We shall elaborate on this in Section 4.

2 Preliminaries

Notations. We let λ denote the security parameter, which will be provided as input to the considered cryptographic algorithms in unary. We use $\mathsf{poly}(\lambda)$ to denote some polynomial in the security parameter. Many of the quantities we consider will be polynomials in λ. We let $\mathsf{negl}(\lambda)$ denote any function $f(\lambda)$ such that $f(\lambda) = O(\lambda^{-c})$ for every constant c. Likewise, we use $\mathsf{non\text{-}negl}(\lambda)$ to denote any function that doesn't satisfy the above property. Furthermore, all the cryptographic algorithms considered in this work will be non-uniform algorithms unless stated otherwise. In other words, they are provided with some $\mathsf{poly}(\lambda)$-size advice state as input, which depends only on the security parameter. In general, we allow the advice to be a quantum state. We let $R_{\mathcal{L}}(x)$ denote the set of all witnesses for a statement $x \in \mathcal{L}$, where $\mathcal{L}$ is a language in NP. We often omit the prefix "quantum" when referring to protocols, assumptions, etc., when it is clear from the context. We let $\mathsf{Adv}[\mathsf{M}]$ denote the advantage

of algorithm M in some experiment, i.e., the absolute difference between its probability of outputting 1 in two different cases. We use the acronym PPT for polynomial-time algorithms (in λ) and the acronym QPT for quantum polynomial-time algorithms. For two classical random variables X and Y, we use the notation $X \approx Y$ to denote that X and Y are computationally indistinguishable by non-uniform QPT distinguishers with quantum advice.

Quantum Information. An m-qubit register X corresponds to a Hilbert space $\mathbb{C}^{2^m}$. A pure state on register X is a unit vector $|\psi\rangle_X \in \mathbb{C}^{2^m}$. A density matrix is a positive semi-definite Hermitian matrix $\rho_X \in \mathbb{C}^{2^m \times 2^m}$ with trace 1. Let $\mathrm{Dens}(\mathbb{C}^{2^m})$ denote the set of all such density matrices. A probability distribution over pure states is captured by the notion of a mixed quantum state, which is represented by a density matrix. The evolution of a quantum state $|\psi\rangle_X$ is captured using a unitary transformation $U|\psi\rangle_X$, where U is a matrix satisfying $U^\dagger U = UU^\dagger = \mathbb{I}_{2^m}$. A measurement performed on a quantum system is described by a positive operator-valued measure (POVM), which is a set of matrices $\{M_i\}_i$ such that $\sum_i M_i^\dagger M_i = \mathbb{I}$. On performing the measurement on mixed state ρ, the probability of obtaining outcome i is given by the trace $\mathrm{Tr}\Big(M_i \rho M_i^\dagger\Big)$. The corresponding post-measurement state is given by $\rho' = \frac{M_i \rho M_i^\dagger}{Tr(M_i \rho M_i^\dagger)}$. The trace distance between mixed states ρ and σ is denoted as $\mathsf{TD}(\rho, \sigma)$, and is defined as follows:

$$\mathsf{TD}(\rho, \sigma) = \frac{1}{2}||\rho - \sigma||_1 = \frac{1}{2}\,\mathrm{Tr}\left[\sqrt{(\rho-\sigma)^\dagger(\rho-\sigma)}\right]$$

Intuitively, this quantity is a measure of the distinguishability between states ρ and σ. The larger the trace distance, the more accurately the states can be distinguished by an unbounded quantum algorithm.

2.1 Hard Languages

The impossibility result presented in this work holds for *hard languages* as defined in this section. Moreover, assuming the existence of post-quantum secure one-way functions (OWFs), every NP-complete language satisfies this notion. Note however that these languages need not necessarily be NP-complete. We describe this notion in the following definition:

Definition 1 (Hard Language). *A language $\mathcal{L} \in \mathsf{NP}$ is called a hard language if it satisfies the following conditions:*

1. *There exists an efficiently sampleable distribution $Z_{\mathcal{L}}$ over statement-witness tuples (x, w) where $x \in \mathcal{L}$ and $w \in R_{\mathcal{L}}(x)$. Let $X_{\mathcal{L}}$ denote the corresponding marginal distribution over the statements.*
2. *There exists a possibly inefficient to sample distribution $\widetilde{X}_{\mathcal{L}}$ over statements $\widetilde{x} \notin \mathcal{L}$.*
3. *The distributions $X_{\mathcal{L}}$ and $\widetilde{X}_{\mathcal{L}}$ are indistinguishable by non-uniform QPT algorithms with quantum advice.*

It is easy to see that such languages exists based on the existence of post-quantum OWFs w.r.t. non-uniform algorithms with quantum advice. To see this, consider the distribution $X_{\mathcal{L}}$ to correspond to the output of a pseudo-random generator (PRG) on a random seed, and the distribution $\widetilde{X}_{\mathcal{L}}$ to correspond to uniformly random values among the outputs that are not in the range of the PRG. The indistinguishability then follows by the security of a PRG.

2.2 Quantum Falsifiable Assumptions

Falsifiable assumptions [Nao03, GW11] categorize most cryptographic assumptions and are considered more desirable than other types of assumptions. Intuitively, these assumptions involve an interactive game between an efficient challenger, denoted as C, and an adversary, denoted as $\mathcal{A}$. At the conclusion of the game, the challenger determines whether the adversary successfully broke the assumption. To break the assumption, $\mathcal{A}$ must run in polynomial time and achieve a success probability significantly greater than a given threshold, denoted as t. For decision assumptions, $t = 1/2$, while for search assumptions, $t = 0$. The challenger C is a QPT machine that is additionally allowed to obtain some non-uniform quantum advice. This captures assumptions such as pseudo-random states [JLS18].

Definition 2 (Quantum Falsifiable Assumption). *A quantum falsifiable assumption (C, t) is defined by an interactive game between a non-uniform QPT algorithm C and an algorithm $\mathcal{A}$, where C outputs a bit at the end of the interaction. The interaction may involve quantum communication. Let $\langle \mathsf{C}, \mathcal{A} \rangle(1^\lambda)$ denote the output of C. The assumption is said to be true if for all non-uniform QPT algorithms $\mathcal{A}$, the following holds:*

$$\Pr\left[\langle \mathsf{C}, \mathcal{A} \rangle(1^\lambda) = 1\right] \leq t + \mathsf{negl}(\lambda)$$

A possibly inefficient algorithm $\mathcal{A}$ breaks the assumption if there exists some polynomial $p(\lambda)$ such that the following holds for infinitely many $\lambda \in \mathbb{N}$:

$$\Pr\left[\langle \mathsf{C}, \mathcal{A}\rangle(1^\lambda) = 1\right] \geq t + \frac{1}{p(\lambda)}$$

2.3 Quantum Statistical Non-Interactive Zero Knowledge Arguments (S-NIZKs)

We will now provide the definition of a Statistical Non-Interactive Zero Knowledge Argument (S-NIZK) in the quantum setting. The definition is similar to that of an S-NIZK in the classical setting, with some modifications to account for quantum capabilities. Specifically, the prover and verifier are allowed to perform quantum computation, and the proof sent by the prover can be a quantum state. Naturally, the zero-knowledge simulator is also allowed to be quantum. It is important to note that the protocol is defined in the Common Reference String (CRS) model, similar to the classical setting. In this model, both the prover and the verifier are initialized with a common string that is drawn from some efficiently sampleable distribution. This model is considered because NIZKs are impossible to obtain in the plain model for languages outside BQP, based on an analogous classical result for BPP [GO94].

Definition 3 (Quantum S-NIZK). *A quantum S-NIZK protocol Π for a language $\mathcal{L} \in \mathsf{NP}$ consists of three QPT algorithms $(\mathsf{Setup}, \mathsf{P}, \mathsf{V})$ corresponding to the CRS generator, prover, and verifier respectively. The protocol must satisfy the following conditions:*

Completeness: *This property requires that for every true statement, the verifier accepts the proof output by the prover with high probability. Formally, for every $x \in \mathcal{L}$ and $w \in R_{\mathcal{L}}(x)$, the following condition must hold:*

$$\Pr\left[\mathsf{V}\big(1^\lambda, \mathsf{crs}, x, |\pi\rangle\big) = 1 \,\middle|\, \begin{array}{c} \mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda) \\ |\pi\rangle \leftarrow \mathsf{P}(1^\lambda, \mathsf{crs}, x, w) \end{array}\right] \geq 1 - \mathsf{negl}(\lambda)$$

Adaptive Soundness: *This property requires that for every QPT malicious prover that outputs a statement and a proof, the probability that the statement is false and the verifier accepts is negligibly small. The adaptive aspect of the definition allows the malicious prover to choose the statement after seeing the CRS. Formally, for every QPT malicious prover $\mathsf{P}^\star$, the following must hold:*

$$\Pr\left[\mathsf{V}\big(1^\lambda, \mathsf{crs}, \widetilde{x}, |\widetilde{\pi}\rangle\big) = 1 \ \wedge\ \widetilde{x} \notin \mathcal{L} \,\middle|\, \begin{array}{c} \mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda) \\ (\widetilde{x}, |\widetilde{\pi}\rangle) \leftarrow \mathsf{P}^\star(1^\lambda, \mathsf{crs}) \end{array}\right] \leq \mathsf{negl}(\lambda)$$

A possibly inefficient adversary $\mathcal{A}$ breaks the adaptive soundness of Π if there exists a polynomial $p(\lambda)$ such that the following holds for infinitely many $\lambda \in \mathbb{N}$:

$$\Pr\left[\mathsf{V}\big(1^\lambda, \mathsf{crs}, \widetilde{x}, |\widetilde{\pi}\rangle\big) = 1 \ \wedge \ \widetilde{x} \notin \mathcal{L} \,\middle|\, \begin{matrix} \mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda) \\ (\widetilde{x}, |\widetilde{\pi}\rangle) \leftarrow \mathcal{A}(1^\lambda, \mathsf{crs}) \end{matrix}\right] \geq \frac{1}{p(\lambda)}$$

Statistical Adaptive Zero Knowledge *Consider an adaptive QPT adversary $\mathcal{A}$ that picks a statement (and witness) for the prover to prove after observing the CRS. Let D be an unbounded distinguisher that receives state st output by $\mathcal{A}$, along with the proof output by the prover. The adaptive zero-knowledge property requires that there exists a QPT simulator that can produce a view for D that is statistically close to the one above, without access to the witness. Intuitively, this captures that no malicious verifier can learn additional information even in unbounded time, even with the ability to choose arbitrary statements.*

Without loss of generality, we can consider a two part simulator $\mathsf{S} = (\mathsf{S}_1, \mathsf{S}_2)$ where S_1 produces a simulated CRS, and S_2 produces a simulated proof for a given statement. Note that in all generality, S is a non-uniform algorithm that may obtain some polynomial size quantum advice state $|\phi\rangle$ as input. However, the state must be well-defined and can be generated by an unbounded algorithm. Furthermore, $|\phi\rangle$ may depend both on $\mathcal{A}$ and D. Formally, we require that for every non-uniform QPT $\mathcal{A}$ and unbounded D, the following condition holds:

$$\begin{aligned} \mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda) \qquad & (\mathsf{crs}', |\mathsf{aux}\rangle) \leftarrow \mathsf{S}_1(1^\lambda, |\phi\rangle) \\ (x, w, \mathsf{st}) \leftarrow \mathcal{A}(1^\lambda, \mathsf{crs}) \qquad & (x', w', \mathsf{st}') \leftarrow \mathcal{A}(1^\lambda, \mathsf{crs}') \\ |\pi\rangle \leftarrow \mathsf{P}(1^\lambda, \mathsf{crs}, x, w) \qquad & |\pi'\rangle \leftarrow \mathsf{S}_2(1^\lambda, x', |\mathsf{aux}\rangle) \end{aligned}$$

$$\Big|\Pr\big[\mathsf{D}\big(\mathsf{st}, |\pi\rangle\big) = 1 \wedge (x, w) \in R_{\mathcal{L}}(x)\big] - \Pr\big[\mathsf{D}\big(\mathsf{st}', |\pi'\rangle\big) = 1 \wedge x' \in \mathcal{L}\big]\Big| \leq \mathsf{negl}(\lambda)$$

2.4 Quantum Oracle Algorithms

For the sake of our main theorem, our focus will be on QPT algorithms that make black-box use of quantum adversaries. This means that the algorithm relies solely on the input-output interface of the adversary and does not depend on its internal workings. Since the adversaries considered in this work are stateless, we can consider algorithms that simply obtain oracle access to the adversary without loss of generality. It is important to

note that, in general, algorithms may have the ability to utilize rewinding. Furthermore, we do not consider quantum algorithms that may query their oracles in superposition. Therefore, the quantum oracle algorithms considered in this work will have the following form:

$$\mathsf{M}^{\mathcal{O}}(1^{\lambda}) = \mathcal{M} \circ \Big(\big(\mathcal{O} \circ \mathcal{M}_q \big) \circ \cdots \circ \big(\mathcal{O} \circ \mathcal{M}_1 \big) \Big) \left| \psi^0 \right\rangle$$

Here, $\mathcal{M}_1 = \{M_i\}_i$ be an arbitrary efficient POVM acting on initial advice state $\left|\psi^0\right\rangle$ in order to produce outcome c_1, and post-measurement state $M_{c_1} \left|\psi^0\right\rangle$. c_1 is then queried to oracle $\mathcal{O}$ to obtain response $|\phi_1\rangle$. The subsequent internal state if defined as follows: $\left|\psi^1\right\rangle = |\phi_1\rangle \otimes |c_1\rangle \otimes M_{c_1} \left|\psi^0\right\rangle$. Then, another POVM $\mathcal{M}_2$ is adaptively chosen, and applied to $\left|\psi^1\right\rangle$, in order to generate the next oracle query c_2. In a similar vein, after q queries have been made, a binary POVM $\mathcal{M}$ is applied to the state $|\psi^q\rangle$ to obtain the output.

