

A Cryptographic Test of Quantumness and Certifiable Randomness from a Single Quantum Device

Zvika Brakerski* Paul Christiano† Urmila Mahadev‡ Umesh Vazirani§
Thomas Vidick¶

Abstract

We give a protocol for producing certifiable randomness from a single untrusted quantum device that is polynomial-time bounded. The randomness is certified to be statistically close to uniform from the point of view of any computationally unbounded quantum adversary, that may share entanglement with the quantum device. The protocol relies on the existence of post-quantum secure trapdoor claw-free functions, and introduces a new primitive for constraining the power of an untrusted quantum device. We then show how to construct this primitive based on the hardness of the learning with errors (LWE) problem.

The randomness protocol can also be used as the basis for an efficiently verifiable “quantum supremacy” proposal, thus answering an outstanding challenge in the field.

*Weizmann Institute of Science, Israel. Email: zvika.brakerski@weizmann.ac.il.

†OpenAI, USA. Work performed while at UC Berkeley

‡UC Berkeley, USA. Email: mahadev@berkeley.edu

§UC Berkeley, USA. Email: vazirani@cs.berkeley.edu

¶California Institute of Technology, USA. Email: vidick@cms.caltech.edu

1 Introduction

In this paper we propose solutions to two basic tasks: how to generate *certifiably random* strings from a *single untrusted* quantum device (also referred to as a prover), and efficient verification of quantum supremacy. The setting we consider is one where the quantum device is polynomial-time bounded but untrusted, and the verifier is entirely classical and also polynomial-time bounded. The peculiarity of this setting is that it allows the verifier to leverage post-quantum cryptography, i.e. the existence of cryptographic primitives that can be implemented efficiently on a classical computer but that cannot be broken by any efficient quantum computer.

There has been considerable research into certifiable quantum random number generation [Col06, PAM⁺10, FGS11, PM11, VV11, MS16, BKG⁺18]. However, all prior works providing verifiable guarantees have focused on the setting where there are multiple quantum devices that share entanglement, and where the randomness certification relies crucially on the violation of a Bell inequality. By contrast, in the setting that we study there is a single polynomial-time bounded quantum device, and the guarantee we seek is that provided the device is unable to break the post-quantum cryptographic assumption during the execution of the protocol, then the output of the protocol must be statistically indistinguishable from a uniformly random sequence of bits, to any computationally *unbounded* adversary that may share prior entanglement with the computationally bounded device. This information-theoretic guarantee, the same guarantee as that offered in the aforementioned works [VV11, MS16], is stronger than computational pseudo-randomness (that is easily achievable under standard cryptographic assumptions, since the verifier starts with a short random seed).

The specific cryptographic primitive we rely on is the existence of a post-quantum secure trapdoor claw-free (in short, TCF) family of functions $f : \{0,1\}^n \rightarrow \{0,1\}^m$, the post-quantum analog of a notion introduced by Goldwasser, Micali and Rivest in the context of digital signatures [GMR84]. A TCF is a 2-to-1 function f that satisfies the following properties: $f(x)$ is efficiently computable on a classical computer, and if $f(x) = y$, then there is a unique $x' \neq x$ such that $f(x') = y$. Moreover, with knowledge of a secret trapdoor it is possible to efficiently (classically) compute x and x' from y , but without the trapdoor there is no efficient quantum algorithm that can compute such a triple (x, x', y) , for any y .

By contrast, a quantum algorithm can simultaneously hold an image y , as well as a superposition $\frac{1}{\sqrt{2}}(|x\rangle + |x'\rangle)$ over two pre-images of y , simply by evaluating f on a uniform superposition over all inputs and measuring the image y . At this point, measuring the quantum state in the standard basis will yield a random pre-image, x or x' . This is not any stronger than a classical device could do, by first sampling a random x and then computing $y = f(x)$. However, the quantum device can do something different from directly measuring a pre-image. Instead, the device can perform Fourier sampling (Hadamard transform followed by a measurement), which yields a random n -bit string $d : d \cdot (x \oplus x') = 0$, thereby revealing some joint information about both pre-images of y .

The ability to perform *either* of these tasks, identifying a pre-image or sampling an equation, but not both, seems to be a unique quantum capability. But of course a classical verifier cannot access the inner workings of the quantum device: from her viewpoint the device is an untrusted black box that outputs y , and then either a pre-image x , or an equation d . However, assuming the verifier knows the secret trapdoor, she can efficiently compute both pre-images x and x' , and verify that indeed $d \cdot (x \oplus x') = 0$. At a high level, the consideration of a trapdoor family restores some symmetry between the quantum prover (the untrusted quantum device) and the classical verifier, by providing a primitive which allows the quantum capabilities of the prover to play a useful role while at the same time giving the classical verifier a handle, namely the ability to compute both x and x' , that the prover does not have access to.

One might now conjecture that for a generic TCF (e.g. modeled as a random oracle), if the output of any efficient quantum device passes this test with non-negligible advantage over $\frac{1}{2}$, then the pair y, d returned in the equation test must have high min-entropy. Such a strong statement would immediately yield a randomness certification protocol. The difficulty in showing such a statement is that both y and d may be adaptively and adversarially chosen. While we do not know how to prove this, we are still able to design and analyze a simple protocol, that can be based on any TCF family with suitable cryptographic properties, and succeeds in certifiable randomness expansion from a single untrusted quantum (polynomial-time) device.

We will refer to protocol described above as a single round test, and our certified randomness expansion protocol will consist of a large number of such rounds. To facilitate the analysis of the protocol, we elaborate on a specific TCF construction in [Mah17a] based on the learning with errors problem (LWE). Specifically, we achieve a construction of TCF that has additional security guarantees, including an “adaptive hardcore bit” condition that is explained below. This condition helps in going beyond the analysis of a single round of our randomness-generation protocol (this analysis can be performed based on a more generic requirement for the TCF) to guarantee randomness accumulation across multiple rounds, a more stringent requirement. For clarity, in this paper we refer to the specific (relaxed) kind of TCF that we rely on as an NTCF, or post-quantum secure noisy trapdoor claw-free family. We give a construction of an NTCF that rests on the hardness of the Learning With Errors (LWE) problem, introduced by Regev [Reg05], with slightly super-polynomial noise ratio, against quantum polynomial-time attacks with nonuniform quantum advice (where the state of the art classical and quantum complexity scale exponentially with the dimension). This construction is similar to the one used in [Mah17a], albeit with some changes in parameters that allow us to prove the adaptive hardcore bit statement. Roughly, it asserts that it is computationally intractable to sample from any distribution on (y, x, d, b) such that $f(x) = y$ and $\Pr(b = d \cdot (x + x')) \geq 1/2 + \varepsilon$ for non-negligible ε .

We now note that the adaptive hardcore bit statement already has an interesting consequence: passing the single round test constitutes a proof of quantumness of the device. This is because there cannot be any efficient classical algorithm that can reliably answer a random challenge, “pre-image” or “equation”, since such an algorithm could be “rewound” to simultaneously answer both challenges, thus contradicting the adaptive hardcore bit property. This has implications for an important milestone in the experimental realization of quantum computers, namely “quantum supremacy”: a proof that an (untrusted) quantum computing device performs some computational task that cannot be solved classically without impractical resources. While quantum supremacy could be achieved, in principle, by demonstrating quantum factoring, the latter requires quantum resources well beyond the capability of near term experiments. Instead current proposals for quantum supremacy are based on sampling problems (see e.g. [HM17] for a recent survey). The major challenge for these proposals is verifying that the quantum computer did indeed sample from the desired probability distribution, and indeed all existing proposals rely on exponential time classical algorithms for verification. By contrast, our single round test provides a proof of quantumness and can be carried out by a classical verifier in polynomial time. This proposal for quantum supremacy seems promising from a practical viewpoint — indeed, even using off-the-shelf bounds for LWE-based cryptography suggests that a protocol providing 50 bits of security could be implemented with a quantum device of around 2000 qubits (see e.g. [LP11]). It would be worth exploring whether there are clever implementations of this scheme that can lead to a quantum supremacy protocol in the 200 – 500 qubit range. Our protocol is robust to a device that only successfully answers the verifier’s challenges a sufficiently large, but constant, fraction of the time; it would be interesting to explore whether such a device could be implemented without resorting to general fault-tolerance techniques.

Using the single round test as a building block for certified randomness expansion requires a deeper

analysis. Let us focus on a single bit of information that the device outputs in the test: whether the pre-image is x or x' and whether the equation d satisfies $d \cdot (x + x') = 0$ or not. We are able to characterize the quantum state and quantum measurements for this task that are derived from the measurements of any efficient device that passes the single round test. Specifically, we show that the state and measurements of the device are close to the following: the device starts with a qubit initialized to $|0\rangle$, which it measures in the standard basis when the challenge is $C = 0$ and in the Hadamard basis when the challenge is $C = 1$. Note that the fact that an efficient quantum device cannot break the cryptographic assumption has thus been translated into a characterization of the state and actions of the quantum device, which further implies that the output of the device in this single round must contain almost a bit of true (information theoretic) randomness.

Achieving certifiable randomness expansion requires an analysis of a protocol with a long sequence of single round tests, which we describe in the next section. We note that while assuming that LWE is hard for polynomial-time adversaries already implies polynomial randomness expansion, we can stretch the parameters further to achieve exponential randomness expansion by assuming that the problem remains hard even against quantum sub-exponential time adversaries, which is still consistent with the current state of the art.

Recently Aaronson [Aar18] discovered a different connection between quantum supremacy proposals and randomness generation. Aaronson’s scheme has the advantage of being as easy to implement as the supremacy proposal. Unfortunately, it also faces the same limitation in terms of verifiability as existing supremacy proposals, with the verification time scaling exponentially with the number of generated bits.

The idea of using a TCF as a basic primitive in interactions between an efficient quantum prover and a classical verifier has been further developed in recent work by Mahadev [Mah17b], giving the first construction of a quantum fully homomorphic encryption scheme with classical keys. In further follow-up work, Mahadev [Mah18] shows a remarkable use of a NTCF family with adaptive hardcore bit. Namely, that the NTCF can be used to certify that a prover measures a qubit in a prescribed basis (standard or Hadamard). This allows to achieve single prover *verifiability* for quantum computations using a purely classical verifier (but relying on computational assumptions).

Independently of this work, a construction of trapdoor one-way functions with second pre-image resistance based on LWE was recently introduced in [CCKW18], where it is used to achieve delegated computation in the weaker honest-but-curious model for the adversary (i.e. without soundness against provers not following the protocol). The family of functions considered in [CCKW18] is not sufficient for our purposes, as it lacks the essential adaptive hardcore bit property.

We believe that the technique of constraining the power of a quantum device using NTCFs, and with the adaptive hardcore bit put forth in this work promises to be a powerful tool for the field of untrusted quantum devices.

Overview of results and proof techniques

We provide a high level description of the main components of our protocol and the proof of security.

A single round of the protocol: test of quantumness. At the computational heart of our protocol is the single round test. In each round, the device is first required to output an element y in the range of NTCF f . The classical verifier then issues one of two challenges, $C = 0$ (“pre-image”) or $C = 1$ (“equation”). The

device responds by providing either a nontrivial $d : d \cdot (x \oplus x') = 0$ ¹ or a valid pre-image for y , respectively. Both conditions can be efficiently checked by the verifier, the equation check requires the verifier to use the trapdoor for f . An honest quantum device can accomplish this by creating a uniform superposition on all x , computing f in superposition, and then measuring $y = f(x)$. This projects the x -register into the superposition $(1/\sqrt{2})|x\rangle + (1/\sqrt{2})|x'\rangle$, where x, x' are the preimages of y . Now measuring the x -register in the standard basis yields a random preimage x or x' , and measuring it in the Hadamard basis yields a random equation $d : d \cdot (x \oplus x') = 0$.

The fact that passing the single round test constitutes a proof of quantumness of the device relies on the adaptive hardcore bit property, which we sketch here and explain in greater detail later. Roughly, it asserts that it is computationally intractable to sample from any distribution on (y, x, d, b) such that $f(x) = y$ and $\Pr(b = d \cdot (x \oplus x')) \geq 1/2 + \varepsilon$ for non-negligible ε . This property implies that any efficient device that can reliably answer a random challenge (equation or pre-image) must be quantum. This is because any classical algorithm could be “rewound” to simultaneously answer both challenges, which is impossible by the adaptive hard core predicate property.

At the heart of certifiable randomness generation lies the claim that if an efficient device passes the single round test, then with high probability the state of the device and its measurements, up to some coarse-graining, are close to the following: it starts with a qubit initialized to $|0\rangle$, which it measures in the standard basis when the challenge is $C = 0$ and in the Hadamard basis when the challenge is $C = 1$. This characterization of the quantum state of the device and its measurements is where the computational bound on the device translates into true randomness rather than pseudorandomness.

To see how the proof proceeds, notice that in a single round the device must make one of two measurements: either a “pre-image” measurement, or an “equation” measurement. The “pre-image” measurement can be treated as a projection into one of two orthogonal subspaces corresponding to the two pre-images x, x' for the element y that the device has returned to the verifier. The “equation” measurement can similarly be coarse-grained into a projection on one of two orthogonal subspaces, “valid” or “invalid”, i.e. the subspace that corresponds to all measurement outcomes d such that $d \cdot (x \oplus x') = 0$, or the subspace associated with outcomes d such that $d \cdot (x \oplus x') = 1$.

Applying Jordan’s lemma, it is possible to decompose the device’s Hilbert space into a direct sum of one- or two-dimensional subspaces, such that within each two-dimensional subspace the “pre-image” and “equation” measurements each correspond to an orthonormal basis, such that the two bases make a certain angle with each other. We argue that almost all angles must be very close to $\pi/4$. Indeed, whenever the angles are *not* near-maximally unbiased, it is possible to show that by considering the effect of performing the measurements in sequence, one can devise an “attack” on the NTCF of a kind that contradicts the adaptive hardcore bit property of the NTCF — informally, the attack can simultaneously produce a valid pre-image and a valid equation, with non-negligible advantage.

Outline of randomness generation protocol. The correctness of our protocol relies on the claim that in a single round, if an efficient quantum algorithm has the ability to generate a valid equation with probability sufficiently close to 1, then, if instead it is asked for a pre-image, this pre-image must be close to uniformly distributed over the two possibilities. Therefore, our protocol repeatedly asks for pre-images (to generate randomness), while inserting a few randomly located “equation” challenges, to test the device. Each time an “equation” challenge has been answered, we refresh the pseudorandom keys used for the NTCF. This is required to avoid a simple “attack” by the device, which would repeatedly use the same y , preimage x , and

¹The exact condition for the equation is slightly more complicated, due to the formal definition of a NTCF; see Section 4 for a complete description of the protocol.

guessed equation d — succeeding in the protocol with probability $\frac{1}{2}$ without generating any randomness.

Let’s call the sequence of rounds with a particular set of pseudorandom keys an epoch. Intuitively, we would like to claim that if the device passes all the equation challenges, then for most epochs and for most rounds within that epoch, the state of the device and its measurements must be (close to) as characterized above: it starts with a qubit initialized to $|0\rangle$, which it measures in the standard basis when the challenge is $C = 0$ and in the Hadamard basis when the challenge is $C = 1$. To show this we would like to claim that if the device passes all the equation challenges, for most such challenges it must produce a valid equation with probability close to 1. Since each equation challenge occurs at a random round in the epoch, it should follow from the adaptive hardcore bit property that the sequence of bits that the verifier extracts from the device’s answers to “pre-image” challenges during that epoch must look statistically random. We give a martingale based argument to formalize this intuition, though the quantum setting makes it considerably more challenging than it appears at first sight.

There is however a bigger challenge to analyzing the protocol — we must show that the sequence that the verifier extracts from the device’s answers to “pre-image” challenges must look statistically random even to an infinitely powerful quantum adversary, who may share an arbitrary entangled state with the quantum device. If we could assert that each round of the protocol is played with a qubit exactly in state $|0\rangle$, and measured in Hadamard basis on challenge $C = 1$, then this would preclude any entanglement with the quantum adversary, leading to an easy proof that the extracted sequence looks random to the adversary. Unfortunately the characterization of device’s qubits leaves plenty of room for entanglement with the adversary — showing that such entanglement cannot leak too much information about the device’s measurements was the major challenge in previous work on certified randomness through Bell inequality violations [VV11, MS14, AFDF⁺18]. Our cryptographic setting presents a new difficulty, which is that in contrast to the two-player game Bell inequality violation scenarios, in our setting it is not *impossible* for a deterministic device to succeed in the test: it is merely *computationally hard* to do so. This prevents us from directly applying the results in [AFDF⁺18, MS14], and requires us to suitably modify their framework.

In terms of efficiency, for the specific LWE-based NTCF that we construct, our protocol can use as few as $\text{poly}(\log(N))$ bits of randomness to generate $O(N)$ bits that are statistically within negligible distance from uniform. However, this requires assuming that the underlying LWE assumption is hard even for sub-exponential size quantum circuits, with polynomial-size quantum advice (which is consistent with current knowledge). The more conservative assumption that our variant of LWE is only hard for polynomial size quantum circuits requires $O(N^\epsilon)$ bits of randomness for generating the NTCF, for any constant $\epsilon > 0$. The following is an informal description; see Theorem 25 for a more formal statement.

Theorem 1 (Informal). *Let \mathcal{F} be a NTCF family and λ a security parameter. Let $N = \Omega(\lambda^2)$ and assume the quantum hardness of solving lattice problems of dimension λ in time $\text{poly}(N)$. There is an N -round protocol for the interaction between a classical polynomial-time verifier and a quantum polynomial-time device such that the protocol can be executed using $\text{poly}(\log(N), \lambda)$ bits of randomness, and for any efficient device and side information E correlated with the device’s initial state,*

$$H_\infty^{N^\delta}(O|CE)_{\bar{p}} \geq (\xi - o(1))N.$$

Here ξ is a positive constant, δ is a negligible function of λ , and \bar{p} is the final state of the classical output register O , the classical register C containing the verifier’s messages to the device, and the side information E , restricted to transcripts that are accepted by the verifier in the protocol.

Sketch of the security analysis. We first describe the protocol in slightly more detail (see Section 4 for a formal description). The verifier first uses $\text{poly}(\log(N), \lambda)$ bits of randomness to select a pair of functions

$\{f_{k,b}\}_{b \in \{0,1\}}$ from an NTCF family, and sends the public function key k to the quantum device. This pair of functions can be interpreted as a single 2-to-1 function $f_k : (b, x) \mapsto f_{k,b}(x)$. The verifier keeps private the trapdoor information that allows to invert f_k . The protocol then proceeds for N rounds. In each round the device first outputs a point y in the common range of $f_{k,0}$ and $f_{k,1}$. After having received y , the verifier issues one of two challenges: 0 or 1 — pre-image or equation. If the challenge is “pre-image”, then the device must output an x such that $f(x) = y$. If the challenge is “equation” then the device must output a nontrivial binary vector d such that $d \cdot (x_0 + x_1) = 0$, where x_0 and x_1 are the unique pre-images of y under $f_{k,0}$ and $f_{k,1}$ respectively. Since the verifier has the secret key, she can efficiently compute x_0 and x_1 from y , and therefore check the correctness of the device’s response to each challenge. The verifier chooses $\text{poly} \log(N)$ rounds in which to issue the challenge 1, or “equation”, at random. Selecting these rounds requires only $\text{poly} \log(N)$ random bits. At the end of each such round, the verifier samples a new pair of functions from the NTCF family, and communicates the new public key to the device. On each of the remaining $N - \text{poly} \log(N)$ rounds the verifier records a random bit according to whether the device returns the pre-image x_0 , or x_1 (e.g. recording 0 for the lexicographically smaller pre-image). At the end of the protocol the verifier uses a strong quantum-proof randomness extractor to extract $\Omega(N)$ bits of randomness from the recorded string (this requires at most an additional $\text{poly} \log(N)$ random bits of seed).

To guarantee that the extractor produces bits that are statistically close to uniform, we would like to prove that the $N - \text{poly} \log(N)$ random bits recorded by the verifier must have $\Omega(N)$ bits of (smoothed) min-entropy,² even conditioned on the side information available to an infinitely powerful quantum adversary, who may share an arbitrary entangled state with the quantum device. The analysis proceeds as follows. First we assume without loss of generality that the entire protocol is run coherently, i.e. we may assume that the initial state of the quantum device (holding quantum register \mathbf{D}) and the adversary (holding quantum register \mathbf{E}) is a pure state $|\phi\rangle_{\mathbf{DE}}$, since the adversary may as well start with a purification of their joint state. We may also assume that the verifier starts with a cat state on $\text{poly} \log(N)$ qubits, and uses one of the registers of the state, \mathbf{C} , to provide the random bits used to select the test rounds and to issue the challenges in those rounds. (This is for the sake of analysis only, the actual verifier is of course completely classical.) We can similarly arrange that the state remains pure throughout the protocol by using the principle of deferred measurement. Our goal is to show a lower bound on the smooth min-entropy of the output register \mathbf{O} in which the verifier has recorded the device’s outputs, conditioned on the state \mathbf{E} of the adversary, and on the register \mathbf{C} of the cat state (conditioning on the latter represents the fact that the verifier’s choice of challenges may be leaked to the adversary, and we would like security even in this scenario). Intuitively, this amounts to bounding the information accessible to the most powerful adversary quantum mechanics allows, conditioned on the joint state of the verifier and device.

In order to bound the entropy of the final state we need to show that the entropy “accumulates” at each round of the protocol. A general framework to establish entropy accumulation in quantum protocols such as the one considered here was introduced in [AFDF⁺18]. At a high level, the approach consists in reducing the goal of a min-entropy bound to a bound on the appropriate notion of $(1 + \varepsilon)$ quantum conditional Rényi entropy, and then arguing that, under suitable conditions on the process that generates the outcomes recorded in the protocol, entropy accumulates sequentially throughout the protocol.

