

Structure vs. Hardness through the Obfuscation Lens ^{*}

Nir Bitansky, Akshay Degwekar, and Vinod Vaikuntanathan

MIT

Abstract. Much of modern cryptography, starting from public-key encryption and going beyond, is based on the hardness of structured (mostly algebraic) problems like factoring, discrete log or finding short lattice vectors. While structure is perhaps what enables advanced applications, it also puts the hardness of these problems in question. In particular, this structure often puts them in low complexity classes such as $\text{NP} \cap \text{coNP}$ or statistical zero-knowledge (SZK).

Is this structure really necessary? For some cryptographic primitives, such as one-way permutations and homomorphic encryption, we know that the answer is *yes* — they imply hard problems in $\text{NP} \cap \text{coNP}$ and SZK, respectively. In contrast, one-way functions do *not* imply such hard problems, at least not by *fully black-box reductions*. Yet, for many basic primitives such as public-key encryption, oblivious transfer, and functional encryption, we do not have any answer.

We show that the above primitives, and many others, do *not* imply hard problems in $\text{NP} \cap \text{coNP}$ or SZK via fully black-box reductions. In fact, we first show that even the very powerful notion of Indistinguishability Obfuscation (IO) does *not* imply such hard problems, and then deduce the same for a large class of primitives that can be constructed from IO.

Keywords: Indistinguishability Obfuscation, Statistical Zero-knowledge, $\text{NP} \cap \text{coNP}$, Structured Hardness, Collision-Resistant Hashing.

1 Introduction

The last four decades of research in cryptography has produced a host of fantastic objects, starting from one-way functions and permutations to public-key encryption [DH76, RSA78, GM82] and zero-knowledge proofs [GMR85] in the 1980s, all the way to fully homomorphic encryption [RAD78, Gen09, BV11] and indistinguishability obfuscation [BGI⁺01, GGH⁺13a] in the modern day.

The existence of all these objects requires at the very minimum that $\text{NP} \not\subseteq \text{BPP}$, but that is hardly ever enough. While one-way functions (OWFs), the

^{*} MIT CSAIL. E-mail: {nirbitan, akshayd, vinodv}@csail.mit.edu. Research supported in part by NSF Grants CNS-1350619 and CNS-1414119, Alfred P. Sloan Research Fellowship, Microsoft Faculty Fellowship, the NEC Corporation, a Steven and Renee Finn Career Development Chair from MIT. This work was also sponsored in part by the Defense Advanced Research Projects Agency (DARPA) and the U.S. Army Research Office under contracts W911NF-15-C-0226.

most basic cryptographic object, does not seem to require much structure, as we advance up the ranks, we seem to require that certain *structured problems are hard*. For example, conjectured hard problems commonly used in cryptography (especially the public-key kind), such as factoring, discrete logarithms, and shortest (or closest) vectors on lattices all have considerable algebraic structure. On the one hand, it is this structure that enables strong applications such as public-key and homomorphic encryption. On the other hand, this structure is also what puts their hardness in question, and is exactly what algorithms may try to exploit in order to solve these problems. There is of course the fear that this structure will (eventually, if not today) deem these problems *easy*. Or, as Barak says more eloquently [Bar13]:

[...] based on the currently well studied schemes, structure is strongly associated with (and perhaps even implied by) public key cryptography. This is troubling news, since it makes public key crypto somewhat of an “endangered species” that could be wiped out by a surprising algorithmic advance. Therefore the question of whether structure is inherently necessary for public key crypto is not only of mathematical interest but also of practical importance as well.

Thus, a fundamental question in cryptography is *what type of structure is necessary for different primitives?* Indeed, the answer to this question may be crucial to our understanding of what are the minimal assumptions required to construct these primitives. While there may be different ways of approaching this question, one main approach, which is also taken in this work, has been through the eyes of complexity theory. That is, we wish to understand which cryptographic primitives require hardness in low (and so called structured) complexity classes such as $\text{NP} \cap \text{coNP}$, TFNP (the class of total NP search problems), or SZK (the class of problems with statistical zero-knowledge proofs).

Aiming to answer this question, one line of research demonstrates that, for some cryptographic primitives, hardness in structured complexity classes is indeed necessary. The existence of one-way permutations (OWPs) requires a hard problem in $\text{NP} \cap \text{coNP}$ [Bra79]; the same holds for restricted cases of public-key encryption schemes satisfying specific structural properties (e.g ciphertext certification) [Bra79, GG98]; homomorphic encryption schemes and non-interactive computational private information retrieval schemes imply hard problems in SZK [BL13, LV16]; and indistinguishability obfuscation schemes imply a hard problem in $\text{PPAD} \subseteq \text{TFNP}$ (assuming $\text{NP} \not\subseteq \text{ioBPP}$) [BPR15].

Yet, for many primitives such hardness is not known to be inherent. While this is perhaps expected for OWFs, it is also the case for seemingly structured primitives such as collision-resistant hash functions, oblivious transfer, and general public-key encryption schemes. *Do these primitives require hardness in structured complexity classes? Can we prove that they do or that they don't?*

Black-Box Separations. Formalizing this question in a meaningful way requires care. Indeed, it may be easy to formalize a statement of the form “the existence of crypto primitive \mathcal{P} implies hardness in a complexity class \mathcal{C} ”: one just needs to show a reduction from breaking \mathcal{P} to solving problems in \mathcal{C} . However, it is not

clear how to prove statements of the form “the existence of crypto primitive \mathcal{P} does *not* imply hardness in a complexity class \mathcal{C} ”. For example, it is commonly believed that $\text{NP} \cap \text{coNP}$ *does* contain hard problems. So in a trivial logical sense the existence of such problems is implied by any primitive \mathcal{P} . Instead, we follow the methodology of black-box separations, whose study in cryptography was pioneered by Impagliazzo and Rudich [IR89]. Faced with a similar problem of how to show that a primitive \mathcal{P} (OWFs) cannot be used to construct another primitive \mathcal{P}' (public-key encryption), they prove this cannot be shown through *black-box reductions* — cryptography’s de facto technique for showing such implications.

A bit more elaborately, a *fully black-box reduction* [RTV04] of a primitive (or, in our case, a problem) \mathcal{P}' to a primitive \mathcal{P} consists of a black-box *construction* and a black-box *security reduction*. The construction of \mathcal{P}' from \mathcal{P} does not exploit the actual implementation of primitive \mathcal{P} , but rather just its input-output interface. The security reduction can use any adversary that breaks (or, in our case, solves) \mathcal{P}' to break \mathcal{P} , and is oblivious to the implementation of the adversary (as well as of that of \mathcal{P}).

Following [IR89], there has been a rich study of black-box separations in cryptography (see, e.g., [Rud91, Sim98, KST99, GKM⁺00, GT00, GMR01, BT03, RTV04, HR04, GGKT05, Pas06, GMM07, BM09, HH09, BKSY11, DLMM11, KSS11, GKLM12, DHT12, BBF13, Fis12, Pas13, BB15, HHR15] and many others). Most of this study has been devoted to establishing separations between different cryptographic primitives. (In particular, the most relevant to us are the recent works of Asharov and Segev [AS15, AS16] that study black-box separations for indistinguishability obfuscation, which we elaborate on below.) Some of this study puts limitations on basing cryptographic primitives on NP-hardness [GG98, AGGM06, MX10, HMX10, BL13, BB15, LV16].

Going back to our main question of which primitives require structured hardness, we know the following.

- As described above, OWPs imply a hard problem in $\text{NP} \cap \text{coNP}$ [Bra79], homomorphic encryption and PIR imply hard problems in SZK [BL13, LV16] and IO (with OWFs) implies a hard problem in PPAD [BPR15] via *black-box reductions*.
- On the flip side, we know that there are no black-box reductions from hard problems in $\text{NP} \cap \text{coNP}$ to OWFs [BI87, Rud88], and from hard-on-average problems in SZK to OWPs (corollary from [Ost91, OV08, HHR15]).

For more advanced primitives, most notably (general) public-key encryption, we do not have results in either direction. In fact, many existing constructions are based on problems in $\text{NP} \cap \text{coNP}$ or SZK. We are thus left with (quite basic) primitives at an unclear state; as far as we know, they may very well imply hard problems in structured complexity classes, even by black-box reductions.

1.1 Our Results

We revisit the relationship between two structured complexity classes, statistical zero-knowledge (SZK) and $\text{NP} \cap \text{coNP}$, and cryptographic primitives. In broad strokes, we show that there are no fully black-box reductions of hard problems in these classes to any one of a variety of cryptographic primitives, including (general) public-key encryption, oblivious transfer, deniable encryption, and functional encryption. More generally, we separate SZK and $\text{NP} \cap \text{coNP}$ from indistinguishability obfuscation (IO). Then, leveraging on the fact that IO can be used to construct a wide variety of cryptographic primitives in a black-box way, we derive corresponding separations for these primitives.¹ One complexity-theoretic corollary of this result is a separation between SZK and $\text{NP} \cap \text{coNP}$ from the class PPAD [MP91] that captures the complexity of computing Nash Equilibria.

On the positive side, we construct collision-resistant hash functions from a strong form of SZK-hardness and IO. It was previously known [AS15] that IO by itself does not imply collision-resistant hashing in a black-box way; we show that it does if one adds SZK-hardness as a “catalyst”.

We now go into more detail on each of the results.

Statistical Zero-Knowledge and Cryptography. The notion of statistical zero-knowledge proofs was introduced in the seminal work of Goldwasser, Micali and Rackoff [GMR85]. The class of *promise problems* with statistical zero-knowledge proofs (SZK) can be characterized by several complete problems, such as *statistical difference* [SV03] and *entropy difference* [GV99]. SZK hardness is known to follow from various number-theoretic problems that are commonly used in cryptography, such as Discrete Logarithms [GK93], Quadratic Residuosity [GMR85], Lattice Problems [GG98, MV03] as well as problems like Graph Isomorphism [GMW91]. As mentioned, we also know that a handful of cryptographic primitives such as homomorphic encryption [BL13], private information retrieval [LV16] and rerandomizable encryption imply hardness in SZK. (On the other hand, $\text{SZK} \subseteq \text{AM} \cap \text{coAM}$ [For89, AH91], and thus, SZK cannot contain NP-hard problems, unless the polynomial hierarchy collapses [BHZ87].)

We ask more generally which cryptographic primitives can be shown to imply such hardness, with the intuition that such primitives are *structured* in a certain way. In particular, whereas one may not expect a seemingly unstructured object like OWFs to imply such hardness, what can we say for instance about OWPs, public-key encryption, or even IO (which has proven to be powerful enough to yield almost any known cryptographic goal)?

We prove that none of these primitives imply such hardness through black-box reductions.

¹ More accurately, these primitives follow from IO and OWFs (OWFs), and accordingly our separation addresses IO and OWFs in conjunction. The concept of a black-box reduction from IO and OWF requires clarification and discussion. Here we will follow the framework of Asharov and Segev [AS15]. We elaborate below.

