

Correcting Subverted Random Oracles

Alexander Russell* Qiang Tang† Moti Yung‡ Hong-Sheng Zhou§ Jiadong Zhu¶

October 14, 2020

Abstract

The random oracle methodology has proven to be a powerful tool for designing and reasoning about cryptographic schemes. In this paper, we focus on the basic problem of correcting faulty—or adversarially corrupted—random oracles, so that they can be confidently applied for such cryptographic purposes.

We prove that a simple construction can transform a “subverted” random oracle—which disagrees with the original one at a small fraction of inputs—into an object that is *indifferentiable* from a random function, even if the adversary is made aware of all randomness used in the transformation. Our results permit future designers of cryptographic primitives in typical kleptographic settings (i.e., those permitting adversaries that subvert or replace basic cryptographic algorithms) to use random oracles as a trusted black box.

1 Introduction

The random oracle methodology of Bellare and Rogaway [BR93] has proven to be a powerful tool for designing and reasoning about cryptographic schemes. It consists of the following two steps: (i) design a scheme Π in which all parties (including the adversary) have oracle access to a common truly random function, and establish the security of Π in this favorable setting; (ii) instantiate the random oracle in Π with a suitable cryptographic hash function (such as SHA256) to obtain an instantiated scheme Π' . The random oracle heuristic states that if the original scheme Π is secure, then the instantiated scheme Π' is also secure. While this heuristic can fail in various settings [CGH98] the basic framework remains a fundamental design and analytic tool. In this work we focus on the problem of correcting faulty—or adversarially corrupted—random oracles so that they can be confidently applied for such cryptographic purposes.¹

Specifically, given a function \tilde{h} drawn from a distribution which agrees in most places with a uniform function, we would like to produce a corrected version which appears uniform to adversaries with a polynomially bounded number of queries. This model is partially motivated by the traditional study of “program checking and self-correcting” [Blu88, BK89, BLR90]: the goal in this theory is to transform a program that is faulty at a small fraction of inputs (modeling an evasive adversary) to a program that is correct at all points with overwhelming probability. Our setting intuitively adapts this classical theory of self-correction to the study of “self-correcting a probability distribution.” Notably, the functions to be corrected are structureless, instead of heavily structured. Despite that, the basic procedure for correction and portions of the technical development are analogous.

One particular motivation for correcting random oracles in a cryptographic context arises from recent work studying design and security in the subversion (i.e., *kleptographic*) setting. In this setting, the various components of a

*University of Connecticut, acr@cse.uconn.edu

†New Jersey Institute of Technology, qiang@njit.edu

‡Google Inc and Columbia University, moti@cs.columbia.edu

§Virginia Commonwealth University, hszhou@vcu.edu

¶University of Connecticut, jiadong.zhu@uconn.edu

¹This work expands on the preliminary conference version of this article which appeared in CRYPTO’18. Several improvements are made to the conference version in this paper. First, this paper fills some gaps in the security proof of the conference paper. Second, this paper applies more powerful analysis tools and establishes significantly tighter results. Third, the security proof in the conference version indicated how the major technical hurdles can be overcome without developing the full details. This article fills those details by game transitions from the construction to the perfect random function.

cryptographic scheme may be subverted by an adversary, so long as the tampering cannot be detected via blackbox testing. This is a challenging framework (as highlighted by Bellare, Paterson, and Rogaway [BPR14]), because many basic cryptographic techniques are not directly available: in particular, the random oracle paradigm is directly undermined. In terms of the discussion above, the random oracle—which is eventually to be replaced with a concrete function—is subject to adversarial subversion which complicates even the first step of the random oracle methodology above. Our goal is to provide a generic approach that can rigorously “protect” the usage of random oracles from subversion and, essentially, establish a “crooked” random oracle methodology.

1.1 Our contributions

We first describe two concrete scenarios where hash functions are subverted in the kleptographic setting. We then express the security properties by adapting the successful framework of *indifferentiability* [MRH04, CDMP05] to our setting with adversarial subversion. This framework provides a satisfactory guarantee of modularity—that is, that the resulting object can be directly employed by other constructions demanding a random oracle. We call this new notion “crooked” indifferentiability to reflect the role of adversary in the modeling; see below. (A formal definition appears in Section 2.)

We prove that a simple construction involving only *public* randomness can boost a “subverted” random oracle into a construction that is indifferentiable from a random function. (Section 3, 4). We expand on these contributions below.

Consequences of kleptographic hash subversion. We first illustrate the security failures that can arise from use of hash functions that are subverted at only a negligible fraction of inputs with two concrete examples:

(1) *Chain take-over attack on blockchain.* For simplicity, consider a proof-of-work blockchain setting where miners compete to find a solution s to the “puzzle” $h(\text{pre}||\text{transactions}||s) \leq d$, where pre denotes the hash of the previous block, transactions denotes the set of valid transactions in the current block, and d denotes the difficulty parameter. Here h is intended to be a strong hash function. Note that the mining machines use a program $\tilde{h}(\cdot)$ (or a dedicated hardware module) which could be designed by a clever adversary. Now if \tilde{h} has been subverted so that $\tilde{h}(*||z) = 0$ for a randomly chosen z —and $\tilde{h}(x) = h(x)$ in all other cases—this will be difficult to detect by prior black-box testing; on the other hand, the adversary who created \tilde{h} has the luxury of solving the proof of work without any effort for any challenge, and thus can completely control the blockchain. (A fancier subversion can tune the “backdoor” z to other parts of the input so that it cannot be reused by other parties; e.g., $\tilde{h}(w||z) = 0$ if $z = f(w)$ for a secret pseudorandom function known to the adversary.)

(2) *System sneak-in attack on password authentication.* In Unix-style system, during system initialization, the root user chooses a master password α and the system stores the digest $\rho = h(\alpha)$, where h is a given hash function normally modeled as a random oracle. During login, the operating system receives input x and accepts this password if $h(x) = \rho$. An attractive feature of this practice is that it is still secure if ρ is accidentally leaked. In the presence of kleptographic attacks, however, the module that implements the hash function h may be strategically subverted, yielding a new function \tilde{h} which destroys the security of the scheme above: for example, the adversary may choose a relatively short random string z and define $\tilde{h}(y) = h(y)$ unless y begins with z , in which case $\tilde{h}(zx) = x$. As above, h and \tilde{h} are indistinguishable by black-box testing; on the other hand, the adversary can login as the system administrator using ρ and its knowledge of the backdoor z (without knowing the actual password α , presenting $z\rho$ instead).

The model of “crooked” indifferentiability. The problem of cleaning defective randomness has a long history in computer science. Our setting requires that the transformation must be carried out by a local rule and involve an exponentially small amount of public randomness (in the sense that we wish to clean a defective random function $h : \{0, 1\}^n \rightarrow \{0, 1\}^n$ with only a polynomial length random string). The basic framework of correcting a subverted random oracle is the following:

First, a function $h : \{0, 1\}^n \rightarrow \{0, 1\}^n$ is drawn uniformly at random. Then, an adversary may *subvert* the function h , yielding a new function \tilde{h} . The subverted function $\tilde{h}(x)$ is described by an adversarially-chosen (polynomial-time) algorithm $\tilde{H}^h(x)$, with oracle access to h . We insist that $\tilde{h}(x) \neq h(x)$ only at a negligible fraction of inputs.² Next, the

²We remark that tampering with even a negligible fraction of inputs can have devastating consequences in many settings of interest: e.g., the blockchain and password examples above. Additionally, the setting of negligible subversion is precisely the desired parameter range for existing

function \tilde{h} is “publicly corrected” to a function \tilde{h}_R (defined below) that involves some public randomness R selected after \tilde{h} is supplied.³

We wish to show that the resulting function (construction) is “as good as” a random oracle, in the sense of indifferentiability. The successful framework of indifferentiability asserts that a construction C^H (having oracle access to an ideal primitive H) is indifferentiable from another ideal primitive \mathcal{F} if there exists a simulator \mathcal{S} so that (C^H, H) and $(\mathcal{F}, \mathcal{S})$ are indistinguishable to any distinguisher \mathcal{D} .

To reflect our setting, an *H-crooked-distinguisher* $\widehat{\mathcal{D}}$ is introduced; the *H-crooked-distinguisher* $\widehat{\mathcal{D}}$ first prepares the subverted implementation \tilde{H} (after querying H first); then a fixed amount of (public) randomness R is drawn and published; the construction C uses only subverted implementation \tilde{H} and R . Now following the indifferentiability framework, we will ask for a simulator \mathcal{S} , such that $(C^{\tilde{H}^H}(\cdot, R), H)$ and $(\mathcal{F}, \mathcal{S}^{\tilde{H}}(R))$ are indistinguishable to any *H-crooked-distinguisher* $\widehat{\mathcal{D}}$ (even one who knows R). A similar security preserving theorem [MRH04, CDMP05] also holds in our model. See Section 2 for details.

The construction. The construction depends on a parameter $\ell = \text{poly}(n)$ and public randomness $R = (r_1, \dots, r_\ell)$, where each r_i is an independent and uniform element of $\{0, 1\}^n$. For simplicity, the construction relies on a family of independent random oracles $h_i(x)$, for $i \in \{0, \dots, \ell\}$. (Of course, these can all be extracted from a single random oracle with slightly longer inputs by defining $\tilde{h}_i(x) = \tilde{h}(i, x)$ and treating the output of $h_i(x)$ as n bits long.) Then we define

$$\tilde{h}_R(x) = \tilde{h}_0 \left(\bigoplus_{i=1}^{\ell} \tilde{h}_i(x \oplus r_i) \right) = \tilde{h}_0(\tilde{g}_R(x)).$$

Note that the adversary is permitted to subvert the function(s) h_i by choosing an algorithm $H^{h_i}(x)$ so that $\tilde{h}_i(x) = H^{h_i}(i, x)$. Before diving into the analysis, let us first quick demonstrate how some simpler constructions fail.

Simple constructions and their shortcomings. During the stage when the adversary manufactures the hash functions $\tilde{h}_* = \{\tilde{h}_i\}_{i=0}^{\ell}$, the randomness $R := r_1, \dots, r_\ell$ are not known to the adversary; they become public in the second query phase. If the “mixing” operation is not carefully designed, the adversary could choose inputs accordingly, trying to “peel off” R . We discuss a few examples:

1. $\tilde{h}_R(x)$ is simply defined as $\tilde{h}_1(x \oplus r_1)$. A straightforward attack is as follows: the adversary can subvert h_1 in a way that $\tilde{h}_1(m) = 0$ for a random input m ; the adversary then queries $m \oplus r_1$ on $\tilde{h}_R(\cdot)$ and can trivially distinguish \tilde{h} from a random function.
2. $\tilde{h}_R(x)$ is defined as $\tilde{h}_1(x \oplus r_1) \oplus \tilde{h}_2(x \oplus r_2)$. Now a slightly more complex attack can still succeed: the adversary subverts h_1 so that $\tilde{h}_1(x) = 0$ if $x = m||*$, that is, when the first half of x equals to a randomly selected string m with length $n/2$; likewise, h_2 is subverted so that $\tilde{h}_2(x) = 0$ if $x = *||m$, that is, the second half of x equals m . Then, the adversary queries $m_1||m_2$ on $\tilde{h}_R(\cdot)$, where $m_1 = m \oplus r_{1,0}$, and $m_2 = m \oplus r_{2,1}$, and $r_{1,0}$ is the first half of r_1 , and $r_{2,1}$ is the second half of r_2 . Again, trivially, it can be distinguished from a random function.

This attack can be generalized in a straightforward fashion to any $\ell \leq n/\lambda$: the input can be divided in into consecutive substrings each with length λ , and the “trigger” substrings can be planted in each chunk.

Challenges in the analysis. To analyze security in the “crooked” indifferentiability framework, our simulator needs to ensure consistency between two ways of generating output values: one is directly from the construction $C^{\tilde{H}^H}(x, R)$; the other calls for an “explanation” of F —a truly random function—via reconstruction from related queries to H (in a way consistent with the subverted implementation \tilde{H}). To ensure a correct simulation, the simulator must suitably answer related queries (defining one value of $C^{\tilde{H}^H}(x, R)$). Essentially, the proof relies on an unpredictability property of the internal function $\tilde{g}_R(x)$ to guarantee the success of simulation. In particular, for any input x (if not yet “fully decided”

models of kleptographic subversion and security. In these models, when an oracle is non-negligibly defective, this can be easily detected by a watchdog using a simple sampling and testing regimen, see e.g., [RTYZ16].

³We remark that in many settings, e.g., the model of classical self-correcting programs, we are permitted to sample fresh and “private” randomness for each query; in our case, we may only use a single polynomial-length random string for all points. Once R is generated, it is made public and fixed, which implicitly defines our corrected function $\tilde{h}_R(\cdot)$. This latter requirement is necessary in our setting as random oracles are typically used as a public object—in particular, our attacker must have full knowledge of R .

by previous queries), the output of $\tilde{g}_R(x)$ is unpredictable to the distinguisher even if she knows the public randomness R (even conditioned on adaptive queries generated by $\tilde{\mathcal{D}}$).

Section 4 develops the detailed security analysis. The proof of correctness for this construction is complicated by the fact that the “defining” algorithm \tilde{H} is permitted to make adaptive queries to h during the definition of \tilde{h} ; in particular, this means that even when a particular “constellation” of points (the $h_i(x \oplus r_i)$) contains a point that is left alone by \tilde{H} —which is to say that it agrees with $h_i(\cdot)$ —there is no guarantee that $\bigoplus_i h_i(x \oplus r_i)$ is uniformly random. This suggests focusing the analysis on demonstrating that the constellation associated with every $x \in \{0, 1\}^n$ will have at least one “good” component, which is (i.) not queried by $\tilde{H}^h(\cdot)$ when evaluated on the other terms, and (ii.) answered honestly. Unfortunately, actually identifying such a good point with certainty appears to require that we examine *all* of the points in the constellation for x , and this interferes with the standard “exposure martingale” proof that is so powerful in the random oracle setting (which capitalizes on the fact that “unexamined” values of h can be treated as independent and uniform values). Thus, even such elementary aspects of the internal function require some care and must be handled with appropriate union bounds. This points to a second difficulty, which is that structural properties must generally be established with exponentially small (that is 2^{-Cn} for $C > 1$) error probability, to permit a union bound over all inputs, or must be shown to fail with negligible probability conditioned on a typical transcript of the distinguisher; thus one must either establish very small error or carry significant conditioning through the proof. In general, we aim for the first of these options in order to simplify the presentation. For example, part of analysis is supported with a Fourier-analytic argument that establishes the near uniformity of the random variable $x_1 \oplus \dots \oplus x_\ell$ where each $x_i \in \mathbb{Z}_2^n$ is drawn from a sufficiently rich multiset. This provides a simple criterion for the initial random function h from which follows a number of strong properties that can conveniently avoid the bookkeeping associated with transcript conditioning.

Applications: Our correction function can be easily applied to save the faulty hash implementation in several important application scenarios.

- Immediate applications, as explained in the motivational examples, one may apply our technique directly in the following scenarios to defend against backdoors in hash implementation.

(1) *Salvaging PoW blockchain against hash subversion.* For proof-of-work based blockchains, as discussed above, miners may rely on a common library \tilde{h} for the hash evaluation, perhaps cleverly implemented by an adversary. Here \tilde{h} is determined before the chain has been deployed. We can then prevent the adversary from capitalizing on this subversion by applying our correction function. In particular, the public randomness R can be embedded in the genesis block; the function $\tilde{h}_R(\cdot)$ is then used for mining (and verification) rather than \tilde{h} .

(2) *Preventing system sneak-in.* The system sneak-in can also be resolved immediately by applying our correcting random oracle. During system initialization (or even when the operating system is released), the system administrator generates some randomness R and wraps the hash module \tilde{h} (potentially subverted) to define $\tilde{h}_R(\cdot)$. The password α then gives rise to the digest $\rho = \tilde{h}_R(\alpha)$ together with the randomness R . Upon receiving input x , the system first “recovers” $\tilde{h}_R(\cdot)$ based on the previously stored R , and then tests if $\rho = \tilde{h}_R(x)$. The access will be enabled if the test is valid. As the corrected random oracle ensures the output to be uniform for every input point, this remains secure in the face of subversion.⁴

- Extended applications:

(3) *Application to cryptographically secure digital signatures.* In a follow-up work, [CRT⁺19], we constructed the first digital signature scheme that is secure against kleptographic attacks with only an offline watchdog via *corrected* full domain hash. There, a further interpretation of our result is used that the corrected random oracle is indistinguishable also to a *keyed* random oracle.

(4) *Relation to random oracle against pre-processing attacks.* It is easy to see that our model is strictly stronger than the pre-processing model considered in [DGK17, CDGS18]. There, the adversary made some random oracle queries first and compress the transcripts as an auxiliary input. While in our model, besides this auxiliary input (which can be considered as part of the backdoor), the adversary is allowed to further subvert the implementation of the hash! It follows that our construction can also be directly used to defend the pre-processing attack.

⁴Typical authentication of this form also uses password “salt,” but this doesn’t change the structure of the attack or the solution.

1.2 Related Work

Related work on indistinguishability. The notion of indistinguishability was proposed in the elegant work of Maurer et al. [MRH04]; this notably extends the classical concept of indistinguishability to circumstances where one or more of the relevant oracles are publicly available (such as a random oracle). It was later adapted by Coron et al. [CDMP05]; several other variants were proposed and studied in [DP06, DP07]. A line of notable work applied the framework to the ideal cipher problem: in particular the Feistel construction (with a small constant number of rounds) is indistinguishable from an ideal cipher, see [CHK⁺16, DKT16, DS16]. Our work adapts the indistinguishability framework to the setting where the construction uses only a subverted implementation and the construction aims to be indistinguishable from a clean random oracle.

Related work on self-correcting programs. The theory of program self-testing, and self-correcting, was pioneered by the work of Blum et al. [Blu88, BK89, BLR90]. This theory addresses the basic problem of program correctness by verifying relationships between the outputs of the program on related, but randomly selected, inputs; additionally, the theory studies transformations of faulty programs that are almost correct (faulty at negligible fraction of inputs) into ones that are correct at every point with an overwhelming probability. See Rubinfeld’s thesis [Rub91] for great reading. Our results can be seen as a distributional version of this theory: (i.) we “correct” independent distributions rather than structured functions, (ii) we insist on using only an exponentially small amount of public randomness.

Related work on random oracles. The random oracle methodology/heuristic was first explicitly introduced by Bellare and Rogaway [BR93]; this methodology can significantly simplify both cryptographic constructions and proofs, even though there exist schemes which are secure using random oracles, but cannot be instantiated in the standard model [CGH98]. Soon after, many separations have been shown for new classes of cryptographic tasks [Nie02, Pas03, GK03, CGH04, BBP04, DOP05, LN09, KP09, BSW11, GKMZ16]. In addition, efforts have been made to identify instantiable assumptions/models in which we may analyze interesting cryptographic tasks [Can97, CMR98, CD08, BCFW09, BF05, BF06, KOS10, BHK13]. Also, we note that research efforts have also been made to investigate *weakened* idealized models [NIT08, KNTX10, Lis07, KLT15]. Finally, there are several recent approaches that study random oracles in the auxiliary input model (or with pre-processing) [DGK17, CDGS18]. Our model is strictly stronger than the pre-processing model: besides pre-processing queries, for example, the adversary may embed some preprocessed information into the subverted implementation; furthermore, our subverted implementation can further misbehave in ways that cannot be captured by any single-shot polynomial-query adversary because the subversion at each point is determined by a local adaptive computation.