In a little more detail, the first step on getting a handle on the smooth min-entropy is to use the quantum asymptotic equipartition property (QAEP) [TCR09] to relate it to the $(1 + \varepsilon)$ conditional Rényi entropy, for suitably small ε . The second step uses a duality relation for the conditional Rényi entropy to relate the $(1 + \varepsilon)$ conditional Rényi entropy of the output register \mathbf{O} , conditioned on the adversary side information in \mathbf{R} and the register \mathbf{C} of the cat state, to a quantity analogous to the $(1 - \varepsilon')$ conditional Rényi entropy

²We refer to Section 2 for definitions of entropic quantities.

of the output register, conditioned on the register E for the device, and a purifying copy of the register C of the cat state. The latter quantity, a suitable conditional entropy of the output register conditioned on the challenge register and the state of the device, is the quantity that we ultimately aim to bound. Note what these transformations have achieved for us: it is now sufficient to consider as side information only “known” quantities in the protocol, the verifier’s choice of challenges and the device’s state; the information held by the adversary plays no other role than that of a purifying register.

As mentioned earlier, our cryptographic setting presents the additional difficulty that our guarantee is only that it is computationally hard for a deterministic device to succeed in the protocol. The results in [AFDF⁺18, MS14] crucially rely on the fact that the process that generates the randomness does so irrespective of the quantum state in which it is initialized (as long as the output of the process satisfies the test’s success criterion). This requirement comes about from conditioning that is performed in order to show that entropy accumulates; in our setting, conditioning is more delicate as it can in principle induce non-computationally efficient states for the device.

Recall that we argued that for a single round of the protocol, we can decompose the device’s Hilbert space into a direct sum of one- or two-dimensional subspaces, such that within most two-dimensional subspace the “pre-image” and “equation” measurements correspond to orthonormal bases that make an angle close to $\pi/4$ with each other. Showing that the Rényi entropy accumulates in each round requires a device in which *all* angles are close to $\pi/4$, not “almost all”. To accommodate for this we “split” the state of the device into its component on the good subspace, where the angles are unbiased, and the bad subspace, where the measurements may be aligned. The fact that the distinction between good and bad subspace is not measured in the protocol, but is only a distinction made for the analysis, requires us to apply a fairly delicate martingale based argument that takes into account possible interference effects and bounds those “branches” where the state has gone through the bad subspace an improbably large number of times. Whenever the state lies in the good subspace, we can appeal to an uncertainty principle from [MS14] to show that the device’s measurement increases the conditional Rényi entropy of the output register by a small additive constant. Pursuing this approach across all N rounds, we obtain a linear lower bound on the conditional Rényi entropy of the output register, conditioned on the state of the device. As argued above this in turn translates into a linear lower bound on the smooth conditional min-entropy of the output, conditioned on the state of the adversary and the verifier’s choice of challenges. It only remains to apply a quantum-proof randomness extractor to the output, using a poly-logarithmic number of additional bits of randomness, to obtain the final result.

Our NTCF Family. Our goal is to construct a family of pairs of injective functions f_0, f_1 with the same image such that it is hard to find a collision x_0, x_1 with $f_0(x_0) = f_1(x_1)$, but so that given a suitable trapdoor it is possible to recover, for any y , values x_0, x_1 such that $f_0(x_0) = f_1(x_1) = y$. We do this by relying on the hardness of the Learning with Errors (LWE) problem [Reg05]. LWE states that given a public uniformly random matrix $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ for $m \gg n$, it is intractable to distinguish between $\mathbf{u} = \mathbf{A}\mathbf{s} + \mathbf{e} \pmod{q}$ and a uniform vector, for a uniform vector \mathbf{s} and small discrete Gaussian vector \mathbf{e} (all arithmetic from here on is performed modulo q ; we use \oplus to denote binary XOR). Inspired by [Mah17a], our function pair will be characterized by $(\mathbf{A}, \mathbf{u} = \mathbf{A}\mathbf{s} + \mathbf{e})$, but for a *binary* vector \mathbf{s} .³ The trapdoor for our function is a lattice trapdoor for \mathbf{A} that allows to recover \mathbf{s}, \mathbf{e} given a vector of the form $\mathbf{A}\mathbf{s} + \mathbf{e}$ (it is possible to generate \mathbf{A} together with a trapdoor such that \mathbf{A} is indistinguishable from uniform, as originally shown by Ajtai [Ajt99]).

³It is known that LWE is hard even with binary secrets. We do not use this property explicitly but rather employ the respective techniques in our proof.

The structure of the LWE problem motivates us to consider functions f_0, f_1 that range over probability distributions. Specifically, we define the distribution $f_b(\mathbf{x})$ as $f_b(\mathbf{x}) = \mathbf{A}\mathbf{x} + b \cdot (\mathbf{A}\mathbf{s}) + \mathbf{e}'$ where \mathbf{e}' is a discrete Gaussian random variable with a sufficiently wide Gaussian parameter. Observe that the functions have overlapping images in the following sense: $f_0(\mathbf{x}) = f_1(\mathbf{x} - \mathbf{s})$. Moreover, since $\mathbf{A}\mathbf{s} + \mathbf{e}'$ is statistically indistinguishable from $\mathbf{u} + \mathbf{e}'$, we can efficiently sample from the distribution $f_b(\mathbf{x})$ up to negligible statistical distance. This probabilistic notion complicates the definition and use of the family, but the principles are similar to the deterministic version.⁴

Finally, we need to show the adaptive hardcore bit property. Formally, letting $\mathbf{x}_0, \mathbf{x}_1 \in \mathbb{Z}_q^n$ be a collision, and letting $z_0, z_1 \in \{0, 1\}^{n \log q}$ be their binary representations, respectively, we need to show that it is intractable to come up with a pair (b, z_b) , for some $b \in \{0, 1\}$, together with a nontrivial vector d and with the value $d \cdot (z_0 \oplus z_1)$, with probability noticeably better than $\frac{1}{2}$. “Nontrivial” here means belonging to a well defined and efficiently recognizable set D with density ≈ 1 in $\{0, 1\}^{n \log q}$ (e.g. the zero vector is obviously excluded). Assume for the sake of this overview that we get a tuple (z_0, d, c) . We first notice that since $\mathbf{x}_0, \mathbf{x}_1$ is a collision, then $\mathbf{x}_1 = \mathbf{x}_0 - \mathbf{s} \pmod{q}$. We now use the fact that \mathbf{s} is a binary vector to show, using simple arithmetic, that $z_0 \oplus z_1$ can be expressed as a linear function of the bits of \mathbf{s} , so that $d \cdot (z_0 \oplus z_1) = \hat{d} \cdot \mathbf{s} \pmod{2}$, for some $\hat{d} \in \{0, 1\}^n$. (the description of this transformation will effect our choice of the set D). We thus need to show that it is intractable, given the instance $(\mathbf{A}, \mathbf{u} = \mathbf{A}\mathbf{s} + \mathbf{e})$, to come up with $\hat{d}, \hat{d} \cdot \mathbf{s} \pmod{2}$. To prove this we use the lossiness technique used in [GKPV10] and show that this is equivalent to coming up with $\hat{d}, \hat{d} \cdot \mathbf{s} \pmod{2}$ given $\mathbf{B}, \mathbf{B}\mathbf{s}$ where $\mathbf{B} \in \mathbb{Z}_q^{k \times n}$ is now a highly shrinking function, even for binary inputs, i.e. $k \log q \ll n$. This seems like an easy task since the adversary now doesn’t have the complete information about \mathbf{s} so it shouldn’t be able to compute $\hat{d} \cdot \mathbf{s} \pmod{2}$ for any reasonable \hat{d} , except \hat{d} might depend on \mathbf{B} itself (recall that \hat{d} is chosen adversarially). We prove via Fourier analysis that if \mathbf{B} is sufficiently shrinking, then there is no \hat{d} that can take advantage of the dependence on \mathbf{B} , which completes the proof.

Organization. We start with some notation and preliminaries in Section 2. Section 3 contains the definition of a noisy trapdoor claw-free family (NTCF). Our construction for such a family is given in Section 8 (with Appendix A containing relevant preliminaries on the learning with errors problem). The randomness generation protocol is described in Section 4. In Section 5 we introduce our formalism for modeling the actions of an arbitrary prover, or device, in the protocol. In Section 6 we analyze a single round of the protocol, and in Section 7 we show that randomness accumulates across multiple rounds.

Acknowledgments. Zvika Brakerski is supported by the Israel Science Foundation (Grant No. 468/14), Binational Science Foundation (Grants No. 2016726, 2014276), and by the European Union Horizon 2020 Research and Innovation Program via ERC Project REACT (Grant 756482) and via Project PROMETHEUS (Grant 780701). Paul Christiano and Urmila Mahadev are supported by a Templeton Foundation Grant 52536, ARO Grant W911NF-12-1-0541, and NSF Grant CCF-1410022. Umesh Vazirani is supported by Templeton Foundation Grant 52536, ARO Grant W911NF-12-1-0541, NSF Grant CCF-1410022, MURI Grant FA9550-18-1-0161 and a Vannevar Bush Faculty Fellowship. Thomas Vidick is supported by NSF CAREER Grant CCF-1553477, AFOSR YIP award number FA9550-16-1-0495, MURI Grant FA9550-18-1-0161, a CIFAR Azrieli Global Scholar award, and the IQIM, an NSF Physics Frontiers Center (NSF Grant PHY-1125565) with support of the Gordon and Betty Moore Foundation (GBMF-12500028).

⁴Another possible variant is to define $f_b(\mathbf{x}) = \lfloor \mathbf{A}\mathbf{x} + b \cdot (\mathbf{A}\mathbf{s}) \rfloor$ where $\lfloor \cdot \rfloor$ is a rounding function that truncates “many” of the least significant bits of its operand. However, we remain with the Gaussian variant which is easier to analyze.

2 Preliminaries

2.1 Notation

\mathbb{Z} is the set of integers, and \mathbb{N} the set of natural numbers. For any $q \in \mathbb{N}$ such that $q \geq 2$ we let \mathbb{Z}_q denote the ring of integers modulo q . We generally identify an element $x \in \mathbb{Z}_q$ with its unique representative $[x]_q \in (-\frac{q}{2}, \frac{q}{2}] \cap \mathbb{Z}$. For $x \in \mathbb{Z}_q$ we define $|x| = |[x]_q|$. When considering an $s \in \{0, 1\}^n$ we sometimes also think of s as an element of \mathbb{Z}_q^n , in which case we write it as \mathbf{s} .

We use the terminology of polynomially bounded and negligible functions. A function $n : \mathbb{N} \rightarrow \mathbb{R}_+$ is *polynomially bounded* if there exists a polynomial p such that $n(\lambda) \leq p(\lambda)$ for all $\lambda \in \mathbb{N}$. A function $n : \mathbb{N} \rightarrow \mathbb{R}_+$ is *negligible* if for every polynomial p , $p(\lambda)n(\lambda) \rightarrow_{\lambda \rightarrow \infty} 0$. We write $\text{negl}(\lambda)$ to denote an arbitrary negligible function of λ . For two parameters κ, λ we write $\kappa \ll \lambda$ to express the constraint that κ should be “sufficiently smaller than” λ , meaning that there exists a small universal constant $c > 0$ such that $\kappa \leq c\lambda$.

\mathcal{H} always denotes a finite-dimensional Hilbert space. We use indices $\mathcal{H}_A, \mathcal{H}_B$, etc., to refer to distinct spaces. $\text{Pos}(\mathcal{H})$ is the set of positive semidefinite operators on \mathcal{H} , and $\text{D}(\mathcal{H})$ the set of density matrices, i.e. the positive semidefinite operators with trace 1. For an operator X on \mathcal{H} , we use $\|X\|$ to denote the operator norm (largest singular value) of X , and $\|X\|_{tr} = \frac{1}{2}\|X\|_1 = \frac{1}{2}\text{Tr}\sqrt{XX^\dagger}$ for the trace norm.

2.2 Distributions

We generally use the letter D to denote a distribution over a finite domain X , and f for a density on X , i.e. a function $f : X \rightarrow [0, 1]$ such that $\sum_{x \in X} f(x) = 1$. We often use the distribution and its density interchangeably. We write U for the uniform distribution. We write $x \leftarrow D$ to indicate that x is sampled from distribution D , and $x \leftarrow_U X$ to indicate that x is sampled uniformly from the set X . We write \mathcal{D}_X for the set of all densities on X . For any $f \in \mathcal{D}_X$, $\text{SUPP}(f)$ denotes the support of f ,

$$\text{SUPP}(f) = \{x \in X \mid f(x) > 0\}.$$

For two densities f_1 and f_2 over the same finite domain X , the Hellinger distance between f_1 and f_2 is

$$H^2(f_1, f_2) = 1 - \sum_{x \in X} \sqrt{f_1(x)f_2(x)}. \quad (1)$$

The total variation distance between f_1 and f_2 is

$$\|f_1 - f_2\|_{TV} = \frac{1}{2} \sum_{x \in X} |f_1(x) - f_2(x)| \leq \sqrt{2H^2(f_1, f_2)}. \quad (2)$$

The following immediate lemma relates the Hellinger distance and the trace distance of superpositions.

Lemma 2. *Let X be a finite set and $f_1, f_2 \in \mathcal{D}_X$. Let*

$$|\psi_1\rangle = \sum_{x \in X} \sqrt{f_1(x)}|x\rangle \quad \text{and} \quad |\psi_2\rangle = \sum_{x \in X} \sqrt{f_2(x)}|x\rangle.$$

Then

$$\| |\psi_1\rangle - |\psi_2\rangle \|_{tr} = \sqrt{1 - (1 - H^2(f_1, f_2))^2}.$$

Definition 3. We say that two families of distributions $\{D_{0,\lambda}\}_{\lambda \in \mathbb{N}}$ and $\{D_{1,\lambda}\}_{\lambda \in \mathbb{N}}$ on the same finite set $\{X_\lambda\}$ are *computationally indistinguishable* if for all quantum polynomial-time procedures $\mathcal{A} : X_\lambda \rightarrow \{0, 1\}$ there exists a negligible function $\mu(\cdot)$ such that for all $\lambda \in \mathbb{N}$,

$$\left| \Pr_{x \leftarrow D_{0,\lambda}} [\mathcal{A}(x) = 0] - \Pr_{x \leftarrow D_{1,\lambda}} [\mathcal{A}(x) = 0] \right| \leq \mu(\lambda). \quad (3)$$

2.3 Entropies

For $p \in [0, 1]$ we write $H(p) = -p \log p - (1-p) \log(1-p)$ for the binary Shannon entropy. We measure randomness using Rényi conditional entropies. For a positive semidefinite matrix $\sigma \in \text{Pos}(\mathcal{H})$ and $\varepsilon \geq 0$, let

$$\langle \sigma \rangle_{1+\varepsilon} = \text{Tr}(\sigma^{1+\varepsilon}).$$

This quantity satisfies the following approximate linearity relations:

$$\forall \varepsilon \in [0, 1], \quad \langle \sigma \rangle_{1+\varepsilon} + \langle \tau \rangle_{1+\varepsilon} \leq \langle \sigma + \tau \rangle_{1+\varepsilon} \leq (1 + O(\varepsilon)) (\langle \sigma \rangle_{1+\varepsilon} + \langle \tau \rangle_{1+\varepsilon}). \quad (4)$$

In addition, for positive semidefinite $\sigma, \rho \in \text{Pos}(\mathcal{H})$ such that the support of ρ is included in the support of σ , and $\varepsilon \geq 0$, let

$$\tilde{Q}_{1+\varepsilon}(\rho \parallel \sigma) = \langle \sigma^{-\frac{\varepsilon}{2(1+\varepsilon)}} \rho \sigma^{-\frac{\varepsilon}{2(1+\varepsilon)}} \rangle_{1+\varepsilon}. \quad (5)$$

Quantum analogues of the conditional Rényi entropies can be defined as follows.

Definition 4. Let $\rho_{AB} \in \text{Pos}(\mathcal{H}_A \otimes \mathcal{H}_B)$ be positive semidefinite. Given $\varepsilon > 0$, the $(1 + \varepsilon)$ Rényi entropy of A conditioned on B is defined as

$$H_{1+\varepsilon}(A|B)_\rho = \sup_{\sigma \in \text{D}(\mathcal{H}_B)} H_{1+\varepsilon}(A|B)_{\rho|\sigma},$$

where for any $\sigma_B \in \text{D}(\mathcal{H}_B)$,

$$H_{1+\varepsilon}(A|B)_{\rho|\sigma} = -\frac{1}{\varepsilon} \log \tilde{Q}_{1+\varepsilon}(\rho \parallel \sigma).$$

Rényi entropies are used in the proofs because they have better “chain-rule-like” properties than the min-entropy, which is the most appropriate measure for randomness quantification.

Definition 5. Let $\rho_{AB} \in \text{Pos}(\mathcal{H}_A \otimes \mathcal{H}_B)$ be positive semidefinite. Given a density matrix the *min-entropy* of A conditioned on B is defined as

$$H_\infty(A|B)_\rho = \sup_{\sigma \in \text{D}(\mathcal{H}_B)} H_\infty(A|B)_{\rho|\sigma},$$

where for any $\sigma_B \in \text{D}(\mathcal{H}_B)$,

$$H_\infty(A|B)_{\rho|\sigma} = \max \{ \lambda \geq 0 \mid 2^{-\lambda} \text{Id}_A \otimes \sigma_B \geq \rho_{AB} \}.$$

It is often convenient to consider the *smooth* min-entropy, which is obtained by maximizing the min-entropy over all positive semidefinite operators matrices in an ε -neighborhood of ρ_{AB} . The definition of neighborhood depends on a choice of metric; the canonical choice is the “purified distance”. Since this choice will not matter for us we defer to [Tom15] for a precise definition.

Definition 6. Let $\varepsilon \geq 0$ and $\rho_{AB} \in \text{Pos}(\mathcal{H}_A \otimes \mathcal{H}_B)$ positive semidefinite. The ε -smooth min-entropy of A conditioned on B is defined as

$$H_\infty^\varepsilon(A|B)_\rho = \sup_{\sigma_{AB} \in \mathcal{B}(\rho_{AB}, \varepsilon)} H_\infty(A|B)_\sigma,$$

where $\mathcal{B}(\rho_{AB}, \varepsilon)$ is the ball of radius ε around ρ_{AB} , taken with respect to the purified distance.

The following theorem relates the min-entropy to the Rényi entropies introduced earlier. The theorem expresses the fact that, up to a small amount of “smoothing” (the parameter δ in the theorem), all these entropies are of similar order.

Theorem 7 (Theorem 4.1 [MS14]). *Let $\rho_{XE} \in \text{Pos}(\mathcal{H}_X \otimes \mathcal{H}_E)$ be positive semidefinite of the form $\rho_{XE} = \sum_{x \in \mathcal{X}} |x\rangle\langle x| \otimes \rho_E^x$, where \mathcal{X} is a finite alphabet. Let $\sigma_E \in \text{D}(\mathcal{H}_E)$ be an arbitrary density matrix. Then for any $\delta > 0$ and $0 < \varepsilon \leq 1$,*

$$H_\infty^\delta(X|E)_\rho \geq -\frac{1}{\varepsilon} \log \left(\sum_x \tilde{Q}_{1+\varepsilon}(\rho_E^x \| \sigma_E) \right) - \frac{1 + 2 \log(1/\delta)}{\varepsilon}.$$

3 Trapdoor claw-free hash functions

Let λ be a security parameter, and \mathcal{X} and \mathcal{Y} be finite sets (depending on λ). For our purposes an ideal family of functions \mathcal{F} would have the following properties. For each public key k , there are two functions $\{f_{k,b} : \mathcal{X} \rightarrow \mathcal{Y}\}_{b \in \{0,1\}}$ that are both injective and have the same range (equivalently, $(b, x) \mapsto f_{k,b}(x)$ is 2-to-1), and are invertible given a suitable trapdoor t_k (i.e. t_k can be used to compute x given b and $y = f_{k,b}(x)$). Furthermore, the pair of functions should be claw-free: it must be hard for an attacker to find two pre-images $x_0, x_1 \in \mathcal{X}$ such that $f_{k,0}(x_0) = f_{k,1}(x_1)$. Finally, the functions should satisfy an adaptive hardcore bit property, which is a stronger form of the claw-free property: assuming for convenience that $\mathcal{X} = \{0,1\}^w$, we would like that it is computationally infeasible to simultaneously generate an $(b, x_b) \in \{0,1\} \times \mathcal{X}$ and a $d \in \{0,1\}^w \setminus \{0^w\}$ such that with non-negligible advantage over $\frac{1}{2}$ the equation $d \cdot (x_0 \oplus x_1) = 0$, where x_{1-b} is defined as the unique element such that $f_{k,1-b}(x_{1-b}) = f_{k,b}(x_b)$, holds.

Unfortunately, we do not know how to construct a function family that exactly satisfies all these requirements under standard cryptographic assumptions. Instead, we construct a family that satisfies slightly relaxed requirements, that we will show still suffice for our purposes, based on the hardness of the learning with errors (LWE) problem. The requirements are relaxed as follows. First, the range of the functions is no longer a set \mathcal{Y} ; instead, it is $\mathcal{D}_\mathcal{Y}$, the set of probability densities over \mathcal{Y} . That is, each function returns a density, rather than a point. The trapdoor injective pair property is then described in terms of the support of the output densities: these supports should either be identical, for a colliding pair, or be disjoint, in all other cases.