Theorem 1.1 (Informal). *There is no fully black-box reduction of any (even worst-case) hard problem in SZK to IO and OWPs.*

Corollary 1.2 (from [SW14, Wat15], Informal). *There is no such reduction to (general) public-key encryption, oblivious transfer, deniable encryption, functional encryption, or any other object that has a black-box reduction to IO and OWPs.*

We would like to elaborate a bit more on what a black-box construction of a hard problem in SZK means. We shall focus on the characterization of SZK by the *statistical difference* promise problem [SV03]. In this problem, an instance is a pair of circuit samplers $C_0, C_1 : \{0, 1\}^n \rightarrow \{0, 1\}^m$ which induce distributions \mathcal{C}_0 and \mathcal{C}_1 where the distribution \mathcal{C}_b obtained by evaluating the circuit C_b on a uniformly random input. The promise is that the statistical distance $s = \Delta(\mathcal{C}_0, \mathcal{C}_1)$ of the corresponding distributions is either large (say, $s \geq 2/3$) or small (say, $s \leq 1/3$). The problem, named $\mathbf{SD}^{1/3, 2/3}$ (or just \mathbf{SD}), is to decide which is the case.

Let us look at a specific example of the construction of such a problem from *rerandomizable encryption*. In a (say, symmetric-key) rerandomizable encryption scheme, on top of the usual encryption and decryption algorithms (Enc, Dec) there is a ciphertext rerandomization algorithm ReRand that can statistically refresh ciphertexts. Namely, for any ciphertext CT encrypting a bit b , $\text{ReRand}(\text{CT})$ produces a ciphertext that is statistically close to a fresh encryption $\text{Enc}_{\text{sk}}(b)$. This immediately gives rise to a hard statistical difference problem [BL13]: given a pair of ciphertexts $(\text{CT}_0, \text{CT}_1)$, decide whether the corresponding rerandomized distributions given by the circuits $(C_0(\cdot), C_1(\cdot)) := (\text{ReRand}(\text{CT}_0; \cdot), \text{ReRand}(\text{CT}_1; \cdot))$ are statistically far or close. Indeed, this corresponds to whether they encrypt the same bit or not, which is hard to decide by the security of the encryption scheme.

A feature of this reduction of hard statistical difference instances to rerandomizable encryption is that, similarly to most reductions in cryptography, it is *fully black-box* [RTV04] in the sense that the circuits C_0, C_1 only make black-box use of the encryption scheme’s algorithms, and can in fact be represented as oracle-aided circuits $(C_0^{\text{ReRand}(\cdot)}, C_1^{\text{ReRand}(\cdot)})$. Furthermore, “hardness” can be shown by a black-box security proof that can use any decider for the problem in a black-box way to break the underlying encryption scheme. More generally, one can consider the statistical difference problem relative to different oracles implementing different cryptographic primitives and ask when can hardness be shown based on a black-box reduction. Theorem 1.1 rules out such reductions relative to IO and OWPs (and everything that follows from these in a fully black-box way). For more details, see Section 1.2 and the full version.

NP \cap coNP and Cryptography. Hard (on average) problems in $\text{NP} \cap \text{coNP}$ are known to follow based on several number-theoretic problems in cryptography, such as Discrete Log, Factoring and Lattice Problems [Has88, LLJS90, AR04]. As in the previous section for SZK, we are interested in understanding which

cryptographic primitives would imply such hardness, again with the intuition that this implies structure. For instance, it is known [Bra79] that any OWP $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$ implies a hard problem in $\text{NP} \cap \text{coNP}$, e.g. given an index $i \in [n]$ and an image $f(x)$ find the i th preimage bit x_i . In contrast, Blum and Impagliazzo [BI87] and Rudich [Rud88] proved that seemingly unstructured objects like OWFs do not imply hardness in $\text{NP} \cap \text{coNP}$ by fully black-box reductions. In this context, a fully black-box reduction essentially means that the non-deterministic verifiers only make black-box use of the OWF (or OWP in the previous example) and the reduction establishing the hardness is also black-box (in both the decider and the OWF).²

But what about more structured primitives such as public-key encryption, oblivious transfer, or even IO? We rule out fully black-box reductions from OWFs (or even *injective OWFs*) and IO to hard problems in $\text{NP} \cap \text{coNP}$. Hence, also for the other primitives, which can be constructed from IO (with OWFs) in a fully black-box way.

Theorem 1.3 (Informal). *There is no fully black-box reduction of any (even worst-case) hard problem in $\text{NP} \cap \text{coNP}$ to IO and OWFs.*

Corollary 1.4 (from [SW14, Wat15], Informal). *There is no such reduction to (general) public-key encryption, oblivious transfer, deniable encryption, functional encryption, or any other object that has a black-box reduction to IO and OWFs.*

Our approach also gives a new (rather different) proof to the original separation between OWFs and $\text{NP} \cap \text{coNP}$ [BI87, Rud88]. For more details, see Section 1.2 and the full version.

We remark that unlike our result for SZK (which ruled out hard *promise problems*), the above result only rules out hard *languages* in $\text{NP} \cap \text{coNP}$. Indeed, Even, Selman, and Yacobi [ESY84] demonstrated promise problems in $\text{NP} \cap \text{coNP}$ that are NP-hard. Hence even the assumption $\text{P} \neq \text{NP}$ (let alone OWFs) gives us hard promise problems in $\text{NP} \cap \text{coNP}$. (See [Gol06] for further reading.)

Relation to the Work of Asharov and Segev. The flood of IO applications following, starting from [GGH⁺13b, SW14], has lead many to conjecture that IO may be “complete for cryptography” (assuming also OWFs, or just $\text{NP} \not\subseteq \text{ioBBP}$ [KMN⁺14]). Nevertheless, some cryptographic goals could not be constructed based on IO.

Asharov and Segev [AS15, AS16] were the first to initiate a formal study to understand *the limits of IO*. Our separations for IO are based on their framework [AS15]. We aim to draw the complexity-theoretic boundaries of IO. Indeed,

² Roughly speaking, [BI87] rule out *perfectly correct constructions*, where the $\text{NP} \cap \text{coNP}$ structure is guaranteed for any implementation of the OWF oracle. In [Rud88], this is generalized also to *almost perfectly correct constructions* that only work for an overwhelming fraction of OWF oracles. We also rule out constructions that are perfectly correct.

black-box separations from IO require some care, given that the typical use of IO makes non-black-box use of the circuits it obfuscates and thus any associated cryptographic primitive such as OWFs. The Asharov-Segev framework considers obfuscators that take as input circuits with OWF (or OWP) gates. They observe, most known IO-based constructions fall into this category. Thus, a separation in this model allows deriving the corresponding separations between SZK or $\text{NP} \cap \text{coNP}$ and a wide variety of cryptographic primitives. See Section 1.2 for more details.

In terms of results, they show that collision-resistant hashing and (domain invariant) OWPs do not have black-box reductions to IO (and OWFs). Our separation of IO and $\text{NP} \cap \text{coNP}$ is more general and implies their previous result for OWPs (and gives a rather different proof for this fact). Their result for collision-resistant hashing is not captured by our results (indeed collision-resistance is not known to imply hardness in either SZK or $\text{NP} \cap \text{coNP}$). We also stress that our separation of SZK from IO and OWPs does not follow from their results; indeed, SZK-hardness is not known to imply collision-resistance.³

Indistinguishability Obfuscation: Perspective. Since the breakthrough of [GGH⁺13b], the notion of IO has been extensively studied. While we already understand that IO has far reaching implications, our understanding of how it can be constructed and under what assumptions is still at an early stage. Indeed, basing IO on solid foundations is one of cryptography’s greatest challenges today. In this context, we stress that the results presented in this work hold regardless of the state of existing candidates. In fact, even if it turned out that there is no secure realization of IO, the separation of SZK and $\text{NP} \cap \text{coNP}$ from primitives such as public-key encryption, which follow from IO, still holds. The expressiveness of IO (established in [GGH⁺13b, SW14] and onwards) allows us to prove many separations in one shot. (Indeed, three years ago we would have probably addressed each primitive separately.)

As for the search for candidates itself, while at this point candidates are based on lattice-related problems that do break in SZK, our work suggests the theoretical possibility that IO candidates may not require such structure. A similar conclusion is true of course for the much more basic and long-studied question of public-key encryption. Almost all known public-key encryption candidates rely on very algebraic assumptions (that do break in SZK or $\text{NP} \cap \text{coNP}$). Constructing public key encryption from less structured assumptions remains a fascinating open question. While there has been initial steps trying to diverge from such structure [Ale03, ABW10], there is yet a long way to go.

On TFNP vs. $\text{NP} \cap \text{coNP}$. One of the corollaries of our result is a separation between SZK and $\text{NP} \cap \text{coNP}$ from the complexity class PPAD. PPAD, a subclass

³ We note that previous work [Ost91, OV08] does imply that constant-round statistically-hiding commitments have a black-box reduction to any *hard-on-average* SZK problem. However, [AS15] do not rule these out (but only collision-resistant hashing). We also note that in any case, our result also rules out constructions of worst-case hard SZK problems (rather than average-case hard problems).

of total NP search problems called TFNP [MP91], was defined by Papadimitriou [Pap94] and has been shown to capture the complexity of computing Nash equilibria [DGP06, CDT09]. It was recently shown [BPR15] that IO and injective OWFs can be used (in a black-box way) to construct hard problems in PPAD. Put together with our separation, we get that there is no black-box construction of an SZK (resp. $\text{NP} \cap \text{coNP}$) hard problem from PPAD-hardness.⁴

Given that TFNP, which contains PPAD, is commonly thought of as a search version of $\text{NP} \cap \text{coNP}$, it is interesting to note that the result shows that hardness in $\text{NP} \cap \text{coNP}$ (of decisional problems) does not follow from hardness in TFNP (aka, hardness of search problems) in a black-box way. Namely, there is no black-box “search-to-decision reduction” between these classes.

The Positive Result: Collision-Resistant Hashing from Strong SZK-Hardness. We end our paper with a positive result. While most of our focus has been on showing that hardness in SZK and $\text{NP} \cap \text{coNP}$ does *not* follow from cryptography, here we ask the “inverse question”, namely whether certain cryptographic primitives can be built from other cryptographic primitives together with hardness in certain structured complexity classes. Little is known in this direction with the exception of the beautiful work of Ostrovsky [Ost91] which constructs a OWF from average-case SZK-hardness, and the recent work of Applebaum and Raykov [AR16] who showed that average-case hardness in the subclass $\text{PRE} \subseteq \text{SRE} \subseteq \text{SZK}$ of languages with a perfect randomized encoding gives us collision-resistant hashing.

We construct collision-resistant hashing from a strong form of SZK-hardness and IO. It was previously known [AS15] that IO by itself does not imply collision-resistant hashing in a black-box way; we show that it does if one adds SZK-hardness as a “catalyst”. Slightly more precisely, in the SZK-complete problem $\text{SD}^{1/3, 2/3}$ is required to distinguish between distributions that are 1/3-close from ones that are 2/3-far. We show that IO together with average-case hardness of $\text{SD}^{0,1}$ (a stronger assumption) implies collision-resistant hashing.

Theorem 1.5 (Informal). *Assuming average-case hardness of $\text{SD}^{0,1}$ and the existence of IO, there is a collision-resistant hashing scheme.*

Organization. Due to the paucity of space, most of the proofs are deferred to the full version. We give an overview of the methodology and techniques used in the following Section 1.2. The black-box separation between SZK and IO (plus OWPs) is stated in Section 2. The separation between $\text{NP} \cap \text{coNP}$ and IO (plus injective OWFs) is described in Section 3.

⁴ We note that in concurrent and independent work, Rosen, Shahaf, and Segev [RSS16] show that one-way functions do not have black-box reductions to PPAD-hardness, which combined with [Ost91], also yields a separation between SZK and PPAD.