Related work on kleptographic security. Kleptographic attacks were originally introduced by Young and Yung [YY96, YY97]: In such attacks, the adversary provides subverted implementations of the cryptographic primitive, trying to learn secret without being detected. In recent years, several remarkable allegations of cryptographic tampering [PLS13, Men13], including detailed investigations [CNE⁺14, CCG⁺16], have produced a renewed interest in both kleptographic attacks and in techniques for preventing them [BBCL13, BPR14, DGG⁺15, MS15, DMS16, DFP15, AMV15, BJK15, SFKR15, BH15, AAB⁺15, DPSW16, RTYZ16, RTYZ17, CDL17]. None of those work considered how to actually correct a subverted random oracle.

Similar constructions in other context. Our construction follows a simple design approach, applying the hash function to the XOR of multiple hash values. The construction calls to mind many classical constructions for other purposes, e.g., hardness amplification such as the Yao XOR lemma, weak PRF [Mye01], and randomizers in the bounded storage model [DM04]. Our construction has to have an “extra” layer of hash application (h_0 in our parlance) to wrap the XOR of terms, and our analysis is of course very different from these classical results due to our starting point of a subverted implementation.

Relationship to the preliminary version. This paper expands on the preliminary conference version of this article which appeared in CRYPTO ’18 [RTYZ18]. The conference version provided a proof sketch which indicated how the major technical hurdles can be overcome without developing the full details. This article, while filling in those details, additionally establishes somewhat tighter and simpler results. Specifically, in many cases, we are able to improve the analysis to show that certain events of interest fail with exponential, rather than negligible, probability; this achieves tighter bounds and also simplifies the presentation.

Other follow-up work. In [BNR20] Bhattacharyya et al. pointed out a gap in the security sketch in our preliminary

version [RTYZ18] and provide an alternate (and independent) approach to prove that the construction of [RTYZ18] is secure. They also explore an alternate, sponge-based construction that reduces the number of external random bits to linear.

2 The Model: Crooked Indifferentiability

2.1 Preliminary: Indifferentiability

The notion of indifferentiability proposed in the elegant work of Maurer et al. [MRH04] has proven to be a powerful tool for studying the security of hash function and many other primitives. The notion extends the classical concept of indistinguishability to the setting where one or more oracles involved in the construction are publicly available. The indifferentiability framework of [MRH04] is built around random systems providing interfaces to other systems. Coron et al. [CDMP05] demonstrate a strengthened⁵ indifferentiability framework built around Interactive Turing Machines (as in [Can01]). Our presentation borrows heavily from [CDMP05]. In the next subsection, we will introduce our new notion, *crooked indifferentiability*.

Defining indifferentiability. An *ideal primitive* is an algorithmic entity which receives inputs from one of the parties and returns its output immediately to the querying party. We now proceed to the definition of indifferentiability [MRH04, CDMP05]:

Definition 1 (Indifferentiability [MRH04, CDMP05]). A Turing machine C with oracle access to an ideal primitive \mathcal{G} is said to be $(t_{\mathcal{D}}, t_{\mathcal{S}}, q, \epsilon)$ -indifferentiable from an ideal primitive \mathcal{F} , if there is a simulator \mathcal{S} , such that for any distinguisher \mathcal{D} , it holds that :

$$\left| \Pr[\mathcal{D}^{C, \mathcal{G}}(1^n) = 1] - \Pr[\mathcal{D}^{\mathcal{F}, \mathcal{S}}(1^n) = 1] \right| \leq \epsilon.$$

The simulator \mathcal{S} has oracle access to \mathcal{F} and runs in time at most $t_{\mathcal{S}}$. The distinguisher \mathcal{D} runs in time at most $t_{\mathcal{D}}$ and makes at most q queries. Similarly, $C^{\mathcal{G}}$ is said to be (computationally) indifferentiable from \mathcal{F} if ϵ is a negligible function of the security parameter n (for polynomially bounded $t_{\mathcal{D}}$ and $t_{\mathcal{S}}$). See Figure 1.




Figure 1: The indifferentiability notion: the distinguisher \mathcal{D} either interacts with algorithm C and ideal primitive \mathcal{G} , or with ideal primitive \mathcal{F} and simulator \mathcal{S} . Algorithm C has oracle access to \mathcal{G} , while simulator \mathcal{S} has oracle access to \mathcal{F} .

As illustrated in Figure 1, the role of the simulator is to simulate the ideal primitive \mathcal{G} so that no distinguisher can tell whether it is interacting with C and \mathcal{G} , or with \mathcal{F} and \mathcal{S} ; in other words, the output of \mathcal{S} should look “consistent” with what the distinguisher can obtain from \mathcal{F} . Note that the simulator does not observe the distinguisher’s queries to \mathcal{F} ; however, it can call \mathcal{F} directly when needed for the simulation. Note that, in some sense, the simulator must “reverse engineer” the construction C so that the simulated oracle appropriately induces \mathcal{F} and, of course, possesses the correct marginal distribution.

⁵Technically, the quantifiers in the security definitions in the original [MRH04] and in the followup [CDMP05] are different; in the former, a simulator needs to be constructed for each adversary, while in the latter a simulator needs to be constructed for *all* adversaries.




Figure 2: The environment \mathcal{E} interacts with cryptosystem \mathcal{P} and attacker \mathcal{A} . In the \mathcal{G} model (left), \mathcal{P} has oracle access to \mathcal{C} whereas \mathcal{A} has oracle access to \mathcal{G} . In the \mathcal{F} model, both \mathcal{P} and $\mathcal{S}_{\mathcal{A}}$ have oracle access to \mathcal{F} .

Replacement. It is shown in [MRH04] that if $\mathcal{C}^{\mathcal{G}}$ is indistinguishable from \mathcal{F} , then $\mathcal{C}^{\mathcal{G}}$ can replace \mathcal{F} in any cryptosystem, and the resulting cryptosystem is at least as secure in the \mathcal{G} model as in the \mathcal{F} model.

We use the definition of [MRH04] to specify what it means for a cryptosystem to be at least as secure in the \mathcal{G} model as in the \mathcal{F} model. A cryptosystem is modeled as an Interactive Turing Machine with an interface to an adversary \mathcal{A} and to a public oracle. The cryptosystem is run by an environment \mathcal{E} which provides a binary output and also runs the adversary. In the \mathcal{G} model, cryptosystem \mathcal{P} has oracle access to \mathcal{C} whereas attacker \mathcal{A} has oracle access to \mathcal{G} . In the \mathcal{F} model, both \mathcal{P} and \mathcal{A} have oracle access to \mathcal{F} . The definition is illustrated in Figure 2.

Definition 2. A cryptosystem is said to be at least as secure in the \mathcal{G} model with algorithm \mathcal{C} as in the \mathcal{F} model, if for any environment \mathcal{E} and any attacker \mathcal{A} in the \mathcal{G} model, there exists an attacker $\mathcal{S}_{\mathcal{A}}$ in the \mathcal{F} model, such that:

$$\Pr[\mathcal{E}(\mathcal{P}^{\mathcal{C}}, \mathcal{A}^{\mathcal{G}}) = 1] - \Pr[\mathcal{E}(\mathcal{P}^{\mathcal{F}}, \mathcal{S}_{\mathcal{A}}^{\mathcal{F}}) = 1] \leq \epsilon.$$

where ϵ is a negligible function of the security parameter n . Similarly, a cryptosystem is said to be computationally at least as secure, etc., if \mathcal{E} , \mathcal{A} and $\mathcal{S}_{\mathcal{A}}$ are polynomial-time in n .

We have the following security preserving (replacement) theorem, which says that when an ideal primitive is replaced by an indistinguishable one, the security of the “big” cryptosystem remains:

Theorem 1 ([MRH04, CDMP05]). Let \mathcal{P} be a cryptosystem with oracle access to an ideal primitive \mathcal{F} . Let \mathcal{C} be an algorithm such that $\mathcal{C}^{\mathcal{G}}$ is indistinguishable from \mathcal{F} . Then cryptosystem \mathcal{P} is at least as secure in the \mathcal{G} model with algorithm \mathcal{C} as in the \mathcal{F} model.

2.2 Crooked indistinguishability

The ideal primitives that we focus on in this paper are random oracles. A random oracle [BR93] is an ideal primitive which provides an independent random output for each new query. We next formalize a new notion called *crooked indistinguishability* to reflect the challenges in our setting with subversion; our formalization is for random oracles, but the formalization can be naturally extended to other ideal primitives.

Crooked indistinguishability for random oracles. As mentioned in the introduction, we consider the problem of “repairing” a subverted random oracle in such a way that the corrected construction can be used as a drop-in replacement for an unsubverted random oracle. This immediately suggests invoking and appropriately adapting the indistinguishability notion. Specifically, we need to adjust the notion to reflect subversion.

We model the act of *subversion of a (hash) function H* as creation of an “implementation” \tilde{H} of the new, subverted (hash) function; in practice, this would be the source code of the subverted version of the function H . In our setting,

however, where H is modeled as a random oracle, we define \tilde{H} as a polynomial-time algorithm with oracle access to H ; thus the subverted function is $x \mapsto \tilde{H}^H(x)$. We proceed to survey the main modifications between crooked indifferenciability and the original notion of indifferenciability.

1. The deterministic construction will have oracle access to the random oracle only via the subverted implementation \tilde{H} but not via the ideal primitive H . (Operationally, the construction has oracle access to the function $x \mapsto \tilde{H}^H(x)$.) The construction depends on access to trusted, but public, randomness R .
2. The simulator is provided, as input, the subverted implementation \tilde{H} (a Turing machine) and the public randomness R ; it has oracle access to the target ideal functionality (\mathcal{F}) .

Point (2) is necessary, and desirable, as it is clearly impossible to achieve indifferenciability using a simulator that has no access to \tilde{H} (the distinguisher can simply query an input such that C will use a value that is modified by \tilde{H} while \mathcal{S} has no way to reproduce this). More importantly, we will show below that security will be preserved by replacing an ideal random oracle with a construction satisfying our definition (with an augmented simulator). Specifically, we prove a security preserving (i.e., replacement) theorem akin to those of [MRH04] and [CDMP05] for our adapted notions.

Definition 3 (H -crooked indifferenciability). *Consider a distinguisher $\hat{\mathcal{D}}$ and the following multi-phase real execution. Initially, the distinguisher $\hat{\mathcal{D}}$ commences the first phase: with oracle access to ideal primitive H the distinguisher constructs and publishes a subverted implementation of H ; this subversion is described as a deterministic polynomial time algorithm denoted \tilde{H} . (Recall that the algorithm \tilde{H} implicitly defines a subverted version of H by providing H to \tilde{H} as an oracle—thus $\tilde{H}^H(x)$ is the value taken by the subverted version of H at x .) Then, a uniformly random string R is sampled and published. Then the second phase begins involving a deterministic construction C : the construction C requires the random string R as input and has oracle access to \tilde{H} (the crooked version of H); explicitly this is the oracle $x \mapsto \tilde{H}^H(x)$. Finally, the distinguisher $\hat{\mathcal{D}}$, now with random string R as input and full oracle access to the pair (C, H) , returns a decision bit b . Often, we call $\hat{\mathcal{D}}$ the H -crooked-distinguisher.*

Consider now the corresponding multi-phase ideal execution with the same H -crooked-distinguisher $\hat{\mathcal{D}}$. The ideal execution introduces a simulator \mathcal{S} responsible for simulating the behavior of H ; the simulator is provided full oracle access to the ideal object \mathcal{F} . Initially, the simulator must answer any queries made to H by $\hat{\mathcal{D}}$ in the first phase. Then the simulator is given the random string R and the algorithm $\langle \tilde{H} \rangle$ (generated by $\hat{\mathcal{D}}$ at the end of the first phase) as input. In the second phase, the H -crooked-distinguisher $\hat{\mathcal{D}}$, now with random string R as input and oracle access to the alternative pair $(\mathcal{F}, \mathcal{S})$, returns a decision bit b .

We say that construction C is $(n_{\text{source}}, n_{\text{target}}, q_{\hat{\mathcal{D}}}, q_{\tilde{H}}, r, \epsilon)$ - H -crooked-indifferenciability from ideal primitive \mathcal{F} if there is an efficient simulator \mathcal{S} so that for any H -crooked-distinguisher $\hat{\mathcal{D}}$ making no more than $q_{\hat{\mathcal{D}}}(n)$ queries and producing a subversion \tilde{H} making no more than $q_{\tilde{H}}(n)$ queries, the real execution and the ideal execution are indistinguishable. Specifically,

$$\left| \Pr_{u, R, H} \left[\tilde{H} \leftarrow \hat{\mathcal{D}}^H(1^n); \hat{\mathcal{D}}^{C^{\tilde{H}}(R), H}(1^n, R) = 1 \right] - \Pr_{u, R, \mathcal{F}} \left[\tilde{H} \leftarrow \hat{\mathcal{D}}^H(1^n); \hat{\mathcal{D}}^{\mathcal{F}, \mathcal{S}^{\mathcal{F}}(R, \langle \tilde{H} \rangle)}(1^n, R) = 1 \right] \right| \leq \epsilon(n).$$

Here R denotes a random string of length $r(n)$ and both $H : \{0, 1\}^{n_{\text{source}}} \rightarrow \{0, 1\}^{n_{\text{source}}}$ and $\mathcal{F} : \{0, 1\}^{n_{\text{target}}} \rightarrow \{0, 1\}^{n_{\text{target}}}$ denote random functions where $n_{\text{source}}(n)$ and $n_{\text{target}}(n)$ are polynomials in the security parameter n . We let u denote the random coins of $\hat{\mathcal{D}}$. The simulator is efficient in the sense that it is polynomial in n and the running time of the supplied algorithm \tilde{H} (on inputs of length n_{source}). See Figure 3 for detailed illustration of the last phase in both real and ideal executions. (While it is not explicitly captured in the description above, the distinguisher $\hat{\mathcal{D}}$ is permitted to carry state from the first phase to the second phase.) The notation $C^{\tilde{H}}(R)$ denotes oracle access to the function $x \mapsto \tilde{H}(x)$.

Our main security proof will begin by demonstrating that in our particular setting, security in a simpler model suffices: this is the *abbreviated crooked indifferenciability* model, articulated below. We then show that—in light of the special structure of our simulator—it can be effectively lifted to the full model above.

Definition 4 (Abbreviated H -crooked indifferenciability). *The abbreviated model calls for the distinguisher to provide the subversion algorithm \tilde{H} at the outset (without the advantage of any preliminary queries to H). Thus, the abbreviated model consists only of the last phase of the full model. Formally, in the abbreviated model the distinguisher is provided*




Figure 3: The H -crooked indistinguishability notion: the distinguisher $\widehat{\mathcal{D}}$, in the first phase, manufactures and publishes a subverted implementation denoted as \tilde{H} , for ideal primitive H ; then in the second phase, a random string R is published; after that, in the third phase, algorithm C , and simulator S are developed; the H -crooked-distinguisher $\widehat{\mathcal{D}}$, in the last phase, either interacting with algorithm C and ideal primitive H , or with ideal primitive \mathcal{F} and simulator S , return a decision bit. Here, algorithm C has oracle access to \tilde{H} , while simulator S has oracle access to \mathcal{F} and \tilde{H} .

as a pair $(\widehat{\mathcal{D}}, \tilde{H})$, the random string R is drawn (as in the full model), and insecurity is expressed as the difference between the behavior of $\widehat{\mathcal{D}}$ on the pair $(C^{\tilde{H}}(R), H)$ and the pair $(\mathcal{F}, S^{\mathcal{F}}(R, \langle \tilde{H} \rangle))$. Specifically, the construction C is $(n_{\text{source}}, n_{\text{target}}, q_{\widehat{\mathcal{D}}}, q_{\tilde{H}}, r, \epsilon)$ -Abbreviated- H -crooked-indistinguishable from ideal primitive \mathcal{F} if there is an efficient simulator S so that for any H -crooked-distinguisher $\widehat{\mathcal{D}}$ making no more than $q_{\widehat{\mathcal{D}}}(n)$ queries and subversion algorithm \tilde{H} making no more than $q_{\tilde{H}}(n)$ queries, the real execution and the ideal execution are indistinguishable:

$$\left| \Pr_{u, R, H} \left[\widehat{\mathcal{D}}^{C^{\tilde{H}}(R), H}(1^n, R) = 1 \right] - \Pr_{u, R, \mathcal{F}} \left[\widehat{\mathcal{D}}^{\mathcal{F}, S^{\mathcal{F}}(R, \langle \tilde{H} \rangle)}(1^n, R) = 1 \right] \right| \leq \epsilon(n).$$

Here R denotes a random string of length $r(n)$ and both $H : \{0, 1\}^{n_{\text{source}}} \rightarrow \{0, 1\}^{n_{\text{source}}}$ and $\mathcal{F} : \{0, 1\}^{n_{\text{target}}} \rightarrow \{0, 1\}^{n_{\text{target}}}$ denote random functions where $n_{\text{source}}(n)$ and $n_{\text{target}}(n)$ are polynomials in the security parameter n . We let u denote the random coins of $\widehat{\mathcal{D}}$. The simulator is efficient in the sense that it is polynomial in n and the running time of the supplied algorithm \tilde{H} (on inputs of length n_{source}).

Observe that while the abbreviated simulator is a fixed algorithm, its running time may depend on the running time of \tilde{H} —in particular, the definition permits S sufficient running time to simulate \tilde{H} on a polynomial number of inputs.

Regarding the difference between these notions, observe that the distinguisher can “compile into” the subversion algorithm \tilde{H} any queries and pre-computation that might have been advantageous to carry out in phase I; such queries and pre-computation can also be mimicked by the distinguisher itself. This technique can effectively simulate the two phase execution with a single phase. Nevertheless, the models do make slightly different demands on the simulator which must be prepared to answer some queries (in Phase I) prior to knowledge of R and \tilde{H} .

Replacement with crooked indistinguishability. The notion of crooked indistinguishability is formalized for a specific ideal primitive, i.e., random oracles. We note that the formalization can be trivially generalized for all ideal primitives.

Security preserving (replacement) has been shown in the indistinguishability framework [MRH04]: if $C^{\mathcal{G}}$ is indistinguishable from \mathcal{F} , then $C^{\mathcal{G}}$ can replace \mathcal{F} in any cryptosystem, and the resulting cryptosystem in the \mathcal{G} model is at least as secure as that in the \mathcal{F} model. We next show that the replacement property also holds in the crooked indistinguishability framework.

Recall that, in the “standard” indistinguishability framework [MRH04, CDMP05], a cryptosystem can be modeled as an Interactive Turing Machine with an interface to an adversary \mathcal{A} and to a public oracle. There the cryptosystem is run by a “standard” environment \mathcal{E} . In our “crooked” indistinguishability framework, a cryptosystem has the interface to an adversary \mathcal{A} and to a public oracle. However, now the cryptosystem is run by a crooked-environment $\widehat{\mathcal{E}}$.