The consideration of functions that return densities gives rise to an additional requirement of efficiency: there should exist a quantum polynomial-time procedure that efficiently prepares a superposition over the range of the function, i.e. for any key k and $b \in \{0,1\}$, the procedure can prepare the state

$$\frac{1}{\sqrt{|\mathcal{X}|}} \sum_{x \in \mathcal{X}, y \in \mathcal{Y}} \sqrt{f_{k,b}(x)(y)} |x\rangle |y\rangle. \quad (6)$$

In our instantiation based on LWE, it is not possible to prepare (6) perfectly, but it is possible to create a superposition with coefficients $\sqrt{f'_{k,b}(x)}$, such that the resulting state is within negligible trace distance

of (6). The density $f'_{k,b}(x)$ is required to satisfy two properties used in our protocol. First, it must be easy to check, without the trapdoor, if an $y \in \mathcal{Y}$ lies in the support of $f'_{k,b}(x)$. Second, the inversion algorithm should operate correctly on all y in the support of $f'_{k,b}(x)$.

We slightly modify the adaptive hardcore bit requirement as well. Since the set \mathcal{X} may not be a subset of binary strings, we first assume the existence of an injective, efficiently invertible map $J : \mathcal{X} \rightarrow \{0,1\}^w$. Next, we only require the adaptive hardcore bit property to hold for a subset of all nonzero strings, instead of the set $\{0,1\}^w \setminus \{0^w\}$. Finally, membership in the appropriate set should be efficiently checkable, given access to the trapdoor.

A formal definition follows.

Definition 8 (NTCF family). Let λ be a security parameter. Let \mathcal{X} and \mathcal{Y} be finite sets. Let $\mathcal{K}_{\mathcal{F}}$ be a finite set of keys. A family of functions

$$\mathcal{F} = \{f_{k,b} : \mathcal{X} \rightarrow \mathcal{D}_{\mathcal{Y}}\}_{k \in \mathcal{K}_{\mathcal{F}}, b \in \{0,1\}}$$

is called a **noisy trapdoor claw free (NTCF) family** if the following conditions hold:

1. **Efficient Function Generation.** There exists an efficient probabilistic algorithm $\text{GEN}_{\mathcal{F}}$ which generates a key $k \in \mathcal{K}_{\mathcal{F}}$ together with a trapdoor t_k :

$$(k, t_k) \leftarrow \text{GEN}_{\mathcal{F}}(1^\lambda).$$

2. **Trapdoor Injective Pair.** For all keys $k \in \mathcal{K}_{\mathcal{F}}$ the following conditions hold.
 - (a) *Trapdoor:* For all $b \in \{0,1\}$ and $x \neq x' \in \mathcal{X}$, $\text{SUPP}(f_{k,b}(x)) \cap \text{SUPP}(f_{k,b}(x')) = \emptyset$. Moreover, there exists an efficient deterministic algorithm $\text{INV}_{\mathcal{F}}$ such that for all $b \in \{0,1\}$, $x \in \mathcal{X}$ and $y \in \text{SUPP}(f_{k,b}(x))$, $\text{INV}_{\mathcal{F}}(t_k, b, y) = x$.
 - (b) *Injective pair:* There exists a perfect matching $\mathcal{R}_k \subseteq \mathcal{X} \times \mathcal{X}$ such that $f_{k,0}(x_0) = f_{k,1}(x_1)$ if and only if $(x_0, x_1) \in \mathcal{R}_k$.

3. **Efficient Range Superposition.** For all keys $k \in \mathcal{K}_{\mathcal{F}}$ and $b \in \{0,1\}$ there exists a function $f'_{k,b} : \mathcal{X} \mapsto \mathcal{D}_{\mathcal{Y}}$ such that
 - (a) For all $(x_0, x_1) \in \mathcal{R}_k$ and $y \in \text{SUPP}(f'_{k,b}(x_b))$, $\text{INV}_{\mathcal{F}}(t_k, b, y) = x_b$ and $\text{INV}_{\mathcal{F}}(t_k, b \oplus 1, y) = x_{b \oplus 1}$.
 - (b) There exists an efficient deterministic procedure $\text{CHK}_{\mathcal{F}}$ that, on input $k, b \in \{0,1\}$, $x \in \mathcal{X}$ and $y \in \mathcal{Y}$, returns 1 if $y \in \text{SUPP}(f'_{k,b}(x))$ and 0 otherwise. Note that $\text{CHK}_{\mathcal{F}}$ is not provided the trapdoor t_k .
 - (c) For every k and $b \in \{0,1\}$,

$$\mathbf{E}_{x \leftarrow \mathcal{X}} [H^2(f_{k,b}(x), f'_{k,b}(x))] \leq \mu(\lambda),$$

for some negligible function $\mu(\cdot)$. Here H^2 is the Hellinger distance; see (1). Moreover, there exists an efficient procedure $\text{SAMP}_{\mathcal{F}}$ that on input k and $b \in \{0,1\}$ prepares the state

$$\frac{1}{\sqrt{|\mathcal{X}|}} \sum_{x \in \mathcal{X}, y \in \mathcal{Y}} \sqrt{(f'_{k,b}(x))(y) |x\rangle |y\rangle}. \quad (7)$$

4. **Adaptive Hardcore Bit.** For all keys $k \in \mathcal{K}_{\mathcal{F}}$ the following conditions hold, for some integer w that is a polynomially bounded function of λ .

- (a) For all $b \in \{0, 1\}$ and $x \in \mathcal{X}$, there exists a set $G_{k,b,x} \subseteq \{0, 1\}^w$ such that $\Pr_{d \leftarrow_U \{0, 1\}^w} [d \notin G_{k,b,x}]$ is negligible, and moreover there exists an efficient algorithm that checks for membership in $G_{k,b,x}$ given k, b, x and the trapdoor t_k .
- (b) There is an efficiently computable injection $J : \mathcal{X} \rightarrow \{0, 1\}^w$, such that J can be inverted efficiently on its range, and such that the following holds. If

$$\begin{aligned} H_k &= \{(b, x_b, d, d \cdot (J(x_0) \oplus J(x_1))) \mid b \in \{0, 1\}, (x_0, x_1) \in \mathcal{R}_k, d \in G_{k,0,x_0} \cap G_{k,1,x_1}\},^5 \\ \overline{H}_k &= \{(b, x_b, d, c) \mid (b, x, d, c \oplus 1) \in H_k\}, \end{aligned}$$

then for any quantum polynomial-time procedure \mathcal{A} there exists a negligible function $\mu(\cdot)$ such that

$$\left| \Pr_{(k,t_k) \leftarrow \text{GEN}_{\mathcal{F}}(1^\lambda)} [\mathcal{A}(k) \in H_k] - \Pr_{(k,t_k) \leftarrow \text{GEN}_{\mathcal{F}}(1^\lambda)} [\mathcal{A}(k) \in \overline{H}_k] \right| \leq \mu(\lambda). \quad (8)$$

4 Protocol description

We introduce two protocols. The first we call the (*general*) *randomness expansion protocol*, or Protocol 1. This is our main randomness expansion protocol. It is introduced in Section 4.1, and summarized in Figure 1. The protocol describes the interaction between a *verifier* and *prover*. Ultimately, we aim to obtain the guarantee that any computationally bounded prover that is accepted with non-negligible probability by the verifier in the protocol must generate transcripts that contain information-theoretic randomness.

The second protocol is called the *simplified protocol*, or Protocol 2. It is introduced in Section 4.2, and summarized in Figure 2. This protocol abstracts some of the main features Protocol 1, and will be used as a tool in the analysis (it is not meant to be executed literally).

4.1 The randomness expansion protocol

Our randomness expansion protocol, Protocol 1, is described in Figure 1. The protocol is parametrized by a choice of security parameter λ . All other parameters are assumed to be specified as a function of λ : the number of rounds N , the error tolerance parameter $\gamma \geq 0$, and the testing parameter $q \in (0, 1]$. For intuition, γ can be thought of as a small constant, q as a parameter that scales as $\text{poly}(\lambda)/N$ (for example, $q = \lambda/N$), and N as a function that may grow super-polynomially, or even exponentially, with λ . In addition, the protocol depends on an NTCF family \mathcal{F} (see Definition 8) that is known to both the verifier and the prover.

At the start of the protocol, the verifier executes $(k, t_k) \leftarrow \text{GEN}_{\mathcal{F}}(1^\lambda)$ to obtain the public key k and trapdoor t_k for a pair of functions $\{f_{k,b} : \mathcal{X} \rightarrow \mathcal{D}_{\mathcal{Y}}\}_{b \in \{0,1\}}$ from the NTCF family. The verifier sends the key k to the prover and keeps the associated trapdoor private.

In each of the N rounds of the protocol, the prover is first required to provide a value $y \in \mathcal{Y}$. For each $b \in \{0, 1\}$, the verifier uses the trapdoor to compute $\hat{x}_b \leftarrow \text{INV}_{\mathcal{F}}(t_k, b, y)$. (If the inversion procedure fails, the verifier requests another sample from the prover.) For convenience, introduce a set

$$\hat{G}_y = G_{k,0,x_0} \cap G_{k,1,x_1}. \quad (9)$$

⁵Note that although both x_0 and x_1 are referred to to define the set H_k , only one of them, x_b , is explicitly specified in any 4-tuple that lies in H_k .

The verifier then chooses a round type $G \in \{0, 1\}$ according to a biased distribution: either a *test round*, $G = 0$, chosen with probability $\Pr(G = 0) = q$, or a *generation round*, $G = 1$, chosen with the remaining probability $\Pr(G = 1) = 1 - q$. The former type of round is less frequent, as the parameter q will eventually be set to a very small value, that goes to 0 with the number of rounds of the protocol. The prover is not told the round type.

Depending on the round type, the verifier chooses a challenge $C \in \{0, 1\}$ that she sends to the prover. In the case of a test round the challenge is chosen uniformly at random; in the case of a generation round the challenge is always $C = 1$. In case $C = 0$ the prover is asked to return a pair $(m, d) \in \{0, 1\} \times \{0, 1\}^w$. The pair is called valid if $m = d \cdot (J(\hat{x}_0) \oplus J(\hat{x}_1))$ and $d \in \hat{G}_y$. If $d \in \hat{G}_y$, the verifier sets a decision bit $W = 1$ if the answer is valid, and $W = 0$ if not. If $d \notin \hat{G}_y$, the verifier sets the decision bit $W \in \{0, 1\}$ uniformly at random.⁶ In case $C = 1$, the prover should return a pair $(b, x) \in \{0, 1\} \times \mathcal{X}$. The pair is called valid if $f_{k,b}(x) = y$. The verifier sets a decision bit $W = 1$ in case the pair is valid, and $W = 0$ otherwise. The set of valid pairs on challenge $C = c \in \{0, 1\}$ is denoted $V_{y,c}$.

After each test round the verifier samples a fresh $(k, t_k) \leftarrow \text{GEN}_{\mathcal{F}}(1^\lambda)$ and communicates the new public key k to the prover.

At the end of the protocol, the verifier computes the fraction of test rounds in which the decision bit has been set to 1. If this fraction is smaller than $(1 - \gamma)$, the verifier aborts. Otherwise, the verifier returns the concatenation of the bits b obtained from the prover in generation rounds. (These bits are recorded in the verifier's output string $O_1 \cdots O_N$, such that $O_i = 0$ whenever the round is a test round.)

4.2 The simplified protocol

For purposes of analysis only we introduce a simplified variant of Protocol 1, which is specified in Figure 2. We call it the *simplified protocol*, or Protocol 2. The protocol is very similar to the randomness expansion protocol described in Figure 1, except that the prover's answers and the verifier's checks are simplified, and in test rounds there is an additional challenge bit $T \in \{0, 1\}$. This new challenge asks the prover to perform a projective measurement on its private space that indicates whether the state lies in a "good subspace" (indicated by an outcome $K = 0$) or in the complementary "bad subspace" (outcome $K = 1$). The "good" and "bad" subspace represent portions of space where the device's other two measurements, M and Π are anti-aligned and aligned respectively; see the definition of a simplified device in Section 5.2 for details.

For the case of a challenge $C = 0$, in Protocol 1 the prover returns an equation (d, m) . In the simplified protocol the prover returns a single bit $u \in \{0, 1\}$ that is meant to directly indicate the verifier's decision (i.e. the bit W). If moreover $T = 1$ the prover is required to reply with an additional bit $k \in \{0, 1\}$. In this case, the verifier makes the decision to accept, i.e. sets $W = 1$, if and only if $u = 1$ and $k = 0$. For the case of a challenge $C = 1$, in Protocol 1 the prover returns a pair (b, x) . In the simplified protocol the prover returns a value $v \in \{0, 1, 2\}$ that is such that $v = b$ in case (b, x) is valid, i.e. $(b, x) \in V_{y,1}$, and $v = 2$ otherwise.

Note that this "honest" behavior for the prover is not necessarily efficient. Moreover, it is easy for a "malicious" prover to succeed in Protocol 2, e.g. by always returning $u = 1$ (valid equation), $k = 0$ (good subspace) and $v \in \{0, 1\}$ (valid pre-image). Our analysis will not consider arbitrary provers in Protocol 2, but instead provers whose measurements satisfy certain constraints that arise from the analysis of Protocol 1. For such provers, it will be impossible to succeed in the simplified protocol without generating randomness. Further details are given in Section 6.

⁶This choice is made for technical reasons that have to do with the definition of the adaptive hardcore bit property; see Section 6 and the proof of Proposition 16 for details.

Let λ be a security parameter. Let N be a polynomially bounded function of λ , and $\gamma, q > 0$ functions of λ . Let \mathcal{F} be an NTCF family.

At the start of the protocol, the verifier communicates N to the device. In addition, the verifier samples an initial key $(k, t_k) \leftarrow \text{GEN}_{\mathcal{F}}(1^\lambda)$, sends k to the prover and keeps the trapdoor information t_k private.

1. For $i = 1, \dots, N$:
 - (a) The prover returns a $y \in \mathcal{Y}$ to the verifier. For $b \in \{0, 1\}$ the verifier uses the trapdoor to compute $\hat{x}_b \leftarrow \text{INV}_{\mathcal{F}}(t_k, b, y)$.
 - (b) The verifier selects a round type $G_i \in \{0, 1\}$ according to a Bernoulli distribution with parameter q : $\Pr(G_i = 0) = q$ and $\Pr(G_i = 1) = 1 - q$. In case $G_i = 0$ (*test round*), she chooses a challenge $C_i \in \{0, 1\}$ uniformly at random. In case $G_i = 1$ (*generation round*), she sets $C_i = 1$. The verifier keeps G_i private, and sends C_i to the prover.
 - i. In case $C_i = 0$ the prover returns $(m, d) \in \{0, 1\} \times \{0, 1\}^w$. If $d \notin \hat{G}_y$, the set defined in (9), the verifier sets W to a uniformly random bit. Otherwise, the verifier sets $W = 1$ if $d \cdot (J(\hat{x}_0) \oplus J(\hat{x}_1)) = m$ and $W = 0$ if not.
 - ii. In case $C_i = 1$ the prover returns $(b, x) \in \{0, 1\} \times \mathcal{X}$. The verifier sets W as the value returned by $\text{CHK}_{\mathcal{F}}(k, b, x, y)$.
 - (c) In case $G_i = 1$, the verifier sets $O_i = b$. In case $G_i = 0$, she sets $W_i = W$.
 - (d) In case $G_i = 0$, the verifier samples a new key $(k, t_k) \leftarrow \text{GEN}_{\mathcal{F}}(1^\lambda)$. She sends k to the prover and keeps the trapdoor information t_k private. This key will be used until the next test round, included.
 2. If $\sum_{i:G_i=0} W_i < (1 - \gamma)qN$, the verifier aborts. Otherwise, she returns the string O obtained by concatenating the bits O_i for all $i \in \{1, \dots, N\}$ such that $G_i = 1$.
-

Figure 1: The randomness expansion protocol, Protocol 1. See Definition 8 for notation associated with the NTCF family \mathcal{F} .

Let λ be a security parameter. Let N be a polynomially bounded function of λ , and $\gamma, \eta, \kappa, q > 0$ functions of λ .

1. For $i = 1, \dots, N$:
 - (a) The verifier selects a round type $G_i \in \{0, 1\}$ according to a Bernoulli distribution with parameter q : $\Pr(G_i = 0) = q$ and $\Pr(G_i = 1) = 1 - q$. In case $G_i = 0$ (*test round*), she chooses $C_i \in \{0, 1\}$ uniformly at random and $T_i \in \{0, 1\}$ such that $\Pr(T_i = 0) = 1 - \kappa$ and $\Pr(T_i = 1) = \kappa$. In case $G_i = 1$ (*generation round*), she sets $C_i = 1$ and $T_i = 0$. The verifier keeps G_i private, and sends (C_i, T_i) to the prover.
 - i. In case $C_i = 0$ the prover returns $u \in \{0, 1\}$. If $T_i = 1$ the prover in addition reports $k \in \{0, 1\}$.⁷ If $T_i = 0$ the verifier sets $W_i = u$. If $T_i = 1$ the verifier sets $W_i = u(1 - k)$.
 - ii. In case $C_i = 1$ the prover returns $v \in \{0, 1, 2\}$. The verifier sets $O_i = v$ and $W_i = 1_{v \in \{0, 1\}}$.
 2. If $\sum_{i: G_i=0 \wedge T_i=1} W_i < (1 - \frac{\gamma}{\kappa} - \eta)\kappa q N$, the verifier rejects the interaction. Otherwise, she returns the string O obtained by concatenating the bits O_i for all $i \in \{1, \dots, N\}$ such that $G_i = 1$.
-

Figure 2: The simplified protocol, Protocol 2.

4.3 Completeness

We describe the intended behavior for the prover in protocol 1. Fix an NTCF family \mathcal{F} and a key $k \in \mathcal{K}_{\mathcal{F}}$. In each round, the “honest” prover performs the following actions.

1. The prover executes the efficient procedure $\text{SAMP}_{\mathcal{F}}$ in superposition to obtain the state

$$|\psi^{(1)}\rangle = \frac{1}{\sqrt{|\mathcal{X}|}} \sum_{x \in \mathcal{X}, y \in \mathcal{Y}, b \in \{0, 1\}} \sqrt{(f'_{k,b}(x))(y)} |b, x\rangle |y\rangle .$$

2. The prover measures the last register to obtain an $y \in \mathcal{Y}$. Using item 2. from the definition of an NTCF, the prover’s re-normalized post-measurement state is

$$|\psi^{(2)}\rangle = \frac{1}{\sqrt{2}} (|0, x_0\rangle + |1, x_1\rangle) |y\rangle .$$

- (a) In case $C_i = 0$, the prover evaluates the function J on the second register, containing x_b , and then applies a Hadamard transform to all $w + 1$ qubits in the first two registers. Tracing out the register that contains y , this yields the state

$$\begin{aligned} |\psi^{(3)}\rangle &= 2^{-\frac{w+2}{2}} \sum_{d, b, m} (-1)^{d \cdot J(x_b) \oplus mb} |m\rangle |d\rangle \\ &= (-1)^{J(x_0)} 2^{-\frac{w}{2}} \sum_{d \in \{0, 1\}^w} |d \cdot (J(x_0) \oplus J(x_1))\rangle |d\rangle . \end{aligned}$$

The prover measures both registers to obtain an (m, d) that it sends back to the verifier.

- (b) In case $C_i = 1$, the prover measures the first two registers of $|\psi^{(2)}\rangle$ in the computational basis, and returns the outcome (b, x_b) to the verifier.

Lemma 9. *For any λ and $k \leftarrow \text{GEN}_{\mathcal{F}}(1^\lambda)$, the strategy for the honest prover (on input k) in one round of the protocol can be implemented in time polynomial in λ and is accepted with probability negligibly close to 1.*

Proof. Both efficiency and correctness of the prover follow from the definition of an NTCF (Definition 8). The prover fails only if he obtains an outcome $d \notin \hat{G}_y$, which by item 4(a) in the definition happens with negligible probability. \square

5 Devices

We model an arbitrary prover in the randomness expansion protocol (Protocol 1 in Figure 1) as a *device* that implements the actions of the prover: the device first returns an $y \in \mathcal{Y}$; then, depending on the challenge $C \in \{0, 1\}$, it either returns an equation (m, d) (case $C = 0$), or a candidate pre-image (b, x) (case $C = 1$). For simplicity we assume that the device makes the same set of measurements in each round of the protocol. This is without loss of generality, as we allow the state of the device to change from one round to the next; in particular the device is allowed to use a quantum memory as a control register for the measurements.

In Section 5.1 we introduce our notation for modeling provers in Protocol 1 as devices. In Section 5.2 we consider a simplified form of device, that is appropriate for modeling a prover in the simplified protocol, Protocol 2. In Section 6 we give a reduction showing how to associate a specific simplified device to any computationally efficient general device, such that the randomness generation properties of the two devices can be related to each other (this is done in Section 7).

For the remainder of this section we fix an NTCF family \mathcal{F} satisfying the conditions of Definition 8, and use notation introduced in the definition.

5.1 General devices

The following notion of device models the behavior of an arbitrary prover in the randomness expansion protocol, Protocol 1 (Figure 1).

Definition 10. Given $k \in \mathcal{K}_{\mathcal{F}}$, a device $D = (\phi, \Pi, M)$ (implicitly, compatible with k) is specified by the following:

1. A (not necessarily normalized) positive semidefinite $\phi \in \text{Pos}(\mathcal{H}_D \otimes \mathcal{H}_Y)$. Here \mathcal{H}_D is an arbitrary space private to the device, and \mathcal{H}_Y is a space of the same dimension as the cardinality of the set \mathcal{Y} , also private to the device. For every $y \in \mathcal{Y}$, define

$$\phi_y = (\text{Id}_D \otimes \langle y|_Y) \phi (\text{Id}_D \otimes |y\rangle_Y) \in \text{Pos}(\mathcal{H}_D).$$

Note that ϕ_y is not normalized, and $\sum_{y \in \mathcal{Y}} \text{Tr}(\phi_y) = \text{Tr}(\phi)$.

