Limits of provable security for homomorphic encryption

Andrej Bogdanov[1]* and Chin Ho Lee[2]**

[1] Dept. of Computer Science and Engineering and Institute for Theoretical Computer Science and Communications, Chinese University of Hong Kong
[2] Dept. of Computer Science and Engineering, Chinese University of Hong Kong

Abstract. We show that public-key bit encryption schemes which support weak (i.e., compact) homomorphic evaluation of any sufficiently "sensitive" collection of functions cannot be proved message indistinguishable beyond AM ∩ coAM via general (adaptive) reductions, and beyond statistical zero-knowledge via reductions of constant query complexity. Examples of sensitive collections include parities, majorities, and the class consisting of all AND and OR functions.

We also give a method for converting a strong (i.e., distribution-preserving) homomorphic evaluator for essentially any boolean function (except the trivial ones, the NOT function, and the AND and OR functions) into a rerandomization algorithm: This is a procedure that converts a ciphertext into another ciphertext which is statistically close to being independent and identically distributed with the original one. Our transformation preserves negligible statistical error.

1 Introduction

In this work we revisit the question of basing cryptography on NP-hardness. If P equals NP then computationally secure encryption is impossible. Is the converse true? Despite considerable efforts, there is no candidate encryption scheme whose security can be plausibly reduced to the worst-case hardness of some NP-complete problem. Neither is there conclusive evidence that rules out constructions of secure encryption schemes from NP-complete problems, although several obstacles have been pointed out over the years.

Restricting the encryption Brassard [Bra79] shows that no public-key encryption scheme can be proved secure beyond NP ∩ coNP, but under the implicit assumption that every public key-ciphertext pair (queried by the reduction) can be decrypted uniquely. Goldreich and Goldwasser [GG98] argue that this assumption is unrealistic by giving examples of encryption schemes that do not satisfy it. They show that the conclusion holds under the relaxed assumption

* andrejb@cse.cuhk.edu.hk. Work supported by grants RGC GRF CUHK410111 and CUHK410112.

** chlee@cse.cuhk.edu.hk.

that invalid queries to the decryption oracle can be efficiently certified as such. (If the reduction is randomized, the limitation weakens to $\mathrm{AM} \cap \mathrm{coAM}$.)

Goldreich and Goldwasser warn that these assumptions are unrealistic as they do not apply to many known proofs of security. Bogdanov and Trevisan [BT06] point out the following example of Even and Yacobi [EY80]. They construct a public key encryption scheme and show how to solve an NP-hard problem using a distinguishing oracle. Their notion of security is unrealistic, as they require a perfect distinguishing oracle. However, their example illustrates that the restrictions imposed by Brassard and Goldreich and Goldwasser do not capture the difficulty of basing cryptography on NP hardness.

Akavia, Goldreich, Goldwasser, and Moshkovitz [AGGM06] rule out reductions from NP-complete problems to inverting one-way functions (the basis of private-key encryption) assuming that sizes of preimage sets are worst-case certifiable in NP. The same considerations apply to their argument. There are natural examples of conjectured one-way functions (for example, Goldreich's function [Gol00]) not known to satisfy the aforementioned assumptions.

Restricting the reduction Another line of works makes restrictive assumptions about the type of reduction used to prove NP-hardness. Feigenbaum and Fortnow [FF93] show that a decision problem cannot be proven NP-hard on average (unless the polynomial hierarchy collapses) by a reduction that is non-adaptive and each of its queries is uniformly distributed. Bogdanov and Trevisan [BT06] obtain the same conclusion without restricting the distribution of queries, but still under non-adaptive reductions. More precisely, they show that if there is a non-adaptive reduction from a decision problem L to a problem in distributional NP, then L must be in $\mathrm{AM/poly} \cap \mathrm{coAM/poly}$. In particular their result applies to the problem of inverting a one-way function. For this important case, Akavia et al. improve the limitation to $\mathrm{AM} \cap \mathrm{coAM}$, also assuming the reduction is non-adaptive.

Haitner, Mahmoody, and Xiao [HMX10] show that collision resistant hash functions and statistically hiding commitments cannot be proved secure beyond $\mathrm{AM} \cap \mathrm{coAM}$ via reductions that make a constant number of rounds of calls to the adversary.

Lattice-based cryptography provides examples of encryption schemes whose insecurity would imply worst-case solutions to conjectured hard problems, like finding short vectors in lattices [Ajt96]. The reduction of Regev [Reg09], which gives the most efficient cryptosystems of this kind with a proof of security (against quantum algorithms), is adaptive. For certain settings of parameters, these cryptosystems support homomorphic evaluation of a bounded class of functionalities (and general functionalities under additional security assumptions) [Gen09,vDGHV10,BV11].

Our results

We say a public-key encryption scheme supports weak (i.e. compact) homomorphic evaluation of a function $f\colon \{0,1\}^* \to \{0,1\}$ if for every n and $x_1 \dots x_n \in$

$\{0,1\}^n$ takes as inputs the public key and encryptions of the bits $x_1, \dots, x_n$ and produces an output of length polynomially bounded in the security parameter that decrypts to $f(x_1 \dots x_n)$. See Section 2 for a formal definition.

Our main theorem (Theorem 1) shows that any public key encryption scheme that supports efficient weak homomorphic evaluation of any sufficiently "sensitive" collection of functions cannot be proved message indistinguishable beyond AM ∩ coAM, even under adaptive reductions. Examples of such functions are parities, majorities, and the collection of all AND and OR functions.

Examples of encryption schemes that our result applies to include El Gamal encryption [Gam85], Paillier encryption [Pai99], as well as the more recent somewhat and fully homomorphic encryption schemes of Gentry [Gen09], Van Dijk et al. [vDGHV10], and Brakerski and Vaikuntanathan [BV11] (which build upon the lattice-based cryptosystems of Regev [Reg09] and Peikert [Pei09]).

In Theorem 2 we show that if the reduction has constant query complexity, then message indistinguishability cannot be proved beyond statistical zero knowledge (SZK), which is a subclass of AM ∩ coAM.

The reductions we consider are randomized and meet the following definition: Given an input, the reduction makes arbitrary (adaptive) queries to a distinguishing oracle for bit encryptions. We require that for any (not necessarily efficient) distinguishing oracle, which may depend on the input to the reduction, the reduction outputs the correct answer. We do not know of any cryptographic reductions that treat the adversary as a black box which fall outside our definition.

Lemma 5, which is used in the proofs of Theorems 1 and 2, gives a way to obtain rerandomization of ciphertexts from any homomorphic evaluator for the function of interest. While rerandomization has been used in constructions of homomorphic evaluators [Gen09,vDGHV10], it is not a priori clear that it is necessary for homomorphic evaluation. Homomorphic evaluation may be implemented deterministically while rerandomization requires randomness.

The statistical error of the rerandomization in Lemma 5 is noticeable. While this is sufficient for our main application, a negligible error would be desirable for most applications of rerandomization in cryptography. In Theorem 3 we show a transformation of a strong homomorphic evaluator for almost any function into a rerandomization that preserves negligible statistical error. Essentially the only exceptions to which our result does not apply are that AND, OR, and NOT functions.