3 Framework for Quantum Cryptographic Reductions

A cryptographic reduction demonstrates that a target primitive $\mathcal{P}$ can be constructed using another cryptographic primitive $\mathcal{Q}$. Such a relation between the primitives is referred to as "$\mathcal{P}$ reduces to $\mathcal{Q}$", and is denoted by $\mathcal{P} \to \mathcal{Q}$. For example, a foundational result shows that pseudo-random generators ($\mathcal{P}$) can be built using one-way functions ($\mathcal{Q}$) (i.e. PRGs reduces to OWFs) [HILL99].

The primary objective of demonstraing such a reduction is to establish the security of a construction G of $\mathcal{P}$ that utilizes an implementation f of $\mathcal{Q}$, assuming the hardness of $\mathcal{Q}$. This is accomplished by developing a security reduction R that successfully breaks f of $\mathcal{Q}$ using any attack on G of $\mathcal{P}$ as a building block. In the quantum setting, either or both of $\mathcal{P}$ and $\mathcal{Q}$ may be quantum primitives. For example, private-key quantum money reduces to pseudo-random states (PRS) [JLS18], and zero-knowledge proofs for QMA reduce to bit commitments [BJSW16]. Moreover, different components of the construction and security reduction can be empowered by quantum capabilities. Consequently, the variety of possible reductions motivates further study.

In this regard, we will begin by defining some building blocks, before describing the various kinds of quantum reductions. Firstly, we require a notion of a primitive. The recent work of Coladangelo et al. [CM24] defines primitives among other fundamental notions in the quantum setting. They define a primitive as a set of quantum channels. However, we believe this is not without loss of generality, especially when considering

stateful primitives being accessed by other constructions. Consequently, we will make use of the notion of a quantum strategy as introduced by Gutoski and Watrous [GW07], and define a primitive using a set of quantum strategies.

Remark 1. One might consider a quantum channel with an internal register that maintains stateful information to be sufficient to capture these cases. However, we believe this is not formal enough, and that the notion of a quantum strategy precisely captures this intuition.

We will begin by recalling the notion of a quantum strategy.

Definition 4 (Quantum Strategy). *A quantum strategy f with N turns corresponding to input Hilbert spaces $\mathcal{X}_1, \ldots, \mathcal{X}_N$ and output Hilbert spaces $\mathcal{Y}_1, \ldots, \mathcal{Y}_N$ consists of the following:*

1. *Hilbert spaces $\mathcal{Z}_1, \ldots, \mathcal{Z}_N$ called as memory spaces.*
2. *An N-tuple of quantum channels $(f_1, \ldots, f_N)$ of the form:*

$$f_1 : L(\mathcal{X}_1) \mapsto L(\mathcal{Y}_1 \otimes \mathcal{Z}_1)$$
$$f_i : L(\mathcal{X}_i \otimes \mathcal{Z}_{i-1}) \mapsto L(\mathcal{Y}_i \otimes \mathcal{Z}_i) \quad (2 \leq i \leq N)$$

Intuitively, the N quantum channels and the N memory spaces capture a stateful machine. The i-th channel acts on a memory state output by the $(i-1)$-th channel along with an input state to produce a bipartite output state over the i-th output and memory registers. We now define a non-uniform quantum strategy which will be helpful in the context of efficient quantum strategies.

Definition 5 (Non-Uniform Quantum Strategy). *A non-uniform quantum strategy g is defined by a tuple (σ, f), where σ is a quantum advice state in some space $\mathcal{Z}_0$, and f is a quantum strategy as defined in Definition 4, except that the channel f_1 is defined as follows:*

$$f_1 : L(\mathcal{X}_1 \otimes \mathcal{Z}_0) \mapsto L(\mathcal{Y}_1 \otimes \mathcal{Z}_1)$$

We now define the notion of a cryptographic primitive as a tuple where the first element of the tuple captures the functionality of the primitive, while the second element captures the security requirement.

Definition 6 (Cryptographic Primitive). *A cryptographic primitive $\mathcal{P}$ is defined by a tuple $(\mathcal{F}_\mathcal{P}, \mathcal{R}_\mathcal{P})$ where $\mathcal{F}_\mathcal{P}$ is a set of non-uniform quantum strategies and $\mathcal{R}_\mathcal{P}$ is a set of tuples of the form $(f, \mathcal{A})$ such that $f \in \mathcal{F}_\mathcal{P}$ and $\mathcal{A}$ is a non-uniform quantum strategy.*

Intuitively, the tuple $(f, \mathcal{A})$ is included in $\mathcal{R}_{\mathcal{P}}$ if adversary $\mathcal{A}$ breaks the security definition for primitive $\mathcal{P}$ for the instantiation f. Next, in order to define what it means to realize a primitive, we provide the following definitions to capture the efficiency and security requirements of an implementation.

Definition 7 (Efficient Strategy). *A non-uniform quantum strategy (σ, f) as defined in Definition 5 is an efficient strategy if the following conditions hold:*

1. *The advice state σ has $\mathsf{poly}(\lambda)$ qubits.*
2. *f is a strategy with $N = \mathsf{poly}(\lambda)$ turns.*
3. *For every channel f_i, there exists a QPT algorithm that implements f_i.*

Definition 8 (Secure Implementation of $\mathcal{P}$). *A non-uniform quantum strategy G is a secure implementation of primitive $\mathcal{P}$ if $\mathsf{G} \in \mathcal{F}_{\mathcal{P}}$ and for every efficient strategy $\mathcal{A}$, $(\mathsf{G}, \mathcal{A}) \notin \mathcal{R}_{\mathcal{P}}$.*

Definition 9 (Efficient Implementation of $\mathcal{P}$). *A quantum strategy G is an efficient implementation of a primitive $\mathcal{P}$ if it is a secure implementation of $\mathcal{P}$ and an efficient strategy.*

We will now describe what it means to provide a strategy with black-box quantum access to another strategy. This captures both the access a construction gets to a primitive, as well as a reduction's access to an adversary. For this, we will utilize the co-strategy formalism of Gutoski and Watrous [GW07], defined as follows:

Definition 10 (Co-Strategy). *A quantum co-strategy G with N turns corresponding to output Hilbert spaces $\mathcal{X}_1, \ldots, \mathcal{X}_N$ and input Hilbert spaces $\mathcal{Y}_1, \ldots, \mathcal{Y}_N$ consists of the following:*

1. *Hilbert spaces $\mathcal{W}_0, \ldots, \mathcal{W}_N$ called as memory spaces.*
2. *An initial state $\rho_0 \in D(\mathcal{X}_1 \otimes \mathcal{W}_0)$*
3. *An N-tuple of quantum channels $(G_1, \ldots, G_N)$ of the form:*

$$\begin{aligned} G_i &: L(\mathcal{Y}_i \otimes \mathcal{W}_{i-1}) \mapsto L(\mathcal{X}_{i+1} \otimes \mathcal{W}_i) \quad (1 \leq i \leq N-1) \\ G_N &: L(\mathcal{Y}_N \otimes \mathcal{W}_{N-1}) \mapsto L(\mathcal{W}_N) \end{aligned}$$

Similar to a strategy, one can analogously define non-uniform and efficient co-strategies.

Definition 11 (Black-Box Access to Strategy). *Let* G *be a quantum co-strategy with black-box quantum access to a quantum strategy* f*, where* G *and* f *are as defined in Definitions 10 and 4 respectively. Such access is denoted by the notation* $\mathsf{G}^{|f\rangle}$*. Then, the* i*-th query of* G *to* f *may be an arbitrary quantum state on register* $\mathcal{X}_i$*.* f *then produces an output state on register* $\mathcal{Y}_i$ *using its internal state* $\mathcal{Z}_{i-1}$*. Consequently,* G *uses this output along with its internal state on register* $\mathcal{W}_{i-1}$ *to produce the next query on register* $\mathcal{X}_{i+1}$*. Analogously,* G *has classical black-box access to* f *(denoted by* G^f*), if the input queries on registers* $\mathcal{X}_i$ *are classical, i.e., are diagonal density matrices.*

Remark 2. Although the above formalism restricts G to query f sequentially, this is not without loss of generality as multiple simultaneous queries (For e.g., to a OWF) can be captured by sequential queries to an appropriately defined stateful machine.

Next, we define non-black-box quantum and classical access. It is important to note that in the non-black-box case, G obtains black-box access as well as the description of the strategy f. This is to ensure that even in the case of inefficient implementations, such access remains more powerful than its black-box counterpart. Consequently, the quantum or classical access refers solely to the nature of the black-box access.

Definition 12 (Non-Black-Box Access to Strategy). *Let* G *be a quantum co-strategy with non-black-box quantum access to a quantum strategy* f*, where* G *and* f *are as defined in Definition 10 and Definition 4. Then,* G *obtains black-box quantum access* $\mathsf{G}^{|f\rangle}$ *as defined in Definition 11. In addition,* G *obtains the description of each of the finite number of channels* $f_1, \cdots, f_N$*. Analogously,* G *has classical non-black-box access to* f *if the black-box access is classical as per Definition 11.*

Note that for the purposes of this work, we do not distinguish between an algorithm obtaining different kinds of descriptions of a quantum strategy.

On Strong-Black-Box Access: We will now informally discuss a notion we call strong-black-box access which works only for quantum strategies where each of the channels is a unitary map. We refer to these as unitary strategies. Moreover, access to the *inverse* of the strategy is also provided. Intuitively, if the strategy is defined by the unitary sequence $(\mathsf{U}_1, \ldots, \mathsf{U}_N)$, the inverse strategy is defined by the unitary sequence $(\mathsf{U}_N^\dagger, \ldots, \mathsf{U}_1^\dagger)$. This is motivated by the fact that such access is

commonly utilized in quantum cryptographic proofs, especially in the context of rewinding [CMSZ22,LMS22]. However, since the classical black-box notion does not allow such access, we choose to classify it separately, and denote such access as $\mathsf{G}^{|f,\dagger\rangle}$.

Next, we shall recap the way classical reductions were categorized in previous works, before discussing the quantum generalizations.

3.1 Recap: Classical CAP Notation

Previous works [RTV04, BBF13] have built frameworks that classify various kinds of classical reductions, with Baecher et al. [BBF13] introducing a convenient notation, *CAP-notation*, for this purpose. The notation uses three characters, each representing either 'B' for *black-box* or 'N' for *non-black-box access* [1] to the construction ($\underline{\mathbf{C}}$AP), adversary (C$\underline{\mathbf{A}}$P) and primitive (CA$\underline{\mathbf{P}}$) respectively.

In this notation, a fully black-box reduction from $\mathcal{P}$ to $\mathcal{Q}$ is represented as $(\mathcal{P} \to \mathcal{Q}) - \mathrm{BBB}$, where the first 'B' indicates that $\mathcal{P}$ is constructed using an implementation of $\mathcal{Q}$ as a black-box; the second 'B' denotes that the security reduction algorithm R makes black-box use of an attacker $\mathcal{A}$ breaking the construction of $\mathcal{P}$; the last 'B' represents that the security reduction R is restricted to black-box access to the implementation of $\mathcal{Q}$. Note that the access of $\mathcal{A}$ to $\mathcal{Q}$ is not captured by the framework, and needs to be specified separately.

Next, we will introduce our framework for the quantum setting by building on the classical framework and utilizing the previously introduced quantum definitions.

3.2 Quantum Reductions and a Quantum CAP system

A (contrived) Example In this work, we develop a quantum CAP system aimed at unifying quantum reductions into a single framework. We refer to this as the *Q-CAP-notation*. To better understand this new system, we first consider the following (contrived) example (illustrated in Fig. 2) for the sake of exposition:

$$(\mathcal{P} \to \mathcal{Q}) - |\mathrm{B}^{c}_{c}\mathrm{SN}\rangle$$

The above reduction is described by the following aspects:

[1] Black-box access refers to using a component solely through its input-output interface, while non-black-box access involves utilizing the component's internal workings.

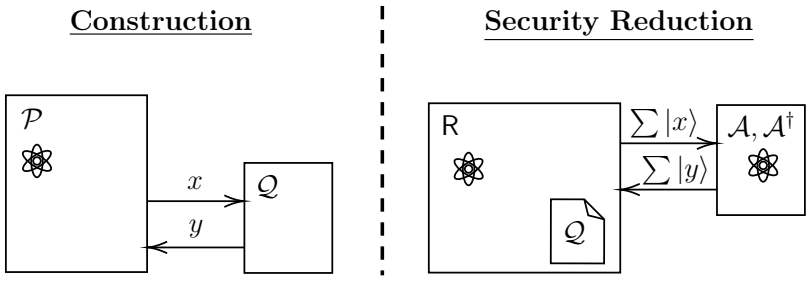

Fig. 2. $(\mathcal{P} \to \mathcal{Q}) - |\mathrm{B}^c_c\mathrm{SN}\rangle$ Reduction: (1) the quantum construction of $\mathcal{P}$ has classical query access to a classical implementation of $\mathcal{Q}$; (2) the quantum reduction R has quantum query access to the quantum adversary $\mathcal{A}$ and its reverse implementation $\mathcal{A}^\dagger$; (3) the reduction R is provided with the code of $\mathcal{Q}$ as non-black-box access. For simplicity, we have not indicated black-box access to $\mathcal{Q}$ by R, nor that R also works for quantum implementations $\mathcal{Q}$, despite the construction only working for classical $\mathcal{Q}$.