2. For every $y \in \mathcal{Y}$, a projective measurement $\{M_y^{(m,d)}\}$ on \mathcal{H}_D , with outcomes $(m, d) \in \{0, 1\} \times \{0, 1\}^w$.

⁷The bit k should not be confused with the public key k for the NTCF that is used in Protocol 1. In Protocol 2, there is no NTCF, and no key.

3. For every $y \in \mathcal{Y}$, a projective measurement $\{\Pi_y^{(b,x)}\}$ on \mathcal{H}_D , with outcomes $(b, x) \in \{0, 1\} \times \mathcal{X}$. For each y , this measurement has two designated outcomes $(0, x_0)$ and $(1, x_1)$, which are the answers that are accepted on challenge $C = 1$ in the protocol; recall that we use the notation $V_{y,1}$ for this set. For $b \in \{0, 1\}$ we use the shorthand $\Pi_y^b = \Pi_y^{(b,x_b)}$, $\Pi_y = \Pi_y^0 + \Pi_y^1$, and $\Pi_y^2 = \text{Id} - \Pi_y^0 - \Pi_y^1$.

By Naimark's theorem, up to increasing the dimension of \mathcal{H}_D the assumption that $\{\Pi_y^{(b,x)}\}$ and $\{M_y^{(m,d)}\}$ are projective is without loss of generality.

We explain the connection between the notion of device in Definition 10 and a prover in Protocol 1. Given a device $D = (\phi, \Pi, M)$, we can define actions for the prover in Protocol 1 as follows. The prover is initialized in state ϕ . When a round of the protocol is initiated, the prover measures register Y in the computational basis and returns the outcome $y \in \mathcal{Y}$. We always assume that the prover directly measures the register, as any pre-processing unitary can be incorporated in the definition of the state ϕ . When sent challenge $C = 0$ (resp. $C = 1$), the prover measures register D using the device's projective measurement $\{M_y^{(m,d)}\}$ (resp. $\{\Pi_y^{(b,x)}\}$), and returns the outcome to the verifier.

Definition 11. We say that a device $D = (\phi, \Pi, M)$ is *efficient* if

1. There is a polynomial-size circuit to prepare ϕ , given the NTCF key k ;
2. For every $y \in \mathcal{Y}$, the measurements $\{M_y^{(m,d)}\}$ and $\{\Pi_y^{(b,x)}\}$ can be implemented by polynomial-size circuits.

Using the definition of an NTCF family (Definition 8), it is straightforward to verify that the device associated with the ‘‘honest’’ prover described in Section 4.3 is efficient.

5.2 Simplified devices

Next we introduce a simplified model of device, that can be used to model the actions of a prover in the simplified protocol, Protocol 2 (Figure 2).

Definition 12. A *simplified device* is a tuple (ϕ, Π, M, K) such that:

1. $\phi = \{\phi_y\}_{y \in \mathcal{Y}} \subseteq \text{Pos}(\mathcal{H}_D)$ is a family of positive semidefinite operators on an arbitrary space \mathcal{H}_D ;
2. For each $y \in \mathcal{Y}$, $\{M_y^0, M_y^1 = \text{Id} - M_y^0\}$, $\{\Pi_y^0, \Pi_y^1, \Pi_y^2 = \text{Id} - \Pi_y^0 - \Pi_y^1\}$, and $\{K_y^0, K_y^1 = \text{Id} - K_y^0\}$ are projective measurements on \mathcal{H}_D ;
3. For each $y \in \mathcal{Y}$, the measurement operators K_y commute with the M_y and with the Π_y . (M_y and Π_y do not necessarily commute with each other.)

To any simplified device $D = (\phi, \Pi, M, K)$ we associate the post-measurement states

$$\begin{aligned} \forall t, u, k \in \{0, 1\}, \quad \phi_{0t}^{uk} &= \sum_{y \in \mathcal{Y}} |y\rangle\langle y| \otimes (1_{t=1} K_y^k) M_y^u \phi_y M_y^u (1_{t=1} K_y^k), \\ \forall v \in \{0, 1, 2\}, \quad \phi_1^v &= \sum_{y \in \mathcal{Y}} |y\rangle\langle y| \otimes \Pi_y^v \phi_y \Pi_y^v. \end{aligned} \tag{10}$$

A simplified device can be used in the simplified protocol in a straightforward way: upon receipt of a challenge $C = 0$ (resp. $C = 1$), the device first samples an $y \in \mathcal{Y}$ according to the distribution with weights $\text{Tr}(\phi_y)$ (the weights may need to be re-normalized). It then performs the projective measurement $\{M_y^0, M_y^1\}$ followed by, if $T = 1$, $\{K_y^0, K_y^1\}$ (resp. $\{\Pi_y^0, \Pi_y^1, \Pi_y^2\}$) on ϕ_y , and returns the outcomes $u, k \in \{0, 1\}$ (resp. $v \in \{0, 1, 2\}$) to the verifier.

We introduce a quantity called *overlap* that measures how “incompatible” a simplified device’s measurements are. This measure is analogous to the measure of overlap used to quantify incompatibility in the derivation of entropic uncertainty relations (see e.g. [MU88]).

Definition 13. Given a simplified device $D = (\phi, \Pi, M, K)$, the *overlap* of D is

$$\Delta(D) = \max_{y \in \mathcal{Y}} \|K_y^0(\Pi_y^0 M_y^1 \Pi_y^0 + \Pi_y^1 M_y^1 \Pi_y^1)\|.$$

Note that the overlap only quantifies the measurement incompatibility in the “good subspace” K_y^0 .

6 Single-round analysis

In this section we consider the behavior of an arbitrary device D in a single round of the randomness expansion protocol, Protocol 1 in Figure 1. Our goal is to introduce a simplified device D' such that analyzing the randomness generation properties of D' is easier than for D , and such that bounds on the amount of randomness generated by D' in the simplified protocol, Protocol 2 in Figure 2, imply bounds on the amount of randomness generated by D in Protocol 1. Throughout the section we fix an NTCF family \mathcal{F} (Definition 8) and a key $k \in \mathcal{K}_{\mathcal{F}}$ sampled according to $\text{GEN}(1^\lambda)$, for a parameter λ that plays the role of security parameter.

In order to make the comparison between D and D' we introduce notation related to the post-measurement states generated by $D = (\phi, M, \Pi)$ in Protocol 1. An execution of Protocol 1 involves a choice of round types $g \in \{0, 1\}^N$ and of challenges $c \in \{0, 1\}^N$ by the verifier, and a sequence of outputs $o \in \{0, 1, 2\}^N$ computed by the verifier as a function of the answers provided by the device. Here, in case $g = 0$ (test round) we use $o \in \{0, 1\}$ to denote the outcome of the test (called W in the protocol description), and in case $g = 1$ (generation round) we use $o \in \{0, 1, 2\}$ such that $o = 2$ in case $W = 0$, and $o = O$ as recorded by the verifier in case $W = 1$. We call the tuple (c, o) the transcript of the protocol; it contains all the information relevant to the verifier’s final acceptance decision, and to the extraction of randomness. Additional information, such as the choice of public keys, the prover’s equations, etc., is discarded for ease of presentation. We let ACC denote the set of tuples (g, c, o) that are accepted by the verifier in the last step of the protocol, i.e. such that $\sum_{i: g_i=0} o_i \geq (1 - \gamma)qN$. The joint state of the transcript and the device at the end of the N rounds (but before the verifier’s decision to abort) of the protocol can then be written as

$$\phi_{\text{COD}}^{(N)} = \sum_{g, c, o} q(g, c) |c\rangle\langle c|_{\text{C}} \otimes |o\rangle\langle o|_{\text{O}} \otimes \phi_{\text{D}}^{co}, \quad (11)$$

where $q(g, c)$ is the probability that the sequence of round types and challenges (g, c) is chosen by the verifier in the protocol (this is a simple function of the parameter q that we do not yet need to make explicit), and ϕ_{D}^{co} is the post-measurement state of the device, conditioned on having received challenges c and returned outcomes o . Finally, we recall that we may purify the initial state ϕ_{D} of the device into a pure state $|\phi\rangle_{\text{DE}}$, where we think of the register E as being held by an arbitrary adversary. We write ρ_{E}^{co} for the post-measurement state of the adversary, conditioned on the transcript being (c, o) .

The preceding paragraph defines post-measurement states $\{\phi^{co}\}$ associated with a device D . Given a simplified device $D' = (\phi', M', \Pi', K)$, ignoring for now the measurement K it is possible to analogously define post-measurement states $\{(\phi')^{co}\}$ that are the post-measurement states obtained by making the measurement M if $c = 1$ and Π if $c = 0$, and recording the outcomes in o . These ensembles of states provide a means to meaningfully compare a device D and a simplified device D' . We record this in the following definition.

Definition 14. Let $D = (\phi, M, \Pi)$ be a device and $D' = (\phi', M', \Pi', K)$ a simplified device. We say that D' *simulates* D if for every $(c, o) \in \{0, 1\}^N \times \{0, 1, 2\}^N$ the states ϕ^{co} and $(\phi')^{co}$ are identical.

We start with a lemma that leverages the adaptive hardcore bit property to argue computational indistinguishability between different post-measurement states. Recall that for any $y \in \mathcal{Y}$, $V_{y,0}$ denotes the set of valid answers for the prover on challenge $C = 0$, and $\hat{G}_y \subseteq \{0, 1\}^w$ is the set used by the verifier on challenge $C = 0$.

Lemma 15. Let $D = (\phi, \Pi, M)$ be an efficient device. Define a sub-normalized density

$$\begin{aligned} \rho_{YBXD} &= \sum_{y \in \mathcal{Y}} |y\rangle\langle y|_Y \otimes \sum_{b \in \{0,1\}} |b, x_b\rangle\langle b, x_b|_{BX} \otimes \Pi_y^{(b, x_b)} \phi_y \Pi_y^{(b, x_b)} \\ &= \sum_{b \in \{0,1\}} |b, x_b\rangle\langle b, x_b|_{BX} \otimes \rho_{YD}^{(b)}. \end{aligned} \quad (12)$$

Then ρ can be efficiently prepared (with success probability $\text{Tr}(\rho)$). Let

$$\begin{aligned} \sigma_0 &= \sum_{b \in \{0,1\}} |b, x_b\rangle\langle b, x_b|_{BX} \otimes \sum_{(m,d) \in V_{y,0}} |d, m\rangle\langle d, m|_{DM} \otimes (\text{Id}_Y \otimes M_y^{(m,d)}) \rho_{YD}^{(b)} (\text{Id}_Y \otimes M_y^{(m,d)}), \\ \sigma_1 &= \sum_{b \in \{0,1\}} |b, x_b\rangle\langle b, x_b|_{BX} \otimes \sum_{(m,d) \notin V_{y,0}} \mathbf{1}_{d \in \hat{G}_y} |d, m\rangle\langle d, m|_{DM} \otimes (\text{Id}_Y \otimes M_y^{(m,d)}) \rho_{YD}^{(b)} (\text{Id}_Y \otimes M_y^{(m,d)}). \end{aligned} \quad (13)$$

Then the states σ_0 and σ_1 are computationally indistinguishable.

Proof. Suppose for contradiction that there exists an efficiently implementable observable O such that

$$\text{Tr}(O(\sigma_0 - \sigma_1)) \geq \kappa, \quad (14)$$

for some non-negligible function $\kappa(\lambda)$. For $u \in \{0, 1\}$ consider the following procedure. The procedure first prepares the state ρ in (12). This can be done efficiently by first preparing ϕ_{YD} , then measuring an $y \in \mathcal{Y}$, then applying the measurement $\{\Pi_y^{(b,x)}\}$ to ϕ_y , and rejecting if the outcome is invalid, i.e. $\text{CHK}_{\mathcal{F}}(k, b, x, y) = 0$.

The procedure then applies the measurement $\{M_y^{(m,d)}\}$ to ρ , obtaining an outcome (d, m) . At this point, if $d \in \hat{G}_y$ then the procedure has either prepared σ_0 or σ_1 . Finally, the procedure measures O to obtain a bit u , and returns $(b, x, d, u \oplus m)$. This defines an efficient procedure. Moreover, using (14) it follows that the procedure violates the hardcore bit property (8) (that only considers those outcome tuples from the adversary that are such that $d \in \hat{G}_y$). \square

Recall the definitions of a simplified device (Definition 12), of the overlap of a simplified device (Definition 13), and of a simplified device simulating a device (Definition 14).

Proposition 16. *Let $D = (\phi, \Pi, M)$ be an efficient device and $\frac{1}{2} < \omega \leq 1$. Then there is a (not necessarily efficient) simplified device $D' = (\phi', \Pi', M', K)$ such that the following hold. First, D' has overlap $\Delta(D') \leq \omega$. Second, for any state ϕ' (that is independent from the key $k \in \mathcal{K}_{\mathcal{F}}$) it holds that*

$$\sum_y \text{Tr}(K_y^1 \phi'_y) \leq C \sqrt{\sum_y \text{Tr}((M')_y^0 \phi'_y)} + \text{negl}(\lambda), \quad (15)$$

where $C > 0$ is a universal constant. Finally, the simplified device D' simulates the device D .

Proof. For each $y \in \mathcal{Y}$ we argue the existence of a “good subspace”, identified by a partial isometry U_y , such that the devices’ measurements are strongly incompatible, when restricted to the good subspace. Informally, incompatibility of the devices’ measurements on most of space (as measured with respect to any side information ϕ' that is independent from the key k) is necessary for the device to *not* be able to break the adaptive hardcore bit property of the NTCF: “aligned” measurements can be performed in sequence, and this would allow the device to efficiently determine both an (b, x_b) and an (d, c) that violate (8). We formalize this intuition below.

For the remainder of the proof we fix an $y \in \mathcal{Y}$ and for clarity suppress dependence of all operators on y . Let

$$\hat{M} = \sum_{(m,d): d \notin G_y} M_y^{(m,d)}, \quad M = \sum_{(m,d) \in V_{y,0}} M_y^{(m,d)} + \frac{1}{2} \hat{M},$$

and for $b \in \{0, 1\}$, $\Pi^b = \Pi_y^{(b, x_b)}$. By introducing an isometry $U : \mathcal{H}_D \rightarrow \mathcal{H}_{D'}$ into a larger space, we can embed M into a projection \overline{M} such that $M = U^\dagger \overline{M} U$. Let ρ be as defined in (12), and $\overline{\phi}' = U \phi' U^\dagger$. Recall that $\Pi = \Pi^0 + \Pi^1$. Let $\tilde{\Pi}^0 = \overline{\Pi}^0$ and $\tilde{\Pi}^1 = \text{Id} - \overline{\Pi}^0$ be projection operators on $\mathcal{H}_{D'}$, where $\overline{\Pi}^0$ is such that $U_y^\dagger \overline{\Pi}^0 U_y = \Pi^0$. Note that the measurement $\{\tilde{\Pi}^0, \tilde{\Pi}^1\}$ can be implemented efficiently. Using Jordan’s lemma we find a basis of $\mathcal{H}_{D'}$ in which

$$\overline{M} = \oplus_j \begin{pmatrix} c_j^2 & c_j s_j \\ c_j s_j & s_j^2 \end{pmatrix} \quad \text{and} \quad \tilde{\Pi}^0 = \oplus_j \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}, \quad (16)$$

where $c_j = \cos \theta_j$, $s_j = \sin \theta_j$, for some angles θ_j . In addition there may be 1-dimensional blocks in the Jordan decomposition, but up to adding a few dimensions these can be identified with two-dimensional blocks such that $c_j^2 \in \{0, 1\}$. For $\frac{1}{2} < \omega \leq 1$ a parameter, let K be the orthogonal projection on those 2-dimensional blocks such that $\min(c_j^2, s_j^2) \geq 1 - \omega$. Note that K commutes with both \overline{M} and $\{\tilde{\Pi}^0, \tilde{\Pi}^1\}$, but not necessarily with $\overline{\phi}'$.

Let $\gamma = 1 - \text{Tr}(\overline{M} \overline{\phi}')$. From the computational indistinguishability of the states σ'_0 and σ'_1 expressed in Lemma 15 it follows that $\text{Tr}(\Pi^0 \overline{M} \Pi^0 \overline{\phi}') + \text{Tr}(\Pi^1 \overline{M} \Pi^1 \overline{\phi}')$ is within negligible of $\frac{1}{2}$. Suppose first that $\gamma = 0$. Then $\overline{\phi}'$ is supported on the range of \overline{M} , and it follows from the block decomposition in (16) that $E_j(c_j^4 + s_j^4)$ is within negligible of $\frac{1}{2}$, where the expectation is taken according to a block index j sampled proportionally to the overlap of $\overline{\phi}'$ on that block. Using that $c_j^4 + s_j^4 \geq 1/2$ always, it follows that for any value of $\omega < 1/2$ we have that $\text{Tr}((\text{Id} - K) \overline{\phi}')$ is negligible.

Next consider the case where $\gamma > 0$. Let $\overline{\phi}'' = \overline{M} \overline{\phi}' \overline{M} / \text{Tr}(\overline{M} \overline{\phi}')$. By the gentle measurement lemma (see e.g. [Wil11, Lemma 9.4.1]),

$$\|\overline{\phi}' - \overline{\phi}''\|_1 \leq 2\sqrt{\gamma}. \quad (17)$$

Using again that $\text{Tr}(\Pi^0 \bar{M} \Pi^0 \bar{\phi}') + \text{Tr}(\Pi^1 \bar{M} \Pi^1 \bar{\phi}')$ is within negligible of $\frac{1}{2}$, it follows that the same expression, evaluated on $\bar{\phi}''$, is at most $\frac{1}{2} + 4\sqrt{\gamma}$. Thus, under the distribution on blocks obtained by sampling an index j proportionally to the overlap of $\bar{\phi}''$ on that block, it holds that $\mathbb{E}_j(c_j^4 + s_j^4) \leq \frac{1}{2} + 4\sqrt{\gamma}$. Using that $c_j^4 + s_j^4 > 1 - 2\omega(1 - \omega)$ whenever $\min(c_j^2, s_j^2) < 1 - \omega$, applying Markov's inequality we get

$$\text{Tr}((\text{Id} - K)\bar{\phi}'') \leq \frac{4\sqrt{\gamma}}{\frac{1}{2} - 2\omega(1 - \omega)} = O(\sqrt{\gamma}).$$

Finally, the same bound transfers to $\bar{\phi}'$, up to an additional loss of $2\sqrt{\gamma}$, using again (17).

To conclude the proof of the proposition we define the device D' . This is done as follows. The device D' first measures an $y \in \mathcal{Y}$ exactly as D would. It then applies the isometry U_y . This defines the states $\{\phi_y\}$.

- The measurement $\{(M')^0, (M')^1\}$ is defined as follows. The device coherently performs the measurement $\{\bar{M}_y^{(d,m)}\}$. If $d \notin \hat{G}_y$ the device returns a random outcome. Otherwise, if $(d, m) \in V_{y,0}$ it returns a 0, and 1 if not.
- The measurement $\{(\Pi')^0, (\Pi')^1, (\Pi')^2\}$ is defined as follows. The device first coherently performs the measurement $\{\bar{\Pi}_y^{(b,x)}\}$. If an outcome (b, x) such that $f_{k,b}(x) = y$ is obtained, the device returns $v = b$. Otherwise, if $f_{k,b}(x) \neq y$ the device returns $v = 2$.
- The measurement $\{K_y^0, K_y^1\}$ is defined by setting $K_y^0 = K_y$ and $K_y^1 = \text{Id} - K_y$.

The three conditions on D' claimed in the lemma follow by definition. (For the simulation property, note that it is possible for D' to further measure the post-measurement states to locally obtain an equation, or a pre-image, as D would have; this guarantees that the post-measurement states of the two devices are identical in each round.) \square

7 Accumulating randomness across multiple rounds

To analyze the randomness generated by a device in the randomness expansion protocol we proceed in two steps. First, we show that the randomness generated by the device can be related to the randomness generated by the simplified device D' that is associated to it by Proposition 16, when it is used as a device in the simplified protocol, Protocol 2. This is done in Section 7.1. Then, in Section 7.2 we analyze the randomness generated in a single round of the simplified protocol, and in Section 7.3 we analyze multiple rounds of the protocol.

7.1 Reduction to the simplified protocol

Let $D = (\phi, \Pi, M, K)$ be a simplified device. Recall the definition of the states $\{\phi^{c^0}\}_{(c,\rho) \in \{0,1\}^N \times \{0,1,2\}^N}$ associated with the device, as described in the beginning of Section 6 (in that section the simplified device is called D' , and the states are marked with a prime), and the pure state $|\phi^{c^0}\rangle$ from which each density is obtained by taking the partial trace over the adversary's system \mathbb{E} .

The main difference between the behavior of the simplified device and the original device it is derived from is that the simplified device (sometimes) performs an additional projective measurement $\{K, \text{Id} - K\}$, in addition to the ‘‘equation’’ measurement $\{M, \text{Id} - M\}$. (Recall that in protocol 2, the device performs

the measurement whenever the verifier selects a challenge bit $T = 1$, which happens with probability $\Pr(K = 1) = \kappa$ in the test rounds.)

We will see in later sections that, in order to analyze the randomness generated by the original device in Protocol 1, it is convenient to obtain the guarantee that, in most rounds of Protocol 1, the state of the device lies largely within the “good subspace” K . However, it is not possible to measure $\{K, \text{Id} - K\}$ repeatedly, as the post-measurement state of the device, hence its behavior in subsequent rounds, would be affected by such a measurement.

This leads us to express the final state $|\phi^{co}\rangle$ of the device as a coherent superposition over exponentially many branches,

$$|\phi^{co}\rangle = \sum_k P_{ct}^{ok} |\phi\rangle, \quad (18)$$

where P_{ct}^{ko} is notation for the operator that corresponds to applying the device’s (projective) measurement operators M and K indicated by c and t respectively, and obtaining the sequence of outcomes o and k respectively.