Consider an ideal primitive \mathcal{G} . Similar to the \mathcal{G} -crooked-distinguisher, we can define the \mathcal{G} -crooked-environment $\widehat{\mathcal{E}}$ as follows: Initially, the crooked-environment $\widehat{\mathcal{E}}$ manufactures and then publishes a subverted implementation of the

ideal primitive \mathcal{G} , and denotes it $\tilde{\mathcal{G}}$. Then $\hat{\mathcal{E}}$ runs the attacker \mathcal{A} , and the cryptosystem \mathcal{P} is developed. In the \mathcal{G} model, cryptosystem \mathcal{P} has oracle access to C whereas attacker \mathcal{A} has oracle access to \mathcal{G} ; note that, C has oracle access to $\tilde{\mathcal{G}}$, not to directly \mathcal{G} . In the \mathcal{F} model, both \mathcal{P} and \mathcal{A} have oracle access to \mathcal{F} . Finally, the crooked-environment $\hat{\mathcal{E}}$ returns a binary decision output. The definition is illustrated in Figure 4.

Definition 5. Consider ideal primitives \mathcal{G} and \mathcal{F} . A cryptosystem \mathcal{P} is said to be at least as secure in the \mathcal{G} -crooked model with algorithm C as in the \mathcal{F} model, if for any \mathcal{G} -crooked-environment $\hat{\mathcal{E}}$ and any attacker \mathcal{A} in the \mathcal{G} -crooked model, there exists an attacker $\mathcal{S}_{\mathcal{A}}$ in the \mathcal{F} model, such that:

$$\Pr[\hat{\mathcal{E}}(\mathcal{P}^{C^{\tilde{\mathcal{G}}}}, \mathcal{A}^{\mathcal{G}}) = 1] - \Pr[\hat{\mathcal{E}}(\mathcal{P}^{\mathcal{F}}, \mathcal{S}_{\mathcal{A}}^{\mathcal{F}}) = 1] \leq \epsilon.$$

where ϵ is a negligible function of the security parameter n .

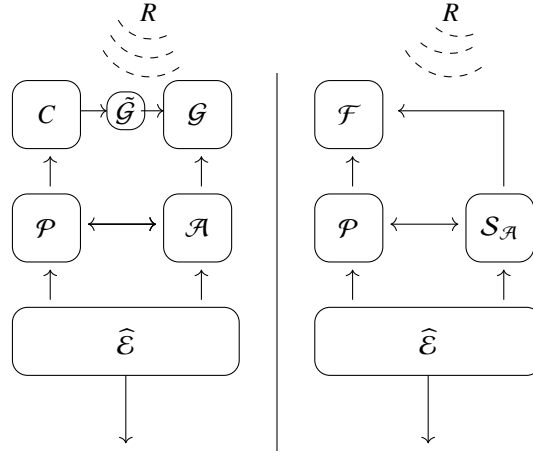


Figure 4: The environment $\hat{\mathcal{E}}$ interacts with cryptosystem \mathcal{P} and attacker \mathcal{A} : In the \mathcal{G} model (left), \mathcal{P} has oracle accesses to C whereas \mathcal{A} has oracle accesses to \mathcal{G} ; the algorithm C has oracle accesses to the subverted $\tilde{\mathcal{G}}$. In the \mathcal{F} model, both \mathcal{P} and $\mathcal{S}_{\mathcal{A}}$ have oracle accesses to \mathcal{F} . In addition, in both \mathcal{G} and \mathcal{F} models, randomness R is publicly available to all entities.

We now demonstrate the following theorem which shows that security is preserved when replacing an ideal primitive by a crooked-indifferentiable one:

Theorem 2. Consider an ideal primitive \mathcal{G} and a \mathcal{G} -crooked-environment $\hat{\mathcal{E}}$. Let \mathcal{P} be a cryptosystem with oracle access to an ideal primitive \mathcal{F} . Let C be an algorithm such that $C^{\mathcal{G}}$ is \mathcal{G} -crooked-indifferentiable from \mathcal{F} . Then cryptosystem \mathcal{P} is at least as secure in the \mathcal{G} -crooked model with algorithm C as in the \mathcal{F} model.

Proof. The proof is very similar to that in [MRH04, CDMP05]. Let \mathcal{P} be any cryptosystem, modeled as an Interactive Turing Machine. Let $\hat{\mathcal{E}}$ be any crooked-environment, and \mathcal{A} be any attacker in the \mathcal{G} -crooked model. In the \mathcal{G} -crooked model, \mathcal{P} has oracle access to C (who has oracle access to $\tilde{\mathcal{G}}$, not to directly \mathcal{G} .) whereas \mathcal{A} has oracle access to ideal primitive \mathcal{G} ; moreover crooked-environment $\hat{\mathcal{E}}$ interacts with both \mathcal{P} and \mathcal{A} . This is illustrated in Figure 5 (left part).

Since C is crooked-indifferentiable from \mathcal{F} (see Figure 3), one can replace $(C^{\tilde{\mathcal{G}}}, \mathcal{G})$ by $(\mathcal{F}, \mathcal{S})$ with only a negligible modification of the crooked-environment $\hat{\mathcal{E}}$'s output distribution. As illustrated in Figure 5, by merging attacker \mathcal{A} and simulator \mathcal{S} , one obtains an attacker $\mathcal{S}_{\mathcal{A}}$ in the \mathcal{F} model, and the difference in $\hat{\mathcal{E}}$'s output distribution is negligible. \square

Restrictions (of using crooked indifferentiability). Ristenpart et al. [RSS11] has demonstrated that the replacement/composition theorem (Theorem 1) in the original indifferentiability framework only holds in single-stage settings.

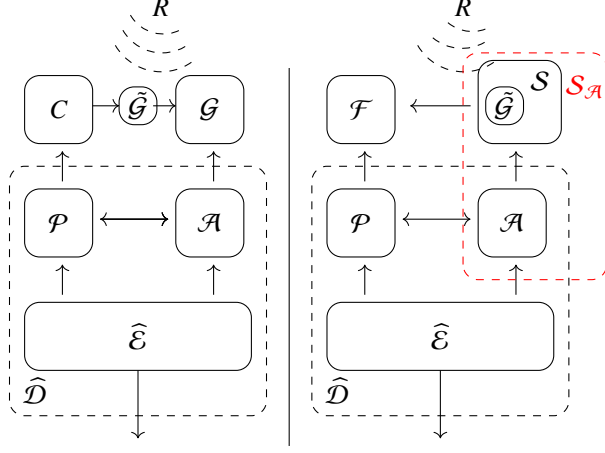


Figure 5: Construction of attacker $\mathcal{S}_{\mathcal{A}}$ from attacker \mathcal{A} and simulator \mathcal{S}

We remark that, the same restriction also applies to our replacement/composition theorem (Theorem 2). We leave it as our future work to extend our crooked indistinguishability to the multi-stage settings where disjoint adversaries are split over several stages.

3 The Construction

From this point on, we use \mathcal{D} rather than $\widehat{\mathcal{D}}$ to denote the distinguisher in our crooked indistinguishability model. For a security parameter n and a (polynomially related) parameter ℓ , the construction depends on public randomness $R = (r_1, \dots, r_\ell)$, where each r_i is an independent and uniform element of $\{0, 1\}^n$.

The source function of the construction is expressed as a family of $\ell + 1$ independent random oracles:

$$\begin{aligned} h_0 &: \{0, 1\}^n \rightarrow \{0, 1\}^{3n}, \\ h_i &: \{0, 1\}^{3n} \rightarrow \{0, 1\}^n, \quad \text{for } i \in \{0, \dots, \ell\}. \end{aligned}$$

These can be realized as slices of a single random function $H : \{0, 1\}^{n'} \rightarrow \{0, 1\}^{n'}$, with $n' = 3n + \lceil \log \ell + 1 \rceil$ by an appropriate convention for embedding and extracting inputs and values. In the few cases where we need to be precise, we write $\{0, \dots, L\} \times \{0, 1\}^{3n} = \{0, 1\}^{n'}$, where $\ell + 1 \leq L = 2^{\lceil \log \ell + 1 \rceil}$, and let $[i, x]$ denote a unique element of $\{0, \dots, L\} \times \{0, 1\}^{3n}$ to implicitly correspond to the element x in the domain for h_i : for concreteness, define $[0, x] = (0, 0^{2n}x)$ and $[i, x] = (i, x)$ for $i > 0$. The output of H is treated as a string of the correct length without any special indication. Given subverted implementations $\{\tilde{h}_i\}_{i=0, \dots, \ell}$ (defined as above by the adversarially-defined algorithm \tilde{H}), the corrected function is defined as:

$$C^{\tilde{H}^H}(x) \stackrel{\text{def}}{=} \tilde{h}_0 \left(\bigoplus_{i=1}^{\ell} \tilde{h}_i(x \oplus r_i) \right),$$

where $R = (r_1, \dots, r_\ell)$ is sampled uniformly after \tilde{h} is provided (and then revealed to the public). We refer to this function C as $\tilde{h}_R(x)$ and, for the purposes of analysis, also give a name to the ‘‘internal function’’ $\tilde{g}_R(\cdot)$:

$$\tilde{g}_R(x) = \bigoplus_{i=1}^{\ell} \tilde{h}_i(x \oplus r_i).$$

Thus

$$\tilde{h}_R(x) \stackrel{\text{def}}{=} \tilde{h}_0(\tilde{g}_R(x)) = \tilde{h}_0 \left(\bigoplus_{i=1}^{\ell} \tilde{h}_i(x \oplus r_i) \right).$$

We wish to show that such a construction is indiffereniable from an actual random oracle (with the proper input/output length). This implies, in particular, that values taken by $\tilde{h}_R(\cdot)$ at inputs that have not been queried have negligible distance from the uniform distribution.

Theorem 3. *We treat a function $h : \{0, 1\}^{n'} \rightarrow \{0, 1\}^{n'}$, with $n' = 3n + \lceil \log \ell + 1 \rceil$, as implicitly defining a family of random oracles*

$$\begin{aligned} h_0 &: \{0, 1\}^{3n} \rightarrow \{0, 1\}^n, & \text{and} \\ h_i &: \{0, 1\}^n \rightarrow \{0, 1\}^{3n}, & \text{for } i > 0, \end{aligned}$$

by treating $\{0, 1\}^{n'} = \{0, \dots, L-1\} \times \{0, 1\}^{3n}$ and defining $h_i(x) = h(i, \cdot)$, for $i = 0, \dots, \ell \leq L-1$. (Output lengths are achieved by removing the appropriate number of trailing symbols). We will use the setting $\ell > n + 4$. Consider a (subversion) algorithm \tilde{H} so that $\tilde{H}^h(x)$ defines a subverted random oracle \tilde{h} . Assume that for every h (and every i),

$$\Pr_{x \in \{0, 1\}^{n'}} [\tilde{h}(i, x) \neq h(i, x)] \leq \epsilon(n) = \text{negl}(n). \quad (1)$$

The construction $\tilde{h}_R(\cdot)$ is $(n', n, q_{\mathcal{D}}, q_{\tilde{H}}, r, \epsilon')$ -indiffereniable from a random oracle $F : \{0, 1\}^n \rightarrow \{0, 1\}^n$, where $\epsilon' = O(\ell q_{\tilde{D}} q_{\tilde{H}} / \sqrt{2^n} + \sqrt{\ell} q_{\tilde{D}} \epsilon^{1/16})$, $q_{\mathcal{D}}$ is the number of queries made by the distinguisher \mathcal{D} and $q_{\tilde{H}}$ is the number of queries made by \tilde{H} as in Definition 3. $q_{\mathcal{D}}$ and $q_{\tilde{H}}$ are both polynomial functions of n .

Domain extension. It is shown in [CDMP05] that an arbitrary-length random oracle can be constructed from a fixed-length one in the indiffereniability framework, thus our result can be easily generalized to handle the case for correcting a subverted arbitrary input length hash.

Roadmap for the proof. We first describe the simulator algorithm. The main challenge for the simulator is to ensure the consistency of two ways of generating the output values of the construction. The idea for simulation is fairly simple: the $h_i(\cdot)$, for $i > 0$, are treated as random functions; h_0 is programmed to suitably agree with F . Specifically, once the value $\tilde{g}_R(x)$ is determined, the value $F(x)$ is used to program h_0 on input $\tilde{g}_R(x)$. Our simulator eagerly “completes” any constellation of related points once any of its constituent points have been queried by the distinguisher; in particular, any query to $h_i(x \oplus r_i)$ prompts the simulator to evaluate $\tilde{h}_j(x \oplus r_j)$ for all j ; this provides enough information to properly program $h_0(\cdot)$ at the desired location $\tilde{g}_R(x)$ to ensure consistency with $F(\cdot)$. This convention leads to a relatively simple invariant maintained during a typical interaction with the distinguisher: all constellations thus far touched by the distinguisher are correctly programmed. Of course, proving that this simulation doesn’t run into consistency problems—for example, a circumstance where $h_0(y)$ has unfortunately already been assigned a value at the moment it would be appropriate to program it—is part of the proof of correctness.

There are two fundamental obstacles that hinder the simulation: (1) for some x , after completing the constellation associated with x , the simulator finds that h_0 has been previously queried on $\tilde{g}_R(x)$ —thus the simulator’s hands are tied when it comes time to program this value of h_0 (to be consistent with F); (2) the distinguisher queries some input x such that $\tilde{g}_R(x)$ falls into the incorrect (subverted) portion of \tilde{h}_0 . It’s convenient to further separate various specific patterns of behavior that can give rise to collisions as described in (1) above. These are reflected in a sequence of four “hybrid games” that interpolate between the two interactions of interest—the result of the distinguisher when interacting with the construction and the result when interacting with the simulator.

To further simplify the proof, we initially focus on the abbreviated variant of crooked indiffereniability where the distinguisher does not have the luxury of making preliminary queries to the hash function while it formulates its subversion algorithm \tilde{H} . This places all queries of the distinguisher on the same footing and somewhat simplifies bookkeeping. We then return at the conclusion of the main proof show that the same techniques can be applied to the full setting.

4 Security Proof

We begin with an abstract formulation of the properties of our construction, and then transition to the detailed description of the simulator algorithm and its effectiveness.

Anticipating the full description of the simulator, we formally set down the notion of a “constellation” of points.

Definition 6 (Constellation). For a fixed $R = (r_1, \dots, r_\ell)$ and an element $x \in \{0, 1\}^n$, the constellation of x is the set of points $\{[i, x \oplus r_i] \mid 0 < i \leq \ell\}$; that is, the set includes, for each $i > 0$, the point in the domain of (the implicitly defined) h_i defining $h_i(x \oplus r_i)$. While this depends on R , we suppress this as it can always be inferred from context.

4.1 The simulator

The major task of the simulator is to ensure the answers to appropriate h_0 queries are consistent with the value of $F(\cdot)$. In general, queries to the h_i , for $i > 0$, are simply treated as random functions and lazily evaluated and cached as usual. The function h_0 must be given special treatment: ideally, the simulator would like to ensure the equality $\tilde{h}_0(\tilde{g}_R(x)) = F(x)$. This poses some challenges because the simulator cannot typically identify—for a particular y in question—an x for which $y = \tilde{g}_R(x)$; thus there will be circumstances where the simulator simply cannot correctly program $h_0(y)$ with the corresponding $F(\cdot)$ value. A further challenge is that values of $\tilde{h}_0(\cdot)$ cannot, in general, be “programmed” to agree with F in circumstances where \tilde{h}_0 has been subverted. To minimize the chance that one of the simulator’s “inconsistencies” is observable by the distinguisher, the simulator aggressively computes $\tilde{g}_R(x)$ for any x whose constellation $\{h_i(x \oplus r_i)\}$ the distinguisher (even partially) examines. In particular, if the distinguisher queries $h_i(x \oplus r_i)$, the simulator will immediately determine the values for all $\tilde{h}_j(x \oplus r_j)$ and, if possible, correctly program $h_0(\cdot)$ at $\tilde{g}_R(x)$. In this way, the simulator typically guarantees that constellations which the distinguisher has examined are correctly programmed. Of course, the distinguisher may speculatively query h_0 on other locations but, with this simulator, it will not be easy for the distinguisher to force a lie at a location that will result in a future inconsistency (so long as the game only runs for polynomially many steps). A similar difficulty arises in cases where the distinguisher queries an $h_i(x \oplus r_i)$ in a constellation for which $\tilde{h}_0(\tilde{g}_R(x))$ has been subverted; again, this situation may interfere with the simulator’s ability to properly program \tilde{h}_0 . Likewise, we will see that such queries are difficult for the distinguisher to discover.

Formally, we define the simulator \mathcal{S} (for the abbreviated model) below. After the proof that this simulator achieves (abbreviated) crooked indistinguishability (Definition 3), we show that the simulator can be lifted to one that achieves full indistinguishability (Definition 4).

THE ABBREVIATED SIMULATOR

The simulator is provided R , the randomness associated with the security game, and $\langle \tilde{H} \rangle$, the subversion algorithm. As in our previous discussions, we implicitly associate H with the family of functions h_i for $i \geq 0$ with particular conventions for padding or embedding inputs and outputs to realize the variety of domains and ranges for the h_i ; for simplicity we assume that all queries (made by \mathcal{D} or \tilde{H}) to H have a unique interpretation as a query to an h_i . Roughly, \mathcal{S} “lazily” maintains the random oracles via tables that are populated according to conventions that minimize the risk of a conflict arising from its handling of the function h_0 .

h_0		h_1		...	h_ℓ	
Query	$h_0(x_i)$	Query	$h_1(x_i)$		Query	$h_\ell(x_i)$
x_1	$v_{1,0}$	x_1	$v_{1,1}$		x_1	$v_{1,\ell}$
\vdots	\vdots	\vdots	\vdots		\vdots	\vdots
x_{q_0}	$v_{q_0,0}$	x_{q_1}	$v_{q_1,1}$		x_{q_ℓ}	$v_{q_\ell,\ell}$

Table 1: Tables maintained by \mathcal{S} .

\mathcal{S} responds to queries from \mathcal{D} with following procedure, which also describes the mechanism for maintaining the value tables.

- All tables are initially empty.
- A query x to $h_i(\cdot)$ is answered as follows. If $i = 0$ and the table for h_0 holds a value for x , this is returned and no further action is taken. If x does not appear in the table, a uniformly random value v is drawn from

$\{0, 1\}^n$, added to the table for h_0 , and returned. Otherwise, $i > 0$ and the situation is more complicated, as \mathcal{S} wishes to compute the value \tilde{g} associated with this query. In this case, the simulator then proceeds as follows.

1. \mathcal{S} determines the “adjusted” query $x' := x \oplus r_i$, and ensures that all points in the constellation—that is $h_j(x' \oplus r_j)$ for $j > 0$ —have assigned values: specifically, for each j , if the result of $h_j(x' \oplus r_j)$ has not yet been determined, this value is drawn uniformly and added to the table.
2. The next task is to evaluate $\tilde{g}_R(x')$, the corresponding query to h_0 , which would ideally be programmed to agree with F . For this purpose, \mathcal{S} then runs the implementation \tilde{h}_j (using \tilde{H}) on $x' \oplus r_j = x \oplus r_j \oplus r_i$, for all $j \in [\ell]$, to derive the value

$$\tilde{g}_R(x') = \bigoplus_{j=1}^{\ell} \tilde{h}_j(x' \oplus r_j),$$

where $[\ell]$ is used to denote the set $\{1, \dots, \ell\}$.

During the execution of \tilde{H} on those inputs, \mathcal{S} must determine any additional random oracle queries issued by the implementation \tilde{H}^{h^*} . These are answered by a less aggressive convention: \mathcal{S} first checks whether the query, say $h_k(y)$, appears in the table; if so, the stored value is used; if not, \mathcal{S} uses a uniformly random value u as answer and records it in the table. (Note that this does not lead to forced evaluation of the full constellations in which these queries lie.)