Our proof

From homomorphic evaluation to rerandomization (Section 4) To begin with let's assume that we have a *strong* (i.e., distribution-preserving) homomorphic evaluator H for the majority function maj_n on n inputs. This is an algorithm that takes as inputs independent encryptions of $x_1, \dots, x_n$ and outputs a ciphertext which is statistically close to an encryption of $maj_n(x_1, \dots, x_n)$. We show that H can be used to obtain an approximate *rerandomization* **Rer**: This is a procedure that takes an encryption as its input and produces an independent and identically distributed encryption as its output. Our rerandomization will be approximate in the sense that the input and output of **Rer** will be only statistically close to independent.

One way to obtain rerandomization is as follows: Given a ciphertext C, generate $(n-1)/2$ independent encryptions of 0, $(n-1)/2$ independent encryptions of 1, randomly shuffle them together with C and feed the n resulting ciphertexts to the homomorphic evaluator for majority. By the strong homomorphic property, the output of the homomorphic evaluator will be identically distributed with C. But why should they be independent? From the point of view of the homomorphic evaluator, if C is an encryption of b, then it is indistinguishable from the other $(n-1)/2$ encryptions of b. Since the output of the homomorphic evaluator is bounded in length, the evaluator must "forget" most of the information about most of the ciphertexts it is given as inputs, including C as it is indistinguishable from the others. Therefore the output is forced to look almost statistically independent of C.

In Lemma 5 we generalize this argument to a much wider class of functions which we call *sensitive* (see Section 2) and to weak (i.e., compact) homomorphic evaluators, in which case we obtain a weaker notion of rerandomization.

A strong rerandomization procedure can be used to distinguish encryptions in statistical zero-knowledge by reduction to the "statistical distance" problem: A rerandomized encryption of 0 is statistically close to an encryption of 0, but statistically far from an encryption of 1. Mahmoody and Xiao's simulation of $\mathrm{BPP}^{\mathrm{SZK}}$ in AM [MX10] can then be used to emulate the reduction by a proof system. When only weak one-sided rerandomization is available, it is not clear that encryptions are distinguishable in statistical zero-knowledge, and we construct a somewhat different proof system. For the sake of clarity, however, in the rest of this discussion we will assume the availability of strong rerandomization.

From rerandomization to a distinguishing protocol (Section 5) To turn a reduction from distinguishing encryptions to L into a proof system for $\overline{L}$, we proceed as in previous works: The verifier plays the role of the reduction and the prover plays the role of the distinguishing oracle. The challenge is to force the prover to give answers that are consistent with a specific, fixed distinguishing oracle.

To illustrate the difficulties in the context of public key encryption, let us point out the deficiencies of some naive proof systems. Suppose the verifier submits a public key-ciphertext query (PK, C) to the prover, who is supposed to act as a distinguishing oracle. A natural attempt is to ask the prover to provide the message m and randomness R such that C is an encryption of m under public key PK with randomness R. This fails to account for the possibility that C may not be a valid ciphertext at all: Perhaps there is no pair (m, R) that encrypts to C under PK. It is not clear how a prover can certify such a statement. Another attempt would be to ask the prover for the secret key SK associated to PK. Again, it is not clear how to achieve completeness in case the public key is invalid and there is no corresponding secret key, or soundness in case the public key can be paired with several different secret keys (the choice of which may affect how different invalid ciphertexts decrypt).

Our protocol works as follows: Given a query (PK, C), the verifier asks the prover for the value b that encrypts to C, together with a proof that the rerandomization of C is statistically close to encryptions of b but statistically far from encryptions of $\overline{b}$. If the pair (PK, C) is properly distributed, this forces the prover to give a unique correct answer. But since statistical closeness and statistical farness are both efficiently verifiable [BBM11,SV03], the prover can now also certify that a pair (PK, C) is *not* a valid public key-ciphertext pair. We call this protocol DP (the distinguishing protocol).

One important detail is that the protocols for statistical closeness and statistical farness are only guaranteed to solve promise versions of these problems: For a given gap $[\ell, r)$, they can distinguish distributions that are within statistical distance ℓ from those that are at distance at least r, but give no guarantee about the outcome for instances that fall inside the gap. Therefore DP is only complete and sound provided that none of the underlying instances fall inside the respective gaps.

The proof system (Section 7) Given a reduction R from a decision problem L to distinguishing encryptions, this suggests the following constant-round proof system for $\overline{L}$: On a given input, the verifier chooses randomness for the reduction and sends this randomness to the prover. The prover sends back a transcript of the reduction interacting with a distinguishing oracle, which includes a list of queries (PK_i, C_i) made by the reduction together with an answer a_i saying if C_i encrypts 0 or 1 under PK_i, or the pair (PK_i, C_i) is invalid $(\perp)$. The verifier and prover then apply the DP protocol to certify that all the answers a_i are correct.

This proof system is complete and sound, given that all inputs (PK_i, C_i, a_i) to the DP protocol satisfy its promise. But in general the verifier does not know in advance if the promise is satisfied or not. We resolve this issue by choosing the width of the gaps $[\ell, r)$ to be sufficiently small and by having the verifier randomize the location of the gaps. This should make it unlikely for any of the queries to fall inside the promise gap of DP.

This approach was also used by Bogdanov and Trevisan [BT06] in the context of non-adaptive reductions. An additional twist is required when the reduction is adaptive because the location of the gaps may affect the answers of the honest prover. For example, imagine an adaptive reduction that does a "binary search" for the gap $[\ell, r)$: If the first answer a is to the right of r, its next query will be $a/2$, and so on until it hits the gap. To handle such reductions, we want to make the location of the gaps in each round independent of the answers of the honest prover in the previous rounds. On the other hand, the locations of these gaps must be consistent with a specific, fixed distinguishing oracle that the prover is required to emulate.

To achieve both objectives we specify a randomized family of distinguishing oracles, where for each query to the oracle the gap location is random, and the gap locations among the various queries are q-wise independent, where q is an upper bound on the number of queries performed by the reduction. In the first round of the reduction the verifier chooses a random oracle from this family and sends its (polynomial length) description to the prover. The honest prover is then expected to give answers that are consistent with this instantiation of the distinguishing oracle. By independence, the probability that any of the queries made by the honest prover falls inside the gap will be small. In Section 6.1 we develop the relevant complexity-theoretic framework and we prove Theorem 1 in Section 7.1.

To prevent any of the queries from falling into the gaps $[\ell, r)$, the size of the gaps needs to be inverse proportional to the number of queries made by the reduction. Unless the reduction makes a bounded number of queries, this requires protocols for statistical closeness and statistical farness where the verifier runs in time inverse polynomial to the size of the gap and the gap can be at an arbitrary location. Such protocols were developed by Bhatnagar, Bogdanov, and Mossel [BBM11][3] and we use them in the proof of Theorem 1.[4]

For reductions that make a constant number of queries, it is sufficient to have statistical closeness/farness protocols over a constant number of disjoint gaps $[\ell, r)$. Sahai and Vadhan [SV03] give implementations of such protocols in SZK. Using their protocols and the closure properties of SZK which we recall in Section 6.2, we prove Theorem 2 in Section 7.2.