Our goal is to bound the contribution to (18) of terms $P_{ct}^{ok} |\phi\rangle$ that correspond to a large fraction of $(\text{Id} - K)$ (“bad subspace”) outcomes, i.e. such that the string $|k \wedge t|$ has large Hamming weight. Due to possible interference between the branches, establishing the right bound is a little more delicate than a generic martingale bound. We first state and prove a general lemma that achieves this in a general setting, and then show how the lemma can be applied in our context.

Lemma 17. *Let $|\psi\rangle_{AB}$ be a bipartite state. Let $\{M_1, \dots, M_N\}$ and $\{K_1, \dots, K_N\}$ be two sequences of projections on \mathcal{H}_A such that for each $i \in \{1, \dots, N\}$, M_i and K_i commute. For $i \in \{1, \dots, N\}$, write $M_i^0 = M_i$ and $M_i^1 = \text{Id} - M_i$, and $K_i^0 = K_i$ and $K_i^1 = \text{Id} - K_i$. For $1 \leq i \leq N$ and $t, m, k \in \{0, 1\}^i$ let*

$$|\psi_t^{mk}\rangle = M_i^{m_i}(\mathbf{1}_{t_i=1} K_i^{k_i}) \cdots M_1^{m_1}(\mathbf{1}_{t_1=1} K_1^{k_1}) |\psi\rangle,$$

and write $|\psi^m\rangle$ for the case $k = t = 0^i$. Assume that there is a monotone concave function $g : [0, 1] \rightarrow [0, 1]$ such that $g(0) = 0$ and for any $i \in \{1, \dots, N - 1\}$ and any sequences $t, m, k \in \{0, 1\}^i$ it holds that

$$\|(\text{Id} - K_{i+1}) |\psi_t^{mk}\rangle\|^2 \leq g\left(\|M_{i+1} |\psi_t^{mk}\rangle\|^2\right). \quad (19)$$

Then for any $0 < \gamma, \kappa, \eta < 1$ such that $\gamma/\kappa \ll \eta$ and $g(\kappa) \ll \eta$,

$$\sum_t \kappa(t) \sum_{m: |m|_H \geq (1-\gamma)N} \left\| |\psi^m\rangle - \sum_{k: |k \wedge t|_H \leq \eta \kappa N} |\psi_t^{mk}\rangle \right\|^2 \leq C_0 2^{-\kappa n}, \quad (20)$$

where $\kappa(t) = \prod_i \kappa^{t_i} (1 - \kappa)^{1-t_i}$, $|\cdot|_H$ denotes the Hamming weight, and $C_0 > 0$ is a constant depending on η, κ, δ .

The proof of Lemma 17 uses a version of Bennett’s inequality for martingales from [FGL12] that we quote in a form tailored to our use.

Theorem 18 (Theorem 2.1 in [FGL12]). *Let $(\xi_i, \mathcal{F}_i)_{0 \leq i \leq N}$ be a supermartingale difference sequence such that $\xi_0 = 0$ and $\xi_i \leq 1$ for each $i \in \{1, \dots, N\}$. Let*

$$X_N = \sum_{i=1}^N \xi_i \quad \text{and} \quad \langle X \rangle_N = \sum_{i=1}^N \mathbb{E}[\xi_i^2 | \mathcal{F}_{i-1}].$$

Then for any $t \geq 0$ and $v > 0$,

$$\Pr \left(X_N \geq tN \text{ and } \langle X \rangle_N \leq v^2 N \right) \leq e^{-\frac{t}{2} \ln \left(\frac{t}{v^2} \right) N + tN}.$$

Using the tail bound from Theorem 18, we give the proof of Lemma 17.

Proof of Lemma 17. We reduce the proof of (20) to a sequence of martingale tail bounds. Define a probability on $3N$ -bit strings $(M, T, K) \in \{0, 1\}^N$ by defining, for any $(m, t, k) \in \{0, 1\}^N$,

$$\Pr \left((M, T, K) = (m, t, k) \right) = \kappa(t) \left\| |\psi_t^{mk}\rangle \right\|^2.$$

That this is a normalized distribution follows from the definition of the states $|\psi_t^{mk}\rangle$. Define a filtration $(\mathcal{F}_1, \dots, \mathcal{F}_i, \dots, \mathcal{F}_N)$ where \mathcal{F}_i is the σ -algebra generated by (M_i, T_i, K_i) . Let $\mathcal{F}_{<i} = \cap_{j < i} \mathcal{F}_j$. Define events

$$\mathcal{G} = \left\{ m \in \{0, 1\}^N : |m|_H \geq (1 - \gamma)N \right\}, \quad \mathcal{B} = \left\{ (t, k) \in \{0, 1\}^{2N} : |t \wedge k| \geq \eta \kappa N \right\}.$$

The proof proceeds in 3 steps.

First step: conditional expectations of M . [Uses the assumption: $\delta_1 \gg \kappa$.]

For $i \in \{1, \dots, N\}$ let $Z_i = M_i - \mathbb{E}[M_i | \mathcal{F}_{<i}]$ and $W_i = Z_1 + \dots + Z_i$. Then the sequence (W_1, \dots, W_N) is a martingale such that $|W_i - W_{i-1}| \leq 1$. Applying Azuma's inequality, it follows that there exists a constant $C_1 > 0$ such that for any $\delta_1 > 0$,

$$\Pr \left(\left| \sum_{i=1}^N Z_i \right| \geq \delta_1 N \right) \leq e^{-C_1 \delta_1 N}. \quad (21)$$

Let $\delta_1 > 0$ be large enough such that the right-hand side of (21) is less than $2^{-2\kappa N}$. Let $\delta'_1 = \delta_1 + \gamma$ and

$$\mathcal{B}' = \left\{ (m, t, k) : \sum_{i=1}^N \mathbb{E}[M_i | (M, T, K)_{<i} = (m, t, k)_{<i}] \leq (1 - \delta'_1)N \right\}.$$

Then,

$$\begin{aligned} \sum_{m \in \mathcal{G}} \sum_t \Pr(T = t) & \left(\sum_{k: (m, t, k) \in \mathcal{B}'} \sqrt{\Pr((M, K) = (m, k) | T = t)} \right)^2 \\ & \leq 2^{\kappa N} \left(\sum_{m \in \mathcal{G}} \sum_{t, k: (m, t, k) \in \mathcal{B}'} \Pr((M, T, K) = (m, t, k)) \right) \\ & \leq 2^{-\kappa N}, \end{aligned} \quad (22)$$

where the last inequality follows from (21) since the event that $m \in \mathcal{G}$ and $(m, t, k) \in \mathcal{B}'$ implies $|\sum z_i| \geq \delta_1 N$.

Second step: conditional expectations of MT . [Uses the assumption: $\delta'_1 \ll \delta_2$.]

For $i \in \{1, \dots, N\}$ let $Z'_i = M_i T_i - \mathbb{E}[M_i T_i | \mathcal{F}_{<i}]$ and $W'_i = Z'_1 + \dots + Z'_i$. Then the sequence (W'_1, \dots, W'_N) is a martingale such that $|W'_i - W'_{i-1}| \leq 1$. Let $v_{Z'}^2 = \sum_i \mathbb{E}[|Z'_i|^2 | \mathcal{F}_{<i}]$. Applying Theorem 18, there exists a constant $C_2 > 0$ such that for any $\delta_2, v > 0$,

$$\Pr \left(\left| \sum Z'_i \right| \geq \delta_2 \kappa N \wedge v_{Z'}^2 \leq v^2 N \right) \leq e^{-C_2 \delta_2 \kappa \ln(\delta_2 \kappa / v^2) N + \delta_2 \kappa N}. \quad (23)$$

For $(m, t, k) \notin \mathcal{B}'$, using that T_i is independent from M_i and $\mathbb{E}[T_i] = \kappa$ it holds that $v_{Z'}^2 \leq \delta'_1 \kappa N$. Let $v^2 = \delta'_1 \kappa N$ and assume δ'_1 small enough, as a function of δ_2 , such that the right-hand side in (23) is less than $2^{-2\kappa N}$. Let $\delta'_2 = \delta_2 + \gamma/\kappa$ and

$$\mathcal{B}'' = \left\{ (m, t, k) \notin \mathcal{B}' : \sum_i \mathbb{E}[M_i T_i | (M, T, K)_{<i} = (m, t, k)_{<i}] \leq (1 - \delta'_2) \kappa N \right\}.$$

Then similarly to (22) we get

$$\sum_{m \in \mathcal{G}} \sum_t \Pr(T = t) \left(\sum_{k: (m, t, k) \in \mathcal{B}''} \sqrt{\Pr((M, K) = (m, k) | T = t)} \right)^2 \leq 2^{-\kappa N}. \quad (24)$$

Third step: conditional expectations of $(1 - M)TK$. [Uses the assumption: $g(\delta'_2) \ll \delta_3$.]

Using assumption (19), the fact that the measurement operators M_i and K_i commute, and concavity of g , for any $(m, t, k) \notin (\mathcal{B}'' \cup \mathcal{B}')$ it holds that

$$\sum_i \mathbb{E}[(1 - M_i) T_i K_i | (M, T, K)_{<i} = (m, t, k)_{<i}] \leq g(\delta'_2) \kappa N. \quad (25)$$

For $i \in \{1, \dots, N\}$ let $Z_i'' = (1 - M_i) T_i K_i - \mathbb{E}[(1 - M_i) T_i K_i | \mathcal{F}_{<i}]$ and $W_i'' = Z_1'' + \dots + Z_i''$. Then the sequence (W_1'', \dots, W_N'') is a martingale such that $|W_i'' - W_{i-1}''| \leq 1$ and by (25), $v_{Z''}^2 = \sum_i \mathbb{E}[|Z_i''|^2 | \mathcal{F}_{<i}] \leq g(\delta'_2) \kappa N$. Applying Theorem 18 once more, for any $\delta_3 > 0$,

$$\Pr \left(\left| \sum Z_i'' \right| \geq \delta_3 \kappa N \wedge \overline{\mathcal{B}'' \cup \mathcal{B}'} \right) \leq e^{-C_2 \delta_3 \kappa \ln(\delta_3 / g(\delta'_2)) N + \delta_3 \kappa N}.$$

Assume δ'_2 small enough, as a function of δ_3 , such that the right-hand side is less than $2^{-2\kappa N}$. Assume further that $\delta_3 + \gamma/\kappa \leq \eta$. Let $\mathcal{B}''' = \overline{\mathcal{B}'' \cup \mathcal{B}'} \cap \mathcal{B}$. Then it follows as in (22), (24) that

$$\sum_{m \in \mathcal{G}} \sum_t \Pr(T = t) \left(\sum_{k: (m, t, k) \in \mathcal{B}'''} \sqrt{\Pr((M, K) = (m, k) | T = t)} \right)^2 \leq 2^{-\kappa N}. \quad (26)$$

Combining (22), (24) and (26) with the triangle inequality proves the lemma. \square

To apply Lemma 17 to our context, for strings $k, t \in \{0, 1\}^N$ we define $|\phi^{ctok}\rangle = P_{ct}^{ok} |\phi\rangle$, where recall that P_{ct}^{ok} is notation for the operator that corresponds to applying the device's measurement operators indicated by c, t and obtaining outcomes o, k . For a parameter $\eta > 0$, for any t let $|\tilde{\phi}^{cto}\rangle = \sum_{k: |k \wedge t| \leq \eta \kappa q_N} |\phi^{ctok}\rangle$, and $\tilde{\phi}^{cto}$ the sub-normalized density obtained by taking partial trace of $|\tilde{\phi}^{cto}\rangle$ over the adversary.

Proposition 19. *Let $D = (\phi, \Pi, M, K)$ be a simplified device such that condition (15) from Proposition 16 holds. Then for any $0 < \gamma, \kappa, \eta < 1$ such that $\gamma/\kappa \ll 1$ and $\kappa \ll \eta^2$,*

$$\sum_{g, c} q(g, c) \sum_t \kappa(t) \sum_{o: (g, c, o) \in \text{Acc}} \|\phi^{co} - \tilde{\phi}^{cto}\|_1 = O(2^{-\kappa N}). \quad (27)$$

Proof. The state $|\phi\rangle_{\text{DE}}$ and the sequence of measurements $\{M_1^0, \dots, M_{q_N}^0\}$ and $\{K_1^0, \dots, K_{q_N}^0\}$ obtained by restricting to those indices $i \in \{0, \dots, N\}$ such that $c_i = 0$, together with the function $g(x) = C\sqrt{x}$ implied by (15), satisfy the assumptions of Lemma 17. The conclusion (20) of the lemma gives 27. \square

The next step is to relate the randomness in the states $\tilde{\phi}^{cto}$ to randomness in the states $\tilde{\phi}^{ctok}$, for k such that $|k \wedge t|_H \leq \eta\kappa qN$, as these are the post-measurement states associated with the simplified device in Protocol 2. This relies on the following variant of the Cauchy-Schwarz inequality.

Lemma 20. *Let $\ell \geq 1$ be an integer and $|v_1\rangle, \dots, |v_\ell\rangle$ arbitrary vectors in \mathbb{C}^d . Then*

$$\left(\sum_{i=1}^{\ell} |v_i\rangle \right) \left(\sum_{i=1}^{\ell} |v_i\rangle \right)^\dagger \leq \ell \sum_{i=1}^{\ell} |v_i\rangle \langle v_i|.$$

Using the lemma, we can show the following.

Proposition 21. *Let D be a simplified device, and $\tilde{\phi}_{cto}$ the ensemble of states associated with D as described in the beginning of Section 7.1. Then*

$$\sum_{g,c,t} q(g,c)\kappa(t) \sum_{o: (g,c,o) \in \text{ACC}} \langle \tilde{\phi}^{cto} \rangle_{1+\varepsilon} \leq 2^{O(H(\eta)\kappa qN)} \sum_{g,c,t} q(g,c)\kappa(t) \sum_{k,o: (g,c,t,o,k) \in \text{ACC}_2} \langle \tilde{\phi}^{ctok} \rangle_{1+\varepsilon},$$

where ACC_2 denotes the set of transcripts that are accepted by the verifier in Protocol 2.

Proof. The proposition follows from the definition of $\tilde{\rho}^{ctok}$, Lemma 20, and the fact that for any t such that $|t| \leq \kappa qN$ there are at most $2^{H(\eta)\kappa qN}$ sequences k such that $k_i = 1$ whenever $t_i = 0$ and $|k \wedge t|_H \leq \eta\kappa qN$. Note that the conditions that $(g,c,o) \in \text{ACC}$ and $|k \wedge t| \leq \eta\kappa qN$ imply $(g,c,t,o,k) \in \text{ACC}_2$. \square

7.2 Randomness accumulation in the simplified protocol

In this section we consider the behavior of a simplified device $D = (\phi, \Pi, M, K)$ in a single round of Protocol 2. The following lemma shows that, provided the device has overlap $\Delta(D)$ bounded away from 1, then if the state ϕ of the device has high overlap with the projection operator M^1 , necessarily performing a measurement of $\{\Pi^0, \Pi^1, \Pi^2\}$ on ϕ perturbs the state (and hence generates randomness). The proof is based on a ‘‘measurement-disturbance trade-off’’ from [MS14], itself a consequence of uniform convexity for certain matrix p -norms.

Lemma 22. *Let $D = (\phi, \Pi, M, K)$ be a simplified device with overlap $\Delta(D) \leq \omega$, for some $\omega < 1$. Let $0 \leq \varepsilon \leq \frac{1}{2}$ and*

$$t = \frac{\langle \phi_G \rangle_{1+\varepsilon}}{\langle \phi \rangle_{1+\varepsilon}}, \quad \text{where } G = \frac{1}{2}(\Pi^0 + \Pi^1) + \frac{1}{2}M^1K^0 \quad \text{and} \quad \phi_G = \sqrt{G}\phi\sqrt{G}. \quad (28)$$

Then

$$\frac{\langle \phi_1^0 \rangle_{1+\varepsilon} + \langle \phi_1^1 \rangle_{1+\varepsilon} + \langle \phi_1^2 \rangle_{1+\varepsilon}}{\langle \phi \rangle_{1+\varepsilon}} \leq 2^{-\varepsilon\lambda_\omega(t)} + O(\varepsilon^2),$$

where the post-measurement states ϕ_1^v , $v \in \{0, 1, 2\}$, are introduced in (10), and

$$\lambda_\omega(t) = 2 \log(e) \left(t - \frac{1}{2} - \frac{\omega}{2} \right)^2 \quad (29)$$

if $t \geq \frac{1}{2} + \frac{\omega}{2}$, and 0 otherwise.

Proof. The proof uses ideas from [MS14]. Let ϕ be as in the lemma and $\phi' = \sum_v \Pi^v \phi \Pi^v$. Then for any i ,

$$\begin{aligned} \langle \sum_v \sqrt{G} \Pi^v \phi \Pi^v \sqrt{G} \rangle_{1+\varepsilon} &\leq \sum_v \langle \phi^{1/2} \Pi^v G \Pi^v \phi^{1/2} \rangle_{1+\varepsilon} + O(\varepsilon) \\ &\leq \left(\frac{1}{2} + \frac{\omega}{2} \right) \langle \phi^{1/2} (\Pi^0 + \Pi^1) \phi^{1/2} \rangle_{1+\varepsilon} + \frac{1}{2} \langle \phi^{1/2} \Pi^2 \phi^{1/2} \rangle_{1+\varepsilon} + O(\varepsilon) \\ &\leq \left(\frac{1}{2} + \frac{\omega}{2} \right) \langle \phi' \rangle_{1+\varepsilon} + O(\varepsilon), \end{aligned}$$

where the first and last lines use the approximate linearity relations (4), and the second line uses the definition of K and of G . This allows us to proceed as in the proof of [MS14, Theorem 6.3] to obtain

$$\langle \phi - \phi' \rangle_{1+\varepsilon} \geq 2 \left(t - \frac{1}{2} - \frac{\omega}{2} \right) \langle \phi \rangle_{1+\varepsilon} - O(\varepsilon),$$

and conclude by applying [MS16, Proposition 5.3]. \square

Using Lemma 22 we proceed to quantify the accumulation of randomness across multiple rounds of the simplified protocol, when it is executed with a simplified device that has overlap bounded away from 1. The following proposition provides a measure of the randomness present in the transcript, conditioned on the verifier not aborting the protocol at the end, i.e. on $(g, c, t, o, k) \in \text{ACC}_2$. (To see the connection with entropy, recall the definition of the $(1 + \varepsilon)$ conditional Rényi entropy in Definition 4. The connection will be made precise in Section 7.3.)

Proposition 23. *Let $D = (\phi, \Pi, M, K)$ be a simplified device such that $\Delta(D) \leq \omega$ for some $\omega < 1$. Let $0 < \varepsilon \leq \frac{1}{2}$. Let $\gamma, \eta, \kappa, q > 0$ and N an integer be parameters for an execution of Protocol 2 (Figure 2) with D . Then, using notation introduced around (11), for any $\varepsilon > 0$,*

$$-\frac{1}{\varepsilon N} \log \left(\frac{\sum_{(g,c,t,o,k) \in \text{ACC}_2} q(g, c) \kappa(t) \langle \phi^{ctok} \rangle_{1+\varepsilon}}{\langle \phi \rangle_{1+\varepsilon}} \right) \geq \lambda_\omega \left(1 - \frac{\gamma}{\kappa} - \eta \right) - O\left(\kappa q + \frac{\varepsilon}{\kappa q} \right), \quad (30)$$

where λ_ω is the function defined in (29).