3. \mathcal{S} checks whether $\tilde{g}_R(x')$ has been previously queried in h_0 ; if not, \mathcal{S} properly programs h_0 : specifically, \mathcal{S} queries F on x' , collects the response $F(x')$, and assigns $h_0(x') = F(x')$. Finally, \mathcal{S} returns to the distinguisher the value appearing in the appropriate table to the original query x . Note that a value always appears at this point in light of step (1) above. (N.b., in many cases, for convenience, we assume that \mathcal{S} returns to the distinguisher all values in the constellation of x —that is, all $h_j(x' \oplus r_i)$ for $j > 0$.)

Probability analysis. We prove indistinguishability by introducing a sequences of four games that connect the real construction and the simulator. In Game 2, we define three crisis events Subv, Pred and Selfref, which play a central role in distinguishing Game 1, the real construction, from Game 4, the simulator interaction game. According to Theorem 4, the gap between (C, H) and $(\mathcal{F}, \mathcal{S})$ is bounded by $\Pr[\text{Subv}] + \Pr[\text{Pred}] + \Pr[\text{Selfref}]$.

Lastly, the three events are shown to be negligible in Theorem 24, Theorem 20, and Theorem 23.

4.2 The hybrid games; the crisis events

4.2.1 The games

We introduce the security games used to organize the main proof. Game 1 corresponds to interaction with the construction, while Game 4 corresponds to interaction with the simulator described above. The intermediate games correspond to hybrid settings that are convenient in the proof. The main “crisis events” that reflect circumstances where the simulator could fail to provide consistency are defined with respect to Game 2.

The description of each game indicates how queries to the two oracles with which \mathcal{D} interacts are answered. We adopt the notational convention that \mathcal{A} denotes the rule for answering oracle queries to F —corresponding to the construction and the ideal functionality—while \mathcal{B} denotes the rule for answering oracle queries to h_i —corresponding to H and the simulator. For expository purposes, several variants are presented for each game; these variants all yield precisely the same transcript with the distinguisher but reflect differing internal conventions for generating the responses.

Game 1 (The construction). $\mathcal{A} = C^H, \mathcal{B} = H$. To most easily compare with the following games, we explicitly express the game as a “data preparation step” and an “interaction” step:

1. Data preparation: Uniformly select the value $h_i(x)$ for each $0 \leq i \leq \ell$ and each x in the domain.
2. Interaction: If \mathcal{D} queries \mathcal{A} for x , \mathcal{A} returns $\tilde{h}_0(\tilde{g}_R(x))$; if \mathcal{D} queries \mathcal{B} for (i, x) , \mathcal{B} returns $h_i(x)$.

Some conventions and terminology. In general, we say that “the distinguisher \mathcal{D} queries the constellation of x ” (for $x \in \{0, 1\}^n$) when \mathcal{D} queries any $h_i(x \oplus r_i)$ term for $i > 0$. (This is consistent with the definition of constellation above.) Beginning with Game 2.2, we consider procedures for maintaining function values by lazy evaluation. In this setting we define the notion of *completing* the constellation of x as the following procedure: $h_i(x \oplus r_i)$ is evaluated for all elements in the constellation—that is, they are assigned to uniformly random values if they have not yet been assigned—followed by evaluation of $\tilde{h}_i(x \oplus r_i)$ (using the subversion algorithm) for each $x \oplus r_i$ in the constellation. The evaluation of $\tilde{h}_i(x \oplus r_i)$ will in general require evaluation of $h_i(\cdot)$ at a collection of further points, which are likewise assigned to new uniformly random values unless they have already been assigned. We also use the notation $\tilde{g}_R(x)$ to denote the sum of $\tilde{h}_j(x \oplus r_j)$ for $0 < j \leq \ell$. This is directly analogous to the “aggressive” constellation evaluation discussed in the formal discussion of the simulator above.

The notation $\tilde{g}_R(\cdot)$ described above deserves special consideration. In Game 1.1, $\tilde{g}_R(x)$ can be unambiguously defined for all x (as $\tilde{h}_i(\cdot)$ can be unambiguously defined); the same can be said of Game 2.1, below. For subsequent games, there is no immediate way to define $\tilde{g}_R(x)$ for all x , as these values may depend on the queries made by \mathcal{D} ; however, the syntactic definition of $\tilde{g}_R(x)$ (as the sum of appropriate $\tilde{h}_j(x \oplus r_j)$) is well-defined for all constellations that have been completed by \mathcal{S} ; thus, we continue to use this notation to indicate this sum, with the understanding that it is only meaningful in settings where the constellation has been completed.

For convenience, we assume the distinguisher does not make repetitive queries.

Game 2 (The construction with unsubverted h_0). Game 2 has two variants. Game 2.1 is identical to Game 1 with the exception that the construction is replaced with an idealized one that uses the unsubverted h_0 . Game 2.2 merely shifts the perspective of Game 2.1 to one involving tables.

- **Game 2.1:**

1. Data preparation: Uniformly select the value $h_i(x)$ for each $0 \leq i \leq \ell$ and each x in the domain.
2. Interaction: If \mathcal{D} queries \mathcal{A} for x , \mathcal{A} returns $h_0(\tilde{g}_R(x))$; if \mathcal{D} queries \mathcal{B} for (i, x) , \mathcal{B} returns $h_i(x)$.

- **Game 2.2:** \mathcal{A} and \mathcal{B} maintain tables, T_F and T_H , respectively, storing the point evaluations of F and H arising during the interaction. Initially, the tables are empty. Interaction with \mathcal{D} proceeds as follows:

- If \mathcal{D} queries \mathcal{B} for (i, x) : if $h_i(x)$ appears in the table T_H , return the recorded value; otherwise, draw a uniform value, assign $h_i(x)$ to the value by adding this entry to the table T_H , and return it. Additionally, if $i > 0$, complete the constellation of $x \oplus r_i$ and uniformly assign $h_0(\tilde{g}_R(x \oplus r_i))$ in T_H unless it has been previously assigned. **If $F(x)$ is unassigned in T_F , it is assigned to the value $h_0(\tilde{g}_R(x \oplus r_i))$.** Otherwise $F(x)$ was already assigned and this conflict is not explicitly managed.
- If \mathcal{D} queries \mathcal{A} for x : if $F(x)$ appears in the table T_F , return the recorded value; otherwise, complete the constellation of x and **properly program $F(x)$ at the value $h_0(\tilde{g}_R(x))$** . Specifically, completion determines the value $\tilde{g}_R(x)$; if $h_0(\tilde{g}_R(x))$ is unassigned, it is assigned a random value in T_H . Finally, \mathcal{A} returns $h_0(\tilde{g}_R(x))$ and $F(x)$ is assigned to this value in T_F .

For technical reasons later, we fill in the empty cells of T_F and T_H in Game 2.2 after the interaction. Empty cells of T_H are filled in by uniformly chosen values while empty cells of T_F are filled in by $F(x) := h_0(\tilde{g}_R(x))$. The resulting T_H is denoted by h_* . Also, in Game 2.2, we say $y \in \{0, 1\}^{3n}$ is in the *subverted area* of h_0 if, in h_* , $\tilde{h}_0(y) \neq h_0(y)$.

Game 3. Game 3 also has two variants. Game 3.1 is a small pivot from Game 2.2 which handles conflicts by a different convention. Game 3.2 is identical to Game 3.1 with all randomness selected prior to the game.

- **Game 3.1:** \mathcal{A} and \mathcal{B} maintain tables, T_F and T_H respectively, storing the evaluated terms during the interaction. Initially, the data sets are empty. Interaction with \mathcal{D} proceeds as follows:

- If \mathcal{D} queries \mathcal{B} for (i, x) : if $h_i(x)$ appears in the table T_H , return the recorded value; otherwise, draw a uniform value, assign $h_i(x)$ to the value by adding this entry to the table T_H , and return it. Additionally, if $i > 0$, complete the constellation of $x \oplus r_i$ and uniformly assign $F(x)$ in T_F unless it has been previously assigned. **If $h_0(\tilde{g}_R(x \oplus r_i))$ is unassigned in T_H , it is assigned to the value $F(x)$ and we say that h_0 has been programmed.** Otherwise h_0 was already assigned and this conflict is not explicitly managed.

- If \mathcal{D} queries \mathcal{A} for x : if $F(x)$ appears in the table T_F , return the recorded value; otherwise, uniformly select assign $F(x)$ and return this value, complete the constellation of x and **properly program $h_0(\tilde{g}_R(x))$ at the value $F(x)$** . Specifically, completion determines the value $\tilde{g}_R(x)$; if $h_0(\tilde{g}_R(x))$ is unassigned, it is assigned to $F(x)$ and we say that h_0 has been *programmed*; otherwise h_0 was already assigned and this conflict is not explicitly managed.

- **Game 3.2:** Game 3.2 is **identical to Game 3.1, but all randomness is selected a priori**; note that responses are still defined in terms of tables as defined in 3.1. Two data sets, DS_1 and DS_2 , are drawn: DS_1 contains uniformly selected values for each $F(x)$ and $h_i(x)$ for all $0 < i \leq \ell$ and all $x \in \{0, 1\}^n$; DS_2 contains uniformly selected values $h_0(y)$ for all $y \in \{0, 1\}^{3n}$. (Note that all entries are selected independently and uniformly.) The game is a straightforward adaptation of Game 3.1: When \mathcal{A} or \mathcal{B} generate values (for table insertion) for T_F or T_H (for h_i for $i > 0$), these values are drawn from DS_1 . The function h_0 is treated specially: If h_0 is programmed at y , the value $h_0(y)$ is drawn from (the table for F in) DS_1 . Otherwise the value is drawn from DS_2 .

Game 4 (The abbreviated simulator). Observe that queries to \mathcal{A} (that is \mathcal{F}) in the interaction between \mathcal{D} and $(\mathcal{F}, \mathcal{S}^{\mathcal{F}})$ do not result in associated queries to \mathcal{S} (cf. Game 3). Game 4 is **identical to the interaction between \mathcal{D} and $(\mathcal{F}, \mathcal{S}^{\mathcal{F}})$, but all randomness is selected a priori**. Prepare two data sets DS_1 and DS_2 : DS_1 contains uniformly selected values for $F(x)$ and $h_i(x)$ for all $0 < i \leq \ell$ and all $x \in \{0, 1\}^n$; DS_2 contains uniformly selected $h_0(y)$ for all $y \in \{0, 1\}^{3n}$. (Note that all data are selected independently and uniformly.) Game 4.1 is then carried out with the following conventions for randomness: All queries to F (either direct responses from \mathcal{A} or queries by $\mathcal{B} = \mathcal{S}$, the simulator) are determined by DS_1 . When \mathcal{B} evaluates h_i for $i > 0$ and records these assignments in T_H , values are likewise drawn from DS_1 . When \mathcal{B} assigns $h_0(y)$ in T_H to a “fresh random value,” this is drawn from DS_2 ; when $h_0(y)$ is programmed in T_H to coincide with F , this value is drawn from DS_1 .

We define the following *crisis events* with respect to a distinguisher \mathcal{D} in Game 2.2.

$$\begin{aligned} \text{Pred} &= \left\{ \text{for some } x, \text{ the distinguisher queries the constellation of } x \text{ for the first time and } h_0(y), \text{ where} \right. \\ &\quad \left. y = \tilde{g}_R(x), \text{ has already been assigned a value by the simulator} \right\}, \\ \text{Subv} &= \{ \text{the distinguisher queries a constellation } x \text{ such that } \tilde{g}_R(x) \text{ falls into the subverted area of } h_0 \}, \\ \text{Selfref} &= \left\{ \text{for some } x, \text{ the distinguisher queries a constellation } x \text{ such that } h_0 \text{ at } y = \tilde{g}_R(x) \text{ is evaluated} \right. \\ &\quad \left. \text{during completion of the constellation } x \right\}. \end{aligned}$$

Note that the crisis event *Pred* involves the *first* query made to the constellation of x . (This distinction can be ignored under the assumption that the simulator returns all constellation values to the distinguisher which are treated as regular distinguisher queries and so will not be repeated.)

Definition 7 (Transcript of a game). *For one of the games above, a distinguisher \mathcal{D} and randomness R , we define the ***R-transcript*** (or ***transcript***) of the game as the random variable given by R and the ordered sequence of (query, answer) pairs. We encode (query, answer) pairs as tuples of the form (x, F, y) or (x, h_i, y) , which indicates that $F(x) = y$ and $h_i(x) = y$, respectively. For the transcript α , we let $\alpha[k]$ denote the first k pairs in α . As a reminder that a transcript α determines R , we let $\alpha[0] = R$.*

The notational convention that $\alpha[0] = R$ is occasionally useful in induction proofs of equality between two transcripts—equality at zero indicates that they correspond to the same randomness.

The analysis focuses on the transcripts of a fixed deterministic distinguisher arising from its interactions in Game 1 through 4. As mentioned above, and discussed more fully below, the variants of each game (e.g., Games 2.1 and 2.2) yield precisely the same distributions of transcripts. (Indeed, for the same random choices these games yield identical transcripts.) To reflect this, for a fixed deterministic distinguisher we adopt the notation α_i , for $i = 1, 2, 3, 4$, to denote the transcript arising from interaction given by Game i .

Without loss of generality, we assume at the end of the interaction there exists no $x \in \{0, 1\}^n$ such that $F(x)$ is queried but the constellation of x is not. Observe that any distinguisher may be placed in this normal form by appending any necessary extra queries at the end of its native sequence of queries to appropriately query every constellation associated with a query to F .

Theorem 4. For any distinguisher \mathcal{D} , $\|\alpha_1 - \alpha_4\|_{t.v.} \leq \Pr[\text{Subv}] + \Pr[\text{Pred}] + \Pr[\text{Selfref}]$.

Theorem 4 is proved by the sequence of lemmas below. We recall, at the outset, that if $X : \Omega \rightarrow V$ and $Y : \Omega \rightarrow V$ are two random variables defined on the same probability space Ω and $\Pr[X \neq Y] = \epsilon$, then $\|X - Y\|_{tv} \leq \epsilon$.

Lemma 5. The transcripts of Game 2.1 and Game 2.2 have the same distribution.

Proof. It suffices to notice that the distributions of the query in Game 2.1 and Game 2.2 are identical because the rule of Game 2.2 guarantees all the $h_i(\alpha)$ for $0 \leq i \leq \ell$, like those in Game 2.1, are chosen uniformly and independently. \square

Notice that the three crisis events defined in Game 2.2 can also be properly defined in 2.1. The proof of Lemma 5 reveals that their probabilities are the same for Game 2.1 and 2.2 since all the three versions of Game 2 have the same distribution of h_* . In the rest of the paper, we abuse the notation and use Subv, Pred and Selfref to denote the crisis events in all the three versions of Game 2.

Lemma 6. For any distinguisher \mathcal{D} , $\|\alpha_1 - \alpha_2\|_{tv} \leq \Pr[\text{Subv}]$.

Proof. Notice that if the data preparation step prepares the same h_* for Game 1 and Game 2.1, the only way to distinguish them is to find a $\tilde{g}_R(x)$ falling into the subverted area of h_0 . That is to say, the transcripts in the two games are different only when Subv occurs. \square

Lemma 7. For any distinguisher \mathcal{D} , $\|\alpha_2 - \alpha_3\|_{tv} \leq \Pr[\text{Pred}]$.

Proof. Imagine the two parties $(\mathcal{A}, \mathcal{B})$ use the same random bits in Game 2.2 and Game 3.1. Then, Game 2.2 and Game 3.1 have the same query result unless Pred occurs in Game 2.2, which has the probability $\Pr[\text{Pred}]$. \square

Lemma 8. For any distinguisher \mathcal{D} the transcripts of Game 3.1 and Game 3.2 have the same distribution.

Proof. The distributions of the entrees of T_F and T_H in Game 3.1 are same as those in Game 3.2, these games have the same transcript distributions. \square

In the next three lemmas, we denote by $\beta_3[k]$ (and $\beta_4[k]$, respectively) the data in T_H in Game 3.2 (and Game 4, respectively) after k th query and answer. Similar to the conventions in α , we let $\beta_3[0](\beta_4[0])$ be R . Elements in β , DS_1 , and DS_2 are given the same tuple syntax as those in the transcript. For any positive k , we say $\beta_3[k]$ agrees with $\beta_4[k]$ if, for any $m \leq k$, $(x, h_0, a) \in \beta_3[m]$ and $(x, h_0, b) \in \beta_4[m]$, we have $a = b$. For $(x, h_0, a) \in \beta_3[k]$ (or $\beta_4[k]$), we say the pair (x, h_0) is a *free term* in $\beta_3[k]$ (or $\beta_4[k]$) if there is a $(x, h_0, b) \in DS_2$ such that $a = b$. If the term (x, h_0) is not free, we say that it is *programmed*. (These notions only apply to h_0 , reflecting the two mechanisms for assigning h_0 values in the games above.)

Lemma 9. Consider the interaction of a distinguisher \mathcal{D} with Games 3.2 and 4 given by the same randomness DS_1 , DS_2 and R . For any $k > 0$, if $\beta_3[k]$ agrees with $\beta_4[k]$ then $\alpha_3[k] = \alpha_4[k]$ and the $k + 1$ st queries of the two games are identical.

Remark: Note that agreement, in the sense defined above, is much weaker than equality or setwise equality. However, Lemma 9 implies that to prove $\alpha_3 = \alpha_4$, it suffices to show that $\beta_3[k]$ agrees with $\beta_4[k]$ for any k .

Proof. We will prove by induction that $\alpha_3[i] = \alpha_4[i]$ for each $0 \leq i \leq k$. Of course $\alpha_3[0] = \alpha_4[0]$, as these are both R . Suppose that $\alpha_3[i - 1] = \alpha_4[i - 1]$ for an $1 \leq i \leq k$. Since the first $i - 1$ queries have been identical and have received the same responses, the i th query in Game 3.2 and Game 4 are identical. If the i th query is made to F or h_i for $i > 0$, the responses are identical because the datasets are equal. Otherwise, the query is made to h_0 , in which case the responses are identical because $\beta_3[k] = \beta_4[k]$. Of course this also implies that the $k + 1$ st queries are identical. \square

Lemma 10. Consider the interaction of a distinguisher \mathcal{D} with Games 3.2 and 4 given by the same randomness DS_1 , DS_2 and R . For any $k > 0$, if $\beta_3[k]$ agrees with $\beta_4[k]$, $\beta_4[k] \subset \beta_3[k]$.