Better rerandomization from strong homomorphic evaluation The rerandomization procedure we described above comes with a non-negligible statistical error. It is not difficult to construct examples showing that this error is inherent, even if the homomorphic evaluation is perfect, i.e. it induces no statistical error. In Section 8 we show that the statistical error can be reduced exponentially by iteratively applying the rerandomization on its output, provided f is not "exceptional". This proves Theorem 3.

2 Definitions

Homomorphic evaluation and rerandomization Let $(\mathbf{Gen}, \mathbf{Enc}, \mathbf{Dec})$ be a bit encryption scheme. Fix a security parameter s and let $(PK, SK) \sim \mathbf{Gen}(1^s)$ the distribution on key pairs. (We will assume that s is implicit in the public and secret keys.)

Definition 1. *Let $f\colon \{0,1\}^n \to \{0,1\}$ be a boolean function. We say H is a* strong homomorphic evaluator *for f with error ε if for all m in the domain of f, the random variables $(PK, H_{PK}(\mathbf{Enc}_{PK}(m_1), \ldots, \mathbf{Enc}_{PK}(m_n)))$ and $(PK, \mathbf{Enc}_{PK}(f(m)))$ (where all encryptions are independent) are within statistical distance ε.*

[3] Technically their statement is not as strong as the one we need here, but their proof can be easily adapted. We provide the details in the full version.

[4] Similar issues arise in the work of Mahmoody and Xiao [MX10]. They work with the SZK-complete problem entropy difference. While their proof can be adapted to our setting, we find it more natural to work directly with instances of statistical difference.

This definition extends to functions from $\{0,1\}^* \to \{0,1\}$ in a straightforward way. We omit the details.

Definition 2. *Let $f\colon \{0,1\}^* \to \{0,1\}$ be a boolean function. We say H is a* weak homomorphic evaluator *for f with error ε if (1) the output length of H is bounded by a function that depends only on the security parameter and (2) for all n and $m \in \{0,1\}^n$ in the domain of f,*

$$\Pr[\mathbf{Dec}_{SK}(PK, H_{PK}(\mathbf{Enc}_{PK}(m_1), \ldots, \mathbf{Enc}_{PK}(m_n))) = f(m)] \geq 1 - \varepsilon,$$

where all encryptions are independent.[5]

A bit encryption scheme is *efficient* if $\mathbf{Gen}, \mathbf{Enc}, \mathbf{Dec}$ all run in time polynomial in the security parameter s. A homomorphic evaluator H is efficient if it is computable in time polynomial in s and n and its output length is polynomially bounded in s.

Definition 3. *Let $\mathbf{Rer}$ be a randomized function that takes as input a public key and a ciphertext. In the following definitions R and R' are independent choices of randomness for $\mathbf{Rer}$.*

- *We say $\mathbf{Rer}$ is a* strong rerandomization *with error ε if for every $m \in \{0,1\}$, the random variables $(PK, E, \mathbf{Rer}_{PK}(E, R))$ and (PK, E, E') where $E, E' \sim \mathbf{Enc}_{PK}(m)$ are independent are within statistical distance ε.*
- *For $b \in \{0,1\}$, we say $\mathbf{Rer}^b$ is a* one-sided weak rerandomization *with decryption error ε and rerandomization error ρ if for every $m \in \{0,1\}$, $\Pr[\mathbf{Dec}_{SK}(\mathbf{Rer}^b_{PK}(\mathbf{Enc}_{PK}(m))) = m] \geq 1 - \varepsilon$ and the random variables $(PK, \mathbf{Rer}^b_{PK}(E, R), \mathbf{Rer}^b_{PK}(E, R'))$ and $(PK, \mathbf{Rer}^b_{PK}(E, R), \mathbf{Rer}^b_{PK}(E', R'))$ where $E, E' \sim \mathbf{Enc}_{PK}(b)$ are independent are within statistical distance ρ.*

We say the rerandomization is *efficient* if it can be evaluated in time polynomial in the security parameter.

Sensitivity of boolean functions We will use the following notion of sensitivity for boolean functions. For $x \in \{0,1\}^k$ let $x|_i$ be the string obtained by flipping the i-th bit of x and leaving the others unchanged. Let $f\colon \{0,1\}^k \to \{0,1\}$ be a boolean function and $b \in \{0,1\}$. We say f has *b-sensitivity* at least s if there exists an input $x \in \{0,1\}^k$ and a set $S \subseteq [k]$ of size s such that $f(x) = b$, $x_i = b$ for every $i \in S$, and $f(x|_i) = \overline{b}$ for every $i \in S$. We call (x, S) a witness that f has b-sensitivity at least s.

We say a family of functions $f = \{f_k\colon \{0,1\}^k \to \{0,1\}\}$ has *certifiable polynomial b-sensitivity* if there exists a constant $\alpha > 0$ so that on input k we can compute in time polynomial in k a witness that f_k has b-sensitivity at least k^α.

[5] Some works adopt the terms “distribution preserving” and ”compact” homomorphic evaluation. We prefer the terms “strong” and “weak” for this work, as we are concerned with questions of computational complexity.

Examples of functions that have certifiable polynomial 0-sensitivity and 1-sensitivity include parity and majority. The AND function has certifiable polynomial 0-sensitivity while the OR function has certifiable polynomial 1-sensitivity.

Examples of functions whose 0-sensitivity and 1-sensitivity is less than s are functions that depend on at most $s-1$ of their inputs, i.e. $(s-1)$-juntas. Simon [Sim82] gives an example of a function on k bits that depends on all its inputs but has 0-sensitivity and 1-sensitivity $O(\log k)$.

3 The main theorems

We say $(\mathbf{Gen}, \mathbf{Enc}, \mathbf{Dec})$ supports weak homomorphic evaluation of f with error ε if it has an efficient homomorphic evaluator for f with error ε.

A γ-*distinguishing oracle* for $(\mathbf{Gen}, \mathbf{Enc}, \mathbf{Dec})$ is a function D such that

$$\Pr[D(PK, \mathbf{Enc}_{PK}(0)) \text{ accepts}] - \Pr[D(PK, \mathbf{Enc}_{PK}(1)) \text{ accepts}] > \gamma.$$

A *reduction* from a decision problem L to γ-distinguishing encryptions is an efficient randomized oracle algorithm $R^?$ such that for every valid input x there exists a γ-distinguishing oracle D such that $R^D(x) = L(x)$ with probability at least 8/9. (For our results the exact constant won't matter, as long as it is strictly greater than 1/2.)

Theorem 1. *Let f_0 and f_1 be functions with certifiable polynomial 0-sensitivity and 1-sensitivity respectively (possibly the same function). Let $\varepsilon \in (0, 1/18)$ be any constant and $\delta \geq 2\sqrt{\varepsilon}$. Let $(\mathbf{Gen}, \mathbf{Enc}, \mathbf{Dec})$ be a public key encryption scheme that supports efficient homomorphic evaluations of both f_0 and f_1 with error at most ε. If there is a reduction from L to $(1-\delta)$-distinguishing $(\mathbf{Gen}, \mathbf{Enc}, \mathbf{Dec})$, then L is in $\mathrm{AM} \cap \mathrm{coAM}$.*

We will assume that the reduction is *query length regular*: On input x, the reduction first computes a query length $m \geq |x|$ and only makes queries of length m. The theorem can be proved without this assumption, but we make it for notational convenience.