Proof. Let $t = \frac{\langle \phi_G \rangle_{1+\varepsilon}}{\langle \phi \rangle_{1+\varepsilon}}$ be as defined in Lemma 22. Recall the notation for the post-measurement states introduced in (10). After one round of Protocol 2 is executed, the post-measurement state of the device can be decomposed into three components. First, in case $G_i = 1$, which happens with probability $(1 - q)$, the round is a generation round. The randomness generated in such round is captured by the bound from Lemma 22,

$$(1 - q) (\langle \phi_1^0 \rangle_{1+\varepsilon} + \langle \phi_1^1 \rangle_{1+\varepsilon} + \langle \phi_1^2 \rangle_{1+\varepsilon}) \leq (1 - q) (1 - \ln(2) \varepsilon \lambda_\omega(t) + O(\varepsilon^2)) \langle \phi \rangle_{1+\varepsilon}. \quad (31)$$

The second case corresponds to $G_i = 0$, which happens with probability q . In this case, for reasons that will become clear later in this proof we weigh the “success” and “failure” components of the post-measurement state differently. For the “failure” part we simply write

$$\frac{q}{2} ((1 - \kappa) \langle \phi_{00}^0 \rangle_{1+\varepsilon} + \kappa \langle \phi_{01}^{00} \rangle_{1+\varepsilon} + \kappa \langle \phi_{01}^{01} \rangle_{1+\varepsilon} + \kappa \langle \phi_{01}^{11} \rangle_{1+\varepsilon} + \langle \phi_1^2 \rangle_{1+\varepsilon}), \quad (32)$$

where the states are labeled as ϕ_{ct}^{ok} (with the outcome k omitted in case $t = 0$). Finally, for the “success” part we add a weight of $2^{\frac{\varepsilon s}{\kappa q}}$, where s is a real parameter to be determined later, to the cases where $T_i = 1$:

$$\begin{aligned} & \frac{(1-\kappa)q}{2} (\langle \phi_{00}^1 \rangle_{1+\varepsilon} + \langle \phi_1^0 \rangle_{1+\varepsilon} + \langle \phi_1^1 \rangle_{1+\varepsilon}) + \frac{\kappa q}{2} 2^{\frac{\varepsilon s}{\kappa q}} (\langle \phi_{01}^{10} \rangle_{1+\varepsilon} + \langle \phi_1^1 \rangle_{1+\varepsilon} + \langle \phi_1^2 \rangle_{1+\varepsilon}) \\ & \leq \frac{(1-\kappa)q}{2} (\langle \phi_{00}^1 \rangle_{1+\varepsilon} + \langle \phi_1^0 \rangle_{1+\varepsilon} + \langle \phi_1^1 \rangle_{1+\varepsilon}) + \kappa q \left(1 + \ln(2) \frac{\varepsilon s}{\kappa q} + O\left(\frac{\varepsilon^2}{\kappa^2 q^2}\right) \right) t \langle \phi \rangle_{1+\varepsilon}, \end{aligned} \quad (33)$$

where the inequality follows directly from the definition of t . Using the first inequality in (4) and regrouping terms, the sum of the left-hand sides of (31), (32) and (33) is at most

$$(31) + (32) + (33) \leq \left(1 - \varepsilon \ln(2) \left(\lambda_\omega(t) - st + O\left(\kappa q + \frac{\varepsilon}{\kappa q}\right) \right) \right) \langle \phi \rangle_{1+\varepsilon}. \quad (34)$$

A convenient choice of s is to take the derivative $s = \lambda'_\omega(r)$ for some $r \in [0, 1]$ to be determined. With this choice, using that λ_ω is convex it follows that $\min_{t \in [0, 1]} \lambda_\omega(t) - st = \lambda_\omega(r) - \lambda'_\omega(r)r$. By chaining the inequality (34) N times, where at each step the density ϕ is updated with the one obtained from the previous round, and using that ACC_2 contains those sequences (g, c, t, o, k) such that the number of occurrences of $(c, t, o, k) \in \{(0, 1, 1, 0), (1, *, 0, *), (1, *, 1, *)\}$ is at least $(1 - \gamma/\kappa - \eta)\kappa q N$ we obtain

$$\begin{aligned} -\frac{1}{\varepsilon N} \log \left(\frac{\sum_{(g,c,t,o,k) \in \text{ACC}_2} q(g,c) \kappa(t) \langle \phi^{ctok} \rangle_{1+\varepsilon}}{\langle \phi \rangle_{1+\varepsilon}} \right) & \geq (\lambda_\omega(r) - \lambda'_\omega(r)r) + \left(1 - \frac{\gamma}{\kappa} - \eta \right) \lambda'_\omega(r) \\ & \quad - O\left(\kappa q + \frac{\varepsilon}{\kappa q}\right), \end{aligned}$$

with the term $(1 - \frac{\gamma}{\kappa} - \eta)\lambda'_\omega(r)$ coming from the acceptance criterion. Choosing $r = (1 - \frac{\gamma}{\kappa} - \eta)$ completes the proof. \square

7.3 Randomness accumulation in the general protocol

In this section we combine the results obtained in the previous two sections to analyze the randomness generated in Protocol 1. The main step is given in the following proposition.

Proposition 24. *Let $D = (\phi, \Pi, M)$ be an efficient device. Let $|\phi\rangle_{DE}$ denote an arbitrary purification of ϕ_D , and $\bar{\rho}_{COE}$ the joint state of the verifier’s choice of challenges, the outputs computed by the verifier, and the adversary’s system E , restricted to transcripts that are accepted by the verifier in the protocol.⁸ Then there is an $\eta \in \text{negl}(\lambda)$ and a constant $C > 0$ such that for any $\delta > 0$,*

$$\frac{1}{N} H_\infty^{N\eta+\delta}(O|CE)_{\bar{\rho}} \geq \lambda_\omega(1 - C\sqrt{\gamma}) - O\left(q + \frac{1 + \log(2/\delta)}{\gamma^{1/4}qN}\right). \quad (35)$$

Proof. Let $\tilde{D} = (\phi, \tilde{\Pi}, \tilde{M}, K)$ be the elementary device obtained by applying Proposition 16 to the device D , for a choice of $\omega = \frac{3}{4}$. Using that the bound in (35) only considers registers C and O (the transcript) and E , by Proposition 19 for any choice of parameters κ, η such that $\kappa \ll \eta^2$ and $\eta\kappa \gg \gamma$, up to an additional error $O(2^{-\kappa q N})$ in the smoothing parameter it suffices to show the bound claimed in (35) for the states $\{\bar{\rho}_E^{cto}\}$ from Proposition 19, where t is an additional input that can be conditioned on. Applying Proposition 21, it

⁸The state $\bar{\rho}$ is sub-normalized.

suffices to prove an analogous bound for the states $\{\tilde{\rho}^{ctok}\}$ generated by \tilde{D} in Protocol 2, at the cost of the prefactor of $2^{H(\eta)\kappa qN}$, that will induce a corresponding loss in entropy.

Let $\tilde{\phi} = \phi^{\frac{1}{1+\varepsilon}}$, where $\varepsilon > 0$ is a small parameter to be specified later. We apply Proposition 23 to \tilde{D} , with ϕ replaced by $\tilde{\phi}$. Then (30) gives

$$-\frac{1}{\varepsilon N} \log \left(\frac{\sum_{(g,c,t,o,k) \in \text{ACC}_2} q(g,c)\kappa(t) \langle \tilde{\phi}^{ctok} \rangle_{1+\varepsilon}}{\langle \tilde{\phi} \rangle_{1+\varepsilon}} \right) \geq \lambda_\omega \left(1 - \frac{\gamma}{\kappa} - \eta \right) - O\left(\kappa q + \frac{\varepsilon}{\kappa q} \right). \quad (36)$$

Note that all states $\tilde{\phi}^{ctko}$ that arise from the combined application of Proposition 19 and Proposition 21 satisfy the acceptance criterion for Protocol 2, so the condition $(g,c,t,o,k) \in \text{ACC}_2$ on the left-hand side of (36) is redundant. From Proposition 21, (36) gives

$$\mathbb{E}_t \left(-\frac{1}{\varepsilon N} \log \left(\frac{\sum_{(g,c,o) \in \text{ACC}} q(g,c) \langle \tilde{\phi}^{cto} \rangle_{1+\varepsilon}}{\langle \tilde{\phi} \rangle_{1+\varepsilon}} \right) \right) \geq \lambda_\omega \left(1 - \frac{\gamma}{\kappa} - \eta \right) - O\left(H(\eta)\kappa \frac{q}{\varepsilon} + \kappa q + \frac{\varepsilon}{\kappa q} \right), \quad (37)$$

where the expectation over t was taken out of the $-\log$ using convexity. We make one ultimate re-writing step. For any fixed t , the post-measurement state $\tilde{\phi}^{cto}$ can be expressed as

$$P_N \cdots P_1 \tilde{\phi} P_1 \cdots P_N,$$

where P_i is the measurement operator associated with challenge c_i and outcome o_i . Using $\langle XX^* \rangle_{1+\varepsilon} = \langle X^* X \rangle_{1+\varepsilon}$ for any X , and recalling the definition of $\tilde{\phi} = \phi^{\frac{1}{1+\varepsilon}}$,

$$\langle P_N \cdots P_1 \tilde{\phi} P_1 \cdots P_N \rangle_{1+\varepsilon} = \langle \phi^{\frac{-\varepsilon}{2(1+\varepsilon)}} \phi^{\frac{1}{2}} P_1 \cdots P_N^2 \cdots P_1 \phi^{\frac{1}{2}} \phi^{\frac{-\varepsilon}{2(1+\varepsilon)}} \rangle_{1+\varepsilon}.$$

Introduce a sub-normalized density

$$\rho_{\mathbb{E}}^{cto} = \phi^{\frac{1}{2}} P_1 \cdots P_N^2 \cdots P_1 \phi^{\frac{1}{2}},$$

that corresponds to the post-measurement state of register \mathbb{E} (recall we assumed a purification $|\phi\rangle_{\text{DE}}$ of ϕ) at the end of Protocol 1, for a given transcript (c,o) for the interaction.

We are in a position to apply Theorem 7, with

$$\rho_{\text{CTOE}}^o = \sum_{(g,c,t): (g,c,o) \in \text{ACC}} q(g,c)\kappa(t) |c,t\rangle\langle c,t|_{\text{CT}} \otimes |o\rangle\langle o|_{\text{O}} \otimes \rho_{\mathbb{E}}^{cto},$$

and $\sigma_{\text{CTE}} = \sum_{(g,c,t)} q(g,c)\kappa(t) |c,t\rangle\langle c,t| \otimes \phi$. Applying the theorem and using (36) and $\langle \tilde{\phi} \rangle_{1+\varepsilon} = 1$ by definition, we get that for any $\delta > 0$,

$$\frac{1}{N} H_\infty^\delta(\text{O}|\text{CTE})_{\tilde{\rho}} \geq \lambda_\omega \left(1 - \frac{\gamma}{\kappa} - \eta \right) - O\left(H(\eta)\kappa \frac{q}{\varepsilon} + \kappa q + \frac{\varepsilon}{\kappa q} \right) - \frac{1 + 2 \log(1/\delta)}{\varepsilon N}.$$

Choose κ, η to be sufficiently large constant multiples of $\sqrt{\gamma}$ and ε to be a sufficiently small constant multiple of $\gamma^{3/4}q$. Adding the contribution of $O(2^{-\kappa q N})$ to the smoothing parameter that comes from the initial application of Proposition 19, and discarding the register T , gives the result. \square

Combining Proposition 24 with Proposition 16 we obtain our main result.

Theorem 25. Let \mathcal{F} be an NTCF family and λ a security parameter. Let N be a polynomially bounded function of λ such that $N = \Omega(\lambda^2)$. Set $q = \lambda/N$. Then there is a negligible (as a function of λ) δ such that for any small enough $\gamma > 0$, any efficient prover, and side information E correlated with the prover's initial state,

$$H_{\infty}^{N\delta}(O|CE)_{\bar{p}} \geq (\xi - o(1))N,$$

where \bar{p} is the final state of the output, challenge, and adversary registers, restricted to transcripts that are accepted by the verifier in the protocol, ξ is a positive constant,⁹ and $o(1)$ is a function that goes to 0 as $\lambda \rightarrow \infty$.

Assume that an execution of $\text{GEN}(1^\lambda)$ requires $O(\lambda^r)$ bits of randomness, for some constant r . (For example, for the case of our construction of a NTCF family based on LWE, we have $r = 2$.) Then an execution of the protocol using the parameters in Theorem 25 requires only $\text{poly}(\lambda, \log N)$ bits of randomness for the verifier to generate the key k and select the challenges. Taking N to be slightly sub-exponential in λ , e.g. $N = 2^{\sqrt{\lambda}}$, yields sub-exponential randomness expansion.

Proof of Theorem 25. Let D be a device that is accepted with non-negligible probability in the protocol, where the parameters are as stated in the theorem. Applying Proposition 24 to D and choosing δ to be a negligible function of N such that δ^{-1} is sub-exponential gives the result. \square

8 A Trapdoor Claw-Free family based on LWE

In this section we present our LWE-based construction of an NTCF. For LWE-related preliminaries and definitions see Appendix A. Let λ be a security parameter. All other parameters are functions of λ . Let $q \geq 2$ be a prime integer. Let $\ell, n, m, w \geq 1$ be polynomially bounded functions of λ and B_L, B_V, B_P be positive integers such that the following conditions hold:

1. $n = \Omega(\ell \log q)$ and $m = \Omega(n \log q)$,
2. $w = n \lceil \log q \rceil$,
3. $B_P = \frac{q}{2C_T \sqrt{mn \log q}}$, for C_T the universal constant in Theorem 36, (38)
4. $2\sqrt{n} \leq B_L < B_V < B_P$,
5. The ratio $\frac{B_P}{B_V}$ and $\frac{B_V}{B_L}$ are both super-polynomial in λ .

Given a choice of parameters satisfying all 5 conditions in (38), we describe the function family \mathcal{F}_{LWE} . Let $\mathcal{X} = \mathbb{Z}_q^n$ and $\mathcal{Y} = \mathbb{Z}_q^m$. The key space is $\mathcal{K}_{\mathcal{F}_{\text{LWE}}} = \mathbb{Z}_q^{m \times n} \times \mathbb{Z}_q^m$. For $b \in \{0, 1\}$, $x \in \mathcal{X}$ and key $k = (\mathbf{A}, \mathbf{A}s + \mathbf{e})$, the density $f_{k,b}(x)$ is defined as

$$\forall y \in \mathcal{Y}, \quad (f_{k,b}(x))(y) = D_{\mathbb{Z}_q^m, B_P}(y - \mathbf{A}x - b \cdot \mathbf{A}s), \quad (39)$$

where the density $D_{\mathbb{Z}_q^m, B_P}$ is defined in (73). Note that $f_{k,b}$ is well-defined given $k = (\mathbf{A}, \mathbf{A}s + \mathbf{e})$, as for our choice of parameters k uniquely identifies s .

The four properties required for a noisy trapdoor claw-free family, as specified in Definition 8, are verified in the following subsections, providing a proof of the following theorem. Recall the definition of the hardness assumption $\text{LWE}_{n,q,\chi}$ given in Definition 35.

⁹The constant ξ is at least some positive universal constant of order $1/10$, for all small enough γ .

Theorem 26. For any choice of parameters satisfying the conditions (38), the function family \mathcal{F}_{LWE} is a noisy trapdoor claw free family under the hardness assumption $\text{LWE}_{\ell, q, D_{\mathbb{Z}_q, B_L}}$.

8.1 Efficient Function Generation

$\text{GEN}_{\mathcal{F}_{\text{LWE}}}$ is defined as follows. First, the procedure samples a random $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$, together with trapdoor information $t_{\mathbf{A}}$. This is done using the procedure $\text{GENTRAP}(1^n, 1^m, q)$ from Theorem 36. The trapdoor allows the evaluation of an inversion algorithm INVERT that, on input \mathbf{A} , $t_{\mathbf{A}}$ and $b = \mathbf{A}\mathbf{s} + \mathbf{e}$ returns \mathbf{s} and \mathbf{e} as long as $\|\mathbf{e}\| \leq \frac{q}{C_T \sqrt{n \log q}}$. Moreover, the distribution on matrices \mathbf{A} returned by GENTRAP is negligibly close to the uniform distribution on $\mathbb{Z}_q^{m \times n}$.

Next, the sampling procedure selects $s \in \{0, 1\}^n$ uniformly at random, and a vector $\mathbf{e} \in \mathbb{Z}_q^m$ by sampling each coordinate independently according to the distribution $D_{\mathbb{Z}_q, B_V}$ defined in (73). $\text{GEN}_{\mathcal{F}_{\text{LWE}}}$ returns $k = (\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e})$ and $t_k = t_{\mathbf{A}}$.

8.2 Trapdoor Injective Pair

- (a) *Trapdoor.* It follows from (39) and the definition of the distribution $D_{\mathbb{Z}_q^m, B_P}$ in (73) that for any key $k = (\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e}) \in \mathcal{K}_{\mathcal{F}_{\text{LWE}}}$ and for all $x \in \mathcal{X}$,

$$\text{SUPP}(f_{k,0}(x)) = \{\mathbf{A}x + \mathbf{e}_0 \mid \|\mathbf{e}_0\| \leq B_P \sqrt{m}\}, \quad (40)$$

$$\text{SUPP}(f_{k,1}(x)) = \{\mathbf{A}x + \mathbf{A}\mathbf{s} + \mathbf{e}_0 \mid \|\mathbf{e}_0\| \leq B_P \sqrt{m}\}. \quad (41)$$

The procedure $\text{INV}_{\mathcal{F}_{\text{LWE}}}$ takes as input the trapdoor $t_{\mathbf{A}}$, $b \in \{0, 1\}$, and $y \in \mathcal{Y}$. It uses the algorithm INVERT to determine $\mathbf{s}_0, \mathbf{e}_0$ such that $y = \mathbf{A}\mathbf{s}_0 + \mathbf{e}_0$, and returns the element $\mathbf{s}_0 - b \cdot \mathbf{s} \in \mathcal{X}$. Using Theorem 36, this procedure returns the unique correct outcome provided $y = \mathbf{A}\mathbf{s}_0 + \mathbf{e}_0$ for some \mathbf{e}_0 such that $\|\mathbf{e}_0\| \leq \frac{q}{C_T \sqrt{n \log q}}$. This condition is satisfied for all $y \in \text{SUPP}(f_{k,b}(x))$ provided B_P is chosen so that

$$B_P \leq \frac{q}{C_T \sqrt{mn \log q}}. \quad (42)$$

- (b) *Injective Pair.* We let \mathcal{R}_k be the set of all pairs (x_0, x_1) such that $f_{k,0}(x_0) = f_{k,1}(x_1)$. By definition this occurs if and only if $x_1 = x_0 - \mathbf{s}$, and so \mathcal{R}_k is a perfect matching.

8.3 Efficient Range Superposition

For $k = (\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e}) \in \mathcal{K}_{\mathcal{F}_{\text{LWE}}}$, $b \in \{0, 1\}$ and $x \in \mathcal{X}$, let

$$(f'_{k,b}(x))(y) = D_{\mathbb{Z}_q^m, B_P}(y - \mathbf{A}x - b \cdot (\mathbf{A}\mathbf{s} + \mathbf{e})). \quad (43)$$

Note that $f'_{k,0}(x) = f_{k,0}(x)$ for all $x \in \mathcal{X}$. The distributions $f'_{k,1}(x)$ and $f_{k,1}(x)$ are shifted by \mathbf{e} . Given the key k and $x \in \mathcal{X}$, the densities $f'_{k,0}(x)$ and $f'_{k,1}(x)$ are efficiently computable. For all $x \in \mathcal{X}$,

$$\text{SUPP}(f'_{k,0}(x)) = \text{SUPP}(f_{k,0}(x)), \quad (44)$$

$$\text{SUPP}(f'_{k,1}(x)) = \{\mathbf{A}x + \mathbf{e}_0 + \mathbf{A}\mathbf{s} + \mathbf{e} \mid \|\mathbf{e}_0\| \leq B_P \sqrt{m}\}. \quad (45)$$

- (a) Using that $B_V < B_P$, it follows that the norm of the term $\mathbf{e}_0 + \mathbf{e}$ in (45) is always at most $2B_P\sqrt{m}$. Therefore, the inversion procedure $\text{INV}_{\mathcal{F}_{\text{LWE}}}$ can be guaranteed to return x on input $t_{\mathbf{A}}, b \in \{0, 1\}$, $y \in \text{SUPP}(f'_{k,b}(x))$ if we strengthen the requirement on B_P given in (42) to

$$B_P \leq \frac{q}{2C_T \sqrt{mn \log q}}. \quad (46)$$

This strengthened trapdoor requirement also implies that for all $b \in \{0, 1\}$, $(x_0, x_1) \in \mathcal{R}_k$, and $y \in \text{SUPP}(f'_{k,b}(x_b))$, $\text{INV}_{\mathcal{F}_{\text{LWE}}}(t_{\mathbf{A}}, b \oplus 1, y) = x_{b \oplus 1}$.