Proof. Assume that $\beta_3[k]$ agrees with $\beta_4[k]$; we proceed by induction to show that $\beta_4[j] \subset \beta_3[j]$ for all $j \leq k$. Of course $\beta_4[0] \subset \beta_3[0]$ as these are both R . For an j in the range $1 \leq j \leq k$, suppose $\beta_4[j - 1] \subset \beta_3[j - 1]$. Note that the j th query and response in Game 3.2 and Game 4 are identical by Lemma 9. We consider the various cases separately:

- If the j th query is (x, F) for some x : Then $\beta_4[j] = \beta_4[j - 1] \subset \beta_3[j - 1] \subset \beta_3[j]$, as desired. Note, in general, that $\beta_3[j - 1] \neq \beta_3[j]$ for such a query.
- If the j th query is (x, h_0) for some x : Note that $\beta_4[j]/\beta_4[j - 1] \subset \{(x, h_0, y)\}$ for some y and that $(x, h_0, y) \in \beta_3[j]$.
- If the j th query is (x, h_i) for some x and positive i : In both games, this calls for the constellation to be completed and $h_0(\tilde{g}_R(x \oplus r_i))$ to be evaluated. As the datasets are equal in the two games, the only possible disagreement that could arise during completion of the constellations must occur on a query to h_0 : however, this is excluded by agreement of $\beta_3[k]$ and $\beta_4[k]$. It follows that the h_i are evaluated at exactly the same sequence of inputs, yielding the same results, by the two completions and hence that any table assignment added to β_4 must also appear in β_3 ; additionally, $\tilde{g}(x \oplus r_i)$ takes the same value in the two games. Finally, $h_0(\tilde{g}(x \oplus r_i))$ is evaluated in both games which, by the same considerations, must yield the same value and must be included in both β_3 and β_4 . \square

Next, we introduce two further events in Game 3.2 and proceed to our main result.

- **Forward Prediction**, denoted $\text{Pred}^{\rightarrow}$: For some y , the constellation of y is completed for the first time at which point it is discovered that $h_0(\tilde{g}(y))$ has previously been evaluated.
- **Backward Prediction**, denoted Pred^{\leftarrow} : For some x , $h_0(\tilde{g}_R(x))$ is queried, either directly by \mathcal{D} or during completion, after \mathcal{D} has already queried $F(x)$ (to \mathcal{A}) but prior to \mathcal{D} making any query to the constellation x .

Lemma 11. *Consider the interaction of a distinguisher \mathcal{D} with Games 3.2 and 4 given by the same randomness DS_1 , DS_2 and R . For any $k > 0$, if $\text{Pred}^{\rightarrow}$ and Pred^{\leftarrow} do not occur in Game 3.2, $\beta_3[k]$ agrees with $\beta_4[k]$.*

Proof. We proceed by induction. First, $\beta_3[0]$ agrees with $\beta_4[0]$. Suppose now that $\beta_4[k - 1]$ agrees with $\beta_3[k - 1]$: we analyze three cases to ensure that $\beta_4[k]$ agrees with $\beta_3[k]$. Note, at the outset, that from Lemma 9 the k th queries made by \mathcal{D} are identical in the two games.

- If \mathcal{D} queries (x, F) for some x : $\beta_4[k] = \beta_4[k - 1]$ and hence agrees with $\beta_3[k]$.
- If \mathcal{D} queries for (x, h_0) for some x :
 1. if (x, h_0) is not evaluated yet in $\beta_3[k - 1]$, by Lemma 2, it is not in $\beta_4[k - 1]$, either. Therefore, $\beta_4[k]/\beta_4[k - 1] = \beta_3[k]/\beta_3[k - 1] = (x, h_0, a)$ for some a (given by DS_2). $\beta_4[k]$ and $\beta_3[k]$ agree.
 2. if (x, h_0) is evaluated in both $\beta_3[k - 1]$ and $\beta_4[k - 1]$. $\beta_4[k] = \beta_4[k - 1]$ agrees with $\beta_3[k] = \beta_3[k - 1]$.
 3. if (x, h_0) is evaluated in $\beta_3[k - 1]$ but not in $\beta_4[k - 1]$. We have two subcases here. First, when (x, h_0) is free in $\beta_3[k - 1]$: In this case it will be free in $\beta_4[k]$ and hence $\beta_4[k]$ agrees with $\beta_3[k]$. Second, (x, h_0) is programmed in $\beta_3[k - 1]$: We show this is actually impossible. Notice that no constellation y for which $\tilde{g}(y) = x$ has been queried by \mathcal{D} in Game 4 since (x, h_0) is not in $\beta_4[k - 1]$. Since $\beta_4[k - 1]$ agrees with $\beta_3[k - 1]$, no such constellation has been queried in Game 3.2 either. Therefore, in this case the event Pred^{\leftarrow} would have occurred.
- If \mathcal{D} queries a constellation x : In both games, this leads to completion of the constellation of x , including evaluation of the last term $h_0(\tilde{g}_R(x))$.
 1. We first analyze the behaviour of the evaluations prior to the last term. We need to show the sequence of evaluations in Game 3.2 coincides with that in Game 4. No possible disagreement can arise on any term that is free in both games, as these are both drawn from DS_1 ; thus any disparity must occur at a term programmed in at least one game. If a term is programmed in only in Game 3.2, the event Pred^{\leftarrow} occurs, contrary to assumption. Observe that if a term is programmed in Game 4 it must also be in Game 3.2, as $\beta_4[k - 1] \subset \beta_3[k - 1]$ from Lemma 10 and these agree. Finally, if a term is programmed in both games they are given the same value since $\beta_4[k - 1]$ agrees with $\beta_3[k - 1]$.

2. Now we proceed to the last term $h_0(\tilde{g}(x))$. If it is not in $\beta_3[k-1]$, then it is not in $\beta_4[k-1]$ by Lemma 10. Thus, the value of the term is either programmed in both games to be consistent with F or was in fact freely assigned in both games during completion of the constellation itself. In either case, the values in two games are equal. If the last term is in $\beta_3[k-1]$, the event Pred^\rightarrow occurs.

In summary, $\beta_3[k]$ agrees with $\beta_4[k]$. \square

Lemma 12. $\|\alpha_3 - \alpha_4\|_{\text{TV}} \leq \Pr[\text{Pred}^\leftarrow] + \Pr[\text{Pred}^\rightarrow] < 2 \Pr[\text{Pred}]$, where Pred^\leftarrow and Pred^\rightarrow are events in Game 3 and Pred is in Game 2.

Proof. We view the two pairs of $(\mathcal{A}, \mathcal{B})$ in Game 2.3 and Game 3.1 as two probabilistic Turing machines sharing the same randomness r for any long enough string r . Given the fixed distinguisher \mathcal{D} from Lemma 6 to 11 and r above, Pred occurs in Game 2.1 if Pred^\leftarrow or Pred^\rightarrow occurs in Game 3.2. Since r and \mathcal{D} are arbitrary, we have $\Pr[\text{Pred}^\leftarrow] < \Pr[\text{Pred}]$ and $\Pr[\text{Pred}^\rightarrow] < \Pr[\text{Pred}]$, which implies the lemma. \square

Theorem 4 follows by applying the triangle inequality to Lemma 6, Lemma 7, and Lemma 12.

4.3 Establishing pointwise unpredictability and subversion-freedom

We first prove that the events Pred and Subv are negligible. Throughout we let ϵ denote the equality guarantee of the subversion algorithm \tilde{H} : (for all i) $h_i(x) \neq \tilde{h}_i(x)$ with probability no more than ϵ (in uniform choice of x).

Definition 8 (Good term). Fix a subversion algorithm \tilde{H} . For any $i \in [\ell]$ and $x \in \{0, 1\}^n$, we say (x, h_i) is good if $\Pr_{h_*}[h_i(x) \neq \tilde{h}_i(x)] < \sqrt{\epsilon}$.

Lemma 13. For any \tilde{H} , any $i \in [\ell]$, and uniform x , $\Pr_x[(x, h_i) \text{ is good}] > 1 - \sqrt{\epsilon}$.

Proof. Notice that $\mathbb{E}_x[\Pr_{h_*}[h_i(x) \neq \tilde{h}_i(x)]] = \Pr_{x, h_*}[h_i(x) \neq \tilde{h}_i(x)] < \epsilon$. The lemma therefore follows directly from Markov's inequality. \square

In several circumstances, it is convenient to determine a hash function Rh_* by “resampling” a particular entry of a hash function h_* . To make this precise, for a hash function $H : \{0, 1\}^{3n} \rightarrow \{0, 1\}^{3n}$, we define $R[z; s]H$ to be the function given by the rule

$$R[z; s]H(x) = \begin{cases} H(x) & \text{if } x \neq z, \\ s & \text{if } x = z, \end{cases}$$

We typically apply this when s is drawn uniformly at random in $\{0, 1\}^{3n}$; we then say that the value of H has been “resampled” at z . In this case, $R[z; s]H$ is a random variable which we write $R[z]H$ (obtained by selecting s uniformly); when z is explicitly identified by context, we simply write RH . Finally, we apply this notation consistently with our conventions for rendering the family h_* from H . Thus $R[i, x]h_*$ denotes the family of functions resulting from resampling (x, h_i) .

Fact 14 (Invariance under resampling). Consider an (unbounded) algorithm A with oracle access to $H : \{0, 1\}^m \rightarrow \{0, 1\}^m$, a uniformly random function. Suppose, further, that A^H carries out a collection of (adaptively chosen) queries to H and with probability 1 returns an element $\text{result}(A^H) \in \{0, 1\}^m$ on which it has not queried H . Then the random variable $R[\text{result}(A^H)]H$ is uniform; that is, the distribution arising from sampling H , querying according to A , and resampling H at $\text{result}(A^H)$ is uniform.

Proof. Observe that the resampling does not change the distribution of H conditioned on the responses the queries of A . \square

Definition 9 (Honest term). Fix a subversion algorithm \tilde{H} . Then for a fixed h_* , an index $i \in [\ell]$, and an $x \in \{0, 1\}^n$, we say (x, h_i) is honest if for any $j \in [\ell]$ and $y \in \{0, 1\}^n$,

$$\Pr_s[Rh_i(x) \neq \widetilde{Rh}_i(x)] < \epsilon^{1/4},$$

where the randomness is over the resampling of $h_j(y)$ (so the resampling operator is $R[j, y; s]$ in both cases). A term is dishonest if it is not honest. (N.b. the function(s) \widetilde{Rh}_i are defined by \tilde{H}^{Rh_*} and hence may be very different from h_{*} .)

Let us consider the following random variables defined by (uniform) random selection of h_* :

$$d_i(x) = \begin{cases} 1 & \text{if } (x, h_i) \text{ is dishonest,} \\ 0 & \text{otherwise.} \end{cases}$$

Lemma 15. $\Pr_{h_*}[\exists i \in [\ell], h_i \text{ has more than } 2^n(q_{\bar{H}} + 1)\epsilon^{1/8} \text{ dishonest inputs}] \leq \ell\epsilon^{1/8}$.

Proof. If (x, h_i) is good, then

$$\begin{aligned} \mathbb{E}[d_i(x)] &= \Pr_{h_*}[\exists(j, y), \Pr_s[\mathbf{R}[j, y; s]h_i(x) \neq \widetilde{\mathbf{R}}[j, y; s]h_i(x)] > \epsilon^{1/4}] \\ &\leq \sum_{k=1, \dots, q_{\bar{H}}} \Pr_{h_*}[\Pr_s[\mathbf{R}[j, y; s]h_i(x) \neq \widetilde{\mathbf{R}}[j, y; s]h_i(x)] > \epsilon^{1/4}, \\ &\quad \text{where } h_j(y) \text{ is the } k\text{-th term queried by the evaluation of } \tilde{h}_i(x)] \\ &\leq q_{\bar{H}} \cdot \Pr_{h_*}[h_i(x) \neq \tilde{h}_i(x)] \leq q_{\bar{H}}\epsilon^{1/4} \end{aligned}$$

by the union bound and invariance under resampling. Therefore,

$$\sum_{x \in \{0,1\}^n} \mathbb{E}[d_i(x)] < 2^n(q_{\bar{H}}\epsilon^{1/4} + \sqrt{\epsilon}) < 2^n(q_{\bar{H}} + 1)\epsilon^{1/4}.$$

Applying Markov's inequality and union bound again, we conclude

$$\Pr_{h_*}[\exists i \in [\ell], h_i \text{ has more than } 2^n(q_{\bar{H}} + 1)\epsilon^{1/8} \text{ dishonest inputs}] \leq \ell\epsilon^{1/8}. \quad \square$$

Definition 10 (Invisible term). *Given (R, h_*) , we say a term (x, h_i) is invisible if for any $\tilde{h}_j(y)$ that queries (x, h_i) , $\tilde{h}_j(y) = h_j(y)$ and*

$$\Pr_s[\mathbf{R}[i, x; s]h_j(y) \neq \widetilde{\mathbf{R}}[i, x; s]h_j(y)] < \epsilon^{1/4}.$$

Lemma 16. $\Pr[\text{There exists a constellation with fewer than } \ell - n \text{ invisible terms}] = O(\ell\epsilon^{1/8})$.

Proof. We say h_* is stealthy if, for all $i \in [\ell]$, there are no more than $\ell 2^n(q_{\bar{H}} + 1)\epsilon^{1/8}$ dishonest terms in h_i . If h_* is stealthy, the dishonest terms in h_* make at most $\ell 2^n q_{\bar{H}}(q_{\bar{H}} + 1)\epsilon^{1/8}$ queries via the subversion algorithm. In other words, each column contains at least $2^n - \ell 2^n q_{\bar{H}}(q_{\bar{H}} + 1)\epsilon^{1/8}$ terms that are not queried by any dishonest term.

Notice that a term that is only queried by honest terms is invisible. Therefore,

$$\begin{aligned} &\Pr_{R, h_*}[\text{There exists a constellation with more than } n \text{ visible terms.}] \\ &< \Pr_{R, h_*}[\text{There exists a constellation with more than } n \text{ visible terms} \mid h_* \text{ is stealthy}] + \Pr_{h_*}[h_* \text{ is stealthy}] \\ &< \Pr_{R, h_*} \left[\begin{array}{l} \text{There exists a constellation with more than } n \text{ terms that} \\ \text{are queried by a dishonest or subverted term} \end{array} \mid h_* \text{ is stealthy} \right] + \ell\epsilon^{1/8} \\ &< 2^n \binom{\ell}{n} (\ell q_{\bar{H}}(q_{\bar{H}} + 1)\epsilon^{1/8} + \ell q_{\bar{H}}\epsilon)^n + \ell\epsilon^{1/8} = O(\ell\epsilon^{1/8}), \end{aligned}$$

as long as $\ell = O(n)$. In the last equality, we use the fact that $q_{\bar{H}}$ is polynomial in n , and for sufficiently large n ,

$$2^n \binom{\ell}{n} (\ell q_{\bar{H}}(q_{\bar{H}} + 1)\epsilon^{1/8} + \ell q_{\bar{H}}\epsilon)^n = O(2^n 4^n \epsilon^{n/16}) = O((8\epsilon^{1/16})^n) = O(\epsilon^{1/8}). \quad \square$$

Definition 11 (Normal randomness). *We say (R, h_*) is normal if every constellation has at least $\ell - n$ invisible terms.*

For convenience, we focus on a fixed distinguisher in Game 2 for the rest of the section.

Definition 12 (Regular transcript). For any $0 < k < q_{\hat{\mathcal{D}}}$, we say an R -transcript $\alpha_2[k]$ is regular if

$$\Pr_{h_*}[(R, h_*) \text{ is normal} \mid \alpha_2[k]] > 1 - \sqrt{\ell}\epsilon^{1/16}.$$

Otherwise, we say $\alpha_2[k]$ is irregular.

Lemma 17. For any $0 < k < q_{\hat{\mathcal{D}}}$.

$$\Pr[\alpha_2[k] \text{ is irregular}] < \sqrt{\ell}\epsilon^{1/16}.$$

Proof.

$$\sum_{\alpha_2[k]} \Pr[(R, h_*) \text{ is normal} \mid \alpha_2[k]] \cdot \Pr[\alpha_2[k]] = \Pr[(R, h_*) \text{ is normal}] = 1 - O(\ell\epsilon^{1/8}).$$

Using Markov's inequality, we have, with probability at least $1 - \sqrt{\ell}\epsilon^{1/16}$ in the choice of $\alpha_2[k]$,

$$\Pr[(R, h_*) \text{ is not normal} \mid \alpha_2[k]] < \sqrt{\ell}\epsilon^{1/16}. \quad \square$$

Definition 13 (Consistency). For any $0 < k < q_{\hat{\mathcal{D}}}$ and any two functions h_*^1, h_*^2 , we say h_*^1 is k -consistent with h_*^2 if for some R in Game 2, $h_* = h_*^1$ and $h_* = h_*^2$ give the same R -transcript $\alpha_2[k]$ at the end of k -th interaction.

Lemma 18. For any $0 < k < q_{\hat{\mathcal{D}}}$ and two k -consistent functions h_*^1, h_*^2 with common R -transcript α' after k -th interaction,

$$\Pr[h_* = h_*^1 \mid \alpha_2[k] = \alpha'] = \Pr[h_* = h_*^2 \mid \alpha_2[k] = \alpha'],$$

where the randomness is over Game 2.

Proof. Observe that, for a fixed distinguisher, the transcript $\alpha_2[k]$ is a deterministic function of h . As h is uniformly distributed, the distribution obtained by conditioning on the value of any fixed function of h is also uniform. \square

4.3.1 Unpredictability

In this section we bound the probability of Pred.

For $0 \leq k < q_{\hat{\mathcal{D}}}$, we define the event $\text{Pred}[k]$ to be: in Game 2.2, Pred occurs before the end of k -th query and answer. In the rest of section 4.4, we denote by $T[k]$ the data in T_H and T_F in Game 2.2 after k -th query and answer.

Corollary 19. For any $0 < k < q_{\hat{\mathcal{D}}}$, $\Pr[\text{Pred}[k] \wedge \neg \text{Pred}[k-1]] = O(\ell^2(k-1)q_{\hat{H}}2^{-3n} + \sqrt{\ell}\epsilon^{1/16})$.