In the case when the reduction has constant query complexity, a stronger conclusion can be obtained.

Theorem 2. *Let f_0 and f_1 be functions with certifiable polynomial 0-sensitivity and 1-sensitivity respectively (possibly the same function). Let q be any constant, $\delta > 0$, and $\varepsilon = \varepsilon(q, \delta)$ a sufficiently small constant. Let $(\mathbf{Gen}, \mathbf{Enc}, \mathbf{Dec})$ be a public key encryption scheme that supports efficient homomorphic evaluations of f_0 and f_1 with error at most ε. If there is a reduction from L to $(1-\delta)$-distinguishing $(\mathbf{Gen}, \mathbf{Enc}, \mathbf{Dec})$ that makes at most q queries, then L is in statistical zero-knowledge.*

In particular, Theorems 1 and 2 apply to the following cases: (1) f_0 and f_1 are the parity function; (2) f_0 and f_1 are the majority function; (3) f_0 is OR and f_1 is AND.

Ron Rothblum [Rot11] shows how to turn a private-key encryption scheme into a public-key one using a homomorphic evaluator for parity. Combining the two results, the conclusions of Theorems 1 and 2 can be extended to private-key encryption schemes that support homomorphic evaluation of parity.

Our last result shows how to obtain strong rerandomization given a homomorphic evaluator for almost any function. We call a function $f\colon \{0,1\}^n \to \{0,1\}$ *exceptional* if it is one of the following functions of the inputs that it depends on: the constant 0, the constant 1, the identity, the NOT function, the AND function, the OR function.

Theorem 3. *Assume $f\colon \{0,1\}^n \to \{0,1\}$ is not exceptional. If $(\mathbf{Gen}, \mathbf{Enc}, \mathbf{Dec})$ is a public key encryption scheme that supports efficient strong homomorphic evaluation of f with negligible error, then $(\mathbf{Gen}, \mathbf{Enc}, \mathbf{Dec})$ has an efficient strong rerandomization with negligible error.*

4 One-sided rerandomization from homomorphic evaluation

In this section we show how to convert a homomorphic evaluation algorithm for a sensitive function into a one-sided rerandomization. In Section 8 we extend these ideas to obtain stronger notions of rerandomization under stronger assumptions. Let H denote entropy and I denote mutual information.

Claim 4. *Let $X_1, \dots, X_n$ be i.i.d. random variables and $I \sim \{1, \dots, n\}$ a uniformly random index. Let F, G, G' be random variables such that (1) F and G are independent conditioned on X_I, (2) F is independent of I, (3) G and G' are identically distributed and (4) F and G' are independent. Then the random variables (F, G) and (F, G') are within statistical distance $\sqrt{2\,\mathrm{H}(F)/n}$.*

Proof.

$$\begin{aligned}\mathrm{H}(X_I \mid F) &\geq \mathrm{H}(X_I \mid F, I) = \frac{1}{n}\sum_{i=1}^{n} \mathrm{H}(X_i \mid F)\\ &\geq \frac{1}{n}\mathrm{H}(X_1, \dots, X_n \mid F) \geq \frac{1}{n}(\mathrm{H}(X_1, \dots, X_n) - \mathrm{H}(F)) = \mathrm{H}(X_I) - \frac{\mathrm{H}(F)}{n}.\end{aligned}$$

Since F and G are conditionally independent of X_I, $\mathrm{I}(F; G) \leq \mathrm{I}(F; X_I)$ and so

$$\mathrm{I}(F; G) \leq \mathrm{I}(F; X_I) = \mathrm{H}(X_I) - \mathrm{H}(X_I \mid F) \leq \frac{\mathrm{H}(F)}{n}.$$

The conclusion follows by Pinsker's inequality [Pin64]. □

The following lemma shows how to obtain one-sided rerandomization from homomorphic evaluation of a sensitive function. This lemma will be used in the proofs of Theorems 1 and 2. In Section 8 we give a version of this lemma that applies to a more restricted class of functions but allows us to achieve a stronger notion of rerandomization. That version will be used for the proof of Theorem 3.

Lemma 5. *Assume f has certifiable polynomial b-sensitivity and let δ be any function inverse polynomial in the security parameter. If* $(\mathbf{Gen}, \mathbf{Enc}, \mathbf{Dec})$ *has a weak efficient homomorphic evaluator for f with error ε, then it has a one-sided weak rerandomization* $\mathbf{Rer}^b$ *with decryption error ε and rerandomization error δ.*

Proof. Suppose f_k has b-sensitivity k^α. Choose $k = (2c/\delta^2)^{1/\alpha}$, where c is the length of ciphertexts (for the given security parameter). Let (x, S) be the witness for b-sensitivity of f_k. Given public key PK and ciphertext E define $\mathbf{Rer}^b$ as follows:

1. Choose a random $I \sim S$.
2. Let
$$X_i = \begin{cases} \mathbf{Enc}_{PK}(x_i, R_i) & \text{if } i \neq I, \\ E & \text{if } i = I. \end{cases}$$
3. Output $F = H_{PK}(X_1, \ldots, X_k)$.

We first condition on the choice of the public key PK, letting ε_{PK} denote the statistical distance between the two distributions in the definition of strong homomorphic evaluator conditioned on PK.

The decryption error of $\mathbf{Rer}^b$ follows directly from the definition. We now show the rerandomization error is at most δ. Let F, G be two independent instantiations of $\mathbf{Rer}^b$ on the same input E. Conditioned on PK, the random variables $X_i \colon i \in S$ and F, G satisfy the assumptions of Claim 4. It follows that (F, G) and (F, G'), where G' is i.i.d with G and therefore with F, are within statistical distance $\sqrt{2c/k^\alpha}$, which is at most δ by our choice of parameters. Averaging over ε_{PK} we prove the lemma. □

5 The distinguishing protocol

In this section we describe a constant-round interactive proof system DP that, given input (PK, C, b), certifies that C is an encryption of b under PK when $b \in \{0, 1\}$ and that (PK, C) is an invalid pair when $b = \perp$. The proof system is parametrized by two gaps $[\ell, r)$ and $[\ell', r')$, which describe a promise on the inputs.

We will assume we have the following constant-round protocols for statistical closeness ($SC_{[\ell,r)}$) and statistical farness ($SF_{[\ell,r)}$), where $0 \leq \ell < r \leq 1$. The protocols take as inputs a pair of sampler circuits D, D' producing distributions over the same set $\{0, 1\}^m$ with the following completeness / soundness properties:

- If D, D' are at statistical distance less than ℓ / at least r, $SC_{[\ell,r)}(D, D')$ accepts / rejects with probability $1 - \sigma$.
- If D, D' are at statistical distance at least r / less than ℓ, $SF_{[\ell,r)}(D, D')$ accepts / rejects with probability $1 - \sigma$.