(1) Modes of access. In our new notation, in addition to 'B' (black-box) and 'N' (non-black-box), we further introduce the letter 'S' to capture the case where, in quantum settings, query access to the inverse of a unitary is also provided. We refer to it as *strong-black-box access*. In the example above, the 'S' specifically denotes that the reduction has access to both the unitary equivalent U of an adversary, and its inverse $U^\dagger$. We note that this is not without loss of generality, and that the strong-black-box characterization is only appropriate for adversary's that have a unitary equivalent.

(2) Indication of quantum components. In this example, to indicate that both the construction and reduction are quantum, we place $|\cdots\rangle$ around all three characters. Some other configurations are $|\mathrm{B}\rangle\,\mathrm{SN}$ (i.e. only construction is quantum) and $\mathrm{B}\,|\mathrm{SN}\rangle$ (i.e only reduction is quantum). Note that wrapping only the second or third character is disallowed (e.g. $\mathrm{B}\,|\mathrm{S}\rangle\,\mathrm{N}$ and $\mathrm{BS}\,|\mathrm{N}\rangle$), as both are related to the reduction and must be quantum/classical 'together'.

(3) Superscripts and subscripts. Since the construction and reduction can be quantum algorithms, they can potentially make superposition queries to the adversary or the implementation of the assumption. To clarify this, we introduce superscripts 'q' or 'c' for each character. On the other hand, the queried party can be implemented either quantumly or classically, as specified by a subscript 'q' or 'c' on each character. We emphasize that there could be cases that a component doesn't work properly when the queried party is quantum, as the queried party can be rewound in the

middle of a multi-round interaction. The subscript is useful for capturing such quantum issues. In our example, the superscript in $|\mathrm{B}^{c}_{c}\mathrm{SN}\rangle$ indicates that the construction makes only classical queries to the implementation of $\mathcal{Q}$; while the subscript in $|\mathrm{B}^{c}_{c}\mathrm{SN}\rangle$ shows that the implementation of $\mathcal{Q}$ is a classical algorithm, as the construction might fail if the underlying implementation were quantum. For simplicity, we set 'quantum' as the default and omit the notation (i.e. $|\mathrm{B}^{c}_{c}\mathrm{S}^{\cdot}_{\cdot}\mathrm{N}^{\cdot}_{\cdot}\rangle$).

Remark 3. For the case of a classical construction $\mathrm{B}\,|\mathrm{BB}\rangle$ or a classical reduction $|\mathrm{B}\rangle\,\mathrm{BB}$, we clearly do not need to specify classical or quantum access corresponding to the classical parts. However, we still allow different subscripts. This is because a classical construction may or may not work based on whether its building block is implemented classically or quantumly.

Remark 4. We do not consider potential issues regarding obtaining superposition access of classical adversaries and implementations in this work, and assume that the reductions and constructions are granted such superposition access.

Formal definitions We will now present formal definitions of some quantum reductions. Although this list is far from being exhaustive, we believe it captures the essence of the framework and that the other definitions can be inferred from these ones.

Definition 13 (**$(\mathcal{P}\to\mathcal{Q})$−$|\mathrm{N}^{c}\mathrm{BN}^{c}\rangle$ Reduction).** *There exists an $|\mathrm{N}^{c}\mathrm{BN}^{c}\rangle$ reduction from $\mathcal{P}$ to $\mathcal{Q}$ if for every $f\in\mathcal{F}_{\mathcal{Q}}$, there exists an efficient quantum co-strategy G s.t.:*
Correctness: *It holds that $\mathsf{G}^{f}\in\mathcal{F}_{\mathcal{P}}$.*
Security: *There exists an efficient quantum co-strategy R such that for every strategy $\mathcal{A}$, if $(\mathsf{G}^{f},\mathcal{A}^{|f\rangle})\in\mathcal{R}_{\mathcal{P}}$, then $(f,\mathsf{R}^{|\mathcal{A}\rangle,f})\in\mathcal{R}_{\mathcal{Q}}$.*

Note that in the above definition, note that the notation $\mathsf{R}^{|\mathcal{A}\rangle,f}$ denotes that R obtains quantum access to $\mathcal{A}$ and classical access to f.

Definition 14 (**$(\mathcal{P}\to\mathcal{Q})$−$\mathrm{B}\,|\mathrm{B}_{c}\mathrm{B}_{c}\rangle$ Reduction).** *There exists a $\mathrm{B}\,|\mathrm{B}_{c}\mathrm{B}_{c}\rangle$ reduction from $\mathcal{P}$ to $\mathcal{Q}$ if there exists an efficient classical co-strategy G and an efficient quantum co-strategy R s.t.:*
Correctness: *For every $f\in\mathcal{F}_{\mathcal{Q}}$, it holds that $\mathsf{G}^{f}\in\mathcal{F}_{\mathcal{P}}$.*
Security: *For every classical strategy $f\in\mathcal{F}_{\mathcal{Q}}$ and every classical strategy $\mathcal{A}$, if $(\mathsf{G}^{f},\mathcal{A}^{|f\rangle})\in\mathcal{R}_{\mathcal{P}}$, then $(f,\mathsf{R}^{|A,f\rangle})\in\mathcal{R}_{\mathcal{Q}}$.*

In the above definition, the notation $\mathsf{R}^{|\mathcal{A},f\rangle}$ denotes that R obtains quantum access to both $\mathcal{A}$ and f.

To demonstrate the generality of the system, we define the contrived example from Fig. 2 formally:

Definition 15 ($(\mathcal{P} \to \mathcal{Q}) - |\mathrm{B}^c_c\mathrm{SN}\rangle$ Reduction). *There exists a $|\mathrm{B}^c_c\mathrm{SN}\rangle$ reduction from $\mathcal{P}$ to $\mathcal{Q}$ if there exists an efficient quantum co-strategy G such that:*
Correctness: *For every classical strategy $f \in \mathcal{F}_\mathcal{Q}$, we have $\mathsf{G}^f \in \mathcal{F}_\mathcal{P}$.*
Security: *For every quantum strategy $f \in \mathcal{F}_\mathcal{Q}$, there exists an efficient quantum co-strategy R such that for every unitary strategy $\mathcal{A}$, if $(\mathsf{G}^f, \mathcal{A}^{|f\rangle}) \in \mathcal{R}_\mathcal{P}$, then $(f, \mathsf{R}^{|\mathcal{A},\dagger,f\rangle}) \in \mathcal{R}_\mathcal{Q}$.*

The notation $\mathsf{R}^{|\mathcal{A},\dagger,f\rangle}$ denotes that R obtains quantum access to both $\mathcal{A}$ and f, but also obtains inverse access to $\mathcal{A}$.

Remark 5. Notice that in the above definition, there is no guarantee that $\mathsf{G}^f \in \mathcal{F}_\mathcal{P}$ for quantum f. But when this holds (and $\mathcal{A}$ is a valid attack), R breaks f.

3.3 Relations

Let us consider reductions of the form $|\mathrm{XYZ}\rangle$, where both the construction and reduction are quantum. We can categorize relations between such reductions into two broad classes. The first comprises relations between black-box and non-black-box variants; the second class includes relations between the quantum versions and their classical counterparts.

1. We use the notation $\mathrm{XYZ} \geq \widetilde{\mathrm{X}}\widetilde{\mathrm{Y}}\widetilde{\mathrm{Z}}$ to indicate that XYZ is greater than or equal to $\widetilde{\mathrm{X}}\widetilde{\mathrm{Y}}\widetilde{\mathrm{Z}}$ in each of the three positions, where $\mathrm{B} > \mathrm{N}$ and $\mathrm{B} > \mathrm{S}$. This convention is used as black-box reductions are more restricted than non-black-box ones, and hence are considered to demonstrate a "stronger" result. Notice that any reduction of the type $|\mathrm{XYZ}\rangle$ implies a reduction of the type $\left|\widetilde{\mathrm{X}}\widetilde{\mathrm{Y}}\widetilde{\mathrm{Z}}\right\rangle$, if $\mathrm{XYZ} \geq \widetilde{\mathrm{X}}\widetilde{\mathrm{Y}}\widetilde{\mathrm{Z}}$. This holds regardless of the 'c' superscripts and subscripts as long as they are the same for both reductions. This is because of the fact that black-box reductions are also non-black-box reductions but not the other way around.

   Remark 6. In the context of *strong-black-box* reductions, we cannot claim for example that a $|\mathrm{BSB}\rangle$ reduction implies a $|\mathrm{BNB}\rangle$ reduction. This is because we require the reduction in the former to work only for adversary's with unitary equivalents, while the latter works for general quantum strategies.

2. Another basic relation is that any reduction of the type $\left|\mathrm{X}^{e}\mathrm{Y}^{f}\mathrm{Z}^{g}\right\rangle$ implies a reduction of the type $|\mathrm{XYZ}\rangle$, where the e, f, g can be either 'c' (classical) or omitted (quantum). This is because classical black-box access is more restricted than quantum access. On the other hand, a reduction of the type $|\mathrm{XYZ}\rangle$ implies a reduction of the type $|\mathrm{X}_u\mathrm{Y}_v\mathrm{Z}_w\rangle$, where the u, v, w can either be 'c' (classical) or omitted (quantum). This is because quantum implementations are potentially problematic to work with due to issues like quantum rewinding. We can also observe relations between reductions which are fully quantum and reductions which are partly classical, in either the construction or the security reduction. Thus, a reduction of the form $\mathrm{X}\,|\mathrm{YZ}\rangle$ implies one of the form $|\mathrm{XYZ}\rangle$ as a classical construction is more restricted.

A few consequences easily follow. A $|\mathrm{BB}^{c}\mathrm{B}\rangle$ reduction implies a $|\mathrm{BNB}\rangle$ reduction. On the other hand, a $|\mathrm{BB}_{c}\mathrm{B}\rangle$ reduction may not imply a $|\mathrm{BBB}\rangle$ reduction, and may not even imply a $|\mathrm{BNB}\rangle$ reduction. For instance, the work of Lombardi et al. [LMQW22] constructs a PRF (among other primitives) based on an interactive proof of quantumness by Brakerski et al. [BCM+21]. The construction is classically-secure but quantumly-insecure under the LWE assumption, due to the difficulty of quantum rewinding. In general, it would be interesting to study such counterexamples to prove these implications are strict. We now illustrate some relations in Fig. 3.

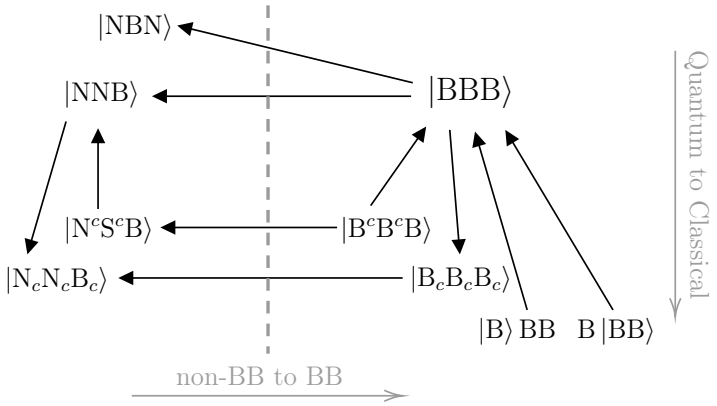

Fig. 3. Q-CAP Relations

3.4 Categorizing Quantum Reductions using Q-CAP

Next we put the Q-CAP system in practice and showcase the relationships between a host of cryptographic primitives in a quantum world (Cf. summary in Fig. 4). By this effort, we can identify clearly the assumptions and qualifications of various results on a common scale, which as a consequence reveals the possible routes of improving them, such as weakening the assumptions or security reductions in feasibility results or ruling out stronger forms of reductions in impossibility results. This also helps gain insights about the strengths and limitations of quantum information processing in cryptography.

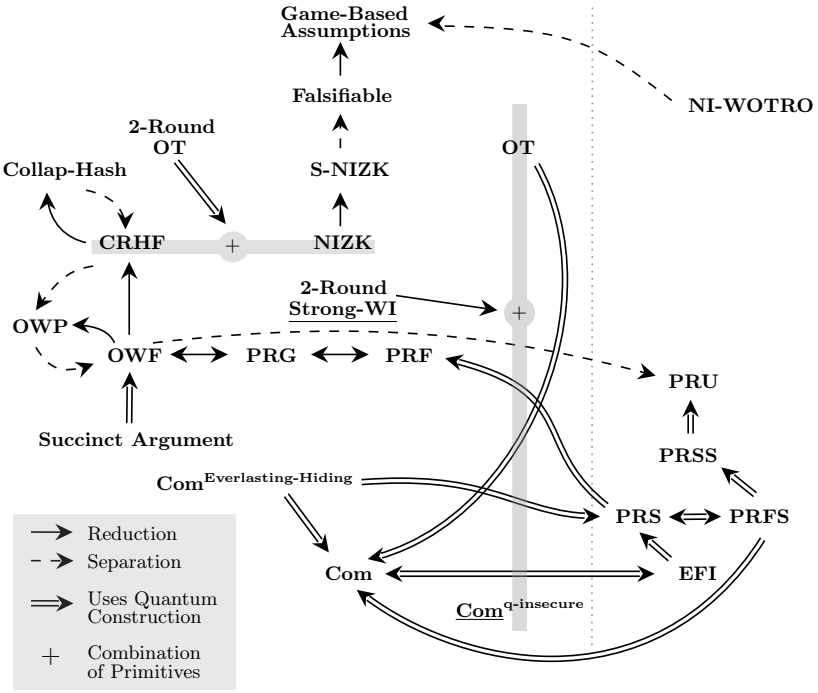

Fig. 4. Representative Reductions in a Quantum World

We consider security against superposition-query by default, e.g., for pseudo-random functions (PRFs) [Zha12]. We denote primitives which are only classically-secure with an underline as $\underline{\mathcal{P}}$. Before proceeding, we informally review a few relatively new primitives depicted in Fig. 4.