- (b) On input $k = (\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e})$, $b \in \{0, 1\}$, $x \in \mathcal{X}$, and $y \in \mathcal{Y}$, the procedure $\text{CHK}_{\mathcal{F}_{\text{LWE}}}$ operates as follows. If $b = 0$, it computes $\mathbf{e}' = y - \mathbf{A}x$. If $\|\mathbf{e}'\| \leq B_P\sqrt{m}$, the procedure returns 1, and 0 otherwise. If $b = 1$, it computes $\mathbf{e}' = y - \mathbf{A}x - (\mathbf{A}\mathbf{s} + \mathbf{e})$. If $\|\mathbf{e}'\| \leq B_P\sqrt{m}$, it returns 1, and 0 otherwise.
- (c) We bound the Hellinger distance between the densities $f_{k,b}(x)$ and $f'_{k,b}(x)$. If $b = 0$ they are identical. If $b = 1$, both densities are shifts of $D_{\mathbb{Z}_q^m, B_P}$, where the shifts differ by \mathbf{e} . Each coordinate of \mathbf{e} is drawn independently from $D_{\mathbb{Z}_q, B_V}$, so $\|\mathbf{e}\| \leq \sqrt{m}B_V$. Applying Lemma 34, we get that

$$H^2(f_{k,1}(x), f'_{k,1}(x)) \leq 1 - e^{-\frac{2\pi m B_V}{B_P}}.$$

Using the assumption that B_P/B_V is super-polynomial, this is negligible, as desired. It remains to describe the procedure $\text{SAMP}_{\mathcal{F}_{\text{LWE}}}$. At the first step, the procedure creates the following superposition

$$\sum_{\mathbf{e}_0 \in \mathbb{Z}_q^m} \sqrt{D_{\mathbb{Z}_q^m, B_P}(\mathbf{e}_0)} |\mathbf{e}_0\rangle. \quad (47)$$

This state can be prepared efficiently as described in [Reg05, Lemma 3.12].¹⁰

At the second step, the procedure creates a uniform superposition over $x \in \mathcal{X}$, yielding the state

$$q^{-\frac{n}{2}} \sum_{\substack{x \in \mathcal{X} \\ \mathbf{e}_0 \in \mathbb{Z}_q^m}} \sqrt{D_{\mathbb{Z}_q^m, B_P}(\mathbf{e}_0)} |x\rangle |\mathbf{e}_0\rangle. \quad (48)$$

At the third step, using the key $k = (\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e})$ and the input bit b the procedure computes

$$q^{-\frac{n}{2}} \sum_{\substack{x \in \mathcal{X} \\ \mathbf{e}_0 \in \mathbb{Z}_q^m}} \sqrt{D_{\mathbb{Z}_q^m, B_P}(\mathbf{e}_0)} |x\rangle |\mathbf{e}_0\rangle |\mathbf{A}x + \mathbf{e}_0 + b \cdot (\mathbf{A}\mathbf{s} + \mathbf{e})\rangle. \quad (49)$$

At this point, observe that \mathbf{e}_0 can be computed from x , the last register, b and the key k . The procedure

¹⁰Specifically, the state can be created using a technique by Grover and Rudolph ([GR02]), who show that in order to create such a state, it suffices have the ability to efficiently compute the sum $\sum_{x=c}^d D_{\mathbb{Z}_q, B_P}(x)$ for any $c, d \in \{-\lfloor \sqrt{B_P} \rfloor, \dots, \lceil \sqrt{B_P} \rceil\} \subseteq \mathbb{Z}_q$ and to within good precision. This can be done using standard techniques used in sampling from the normal distribution.

can then uncompute the register containing \mathbf{e}_0 , yielding

$$\begin{aligned}
& q^{-\frac{n}{2}} \sum_{\substack{x \in \mathcal{X} \\ \mathbf{e}_0 \in \mathbb{Z}_q^m}} \sqrt{D_{\mathbb{Z}_q^m, B_P}(\mathbf{e}_0) |x\rangle | \mathbf{A}x + \mathbf{e}_0 + b \cdot (\mathbf{A}\mathbf{s} + \mathbf{e}) \rangle} \\
&= q^{-\frac{n}{2}} \sum_{x \in \mathcal{X}, y \in \mathcal{Y}} \sqrt{D_{\mathbb{Z}_q^m, B_P}(y - \mathbf{A}x - b \cdot (\mathbf{A}\mathbf{s} + \mathbf{e})) |x\rangle |y\rangle} \\
&= q^{-\frac{n}{2}} \sum_{x \in \mathcal{X}, y \in \mathcal{Y}} \sqrt{(f'_{k,b}(x))(y) |x\rangle |y\rangle}. \tag{50}
\end{aligned}$$

8.4 Adaptive Hardcore Bit

This section is devoted to the proof of the adaptive hardcore bit condition. The main statement is provided in Lemma 32 in Section 8.4.3. The proof of the lemma proceeds in three steps. First, in Section 8.4.1 we establish some preliminary results on the distribution of the inner product $(\hat{d} \cdot s \bmod 2)$, where $\hat{d} \in \{0, 1\}^n$ is a fixed nonzero binary vector and $s \leftarrow_U \{0, 1\}^n$ a uniformly random binary vector, conditioned on $\mathbf{C}\mathbf{s} = \mathbf{v}$ for some randomly chosen matrix $\mathbf{C} \in \mathbb{Z}_q^{\ell \times n}$ and arbitrary $\mathbf{v} \in \mathbb{Z}_q^\ell$. This condition is combined with the LWE assumption in Section 8.4.2 to argue that $(\hat{d} \cdot s \bmod 2)$ remains computationally close to uniform even when the matrix \mathbf{C} is an LWE matrix \mathbf{A} , and the adversary is able to choose \hat{d} after being given access to $\mathbf{A}\mathbf{s} + \mathbf{e}$ for some error vector $\mathbf{e} \in \mathbb{Z}_q^m$. Finally, in Section 8.4.3 the required hardcore bit condition is reduced to the one established in Section 8.4.2 by relating the inner product appearing in the definition of H_k (in condition 4(b) of Definition 8) to an inner product of the form $\hat{d} \cdot s$, where \hat{d} can be efficiently computed from d .

8.4.1 Moderate matrices

The following lemma argues that, provided the matrix $\mathbf{C} \in \mathbb{Z}_q^{\ell \times n}$ is a uniformly random matrix with sufficiently few rows, the distribution $(\mathbf{C}, \mathbf{C}\mathbf{s})$ for arbitrary $s \in \{0, 1\}^n$ does not reveal any parity of s .

Lemma 27. *Let q be a prime, $\ell, n \geq 1$ be integers, and $\mathbf{C} \in \mathbb{Z}_q^{\ell \times n}$ be a uniformly random matrix. With probability at least $1 - q^\ell \cdot 2^{-\frac{n}{8}}$ over the choice of \mathbf{C} the following holds: for a fixed \mathbf{C} , all $\mathbf{v} \in \mathbb{Z}_q^\ell$ and $\hat{d} \in \{0, 1\}^n \setminus \{0^n\}$, the distribution of $(\hat{d} \cdot s \bmod 2)$, where s is uniform in $\{0, 1\}^n$ conditioned on $\mathbf{C}\mathbf{s} = \mathbf{v}$, is within statistical distance $O(q^{\frac{3\ell}{2}} \cdot 2^{-\frac{n}{40}})$ of the uniform distribution over $\{0, 1\}$.*

To prove the lemma we introduce the notion of a *moderate* matrix.

Definition 28. Let $\mathbf{b} \in \mathbb{Z}_q^n$. We say that \mathbf{b} is moderate if it contains at least $\frac{n}{4}$ entries whose unique representative in $(-q/2, q/2]$ has its absolute value in the range $(\frac{q}{8}, \frac{3q}{8}]$. A matrix $\mathbf{C} \in \mathbb{Z}_q^{\ell \times n}$ is moderate if its entire row span (except 0^n) is moderate.

Lemma 29. *Let q be prime and ℓ, n be integers. Then*

$$\Pr_{\mathbf{C} \leftarrow_U \mathbb{Z}_q^{\ell \times n}} [\mathbf{C} \text{ is moderate}] \geq 1 - q^\ell \cdot 2^{-\frac{n}{8}}.$$

Proof. Consider an arbitrary non zero vector \mathbf{b} in the row-span of a uniform \mathbf{C} . Then the marginal distribution of \mathbf{b} is uniform. By Chernoff, \mathbf{b} is moderate with probability at least $1 - e^{-\frac{2n}{16}} \geq 1 - 2^{-\frac{n}{8}}$. Applying the union bound over all at most $q^\ell - 1$ non zero vectors in the row span, the result follows. \square

Lemma 30. Let $\mathbf{C} \in \mathbb{Z}_q^{\ell \times n}$ be an arbitrary moderate matrix and let $\hat{d} \in \{0, 1\}^n \setminus \{0^n\}$ be an arbitrary non zero binary vector. Let s be uniform over $\{0, 1\}^n$ and consider the random variables $\mathbf{v} = \mathbf{C}s \bmod q$ and $z = \hat{d} \cdot s \bmod 2$. Then (\mathbf{v}, z) is within total variation distance at most $q^{\frac{\ell}{2}} \cdot 2^{-\frac{n}{40}}$ of the uniform distribution over $\mathbb{Z}_q^\ell \times \{0, 1\}$.

Proof. Let f be the probability density function of (\mathbf{v}, z) . Interpreting z as an element of \mathbb{Z}_2 , let \hat{f} be the Fourier transform over $\mathbb{Z}_q^\ell \times \mathbb{Z}_2$. Let U denote the density of the uniform distribution over $\mathbb{Z}_q^\ell \times \mathbb{Z}_2$. Applying the Cauchy-Schwarz inequality,

$$\begin{aligned} \frac{1}{2} \|f - U\|_1 &\leq \sqrt{\frac{q^\ell}{2}} \|f - U\|_2 \\ &= \frac{1}{2} \|\hat{f} - \hat{U}\|_2 \\ &= \frac{1}{2} \left(\sum_{(\hat{\mathbf{v}}, \hat{z}) \in \mathbb{Z}_q^\ell \times \mathbb{Z}_2 \setminus \{(0,0)\}} |\hat{f}(\hat{\mathbf{v}}, \hat{z})|^2 \right)^{1/2}, \end{aligned} \quad (51)$$

where the second line follows from Parseval's identity, and for the third line we used $\hat{f}(\mathbf{0}, 0) = \hat{U}(\mathbf{0}, 0) = 1$ and $\hat{U}(\hat{\mathbf{v}}, \hat{z}) = 0$ for all $(\hat{\mathbf{v}}, \hat{z}) \neq (0^\ell, 0)$. To bound (51) we estimate the Fourier coefficients of f . Denoting $\omega_{2q} = e^{-\frac{2\pi i}{2q}}$, for any $(\hat{\mathbf{v}}, \hat{z}) \in \mathbb{Z}_q^\ell \times \mathbb{Z}_2$ we can write

$$\begin{aligned} \hat{f}(\hat{\mathbf{v}}, \hat{z}) &= \mathbb{E}_{\mathbf{s}} \left[\omega_{2q}^{(2 \cdot \hat{\mathbf{v}}^T \mathbf{C} + q \cdot \hat{z} \hat{\mathbf{d}}^T) \mathbf{s}} \right] \\ &= \mathbb{E}_{\mathbf{s}} \left[\omega_{2q}^{\mathbf{w}^T \mathbf{s}} \right] \\ &= \prod_i \mathbb{E}_{s_i} \left[\omega_{2q}^{w_i s_i} \right], \end{aligned} \quad (52)$$

where we wrote $\mathbf{w}^T = 2 \cdot \hat{\mathbf{v}}^T \mathbf{C} + q \cdot \hat{z} \hat{\mathbf{d}}^T \in \mathbb{Z}_{2q}^n$. It follows that $\hat{f}(0^\ell, 1) = 0$, since $(d \cdot s \bmod 2)$ is uniform for s uniform.

We now observe that for all $i \in \{1, \dots, n\}$ such that the representative of $(\hat{\mathbf{v}}^T \mathbf{C})_i$ in $(-q/2, q/2]$ has its absolute value in $(\frac{q}{8}, \frac{3q}{8}]$ it holds that $\frac{w_i}{q} \in (\frac{1}{4}, \frac{3}{4}] \bmod 1$, in which case

$$\left| \mathbb{E}_{s_i} [\omega_{2q}^{w_i s_i}] \right| = \left| \cos \left(\frac{\pi}{2} \cdot \frac{w_i}{q} \right) \right| \leq \cos \left(\frac{\pi}{8} \right) \leq 2^{-\frac{1}{10}}. \quad (53)$$

Since \mathbf{C} is moderate, there are at least $\frac{n}{4}$ such entries, so that from (52) it follows that $|\hat{f}(\hat{\mathbf{v}}, \hat{z})| \leq 2^{-\frac{n}{40}}$ for all $\hat{\mathbf{v}} \neq \mathbf{0}$. Recalling (51), the lemma is proved. \square

We now prove Lemma 27 by generalizing Lemma 30 to adaptive d (i.e. d can depend on $\mathbf{C}, \mathbf{C}s$).

Proof of Lemma 27. We assume \mathbf{C} is moderate; by Lemma 29, \mathbf{C} is moderate with probability at least $1 - q^\ell \cdot 2^{-\frac{n}{8}}$. Let s be uniform over $\{0, 1\}^n$, $D_1 = (\mathbf{C}s, \hat{d} \cdot s \bmod 2)$, and D_2 uniformly distributed over $\mathbb{Z}_q^\ell \times \{0, 1\}$. Using that \mathbf{C} is moderate, it follows from Lemma 30 that

$$\epsilon = \|D_1 - D_2\|_{TV} \leq q^{\frac{\ell}{2}} \cdot 2^{-\frac{n}{40}}. \quad (54)$$

Fix $\mathbf{v}_0 \in \mathbb{Z}_q^\ell$ and let

$$\Delta = \frac{1}{2} \sum_{b \in \{0,1\}} \left| \Pr_{s \leftarrow_U \{0,1\}^n} [\hat{d} \cdot s \bmod 2 = b \mid \mathbf{Cs} = \mathbf{v}_0] - \frac{1}{2} \right|. \quad (55)$$

To prove the lemma it suffices to establish the appropriate upper bound on Δ , for all \mathbf{v}_0 . By definition,

$$\begin{aligned} \epsilon = \|D_1 - D_2\|_{TV} &= \frac{1}{2} \sum_{b \in \{0,1\}, v \in \mathbb{Z}_q^\ell} \left| \Pr[\mathbf{Cs} = \mathbf{v}] \Pr[\hat{d} \cdot s \bmod 2 = b \mid \mathbf{Cs} = \mathbf{v}] - \frac{1}{2q^\ell} \right| \\ &\geq \frac{1}{2} \sum_{b \in \{0,1\}} \left| \Pr[\mathbf{Cs} = \mathbf{v}_0] \Pr[\hat{d} \cdot s \bmod 2 = b \mid \mathbf{Cs} = \mathbf{v}_0] - \frac{1}{2q^\ell} \right| \\ &= \frac{1}{2} \sum_{b \in \{0,1\}} \left| \Pr[\mathbf{Cs} = \mathbf{v}_0] \left(\frac{1}{2} + (-1)^b \Delta \right) - \frac{1}{2q^\ell} \right|, \end{aligned} \quad (56)$$

where all probabilities are under a uniform choice of $s \leftarrow_U \{0,1\}^n$, and the last line follows from the definition of Δ in (55). Applying the inequality $|a| + |b| \geq \max(|a-b|, |a+b|)$, valid for any real a, b , to (56) it follows that

$$\Pr[\mathbf{Cs} = \mathbf{v}_0] \cdot \Delta \leq \epsilon \quad \text{and} \quad \Pr[\mathbf{Cs} = \mathbf{v}_0] \geq \frac{1}{q^\ell} - 2\epsilon. \quad (57)$$

If $q^{3\ell/2} 2^{-\frac{n}{40}} > \frac{1}{3}$ the bound claimed in the lemma is trivial. If $q^{3\ell/2} 2^{-\frac{n}{40}} \leq \frac{1}{3}$, then $\epsilon q^\ell \leq \frac{1}{3}$, so it follows from (57) that $\Delta \leq 3q^\ell \epsilon$, which together with (54) proves the lemma. \square

8.4.2 LWE Hardcore bit

The next step is to use Lemma 27 to obtain a form of the hardcore bit statement that is appropriate for our purposes. In this section we use the notation introduced in Appendix A. We also use the following notation: we write $s \in \{0,1\}^n$ as $s = (s_0, s_1)$, where $s_0, s_1 \in \{0,1\}^{\frac{n}{2}}$ are the $\frac{n}{2}$ -bit prefix and suffix of s respectively (for simplicity, assume n is even; if not, ties can be broken arbitrarily). We will show computational indistinguishability based on the hardness assumption $\text{LWE}_{\ell, q, D_{\mathbb{Z}_q, B_L}}$ (described at the start of Appendix A).

For reasons that will become clear in the next section, we consider procedures that output a tuple $(b, x, d, c) \in \{0,1\} \times \mathcal{X} \times \{0,1\}^w \times \{0,1\}$.

Lemma 31. *Assume a choice of parameters satisfying the conditions (38). Assume the hardness assumption $\text{LWE}_{\ell, q, D_{\mathbb{Z}_q, B_L}}$ holds. Let*

$$\mathcal{A} : \mathbb{Z}_q^{m \times n} \times \mathbb{Z}_q^m \rightarrow \{0,1\} \times \mathcal{X} \times \{0,1\}^w \times \{0,1\}$$

be a quantum polynomial-time procedure. For $b \in \{0,1\}$ and $x \in \mathcal{X}$ let $I_{b,x} : \{0,1\}^w \rightarrow \{0,1\}^n$ be an efficiently computable map. For every $s = (s_0, s_1) \in \{0,1\}^n$ and $(b, x) \in \{0,1\} \times \mathcal{X}$, let $\hat{G}_{s_{b \oplus 1}, b, x} \subseteq \{0,1\}^w$ be a set depending only on b, x and $s_{b \oplus 1}$ and such that for all $d \in \hat{G}_{s_{b \oplus 1}, b, x}$ the first (if $b = 0$) or last (if $b = 1$) $\frac{n}{2}$ bits of $I_{b,x}(d)$ are not all 0. Then the distributions

$$D_0 = ((\mathbf{A}, \mathbf{As} + \mathbf{e}) \leftarrow \text{GEN}_{\mathcal{F}_{\text{LWE}}}(1^\lambda), (b, x, d, c) \leftarrow \mathcal{A}(\mathbf{A}, \mathbf{As} + \mathbf{e}), I_{b,x}(d) \cdot s \bmod 2) \quad (58)$$

and

$$D_1 = ((\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e}) \leftarrow \text{GEN}_{\mathcal{F}_{\text{LWE}}}(1^\lambda), (b, x, d, c) \leftarrow \mathcal{A}(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e}), (\delta_{d \in \hat{G}_{s_b \oplus 1, b, x}} r) \oplus (I_{b, x}(d) \cdot s \pmod{2})), \quad (59)$$

where $r \leftarrow_U \{0, 1\}$ and $\delta_{d \in \hat{G}_{s_b \oplus 1, b, x}}$ is 1 if $d \in \hat{G}_{s_b \oplus 1, b, x}$ and 0 otherwise, are computationally indistinguishable.

Proof. We prove computational indistinguishability by introducing a sequence of hybrids. For the first step we let

$$D^{(1)} = ((\tilde{\mathbf{A}}, \tilde{\mathbf{A}}\mathbf{s} + \mathbf{e}), (b, x, d, c) \leftarrow \mathcal{A}(\tilde{\mathbf{A}}, \tilde{\mathbf{A}}\mathbf{s} + \mathbf{e}), I_{b, x}(d) \cdot s \pmod{2}), \quad (60)$$

where $\tilde{\mathbf{A}} = \mathbf{BC} + \mathbf{F} \leftarrow \text{LOSSY}(1^n, 1^m, 1^\ell, q, D_{\mathbb{Z}_q, B_L})$ is sampled from a lossy sampler (see Definition 37). From the definition, $\mathbf{F} \in \mathbb{Z}_q^{m \times n}$ has entries i.i.d. from the distribution $D_{\mathbb{Z}_q, B_L}$ over \mathbb{Z}_q . To see that D_0 and $D^{(1)}$ are computationally indistinguishable, first note that the distribution of matrices \mathbf{A} generated by $\text{GEN}_{\mathcal{F}_{\text{LWE}}}$ is negligibly far from the uniform distribution (see Theorem 36). Next, by Theorem 38, under the $\text{LWE}_{\ell, q, D_{\mathbb{Z}_q, B_L}}$ assumption a uniformly random matrix \mathbf{A} and a lossy matrix $\tilde{\mathbf{A}}$ are computationally indistinguishable. Note that this step, as well as subsequent steps, uses that \mathcal{A} and $I_{b, x}$ are efficiently computable.

For the second step we remove the term $\mathbf{F}\mathbf{s}$ from the lossy LWE sample $\tilde{\mathbf{A}}\mathbf{s} + \mathbf{e}$ to obtain the distribution

$$D^{(2)} = ((\mathbf{BC} + \mathbf{F}, \mathbf{BC}\mathbf{s} + \mathbf{e}), (b, x, d, c) \leftarrow \mathcal{A}(\mathbf{BC} + \mathbf{F}, \mathbf{BC}\mathbf{s} + \mathbf{e}), I_{b, x}(d) \cdot s \pmod{2}). \quad (61)$$

Using that \mathbf{s} is binary and the entries of \mathbf{F} are taken from a B_L -bounded distribution, it follows that $\|\mathbf{F}\mathbf{s}\| \leq n\sqrt{m}B_L$. Applying Lemma 34, it follows that the statistical distance between $D^{(1)}$ and $D^{(2)}$ is at most

$$\gamma = \sqrt{2} \left(1 - e^{-\frac{2\pi mn B_L}{B_V}} \right)^{1/2}, \quad (62)$$

which is negligible, due to the requirement that $\frac{B_V}{B_L}$ is superpolynomial given in (38).

For the third step, observe that the distribution $D^{(2)}$ in (61) only depends on s_b through $\mathbf{C}\mathbf{s}$ and $I_{b, x}(d) \cdot s$, where \mathbf{C} is uniformly random. It follows from Lemma 27 that provided $\frac{n}{2} = \Omega(\ell \log q)$, with overwhelming probability over the choice of \mathbf{C} , if we fix all variables except for s_b , the resulting distribution of $(I_{b, x}(d) \cdot s \pmod{2})$ is statistically indistinguishable from $r \leftarrow_U \{0, 1\}$ as long as the $\frac{n}{2}$ bits of $I_{b, x}(d)$ associated with s_b are not all 0 (i.e. the first $\frac{n}{2}$ bits if $b = 0$ or the last $\frac{n}{2}$ bits if $b = 1$). Using that since $d \in \hat{G}_{s_b \oplus 1, b, x}$, by assumption the $\frac{n}{2}$ bits of $I_{b, x}(d)$ associated with s_b are not all 0, the distribution $D^{(2)}$ in (61) is statistically indistinguishable from

$$D^{(3)} = ((\mathbf{BC} + \mathbf{F}, \mathbf{BC}\mathbf{s} + \mathbf{e}), (b, x, d, c) \leftarrow \mathcal{A}(\mathbf{BC} + \mathbf{F}, \mathbf{BC}\mathbf{s} + \mathbf{e}), (\delta_{d \in \hat{G}_{s_b \oplus 1, b, x}} r) \oplus (I_{b, x}(d) \cdot s \pmod{2})),$$

where $r \leftarrow_U \{0, 1\}$ and $\delta_{d \in \hat{G}_{s_b \oplus 1, b, x}}$ is 1 if $d \in \hat{G}_{s_b \oplus 1, b, x}$ and 0 otherwise.