Here σ can be any inverse polynomial in the size of the input. The following two theorems state the existence of these protocols. The second one is stronger as it provides statistical zero-knowledge implementation, but makes a stronger assumption about the gaps.

Formally we will view SC and SF as promise problems that take ℓ, r, D, D' as their inputs. Theorem 6 essentially follows from work of Bhatnagar, Bogdanov, and Mossel [BBM11]. We provide the missing details in the full version..

Theorem 6. *For $r > \ell$, the problems SC and SF are in* AM *where the running time of the verifier is polynomial in the size of D, the size of D', and $1/(r - \ell)$.*

Theorem 7 is proved by Sahai and Vadhan [SV03].

Theorem 7. *For $r^2 > \ell$, the problems SC and SF are in* SZK *where the running time of the verifier is polynomial in the size of D, the size of D', and $1/\ell^{1/\log(r^2/\ell)}$.*

The protocol DP will certify that the rerandomization of C is close to an rerandomized encryption of b but far from a rerandomized encryption of $\overline{b}$ when $b \in \{0, 1\}$. When $b = \bot$, it certifies that either the rerandomized encryptions of 0 and 1 are close, or the rerandomized encryption of C is far from both.

Let $Z_{PK,b}$ ($b \in \{0, 1\}$) be the following circuit: On input R, R', output $\mathbf{Rer}^b_{PK}(\mathbf{Enc}_{PK}(b, R), R')$, i.e. a one-sided rerandomized encryption of b.

The distinguishing protocol $DP_{[\ell,r),[\ell',r')}$

```
On input $(PK, C, b)$, where $b \in \{0, 1, \bot\}$:
If $b = 0$ or $b = 1$:
    Verifier and Prover execute $SF_{[\ell,r)}(Z_{PK,0}, Z_{PK,1})$.
    If the protocol rejects, reject. Otherwise:
        Verifier and Prover execute $SC_{[\ell',r')}(Z_{PK,b}, \mathbf{Rer}^b_{PK}(C))$.
        If the protocol accepts, accept, else reject.
If $b = \bot$:
    Verifier and Prover execute $SC_{[\ell,r)}(Z_{PK,0}, Z_{PK,1})$.
    If the protocol accepts, accept. Otherwise:
        Verifier and Prover execute $SF_{[\ell',r')}(Z_{PK,0}, \mathbf{Rer}^0_{PK}(C))$.
        Verifier and Prover execute $SF_{[\ell',r')}(Z_{PK,1}, \mathbf{Rer}^1_{PK}(C))$.
        If both accept, accept, else reject.
```

The distinguishing oracle We define an oracle π that distinguishes between encryptions of 0 and encryptions of 1. This oracle answers $\bot$ on all queries (PK, C) that do not represent valid key-ciphertext pairs and answers bad on all queries that fall inside the gaps of the underlying protocols SC and SF. We then show that for every pair (PK, C) that falls outside the gaps, $b = \pi(PK, C)$ is the unique answer that makes $DP(PK, C, b)$ accept. Owing to lack of space the definition of π, as well as the proof of the following claim which shows π is a distinguishing oracle, are given in the full version.

Claim 8. *Assume* $\mathbf{Rer}^0, \mathbf{Rer}^1$ *are one-sided rerandomizations with decryption error* $\varepsilon < (1-r)^2/2$ *and rerandomization error* $\rho < \ell'^2$*. Then* $\Pr[\pi(PK, \mathbf{Enc}_{PK}(b)) = b] \geq 1 - \sqrt{2\varepsilon} - \sqrt{\rho}$ *for every* $b \in \{0,1\}$*.*

The following claims are immediate from the definitions and the completeness and soundness assumptions on SC and SF.

Claim 9. *(Completeness) Assume* $\ell' < r/2$ *and* $\pi(PK, C) \neq \mathrm{bad}$*. Then* $DP(PK, C, \pi(PK, C))$ *accepts with probability at least* $1 - 3\sigma$*.*

Claim 10. *(Soundness) Assume* $\ell' < r/2$*. If* $DP(PK, C, b)$ *accepts with probability more than* 3σ*, then* $\pi(PK, C) \in \{b, \mathrm{bad}\}$*.*

6 Complexity theoretic setup

In this section we cover the complexity-theoretic framework for the proofs of Theorems 1 and 2. Proof of the claims can be found in the full version.

6.1 Promise oracles for adaptive reductions

Let Ξ be any finite set of values that includes the special symbol bad. An *oracle family* over input length m with size d is a multiset Π of functions $\pi \colon \{0,1\}^m \to \Xi$. We say Π is ε*-bad* if for every input x, $\Pr_{\pi \sim \Pi}[\pi(x) = \mathrm{bad}] \leq \varepsilon$.

Let $F \colon \{0,1\}^m \to [d]$ be a function. The oracle $\Pi_F \colon \{0,1\}^m \to \Xi$ is given by $\Pi_F(z) = \pi_{F(z)}(z)$. In the lemma below F will be a randomized function of the same form.

Lemma 11. *Let* $R^?$ *be a reduction that on an input of length* n*, makes at most* q *queries of length* m*. Let* Π *be an oracle family of size* d*. Assume* d *is a power of two. There exists a randomized function* $F \colon \{0,1\}^m \to [d]$ *such that:*

- *F is computable in time (and hence uses randomness) polynomial in m, q, and d.*
- *For every input x of length n, the probability that $R^{\Pi_F}(x)$ never receives* bad *as an answer to any of its queries is at least $(1-\varepsilon)^q$.*

6.2 Statistical zero-knowledge

We recall some results about the complexity of statistical zero-knowledge SZK. Sahai and Vadhan [SV03] show that the statistical distance problem $SD = SF_{[1/9, 8/9)}$ is complete for SZK under many-one reductions.

We also need the following result of Sahai and Vadhan [SV03] that states the closure of SZK under truth-table reductions.

Theorem 12. *There is a deterministic algorithm that takes as input instances* $x_1, \ldots, x_k$ *of* SD *and a boolean predicate* $P \colon \{0,1\}^k \to \{0,1\}$ *and outputs an instance* x *of* SD *such that* $SD(x) = P(SD(x_1), \ldots, SD(x_k))$*. The running time of the algorithm is polynomial in* 2^k *and the sizes of* $x_1, \ldots, x_k$*.*

We also need the following fact, which says that reductions within SZK can without loss of generality be assumed deterministic.

Claim 13. *If there is a randomized many-one reduction R from L to SD such that $\Pr[SD(R(x)) = L(x)] \geq p$, where p is any constant above $1/2$, then L is in* SZK*.*

Combining Theorem 12 and Claim 13 we get the following corollary.

Corollary 14. *Suppose there is a randomized algorithm A that on input x of length n and randomness r computes inputs $x_1, \ldots, x_k$ and a predicate $P\colon \{0,1\}^k \to \{0,1\}$, where $k = O(\log n)$ and accepts if $P(SD(x_1), \ldots, SD(x_k))$ is true. If $\Pr[A(x) = L(x)] \geq p$, where p is any constant greater than $1/2$, then L is in* SZK*.*

7 Proofs of the main theorems

7.1 Proof of Theorem 1

Let $F_\omega\colon \{0,1\}^m \to [d]$ be the randomized function from Lemma 11, with ω describing the randomness. We set $I_j = \left[\frac{1}{3} + \frac{j-1}{3d}, \frac{1}{3} + \frac{j}{3d}\right)$ and $I'_j = \frac{1}{3} I_j$, where $1 \leq j \leq d$.