– *Weak One-Time Random Oracle (WOTRO) [DLS22].* A WOTRO is a restricted version of a random oracle but remains sufficient for the

Fiat-Shamir transform of Σ-protocols. It is shown in the shared EPR pairs model that non-interactive WOTRO (NI-WOTRO) cannot be based on any game-based cryptographic assumptions.
- *Pseudo-Random State Scramblers (PRSS) [LQS+23].* A PRSS transforms an *arbitrary* initial state into a pseudo-random state. Their work constructs PRSS based on OWFs. PRSSs are implied by pseudo-random unitaries (PRUs), and by the work of Kretschmer [Kre21], are known to be oracle separated from OWFs.
- *Commitments with certified everlasting hiding [BK23].* Such commitments have the additional property that a receiver can provably delete its committed information before receiving an opening. Bartusek et al. [BK23] showed that such commitments could be obtained from standard commitments with the help of quantum information.
- EFI pairs were introduced in the work of Brakerski et al. [BCQ23] as a candidate for the minimal assumption required for computational quantum cryptography. They show that many primitives in the quantum setting imply EFI.

Quantum Security of Classical Primitives. Establishing security of classical cryptographic constructions against quantum adversaries is an important endeavor. We discuss two basic results.

i. **(PRF $\rightarrow$ PRG)—$\mathbf{B}\,|\mathbf{BB}\rangle$** (Zhandry [Zha12]). This reduction utilizes the classical GGM construction [GGM86]. The security reduction needs to simulate the responses for the superposition queries made by the PRF adversary. Hence, it needs to evaluate the PRG on a superposition of different seeds.
ii. **(Succinct Arguments $\rightarrow$ Collapsing Hash)—$\mathbf{B}\,|\mathbf{SB}\rangle$** (Chiesa et al. [CMSZ22]). This proves quantum security of Killian's [Kil92] succinct arguments for NP based on collapsing hash functions. However, due to quantum rewinding, the security reduction here uses strong-black-box quantum access. Since the adversary evaluates the hash function in superposition to break the collapsing property, the reduction needs quantum access to it as well. Classically Killian's result establishes a BBB reduction from succinct arguments to CRHFs. Hence the characterization here reveals an open question of making the reduction fully black-box.

Quantum Advantages. The quantum capability can sometimes help realize what is difficult or even impossible classically. The examples below illustrate different flavors of harnessing the quantum power.

i. **(Succinct Arguments → OWF)**—$|\mathbf{BSB}\rangle$(Gunn et al. [GJMZ23]). This work builds succinct arguments for NP based on OWFs by utilizing quantum communication. Moreover, the protocol runs in three rounds only, whereas the best known classical result requires four rounds [Kil92, CMSZ22]. The protocol is built using a succinct quantum commitment scheme. This succinct commitment involves performing a quantum pseudo-one-time-pad operation in superposition, and hence needs quantum access to the PRG. Once again, strong-black-box access to the adversary is utilized for rewinding.

ii. **(OT → OWF)**—$|\mathbf{N}^{\mathbf{c}}_{\boldsymbol{c}}\mathbf{B}\mathbf{N}^{\boldsymbol{c}}_{\boldsymbol{c}}\rangle$ (Grilo et al. [GLSV21]). The construction here is a protocol for oblivious transfer (OT) that requires quantum communication. It makes non-black-box use of a bit commitment scheme due to the use of zero-knowledge proofs specific to the scheme. Since the ZK proof used is one for NP statements, the reduction only works for classical implementations of the OWF. In contrast, in the classical world, no construction of OT from OWFs is known. Moreover, the work of Impagliazzo and Rudich [IR89] rules out BNN reductions from OT to OWFs. We would like to mention that the work of Bartusek et al. [BCKM21] shows a different reduction that makes black-box use of a commitment scheme to obtain OT.

iii. **(2-Round OT → NIZK + CRHF)**—$|\mathbf{N}_{\boldsymbol{c}}\mathbf{B}\mathbf{N}_{\boldsymbol{c}}\rangle$ (Colisson et al. [CMS23]). Here, a 2-round (round-optimal) quantum protocol for OT is shown using NIZK and CRHFs (with an additional hiding property). Since NIZK can be obtained in the random oracle (RO) model, this provides a construction of OT from minicrypt primitives like the previous example. The protocol makes non-black-box use of the hash function for the NIZK proof, and also runs the hash function in superposition. As mentioned in the work, such use of the CRHF is not desirable for practical purposes, and it is open if this is necessary. The use of the NIZK proof for NP implies that the reduction only works for classical implementations of the hash function.

Remark 7. Even when a construction or reduction obtains the code of a primitive, we view direct executions of the primitive as black-box access. Hence, the absence of a superscript on the third 'N' denotes that this black-box access is quantum.

iv. **(2-Round Strong-WI → OT + $\underline{\mathbf{Com}}^{\mathbf{q\text{-}insecure}}$)**—$\mathbf{B}\,|\mathbf{B}^{\boldsymbol{c}}_{\boldsymbol{c}}\mathbf{B}^{\boldsymbol{c}}\rangle$(Kalai and Khurana [KK19]). This work demonstrates a novel quantum advantage in the context of complexity leveraging, which enables constructing *classically secure primitives by quantum reductions*. It shows

a classically-secure 2-Round Strong Witness Indistinguishable (Strong-WI) Argument based on quantum-secure OT and classical-secure yet quantum-insecure commitments. The construction is classical, but the security reduction is quantum, and exploits the ability to break the quantum-insecure commitment scheme. In contrast, classical NBN reductions are ruled out [Kiy21].

Quantum Black-Box Separations.

i. **(OWP $\not\to$ OWF)**—**$\mathbf{B_c\,|B_cB_c\rangle}$**(Cao and Xue [CX21]): This separation result rules out reductions comprising of a classical construction and a quantum security reduction. The reduction may obtain quantum access to the adversary and the OWP. In another work by Chung et al. [CLM23], a conditional separation also ruled out quantum constructions, and hence $|\mathrm{B}_c\mathrm{B}_c\mathrm{B}_c\rangle$ reductions.
ii. **(CRHF $\not\to$ OWP)**—**$\mathbf{|B_cB_cB_c\rangle}$** (Hosoyamada and Yamakawa [HY20]): This result rules out fully-black-box reductions of the kind $|\mathrm{B}_c\mathrm{B}_c\mathrm{B}_c\rangle$ from CRHFs to OWPs. Although their work considers quantum security, it is easy to see that their result also rules out classically-secure CRHFs using quantum reductions. Intuitively, this is because the breaking oracle need only be accessed classically to break the CRHF. In the classical world however, a stronger BNN separation is known [Sim98].

Both examples allude to the possibility of stronger separations, and it is interesting to settle down whether that can be achieved.

Quantum Primitives. We now turn to studying some relations involving quantum primitives.

i. **(OWF $\not\to$ PRU)**—**$\mathbf{|BNN\rangle}$** (Kretschmer [Kre21]). The work of Kretschmer [Kre21] shows a quantum oracle relative to which PRUs exist but OWFs do not. As such, it rules out relativizing reductions from OWFs to PRUs. In the work of Reingold et al. [RTV04], it was shown that ruling out relativizing reductions also rules out BNN reductions. It is not hard to see that an analogous claim extends to the quantum case, which in-turn rules out $|\mathrm{BNN}\rangle$ reductions.
ii. **(PRFS $\to$ PRS)**—**$\mathbf{|B^cBB^c\rangle}$** (Ananth et al. [AQY22]). This construction of pseudorandom function-like states involves classical black-box access to the PRS, and is an example of a reduction between two quantum primitives. The security reduction uses classical black-box access

to the PRS and obtains its quantum state outputs. It then performs a measurement on these states and evaluates the PRFS adversary on the post-measurement state.

iii. **(Com $\rightarrow$ PRFS)—$|\mathbf{S^cBS^c}\rangle$** (Ananth et al. [AQY22]). This is an example of a reduction between a classical and a quantum primitive. The bit commitment scheme shown in this work utilizes a procedure to verify the validity of a PRFS output. The verification procedure relies on both the generator of the PRFS and its inverse, i.e., it requires strong-black-box access.

4 Impossibility of S-NIZK in the Quantum Setting

In this section, we will prove our impossibility result that corresponds to a setting where the parties can compute and communicate quantumly. The result shows that quantum black-box reductions cannot be used to prove the security of an S-NIZK protocol for an NP-complete language based on any falsifiable assumption. More specifically, we require the language to satisfy the hard language definition (Definition 1), which holds for any NP-complete language assuming the existence of post-quantum OWFs. Essentially, the result shows that if such a black-box reduction were to exist, then the falsifiable assumption must be false. However, if the falsifiable assumption is false, then a security reduction for the S-NIZK based on the assumption being true is meaningless. The security reduction we refer to here corresponds to security against a malicious prover, i.e., adaptive soundness (Definition 3).

On Adaptive Soundness and Zero-Knowledge: As in the work of Pass, we consider adaptive notions of both soundness and zero-knowledge. In the case of *non-adaptive* soundness, the malicious prover is restricted to choose a statement before the CRS is initialized. Previous works have constructed non-adaptive sound S-NIZKs based on variants of LWE [CCH$^+$19, PS19], and are also known to satisfy post-quantum security [CVZ20]. Interestingly, there are two important variants of adaptive soundness, namely the *penalizing* variant and the weaker *exclusive* one. The one considered in the work of Pass and ours is the penalizing one. The exclusive one is weaker in that the malicious prover is guaranteed to choose false statements, i.e., provers that may sometimes output true statements are not considered in the definition. In the work of Fischlin et al. [FR21], these notions are explored in detail. Their work also mentions that the aforementioned non-adaptive S-NIZK protocols are adaptively sound by

the exclusive definition. Later, we will discuss why the proof breaks down for the exclusive definition.

Notice the impossibility also requires the S-ZK property to be adaptive. We will elaborate later that the two-part simulator is important for the impossibility. However, this could be an artifact of the proof technique as no adaptive sound and non-adaptive S-ZK NIZK is known. We now state the main theorem:

Theorem 2. *Let Π be a non-interactive quantum protocol satisfying the completeness and adaptive statistical zero-knowledge properties for an* NP*-complete language $\mathcal{L}$. Let* B *be a quantum black-box reduction such that, for every attacker $\mathcal{A}$ that breaks the adaptive soundness of Π, $\mathsf{B}^{\mathcal{A}}$ breaks some quantum falsifiable assumption (C, t). Then, the assumption (C, t) is false, assuming the existence of post-quantum secure one-way functions with respect to quantum advice.*

Proof Overview. Our strategy follows along the lines of the classical proof by Pass, which follows the meta-reduction paradigm. This involves showing the existence of an (inefficient) attacker $\mathcal{A}$ that breaks the security of the construction, along with an entity called the emulator $\mathcal{E}$. Unlike $\mathcal{A}$, $\mathcal{E}$ is efficient but does not break the security of the construction. However, $\mathcal{E}$ "mimics" $\mathcal{A}$ so that no QPT algorithm can tell them apart. Consider a black-box reduction that transforms $\mathcal{A}$ into an attack on a falsifiable assumption. When we replace $\mathcal{A}$ with $\mathcal{E}$, the new combination efficiently breaks the assumption. If not, the reduction, in conjunction with the challenger, would be able to distinguish between $\mathcal{A}$ and $\mathcal{E}$, a contradiction. Consequently, our proof can be broken down into two sub-parts:

1. **Existence of $\mathcal{A}$ and $\mathcal{E}$.** At a high level, the emulator $\mathcal{E}$ runs the honest prover's algorithm of the S-NIZK, on a true statement of a hard language. The corresponding attacker $\mathcal{A}$ makes use of the two-part zero-knowledge simulator $(\mathsf{S}_1, \mathsf{S}_2)$, and crucially relies on the *statistical* zero-knowledge property to generate accepting proofs for false statements. We observe that similar constructions of $\mathcal{A}$ and $\mathcal{E}$ as that of previous work suffice for the quantum setting. More precisely, these two algorithms are shown to be computationally indistinguishable by a reduction making a *single* CRS query.

2. **Indistinguishability of $\mathcal{A}$ and $\mathcal{E}$ with (multi-query) oracle access.** To prove this, the work of Pass reduces oracle indistinguishability to single-query indistinguishability, a method also employed in

other classical meta reduction impossibility results [GW11, CLMP12, BDSG+13]. However, the reduction argument does not directly extend to the quantum setting. This is because the classical argument involves programming the random tape of a classical algorithm, which does not have a quantum analogue. To circumvent this issue, we provide an alternate argument, elaborated below.