1.2 Overview of Techniques

We now give an overview of our approach and main ideas. We start by discussing how to capture fully black-box constructions in the context of indistinguishability obfuscation following [AS15]. We then recall the common methodology for ruling out black-box constructions [IR89, RTV04, BBF13], and explain the main ideas behind our impossibility results for SZK and $\text{NP} \cap \text{coNP}$. In the last part of this section, we outline the construction of collision-resistant hashing from indistinguishability obfuscation and SZK-hardness and the main ideas behind it.

Indistinguishability Obfuscation and Black-Box Constructions. Traditionally, when thinking about a *black-box construction* of one cryptographic primitive \mathcal{P}' (e.g., a pseudo-random generator) from a primitive \mathcal{P} (e.g., a one-way function), we mean that all algorithms in the construction of \mathcal{P}' invoke \mathcal{P} as a black-box, oblivious of its actual implementation. This is hardly the case in constructions based on indistinguishability obfuscation where circuits that explicitly invoke the primitive \mathcal{P} may be obfuscated.

Nonetheless, as observed by Asharov and Segev [AS15], in almost all existing constructions, the code implementing \mathcal{P} is used in a very restricted manner. Typically, obfuscated circuits can be implemented as oracle aided circuits $C^{\mathcal{P}}$ that are completely black-box in \mathcal{P} , where \mathcal{P} is some low-level primitive, such as a one-way function. Indeed, in most cases the circuits obfuscated are symmetric-key primitives, such as puncturable pseudo-random functions [SW14], which can be constructed in a black-box way from one-way functions (in some constructions more structured low-level primitives may be used, like injective one-way functions, or one-way permutations). Furthermore, in these constructions, the obfuscator $i\mathcal{O}$ itself is also treated as a black-box.

Accordingly, almost all existing constructions based on indistinguishability obfuscation can be cast into a model in which indistinguishability obfuscation exists for oracle-aided circuits $C^{\mathcal{P}}$, where \mathcal{P} is say a one-way function, and both \mathcal{P} and the obfuscator $i\mathcal{O}$ can only be accessed as black-boxes. On top of that, they can be proven secure in this model by a *black-box reduction* that makes black-box use of $(\mathcal{P}, i\mathcal{O})$ and any attacker against the constructed primitive \mathcal{P}' . Such constructions where both the construction itself and the reduction are black-box are called *fully black-box constructions* [RTV04]. Following Asharov and Segev [AS15, AS16], we shall prove our results in this model, ruling out black-box constructions of hard problems in SZK and $\text{NP} \cap \text{coNP}$ based on indistinguishability obfuscation for oracle-aided circuits. Further details follow.

Ruling out Black-Box Reductions. We prove our results in the model described above following the methodology of oracle separations (see e.g. [IR89, Sim98, RTV04, HR04]). Concretely, to prove that there is no fully black-box construction of a primitive \mathcal{P}' from primitive \mathcal{P} , we demonstrate oracles (Ψ, \mathcal{A}) such that:

- relative to Ψ , there exists a construction $C_{\mathcal{P}}^{\Psi}$ realizing \mathcal{P} that is secure in the presence of \mathcal{A} ,

- but *any* construction $C_{\mathcal{P}'}^{\Psi}$, realizing \mathcal{P}' can be broken in the presence of \mathcal{A} .

Indeed, if such oracles (Ψ, \mathcal{A}) exist, then no efficient reduction will be able to use (as a black-box) the attacker \mathcal{A} against \mathcal{P}' to break \mathcal{P} (as the construction of \mathcal{P} is secure in the presence of \mathcal{A}). In our case, we would like to apply this paradigm rule out black-box constructions of hard instances in either SZK or $\text{NP} \cap \text{coNP}$ from a low-level primitive (e.g. a one-way function) indistinguishability obfuscation for oracle-aided circuits. We next outline the main ideas behind the construction and analysis of the oracles (Ψ, \mathcal{A}) in each of the two cases.

Ruling out Black-Box Constructions of Hard SZK Problems. As explained in the previous section, we focus on the characterization of SZK by its complete problem: the statistical difference problem **SD** [SV03]. We demonstrate oracles (Ψ, \mathcal{A}) such that relative to Ψ there exist constructions of one-way permutations (OWPs) and IO for circuits with OWP gates, and these constructions are secure in the presence of \mathcal{A} . At the same time, \mathcal{A} will decide (in the worst-case) SD^{Ψ} . Since **SD** is complete for SZK in a relativizing manner, deciding SD^{Ψ} suffices to break SZK^{Ψ} . That is, \mathcal{A} will decide *all* instances (C_0^{Ψ}, C_1^{Ψ}) of circuit samplers that only use the IO and OWPs realized by Ψ in a black-box manner. We next explain how each of the two are constructed.

The construction of Ψ follows a general recipe suggested in [AS15, AS16]. The oracle consists of three parts $(f, \mathcal{O}, \text{Eval}^{f, \mathcal{O}})$ where:

1. f is a random permutation, realizing the one-way permutation primitive.
2. \mathcal{O} is a random injective function, realizing the obfuscation algorithm. It takes as input an oracle-aided circuit $C^{(\cdot)}$ along with randomness r and outputs an obfuscation $\widehat{C} = \mathcal{O}(C, r)$.
3. $\text{Eval}^{\mathcal{O}, f}$ realizes evaluation of obfuscated circuits. On input (\widehat{C}, x) , it inverts \mathcal{O} to find (C, r) , and outputs $C^f(x)$. If \widehat{C} is not in the image of \mathcal{O} , it returns \perp .

The above construction readily satisfies the syntactic (or “functionality”) requirements of one-way permutations and indistinguishability obfuscation. Furthermore, using standard techniques, it is not hard to show that relative to Ψ , the function f is one-way and \mathcal{O} satisfies IO indistinguishability requirement. The challenge is to now come up with an oracle \mathcal{A} that, on one hand, will decide SD^{Ψ} , but on the other, will not compromise the security of the latter primitives.

Recall that deciding SD^{Ψ} means that given two oracle-aided circuit samplers (C_0, C_1) such that the statistical distance of the corresponding distributions (C_0^{Ψ}, C_1^{Ψ}) is $s = \Delta(C_0^{\Psi}, C_1^{\Psi}) \in [0, \frac{1}{3}] \cup [\frac{2}{3}, 1]$, the oracle \mathcal{A} must decide in which of the two intervals s lies, whereas if the promise is not satisfied and $s \in (\frac{1}{3}, \frac{2}{3})$, there is no requirement whatsoever. With this in mind, a first naive attempt would be the following. \mathcal{A} will have unbounded access to Ψ , give a query (C_0, C_1) , it would compute $s = \Delta(C_0, C_1)$, and simply say whether $s < \frac{1}{2}$ or $s \geq \frac{1}{2}$. While such an oracle would definitely decide SD^{Ψ} , it is not too hard to show that it is simply too powerful, and would not only break IO and OWPs, but

would, in fact, allow solving any problem in NP^Ψ (or even in PP^Ψ). Other naive attempts such as refusing to answer outside the promise intervals, encounter a similar problem.

At high-level, the problem with such oracles is that solutions to hard problems can be easily correlated with “tiny” differences in the statistical distance of the two input circuits, whereas the above oracle may reflect tiny changes when the statistical distance is close to some threshold ($1/2$ in the above example) on which the oracle changes its behaviour. This motivates our actual definition of \mathcal{A} as a *noisy oracle* that produces its answer, not according to some fixed threshold, but according to a random threshold, chosen afresh for each and every query. Concretely, the oracle, which we call StaDif^Ψ , for any query (C_0, C_1) , chooses a uniformly random threshold $t \leftarrow (\frac{1}{3}, \frac{1}{3})$, and answers accordingly:

$$\text{StaDif}^\Psi(C_0, C_1) = \begin{cases} Y & \text{if } s \geq t \text{ (far distributions)} \\ N & \text{if } s < t \text{ (similar distributions)} \end{cases} .$$

The main challenge in proving that the security of the IO and OWPs realized by \mathcal{A} is not compromised by this oracle is that StaDif^Ψ has the power to query Ψ on exponentially many points in order to compute s . For instance, it may query Ψ on the preimage of a OWP challenge $f(x)$ or of a given obfuscation $\mathcal{O}(C, r)$. The key observation behind the proof is that the oracle’s final answer still does not reflect how Ψ behaves locally on random points.

Intuitively, choosing the threshold t at random, for each query (C_0, C_1) , guarantees that with high probability t is “far” from the corresponding statistical distance $s = \Delta(C_0^\Psi, C_1^\Psi)$. Thus, changing the oracle Ψ on, say, a single input x , such as the preimage of an OWP challenge $f(x)$, should not significantly change s and will not affect the oracle’s answer; that is, unless the circuits query Ψ on x with high probability to begin with. We give a reduction showing that we can always assume that (C_0, C_1) are “smooth”, in the sense that they do not make any specific query to Ψ with too high probability.

Following this intuition, we are able to show that through such local changes that go undetected by StaDif^Ψ , we can move to an ideal world where inverting the OWP or breaking IO can be easily shown to be impossible. We refer the reader to the full version for further details.

Ruling out Black-Box Constructions of Hard $\text{NP} \cap \text{coNP}$ Problems. As mentioned earlier, a fully black-box construction of hard problems in $\text{NP} \cap \text{coNP}$ is actually known assuming one-way permutations (OWPs), and cannot be ruled out as in the case of SZK. Instead, we rule out constructions from (non-surjective) injective one-way functions (IOWFs) and IO for circuits with IOWF gates. This generalizes several previous results by Blum and Impagliazzo [BI87] and Rudich [Rud88], showing that OWFs do not give hardness in $\text{NP} \cap \text{coNP}$, by Matsuda and Matsuura [MM11], showing that IOWFs do not give OWPs (which are a special case of hardness $\text{NP} \cap \text{coNP}$), and by Asharov and Segev [AS16], showing that OWFs and IO for circuits with OWF gates do not give OWPs. In fact, our approach yields a new (and rather different) proof for each one of these results.

We follow a similar methodology to one we used for the case of SZK. That is, we would like to come up with oracles (Ψ, \mathcal{A}) such that Ψ realizes IOWFs and IO for circuits with IOWFs gates, which are both secure in the presence of \mathcal{A} , whereas black-box constructions of problems in $\text{NP} \cap \text{coNP}$ from these primitives can be easily solved by \mathcal{A} . By black-box constructions here we mean a pair of efficient oracle-aided non-deterministic verifiers $V_0^{(\cdot)}, V_1^{(\cdot)}$ that for every oracle Ψ implementing IOWFs and IO, yield co-languages \bar{L}^Ψ, L^Ψ in $\text{NP}^\Psi \cap \text{coNP}^\Psi$.

The requirement that V_0, V_1 give a language in $\text{NP} \cap \text{coNP}$ for *every* oracle implementing IOWFs and IO follows previous modeling [BI87],⁵ and aligns with how we usually think about *correctness* of black-box constructions of cryptographic primitives. For instance, the construction of public-key encryption from trapdoor permutations is promised to be correct, for all oracles implementing the trapdoor permutation. Similarly, the construction of hard $\text{NP} \cap \text{coNP}$ languages from one-way permutations, give an $\text{NP} \cap \text{coNP}$ language for any oracle implementing a permutation.⁶

We stress that a construction where correctness is only guaranteed for particular (even if natural) oracles may definitely exist. This is for example the case if we only consider implementations of IO similar to those presented above in the context of SZK. Indeed, in that construction the implementation of IO has an additional property — it allows identifying *invalid obfuscations* (the Eval oracle would simply return \perp on such obfuscations). This “verifiability” property coupled with the injectivity of obfuscators actually imply a hard problem in $\text{NP} \cap \text{coNP}$ in a black-box way.⁷ Our separation thus leverages the fact that IO need not necessarily be verifiable, and rules out constructions that are required to be correct for any implementation of IO, even a non-verifiable one.