The decision protocol DL: On input x:

V: Compute the oracle query length m. Let d be the smallest power of two above $90q$ where q is an upper bound on the number of queries $R^?(x)$ makes. Choose randomness r for the reduction and randomness ω for F_ω. Send r, d, ω to the prover.

P: Send a sequence $((PK_i, C_i), b_i)$, $1 \leq i \leq q$ of oracle query-answer pairs.

V: Verify that the received query-answer pairs determine an accepting computation of $R^?(x, r)$. If not, reject. For every query i, compute $j = F_\omega(PK_i, C_i)$ and let $[\ell_i, r_i) = I_j$ and $[\ell'_i, r'_i) = I'_j$.

B: Execute in parallel the protocols $DP_{[\ell_i, r_i), [\ell'_i, r'_i)}(PK_i, C_i, b_i)$ for $1 \leq i \leq q$ with completeness/soundness gap $\sigma = 1/9q$. If any of them result in rejection, reject. Otherwise, accept.

Let $\pi_j = \pi_{I_j, I'_j}$ and $\mathit{\Pi}_F$ be the oracle from Lemma 11.

Claim 15. *The oracle family $\{\pi_j\}_{1 \leq j \leq d}$ is at most $3/d$-bad.*

Proof of Theorem 1 It is sufficient to prove that $L \in$ AM. By applying the same argument to its complement $\overline{L}$ we also get $L \in$ coAM.

Assume $(\mathbf{Gen}, \mathbf{Enc}, \mathbf{Dec})$ supports homomorphic evaluation of f with error at most ε and there is a reduction $R^?$ from L to $(1-\delta)$-distinguishing encryptions.

We instantiate the constructions with the following parameters. Let ε be the homomorphic evaluation error and c an upper bound on the length of ciphertexts queried by the reduction on input x. Let $\mathbf{Rer}^b$ be the rerandomization from Lemma 5 with parameters chosen so that the decryption error is ε and the rerandomization error is at most $\rho \leq \varepsilon^2$. The protocol DP is instantiated with the rerandomizations $\mathbf{Rer}^0$ and $\mathbf{Rer}^1$.

Claim 16. *For an appropriate choice of parameters and for every F, Π_F is a $(1-\delta)$-distinguishing oracle.*

By Theorem 6, the verifier for DL can be implemented in polynomial time. Theorem 1 the follows by the next two claims:

Claim 17. *(Completeness) If $x \in L$, there exists a prover that makes $DL(x)$ accept with probability at least $2/3$.*

Claim 18. *(Soundness) If $x \notin L$ then no prover makes $DL(x)$ accept with probability at least $1/3$.*

7.2 Proof of Theorem 2

Let $I_j, 1 \leq j \leq d$ be the following collection of intervals: $I_j = [\ell_j, r_j)$ where $r_1 = 1/2$, $\ell_j = r_j^2/4$, and $r_{j+1} = \ell_j$. Let $I'_j = \frac{1}{3}I_j$. Assume the reduction makes at most q queries on every input and let $d = 27q \cdot 3^q$.

By Theorem 7, for every j the problems $SC_{I_j}, SC_{I'_j}, SF_{I_j}, SF_{I'_j}$ are all in SZK so by Theorem 12 and the completeness of SD, DP_{I_j,I'_j} is also in SZK for every j.

Consider the following algorithm A. On input x, choose randomness r for R and a random $j \sim [d]$ and accept if there exists a sequence of answers $(a_1, \dots, a_q) \in \{0, 1, \perp\}^q$ such that $R(x, r)$ accepts given these oracle answers and $DP_{I_j,I'_j}(Q_i, a_i)$ accepts for all $1 \leq i \leq q$. Since DP_{I_j,I'_j} is in SZK and SD is complete for SZK, A satisfies the assumption of Corollary 14, so if we can prove that $\Pr[A(x) = L(x)] \geq 2/3$, it will follow that L is in SZK.

Say j is bad if $\pi_j = \pi_{I_j,I'_j}$ answers bad on any pair (Q, a) queried by A. Since A makes at most $q3^q$ queries, by Claim 15 and a union bound the probability that A answers bad on any of its queries is at most $1/9$.

Fix an input x. By our choice of parameters, when ε is sufficiently small and $\rho = \varepsilon^2$, Claim 8 guarantees that π_j is a $(1-4\varepsilon)$-decryption oracle for every $1 \leq j \leq d$. So for at least $8/9$ fraction of r, $R^{\pi_j}(x, r) = L(x)$. Therefore with probability at least $7/9$, both $R^{\pi_j}(x, r) = L(x)$ and π_j never answers bad on any of A's queries. By Claims 9 and Claim 10, it must then hold that $a = \pi_j(Q)$ for all query-answer pairs (Q, a) made by A, and so $A(x) = L(x)$.

8 Strong rerandomization from strong homomorphic evaluation

In this Section we prove Theorem 3. We begin by defining "t-symmetric functions". The proofs of the claims in this section can be found in the full version.

t-symmetric functions Let G be a subgroup of the symmetric group S_k and $x \in \{0,1,\star\}^k$ be a string containing exactly one $\star$. Let $t_0(G,x)$ (resp., $t_1(G,x)$) be the number of transpositions $\tau \in G$ that transpose a 0 and a $\star$ (resp., a 1 and a $\star$) when acting on x. Observe that $t_b(G,\sigma x) = t_b(G,x)$ for every $\sigma \in G$.

Let $x|_{\star\to 0}, x|_{\star\to 1}$ be the string obtained when the $\star$ in x is replaced by a 0 and a 1 respectively. We will say a boolean function $f\colon \{0,1\}^k \to \{0,1\}$ is t-symmetric if there exist x and G with $t_0(G,x), t_1(G,x) > t$ and $f(\sigma x|_{\star\to b}) = b$ for every $\sigma \in G$.

For example, the majority function on 3 bits is 2-symmetric: Take $G = S_3$ and let $x = 01\star$. So is parity on 4 bits: Take $G = S_4$ and $x = 110\star$. The DNF $(x_{11} \wedge x_{12}) \vee (x_{21} \wedge x_{22})$ is also 2-symmetric. To see this take x to be the string $x_{11} = \star, x_{12} = 1, x_{21} = 0, x_{22} = 1$ and G to be the "wreath product" $S_2 \wr S_2$, which acts on x by first permuting the inputs in each term of the DNF independently, then permuting the terms.

Proof of Theorem 3 The theorem follows from the next two claims, proved below.