Classically, the reduction executes an oracle distinguisher which expects responses from an inefficient oracle. Since the inefficient oracle cannot be efficiently simulated by the reduction, it is provided with some optimal random coins as advice. More precisely, these random coins include the optimal random tape of the distinguisher, and the corresponding optimal responses of the inefficient oracle. The challenge in the quantum setting is that the distinguisher cannot be forced to run on a specific "random tape", which would determine the queries it makes. However, we realize that it is not necessary to explicitly force the distinguisher to output specific queries and then simulate the responses. Instead, we utilize quantum advice to provide the reduction with an optimal internal quantum state that aligns with the best measurement outcomes and responses from the inefficient oracle. The reduction can then straightforwardly execute the oracle distinguisher in a non-black-box way.

4.1 Proof of Theorem 2

Constructions of the Attacker and the Emulator. As noted previously, we first provide the constructions of the attacker $\mathcal{A}$ and the emulator $\mathcal{E}$. Given a crs as input, the emulator runs as follows:

Emulator $\mathcal{E}(1^\lambda, \mathsf{crs})$:

1. $(x, w) \leftarrow Z_{\mathcal{L}}$, where $Z_{\mathcal{L}}$ is defined as in Definition 1.
2. $|\pi\rangle \leftarrow \mathsf{P}(1^\lambda, \mathsf{crs}, x, w)$
3. Output $(x, |\pi\rangle)$

It is easy to see that $\mathcal{E}$ is efficient, as it samples from the efficiently sampleable distribution $Z_{\mathcal{L}}$, and executes the honest prover of the S-NIZK protocol Π. However, it is clearly not a valid attack against the adaptive soundness of Π, as it only outputs accepting proofs of true statements.

We now turn to describing the corresponding inefficient attacker that $\mathcal{E}$ is supposed to emulate:

Attacker $\mathcal{A}(1^\lambda, \mathsf{crs})$:

1. If $\mathsf{corner\text{-}case}(\mathsf{crs}) = 1$, output $(x, |\pi\rangle) \leftarrow \mathcal{E}(1^\lambda, \mathsf{crs})$;
   Else, proceed to the next step.
2. While $\mathsf{crs}' \neq \mathsf{crs}$:
   - Generate non-uniform advice $|\phi\rangle$.
   - Run $(\mathsf{crs}', |\mathsf{aux}\rangle) \leftarrow \mathsf{S}_1(1^\lambda, |\phi\rangle)$.
   - If unsuccessful after 2^{2^λ} attempts, output $\perp$.
3. $\widetilde{x} \leftarrow \widetilde{X}_{\mathcal{L}}$, where $\widetilde{X}_{\mathcal{L}}$ is defined as in Definition 1.
4. $|\widetilde{\pi}\rangle \leftarrow \mathsf{S}_2(1^\lambda, |\mathsf{aux}\rangle, \widetilde{x})$
5. Output $(\widetilde{x}, |\widetilde{\pi}\rangle)$

where $\mathsf{corner\text{-}case}$ is defined as follows:

$$\mathsf{corner\text{-}case}(\mathsf{crs}) = \left\{ \begin{matrix} 1, \text{ when } \mathsf{TD}\Big(\mathcal{A}'_0(1^\lambda, \mathsf{crs}), \mathcal{E}(1^\lambda, \mathsf{crs})\Big) = \mathsf{non\text{-}negl}(\lambda) \\ 0, \text{ otherwise} \end{matrix} \right\}$$

where $\mathcal{A}'_0$ is an algorithm defined as follows:

Algorithm $\mathcal{A}'_0(1^\lambda, \mathsf{crs})$:

1. While $\mathsf{crs}' \neq \mathsf{crs}$:
   - Generate non-uniform advice $|\phi\rangle$.
   - Run $(\mathsf{crs}', |\mathsf{aux}\rangle) \leftarrow \mathsf{S}_1(1^\lambda, |\phi\rangle)$.
   - If unsuccessful after 2^{2^λ} iterations, output $\perp$.
2. $(x, w) \leftarrow Z_{\mathcal{L}}$, where $Z_{\mathcal{L}}$ is defined as in Definition 1.
3. $|\widetilde{\pi}\rangle \leftarrow \mathsf{S}_2(1^\lambda, |\mathsf{aux}\rangle, x)$
4. Output $(x, |\widetilde{\pi}\rangle)$

Remark 8. In the exclusive variant of adaptive soundness, $\mathcal{A}$ cannot execute $\mathcal{E}$ on "bad" crs, as it is forced to always output false statements.

Remark 9. Since $|\phi\rangle$ is a non-uniform advice state, there exists a quantum circuit for every security parameter λ that generates it. Consequently, an inefficient algorithm like $\mathcal{A}$ or $\mathcal{A}'_0$ can compute such a state uniformly.

We will first describe $\mathcal{A}$ while ignoring the details of the $\mathsf{corner\text{-}case}$ function. The attacker $\mathcal{A}$ first brute-forces through S_1, i.e., repeatedly runs it until its output crs' matches the input crs. This is done so that S_2 can later produce a proof consistent with crs using the corresponding auxiliary state $|\mathsf{aux}\rangle$. Notice that such a search is inefficient as double-exponentially many steps are necessary for locating a particular crs with high probability. Moreover, its initial advice state $|\phi\rangle$ may also be inefficient to produce.

The rest of the steps are straightforward. It samples a false statement from $\widetilde{X}_{\mathcal{L}}$ and then utilizes S_2 to generate a proof.

The corner-case function outputs 1 on "bad" crs strings, for which the S-ZK property does not hold when conditioned on the crs. This is indirectly captured by the trace-distance between $\mathcal{E}$ and $\mathcal{A}'_0$, where $\mathcal{A}'_0$ is similar to $\mathcal{A}$ except the corner-case check and that it samples true statements instead of false ones. Since the reduction does not have to sample crs from Setup, it can query its oracle on such a "bad" crs to distinguish $\mathcal{A}$ from $\mathcal{E}$. Consequently, the attacker $\mathcal{A}$ switches to executing the emulator $\mathcal{E}$ when queried on such a crs. In order to do this, $\mathcal{A}$ estimates the value $\mathsf{TD}\Big(\mathcal{A}'_0(1^\lambda, \mathsf{crs}), \mathcal{E}(1^\lambda, \mathsf{crs})\Big)$ for each crs string by running $\mathcal{A}'_0$ and $\mathcal{E}$ several times and executing a tomography procedure. This is possible in unbounded but finite time as there are a bounded number of possible crs strings. Later on, we will argue that for a crs sampled from Setup, the probability that it is "bad" is negligibly small, which follows from the S-ZK property. Hence, $\mathcal{A}$ still chooses false statements with high probability and breaks adaptive soundness.

Indistinguishability of $\mathcal{A}$ and $\mathcal{E}$: We first show the single query case – no QPT algorithm M making one query, crs, to $\mathcal{A}$ and $\mathcal{E}$ can distinguish between them. This will in-turn be used to establish that $\mathcal{A}$ breaks adaptive soundness, and later to show the multi-query indistinguishability.

Lemma 1 (Single-query Indistinguishability). *Let M be a QPT algorithm that makes a single query to its oracle. Then, M cannot distinguish between the oracles $\mathcal{A}$ and $\mathcal{E}$ except with negligible probability, i.e.,*

$$\left|\Pr\left[\mathsf{M}^{\mathcal{A}}(1^\lambda) = 1\right] - \Pr\left[\mathsf{M}^{\mathcal{E}}(1^\lambda) = 1\right]\right| \leq \mathsf{negl}(\lambda)$$

Proof. We will construct a hybrid algorithm $\mathcal{A}'$ that works similar to the attacker $\mathcal{A}$. Unlike $\mathcal{A}$ which chooses false statements from the hard language, $\mathcal{A}'$ picks true instances, as the emulator $\mathcal{E}$ does.

Hybrid Attacker $\mathcal{A}'(1^\lambda, \mathsf{crs})$:

1. If $\mathsf{corner\text{-}case}(\mathsf{crs}) = 1$, output $(x, |\pi\rangle) \leftarrow \mathcal{E}(1^\lambda, \mathsf{crs})$;
   Else, proceed to the next step.

2. Run Algorithm $\mathcal{A}'_0(1^\lambda, \mathsf{crs})$.

First, we show that $\mathcal{A}'$ and $\mathcal{E}$ are indistinguishable to M. This holds because $\mathcal{A}'$ runs $\mathcal{E}$ when $\mathsf{corner\text{-}case}$ outputs 1, and runs $\mathcal{A}'_0$ otherwise. In the latter case, it is guaranteed that the output of $\mathcal{A}'_0$ is indistinguishable from that of $\mathcal{E}$ by the definition of $\mathsf{corner\text{-}case}$. Consequently, we have the following:

$$\left|\Pr\left[\mathsf{M}^{\mathcal{A}'}(1^\lambda)=1\right]-\Pr\left[\mathsf{M}^{\mathcal{E}}(1^\lambda)=1\right]\right|\leq \mathsf{negl}(\lambda) \tag{1}$$

It is now left to show that $\mathcal{A}$ and $\mathcal{A}'$ are also indistinguishable to M. The sole distinction between these two is that one operates on false statements while the other on true instances. Suppose that M distinguishes between them with $\mathsf{non\text{-}negl}(\lambda)$ advantage. We will now construct a non-uniform QPT distinguisher D that makes use of M in a non-black-box way to break the language hardness property.

Remark 10. The constructed distinguisher D will run the zero-knowledge simulator S_2 in order to break the hardness assumption. Consequently, the impossibility does not hold for unbounded simulators.

Consider M to consist of two parts: (1) Make one query crs to the oracle $\mathcal{A}'$ or $\mathcal{A}$; (2) Output a bit b upon receiving the response $(x_{\mathsf{crs}}, |\pi\rangle_{\mathsf{crs}})$.

Let $|\psi\rangle$ be an optimal internal state of M after part (1). Let M' denote a QPT algorithm that takes the response $(x_{\mathsf{crs}}, |\pi\rangle_{\mathsf{crs}})$ along with $|\psi\rangle$ and outputs b, winning with the same probability as M.

Let y be a statement sampled from $X_{\mathcal{L}}$ or $\widetilde{X}_{\mathcal{L}}$ with probability $1/2$ each. To preserve the efficiency of D, we allow D to have the state $|\mathsf{aux}\rangle$ (the internal state output by S_1 along with crs) and $|\psi\rangle$ as non-uniform quantum advice that corresponds to the crs query made by M.

Language Distinguisher $\mathsf{D}(1^\lambda, y, \mathsf{crs}, |\mathsf{aux}\rangle, |\psi\rangle)$

1. Execute $|\pi_y\rangle \leftarrow \mathsf{S}_2(1^\lambda, |\mathsf{aux}\rangle, y)$
2. Run M' on the response $(y, |\pi_y\rangle)$ using the initial state $|\psi\rangle$; and receive a bit b.
3. Output b.

Observe that D simulates the response of either $\mathcal{A}'$ or $\mathcal{A}$ based on whether y is in $\mathcal{L}$ or not. Thus, it wins the distinguishing game with non-negligible advantage, contradicting the language hardness property. Hence, we have that:

$$\left|\Pr\left[\mathsf{M}^{\mathcal{A}'}(1^\lambda)=1\right]-\Pr\left[\mathsf{M}^{\mathcal{A}}(1^\lambda)=1\right]\right|\leq \mathsf{negl}(\lambda) \tag{2}$$

Remark 11. In the case of *non-adaptive* ZK, there exists a single-part simulator S that takes an input x and outputs $(\mathsf{crs},|\pi\rangle)$. Consequently, the string crs may depend on the input x, which is why it cannot be provided as advice to D.

From Equations (1) and (2), we have that:

$$\left|\Pr\left[\mathsf{M}^{\mathcal{A}}(1^\lambda)=1\right]-\Pr\left[\mathsf{M}^{\mathcal{E}}(1^\lambda)=1\right]\right|\leq \mathsf{negl}(\lambda) \tag{3}$$

□

Now, in order to prove that $\mathcal{A}$ breaks adaptive soundness, we first need the following lemma, which shows $\mathsf{corner\text{-}case}$ outputs 1 with $\mathsf{negl}(\lambda)$ probability for a randomly chosen crs.