Accordingly, the oracles $\Psi = (f, \mathcal{O}, \text{Eval}^{f, \mathcal{O}})$ that we consider are a tweaked version of the oracles considered in the SZK case. Now f is a random injective function that is expanding, rather than a permutation, the oracle \mathcal{O} is defined as before, and the oracle $\text{Eval}^{f, \mathcal{O}}$ is defined as before for valid obfuscations $\hat{C} \in \text{Image}(\mathcal{O})$ but is allowed to act arbitrarily for invalid obfuscations. As for \mathcal{A} , this time it is trivially implemented by an oracle Decide^Ψ that, given input x , simply returns the unique bit b such that $V_b(x) = 1$, namely it just decides the corresponding language L^Ψ .

In the results mentioned above [Rud88, MM11, AS16], it is actually shown that any query to such an oracle can be completely simulated with a small

⁵ Rudich [Rud88] also considered a slight relaxation of constructions that are correct for an overwhelming fraction of oracles rather than all.

⁶ We note that this issue does not come up for black-box constructions of SZK *promise* problems, because the construction is allowed to yield instances that do not obey the promise; there correctness is always guaranteed, and the only question is whether the instances that do satisfy the promise are hard to decide.

⁷ E.g. the language of all valid obfuscations and indices i , such that the i th bit of the obfuscated circuit is 1

number of queries to Ψ .⁸ We do not show such a simulation process. Instead, we take a different approach inspired by our proof for the SZK setting described above. Roughly speaking, we show that somewhat similarly to our statistical difference oracle StaDif^Ψ , the oracle Decide^Ψ is also rather robust to random local changes. The main observation here is that for any fixed yes-instance $x \in L^\Psi$, tweaking Ψ at a random input into a new oracle Ψ' , it is likely that x will still be a yes-instance in $L^{\Psi'}$, as long as Ψ' is in our allowed family of oracles and $L^{\Psi'}$ is indeed in $\text{NP}^{\Psi'} \cap \text{coNP}^{\Psi'}$ (and the same is true for no-instances).

In slightly more detail, fixing a witness w such that $V_1^\Psi(x, w) = 1$, we can show that since V_1 makes a small number of oracle calls, with high probability tweaking the oracle Ψ at a random place will not affect these oracle calls and thus $V_1^{\Psi'}(x, w) = V_1^\Psi(x, w) = 1$. Then, assuming $L^{\Psi'}$ is guaranteed to be in $\text{NP} \cap \text{coNP}$, we can deduce that x must still a yes-instance (other witnesses for this fact may be added or disappear, but this does not change the oracle's answer). In the body, we argue that indeed $L^{\Psi'} \in \text{NP}^{\Psi'} \cap \text{coNP}^{\Psi'}$, where we strongly rely on the fact that arbitrary behavior of Eval is permitted on invalid obfuscations.

Once again, we show that through local changes that go undetected by Decide^Ψ , we can move to an ideal world where inverting the IOWF or breaking IO can be easily shown to be impossible. We refer the reader to Section 3 for further details.

Implied Separations. As a result of the two separations discussed above, we can rule out black-box constructions of hard problems in SZK or $\text{NP} \cap \text{coNP}$ from various cryptographic primitives or complexity classes. This essentially includes all primitives that have fully black-box constructions from OWPs (or IOWFs) and IO for circuits with OWP (or IWOFF) gates. This includes public-key encryption, oblivious transfer, deniable encryption [SW14], functional encryption [Wat15], delegation, [BGL⁺15, CHJV15, KLV15], hard (on-average) PPAD instances [BPR15], and more.

We note that there are a few applications of IO that do not fall under this characterization. For instance, the construction of IO for Turing machines from IO-based succinct randomized encodings [BGL⁺15, CHJV15, KLV15] involves obfuscating a circuit that itself outputs (smaller) obfuscated circuits. To capture this, we would need to extend the above model to IO for circuits that can also make IO oracle calls (on smaller circuits). Another example is the construction of non-interactive witness indistinguishable proofs from IO [BP15]. There an obfuscated circuit may get as input another obfuscated circuit and would have to internally run it; furthermore, in this application, the code of the obfuscator is used in a (non-black-box) ZAP. Extending the above model to account for this type of IO applications is an interesting question that we leave for future exploration.

⁸ More accurately, this is the case for Rudich's result for $\text{NP} \cap \text{coNP}$, whereas for the other results that rule out constructions of one-way permutations, one can simulate an analog of Decide that inverts the permutation.

The Positive Result: Collision-Resistance from IO and SZK-Hardness. We now described the main ideas behind our construction of collision-resistant hash functions. The starting point for the construction is the work of Ishai, Kushilevitz, and Ostrovsky [IKO05] that shows how to construct collision-resistant hash functions from commitments that are additively homomorphic (for simplicity, say over \mathbb{F}_2). The idea is simple: we can hash ℓ bits to m bits, where m is the size of a single bit commitment and ℓ can be arbitrarily longer, as follows. The hash key is a commitment $\gamma := (\text{com}(\beta_1), \dots, \text{com}(\beta_\ell))$ to a random vector $\beta \in \mathbb{F}_2^\ell$, and hashing $x \in \mathbb{F}_2^\ell$, is done by homomorphically computing a commitment to the inner product $\text{CRH}_\gamma(x) = \text{com}(\langle \beta, x \rangle)$. Intuitively, the reason this works is that any collision in CRH_γ reveals a vector that is orthogonal to β and thus leaks information about it and violating the hiding of the commitment.

At a high-level, we aim to mimic the above construction based on obfuscation. As a key for the collision-resistant hash we can obfuscate a program Π_β associated with a random vector β that given x outputs a commitment $\text{com}(\langle \beta, x \rangle)$, where the commitment is derandomized using a PRF.⁹ The obfuscation $i\mathcal{O}(\Pi_\beta)$ can be thought of as the commitment to β , and evaluating this program at x , corresponds to homomorphic evaluation. Despite the clear intuition behind this construction, it is not clear how to prove its security based on IO. In fact, by the work of Asharov and Segev [AS15], it cannot be proven based on a black-box reduction as long as plain statistically-binding commitments are used, as these can be constructed from OWPs in a fully black-box manner, and [AS15] rule out black-box constructions of collision-resistant hashing from OWPs and IO for circuits with OWP gates.

We show, however, that relying on a relaxed notion of perfectly-hiding commitments, as well as subexponential hardness of IO and puncturable PRFs, the construction can be proven secure. The perfect hiding of the commitment is leveraged in a probabilistic IO argument [CLTV15] that involves a number of hybrids larger than the overall number of commitments. We then observe that these relaxed commitments follow from average-case hardness of the polar statistical difference problem $\text{SD}^{0,1}$.¹⁰

2 One-Way Permutations, Indistinguishability Obfuscation, and Hardness in SZK

In this section, we ask which cryptographic primitives imply hardness in the class statistical zero-knowledge (SZK). Roughly speaking, we show that one-way permutations (OWPs) and indistinguishability obfuscation (IO), for circuits with OWP-gates, do not give rise to a black-box construction of hard problems in SZK. This, in turn implies that many cryptographic primitives (e.g., public-key

⁹ In the body, we describe a slightly more abstract construction where inner product is replaced by an arbitrary 2-universal hash function.

¹⁰ Similar SZK-hardness is known to imply statistically-hiding commitments against malicious receivers, but with a larger (constant) number of rounds [OV08].

encryption, functional encryption, and delegation), and hardness in certain low-level complexity classes (e.g. PPAD), also do not yield black-box constructions of hard problems in SZK.

We first motivate and define a framework of SZK relative to oracles, define fully black-box constructions of hard SZK problems, and then move on to the actual separation.

2.1 SZK and Statistical Difference

The notion of statistical zero-knowledge proofs was introduced in the seminal work of Goldwasser, Micali and Rackoff [GMR85]. The class of promise problems with statistical zero-knowledge proofs (SZK) can be characterized by several complete problems, such as *statistical difference* [SV03] and *entropy difference* [GV99] (see also [Vad99] and references within). We shall focus on the characterization of SZK by the statistical difference problem. Here an instance is a pair of circuit samplers $C_0, C_1 : \{0, 1\}^n \rightarrow \{0, 1\}^m$ with the promise that the statistical distance $s = \Delta(C_0, C_1)$ of the corresponding distributions is either large (say, $s \geq 2/3$) or small (say, $s \leq 1/3$). The problem is to decide which is the case.

Hard Statistical Difference Problems from Cryptography: Motivation. SZK hardness, and in particular hard statistical difference problems, are known to follow from various number-theoretic and lattice problems that are commonly used in cryptography, such as Decision Diffie-Hellman, Quadratic Residuosity, and Learning with Errors. We ask more generally which cryptographic primitives can be shown to imply such hardness, with the intuition that such primitives are *structured* in a certain way. In particular, whereas one would not expect a completely unstructured object like one-way functions to imply such hardness, what can we say for instance about public-key encryption, or even indistinguishability obfuscation (which has proven to be structured enough to yield almost any known cryptographic goal).

We prove that none of these primitives imply such hardness through the natural class of black-box constructions and security reductions. To understand what a black-box construction of a hard statistical difference problem means, let us look at a specific example of the construction of such a problem from *rerandomizable encryption*. In a (say, symmetric-key) rerandomizable encryption scheme, on top of the usual encryption and decryption algorithms (Enc, Dec) there is a ciphertext rerandomization algorithm ReRand that can statistically refresh ciphertexts. Namely, for any ciphertext CT encrypting a bit b , $\text{ReRand}(\text{CT})$ produces a ciphertext that is statistically close to a fresh encryption $\text{Enc}(b)$. Note that this immediately gives rise to a hard statistical difference problem: given a pair of ciphertexts (CT, CT') , decide whether the corresponding rerandomized distributions given by the circuits $(C_0(\cdot), C_1(\cdot)) := (\text{ReRand}(\text{CT}; \cdot), \text{ReRand}(\text{CT}'; \cdot))$ are statistically far or close. Indeed, this corresponds to whether they encrypt the same bit or not, which is hard to decide by the security of the encryption scheme.

A feature of this construction of hard statistical difference instances is that, similarly to most constructions in cryptography, it is *fully black-box* [RTV04] in the sense that the circuits C_0, C_1 only make black-box use of the encryption scheme’s algorithms, and can in fact be represented as oracle-aided circuits $(C_0^{\text{ReRand}(\cdot)}, C_1^{\text{ReRand}(\cdot)})$. Furthermore, “hardness” can be shown by a black-box reduction that can use any decider for the problem in a black-box way to break the underlying encryption scheme. More generally, one can consider the statistical difference problem relative to different oracles implementing different cryptographic primitives and ask when can hardness be shown based on a black-box reduction. We will rule out such reductions relative to IO and OWPs (and everything that follows from these in a fully black-box way).