Claim 19. *Let $f\colon \{0,1\}^k \to \{0,1\}$, $k \geq 2$ be any boolean function that depends on all its inputs and is not one of OR / AND. If $(\mathbf{Gen},\mathbf{Enc},\mathbf{Dec})$ supports efficient strong homomorphic evaluation of f with error ε, then $(\mathbf{Gen},\mathbf{Enc},\mathbf{Dec})$ supports efficient strong homomorphic evaluation of a 2-symmetric function with error at most 12ε.*

Claim 20. *Let $f\colon \{0,1\}^k \to \{0,1\}$ be a 2-symmetric function. If $(\mathbf{Gen},\mathbf{Enc},\mathbf{Dec})$ is a public key encryption scheme that supports efficient strong homomorphic evaluation of f with negligible error, then $(\mathbf{Gen},\mathbf{Enc},\mathbf{Dec})$ has an efficient strong rerandomization with negligible error.*

8.1 Proof of Claim 19

Claim 21. *Let $f\colon \{0,1\}^k \to \{0,1\}$, $k \geq 2$ be a monotone function that depends on all its inputs.*

1. *If f is not the AND function, then f has 0-sensitivity at least 2.*
2. *If f is not the OR function, then f has 1-sensitivity at least 2.*

Let $f\colon \{0,1\}^k \to \{0,1\}$ be a boolean function. We say f is an *extension* of g if there exists a set $S \in [k]$ and $z \in \{0,1\}^{\overline{S}}$ such that g is the restriction of f to S using z, i.e. $f_{S|z}(x) = g(x)$ for every $x \in \{0,1\}^S$.

Claim 22. *Let g be a function with b-sensitivity at least s and f be any extension of g. Let $(\mathbf{Gen},\mathbf{Enc},\mathbf{Dec})$ be a public key encryption scheme. If $(\mathbf{Gen},\mathbf{Enc},\mathbf{Dec})$ supports strong homomorphic evaluation of f with error ε, $(\mathbf{Gen},\mathbf{Enc},\mathbf{Dec})$ supports strong homomorphic evaluation of g with error ε.*

Claim 23. *Let $g\colon \{0,1\}^k \to \{0,1\}$ be a boolean function. For every $i \in [k]$, let $f_i\colon \{0,1\}^{k_i} \to \{0,1\}$ be a boolean function. Let $(\mathbf{Gen},\mathbf{Enc},\mathbf{Dec})$ be a public key encryption scheme. If $(\mathbf{Gen},\mathbf{Enc},\mathbf{Dec})$ supports strong homomorphic evaluation of g with error ε and each of the f_i's with error ε_i, then $(\mathbf{Gen},\mathbf{Enc},\mathbf{Dec})$ supports strong homomorphic evaluation of $g(f_1,\ldots,f_k)$ with error $\varepsilon+\varepsilon_1+\cdots+\varepsilon_k$.*

Proof (of Claim 19). First, we show that $(\mathbf{Gen}, \mathbf{Enc}, \mathbf{Dec})$ supports homomorphic evaluation of f_0 and f_1 with error at most 4ε, where f_b has b-sensitivity 2. Consider the 2-symmetric function $g\colon \{0,1\}^4 \to \{0,1\}$ defined by $g(x_{11}, x_{12}, x_{21}, x_{22}) = f_0(f_1(x_{11}, x_{12}), f_1(x_{21}, x_{22}))$. Since g is a composition of f_0 and f_1, by Claim 23 $(\mathbf{Gen}, \mathbf{Enc}, \mathbf{Dec})$ has a strong homomorphic evaluation of g with error at most 12ε.

Now we show that $(\mathbf{Gen}, \mathbf{Enc}, \mathbf{Dec})$ supports homomorphic evaluation of f_0 and f_1. This follows from Claim 21 and 22 if f is monotone. If f is not monotone, there is an $x \in \{0,1\}^k$ and $i \in [k]$ such that $x_i = 1$, $f(x) = 0$ and $f(x|_i) = 1$. Let g be the restriction of f to the rest of the bits using x_i. Note that g is the NOT function and so by Claim 22 $(\mathbf{Gen}, \mathbf{Enc}, \mathbf{Dec})$ supports homomorphic evaluation of the NOT function with error ε. It is easy to see that one can obtain f_0 and f_1 by composing g with a restriction of f. The rest follows by Claim 23. □

8.2 Proof of Claim 20

We start with the following Corollary of Claim 4 for the special case when $G = X_I$.

Corollary 24. *Let $X_1, \ldots, X_n$ be i.i.d and $I \sim \{1, \ldots, n\}$ a uniformly random index and F be independent of I. Then (F, X_I) and (F, X) are within statistical distance $\sqrt{2\,\mathrm{H}(F)/n}$, where X is i.d. with $X_1, \ldots, X_n$ and independent of F.*

The following lemma shows how to obtain strong rerandomization from any t-symmetric function. The resulting rerandomization error is noticeable. It is similar to Lemma 5 and the proof is given in the full version.

Lemma 25. *Let $f\colon \{0,1\}^k \to \{0,1\}$ be any t-symmetric function. If $(\mathbf{Gen}, \mathbf{Enc}, \mathbf{Dec})$ has a strong efficient homomorphic evaluator for f with error ε, then it has a strong efficient rerandomization $\mathbf{Rer}$ with error at most $\varepsilon + \sqrt{2c/t}$ (resp. decryption error ε and rerandomization error $\sqrt{2c/t}$), where c is the length of ciphertexts.*

We now show that for strong homomorphic evaluation, the error can be reduced and prove Theorem 3.

For a boolean function $f\colon \{0,1\}^k \to \{0,1\}$, Let $f^{(r)}\colon \{0,1\}^{k^r} \to \{0,1\}$ be defined recursively by first applying $f^{(r-1)}$ on k independent tuples of k^{r-1} inputs and then applying f to these k values. For the base case we take $f^{(1)} = f$.

Claim 26. *If f is t-symmetric, then $f^{(r)}$ is t^r-symmetric.*

Proof (of Claim 20). Let $\mathbf{Rer}$ be the rerandomization of f from the proof of Lemma 25. We define $\mathbf{Rer}^{(r)}$ recursively by $\mathbf{Rer}^{(1)} = \mathbf{Rer}$ and

$$\mathbf{Rer}^{(r)}_{PK}(E, (R_1, \ldots, R_r)) = \mathbf{Rer}_{PK}(\mathbf{Rer}^{(r-1)}_{PK}(E, (R_1, \ldots, R_{r-1})), R_r).$$

where $R_1, \ldots, R_r$ are independent random strings. We now argue that $\mathbf{Rer}^{(r)}$ has the desired properties.

Let $\mathbf{Rer}'^{(r)}$ be the rerandomization obtained by applying the construction of Lemma 25 to the function $f^{(r)}$. We claim that the distributions $(PK, E, \mathbf{Rer}^{(r)}_{PK}(E))$ and $(PK, E, \mathbf{Rer}'^{(r)}_{PK}(E))$, where $E \sim \mathbf{Enc}_{PK}(b)$, are within statistical distance at most εk^{r-1}. We show this by induction. The base case $r = 1$ is obvious (the statistical distance is zero).