Lemma 2 (Corner Case).

$$\Pr\left[\mathit{corner\text{-}case}(\mathsf{crs})=1 \;\middle|\; \mathsf{crs}\leftarrow\mathsf{Setup}(1^\lambda)\right]\leq\mathsf{negl}(\lambda)$$

Proof. For the sake of contradiction, suppose that $\mathsf{TD}(\mathcal{A}'_0(1^\lambda,\mathsf{crs}),\mathcal{E}(1^\lambda,\mathsf{crs}))>\frac{1}{p(\lambda)}$ with probability $\frac{1}{p(\lambda)}$, for some polynomial p and infinitely many λ. Then, the trace distance between the following distributions is at least $(\frac{1}{p(\lambda)})^2$:

$$\left\{(\mathsf{crs},x,|\pi\rangle)\;\middle|\;\mathsf{crs}\leftarrow\mathsf{Setup}(1^\lambda);(x,w)\leftarrow Z_{\mathcal{L}};|\pi\rangle\leftarrow\mathsf{P}(1^\lambda,x,w,\mathsf{crs})\right\}$$

$$\left\{(\mathsf{crs},x,|\pi\rangle)\;\middle|\;\mathsf{crs}\leftarrow\mathsf{Setup}(1^\lambda);(x,|\pi\rangle)\leftarrow\mathcal{A}'_0(1^\lambda,\mathsf{crs})\right\}$$

By the statistical zero-knowledge property, the latter distribution is statistically-close to the following:

$$\left\{(\mathsf{crs},x,|\pi\rangle)\;\middle|\;(\mathsf{crs},|\mathsf{aux}\rangle)\leftarrow\mathsf{S}_1(1^\lambda);(x,|\pi\rangle)\leftarrow\mathcal{A}'_0(1^\lambda,\mathsf{crs})\right\}$$

As $\mathcal{A}'_0$ runs S_1 doubly-exponentially many times in order to output the same crs, this is in-turn statistically-close to the following:

$$\left\{(\mathsf{crs}, x, |\pi\rangle) \;\middle|\; (\mathsf{crs}, |\mathsf{aux}\rangle) \leftarrow \mathsf{S}_1(1^\lambda); (x, w) \leftarrow Z_{\mathcal{L}}; |\pi\rangle \leftarrow \mathsf{S}_2(1^\lambda, |\mathsf{aux}\rangle, x)\right\}$$

This contradicts the statistical zero-knowledge property, proving the claim. □

We will now use Lemmas 1 and 2 to show $\mathcal{A}$ is indeed a valid attack, i.e., that it successfully breaks the adaptive soundness of the S-NIZK protocol Π.

Lemma 3 (Valid Attack). *The algorithm $\mathcal{A}$ breaks the adaptive soundness of the protocol Π, i.e., it satisfies the following for some polynomial $p(\lambda)$ and infinitely many $\lambda \in \mathbb{N}$.*

$$\Pr\left[\mathsf{V}\big(1^\lambda, \mathsf{crs}, \widetilde{x}, |\widetilde{\pi}\rangle\big) = 1 \;\wedge\; \widetilde{x} \notin \mathcal{L} \;\middle|\; \begin{array}{c}\mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda) \\ (\widetilde{x}, |\widetilde{\pi}\rangle) \leftarrow \mathcal{A}(1^\lambda, \mathsf{crs})\end{array}\right] \geq \frac{1}{p(\lambda)}$$

Proof. From the completeness property of Π, we have the following:

$$\Pr\left[\mathsf{V}\big(1^\lambda, \mathsf{crs}, x, |\pi\rangle\big) = 1 \;\middle|\; \begin{array}{c}\mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda) \\ (x, |\pi\rangle) \leftarrow \mathcal{E}(1^\lambda, \mathsf{crs})\end{array}\right] \geq 1 - \mathsf{negl}(\lambda)$$

This is because $\mathcal{E}$ samples a true statement x and runs the honest prover algorithm P to generate the proof $|\pi\rangle$. Consider now the following:

$$\Pr\left[\mathsf{V}\big(1^\lambda, \mathsf{crs}, \widetilde{x}, |\widetilde{\pi}\rangle\big) = 1 \;\wedge\; \widetilde{x} \notin \mathcal{L} \;\middle|\; \begin{array}{c}\mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda) \\ (\widetilde{x}, |\widetilde{\pi}\rangle) \leftarrow \mathcal{A}(1^\lambda, \mathsf{crs})\end{array}\right]$$

$$= \Pr\left[\mathsf{V}\big(1^\lambda, \mathsf{crs}, \widetilde{x}, |\widetilde{\pi}\rangle\big) = 1 \;\wedge\; \mathsf{corner\text{-}case}(\mathsf{crs}) = 0 \;\middle|\; \begin{array}{c}\mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda) \\ (\widetilde{x}, |\widetilde{\pi}\rangle) \leftarrow \mathcal{A}(1^\lambda, \mathsf{crs})\end{array}\right]$$

$$= \Pr\left[\mathsf{V}\big(1^\lambda, \mathsf{crs}, \widetilde{x}, |\widetilde{\pi}\rangle\big) = 1 \;\middle|\; \begin{array}{c}\mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda) \\ (\widetilde{x}, |\widetilde{\pi}\rangle) \leftarrow \mathcal{A}(1^\lambda, \mathsf{crs}) \\ \mathsf{corner\text{-}case}(\mathsf{crs}) = 0\end{array}\right]$$

$$\times \Pr\left[\mathsf{corner\text{-}case}(\mathsf{crs}) = 0 \;\middle|\; \mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda)\right]$$

$$= \left(\frac{1}{q(\lambda)}\right) \cdot (1 - \mathsf{negl}(\lambda)) \geq \frac{1}{p(\lambda)}$$

where $q(\lambda)$ and $p(\lambda)$ are some polynomials. The $\frac{1}{q(\lambda)}$ part follows from Equation 4.1, Lemma 1 and Lemma 2. This is because if V were to succeed in distinguishing $\mathcal{A}$ and $\mathcal{E}$ (conditioned on the high probability event $\mathsf{corner\text{-}case}(\mathsf{crs}) = 0$), then it can be used to break Lemma 1. The $(1 - \mathsf{negl}(\lambda))$ part follows directly from Lemma 2. Hence, $\mathcal{A}$ must break adaptive soundness with at least $\frac{1}{p(\lambda)}$ probability, which satisfies the breach of soundness requirement. □

Finally, we will prove by contradiction that no QPT algorithm M can distinguish between $\mathcal{A}$ and $\mathcal{E}$ when granted oracle access, meaning it is allowed to make polynomial number of queries. Specifically, if there is a QPT M succeeding in this setting, it also leads to an efficient quantum distinguisher D that succeeds in the single-query case, thereby violating Lemma 1. Formally, we will show:

Lemma 4 (Emulatable Attack). *Let M be a QPT algorithm. Then, M cannot distinguish between oracles $\mathcal{A}$ and $\mathcal{E}$, except with negligible probability, i.e.:*

$$\left|\Pr\left[\mathsf{M}^{\mathcal{A}}(1^\lambda) = 1\right] - \Pr\left[\mathsf{M}^{\mathcal{E}}(1^\lambda) = 1\right]\right| \leq \mathsf{negl}(\lambda)$$

Proof. Let M make $q = \mathsf{poly}(\lambda)$ queries to its oracle $\mathcal{O}$, which can be either $\mathcal{A}$ or $\mathcal{E}$. We will now consider $q+1$ hybrid oracles $\mathcal{H}_0, \ldots, \mathcal{H}_q$ defined as follows:

Hybrid Oracle $\mathcal{H}_i(1^\lambda, \mathsf{crs})$:

1. Answer the first i queries with the output of $\mathcal{E}$.
2. Answer the next $q - i$ queries with the output of $\mathcal{A}$.

Notice that the oracle $\mathcal{H}_0$ is the same as $\mathcal{A}$, and the oracle $\mathcal{H}_q$ is the same as $\mathcal{E}$. We will now demonstrate that for an arbitrary $j \in \{0, \ldots, q\}$, oracle access to $\mathcal{H}_j$ is indistinguishable from oracle access to $\mathcal{H}_{j+1}$. Using a standard hybrid argument, it then follows that oracle access to $\mathcal{A}$ is indistinguishable from oracle access to $\mathcal{E}$. First, we assume that M is a QPT algorithm that distinguishes $\mathcal{H}_j$ from $\mathcal{H}_{j+1}$ with non-negligible advantage, i.e.,

$$\mathsf{Adv}[\mathsf{M}] = \left|\Pr\left[\mathsf{M}^{\mathcal{H}_j}(1^\lambda) = 1\right] - \Pr\left[\mathsf{M}^{\mathcal{H}_{j+1}}(1^\lambda) = 1\right]\right| \geq \mathsf{non\text{-}negl}(\lambda)$$

We will now construct an efficient D that distinguishes between $\mathcal{A}$ and $\mathcal{E}$ with only one query, by using M as a subroutine. In order for D to use M, it must simulate the oracle responses from $\mathcal{H}_j$ (or $\mathcal{H}_{j+1}$) for the q queries made by M. All the responses for queries to $\mathcal{E}$ are straightforward to simulate by simply executing $\mathcal{E}$ itself. However, in the case of queries directed to $\mathcal{A}$, D cannot simulate the correct responses, as $\mathcal{A}$ is inefficient to execute. As for the $(j+1)$-th query made by M, D simply forwards it to its single-query oracle. Now, in order for D to be able to simulate the responses of $\mathcal{A}$, it must be provided some advice state. In previous classical works [GW11, CLMP12, Pas13, BDSG+13], D is provided with an optimal random tape of M, along with the corresponding optimal responses of $\mathcal{A}$. Consequently, D can run M on this random tape and then provide it with the correct responses of $\mathcal{A}$. However, this argument doesn't extend to the quantum setting, as it is unclear how to program the random tape of a quantum algorithm. To circumvent this issue, we provide D with an optimal internal state of M as quantum advice. We now proceed to explain this formally:

Recall that without loss of generality, the QPT algorithm M is of the form:

$$\mathsf{M}^{\mathcal{O}}(1^\lambda) = \mathcal{M} \circ \Big(\big(\mathcal{O} \circ \mathcal{M}_q\big) \circ \cdots \circ \big(\mathcal{O} \circ \mathcal{M}_1\big) \Big) \left|\psi^0\right\rangle$$

Let $\mathsf{M}_j^{\mathcal{O}}(1^\lambda)$ be an algorithm that acts on the internal state $\left|\psi^j\right\rangle$ as follows:

$$\mathsf{M}_j^{\mathcal{O}}(1^\lambda) = \mathcal{M} \circ \Big(\big(\mathcal{O} \circ \mathcal{M}_q\big) \circ \cdots \circ \big(\mathcal{O} \circ \mathcal{M}_{j+1}\big) \Big) \left|\psi^j\right\rangle$$

Notice that by an averaging argument, there must exist an optimal quantum state $\left|\psi_\star^j\right\rangle$ corresponding to the best measurement outcomes, such that the algorithm $\mathsf{M}_j^{\mathcal{O}}(1^\lambda)$ retains advantage $\mathsf{Adv}[\mathsf{M}]$. Furthermore, $\left|\psi_\star^j\right\rangle$ does not depend on the single-query oracle that the distinguisher D interacts with, as this oracle is not queried until the $(j+1)$-th query. Consequently, this state depends only on $\mathcal{A}$ and M (and the advice $\left|\psi^0\right\rangle$), and is thus a function of the security parameter. Moreover, as it is polynomially large, it can be provided as non-uniform quantum advice to D. Note that Lemma 1 holds against distinguishers with quantum advice. Hence, D simply runs the algorithm M_j instead of M. This is efficient because the $(j+1)$-th oracle query is answered by forwarding the response obtained by the single-query oracle of D, while the remaining responses corresponding to $\mathcal{E}$ are efficiently simulated. This gives us a distinguisher

for the single-query case (Lemma 1) with non-negligible advantage, which is a contradiction. □

We note that the aforementioned issue about simulating the inefficient oracle $\mathcal{A}$ does not arise in previous works [DLS22,DFG13]. This is because their single-query premise satisfies the stronger notion of statistical indistinguishability. This allows the multi-query reduction D to be inefficient, allowing it to run $\mathcal{A}$ by itself.

Finishing the Proof of Theorem 2. Finally, after having demonstrated how to construct $\mathcal{A}$ and $\mathcal{E}$, as well as their indistinguishability with oracle access, we conclude the proof. Recall that $\langle \mathsf{C}, \mathsf{B} \rangle(1^\lambda)$ denotes the output of the challenger C upon interaction with reduction B. As a direct consequence of Lemma 3 and the fact that B is a reduction for the assumption (C, t), we have that for some polynomial $p(\lambda)$ and infinitely many $\lambda \in \mathbb{N}$:

$$\Pr\left[\langle \mathsf{C}, \mathsf{B}^{\mathcal{A}} \rangle(1^\lambda) = 1\right] \geq t + \frac{1}{p(\lambda)}$$

From Lemma 4 and considering M to consist of both C and B, we have:

$$\left|\Pr\left[\langle \mathsf{C}, \mathsf{B}^{\mathcal{A}} \rangle(1^\lambda) = 1\right] - \Pr\left[\langle \mathsf{C}, \mathsf{B}^{\mathcal{E}} \rangle(1^\lambda) = 1\right]\right| \leq \mathsf{negl}(\lambda)$$

Now, we have that for some polynomial $q(\lambda)$ and infinitely many $\lambda \in \mathbb{N}$:

$$\Pr\left[\langle \mathsf{C}, \mathsf{B}^{\mathcal{E}} \rangle(1^\lambda) = 1\right] \geq t + \frac{1}{p(\lambda)} - \mathsf{negl}(\lambda) \geq t + \frac{1}{q(\lambda)}$$

Since B and $\mathcal{E}$ are both QPT algorithms, we have that $\mathsf{B}^{\mathcal{E}}$ is an efficient algorithm that breaks the assumption (C, t), which is a contradiction. □