2.2 Fully Black-Box Constructions of Hard SD Problems from IO and OWPs

We start by defining statistical difference problem relative to oracles. In the following definition, for an oracle-aided (sampler) circuit $C^{(\cdot)}$ with n -bit input and an oracle Ψ , we denote by \mathbf{C}^Ψ the output distribution $C^\Psi(r)$ where $r \leftarrow \{0, 1\}^n$. For two distributions \mathbf{X} and \mathbf{Y} we denote their statistical distance by $\Delta(\mathbf{X}, \mathbf{Y})$.

Definition 2.1 (Statistical difference relative to oracles). *For an oracle Ψ , the statistical difference promise problem relative to Ψ , denoted as $\mathbf{SD}^\Psi = (\mathbf{SD}_Y^\Psi, \mathbf{SD}_N^\Psi)$, is given by*

$$\begin{aligned} \mathbf{SD}_Y^\Psi &= \left\{ (C_0, C_1) \mid \Delta(\mathbf{C}_0^\Psi, \mathbf{C}_1^\Psi) \geq \frac{2}{3} \right\}, \\ \mathbf{SD}_N^\Psi &= \left\{ (C_0, C_1) \mid \Delta(\mathbf{C}_0^\Psi, \mathbf{C}_1^\Psi) \leq \frac{1}{3} \right\}. \end{aligned}$$

We now formally define the class of constructions and reductions ruled out. That is, *fully black-box* constructions of hard statistical distance problems from OWPs and IO for OWP-aided circuits. The definition is similar in spirit to those in [AS15, AS16], adapted to our context of SZK-hardness.

Definition 2.2. *A fully black-box construction of a hard statistical distance problem from OWPs and IO for the class \mathcal{C} of circuits with OWP-gates consists of a collection of oracle-aided circuit pairs $\Pi^{(\cdot)} = \left\{ \Pi_n^{(\cdot)} = \left\{ (C_0^{(\cdot)}, C_1^{(\cdot)}) \in \{0, 1\}^{n \times 2} \right\} \right\}_{n \in \mathbb{N}}$ and a probabilistic oracle-aided reduction \mathcal{R} that satisfy:*

- **Black-box security proof:** *There exist functions $q_{\mathcal{R}}(\cdot), \varepsilon_{\mathcal{R}}(\cdot)$ such that the following holds. Let f be any distribution on permutations and let $i\mathcal{O}$ be any distribution on functions such that $\tilde{C}^f \equiv C^f$ for any $C^{(\cdot)}$ and r , where $\tilde{C}^{(\cdot)} := i\mathcal{O}(C^{(\cdot)}, r)$. Then for any probabilistic oracle-aided \mathcal{A} that decides Π in the worst-case, namely, for all $n \in \mathbb{N}$*

$$\Pr_{f, i\mathcal{O}, \mathcal{A}} \left[\mathcal{A}^{f, i\mathcal{O}}(C_0, C_1) = B \quad \text{for all} \quad \begin{array}{l} (C_0, C_1) \in \Pi_n, B \in \{Y, N\} \\ \text{such that } (C_0, C_1) \in \mathbf{SD}_B^{f, i\mathcal{O}} \end{array} \right] = 1$$

the reduction breaks either f or $i\mathcal{O}$, namely, for infinitely many $n \in \mathbb{N}$ either

$$\Pr_{\substack{x \leftarrow \{0,1\}^n \\ f, i\mathcal{O}, \mathcal{A}}} [\mathcal{R}^{\mathcal{A}, f, i\mathcal{O}}(f(x)) = x] \geq \varepsilon_{\mathcal{R}}(n) ,$$

or

$$\left| \Pr \left[\text{Exp}_{(f, i\mathcal{O}), i\mathcal{O}, \mathcal{C}, \mathcal{R}^{\mathcal{A}}}^{\text{IO}}(n) = 1 \right] - \frac{1}{2} \right| \geq \varepsilon_{\mathcal{R}}(n) ,$$

where in both \mathcal{R} makes at most $q_{\mathcal{R}}(n)$ queries to any of its oracles $(\mathcal{A}, f, i\mathcal{O})$, and any query $(C_0^{(\cdot)}, C_1^{(\cdot)})$ it makes to \mathcal{A} consists of circuits that also make at most $q_{\mathcal{R}}(n)$ queries to their oracles $(f, i\mathcal{O})$. The random variable $\text{Exp}_{(f, i\mathcal{O}), i\mathcal{O}, \mathcal{C}, \mathcal{R}^{\mathcal{A}}}^{\text{IO}}(n)$ represents the reductions winning probability in the IO security game relative to $(f, i\mathcal{O})$.

We make several remarks about the definition:

- **Correctness.** Typically, we also require certain *correctness* from the black-box construction. For instance, in the next section, we shall require that the construction always satisfies the $\text{NP} \cap \text{coNP}$ structure. In the above definition, the construction is allowed to yield instances $(C_0^{f, i\mathcal{O}}, C_1^{f, i\mathcal{O}})$ that do not satisfy the SZK promise; namely $(C_0^{f, i\mathcal{O}}, C_1^{f, i\mathcal{O}}) \notin \text{SD}_Y^{f, i\mathcal{O}} \cup \text{SD}_N^{f, i\mathcal{O}}$. It is natural to think of more stringent definitions that require that the corresponding problem $\Pi^{f, i\mathcal{O}}$ is non-trivial, in the sense that $\Pi^{f, i\mathcal{O}} \cap \text{SD}_Y^{f, i\mathcal{O}} \neq \emptyset$ and $\Pi^{f, i\mathcal{O}} \cap \text{SD}_N^{f, i\mathcal{O}} \neq \emptyset$ (which is the case for known constructions of SZK hardness from cryptographic primitives). Our impossibility is more general and would, in particular, rule out such definitions as well.
- **Worst-Case vs. Average-Case Hardness.** In the above, we address *worst-case hardness*, in the sense that the reduction \mathcal{R} has to break the underlying primitives only given a decider \mathcal{A} that is always correct. One could further ask whether IO and OWPs even imply average-case hardness in SZK (as do many of the algebraic hardness assumptions in cryptography). Ruling out worst-case hardness (as we will do shortly) in particular rules out such average-case hardness as well.
- **IO for Oracle-Aided Circuits.** Following [AS15, AS16], we consider indistinguishability obfuscation for oracle-aided circuits C^f that can make calls to the one-way permutation oracle. This model captures constructions where IO is applied to circuits that use pseudo-random generators, puncturable pseudo-random functions, or injective one-way functions as all of those have fully black-box constructions from one-way permutations (see further discussion in [AS15]). This includes almost all known constructions from IO, including public-key encryption, deniable encryption [SW14], functional encryption [Wat15], delegation [BGL⁺15, CHJV15, KLV15], and hard (on-average) PPAD instances [BPR15]. Accordingly, separating SZK from IO and OWPs in this model, results in a similar separation between SZK and any one of these primitives.

We note that there are a few applications though that do not fall under this model. The first is in applications where the obfuscated circuit might itself output (smaller) obfuscated circuit, for instance in the construction of IO for Turing machines from IO-based succinct randomized encodings [BGL⁺15, CHJV15, KLW15]. To capture such applications, one would have to extend the model to also account for circuits with IO gates (and not only OWP gates). A second example is the construction of non-interactive witness indistinguishable proofs from IO [BP15]. There an obfuscated circuit may get as input another obfuscated circuit and would have to internally run it; furthermore, in this application, the code of the obfuscator is used in a (non-black-box) ZAP. Extending our results (and those of [AS15, AS16]) to these models is an interesting question, left for future work.

- **Security Loss.** In the above definition the functions $q_{\mathcal{R}}$ and $\varepsilon_{\mathcal{R}}$ capture the *security loss* of the reduction. Most commonly in cryptography, the query complexity is polynomial $q_{\mathcal{R}}(n) = n^{O(1)}$ and the probability of breaking the underlying primitive is inverse polynomial $\varepsilon_{\mathcal{R}}(n) = n^{-O(1)}$. Our lower-bounds will in-fact apply for *exponential* $q_{\mathcal{R}}, \varepsilon_{\mathcal{R}}^{-1}$. This allows capturing also constructions that rely on subexponentially secure primitives (e.g., [BGL⁺15, CHJV15, KLW15, BPR15, BPW16]).

Ruling Out Fully Black-Box Constructions: A Road Map. Our main result in this section is that a fully black-box construction of a hard statistical difference problem from IO and OWPs does not exist. Furthermore, this holds even if the latter primitives are exponentially secure.

Theorem 2.3. *Any fully black-box construction of a statistical difference problem Π from OWPs and IO for circuits with OWP gates has an exponential security loss: $\max(q_{\mathcal{R}}(n), \varepsilon_{\mathcal{R}}^{-1}(n)) \geq \Omega(2^{n/12})$.*

The proof of the theorem follows a common methodology (applied for instance in [HR04, HHR15, AS15]). We exhibit two (distributions on) oracles $(\Psi, \text{StaDif}^{\Psi})$, where Ψ realizes OWPs and IO for circuits with OWP gates, and StaDif^{Ψ} that decides SD^{Ψ} , the statistical difference problem relative to Ψ , in the worst case. Since SD is complete for SZK in a relativizing manner, solving SD^{Ψ} suffices to break SZK^{Ψ} . We then show that the primitives realized by Ψ are (exponentially) secure even in the presence of StaDif^{Ψ} . Then viewing StaDif as a worst-case decider \mathcal{A} (as per Definition 2.2) directly implies Theorem 2.3, ruling out fully black-box constructions with a subexponential security loss. We defer the oracle description and the proof to the full version.

3 One-Way Functions, Indistinguishability Obfuscation, and Hardness in $\text{NP} \cap \text{coNP}$

In this section, we show that injective one-way functions (IOWFs) and indistinguishability obfuscation (IO), for circuits with IOWF-gates, do not give rise to a black-box construction of hard problems in $\text{NP} \cap \text{coNP}$. This can be seen as a

generalization of previous separations by Rudich [Rud88], showing that OWFs do not give hardness in $\text{NP} \cap \text{coNP}$, by Matsuda and Matsuura [MM11], showing that IOWFs do not give one-way permutations (which are a special case of hardness $\text{NP} \cap \text{coNP}$), and by Asharov and Segev [AS16], showing that OWFs and IO do not give one-way permutations. As in the previous section, the result implies that many cryptographic primitives and hardness in PPAD, also do not yield black-box constructions of hard problems in $\text{NP} \cap \text{coNP}$.

We first define the framework of $\text{NP} \cap \text{coNP}$ relative to oracles, define fully black-box constructions of hard $\text{NP} \cap \text{coNP}$ problems, and then move on to the actual separation.

3.1 $\text{NP} \cap \text{coNP}$

Throughout, we shall canonically represent languages $L \in \text{NP} \cap \text{coNP}$ by their corresponding non-deterministic poly-time verifiers V_1, V_0 , where

$$\begin{aligned} L &= \{x \in \{0, 1\}^* \mid \exists w : V_1(x, w) = 1\} \text{ ,} \\ \bar{L} &= \{x \in \{0, 1\}^* \mid \exists w : V_0(x, w) = 1\} = \{0, 1\}^* \setminus L \text{ .} \end{aligned}$$

Hardness in $\text{NP} \cap \text{coNP}$ from Cryptography - Motivation. Hard (on average) problems in $\text{NP} \cap \text{coNP}$ are known to follow based on certain number-theoretic problems in cryptography, such as Discrete Log and Factoring. As in the previous section for SZK, we are interested in understanding which cryptographic primitives would imply such hardness, again with the intuition that these should be appropriately structured. For instance, it is known [Bra79] that any one-way permutation $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$ implies a hard problem in $\text{NP} \cap \text{coNP}$, e.g. given an index $i \in [n]$ and an image $f(x)$ find the i -th pre-image bit x_i . In contrast, in his seminal work, Rudich [Rud88] proved that completely unstructured objects like one-way functions cannot construct even worst-case hard instances by fully black-box constructions. Here a fully black-box construction essentially means that the non-deterministic verifiers only make black-box use of the OWF (or OWP in the previous example) and the reduction establishing the hardness is also black-box (in both the adversary and the OWF).