For the inductive step, we can describe $\mathbf{Rer}^{(r)}_{PK}(E)$ as follows: First, choose X by applying a random permutation π to the indices of $x \in \{0, 1, \star\}$. Then $\mathbf{Rer}^{(r)}_{PK}(E) = H_{PK}(e_1, \dots, e_k)$ where $e_i = \mathbf{Enc}_{PK}(X_i)$ when $X_i \neq \star$ and $e_i = \mathbf{Rer}^{(r-1)}_{PK}(E)$ when $X_i = \star$. On the other hand $\mathbf{Rer}'^{(r)}_{PK}(E)$ can be described as follows: First, choose X by applying a random permutation π to the indices of $x \in \{0, 1, \star\}$. Then $\mathbf{Rer}'^{(r)}_{PK}(E) = H_{PK}(e'_1, \dots, e'_k)$ where $e'_i = \mathbf{Rer}'^{(r-1)}_{PK}(\mathbf{Enc}_{PK}(X_i))$ when $X_i \neq \star$ and $e'_i = \mathbf{Rer}'^{(r-1)}_{PK}(E)$ when $X_i = \star$. By inductive assumption, the statistical distance between $(PK, \mathbf{Rer}^{(r-1)}_{PK}(E))$ and $(PK, \mathbf{Rer}'^{(r-1)}_{PK}(E))$ is at most εk^{r-2}. Since H_{PK} has error ε, the statistical distance between $(PK, \mathbf{Enc}_{PK}(b))$ and $(PK, \mathbf{Rer}'^{(r-1)}_{PK}(\mathbf{Enc}_{PK}(b))$ can also be bounded by εk^{r-2} using an inductive argument. Applying a hybrid argument we conclude that the distributions $(PK, e_1, \dots, e_k)$ and $(PK, e'_1, \dots, e'_k)$ are within distance at most εk^{r-1} and therefore so are the distributions $(PK, \mathbf{Rer}^{(r)}_{PK}(E))$ and $(PK, \mathbf{Rer}'^{(r)}_{PK}(E))$.

By Claim 26, $f^{(r)}$ is t^r symmetric. It follows from Claim 23 that the function $H^{(r)}_{PK}$ defined recursively by $H^{(1)}_{PK} = H_{PK}$ and $H^{(r)}_{PK} = H_{PK}(H^{(r-1)}_{PK}, \dots, H^{(r-1)}_{PK})$ is a homomorphic evaluation of $f^{(r)}$ with error at most εk^r. By Lemma 25, $\mathbf{Rer}'^{(r)}$ has error $k^r \varepsilon + \sqrt{2c/t^r}$. Therefore $\mathbf{Rer}^{(r)}$ has error at most $\varepsilon(k^{r-1} + k^r) + \sqrt{2c/t^r}$. Let $\alpha = \log t / \log k$. By choosing $r = 1/(2 + \alpha) \cdot \log(2c/\varepsilon^2)/\log k$ we get that $\mathbf{Rer}^{(r)}$ has error $O(\varepsilon^{\alpha/(2+\alpha)})$, which is negligible when ε is negligible. □

Acknowledgments. We thank Benny Applebaum for explaining the relevance of the work of Mahmoody and Xiao, Alon Rosen for helpful discussions about the power of our rerandomization procedure, and Cheuk Ting Li for supplying a counterexample to a conjectured generalization of Lemma 25.

References

AGGM06. Adi Akavia, Oded Goldreich, Shafi Goldwasser, and Dana Moshkovitz. On basing one-way functions on NP-hardness. In *Proceedings of the 38th ACM Symposium on Theory of Computing*, 2006.

Ajt96. M. Ajtai. Generating hard instances of lattice problems (extended abstract). In *Proceedings of the 28th ACM Symposium on Theory of Computing*, STOC '96, pages 99–108, New York, NY, USA, 1996. ACM.

BBM11. Nayantara Bhatnagar, Andrej Bogdanov, and Elchanan Mossel. The computational complexity of estimating convergence time. In *Proceedings of the 15th International Workshop on Randomization and Computation (RANDOM)*, 2011.

Bra79. Gilles Brassard. Relativized cryptography. In *Proceedings of the 20th IEEE Symposium on Foundations of Computer Science*, pages 383–391, 1979.

BT06. Andrej Bogdanov and Luca Trevisan. On wost-case to average-case reductions for NP problems. *SIAM J. Comp.*, 36(4), 2006.

BV11. Z. Brakerski and V. Vaikuntanathan. Efficient Fully Homomorphic Encryption from (Standard) LWE. In *Proceedings of the 53rd Annual Symposium on Foundations of Computer Science*, 2011.

EY80. Shimon Even and Yacob Yacobi. Cryptography and NP-completeness. In *Proceedings of the 7th ICALP*, volume 85 of *LNCS*, pages 195–207. Springer-Verlag, 1980.

FF93. Joan Feigenbaum and Lance Fortnow. Random-self-reducibility of complete sets. *SIAM Journal on Computing*, 22:994–1005, 1993.

Gam85. T. El Gamal. A public-key cryptosystem and a signature scheme based on discrete logarithms. *IEEE Trans. Info. Theory*, 31(4):469–472, 1985.

Gen09. Craig Gentry. Fully Homomorphic Encryption Using Ideal Lattices. In *STOC*, pages 169–178, 2009.

GG98. Oded Goldreich and Shafi Goldwasser. On the possibility of basing cryptography on the assumption that $P \neq NP$. Unpublished manuscript, 1998.

Gol00. Oded Goldreich. Candidate one-way functions based on expander graphs. *Electronic Colloquium on Computational Complexity (ECCC)*, 7(090), 2000.

HMX10. I. Haitner, M. Mahmoody, and D. Xiao. A new sampling protocol and applications to basing cryptographic primitives on np. In *Proceeedings of 25th IEEE Conference on Computational Complexity (CCC)*, pages 76–87, 2010.

MX10. M. Mahmoody and D. Xiao. On the power of randomized reductions and the checkability of sat. In *Proceeedings of 25th IEEE Conference on Computational Complexity (CCC)*, pages 64–75, 2010.

Pai99. P. Paillier. Public-key cryptosystems based on composite degree residuosity classes. In *Advances in Cryptology – Eurocrypt '99*, pages 223–238, 1999.

Pei09. Chris Peikert. Public-key cryptosystems from the worst-case shortest vector problem: extended abstract. In *Proceedings of the 41th ACM Symposium on Theory of Computing*, pages 333–342, New York, NY, USA, 2009. ACM.

Pin64. M. S. Pinsker. *Information and information stability of random variables and processes*. Holden-Day, 1964.

Reg09. Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6), 2009.

Rot11. Ron Rothblum. Homomorphic encryption: From private-key to public-key. In *Proceedings of the Theory of Cryptography Conference (TCC)*, pages 219–234, 2011.

Sim82. Hans-Ulrich Simon. A tight $\log\log n$-bound on the time for parallel ram's to compute nondegenerated boolean functions. *Information and Control*, 55(1):102 – 107, 1982.

SV03. A. Sahai and S. Vadhan. A complete problem for statistical zero knowledge. *J. ACM*, 50:196–249, 2003.

vDGHV10. M. van Dijk, C. Gentry, S. Halevi, and V. Vaikuntanathan. Fully Homomorphic Encryption from Integers. In *Eurocrypt*, 2010.