4.2 Viewing the Impossibility via the Black-Box Framework

Consider the primitive $\mathsf{S\text{-}NIZK}$, which denotes an S-NIZK proof system for all of NP. Theorem 2 implies that there are no reductions of the type $|\mathrm{NB}^c\mathrm{N}\rangle$, $|\mathrm{N}_c\mathrm{B}^c\mathrm{N}\rangle$, $|\mathrm{NB}^c\mathrm{N}_c\rangle$ or $|\mathrm{N}_c\mathrm{B}^c\mathrm{N}_c\rangle$ from $\mathsf{S\text{-}NIZK}$ to any cryptographic primitive $\mathcal{Q}$ that has a falsifiable security game. The primitive $\mathcal{Q}$ may also be quantum (in which case only the first type of reductions exist anyway). This means that we consider quantum S-NIZK protocols (including communication) and restrict the quantum security reduction to classical black-box access to the quantum adversary. Notice that there is no restriction

on how the primitive is used, either by the construction or the reduction. Inspired by the meta-reduction definitions of Baecher et al. [BBF13], we show that our result exhibits an $((\mathsf{S\text{-}NIZK} \to \mathcal{Q}) - |\mathrm{N}_c\mathrm{B}^c\mathrm{N}_c\rangle) \to \mathcal{Q}$ meta-reduction, where $\mathcal{Q}$ is any primitive with a falsifiable security game. We now define this formally:

Definition 16 (**$((\mathcal{P} \to \mathcal{Q}) - |\mathrm{N}_c\mathrm{B}^c\mathrm{N}_c\rangle) \to \mathcal{Q}$ Meta-Reduction).** *For every classical implementation $g \in \mathcal{F}_{\mathcal{Q}}$, there exists a classical implementation $f \in \mathcal{F}_{\mathcal{Q}}$ such that the following holds:*
For every efficient construction (co-strategy) G and every efficient reduction (co-strategy) R, there exists an adversarial strategy $\mathcal{A}$ and an efficient ***meta-reduction*** *co-strategy M such that:*

$$\left(\left(\mathsf{G}^{|f\rangle}, \mathcal{A}^{|f\rangle}\right) \in \mathcal{R}_{\mathcal{P}} \implies \mathsf{R}^{\mathcal{A},|f\rangle} \in \mathcal{R}_{\mathcal{Q}}\right) \implies \left(g, \mathsf{M}^{|g\rangle}\right) \in \mathcal{R}_{\mathcal{Q}}$$

Notice that this only shows that every classical implementation of $\mathcal{Q}$ is broken when the reduction only works for classical implementations. However, this criterion is sufficient to rule out such reductions. Now, based on the above definition, we have the following as a corollary of Theorem 2. Here, the security reduction of the $\mathsf{S\text{-}NIZK}$ refers to the adaptive-soundness reduction.

Corollary 1. *There exists an $((\mathsf{S\text{-}NIZK} \to \mathcal{Q}) - |\mathrm{N}_c\mathrm{B}^c\mathrm{N}_c\rangle) \to \mathcal{Q}$ Meta-Reduction for every primitive $\mathcal{Q}$ that has a falsifiable security game. As a consequence, if the primitive $\mathcal{Q}$ exists, then there is no $|\mathrm{N}_c\mathrm{B}^c\mathrm{N}_c\rangle$ reduction from $\mathsf{S\text{-}NIZK}$ to $\mathcal{Q}$.*

We note that similar meta-reductions can be defined for the other types of reductions $|\mathrm{NB}^c\mathrm{N}\rangle$, $|\mathrm{NB}^c\mathrm{N}_c\rangle$ and $|\mathrm{N}_c\mathrm{B}^c\mathrm{N}\rangle$. Moreover, it is not hard to see that Theorem 2 implies such meta-reductions as well. This is because the technique used is oblivious to the workings of the construction. Likewise, no part of the proof relies on the reduction's use of the implementation.

4.3 On Reductions Making Quantum Queries

We will now highlight some challenges in lifting the impossibility result for reductions making quantum queries to the attacker. First, notice that the current attacker $\mathcal{A}$ expects a classical input crs. Then, it repeatedly executes S_1 until it outputs the same crs. Let us assume that the reduction R makes a quantum query $|q\rangle = \sum_i |\mathsf{crs}_i\rangle$, and in the case of the emulator $\mathcal{E}$,

obtains the superposition response $\sum_i |\mathsf{crs}_i\rangle |\pi_i\rangle$, where $|\pi_i\rangle = \mathcal{E}(1^\lambda, \mathsf{crs})$. In this case, if $\mathcal{A}$ observes/measures the query $|q\rangle$, it may destroy the superposition, making it easy to distinguish from $\mathcal{E}$. Consequently, we now provide an informal description of a new attacker, before discussing another problem.

The initial goal of the new inefficient attacker is to map $|q\rangle = \sum_i |\mathsf{crs}_i\rangle$ to a state $\sum_i |\mathsf{crs}_i\rangle |\mathsf{aux}_i\rangle$, such that every $(\mathsf{crs}_i, \mathsf{aux}_i)$ is a valid output of S_1. In order to do this, $\mathcal{A}$ repeatedly runs S_1 until it obtains each possible output crs_i sufficiently many (perhaps exponentially many) times. Then, it performs a tomography procedure to learn the corresponding quantum state $|\mathsf{aux}\rangle_i$ (which could be a mixed state in general). This way, $\mathcal{A}$ obtains a possibly inefficient channel that performs the desired mapping. It then executes S_2 on the state $\sum_i |\mathsf{crs}_i\rangle |\mathsf{aux}_i\rangle$ to obtain $\sum_i |\mathsf{crs}_i\rangle |\widetilde{\pi}_i\rangle$, where $|\widetilde{\pi}_i\rangle = \mathsf{S}_2(|\mathsf{aux}\rangle_i, \widetilde{x})$ for some false statement $\widetilde{x}$. Likewise, consider the state $\sum_i |\mathsf{crs}_i\rangle |\pi'_i\rangle$ where $|\pi'_i\rangle = \mathsf{S}_2(|\mathsf{aux}\rangle_i, x)$ for some true statement x, corresponding to the intermediate attacker $\mathcal{A}'$.

This brings us to the next challenge. What is the guarantee that the states $\sum_i |\mathsf{crs}_i\rangle |\pi_i\rangle$ and $\sum_i |\mathsf{crs}_i\rangle |\pi'_i\rangle$ are indistinguishable from one another? The statistical zero-knowledge property only guarantees that this indistinguishability holds for classical crs queries, and it is not clear if the quantum case can be reduced to the classical one. It might be possible that there exists some possibly contrived protocol $(\mathsf{Setup}, \mathsf{P}, \mathsf{V})$ for which the classical indistinguishability holds but not the quantum one. In the recent work of Abdolmaleki et al. [ACE+23], a quantum notion of zero-knowledge is explored where the distinguisher obtains superposition access to the simulator. However, the superposition is over statements, and not common reference strings. Moreover, it is not clear if a stronger notion of ZK that allows such superposition access with respect to the CRS is a meaningful one. We leave open the study of this notion for future work.

Acknowledgment. We would like to thank Fang Song for helpful discussions. CL and NP were supported by the US National Science Foundation grants CCF-2042414, CCF-2054758 (CAREER) and CCF-2224131.

References

Aar09. Scott Aaronson. Quantum copy-protection and quantum money. In *2009 24th Annual IEEE Conference on Computational Complexity*, pages 229–242. IEEE, 2009.

ABDS21. Gorjan Alagic, Zvika Brakerski, Yfke Dulek, and Christian Schaffner. Impossibility of quantum virtual black-box obfuscation of classical circuits. In *Advances in Cryptology–CRYPTO 2021: 41st Annual International Cryptology Conference, CRYPTO 2021, Virtual Event, August 16–20, 2021, Proceedings, Part I 41*, pages 497–525. Springer, 2021.

ABKK23. Amit Agarwal, James Bartusek, Dakshita Khurana, and Nishant Kumar. A new framework for quantum oblivious transfer. In *Advances in Cryptology–EUROCRYPT 2023: 42nd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Lyon, France, April 23-27, 2023, Proceedings, Part I*, pages 363–394. Springer, 2023.

ACE+23. Behzad Abdolmaleki, Céline Chevalier, Ehsan Ebrahimi, Giulio Malavolta, and Quoc-Huy Vu. On quantum simulation-soundness. *Cryptology ePrint Archive*, 2023.

AF07. Masayuki Abe and Serge Fehr. Perfect nizk with adaptive soundness. In *Theory of Cryptography Conference*, pages 118–136. Springer, 2007.

AHY23. Prabhanjan Ananth, Zihan Hu, and Henry Yuen. On the (im)plausibility of public-key quantum money from collision-resistant hash functions. In *Advances in Cryptology – ASIACRYPT 2023: 29th International Conference on the Theory and Application of Cryptology and Information Security, Guangzhou, China, December 4–8, 2023, Proceedings, Part VIII*, page 39–72, Berlin, Heidelberg, 2023. Springer-Verlag.

AMRS20. Gorjan Alagic, Christian Majenz, Alexander Russell, and Fang Song. Quantum-access-secure message authentication via blind-unforgeability. In *Advances in Cryptology–EUROCRYPT 2020: 39th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, May 10–14, 2020, Proceedings, Part III 39*, pages 788–817. Springer, 2020.

AQY22. Prabhanjan Ananth, Luowen Qian, and Henry Yuen. Cryptography from pseudorandom quantum states. In Yevgeniy Dodis and Thomas Shrimpton, editors, *Advances in Cryptology – CRYPTO 2022*, pages 208–236, Cham, 2022. Springer Nature Switzerland.

BB14. Charles H Bennett and Gilles Brassard. Quantum cryptography: Public key distribution and coin tossing. *Theoretical computer science*, 560:7–11, 2014.

BBF13. Paul Baecher, Christina Brzuska, and Marc Fischlin. Notions of black-box reductions, revisited. In *Advances in Cryptology-ASIACRYPT 2013: 19th International Conference on the Theory and Application of Cryptology and Information Security, Bengaluru, India, December 1-5, 2013, Proceedings, Part I 19*, pages 296–315. Springer, 2013.

BCKM21. James Bartusek, Andrea Coladangelo, Dakshita Khurana, and Fermi Ma. One-way functions imply secure computation in a quantum world. In *Advances in Cryptology–CRYPTO 2021: 41st Annual International Cryptology Conference, CRYPTO 2021, Virtual Event, August 16–20, 2021, Proceedings, Part I 41*, pages 467–496. Springer, 2021.

BCM+21. Zvika Brakerski, Paul Christiano, Urmila Mahadev, Umesh Vazirani, and Thomas Vidick. A cryptographic test of quantumness and certifiable randomness from a single quantum device, 2021.

BCQ23. Zvika Brakerski, Ran Canetti, and Luowen Qian. On the computational hardness needed for quantum cryptography. In *14th Innovations in Theoretical Computer Science Conference (ITCS 2023)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2023.

BDSG+13. Nir Bitansky, Dana Dachman-Soled, Sanjam Garg, Abhishek Jain, Yael Tauman Kalai, Adriana López-Alt, and Daniel Wichs. Why "fiat-shamir for proofs" lacks a proof. In Amit Sahai, editor, *Theory of Cryptography*, pages 182–201, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.

BDSMP91. Manuel Blum, Alfredo De Santis, Silvio Micali, and Giuseppe Persiano. Noninteractive zero-knowledge. *SIAM Journal on Computing*, 20(6):1084–1118, 1991.

BFM88. Manuel Blum, Paul Feldman, and Silvio Micali. Non-interactive zero-knowledge and its applications. In *Proceedings of the Twentieth Annual ACM Symposium on Theory of Computing*, STOC '88, page 103–112, New York, NY, USA, 1988. Association for Computing Machinery.

BG89. Mihir Bellare and Shafi Goldwasser. New paradigms for digital signatures and message authentication based on non-interactive zero knowledge proofs. In *Conference on the Theory and Application of Cryptology*, pages 194–211. Springer, 1989.

BI20. Anne Broadbent and Rabib Islam. Quantum encryption with certified deletion. In *Theory of Cryptography: 18th International Conference, TCC 2020, Durham, NC, USA, November 16–19, 2020, Proceedings, Part III 18*, pages 92–122. Springer, 2020.

BJSW16. Anne Broadbent, Zhengfeng Ji, Fang Song, and John Watrous. Zero-knowledge proof systems for qma. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 31–40. IEEE, 2016.

BK23. James Bartusek and Dakshita Khurana. Cryptography with certified deletion. In *Annual International Cryptology Conference*, pages 192–223. Springer, 2023.

BKS23. James Bartusek, Dakshita Khurana, and Akshayaram Srinivasan. Secure computation with shared epr pairs (or: How to teleport in zero-knowledge). In Helena Handschuh and Anna Lysyanskaya, editors, *Advances in Cryptology – CRYPTO 2023*, pages 224–257, Cham, 2023. Springer Nature Switzerland.

BL19. Anne Broadbent and Sébastien Lord. Uncloneable quantum encryption via oracles. *arXiv preprint arXiv:1903.00130*, 2019.

BR90. Mihir Bellare and Phillip Rogaway. Non-interactive perfect zero-knowledge. *Unpublished manuscript, June*, pages 1084–1118, 1990.

BV98. Dan Boneh and Ramarathnam Venkatesan. Breaking rsa may not be equivalent to factoring. In Kaisa Nyberg, editor, *Advances in Cryptology — EUROCRYPT'98*, pages 59–71, Berlin, Heidelberg, 1998. Springer Berlin Heidelberg.

BZ13. Dan Boneh and Mark Zhandry. Secure signatures and chosen ciphertext security in a quantum computing world. In *Advances in Cryptology–CRYPTO 2013: 33rd Annual Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2013. Proceedings, Part II*, pages 361–379. Springer, 2013.

CCH+19. Ran Canetti, Yilei Chen, Justin Holmgren, Alex Lombardi, Guy N Rothblum, Ron D Rothblum, and Daniel Wichs. Fiat-shamir: from practice to theory. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 1082–1090, 2019.