But what about more structured primitives such as public-key encryption, oblivious transfer, or even indistinguishability obfuscation. Indeed, IO (plus OWFs) has been shown to imply hardness in PPAD and more generally in the class TFNP of total search problems, which is often viewed as the search analog of $\text{NP} \cap \text{coNP}$ [MP91]. We will show, however, that fully black-box constructions do not give rise to a hard problem in $\text{NP} \cap \text{coNP}$ from OWFs (or even injective OWFs) and IO for circuits with OWF gates.

3.2 Fully Black-Box Constructions of Hardness in $\text{NP} \cap \text{coNP}$ from IO and IOWFs

We start by defining $\text{NP} \cap \text{coNP}$ relative to oracles [Rud88]. This, in particular, captures black-box constructions of such languages from cryptographic primi-

tives, such as one-way functions in [Rud88] or indistinguishability obfuscation, which we will consider in this work.

Definition 3.1 (NP \cap coNP relative to oracles). Let \mathfrak{S} be a family of oracles and let $V_1^{(\cdot)}, V_0^{(\cdot)}$ be a pair of oracle-aided non-deterministic polynomial-time verifiers. We say that V_1, V_0 define a collection of languages $L^\mathfrak{S} = \{L^\Gamma \mid \Gamma \in \mathfrak{S}\}$ in NP \cap coNP relative to \mathfrak{S} if for any $\Gamma \in \mathfrak{S}$, the machines V_1^Γ, V_0^Γ define a language $L^\Gamma \in \text{NP}^\Gamma \cap \text{coNP}^\Gamma$. That is

$$\begin{aligned} L^\Gamma &= \{x \in \{0, 1\}^* \mid \exists w : V_1^\Gamma(x, w) = 1\} \quad , \\ \bar{L}^\Gamma &= \{x \in \{0, 1\}^* \mid \exists w : V_0^\Gamma(x, w) = 1\} = \{0, 1\}^* \setminus L \quad . \end{aligned}$$

We now formally define the class of constructions and reductions ruled out. That is, *fully black-box* constructions of hard problems in NP \cap coNP from injective one-way functions (IOWFs) and IO for IOWF-aided circuits. The definition is similar in spirit to those in [AS15, AS16] and in the Section 2, adapted to the context of NP \cap coNP hardness.

Definition 3.2. A *fully black-box construction* of a hard NP \cap coNP problem L from IOWFs and IO for the class \mathcal{C} of circuits with IOWF-gates is given by two oracle aided poly-time machines (V_0, V_1) and a probabilistic oracle-aided reduction \mathcal{R} that satisfy:

1. **Structure:** Let \mathfrak{S} be the family of all oracles $(f, i\mathcal{O})$ such that f is injective and $i\mathcal{O}$ is a function such that $\widehat{C}^f \equiv C^f$ for any $C^{(\cdot)} \in \mathcal{C}$, r , and $\widehat{C}^{(\cdot)} := i\mathcal{O}(C, r)$. Then (V_0, V_1) define a language $L^{f, i\mathcal{O}} \in \text{NP}^{f, i\mathcal{O}} \cap \text{coNP}^{f, i\mathcal{O}}$ relative to any oracle $(f, i\mathcal{O}) \in \mathfrak{S}$ (as per Definition 3.1).
2. **Black-box security proof:** There exist functions $q_{\mathcal{R}}(\cdot), \varepsilon_{\mathcal{R}}(\cdot)$ such that the following holds. Let $(f, i\mathcal{O})$ be any distribution supported on the family \mathfrak{S} defined above. Then for any probabilistic oracle-aided \mathcal{A} that decides $L^{f, i\mathcal{O}}$ in the worst-case, namely, for all $n \in \mathbb{N}$

$$\Pr_{f, i\mathcal{O}, \mathcal{A}} \left[\mathcal{A}^{f, i\mathcal{O}}(x) = b \quad \text{for all} \quad \begin{array}{l} x \in \{0, 1\}^n, b \in \{0, 1\} \\ \text{such that } V_b(x) = 1 \end{array} \right] = 1$$

the reduction breaks either f or $i\mathcal{O}$, namely, for infinitely many $n \in \mathbb{N}$ either

$$\Pr_{\substack{x \leftarrow \{0, 1\}^n \\ f, i\mathcal{O}, \mathcal{A}}} [\mathcal{R}^{\mathcal{A}, f, i\mathcal{O}}(f(x)) = x] \geq \varepsilon_{\mathcal{R}}(n) \quad ,$$

or

$$\left| \Pr \left[\text{Exp}_{(f, i\mathcal{O}), i\mathcal{O}, \mathcal{C}, \mathcal{R}^{\mathcal{A}}}^{\text{IO}}(n) = 1 \right] - \frac{1}{2} \right| \geq \varepsilon_{\mathcal{R}}(n) \quad ,$$

where in both \mathcal{R} makes at most $q_{\mathcal{R}}(n)$ queries to any of its oracles $(\mathcal{A}, f, i\mathcal{O})$, and for any query x made to \mathcal{A} , the non-deterministic verifiers $V_0^{f, i\mathcal{O}}(x), V_1^{f, i\mathcal{O}}(x)$ make at most $q_{\mathcal{R}}(n)$ queries to their oracles (for any non-deterministic choice of a witness w). The random variable $\text{Exp}_{(f, i\mathcal{O}), i\mathcal{O}, \mathcal{C}, \mathcal{R}^{\mathcal{A}}}^{\text{IO}}(n)$ represents the reductions winning probability in the IO security game relative to $(f, i\mathcal{O})$.

Remark about Correct Structure. We note that here we explicitly do put a *correctness* requirement, which we refer to as *structure*; namely, that the construction yields a language in $\text{NP} \cap \text{coNP}$ for any implementation of OWPs and IO. This is different from the setting from Definition 2.2 where we considered *promise problems* and allowed the construction not to satisfy the promise occasionally.

Concretely, we require that V_0, V_1 give a language in $\text{NP} \cap \text{coNP}$ for *every* oracle implementing IOWFs and IO. This follows the modeling of [BI87],¹¹ and aligns with how we usually think about *correctness* of black-box constructions of cryptographic primitives. For instance, the construction of public-key encryption from trapdoor permutations is promised to be correct, for all oracles implementing the trapdoor permutation. Similarly, the construction of hard $\text{NP} \cap \text{coNP}$ languages from one-way permutations, give an $\text{NP} \cap \text{coNP}$ language for any oracle implementing a permutation.

We also note that as in Definition 2.2, our definition addresses *worst-case hardness*, which makes our impossibility result stronger. See further discussion after Definition 2.2.

Ruling out Fully Black-Box Constructions: A Road Map. Our main result in this section is that fully black-box constructions of a hard $\text{NP} \cap \text{coNP}$ problem from IO and IOWFs do not exist. Furthermore, this holds even if the latter primitives are exponentially secure.

Theorem 3.3. *Any fully black-box construction of an $\text{NP} \cap \text{coNP}$ problem L from IOWFs and IO for circuits with IOWF gates has an exponential security loss:*

$$\max(q_{\mathcal{R}}(n), \varepsilon_{\mathcal{R}}^{-1}(n)) \geq \Omega(2^{n/6})$$

The proof of the theorem follows a similar methodology to the proof of Theorem 2.3. We exhibit two (distributions on) oracles $(\Psi, \text{Decide}^{\Psi})$, where Ψ realizes IOWFs and IO for circuits with IOWF gates, and Decide^{Ψ} that decides $L^{\Psi} \in \text{NP}^{\Psi} \cap \text{coNP}^{\Psi}$ in the worst case. We then show that the primitives realized by Ψ are (exponentially) secure even in the presence of Decide^{Ψ} . Then viewing Decide as a worst-case decider \mathcal{A} (as per Definition 3.2) directly implies Theorem 3.3, ruling out fully black-box constructions with a subexponential security loss.

We defer the formal treatment to the full version.

4 Collision-Resistance from IO and SZK-Hardness

Asharov and Segev [AS15] showed that collision-resistant hashing cannot be constructed from (even subexponentially hard) indistinguishability obfuscation (IO) and one-way permutations (OWPs) relying on common IO techniques. Slightly more accurately, they rule out fully black-box constructions where (as in previous sections) IO is defined with respect to circuits with OWP oracle gates. In this section, we show that, assuming IO and a strong form of SZK-hardness, there is indeed a construction of collision-resistant hashing (CRH).

¹¹ Rudich [Rud88] also considered a slight relaxation of constructions that are correct for an overwhelming fraction of oracles rather than all.

The High-Level Idea Behind the Construction. The starting point for our construction is the work of Ishai, Kushilevitz, and Ostrovsky [IKO05] that shows how to construct collision-resistant hash functions from commitments that are additively homomorphic (for simplicity, say over \mathbb{F}_2). The idea is simple: we can hash ℓ bits to m bits, where m is the size of a single bit commitment and ℓ can be arbitrarily longer, as follows. The hash key is a commitment $\gamma := (\text{com}(\beta_1), \dots, \text{com}(\beta_\ell))$ to a random vector $\beta \in \mathbb{F}_2^\ell$, and hashing $x \in \mathbb{F}_2^\ell$, is done by homomorphically computing a commitment to the inner product $\text{CRH}_\gamma(x) = \text{com}(\langle \beta, x \rangle)$.

This idea can, in fact, be abstracted to work with any commitment scheme wherein given a commitment $\text{com}(\beta)$ for a random key for a 2-universal hash allows to homomorphically compute a commitment $\text{com}(2\text{UH}_\beta(x))$ to the hash at any point x , so that the resulting commitment is compact in the sense that it depends only on the size of $2\text{UH}_\beta(x)$ and not on the size of x . Intuitively, the reason this works is that any collision in CRH_γ implies a collision in the underlying 2-universal hash 2UH_β , which leaks information about the hash key β (concretely, any fixed x, x' form a collision in a random hash function with small probability) thereby violating the hiding of the commitment.

At a high-level, we aim to mimic the above construction based on obfuscation. As a key for the collision-resistant hash we can obfuscate a program Π_β associated with a secret hash key β that given x outputs a commitment $\text{com}(2\text{UH}_\beta(x))$, where the commitment is derandomized using a PRF. The obfuscation $i\mathcal{O}(\Pi_\beta)$ can be thought of as the commitment to β , and evaluating this program at x , corresponds to homomorphic evaluation. Despite the clear intuition behind this construction, it is not clear how to prove its security based on IO. In fact, by [AS15], it cannot be proven based on a black-box reduction as long as plain statistically-binding commitments are used, as these can be constructed from OWPs in a fully black-box manner.

We show, however, that relying on a relaxed notion of perfectly-hiding commitments, as well as subexponential hardness of IO and puncturable PRFs, the construction can be proven secure. The perfect hiding of the commitment is leveraged in a probabilistic IO argument [CLTV15] that involves a number of hybrids larger than the overall number of commitments. We then observe that these relaxed commitments follow from appropriate average-case hardness of SZK.¹²

Acknowledgements

We thank Gil Segev, Iftach Haitner and Mohammad Mahmoody for elaborately answering our questions regarding existing separation results in cryptography. We also thank the anonymous reviewers for their valuable comments.