Proof. We consider the event $\Pr[\text{Pred}[k] \wedge \neg \text{Pred}[k-1]]$.

$$\begin{aligned} & \Pr[\text{Pred}[k] \wedge \neg \text{Pred}[k-1]] \\ & \leq \Pr[\text{Pred}[k] \wedge \neg \text{Pred}[k-1] \mid \alpha_2[k-1] \text{ is regular}] + \Pr[\alpha_2[k-1] \text{ is not regular}] \\ & \leq \Pr[\text{Pred}[k] \wedge \neg \text{Pred}[k-1] \text{ and } (R, h_*) \text{ is normal} \mid \alpha_2[k-1] \text{ is regular}] \\ & \quad + \Pr[(R, h_*) \text{ is not normal} \mid \alpha_2[k-1] \text{ is regular}] + \sqrt{\ell}\epsilon^{1/16} \\ & \leq \Pr \left[\begin{array}{l} \text{the } k\text{-th query is a constellation } y \text{ that contains} \\ \text{an invisible term and } \tilde{g}_R(y) \in T[k-1] \end{array} \mid \alpha_2[k-1] \text{ is regular} \right] + 2\sqrt{\ell}\epsilon^{1/16} \\ & \leq \sum_{i \in [\ell]} \Pr \left[\begin{array}{l} \text{the } k\text{-th query is a constellation } y, \tilde{g}_R(y) \in T[k-1] \\ \text{and the } i\text{th term } (y, h_i) \text{ is invisible} \end{array} \mid \alpha_2[k-1] \text{ is regular} \right] + 2\sqrt{\ell}\epsilon^{1/16} \end{aligned}$$

Now consider any fixed $\alpha_2[k-1]$ in conjunction with a fixed index i . Recall that fixing $\alpha_2[k-1]$ determines R and hence the constellation determined by the k th query; thus i uniquely determines a term (y, h_i) of the constellation y queried by \mathcal{D} at the k th step. Consider further a ‘‘partial assignment’’ p_* to the functions h_* that provides values at all points in the domain except for (y, h_i) . We say that p_* has an *invisible completion* (w.r.t. $\alpha_2[k-1]$) if there is some

assignment to (y, h_i) so that this term is invisible (for p_* and the R given by $\alpha[k-1]$). With this language, we may further upper bound the expression just above by

$$\leq \sum_{i \in [\ell]} \Pr \left[\begin{array}{l} \tilde{g}_R(y) \in T[k-1] \text{ and the } i\text{th} \\ \text{term } (y, h_i) \text{ is invisible} \end{array} \middle| \begin{array}{l} \alpha_2[k-1] \text{ is regular, the } k\text{-th query is a constellation} \\ y \text{ with } i\text{th term } (y, h_i), \text{ the partial assignment } p_* \\ \text{excluding } (y, h_i) \text{ has an invisible completion} \end{array} \right] + 2\sqrt{\ell}\epsilon^{1/16}. \quad (2)$$

(Here, the partial assignment p_* is simply the function h_* with the value at (y, h_i) excluded.) Observe that for any fixed triple $(\alpha_2[k-1], i, p_*)$ that may arise in the conditioning above, there are exactly 2^{3n} completions of p_* to a full function h_* . Recalling the definition of invisible, it follows that at least a $1 - \epsilon^{1/4}$ fraction of all completions are consistent with $\alpha_2[k-1]$ and, furthermore, maintain the invisibility of (y, h_i) . Note, additionally, that at most $|T[k-1]|$ of these completions can have the property that $\tilde{g}_R(y) \in T[k-1]$ and, additionally, that all such completions have the same conditional probability (Lemma 18). We conclude that the probability of (2) is no more than

$$\leq \frac{\ell}{1 - \epsilon^{1/4}} \frac{|T[k-1]|}{2^{3n}} + 2\sqrt{\ell}\epsilon^{1/16} = O(\ell^2(k-1)q_{\tilde{H}}2^{-3n} + \sqrt{\ell}\epsilon^{1/16}). \quad \square$$

Theorem 20. $\Pr[\text{Pred}] = O(\ell^2 q_{\tilde{D}}^2 q_{\tilde{H}} 2^{-3n} + \sqrt{\ell} q_{\tilde{D}} \epsilon^{1/16})$.

Theorem 20 follows from the union bound to Corollary 19 over k .

4.3.2 Subversion freedom

Now we bound the probability of the event Subv.

Definition 14 (Silent term). *Given (R, h_*) , for any $x \in \{0, 1\}^n$ and $i \in [\ell]$, we say (x, h_i) is silent if (x, h_i) is invisible and is queried by fewer than $2^{2.5n} q_{\tilde{H}}$ terms of \tilde{h}_0 .*

Lemma 21. *For any $0 < k < q_{\tilde{D}}$, assuming that $\ell > n + 4$,*

$$\Pr[\text{there exists } x \text{ such that the constellation } x \text{ has no silent term} \mid \alpha_2[k] \text{ is regular}] = O\left(\frac{\ell^4}{2^n} + \sqrt{\ell}\epsilon^{1/16}\right).$$

Proof. Using Markov's inequality, for each $x \in \{0, 1\}^n$ and $i \in [\ell]$,

$$\Pr_{R, h_*} [h_i(x \oplus r_i) \text{ is queried by } \tilde{h}_0(y) \text{ for more than } 2^{2.5n} q_{\tilde{H}} \text{ many } y\text{'s}] < \frac{1}{\sqrt{2^n}}.$$

Let us call such a $(x \oplus r_i, h_i)$ *thick* if it is queried by so many $\tilde{h}_0(y)$ and *thin* otherwise. (So that a term is silent if it is invisible and thin.) Note that for any fixed constellation x , the events that $(x \oplus r_i, h_i)$ are thin are independent (as they are determined by different r_i). Since $\ell > n + 4$,

$$\begin{aligned} & \Pr[\text{there exists } x \text{ such that the constellation } x \text{ has no silent term} \mid \alpha_2[k] \text{ is regular}] \\ & \leq \Pr_{R, h_*} \left[\begin{array}{l} \text{there exists } x \text{ such that the constellation } x \text{ has no silent} \\ \text{term and has more than } \ell - n \text{ invisible terms} \end{array} \middle| \alpha_2[k] \text{ is regular} \right] \\ & \quad + \Pr_{R, h_*} [\text{there exists a constellation } x \text{ that has more than } n \text{ visible terms} \mid \alpha_2[k] \text{ is regular}] \\ & \leq \Pr_{R, h_*} \left[\begin{array}{l} \text{there exists } x \text{ such that the constellation } x \text{ has more than} \\ \ell - n \text{ invisible terms and all these terms are thick} \end{array} \right] / \Pr[\alpha_2[k] \text{ is regular}] + O(\sqrt{\ell}\epsilon^{1/16}) \\ & \leq \Pr_{R, h_*} \left[\begin{array}{l} \text{there exists } x \text{ such that the constellation } x \text{ has more than} \\ \ell - n \text{ thick terms} \end{array} \right] / \Pr[\alpha_2[k] \text{ is regular}] + O(\sqrt{\ell}\epsilon^{1/16}) \\ & \leq 2^n \binom{\ell}{4} \left(\frac{1}{\sqrt{2^n}}\right)^4 \cdot 2 + O(\sqrt{\ell}\epsilon^{1/16}) = O(\ell^4/2^n + \sqrt{\ell}\epsilon^{1/16}). \quad \square \end{aligned}$$

For $0 \leq k < q_{\tilde{D}}$, we define the event Subv $[k]$ to be: in Game 2.2, Subv occurs before the end of k -th query and answer.

Corollary 22. For any $0 < k < q_{\hat{D}}$, $\Pr[\text{Subv}[k] \wedge \neg \text{Subv}[k-1]] = O(\ell q_{\hat{H}}/\sqrt{2^n} + \sqrt{\ell} \epsilon^{1/16})$.

Proof. We consider the event $\text{Subv}[k] \wedge \neg \text{Subv}[k-1]$.

$$\begin{aligned}
& \Pr[\text{Subv}[k] \wedge \neg \text{Subv}[k-1]] \\
& \leq \Pr[\text{Subv}[k] \wedge \neg \text{Subv}[k-1] \mid \alpha_2[k-1] \text{ is regular}] + \Pr[\alpha_2[k-1] \text{ is not regular}] \\
& \leq \Pr[\text{Subv}[k] \wedge \neg \text{Subv}[k-1] \text{ and every constellation has a silent term} \mid \alpha_2[k-1] \text{ is regular}] \\
& \quad + \Pr[\text{not all constellations have a silent term} \mid \alpha_2[k-1] \text{ is regular}] + \sqrt{\ell} \epsilon^{1/16} \\
& \leq \sum_{i \in [\ell]} \Pr \left[\begin{array}{l} \text{the } i\text{th term } (x, h_i) \text{ of the constellation } x \text{ queried} \\ \text{at step } k \text{ is silent and } h_0 \text{ is subverted at } \tilde{g}_R(x) \end{array} \mid \alpha_2[k-1] \text{ is regular} \right] + \ell^4/2^n + \sqrt{\ell} \epsilon^{1/16}
\end{aligned}$$

Now consider any fixed $\alpha_2[k-1]$ in conjunction with a fixed index i . Recall that fixing $\alpha_2[k-1]$ determines R and hence the constellation determined by the k th query; thus i uniquely determines a term (y, h_i) of the constellation y queried by \mathcal{D} at the k th step. Consider further a ‘‘partial assignment’’ p_* to the functions h_* that provides values at all points in the domain except for (y, h_i) . We say that p_* has a *silent completion* (w.r.t. $\alpha_2[k-1]$) if there is some assignment to (y, h_i) so that this term is silent (for p_* and the R given by $\alpha[k-1]$). With this language, we may further upper bound the expression just above by

$$\begin{aligned}
& \leq \sum_{i \in [\ell]} \Pr \left[\begin{array}{l} (x, h_i) \text{ is silent and } h_0 \\ \text{is subverted at } \tilde{g}_R(x) \end{array} \mid \begin{array}{l} \alpha_2[k-1] \text{ is regular, the partial assignment ex-} \\ \text{cluding the } i\text{th term } (x, h_i) \text{ of the constellation } x \\ \text{queried at step } k \text{ has a silent completion} \end{array} \right] + \ell^4/2^n + \sqrt{\ell} \epsilon^{1/16} \\
& \leq \sum_{i \in [\ell]} \Pr \left[\begin{array}{l} h_0 \text{ is subverted at } \tilde{g}_R(x) \text{ and} \\ \tilde{h}_0(\tilde{g}_R(x)) \text{ does not query } (x, h_i) \end{array} \mid \begin{array}{l} \alpha_2[k-1] \text{ is regular, the partial assignment ex-} \\ \text{cluding the } i\text{th term } (x, h_i) \text{ of the constellation } x \\ \text{queried at step } k \text{ has a silent completion} \end{array} \right] \\
& \quad + \sum_{i \in [\ell]} \Pr \left[\begin{array}{l} h_0 \text{ is subverted at } \tilde{g}_R(x) \text{ and} \\ \tilde{h}_0(\tilde{g}_R(x)) \text{ queries } (x, h_i) \end{array} \mid \begin{array}{l} \alpha_2[k-1] \text{ is regular, the partial assignment ex-} \\ \text{cluding the } i\text{th term } (x, h_i) \text{ of the constellation } x \\ \text{queried at step } k \text{ has a silent completion} \end{array} \right] + \ell^4/2^n + \sqrt{\ell} \epsilon^{1/16}.
\end{aligned}$$

(Here, the partial assignment p_* is simply the function h_* with the value at (y, h_i) excluded.) Observe that for any fixed triple $(\alpha_2[k-1], i, p_*)$ that may arise in the conditioning above, there are exactly 2^{3n} completions of p_* to a full function h_* . Recalling the definition of silent, it follows that at least a $1 - \epsilon^{1/4}$ fraction of all completions are consistent with $\alpha_2[k-1]$ and, furthermore, maintain the silence of (y, h_i) . Note, additionally, that at most $2^{3n}\epsilon$ of these completions can have the property that h_0 is subverted at $\tilde{g}_R(x)$ and $\tilde{h}_0(\tilde{g}_R(x))$ does not query (x, h_i) , while at most $2^{2.5n} q_{\hat{H}}$ of these completions can have the property that h_0 is subverted at $\tilde{g}_R(x)$ and $\tilde{h}_0(\tilde{g}_R(x))$ queries (x, h_i) . Moreover, all such completions have the same conditional probability (Lemma 18). We conclude that the probability above is no more than

$$\leq \ell \left(\frac{2^{3n}\epsilon}{2^{3n}} + \frac{2^{2.5n} q_{\hat{H}}}{2^{3n}} \right) / (1 - \epsilon^{1/4}) + \ell^4/2^n + \sqrt{\ell} \epsilon^{1/16} = O(\ell q_{\hat{H}}/\sqrt{2^n} + \sqrt{\ell} \epsilon^{1/16}). \quad \square$$

Theorem 23. $\Pr[\text{Subv}] = O(\ell q_{\hat{D}} q_{\hat{H}}/\sqrt{2^n} + \sqrt{\ell} q_{\hat{D}} \epsilon^{1/16})$.

Theorem 23 can be proved by applying the union bound to Corollary 22 over k .

4.4 Controlling the self-referential probability

Now we proceed to prove that the crisis event *Selfref* is negligible. We prove this result by establishing a stronger theorem that says the entire data table in Game 2.1 has a negligible chance to have a self-referential constellation.

Theorem 24. In Game 2.1, for any distinguisher \mathcal{D}

$$\Pr_{h,R}[\text{Selfref}] = O(q_{\hat{H}} \ell^{2-2n}).$$

We state a standard Chernoff bound; see [MR95], or [Lev] for this particular formulation.

Lemma 25. (Chernoff Bound) Let X_1, \dots, X_n be discrete, independent random variable such that $E[X_i]=0$ and $|X_i| \leq 1$ for all i . Let $X = \sum_{i=1}^n X_i$ and σ^2 be the variance of X . Then $\Pr[|X| \geq \lambda\sigma] \leq 2e^{-\lambda^2/4}$ for any $0 \leq \lambda \leq 2\sigma$.

In the next two lemmas, we introduce the Fourier transformation to analyze the total variation distance between the distribution of $\tilde{g}_R(x)$ and the uniform distribution.

Definition 15 (Character and Dual group). Suppose G is a finite abelian group. A character χ of G is a group homomorphism from G into the multiplicative group T of complex numbers of norm 1. We define the dual group \hat{G} to be the set of all characters of G ; these make an abelian group under the pointwise product operation $[\chi\sigma](x) = \chi(x) \cdot \sigma(x)$. We let $\mathbf{1} : G \rightarrow T$ denote the trivial character $\mathbf{1} : g \mapsto 1$.

It is a fact that $|G| = |\hat{G}|$ for any finite abelian group G and, furthermore, that the functions of \hat{G} are linearly independent: that is, if $\sum_{\chi} a_{\chi}\chi$ is the zero function for some collection of coefficients $a_{\chi} \in \mathbb{C}$, then all $a_{\chi} = 0$.

Definition 16 (Discrete Fourier transform). Suppose G is a finite abelian group. Let $L^2(G)$ be the set of all complex-valued functions on G . The discrete Fourier transform on $f \in L^2(G)$ is defined by

$$\hat{f}(\chi) = \sum_{a \in G} f(a)\overline{\chi(a)} = \langle f, \chi \rangle, \text{ for } \chi \in \hat{G}.$$

Here $\langle \cdot, \cdot \rangle$ denotes the inner product of complex-valued functions on G . We remark that if $f : G \rightarrow \mathbb{R}$ is a probability distribution then $\hat{f}(\mathbf{1}) = 1$ and $|\hat{f}(\chi)| \leq 1$ for all χ

Lemma 26. Define convolution by

$$(f * g)(x) = \sum_{y \in G} f(y)g(x-y), \text{ for } x \in G.$$

Then

$$\widehat{f * g}(\chi) = \hat{f}(\chi) \cdot \hat{g}(\chi), \text{ for all } \chi \in \hat{G}.$$

Lemma 27. Let Q be a probability distribution and U be the uniform distribution on a finite abelian group G . Then,

$$\|Q - U\|_{\text{tv}} \leq \frac{1}{2} \left(\sum_{\chi \in \hat{G}, \chi \neq \mathbf{1}} |\hat{Q}(\chi)|^2 \right)^{1/2}.$$

In our application of the Fourier transform, G is the group $(\mathbb{Z}/2)^{3n}$. For any fixed $i \in [\ell]$ and $x \in \{0, 1\}^n$, we denote the distribution of $h_i(x \oplus r_i)$, $\tilde{h}_i(x \oplus r_i)$ and $\tilde{g}_R(x)$ by p_i^x , $p'_i{}^x$ and P^x , respectively. These are random variables defined by uniform selection h_* and R . For a fixed value of h_* , we further define p_{i,h_*}^x , $p'_{i,h_*}{}^x$, and $P_{h_*}^x$ to be the distributions of $h_i(x \oplus r_i)$, $\tilde{h}_i(x \oplus r_i)$ and $\tilde{g}_R(x)$ respectively, over the randomness of R . In most cases below we omit x in the notation when there is no ambiguity.

Lemma 28. For any $x, r \in \{0, 1\}^n$ and any $t \in [\ell]$,

$$\Pr_{h_*} \left[\|P_{h_*,r,t}^x - U\|_{\text{tv}} \geq 2^{3n-1}(n2^{-n/2} + \epsilon)^{\ell/2} \right] \leq 2^{-\Omega(n^2\ell)},$$

where $P_{h_*,r,t}^x$ is the distribution of $\tilde{g}_R(x)$, conditioned on h_* , over uniform R with the constraint $r_t = r$. U is the uniform distribution on G and ϵ here is the (negligible) disagreement probability of (1) above.

Proof. Consider an arbitrary $i \in [\ell]$ and nontrivial $\chi \in \hat{G}$ (that is, $\chi \neq \mathbf{1}$). Observe that for any h_* ,

$$\hat{p}_{i,h_*} = \sum_{v \in (\mathbb{Z}/2)^{3n}} \Pr_{r_i} [h_i(x \oplus r_i) = v] \overline{\chi(v)} = \frac{1}{2^n} \sum_{r_i} \chi(h_i(x \oplus r_i)),$$

as $\chi()$ is always real. When h_i is selected uniformly, the $\chi(h_i(x \oplus r_i))$ are i.i.d. random variables taking values in $\{\pm 1\}$ and we may apply Lemma 25 with $\lambda = n$; noting that the variance is $\sigma^2 = 2^n$ this yields

$$\Pr_{h_*} [2^n \hat{p}_{i,h_*}(\chi) \geq n2^{n/2}] \leq 2e^{-n^2/4}.$$

For any fixed h_i and x , by assumption $\Pr_{r_i} [\tilde{h}_i(x \oplus r_i) \neq h_i(x \oplus r_i)] \leq \epsilon$ and hence $|\hat{p}_{i,h_*}(\chi) - \hat{p}'_{i,h_*}(\chi)| \leq \epsilon$ for every χ . In light of this, for any t and χ ,

$$\begin{aligned} \Pr_{h_*} [|\hat{P}_{h_*,t,r}(\chi)| \geq (n2^{-n/2} + \epsilon)^{\ell/2}] &= \Pr_{h_*} \left[\prod_{i=1, i \neq t}^{\ell} |\hat{p}'_{i,h_*}(\chi)| \geq (n2^{-n/2} + \epsilon)^{\ell/2} \right] \\ &\leq \Pr_{h_*} [\exists \text{ half of the coordinates } i \in [\ell] \setminus \{r_t\} \text{ such that } |\hat{p}'_{i,h_*}(\chi)| \geq n2^{-n/2} + \epsilon] \\ &\leq \Pr_{h_*} [\exists \text{ half of the coordinates } i \in [\ell] \setminus \{r_t\} \text{ such that } |\hat{p}_{i,h_*}(\chi)| \geq n2^{-n/2}] \\ &\leq \binom{\ell}{\ell/2} \cdot (2e^{-n^2/4})^{\ell/2} = (8e^{-n^2/4})^{\ell/2}. \end{aligned}$$

Thus,

$$\Pr_{h_*} [\exists \chi \in \hat{G} \setminus \{\mathbf{1}\}, |\hat{P}_{h_*,r,t}(\chi)| \geq (n2^{-n/2} + \epsilon)^{\ell/2}] \leq 2^{6n} (8e^{-n^2/4})^{\ell/2}$$

and we conclude that

$$\begin{aligned} \Pr_{h_*} \left[\|P_{h_*,r,t}^x - U\|_{\text{tv}} \geq \frac{1}{2} 2^{3n} (n2^{-n/2} + \epsilon)^{\ell/2} \right] &= \Pr_{h_*} \left[\frac{1}{2} \sum_{\chi \in \hat{G}, \chi \neq \mathbf{1}} |\hat{P}_{h_*,r,t}(\chi)| \geq 2^{3n-1} (n2^{-n/2} + \epsilon)^{\ell/2} \right] \\ &\leq \Pr_{h_*} [\exists \chi \in \hat{G}, \chi \neq \mathbf{1}, |\hat{P}_{h_*,r,t}(\chi)| \geq (n2^{-n/2} + \epsilon)^{\ell/2}] \\ &\leq 2^{6n} (8e^{-n^2/4})^{\ell/2} = 2^{-\Omega(n^2\ell)}, \end{aligned}$$

where the first equality is Lemma 27. \square

Now we are ready to show that, with overwhelmingly probability, there is no self-referential constellation in the entire table.