CGKS23. Matteo Campanelli, Chaya Ganesh, Hamidreza Khoshakhlagh, and Janno Siim. Impossibilities in succinct arguments: Black-box extraction and more. In *International Conference on Cryptology in Africa*, pages 465–489. Springer, 2023.

CHS20. Nai-Hui Chia, Sean Hallgren, and Fang Song. On basing one-way permutations on NP-hard problems under quantum reductions. *Quantum*, 4:312, aug 2020.

CLM23. Kai-Min Chung, Yao-Ting Lin, and Mohammad Mahmoody. Black-box separations for non-interactive classical commitments in a quantum world. In Carmit Hazay and Martijn Stam, editors, *Advances in Cryptology – EUROCRYPT 2023*, pages 144–172, Cham, 2023. Springer Nature Switzerland.

CLMP12. Kai-Min Chung, Edward Lui, Mohammad Mahmoody, and Rafael Pass. Unprovable security of 2-message zero knowledge. Cryptology ePrint Archive, Paper 2012/711, 2012. `https://eprint.iacr.org/2012/711`.

CM24. Andrea Coladangelo and Saachi Mutreja. On black-box separations of quantum digital signatures from pseudorandom states. *arXiv preprint arXiv:2402.08194*, 2024.

CMS23. Léo Colisson, Garazi Muguruza, and Florian Speelman. Oblivious transfer from zero-knowledge proofs: Or how to achieve round-optimal quantum oblivious transfer and zero-knowledge proofs on quantum states. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 3–38. Springer, 2023.

CMSZ22. Alessandro Chiesa, Fermi Ma, Nicholas Spooner, and Mark Zhandry. Post-quantum succinct arguments: breaking the quantum rewinding barrier. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 49–58. IEEE, 2022.

Cor02. Jean-Sébastien Coron. Optimal security proofs for pss and other signature schemes. In *Advances in Cryptology—EUROCRYPT 2002: International Conference on the Theory and Applications of Cryptographic Techniques Amsterdam, The Netherlands, April 28–May 2, 2002 Proceedings 21*, pages 272–287. Springer, 2002.

CVZ20. Andrea Coladangelo, Thomas Vidick, and Tina Zhang. Non-interactive zero-knowledge arguments for qma, with preprocessing. In Daniele Micciancio and Thomas Ristenpart, editors, *Advances in Cryptology – CRYPTO 2020*, pages 799–828, Cham, 2020. Springer International Publishing.

CX21. Shujiao Cao and Rui Xue. Being a permutation is also orthogonal to one-wayness in quantum world: Impossibilities of quantum one-way permutations from one-wayness primitives. *Theoretical Computer Science*, 855:16–42, 2021.

DFG13. Özgür Dagdelen, Marc Fischlin, and Tommaso Gagliardoni. The fiat–shamir transformation in a quantum world. In *Advances in Cryptology-ASIACRYPT 2013: 19th International Conference on the Theory and Application of Cryptology and Information Security, Bengaluru, India, December 1-5, 2013, Proceedings, Part II 19*, pages 62–81. Springer, 2013.

DH22. Whitfield Diffie and Martin E. Hellman. *New Directions in Cryptography*, page 365–390. Association for Computing Machinery, 2022.

DLS22. Frédéric Dupuis, Philippe Lamontagne, and Louis Salvail. Fiat-shamir for proofs lacks a proof even in the presence of shared entanglement. Cryptology ePrint Archive, Paper 2022/435, 2022. `https://eprint.iacr.org/2022/435`.

DOP05. Yevgeniy Dodis, Roberto Oliveira, and Krzysztof Pietrzak. On the generic insecurity of the full domain hash. In Victor Shoup, editor, *Advances in Cryptology – CRYPTO 2005*, pages 449–466, Berlin, Heidelberg, 2005. Springer Berlin Heidelberg.

DSDCO+01. Alfredo De Santis, Giovanni Di Crescenzo, Rafail Ostrovsky, Giuseppe Persiano, and Amit Sahai. Robust non-interactive zero knowledge. In *Advances in Cryptology—CRYPTO 2001: 21st Annual International Cryptology Conference, Santa Barbara, California, USA, August 19–23, 2001 Proceedings 21*, pages 566–598. Springer, 2001.

FLS99. Uriel Feige, Dror Lapidot, and Adi Shamir. Multiple noninteractive zero knowledge proofs under general assumptions. *SIAM Journal on Computing*, 29(1):1–28, 1999.

FR21. Marc Fischlin and Felix Rohrbach. Single-to-multi-theorem transformations for non-interactive statistical zero-knowledge. In *IACR International Conference on Public-Key Cryptography*, pages 205–234. Springer, 2021.

FS10. Marc Fischlin and Dominique Schröder. On the impossibility of three-move blind signature schemes. In Henri Gilbert, editor, *Advances in Cryptology – EUROCRYPT 2010*, pages 197–215, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg.

GGM86. Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions. *Journal of the ACM (JACM)*, 33(4):792–807, 1986.

GJMZ23. Sam Gunn, Nathan Ju, Fermi Ma, and Mark Zhandry. Commitments to quantum states. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, STOC 2023, page 1579–1588, New York, NY, USA, 2023. Association for Computing Machinery.

GLSV20. Alex B. Grilo, Huijia Lin, Fang Song, and Vinod Vaikuntanathan. Oblivious transfer is in miniqcrypt, 2020.

GLSV21. Alex B Grilo, Huijia Lin, Fang Song, and Vinod Vaikuntanathan. Oblivious transfer is in miniqcrypt. In *Advances in Cryptology–EUROCRYPT 2021: 40th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, October 17–21, 2021, Proceedings, Part II*, pages 531–561. Springer, 2021.

GMMM18. Sanjam Garg, Mohammad Mahmoody, Daniel Masny, and Izaak Meckler. On the round complexity of ot extension. In Hovav Shacham and Alexandra Boldyreva, editors, *Advances in Cryptology – CRYPTO 2018*, pages 545–574, Cham, 2018. Springer International Publishing.

GMR01. Y. Gertner, T. Malkin, and O. Reingold. On the impossibility of basing trapdoor functions on trapdoor predicates. In *Proceedings of the 42nd IEEE Symposium on Foundations of Computer Science*, FOCS '01, page 126, USA, 2001. IEEE Computer Society.

GO94. Oded Goldreich and Yair Oren. Definitions and properties of zero-knowledge proof systems. *Journal of Cryptology*, 7(1):1–32, 1994.

GOS06. Jens Groth, Rafail Ostrovsky, and Amit Sahai. Perfect non-interactive zero knowledge for np. In Serge Vaudenay, editor, *Advances in Cryptology - EUROCRYPT 2006*, pages 339–358, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg.

GW07. Gus Gutoski and John Watrous. Toward a general theory of quantum games. In *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, STOC07. ACM, June 2007.

GW11. Craig Gentry and Daniel Wichs. Separating succinct non-interactive arguments from all falsifiable assumptions. In *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing*, STOC '11, page 99–108, New York, NY, USA, 2011. Association for Computing Machinery.

HILL99. Johan Håstad, Russell Impagliazzo, Leonid A Levin, and Michael Luby. A pseudorandom generator from any one-way function. *SIAM Journal on Computing*, 28(4):1364–1396, 1999.

HY20. Akinori Hosoyamada and Takashi Yamakawa. Finding collisions in a quantum world: quantum black-box separation of collision-resistance and one-wayness. In *Advances in Cryptology–ASIACRYPT 2020: 26th International Conference on the Theory and Application of Cryptology and Information Security, Daejeon, South Korea, December 7–11, 2020, Proceedings, Part I 26*, pages 3–32. Springer, 2020.

IR89. R. Impagliazzo and S. Rudich. Limits on the provable consequences of one-way permutations. In *Proceedings of the Twenty-First Annual ACM Symposium on Theory of Computing*, STOC '89, page 44–61, New York, NY, USA, 1989. Association for Computing Machinery.

JLS18. Zhengfeng Ji, Yi-Kai Liu, and Fang Song. Pseudorandom quantum states. In Hovav Shacham and Alexandra Boldyreva, editors, *Advances in Cryptology – CRYPTO 2018*, pages 126–152, Cham, 2018. Springer International Publishing.

JZM21. Haodong Jiang, Zhenfeng Zhang, and Zhi Ma. On the non-tightness of measurement-based reductions for key encapsulation mechanism in the quantum random oracle model. In *Advances in Cryptology–ASIACRYPT 2021: 27th International Conference on the Theory and Application of Cryptology and Information Security, Singapore, December 6–10, 2021, Proceedings, Part I 27*, pages 487–517. Springer, 2021.

Kil92. Joe Kilian. A note on efficient zero-knowledge proofs and arguments. In *Proceedings of the twenty-fourth annual ACM symposium on Theory of computing*, pages 723–732, 1992.

Kiy21. Susumu Kiyoshima. Black-box impossibilities of obtaining 2-round weak zk and strong wi from polynomial hardness. In Kobbi Nissim and Brent Waters, editors, *Theory of Cryptography*, pages 369–400, Cham, 2021. Springer International Publishing.

KK19. Yael Tauman Kalai and Dakshita Khurana. Non-interactive non-malleability from quantum supremacy. In *Annual International Cryptology Conference*, pages 552–582. Springer, 2019.

Kre21. William Kretschmer. Quantum Pseudorandomness and Classical Complexity. In Min-Hsiu Hsieh, editor, *16th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2021)*, volume 197 of *Leibniz International Proceedings in Informatics (LIPIcs)*,

pages 2:1–2:20, Dagstuhl, Germany, 2021. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.

LMQW22. Alex Lombardi, Ethan Mook, Willy Quach, and Daniel Wichs. Post-quantum insecurity from lwe. In *Theory of Cryptography: 20th International Conference, TCC 2022, Chicago, IL, USA, November 7–10, 2022, Proceedings, Part I*, pages 3–32. Springer, 2022.

LMS22. Alex Lombardi, Fermi Ma, and Nicholas Spooner. Post-quantum zero knowledge, revisited or: How to do quantum rewinding undetectably. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 851–859. IEEE, 2022.

Lo97. Hoi-Kwong Lo. Insecurity of quantum secure computations. *Physical Review A*, 56(2):1154, 1997.

LQS+23. Chuhan Lu, Minglong Qin, Fang Song, Penghui Yao, and Mingnan Zhao. Quantum pseudorandom scramblers. *arXiv preprint arXiv:2309.08941*, 2023.

MP18. Andrew Morgan and Rafael Pass. On the security loss of unique signatures. In *Theory of Cryptography Conference*, pages 507–536. Springer, 2018.

MY23. Tomoyuki Morimae and Takashi Yamakawa. Classically verifiable nizk for qma with preprocessing. In *Advances in Cryptology – ASIACRYPT 2022: 28th International Conference on the Theory and Application of Cryptology and Information Security, Taipei, Taiwan, December 5–9, 2022, Proceedings, Part IV*, page 599–627, Berlin, Heidelberg, 2023. Springer-Verlag.

Nao03. Moni Naor. On cryptographic assumptions and challenges. In *Advances in Cryptology-CRYPTO 2003: 23rd Annual International Cryptology Conference, Santa Barbara, California, USA, August 17-21, 2003. Proceedings 23*, pages 96–109. Springer, 2003.

Pas13. Rafael Pass. Unprovable security of perfect nizk and non-interactive non-malleable commitments. In Amit Sahai, editor, *Theory of Cryptography*, pages 334–354, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.

PS19. Chris Peikert and Sina Shiehian. Noninteractive zero knowledge for np from (plain) learning with errors. In *Annual International Cryptology Conference*, pages 89–114. Springer, 2019.

RSA78. Ronald L Rivest, Adi Shamir, and Leonard Adleman. A method for obtaining digital signatures and public-key cryptosystems. *Communications of the ACM*, 21(2):120–126, 1978.

RTV04. Omer Reingold, Luca Trevisan, and Salil Vadhan. Notions of reducibility between cryptographic primitives. In Moni Naor, editor, *Theory of Cryptography*, pages 1–20, Berlin, Heidelberg, 2004. Springer Berlin Heidelberg.

Sah99. Amit Sahai. Non-malleable non-interactive zero knowledge and adaptive chosen-ciphertext security. In *40th Annual Symposium on Foundations of Computer Science (Cat. No. 99CB37039)*, pages 543–553. IEEE, 1999.

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.

Shm21. Omri Shmueli. Multi-theorem designated-verifier nizk for qma. In *Annual International Cryptology Conference*, pages 375–405. Springer, 2021.

Sho99. Peter W Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM review*, 41(2):303–332, 1999.

Sim98. Daniel R. Simon. Finding collisions on a one-way street: Can secure hash functions be based on general assumptions? In Kaisa Nyberg, editor, *Advances in Cryptology — EUROCRYPT'98*, pages 334–345, Berlin, Heidelberg, 1998. Springer Berlin Heidelberg.

SW14. Amit Sahai and Brent Waters. How to use indistinguishability obfuscation: Deniable encryption, and more. In *Proceedings of the Forty-Sixth Annual ACM Symposium on Theory of Computing*, STOC '14, page 475–484, New York, NY, USA, 2014. Association for Computing Machinery.

Wic13. Daniel Wichs. Barriers in cryptography with weak, correlated and leaky sources. In *Proceedings of the 4th conference on Innovations in Theoretical Computer Science*, pages 111–126, 2013.

Zha12. Mark Zhandry. How to construct quantum random functions. In *2012 IEEE 53rd Annual Symposium on Foundations of Computer Science*, pages 679–687. IEEE, 2012.