¹² Similar SZK-hardness is known to imply statistically-hiding commitments against malicious receivers, but with a larger (constant) number of rounds [OV08].

Bibliography

- [ABW10] Benny Applebaum, Boaz Barak, and Avi Wigderson. Public-key cryptography from different assumptions. In *Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC 2010, Cambridge, Massachusetts, USA, 5-8 June 2010*, pages 171–180, 2010.
- [AGGM06] Adi Akavia, Oded Goldreich, Shafi Goldwasser, and Dana Moshkovitz. On basing one-way functions on np-hardness. In Kleinberg [Kle06], pages 701–710.
- [AH91] William Aiello and Johan Hastad. Statistical zero-knowledge languages can be recognized in two rounds. *Journal of Computer and System Sciences*, 42(3):327–345, 1991.
- [Ale03] Michael Alekhnovich. More on average case vs approximation complexity. In *44th Symposium on Foundations of Computer Science (FOCS 2003), 11-14 October 2003, Cambridge, MA, USA, Proceedings* [DBL03], pages 298–307.
- [AR04] Dorit Aharonov and Oded Regev. Lattice problems in NP cap comp. In *45th Symposium on Foundations of Computer Science (FOCS 2004), 17-19 October 2004, Rome, Italy, Proceedings*, pages 362–371. IEEE Computer Society, 2004.
- [AR16] Benny Applebaum and Pavel Raykov. From private simultaneous messages to zero-information arthur-merlin protocols and back. In Eyal Kushilevitz and Tal Malkin, editors, *Theory of Cryptography - 13th International Conference, TCC 2016-A, Tel Aviv, Israel, January 10-13, 2016, Proceedings, Part II*, volume 9563 of *Lecture Notes in Computer Science*, pages 65–82. Springer, 2016.
- [AS15] Gilad Asharov and Gil Segev. Limits on the power of indistinguishability obfuscation and functional encryption. In *Symposium on the Foundations of Computer Science*, 2015.
- [AS16] Gilad Asharov and Gil Segev. On constructing one-way permutations from indistinguishability obfuscation. In *Theory of Cryptography*, pages 512–541. Springer, 2016.
- [Bar13] Boaz Barak. Structure vs. combinatorics in computational complexity. <http://windowsontheory.org/2013/10/07/structure-vs-combinatorics-in-computational-complexity/>, 2013.
- [BB15] Andrej Bogdanov and Christina Brzuska. On basing size-verifiable one-way functions on np-hardness. In Yevgeniy Dodis and Jesper Buus Nielsen, editors, *Theory of Cryptography - 12th Theory of Cryptography Conference, TCC 2015, Warsaw, Poland, March 23-25, 2015, Proceedings, Part I*, volume 9014 of *Lecture Notes in Computer Science*, pages 1–6. Springer, 2015.
- [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*, pages 296–315, 2013.
- [BGI⁺01] Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil P. Vadhan, and Ke Yang. On the (im)possibility of obfuscating programs. In Joe Kilian, editor, *Advances in Cryptology - CRYPTO 2001, 21st Annual International Cryptology Conference, Santa Barbara, California, USA, August 19-23, 2001, Proceedings*, volume 2139 of *Lecture Notes in Computer Science*, pages 1–18. Springer, 2001.
- [BGL⁺15] Nir Bitansky, Sanjam Garg, Huijia Lin, Rafael Pass, and Sidharth Telang. Succinct randomized encodings and their applications. In *Symposium on Theory of Computing, STOC 2015*, 2015.
- [BHZ87] Ravi B Boppana, Johan Hastad, and Stathis Zachos. Does comp have short interactive proofs? *Information Processing Letters*, 25(2):127–132, 1987.
- [BI87] Manuel Blum and Russell Impagliazzo. Generic oracles and oracle classes. In *Proceedings of the 28th Annual Symposium on Foundations of Computer Science, SFCS '87*, pages 118–126, Washington, DC, USA, 1987. IEEE Computer Society.
- [BKS⁺11] Zvika Brakerski, Jonathan Katz, Gil Segev, and Arkady Yerukhovich. Limits on the power of zero-knowledge proofs in cryptographic constructions. In Ishai [Ish11], pages 559–578.
- [BL13] Andrej Bogdanov and Chin Ho Lee. Limits of provable security for homomorphic encryption. In Ran Canetti and Juan A. Garay, editors, *Advances in Cryptology - CRYPTO 2013 - 33rd Annual Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2013. Proceedings, Part I*, volume 8042 of *Lecture Notes in Computer Science*, pages 111–128. Springer, 2013.
- [BM09] Boaz Barak and Mohammad Mahmoody-Ghidary. Merkle Puzzles Are Optimal - An $O(n^2)$ -Query Attack on Any Key Exchange from a Random Oracle. In Shai Halevi, editor, *Advances in Cryptology - CRYPTO 2009, 29th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 16-20, 2009. Proceedings*, volume 5677 of *Lecture Notes in Computer Science*, pages 374–390. Springer, 2009.
- [BP15] Nir Bitansky and Omer Paneth. Zaps and non-interactive witness indistinguishability from indistinguishability obfuscation. In Yevgeniy Dodis and Jesper Buus Nielsen, editors, *Theory of Cryptography - 12th Theory of Cryptography Conference, TCC 2015, Warsaw, Poland, March 23-25, 2015, Proceedings, Part II*, volume 9015 of *Lecture Notes in Computer Science*, pages 401–427. Springer, 2015.
- [BPR15] Nir Bitansky, Omer Paneth, and Alon Rosen. On the cryptographic hardness of finding a nash equilibrium. In Venkatesan Guruswami, editor, *IEEE 56th Annual Symposium on Foundations of Computer*

- Science, FOCS 2015, Berkeley, CA, USA, 17-20 October, 2015*, pages 1480–1498. IEEE Computer Society, 2015.
- [BPW16] Nir Bitansky, Omer Paneth, and Daniel Wichs. Perfect structure on the edge of chaos - trapdoor permutations from indistinguishability obfuscation. In *Theory of Cryptography - 13th International Conference, TCC 2016-A, Tel Aviv, Israel, January 10-13, 2016, Proceedings, Part I*, pages 474–502, 2016.
- [Bra79] Gilles Brassard. Relativized cryptography. In *20th Annual Symposium on Foundations of Computer Science, San Juan, Puerto Rico, 29-31 October 1979*, pages 383–391. IEEE Computer Society, 1979.
- [BT03] Andrej Bogdanov and Luca Trevisan. On worst-case to average-case reductions for NP problems. In *44th Symposium on Foundations of Computer Science (FOCS 2003), 11-14 October 2003, Cambridge, MA, USA, Proceedings [DBL03]*, pages 308–317.
- [BV11] Zvika Brakerski and Vinod Vaikuntanathan. Efficient fully homomorphic encryption from (standard) LWE. In Rafail Ostrovsky, editor, *FOCS*, pages 97–106. IEEE, 2011. Invited to SIAM Journal on Computing.
- [CDT09] Xi Chen, Xiaotie Deng, and Shang-Hua Teng. Settling the complexity of computing two-player nash equilibria. *J. ACM*, 56(3), 2009.
- [CHJV15] Ran Canetti, Justin Holmgren, Abhishek Jain, and Vinod Vaikuntanathan. Succinct garbling and indistinguishability obfuscation for RAM programs. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC 2015, Portland, OR, USA, June 14-17, 2015*, pages 429–437, 2015.
- [CLTV15] Ran Canetti, Huijia Lin, Stefano Tessaro, and Vinod Vaikuntanathan. Obfuscation of probabilistic circuits and applications. In *TCC*, 2015.
- [Cra12] Ronald Cramer, editor. *Theory of Cryptography - 9th Theory of Cryptography Conference, TCC 2012, Taormina, Sicily, Italy, March 19-21, 2012. Proceedings*, volume 7194 of *Lecture Notes in Computer Science*. Springer, 2012.
- [DBL00] *41st Annual Symposium on Foundations of Computer Science, FOCS 2000, 12-14 November 2000, Redondo Beach, California, USA*. IEEE Computer Society, 2000.
- [DBL03] *44th Symposium on Foundations of Computer Science (FOCS 2003), 11-14 October 2003, Cambridge, MA, USA, Proceedings*. IEEE Computer Society, 2003.
- [DGP06] Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. The complexity of computing a nash equilibrium. In Kleinberg [Kle06], pages 71–78.
- [DH76] Whitfield Diffie and Martin E. Hellman. New directions in cryptography. *IEEE Transactions on Information Theory*, 22(6):644–654, 1976.

- [DHT12] Yevgeniy Dodis, Iftach Haitner, and Aris Tentis. On the instantiability of hash-and-sign RSA signatures. In Cramer [Cra12], pages 112–132.
- [DLMM11] Dana Dachman-Soled, Yehuda Lindell, Mohammad Mahmoody, and Tal Malkin. On the black-box complexity of optimally-fair coin tossing. In Ishai [Ish11], pages 450–467.
- [ESY84] Shimon Even, Alan L. Selman, and Yacov Yacobi. The complexity of promise problems with applications to public-key cryptography. *Information and Control*, 61(2):159–173, 1984.
- [Fis12] Marc Fischlin. Black-box reductions and separations in cryptography. In Aikaterini Mitrokotsa and Serge Vaudenay, editors, *Progress in Cryptology - AFRICACRYPT 2012 - 5th International Conference on Cryptology in Africa, Ifrance, Morocco, July 10-12, 2012. Proceedings*, volume 7374 of *Lecture Notes in Computer Science*, pages 413–422. Springer, 2012.
- [For89] Lance Jeremy Fortnow. *Complexity-theoretic aspects of interactive proof systems*. PhD thesis, Massachusetts Institute of Technology, 1989.
- [Gen09] Craig Gentry. Fully homomorphic encryption using ideal lattices. In *STOC*, pages 169–178, 2009.
- [GG98] Oded Goldreich and Shafi Goldwasser. On the possibility of basing cryptography on the assumption that $\$p \not\equiv \text{np}\$$. *IACR Cryptology ePrint Archive*, 1998:5, 1998.
- [GGH⁺13a] Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *54th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2013, 26-29 October, 2013, Berkeley, CA, USA*, pages 40–49. IEEE Computer Society, 2013.
- [GGH⁺13b] Sanjam Garg, Craig Gentry, Shai Halevi, Amit Sahai, Mariana Raykova, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *FOCS*, 2013.
- [GGKT05] Rosario Gennaro, Yael Gertner, Jonathan Katz, and Luca Trevisan. Bounds on the efficiency of generic cryptographic constructions. *SIAM J. Comput.*, 35(1):217–246, 2005.
- [GK93] Oded Goldreich and Eyal Kushilevitz. A perfect zero-knowledge proof system for a problem equivalent to the discrete logarithm. *Journal of Cryptology*, 6(2):97–116, 1993.
- [GKLM12] Vipul Goyal, Virendra Kumar, Satyanarayana V. Lokam, and Mohammad Mahmoody. On black-box reductions between predicate encryption schemes. In Cramer [Cra12], pages 440–457.
- [GKM⁺00] Yael Gertner, Sampath Kannan, Tal Malkin, Omer Reingold, and Mahesh Viswanathan. The relationship between public key encryption and oblivious transfer. In *41st Annual Symposium on Foundations of Computer Science, FOCS 2000, 12-14 November 2000, Redondo Beach, California, USA* [DBL00], pages 325–335.