Lemma 29. *In Game 2.1, for some R , h_* and $x \in \{0, 1\}^n$, we say constellation x is self-referential if there exists $t \in [\ell]$ such that $h_0(\tilde{g}_R(x))$ is queried by $\tilde{h}_t(x \oplus r_t)$. Then, assuming that $\ell > n$,*

$$\Pr_{R, h_*} [\text{there exists a self-referential constellation}] = O(q_{\tilde{H}} \ell 2^{-2n}).$$

Proof. For any $x, r \in \{0, 1\}^n$ and any $t \in [\ell]$, we say h_* is (x, r, t) -good if $\|P_{h_*,r,t}^x - U\|_{\text{tv}} \leq 2^{3n-1} (n2^{-n/2} + \epsilon)^{\ell/2}$. Otherwise, we say h_* is (x, r, t) -bad. Let $P_{r,t}^x$ be the distribution of $\tilde{g}_R(x)$ over the randomness of h_* and R with the constraint $r_t = r$. Then,

$$\begin{aligned} &\Pr_{R, h_*} [\exists x \in \{0, 1\}^n, \text{ the constellation of } x \text{ is self-referential}] \\ &\leq \sum_{x \in \{0, 1\}^n} \sum_{t=1}^{\ell} \sum_{r \in \{0, 1\}^n} \Pr[r_t = r] \Pr_{R/r_t, h_*} [h_0(\tilde{g}_R(x)) \text{ is queried by } \tilde{h}_t(x \oplus r_t) \text{ with } r_t = r] \\ &\leq \ell 2^n \max_{x, r, t} \left\{ \Pr_{R/r_t, h_*} [h_0(\tilde{g}_R(x)) \text{ is queried by } \tilde{h}_t(x \oplus r) \mid h_* \text{ is } (x, r, t)\text{-good}] \cdot \Pr_{h_*} [h_* \text{ is } (x, r, t)\text{-good}] \right. \\ &\quad \left. + \Pr_{h_*} [h_* \text{ is } (x, r, t)\text{-bad}] \right\} \\ &\leq \ell 2^n \left\{ (2^{3n-1} (n2^{-n/2} + \epsilon)^{\ell/2} + 2^{-3n} \cdot q_{\tilde{H}}) + O(2^{-\Omega(n^2\ell)}) \right\} \\ &= O(q_{\tilde{H}} \ell 2^{-2n}), \end{aligned}$$

so long as $\ell > n$. (Note that for sufficiently large n , we have $(n2^{-n/2} + \epsilon)^{\ell/2} = O(n^{\ell/2}2^{-\Theta(n\ell)} + \epsilon^{\Theta(\ell)})$.) In the last inequality we use the fact that if $\|D_1 - D_2\| \leq \delta$ for two distributions D_1 and D_2 then, for any event E , $\Pr_{D_1}[E] \leq \Pr_{D_2}[E] + \delta$. \square

Theorem 24 follows immediately from Lemma 29.

4.5 Crooked indifferenciability in the full model

Now we show the simulator \mathcal{S} achieving abbreviated crooked indifferenciability can be lifted to a simulator that achieves full indifferenciability (Definition 4).

Theorem 30. *If the construction in Section 3 is $(n', n, q_{\mathcal{D}}, q_{\tilde{H}}, r, \epsilon')$ -Abbreviated- H -crooked-indifferenciability from a random oracle F , it is $(n', n, q_{\mathcal{D}}, q_{\tilde{H}}, r, \epsilon' + O(q_{\mathcal{D}}^2 q_{\tilde{H}} \ell 2^{-3n}))$ - H -crooked-indifferenciability from F .*

Proof. Consider the following simulator \mathcal{S}_F built on \mathcal{S} :


1. Draw two data sets, DS_1 and DS_2 : DS_1 contains uniformly selected values for each $F(x)$ and $h_i(x)$ for all $0 < i \leq \ell$ and all $x \in \{0, 1\}^n$; DS_2 contains uniformly selected values $h_0(y)$ for all $y \in \{0, 1\}^{3n}$.
2. In the first phase, \mathcal{S}_F answers $h_i(x)$ ($0 < i \leq \ell$) and $h_0(y)$ queries according to DS_1 and DS_2 , respectively.
3. The second phase, after which \mathcal{S}_F receives R , is divided into two sub-phases.
 - First, \mathcal{S}_F simulates \mathcal{S} in Game 4 with data sets DS_1 and DS_2 generated above. It then plays the role of the distinguisher, and asks \mathcal{S} all the questions that were actually asked by the the distinguisher in the first phase. \mathcal{S}_F aborts the game if, in this sub-phase, there are two (simulated) queries (x, h_0) and (y, h_i) ($0 < i \leq \ell$) such that $x = \tilde{g}_R(y \oplus r_i)$.
 - Second, \mathcal{S}_F simulates \mathcal{S} and answers the second-phase questions from the distinguisher.

For an arbitrary full model distinguisher \mathcal{D}_F , we construct the an abbreviated model distinguisher \mathcal{D} as follows. The proof will show that, with high probability, the execution that takes place between \mathcal{D} and \mathcal{S} can be “lifted” to an associated execution between \mathcal{D}_F and \mathcal{S}_F .

1. Prior to the game, \mathcal{D} must publish a subversion algorithm \tilde{H} . This program is constructed as follows. To decide how to subvert a certain term $h_i(x)$, \tilde{H} first simulates the first phase of \mathcal{D}_F ; all queries made by this simulation are asked as regular queries by \tilde{H} and, at the conclusion, this first phase of \mathcal{D}_F produces, as output, a subversion algorithm \tilde{H}_F . \tilde{H} then simulates the algorithm \tilde{H}_F on the term $h_i(x)$.
2. In the game, \mathcal{D} simulates the queries of \mathcal{D}_F in \mathcal{D}_F 's first phase. After that, \mathcal{D} continues to simulate \mathcal{D}_F in the second phase. (Note that at the point in \mathcal{D}_F 's game where it produces the subversion algorithm \tilde{H} , this is simply ignored by \mathcal{D} .)

Now we are ready to prove \mathcal{S}_F is secure against the arbitrarily chosen distinguisher \mathcal{D}_F . We organize the proof around four different transcripts:

γ_{FC}	transcript of \mathcal{C} interacting with \mathcal{D}_F
γ_C	transcript of \mathcal{C} interacting with \mathcal{D}
γ_{FS}	transcript of \mathcal{S}_F interacting with \mathcal{D}_F
γ_S	transcript of \mathcal{S} interacting with \mathcal{D}



(Here \mathcal{C} denotes the construction, as usual.) Since

$$\|\gamma_{FC} - \gamma_{FS}\|_{\text{tv}} \leq \|\gamma_{FC} - \gamma_C\|_{\text{tv}} + \|\gamma_C - \gamma_S\|_{\text{tv}} + \|\gamma_S - \gamma_{FS}\|_{\text{tv}},$$

it is sufficient to prove the three terms in the right-hand side of the inequality are all negligible.

- $\|\gamma_{FC} - \gamma_C\|_{\text{tv}} = 0$. This is obvious by observing that $\gamma_{FC} = \gamma_C$ when the underlying values of H are the same.

- $\|\gamma_C - \gamma_S\|_{\text{tv}} = \epsilon'$. This is true because \mathcal{S} achieves abbreviated crooked indifferentiability.
- $\|\gamma_S - \gamma_{FS}\|_{\text{tv}} = O(q_{\mathcal{D}}^2 q_{\tilde{H}} \ell 2^{-3n})$. To prove this statement, we suppose both the full model game and the abbreviated model game select all randomness *a priori* (as in the descriptions above). Suppose the game between \mathcal{S}_F and \mathcal{D}_F and the game between \mathcal{S} and \mathcal{D} share the same data sets DS_1, DS_2 and R . Notice that the two games have same transcripts unless, \mathcal{S}_F aborts the game in the first sub-phase of the second phase. We denote this bad event by **Conflict**, which by the following lemma 31, is negligible. \square

N.b. While the description of the simulator above calls for all randomness to be generated in advance, it is easy to see that the simulator can in fact be carried out lazily with tables.

To explore the failure event above, we define the following game, named **Exp-Many**, played by an unbounded adversary \mathcal{M} .

EXP-MANY
<ol style="list-style-type: none"> 1. Select two data sets DS_1 and DS_2: DS_1 contains uniformly selected values for each $F(x)$ and $h_i(x)$ for all $0 < i \leq \ell$ and all $x \in \{0, 1\}^n$; DS_2 contains uniformly selected values $h_0(y)$ for all $y \in \{0, 1\}^{3n}$. 2. Given DS_1 and DS_2, \mathcal{M} publishes a subversion algorithm \tilde{H} and a sequence of $p(n) (< q_{\mathcal{D}})$ terms in the form of (y, h_i) ($0 < i \leq \ell$) or (x, h_0). 3. R is selected uniformly. 4. Implement Game 4 with DS_1, DS_2, R, \tilde{H} and the queries prepared above. 5. Output 1 if, in Game 4 of the last step, there exist two queries (x, h_0) and (y, h_i) ($i > 0$) such that $x = \tilde{g}_R(y \oplus r_i)$. Otherwise, output 0.

Lemma 31. For any distinguisher \mathcal{M} , $\Pr[\mathbf{Exp-Many} \text{ outputs } 1] = O(q_{\mathcal{D}}^2 q_{\tilde{H}} \ell 2^{-3n})$.

A quick thought reveals that Lemma 31 implies the negligibility of **Conflict**. Suppose there exists a distinguisher \mathcal{M} in the full model such that **Conflict** is non-negligible. Consider the following distinguisher \mathcal{M}' in **Exp-Many**: In step 2 of **Exp-Many**, simulate \mathcal{M} and publish the queries \mathcal{M} publishes. Obviously, by definition of **Exp-Many**, the probability that **Exp-Many** outputs 1 when against \mathcal{M}' is non-negligible.

Now we proceed to prove Lemma 31. We introduce several concepts that are useful in the proof. Given data sets DS_1, DS_2 and a subversion algorithm \tilde{H} , for any term (x, h_i) with $0 < i \leq \ell$ and $x \in \{0, 1\}^n$, we define its *ideal subversion* $I(x, i)$ to be the subverted value of $h_i(x)$ via the subversion algorithm \tilde{H} , using h_0 values in DS_2 and h_i ($i > 0$) values in DS_1 . The trace of (x, h_i) is defined to be the set

$$\text{Tr}(x, i) = \{y \in \{0, 1\}^{3n} \mid (y, h_0) \text{ is queried in the evaluation of } I(x, i)\}.$$

Given DS_1, DS_2, R and \tilde{H} , for any constellation x , its ideal output is

$$I(x) = \bigoplus_{i=1}^{\ell} I(x \oplus r_i, i), \quad \text{and its trace is defined by} \quad \text{Tr}(x) = \bigcup_{i=1}^{\ell} \text{Tr}(x \oplus r_i, i).$$

In **Exp-Many**, we define the event

$$\text{Crossref} = \{\text{In step 4, a constellation } x \text{ is queried such that } I(x) \neq \tilde{g}_R(x)\}.$$

Lemma 32. For any distinguisher \mathcal{M} in **Exp-Many**, $\Pr[\text{Crossref}] = O(q_{\mathcal{D}}^2 q_{\tilde{H}} \ell 2^{-3n})$.

Proof. Notice that, if **Crossref** occurs, there are two queries (x, h_i) and (y, h_j) ($0 < i, j < \ell$) such that $I(x \oplus r_i) \in \text{Tr}(y \oplus r_j)$. We define the event **Two-cross** to be

In **Exp-Many**, among the $p(n)$ queries made by \mathcal{M} , there are two terms (x, h_i) and (y, h_j) ($0 < i, j < \ell$) such that $I(x \oplus r_i) \in \text{Tr}(y \oplus r_j)$.

It is sufficient to show $\Pr[\mathbf{Two-cross}] = O(q_{\mathcal{D}}^2 q_{\tilde{H}} \ell 2^{-3n})$.

For any two terms (x, h_i) and (y, h_j) ($0 < i, j < \ell$) in the queries of \mathcal{M} , whether $I(x \oplus r_i) \in \text{Tr}(y \oplus r_j)$ only depends on DS_1, DS_2, \tilde{H} and R (i.e., it has nothing to do with the queries chosen by \mathcal{M}).

For any $x, r \in \{0, 1\}^n$ and any $t \in [\ell]$, we say (DS_1, DS_2) is (x, r, t) -good if $\|P_{h_{i,r,t}}^x - U\|_{\text{tv}} \leq 2^{3n-1} (n2^{-n/2} + \epsilon)^{\ell/2}$, where $P_{h_{i,r,t}}^x$ is the distribution of $I(x)$, conditioned on DS_1 and DS_2 , over uniform R with the constraint $r_t = r$, and U is the uniform distribution. Using the same proof in Lemma 28, we can show for any $x \in \{0, 1\}^n$ and $t \in [\ell]$,

$$\Pr_{DS_1, DS_2, r} [(DS_1, DS_2) \text{ is } (x, r, t)\text{-good}] \geq 1 - 2^{-\Omega(n^2 \ell)}.$$

By Markov's inequality and the union bound, with probability smaller than $\ell 2^{-\Omega(n^2 \ell)}$ in the choice of (DS_1, DS_2) , we have

$$\Pr_r [(DS_1, DS_2) \text{ is not } (x, r, t)\text{-good} \mid (DS_1, DS_2)] \leq 2^{-\Omega(n^2 \ell)},$$

for any $x \in \{0, 1\}^n$ and $t \in [\ell]$. We say (DS_1, DS_2) is all-good if it satisfies the inequality above.

By a Fourier transform argument similar to that in Lemma 29, if (DS_1, DS_2) is all-good in **Exp-Many**, then for any \tilde{H} , any $0 < i, j \leq \ell$ and any $x, y \in \{0, 1\}^n$, and $\ell > n$,

$$\Pr_R [I(x \oplus r_i) \in \text{Tr}(y \oplus r_j)] \leq \sum_{k=1}^{\ell} \Pr_R [I(x \oplus r_i) \in \text{Tr}(y \oplus r_j, k)] = O(q_{\tilde{H}} \ell 2^{-3n})$$

Finally,

$$\begin{aligned} & \Pr[\mathbf{Two-cross}] \\ & < \Pr[(DS_1, DS_2) \text{ is all-good}] + \Pr[\mathbf{Two-cross} \mid (DS_1, DS_2) \text{ is not all-good}] \\ & < 2^{-\Omega(n^2 \ell)} + \binom{p(n)}{2} O(q_{\tilde{H}} \ell 2^{-3n}) \\ & = O(q_{\mathcal{D}}^2 q_{\tilde{H}} \ell 2^{-3n}). \end{aligned}$$

□

To complete the proof, we consider an even simpler game that focuses attention on a pair of queries.

EXP-ONE
1. Select two data sets DS_1 and DS_2 : DS_1 contains uniformly selected values for each $F(x)$ and $h_i(x)$ for all $0 < i \leq \ell$ and all $x \in \{0, 1\}^n$; DS_2 contains uniformly selected values $h_0(y)$ for all $y \in \{0, 1\}^{3n}$.
2. Given DS_1 and DS_2 , \mathcal{M}' publishes a subversion algorithm \tilde{H} and a triple $((x, h_i), y)$ for some $0 \leq i \leq \ell$, $x \in \{0, 1\}^n$ and $y \in \{0, 1\}^{3n}$.
3. R is selected uniformly.
4. Implement Game 4 with DS_1, DS_2, R, \tilde{H} prepared above and the query (x, h_i) .
5. Output 1 if in the last step of Game 4, $\tilde{g}_R(x) = y$. Otherwise, output 0.

Lemma 33. For any adversary \mathcal{M} in **Exp-Many**, there exists an adversary \mathcal{M}' in **Exp-One** such that

$$\Pr[\mathbf{Exp-One} \text{ outputs 1 against } \mathcal{M}'] > \frac{1}{p^2(n)} (\Pr[\mathbf{Exp-Many} \text{ outputs 1 against } \mathcal{M}] - \Pr[\mathbf{Crossref}])$$

Proof. For an arbitrary adversary \mathcal{M} in **Exp-Many**, consider the following adversary \mathcal{M}' in **Exp-One**:

- In step 2 of **Exp-One**

1. When given DS_1 and DS_2 in step 2 of **Exp-One**, \mathcal{M}' simulates \mathcal{M} , gets a sequence of polynomial many terms in the form of (x, h_i) , and a subversion algorithm \tilde{H} .
2. Write the sequence as the union of two sets A and B , where A consists of elements of the form (x, h_i) ($i > 0$) and B consists of elements of the form (y, h_0) . Uniformly select an element (x, h_i) from A and an element (y, h_0) from B . Output $((x, h_i), y)$ and the subversion algorithm \tilde{H} .

- In step 4 of **Exp-One**

\mathcal{M}' play Game 4 according to DS_1, DS_2, R and \tilde{H} determined in previous steps. In this game, \mathcal{M}' queries (x, h_i) .

For **Exp-Many** with the adversary \mathcal{M} , if Crossref does not occur, then for any term (x, h_i) ($i > 0$) in step 2, we have $I(x \oplus r_i) = \tilde{g}_R(x \oplus r_i)$. If at the same time **Exp-Many** outputs 1, there must exist (y, h_0) and (x, h_i) ($i > 0$) in step 2 such that $y = \tilde{g}_R(x \oplus r_i) = I(x \oplus r_i)$.

$$\begin{aligned} & \Pr[\mathbf{Exp-One} \text{ outputs } 1 \text{ against } \mathcal{M}'] \\ & > \frac{1}{p^2(n)} \Pr[\text{There must exist } (y, h_0) \text{ and } (x, h_i) \text{ (} i > 0 \text{) in the sequence such that } y = \tilde{g}_R(x \oplus r_i) = I(x \oplus r_i)] \\ & > \frac{1}{p^2(n)} (\Pr[\mathbf{Exp-Many} \text{ outputs } 1] - \Pr[\text{Crossref}]). \end{aligned}$$

□

Lemma 34. $\Pr[\mathbf{Exp-One} \text{ outputs } 1] = O(q_{\tilde{H}} \ell 2^{-3n})$.

Proof. It is sufficient to show $\tilde{g}_R(x \oplus r_i)$ is unpredictable for any (x, h_i) ($i > 0$) in DS_1 .

The unpredictability comes from the fact that R is unknown when the adversary publishes the queries. For any (x, h_i) ($i > 0$) in DS_1 , select and fix an arbitrary r_i . The distribution of $\tilde{g}_R(x \oplus r_i)$ is then same as that of the sum of $\tilde{h}_j(x \oplus r_i \oplus r_j)$ for $j \neq i$, which can be proved to be unpredictable by a Fourier transform similar to that in Lemma 29. □

Proof of Lemma 31. Lemma 31 can be derived by simply summing up Lemma 32, Lemma 33 and Lemma 34. □

Proof of Theorem 3. Theorem 3 can be proved by taking the sum of inequalities in Theorem 20, Theorem 23, Theorem 24 and Theorem 30. □

5 Conclusions and Open Problems

In this paper, we initiate the study of correcting subverted random oracles which are adversarially tampered and disagree with the original random oracle at a negligible fraction of inputs. We give a simple construction that can be proven indifferntiable from a random oracle. Our analysis involves, for a given output produced: identifying a good term of the construction which is both honestly evaluated and independent of other terms for its input; and developing a new machinery of rejection resampling lemma (to assure the existence of this term). Our work provides a general tool to transform a buggy implementation of random oracle into a well-behaved one and directly applies to the kleptographic setting.

There are many interesting problems worth further exploring. Here we only list a few. First, a better construction that may tolerate a larger fraction of errors in the subverted random oracle. Second, develop a parallel theory of self-correcting distributions. Third, consider correcting other subverted ideal objectives such as ideal cipher. Fourth, given our extended model of indifferntiability in the presence of crooked elements, consider stronger indifferntiability models (e.g., with a global random oracle, more robust replacement theorem with more subverted components). Fifth, we may consider other type of constructions such as the Feistel structure or sponge construction that maybe more efficient. Last but not least, can we build connections to (or even a unified theory of) error correction, self-correcting programs, randomness extraction and our problem of correcting subverted ideal objects?

Acknowledgement

Alexander Russell is supported in part by NSF grant CNS #1801487, Qiang Tang is supported in part by NSF grant CNS #1801492, and Hong-Sheng Zhou is supported in part by NSF grant CNS #1801470.

References

- [AAB⁺15] Harold Abelson, Ross J. Anderson, Steven M. Bellovin, Josh Benaloh, Matt Blaze, Whitfield Diffie, John Gilmore, Matthew Green, Susan Landau, Peter G. Neumann, Ronald L. Rivest, Jeffrey I. Schiller, Bruce Schneier, Michael A. Specter, and Daniel J. Weitzner. Keys under doormats. *Commun. ACM*, 58(10):24–26, 2015.
- [AMV15] Giuseppe Ateniese, Bernardo Magri, and Daniele Venturi. Subversion-resilient signature schemes. In Indrajit Ray, Ninghui Li, and Christopher Kruegel, editors, *ACM CCS 2015*, pages 364–375. ACM Press, October 2015.
- [BBCL13] Steven M. Bellovin, Matt Blaze, Sandy Clark, and Susan Landau. Going bright: Wiretapping without weakening communications infrastructure. *IEEE Security & Privacy*, 11(1):62–72, 2013.
- [BBP04] Mihir Bellare, Alexandra Boldyreva, and Adriana Palacio. An uninstantiable random-oracle-model scheme for a hybrid-encryption problem. In Christian Cachin and Jan Camenisch, editors, *EUROCRYPT 2004*, volume 3027 of *LNCS*, pages 171–188. Springer, Heidelberg, May 2004.
- [BCFW09] Alexandra Boldyreva, David Cash, Marc Fischlin, and Bogdan Warinschi. Foundations of non-malleable hash and one-way functions. In Mitsuru Matsui, editor, *ASIACRYPT 2009*, volume 5912 of *LNCS*, pages 524–541. Springer, Heidelberg, December 2009.
- [BF05] Alexandra Boldyreva and Marc Fischlin. Analysis of random oracle instantiation scenarios for OAEP and other practical schemes. In Victor Shoup, editor, *CRYPTO 2005*, volume 3621 of *LNCS*, pages 412–429. Springer, Heidelberg, August 2005.
- [BF06] Alexandra Boldyreva and Marc Fischlin. On the security of OAEP. In Xuejia Lai and Kefei Chen, editors, *ASIACRYPT 2006*, volume 4284 of *LNCS*, pages 210–225. Springer, Heidelberg, December 2006.
- [BH15] Mihir Bellare and Viet Tung Hoang. Resisting randomness subversion: Fast deterministic and hedged public-key encryption in the standard model. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part II*, volume 9057 of *LNCS*, pages 627–656. Springer, Heidelberg, April 2015.
- [BHK13] Mihir Bellare, Viet Tung Hoang, and Sriram Keelveedhi. Instantiating random oracles via UCes. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 398–415. Springer, Heidelberg, August 2013.
- [BJK15] Mihir Bellare, Joseph Jaeger, and Daniel Kane. Mass-surveillance without the state: Strongly undetectable algorithm-substitution attacks. In Indrajit Ray, Ninghui Li, and Christopher Kruegel, editors, *ACM CCS 2015*, pages 1431–1440. ACM Press, October 2015.
- [BK89] Manuel Blum and Sampath Kannan. Designing programs that check their work. In *21st ACM STOC*, pages 86–97. ACM Press, May 1989.
- [BLR90] Manuel Blum, Michael Luby, and Ronitt Rubinfeld. Self-testing/correcting with applications to numerical problems. In *22nd ACM STOC*, pages 73–83. ACM Press, May 1990.
- [Blu88] Manuel Blum. Designing programs that check their work, November 1988. Technical Report TR-88-009, International Computer Science Institute. Available at <http://www.icsi.berkeley.edu/pubs/techreports/tr-88-009.pdf>.

- [BNR20] Rishiraj Bhattacharyya, Mridul Nandi, and Anik Raychaudhuri. Crooked indifferenciability revisited. Private Communication, 2020.
- [BPR14] Mihir Bellare, Kenneth G. Paterson, and Phillip Rogaway. Security of symmetric encryption against mass surveillance. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 1–19. Springer, Heidelberg, August 2014.
- [BR93] Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In Dorothy E. Denning, Raymond Pyle, Ravi Ganesan, Ravi S. Sandhu, and Victoria Ashby, editors, *ACM CCS 93*, pages 62–73. ACM Press, November 1993.
- [BSW11] Dan Boneh, Amit Sahai, and Brent Waters. Functional encryption: Definitions and challenges. In Yuval Ishai, editor, *TCC 2011*, volume 6597 of *LNCS*, pages 253–273. Springer, Heidelberg, March 2011.
- [Can97] Ran Canetti. Towards realizing random oracles: Hash functions that hide all partial information. In Burton S. Kaliski Jr., editor, *CRYPTO ’97*, volume 1294 of *LNCS*, pages 455–469. Springer, Heidelberg, August 1997.
- [Can01] Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *42nd FOCS*, pages 136–145. IEEE Computer Society Press, October 2001.
- [CCG⁺16] Stephen Checkoway, Shaanan Cohney, Christina Garman, Matthew Green, Nadia Heninger, Jacob Maskewicz, Eric Rescorla, Hovav Shacham, and Ralf-Philipp Weinmann. A systematic analysis of the Juniper Dual EC incident. In *Proceedings of ACM CCS 2016*, 2016. Full version available at <http://eprint.iacr.org/2016/376>.
- [CD08] Ran Canetti and Ronny Ramzi Dakdouk. Extractable perfectly one-way functions. In Luca Aceto, Ivan Damgård, Leslie Ann Goldberg, Magnús M. Halldórsson, Anna Ingólfssdóttir, and Igor Walukiewicz, editors, *ICALP 2008, Part II*, volume 5126 of *LNCS*, pages 449–460. Springer, Heidelberg, July 2008.
- [CDGS18] Sandro Coretti, Yevgeniy Dodis, Siyao Guo, and John P. Steinberger. Random oracles and non-uniformity. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part I*, volume 10820 of *LNCS*, pages 227–258. Springer, Heidelberg, April / May 2018.
- [CDL17] Jan Camenisch, Manu Drijvers, and Anja Lehmann. Anonymous attestation with subverted TPMs. In Jonathan Katz and Hovav Shacham, editors, *CRYPTO 2017, Part III*, volume 10403 of *LNCS*, pages 427–461. Springer, Heidelberg, August 2017.
- [CDMP05] Jean-Sébastien Coron, Yevgeniy Dodis, Cécile Malinaud, and Prashant Puniya. Merkle-Damgård revisited: How to construct a hash function. In Victor Shoup, editor, *CRYPTO 2005*, volume 3621 of *LNCS*, pages 430–448. Springer, Heidelberg, August 2005.
- [CGH98] Ran Canetti, Oded Goldreich, and Shai Halevi. The random oracle methodology, revisited (preliminary version). In *30th ACM STOC*, pages 209–218. ACM Press, May 1998.
- [CGH04] Ran Canetti, Oded Goldreich, and Shai Halevi. On the random-oracle methodology as applied to length-restricted signature schemes. In Moni Naor, editor, *TCC 2004*, volume 2951 of *LNCS*, pages 40–57. Springer, Heidelberg, February 2004.
- [CHK⁺16] Jean-Sébastien Coron, Thomas Holenstein, Robin Künzler, Jacques Patarin, Yannick Seurin, and Stefano Tessaro. How to build an ideal cipher: The indifferenciability of the Feistel construction. *Journal of Cryptology*, 29(1):61–114, January 2016.
- [CMR98] Ran Canetti, Daniele Micciancio, and Omer Reingold. Perfectly one-way probabilistic hash functions (preliminary version). In *30th ACM STOC*, pages 131–140. ACM Press, May 1998.

- [CNE⁺14] Stephen Checkoway, Ruben Niederhagen, Adam Everspaugh, Matthew Green, Tanja Lange, Thomas Ristenpart, Daniel J. Bernstein, Jake Maskiewicz, Hovav Shacham, and Matthew Fredrikson. On the practical exploitability of dual EC in TLS implementations. In *Proceedings of the 23rd USENIX Security Symposium, San Diego, CA, USA, August 20-22, 2014.*, pages 319–335, 2014.
- [CRT⁺19] Sherman S. M. Chow, Alexander Russell, Qiang Tang, Moti Yung, Yongjun Zhao, and Hong-Sheng Zhou. Let a non-barking watchdog bite: Cliptographic signatures with an offline watchdog. In Dongdai Lin and Kazue Sako, editors, *PKC 2019, Part I*, volume 11442 of *LNCS*, pages 221–251. Springer, Heidelberg, April 2019.
- [DFP15] Jean Paul Degabriele, Pooya Farshim, and Bertram Poettering. A more cautious approach to security against mass surveillance. In Gregor Leander, editor, *FSE 2015*, volume 9054 of *LNCS*, pages 579–598. Springer, Heidelberg, March 2015.
- [DGG⁺15] Yevgeniy Dodis, Chaya Ganesh, Alexander Golovnev, Ari Juels, and Thomas Ristenpart. A formal treatment of backdoored pseudorandom generators. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part I*, volume 9056 of *LNCS*, pages 101–126. Springer, Heidelberg, April 2015.
- [DGK17] Yevgeniy Dodis, Siyao Guo, and Jonathan Katz. Fixing cracks in the concrete: Random oracles with auxiliary input, revisited. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *EUROCRYPT 2017, Part II*, volume 10211 of *LNCS*, pages 473–495. Springer, Heidelberg, April / May 2017.
- [DKT16] Dana Dachman-Soled, Jonathan Katz, and Aishwarya Thiruvengadam. 10-round Feistel is indifferentiable from an ideal cipher. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part II*, volume 9666 of *LNCS*, pages 649–678. Springer, Heidelberg, May 2016.
- [DM04] Stefan Dziembowski and Ueli M. Maurer. Optimal randomizer efficiency in the bounded-storage model. *Journal of Cryptology*, 17(1):5–26, January 2004.
- [DMS16] Yevgeniy Dodis, Ilya Mironov, and Noah Stephens-Davidowitz. Message transmission with reverse firewalls—secure communication on corrupted machines. In Matthew Robshaw and Jonathan Katz, editors, *CRYPTO 2016, Part I*, volume 9814 of *LNCS*, pages 341–372. Springer, Heidelberg, August 2016.
- [DOP05] Yevgeniy Dodis, Roberto Oliveira, and Krzysztof Pietrzak. On the generic insecurity of the full domain hash. In Victor Shoup, editor, *CRYPTO 2005*, volume 3621 of *LNCS*, pages 449–466. Springer, Heidelberg, August 2005.
- [DP06] Yevgeniy Dodis and Prashant Puniya. On the relation between the ideal cipher and the random oracle models. In Shai Halevi and Tal Rabin, editors, *TCC 2006*, volume 3876 of *LNCS*, pages 184–206. Springer, Heidelberg, March 2006.
- [DP07] Yevgeniy Dodis and Prashant Puniya. Feistel networks made public, and applications. In Moni Naor, editor, *EUROCRYPT 2007*, volume 4515 of *LNCS*, pages 534–554. Springer, Heidelberg, May 2007.
- [DPSW16] Jean Paul Degabriele, Kenneth G. Paterson, Jacob C. N. Schuldt, and Joanne Woodage. Backdoors in pseudorandom number generators: Possibility and impossibility results. In Matthew Robshaw and Jonathan Katz, editors, *CRYPTO 2016, Part I*, volume 9814 of *LNCS*, pages 403–432. Springer, Heidelberg, August 2016.
- [DS16] Yuanxi Dai and John P. Steinberger. Indifferentiability of 8-round Feistel networks. In Matthew Robshaw and Jonathan Katz, editors, *CRYPTO 2016, Part I*, volume 9814 of *LNCS*, pages 95–120. Springer, Heidelberg, August 2016.
- [GK03] Shafi Goldwasser and Yael Tauman Kalai. On the (in)security of the Fiat-Shamir paradigm. In *44th FOCS*, pages 102–115. IEEE Computer Society Press, October 2003.

- [GKMZ16] Matthew D. Green, Jonathan Katz, Alex J. Malozemoff, and Hong-Sheng Zhou. A unified approach to idealized model separations via indistinguishability obfuscation. In Vassilis Zikas and Roberto De Prisco, editors, *SCN 16*, volume 9841 of *LNCS*, pages 587–603. Springer, Heidelberg, August / September 2016.
- [KLT15] Jonathan Katz, Stefan Lucks, and Aishwarya Thiruvengadam. Hash functions from defective ideal ciphers. In Kaisa Nyberg, editor, *CT-RSA 2015*, volume 9048 of *LNCS*, pages 273–290. Springer, Heidelberg, April 2015.
- [KNTX10] Akinori Kawachi, Akira Numayama, Keisuke Tanaka, and Keita Xagawa. Security of encryption schemes in weakened random oracle models. In Phong Q. Nguyen and David Pointcheval, editors, *PKC 2010*, volume 6056 of *LNCS*, pages 403–419. Springer, Heidelberg, May 2010.
- [KOS10] Eike Kiltz, Adam O’Neill, and Adam Smith. Instantiability of RSA-OAEP under chosen-plaintext attack. In Tal Rabin, editor, *CRYPTO 2010*, volume 6223 of *LNCS*, pages 295–313. Springer, Heidelberg, August 2010.
- [KP09] Eike Kiltz and Krzysztof Pietrzak. On the security of padding-based encryption schemes - or - why we cannot prove OAEP secure in the standard model. In Antoine Joux, editor, *EUROCRYPT 2009*, volume 5479 of *LNCS*, pages 389–406. Springer, Heidelberg, April 2009.
- [Lev] Kirill Levchenko. Chernoff bound. <http://cseweb.ucsd.edu/~klevchen/techniques/chernoff.pdf>.
- [Lis07] Moses Liskov. Constructing an ideal hash function from weak ideal compression functions. In Eli Biham and Amr M. Youssef, editors, *SAC 2006*, volume 4356 of *LNCS*, pages 358–375. Springer, Heidelberg, August 2007.
- [LN09] Gaëtan Leurent and Phong Q. Nguyen. How risky is the random-oracle model? In Shai Halevi, editor, *CRYPTO 2009*, volume 5677 of *LNCS*, pages 445–464. Springer, Heidelberg, August 2009.
- [Men13] Joseph Menn. Exclusive: Secret contract tied NSA and security industry pioneer. *Reuters*, December 2013.
- [MR95] Rajeev Motwani and Prabhakar Raghavan. *Randomized Algorithms*. Cambridge University Press, 1995.
- [MRH04] Ueli M. Maurer, Renato Renner, and Clemens Holenstein. Indifferentiability, impossibility results on reductions, and applications to the random oracle methodology. In Moni Naor, editor, *TCC 2004*, volume 2951 of *LNCS*, pages 21–39. Springer, Heidelberg, February 2004.
- [MS15] Ilya Mironov and Noah Stephens-Davidowitz. Cryptographic reverse firewalls. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part II*, volume 9057 of *LNCS*, pages 657–686. Springer, Heidelberg, April 2015.
- [Mye01] Steven Myers. Efficient amplification of the security of weak pseudo-random function generators. In Birgit Pfitzmann, editor, *EUROCRYPT 2001*, volume 2045 of *LNCS*, pages 358–372. Springer, Heidelberg, May 2001.
- [Nie02] Jesper Buus Nielsen. Separating random oracle proofs from complexity theoretic proofs: The non-committing encryption case. In Moti Yung, editor, *CRYPTO 2002*, volume 2442 of *LNCS*, pages 111–126. Springer, Heidelberg, August 2002.
- [NIT08] Akira Numayama, Toshiyuki Isshiki, and Keisuke Tanaka. Security of digital signature schemes in weakened random oracle models. In Ronald Cramer, editor, *PKC 2008*, volume 4939 of *LNCS*, pages 268–287. Springer, Heidelberg, March 2008.
- [Pas03] Rafael Pass. On deniability in the common reference string and random oracle model. In Dan Boneh, editor, *CRYPTO 2003*, volume 2729 of *LNCS*, pages 316–337. Springer, Heidelberg, August 2003.

- [PLS13] Nicole Perloth, Jeff Larson, and Scott Shane. N.S.A. able to foil basic safeguards of privacy on web. The New York Times, 2013. <http://www.nytimes.com/2013/09/06/us/nsa-foils-much-internet-encryption.html>.
- [RSS11] Thomas Ristenpart, Hovav Shacham, and Thomas Shrimpton. Careful with composition: Limitations of the indistinguishability framework. In Kenneth G. Paterson, editor, *EUROCRYPT 2011*, volume 6632 of *LNCS*, pages 487–506. Springer, Heidelberg, May 2011.
- [RTYZ16] Alexander Russell, Qiang Tang, Moti Yung, and Hong-Sheng Zhou. Cliptography: Clipping the power of kleptographic attacks. In Jung Hee Cheon and Tsuyoshi Takagi, editors, *ASIACRYPT 2016, Part II*, volume 10032 of *LNCS*, pages 34–64. Springer, Heidelberg, December 2016.
- [RTYZ17] Alexander Russell, Qiang Tang, Moti Yung, and Hong-Sheng Zhou. Generic semantic security against a kleptographic adversary. In Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu, editors, *ACM CCS 2017*, pages 907–922. ACM Press, October / November 2017.
- [RTYZ18] Alexander Russell, Qiang Tang, Moti Yung, and Hong-Sheng Zhou. Correcting subverted random oracles. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part II*, volume 10992 of *LNCS*, pages 241–271. Springer, Heidelberg, August 2018.
- [Rub91] Ronitt A. Rubinfeld. *A Mathematical Theory of Self-checking, Self-testing and Self-correcting Programs*. PhD thesis, University of California at Berkeley, Berkeley, CA, USA, 1991. UMI Order No. GAX91-26752.
- [SFKR15] Bruce Schneier, Matthew Fredrikson, Tadayoshi Kohno, and Thomas Ristenpart. Surreptitiously weakening cryptographic systems. Cryptology ePrint Archive, Report 2015/097, 2015. <http://eprint.iacr.org/2015/097>.
- [YY96] Adam Young and Moti Yung. The dark side of “black-box” cryptography, or: Should we trust capstone? In Neal Koblitz, editor, *CRYPTO’96*, volume 1109 of *LNCS*, pages 89–103. Springer, Heidelberg, August 1996.
- [YY97] Adam Young and Moti Yung. Kleptography: Using cryptography against cryptography. In Walter Fumy, editor, *EUROCRYPT’97*, volume 1233 of *LNCS*, pages 62–74. Springer, Heidelberg, May 1997.