- [GM82] Shafi Goldwasser and Silvio Micali. Probabilistic encryption and how to play mental poker keeping secret all partial information. In Harry R. Lewis, Barbara B. Simons, Walter A. Burkhard, and Lawrence H. Landweber, editors, *Proceedings of the 14th Annual ACM Symposium on Theory of Computing, May 5-7, 1982, San Francisco, California, USA*, pages 365–377. ACM, 1982.
- [GMM07] Yael Gertner, Tal Malkin, and Steven Myers. Towards a separation of semantic and CCA security for public key encryption. In Salil P. Vadhan, editor, *Theory of Cryptography, 4th Theory of Cryptography Conference, TCC 2007, Amsterdam, The Netherlands, February 21-24, 2007, Proceedings*, volume 4392 of *Lecture Notes in Computer Science*, pages 434–455. Springer, 2007.
- [GMR85] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof-systems (extended abstract). In Robert Sedgewick, editor, *Proceedings of the 17th Annual ACM Symposium on Theory of Computing, May 6-8, 1985, Providence, Rhode Island, USA*, pages 291–304. ACM, 1985.
- [GMR01] Yael Gertner, Tal Malkin, and Omer Reingold. On the impossibility of basing trapdoor functions on trapdoor predicates. In *42nd Annual Symposium on Foundations of Computer Science, FOCS 2001, 14-17 October 2001, Las Vegas, Nevada, USA*, pages 126–135. IEEE Computer Society, 2001.
- [GMW91] Oded Goldreich, Silvio Micali, and Avi Wigderson. Proofs that yield nothing but their validity for all languages in NP have zero-knowledge proof systems. *J. ACM*, 38(3):691–729, 1991.
- [Gol06] Oded Goldreich. On promise problems: A survey. In *Theoretical Computer Science, Essays in Memory of Shimon Even*, pages 254–290, 2006.
- [GT00] Rosario Gennaro and Luca Trevisan. Lower bounds on the efficiency of generic cryptographic constructions. In *41st Annual Symposium on Foundations of Computer Science, FOCS 2000, 12-14 November 2000, Redondo Beach, California, USA [DBL00]*, pages 305–313.
- [GV99] Oded Goldreich and Salil P. Vadhan. Comparing entropies in statistical zero knowledge with applications to the structure of SZK. In *Proceedings of the 14th Annual IEEE Conference on Computational Complexity, Atlanta, Georgia, USA, May 4-6, 1999*, page 54, 1999.
- [Has88] Johan Hastad. Dual vectors and lower bounds for the nearest lattice point problem. *Combinatorica*, 8(1):75–81, 1988.
- [HH09] Iftach Haitner and Thomas Holenstein. On the (im)possibility of key dependent encryption. In Omer Reingold, editor, *Theory of Cryptography, 6th Theory of Cryptography Conference, TCC 2009, San Francisco, CA, USA, March 15-17, 2009. Proceedings*, volume 5444 of *Lecture Notes in Computer Science*, pages 202–219. Springer, 2009.

- [HHR15] Iftach Haitner, Jonathan J Hoch, Omer Reingold, and Gil Segev. Finding collisions in interactive protocols—tight lower bounds on the round and communication complexities of statistically hiding commitments. *SIAM Journal on Computing*, 44(1):193–242, 2015.
- [HMX10] Iftach Haitner, Mohammad Mahmoody, and David Xiao. A new sampling protocol and applications to basing cryptographic primitives on the hardness of np. In *Computational Complexity (CCC), 2010 IEEE 25th Annual Conference on*, pages 76–87. IEEE, 2010.
- [HR04] Chun-Yuan Hsiao and Leonid Reyzin. Finding collisions on a public road, or do secure hash functions need secret coins? In *Advances in Cryptology - CRYPTO 2004, 24th Annual International Cryptology Conference, Santa Barbara, California, USA, August 15-19, 2004, Proceedings*, pages 92–105, 2004.
- [IKO05] Yuval Ishai, Eyal Kushilevitz, and Rafail Ostrovsky. Sufficient conditions for collision-resistant hashing. In *Theory of Cryptography, Second Theory of Cryptography Conference, TCC 2005, Cambridge, MA, USA, February 10-12, 2005, Proceedings*, pages 445–456, 2005.
- [IR89] Russell Impagliazzo and Steven Rudich. Limits on the provable consequences of one-way permutations. In *Proceedings of the twenty-first annual ACM symposium on Theory of computing*, pages 44–61. ACM, 1989.
- [Ish11] Yuval Ishai, editor. *Theory of Cryptography - 8th Theory of Cryptography Conference, TCC 2011, Providence, RI, USA, March 28-30, 2011. Proceedings*, volume 6597 of *Lecture Notes in Computer Science*. Springer, 2011.
- [Kle06] Jon M. Kleinberg, editor. *Proceedings of the 38th Annual ACM Symposium on Theory of Computing, Seattle, WA, USA, May 21-23, 2006*. ACM, 2006.
- [KLW15] Venkata Koppula, Allison Bishop Lewko, and Brent Waters. Indistinguishability obfuscation for turing machines with unbounded memory. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC 2015, Portland, OR, USA, June 14-17, 2015*, pages 419–428, 2015.
- [KMN⁺14] Ilan Komargodski, Tal Moran, Moni Naor, Rafael Pass, Alon Rosen, and Eylon Yogev. One-way functions and (im)perfect obfuscation. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*, pages 374–383. IEEE Computer Society, 2014.
- [KSS11] Jeff Kahn, Michael E. Saks, and Clifford D. Smyth. The dual BKR inequality and rudich’s conjecture. *Combinatorics, Probability & Computing*, 20(2):257–266, 2011.
- [KST99] Jeong Han Kim, Daniel R. Simon, and Prasad Tetali. Limits on the efficiency of one-way permutation-based hash functions. In *40th Annual Symposium on Foundations of Computer Science, FOCS '99, 17-18 October, 1999, New York, NY, USA*, pages 535–542. IEEE Computer Society, 1999.

- [LLJS90] Jeffrey C Lagarias, Hendrik W Lenstra Jr, and Claus-Peter Schnorr. Korkin-zolotarev bases and successive minima of a lattice and its reciprocal lattice. *Combinatorica*, 10(4):333–348, 1990.
- [LV16] Tianren Liu and Vinod Vaikuntanathan. On basing private information retrieval on np-hardness. In Eyal Kushilevitz and Tal Malkin, editors, *Theory of Cryptography - 13th International Conference, TCC 2016-A, Tel Aviv, Israel, January 10-13, 2016, Proceedings, Part I*, volume 9562 of *Lecture Notes in Computer Science*, pages 372–386. Springer, 2016.
- [MM11] Takahiro Matsuda and Kanta Matsuura. On black-box separations among injective one-way functions. In *Theory of Cryptography*, pages 597–614. Springer, 2011.
- [MP91] Nimrod Megiddo and Christos H. Papadimitriou. On total functions, existence theorems and computational complexity. *Theor. Comput. Sci.*, 81(2):317–324, 1991.
- [MV03] Daniele Micciancio and Salil P. Vadhan. Statistical zero-knowledge proofs with efficient provers: Lattice problems and more. In Dan Boneh, editor, *Advances in Cryptology - CRYPTO 2003, 23rd Annual International Cryptology Conference, Santa Barbara, California, USA, August 17-21, 2003, Proceedings*, volume 2729 of *Lecture Notes in Computer Science*, pages 282–298. Springer, 2003.
- [MX10] Mohammad Mahmoody and David Xiao. On the power of randomized reductions and the checkability of sat. In *Computational Complexity (CCC), 2010 IEEE 25th Annual Conference on*, pages 64–75. IEEE, 2010.
- [Ost91] Rafail Ostrovsky. One-way functions, hard on average problems, and statistical zero-knowledge proofs. In *Structure in Complexity Theory Conference, 1991., Proceedings of the Sixth Annual*, pages 133–138. IEEE, 1991.
- [OV08] Shien Jin Ong and Salil P. Vadhan. An equivalence between zero knowledge and commitments. In *Theory of Cryptography, Fifth Theory of Cryptography Conference, TCC 2008, New York, USA, March 19-21, 2008.*, pages 482–500, 2008.
- [Pap94] Christos H. Papadimitriou. On the complexity of the parity argument and other inefficient proofs of existence. *J. Comput. Syst. Sci.*, 48(3):498–532, 1994.
- [Pas06] Rafael Pass. Parallel repetition of zero-knowledge proofs and the possibility of basing cryptography on np-hardness. In *21st Annual IEEE Conference on Computational Complexity (CCC 2006), 16-20 July 2006, Prague, Czech Republic*, pages 96–110. IEEE Computer Society, 2006.
- [Pas13] Rafael Pass. Unprovable security of perfect NIZK and non-interactive non-malleable commitments. In *TCC*, pages 334–354, 2013.

- [RAD78] R. Rivest, L. Adleman, and M. Dertouzos. On data banks and privacy homomorphisms. In *Foundations of Secure Computation*, pages 169–177. Academic Press, 1978.
- [RSA78] Ronald L. Rivest, Adi Shamir, and Leonard M. Adleman. A method for obtaining digital signatures and public-key cryptosystems. *Commun. ACM*, 21(2):120–126, 1978.
- [RSS16] Alon Rosen, Gil Segev, and Ido Shahaf. Can PPAD hardness be based on standard cryptographic assumptions? *Electronic Colloquium on Computational Complexity (ECCC)*, 23:59, 2016.
- [RTV04] Omer Reingold, Luca Trevisan, and Salil P. Vadhan. Notions of reducibility between cryptographic primitives. In *Theory of Cryptography, First Theory of Cryptography Conference, TCC 2004, Cambridge, MA, USA, February 19-21, 2004, Proceedings*, pages 1–20, 2004.
- [Rud88] Steven Rudich. *Limits on the Provable Consequences of One-Way Functions*. PhD thesis, University of California, Berkeley, 1988.
- [Rud91] Steven Rudich. The use of interaction in public cryptosystems (extended abstract). In Joan Feigenbaum, editor, *Advances in Cryptology - CRYPTO '91, 11th Annual International Cryptology Conference, Santa Barbara, California, USA, August 11-15, 1991, Proceedings*, volume 576 of *Lecture Notes in Computer Science*, pages 242–251. Springer, 1991.
- [Sim98] Daniel R Simon. Finding collisions on a one-way street: Can secure hash functions be based on general assumptions? In *Advances in Cryptology EUROCRYPT'98*, pages 334–345. Springer, 1998.
- [SV03] Amit Sahai and Salil Vadhan. A complete problem for statistical zero knowledge. *Journal of the ACM (JACM)*, 50(2):196–249, 2003.
- [SW14] Amit Sahai and Brent Waters. How to use indistinguishability obfuscation: deniable encryption, and more. In David B. Shmoys, editor, *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014*, pages 475–484. ACM, 2014.
- [Vad99] Salil Pravin Vadhan. *A study of statistical zero-knowledge proofs*. PhD thesis, Massachusetts Institute of Technology, 1999.
- [Wat15] Brent Waters. A punctured programming approach to adaptively secure functional encryption. In *Advances in Cryptology - CRYPTO 2015 - 35th Annual Cryptology Conference, Santa Barbara, CA, USA, August 16-20, 2015, Proceedings, Part II*, pages 678–697, 2015.