For the fourth step we reinsert the term $\mathbf{F}\mathbf{s}$ to obtain

$$D^{(4)} = ((\tilde{\mathbf{A}}, \tilde{\mathbf{A}}\mathbf{s} + \mathbf{e}), (b, x, d, c) \leftarrow \mathcal{A}(\tilde{\mathbf{A}}, \tilde{\mathbf{A}}\mathbf{s} + \mathbf{e}), (\delta_{d \in \hat{G}_{s_b \oplus 1, b, x}} r) \oplus (I_{b, x}(d) \cdot s \pmod{2})).$$

Statistical indistinguishability between $D^{(3)}$ and $D^{(4)}$ follows similarly as between $D^{(1)}$ and $D^{(2)}$. Finally, computational indistinguishability between $D^{(4)}$ and D_1 follows similarly to between $D^{(1)}$ and D_0 . \square

8.4.3 Adaptive hardcore lemma

We now prove that condition 4 of Definition 8 holds. Recall that $\mathcal{X} = \mathbb{Z}_q^n$ and $w = n \lceil \log q \rceil$. Let $J : \mathcal{X} \rightarrow \{0, 1\}^w$ be such that $J(x)$ returns the binary representation of $x \in \mathcal{X}$. For $b \in \{0, 1\}$, $x \in \mathcal{X}$, and $d \in \{0, 1\}^w$, let $I_{b,x}(d) \in \{0, 1\}^n$ be the vector whose each coordinate is obtained by taking the inner product mod 2 of the corresponding block of $\lceil \log q \rceil$ coordinates of d and of $J(x) \oplus J(x - (-1)^b \mathbf{1})$, where $\mathbf{1} \in \mathbb{Z}_q^n$ is the vector with all its coordinates equal to 1 in \mathbb{Z}_q . For $k = (\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e})$, $b \in \{0, 1\}$ and $x \in \mathcal{X}$, we define the set $G_{k,b,x}$ as

$$G_{k,b,x} = \left\{ d \in \{0, 1\}^w \mid \exists i \in \left\{ b \frac{n}{2}, \dots, b \frac{n}{2} + \frac{n}{2} \right\} : (I_{b,x}(d))_i \neq 0 \right\}.$$

Observe that for all $b \in \{0, 1\}$ and $x \in \mathcal{X}$, if d is sampled uniformly at random, $d \notin G_{k,b,x}$ with negligible probability. This follows simply because for any $b \in \{0, 1\}$, $J(x) \oplus J(x - (-1)^b \mathbf{1})$ is non-zero, since J is injective. Observe also that checking membership in $G_{k,b,x}$ is possible given only b, x . This shows condition 4(a) in the adaptive hardcore bit condition in Definition 8.

Given $(x_0, x_1) \in \mathcal{R}_k$ (where $k = (\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e})$), recall from Section 8.2 that $x_1 = x_0 - \mathbf{s}$. For convenience we also introduce the following set, where $y = f_{k,0}(x_0) = f_{k,1}(x_1)$:

$$\hat{G}_y = G_{k,0,x_0} \cap G_{k,1,x_1}. \quad (63)$$

The following lemma establishes item 4(b) in the adaptive hardcore bit condition of Definition 8.

Lemma 32. *Assume a choice of parameters satisfying the conditions (38). Assume the hardness assumption $\text{LWE}_{\ell,q,D_{\mathbb{Z}_q},B_L}$ holds. Let $s \in \{0, 1\}^n$. Let ¹¹*

$$H_s = \left\{ (b, x, d, d \cdot (J(x) \oplus J(x - (-1)^b \mathbf{s}))) \mid b \in \{0, 1\}, x \in \mathcal{X}, d \in \hat{G}_{f_{k,b}(x)} \right\}, \quad (64)$$

$$\overline{H}_s = \left\{ (b, x, d, c) \mid (b, x, d, c \oplus 1) \in H_s \right\}. \quad (65)$$

Then for any quantum polynomial-time procedure

$$\mathcal{A} : \mathbb{Z}_q^{m \times n} \times \mathbb{Z}_q^m \rightarrow \{0, 1\} \times \mathcal{X} \times \{0, 1\}^w \times \{0, 1\}$$

there exists a negligible function $\mu(\lambda)$ such that

$$\left| \Pr_{(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e}) \leftarrow \text{GEN}_{\mathcal{F}_{\text{LWE}}}(1^\lambda)} [\mathcal{A}(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e}) \in H_s] - \Pr_{(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e}) \leftarrow \text{GEN}_{\mathcal{F}_{\text{LWE}}}(1^\lambda)} [\mathcal{A}(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e}) \in \overline{H}_s] \right| \leq \mu(\lambda). \quad (66)$$

Proof. The proof is by contradiction. Assume that there exists a quantum polynomial-time procedure \mathcal{A} such that the left-hand side of (66) is at least some non-negligible function $\eta(\lambda)$. We derive a contradiction by showing that the two distributions D_0 and D_1 in Lemma 31, for $I_{b,x}$ defined at the start of this section and $\hat{G}_{s_{b \oplus 1}, b, x}$ defined in (63), are computationally distinguishable, giving a contradiction.

Let $(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e}) \leftarrow \text{GEN}_{\mathcal{F}_{\text{LWE}}}(1^\lambda)$ and $(b, x, d, c) \leftarrow \mathcal{A}(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e})$. To link \mathcal{A} to the distributions in Lemma 31 we relate the inner product condition in (64) to an inner product $\hat{d} \cdot s$ of the form appearing in (58), for $\hat{d} = I_{b,x}(d)$ that can be computed efficiently from b, x and d . This is based on the following claim.

¹¹We write the sets as H_s instead of H_k to emphasize the dependence on s . Also, recall the definition of \mathcal{R}_k in Section 8.2 and the definition of \hat{G}_y in (63), that relates it to the set $G_{k,b,x}$ appearing in the definition of the adaptive hardcore bit condition.

Claim 33. For all $b \in \{0, 1\}, x \in \mathcal{X}, d \in \{0, 1\}^w$ and $s \in \{0, 1\}^n$ the following equality holds:

$$d \cdot (J(x) \oplus J(x - (-1)^b \mathbf{s})) = I_{b,x_b}(d) \cdot s. \quad (67)$$

Moreover, the function $I_{b,x}$ is efficiently computable given b, x .

Proof. We do the proof in case $n = 1$ and $w = \lceil \log q \rceil$, as the case of general n follows by linearity. In this case s is a single bit. If $s = 0$ then both sides of (67) evaluate to zero, so the equality holds trivially. It then suffices to define $I_{b,x_b}(d)$ so that the equation holds when $s = 1$. A choice of either of

$$I_{0,x_0}(d) = d \cdot (J(x_0) \oplus J(x_0 - \mathbf{1})), \quad I_{1,x_1}(d) = d \cdot (J(x_1) \oplus J(x_1 + \mathbf{1})) \quad (68)$$

satisfies all requirements. It is clear from the definition of $I_{b,x}$ (given at the start of this section) that it can be computed efficiently given b, x . \square

The procedure \mathcal{A} , the function $I_{b,x}$ defined at the start of this section and the sets \hat{G}_y in (63) fully specify D_0 and D_1 . To conclude we construct a distinguisher \mathcal{A}' between D_0 and D_1 . Consider two possible distinguishers, \mathcal{A}'_u for $u \in \{0, 1\}$. Given a sample $((\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e}), (b, x, d, c), t)$, \mathcal{A}'_u returns 0 if $c = t \oplus u$, and 1 otherwise. First note that:

$$\begin{aligned} & \sum_{u \in \{0,1\}} \left| \Pr_{w=(b,x,d,c) \leftarrow D_0} [\mathcal{A}'_u(w) = 0] - \Pr_{w=(b,x,d,c) \leftarrow D_1} [\mathcal{A}'_u(w) = 0] \right| \quad (69) \\ &= \sum_{u \in \{0,1\}} \left| \Pr_{w=(b,x,d,c) \leftarrow D_0} [\mathcal{A}'_u(w) = 0 \wedge d \in \hat{G}_{f_{k,b}(x)}] - \Pr_{w=(b,x,d,c) \leftarrow D_1} [\mathcal{A}'_u(w) = 0 \wedge d \in \hat{G}_{f_{k,b}(x)}] \right| \end{aligned}$$

since if $d \notin \hat{G}_{f_{k,b}(x)}$, the distributions D_0 and D_1 are identical by definition. Next, if the sample held by \mathcal{A}'_u is from the distribution D_0 and if $(b, x, d, c) \in H_s$, then by the definition of H_s and (67) and it follows that $c = d \cdot (J(x) \oplus J(x - (-1)^b \mathbf{s})) = I_{b,x}(d) \cdot s = t$. If instead $(b, x, d, c) \in \overline{H}_s$ then $c \oplus 1 = d \cdot (J(x) \oplus J(x - (-1)^b \mathbf{s})) = I_{b,x}(d) \cdot s = t$. Thus we obtain that the expression in (70) is equal to:

$$= \left| \Pr_{(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e}) \leftarrow \text{GEN}_{\mathcal{F}_{\text{LWE}}}(1^\lambda)} [\mathcal{A}(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e}) \in H_s] - \frac{1}{2} \Pr_{w=(b,x,d,c) \leftarrow D_1} [d \in \hat{G}_{f_{k,b}(x)}] \right| \quad (71)$$

$$+ \left| \Pr_{(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e}) \leftarrow \text{GEN}_{\mathcal{F}_{\text{LWE}}}(1^\lambda)} [\mathcal{A}(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e}) \in \overline{H}_s] - \frac{1}{2} \Pr_{w=(b,x,d,c) \leftarrow D_1} [d \in \hat{G}_{f_{k,b}(x)}] \right| \quad (72)$$

$$\begin{aligned} & \geq \left| \Pr_{(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e}) \leftarrow \text{GEN}_{\mathcal{F}_{\text{LWE}}}(1^\lambda)} [\mathcal{A}(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e}) \in H_s] - \Pr_{(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e}) \leftarrow \text{GEN}_{\mathcal{F}_{\text{LWE}}}(1^\lambda)} [\mathcal{A}(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e}) \in \overline{H}_s] \right| \\ & \geq \eta \end{aligned}$$

Therefore, at least one of \mathcal{A}'_0 or \mathcal{A}'_1 must successfully distinguish between D_0 and D_1 with advantage at least $\frac{\eta}{2}$, a contradiction with the statement of Lemma 31. \square

A Learning With Errors

In this section we give some background on the Learning with Errors problem (LWE). For a positive real B and positive integers q , the truncated discrete Gaussian distribution over \mathbb{Z}_q with parameter B is supported on $\{x \in \mathbb{Z}_q : \|x\| \leq B\}$ and has density

$$D_{\mathbb{Z}_q, B}(x) = \frac{e^{-\frac{\pi\|x\|^2}{B^2}}}{\sum_{x \in \mathbb{Z}_q, \|x\| \leq B} e^{-\frac{\pi\|x\|^2}{B^2}}}. \quad (73)$$

We note that for any $B > 0$, the truncated and non-truncated distributions have statistical distance that is exponentially small in B [Ban93, Lemma 1.5]. For a positive integer m , the truncated discrete Gaussian distribution over \mathbb{Z}_q^m with parameter B is supported on $\{x \in \mathbb{Z}_q^m : \|x\| \leq B\sqrt{m}\}$ and has density

$$\forall x = (x_1, \dots, x_m) \in \mathbb{Z}_q^m, \quad D_{\mathbb{Z}_q^m, B}(x) = D_{\mathbb{Z}_q, B}(x_1) \cdots D_{\mathbb{Z}_q, B}(x_m). \quad (74)$$

Lemma 34. *Let B be a positive real number and q, m be positive integers. Let $\mathbf{e} \in \mathbb{Z}_q^m$. The Hellinger distance between the distribution $D = D_{\mathbb{Z}_q^m, B}$ and the shifted distribution $D + \mathbf{e}$ satisfies*

$$H^2(D, D + \mathbf{e}) \leq 1 - e^{-\frac{2\pi\sqrt{m}\|\mathbf{e}\|}{B}}, \quad (75)$$

and the statistical distance between the two distributions satisfies

$$\|D - (D + \mathbf{e})\|_{TV}^2 \leq 2\left(1 - e^{-\frac{2\pi\sqrt{m}\|\mathbf{e}\|}{B}}\right). \quad (76)$$

Proof. Let $\tau = \sum_{x \in \mathbb{Z}_q, \|x\| \leq B} e^{-\frac{\pi\|x\|^2}{B^2}}$. We can compute

$$\begin{aligned} \sum_{\mathbf{e}_0 \in \mathbb{Z}_q^m} \sqrt{D_{\mathbb{Z}_q^m, B}(\mathbf{e}_0) D_{\mathbb{Z}_q^m, B}(\mathbf{e}_0 - \mathbf{e})} &= \sum_{\mathbf{e}_0 \in \mathbb{Z}_q^m} \sqrt{D_B(\mathbf{e}_0) D_B(\mathbf{e}_0 - \mathbf{e})} \\ &= \frac{1}{\tau^m} \sum_{\mathbf{e} \in \mathbb{Z}_q^m} e^{-\frac{\pi(\|\mathbf{e}_0\|^2 + \|\mathbf{e}_0 - \mathbf{e}\|^2)}{2B^2}} \\ &\geq \frac{1}{\tau^m} \sum_{\mathbf{e}_0 \in \mathbb{Z}_q^m} e^{-\frac{\pi(\|\mathbf{e}_0\|^2 + (\|\mathbf{e}_0\| + \|\mathbf{e}\|)^2)}{2B^2}} \\ &= \frac{1}{\tau^m} \sum_{\mathbf{e}_0 \in \mathbb{Z}_q^m} e^{-\frac{\pi(\|\mathbf{e}_0\|^2)}{B^2}} e^{-\frac{\pi(2\|\mathbf{e}_0\|\|\mathbf{e}\|)}{2B^2}} e^{-\frac{\pi(\|\mathbf{e}\|^2)}{2B^2}} \\ &\geq e^{-\frac{\pi(\|\mathbf{e}\|^2 + 2\|\mathbf{e}_0\|\|\mathbf{e}\|)}{2B^2}} \frac{1}{\tau^m} \sum_{\mathbf{e}_0 \in \mathbb{Z}_q^m} e^{-\frac{\pi(\|\mathbf{e}_0\|^2)}{B^2}} \\ &= e^{-\frac{\pi(\|\mathbf{e}\|^2 + 2\|\mathbf{e}_0\|\|\mathbf{e}\|)}{2B^2}} \\ &\geq e^{-\frac{2\pi\|\mathbf{e}_0\|\|\mathbf{e}\|}{B^2}}. \end{aligned}$$

Using the fact that for any \mathbf{e}_0 in the support of $D_{\mathbb{Z}_q^m, B}$, $\|\mathbf{e}_0\| \leq B\sqrt{m}$, gives the claimed bound. \square

Definition 35. For a security parameter λ , let $n, m, q \in \mathbb{N}$ be integer functions of λ . Let $\chi = \chi(\lambda)$ be a distribution over \mathbb{Z} . The $\text{LWE}_{n,m,q,\chi}$ problem is to distinguish between the distributions $(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e} \pmod{q})$ and (\mathbf{A}, \mathbf{u}) , where \mathbf{A} is uniformly random in $\mathbb{Z}_q^{n \times m}$, \mathbf{s} is a uniformly random row vector in \mathbb{Z}_q^n , \mathbf{e} is a row vector drawn at random from the distribution χ^m , and \mathbf{u} is a uniformly random vector in \mathbb{Z}_q^m . Often we consider the hardness of solving LWE for any function m such that m is at most a polynomial in $n \log q$. This problem is denoted $\text{LWE}_{n,q,\chi}$. In this paper we make the assumption that no quantum polynomial-time procedure can solve the $\text{LWE}_{n,q,\chi}$ problem with more than a negligible advantage in λ , even when given access to a quantum polynomial-size state depending on the parameters n, m, q and χ of the problem. We refer to this assumption as “the $\text{LWE}_{n,q,\chi}$ assumption”.

As shown in [Reg05, PRS17], for any $\alpha > 0$ such that $\sigma = \alpha q \geq 2\sqrt{n}$ the $\text{LWE}_{n,q,D_{\mathbb{Z}_q,\sigma}}$ problem, where $D_{\mathbb{Z}_q,\sigma}$ is the discrete Gaussian distribution, is at least as hard as approximating the shortest independent vector problem (SIVP) to within a factor of $\gamma = \tilde{O}(n/\alpha)$ in *worst case* dimension n lattices. This is proven using a quantum reduction. Classical reductions (to a slightly different problem) exist as well [Pei09, BLP⁺13] but with somewhat worse parameters. The best known (classical or quantum) algorithm for these problems run in time $2^{\tilde{O}(n/\log \gamma)}$. For our construction we assume hardness of the problem against a quantum polynomial-time adversary in the case that γ is a super polynomial function in n . This is a commonly used assumption in cryptography (for e.g. homomorphic encryption schemes such as [GSW13]).

We use two additional properties of the LWE problem. The first is that it is possible to generate LWE samples $(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e})$ such that there is a trapdoor allowing recovery of \mathbf{s} from the samples.

Theorem 36 (Theorem 5.1 in [MP11]). *Let $n, m \geq 1$ and $q \geq 2$ be such that $m = \Omega(n \log q)$. There is an efficient randomized algorithm $\text{GENTRAP}(1^n, 1^m, q)$ that returns a matrix $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ and a trapdoor $t_{\mathbf{A}}$ such that the distribution of \mathbf{A} is negligibly (in n) close to the uniform distribution. Moreover, there is an efficient algorithm INVERT that, on input $\mathbf{A}, t_{\mathbf{A}}$ and $\mathbf{A}\mathbf{s} + \mathbf{e}$ where $\|\mathbf{e}\| \leq q / (C_T \sqrt{n \log q})$ and C_T is a universal constant, returns \mathbf{s} and \mathbf{e} with overwhelming probability over $(\mathbf{A}, t_{\mathbf{A}}) \leftarrow \text{GENTRAP}$.*

The second property is the existence of a “lossy mode” for LWE. The following definition is Definition 3.1 in [AKPW13].

Definition 37. Let $\chi = \chi(\lambda)$ be an efficiently sampleable distribution over \mathbb{Z}_q . Define a lossy sampler $\tilde{\mathbf{A}} \leftarrow \text{LOSSY}(1^n, 1^m, 1^\ell, q, \chi)$ by $\tilde{\mathbf{A}} = \mathbf{B}\mathbf{C} + \mathbf{F}$, where $\mathbf{B} \leftarrow_{\mathcal{U}} \mathbb{Z}_q^{m \times \ell}$, $\mathbf{C} \leftarrow_{\mathcal{U}} \mathbb{Z}_q^{\ell \times n}$, $\mathbf{F} \leftarrow \chi^{m \times n}$.

Theorem 38 (Lemma 3.2 in [AKPW13]). *Under the $\text{LWE}_{\ell,q,\chi}$ assumption, the distribution of $\tilde{\mathbf{A}} \leftarrow \text{LOSSY}(1^n, 1^m, 1^\ell, q, \chi)$ is computationally indistinguishable from $\mathbf{A} \leftarrow_{\mathcal{U}} \mathbb{Z}_q^{m \times n}$.*

References

- [Aar18] Scott Aaronson. Certified randomness from quantum supremacy. Personal communication, 2018.
- [AFDF⁺18] Rotem Arnon-Friedman, Frédéric Dupuis, Omar Fawzi, Renato Renner, and Thomas Vidick. Practical device-independent quantum cryptography via entropy accumulation. *Nature communications*, 9(1):459, 2018.
- [Ajt99] Miklós Ajtai. Generating hard instances of the short basis problem. In *International Colloquium on Automata, Languages, and Programming*, pages 1–9. Springer, 1999.

- [AKPW13] Joel Alwen, Stephan Krenn, Krzysztof Pietrzak, and Daniel Wichs. Learning with rounding, revisited: New reduction, properties and applications. Cryptology ePrint Archive, Report 2013/098, 2013. <https://eprint.iacr.org/2013/098>.
- [Ban93] Wojciech Banaszczyk. New bounds in some transference theorems in the geometry of numbers. *Mathematische Annalen*, 296(1):625–635, 1993.
- [BKG⁺18] Peter Bierhorst, Emanuel Knill, Scott Glancy, Yanbao Zhang, Alan Mink, Stephen Jordan, Andrea Rommal, Yi-Kai Liu, Bradley Christensen, Sae Woo Nam, et al. Experimentally generated randomness certified by the impossibility of superluminal signals. *arXiv preprint arXiv:1803.06219*, 2018.
- [BLP⁺13] Zvika Brakerski, Adeline Langlois, Chris Peikert, Oded Regev, and Damien Stehlé. Classical hardness of learning with errors. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *Symposium on Theory of Computing Conference, STOC’13, Palo Alto, CA, USA, June 1-4, 2013*, pages 575–584. ACM, 2013.
- [CCKW18] Alexandru Cojocaru, Léo Colisson, Elham Kashefi, and Petros Wallden. Delegated pseudo-secret random qubit generator. *arXiv preprint arXiv:1802.08759*, 2018.
- [Col06] R. Colbeck. *Quantum And Relativistic Protocols For Secure Multi-Party Computation*. PhD thesis, Trinity College, University of Cambridge, November 2006.
- [FGL12] Xiequan Fan, Ion Grama, and Quansheng Liu. Hoeffdings inequality for supermartingales. *Stochastic Processes and their Applications*, 122(10):3545–3559, 2012.
- [FGS11] S. Fehr, R. Gelles, and C. Schaffner. Security and composability of randomness expansion from Bell inequalities. Technical report arXiv:1111.6052, 2011.
- [GKPV10] Shafi Goldwasser, Yael Tauman Kalai, Chris Peikert, and Vinod Vaikuntanathan. Robustness of the learning with errors assumption. In Andrew Chi-Chih Yao, editor, *ICS*, pages 230–240. Tsinghua University Press, 2010.
- [GMR84] Shafi Goldwasser, Silvio Micali, and Ronald L. Rivest. A ”paradoxical” solution to the signature problem (abstract). In G. R. Blakley and David Chaum, editors, *Advances in Cryptology, Proceedings of CRYPTO ’84, Santa Barbara, California, USA, August 19-22, 1984, Proceedings*, volume 196 of *Lecture Notes in Computer Science*, page 467. Springer, 1984.
- [GR02] Lov Grover and Terry Rudolph. Creating superpositions that correspond to efficiently integrable probability distributions, 2002.
- [GSW13] Craig Gentry, Amit Sahai, and Brent Waters. Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. Cryptology ePrint Archive, Report 2013/340, 2013. <http://eprint.iacr.org/2013/340>.
- [HM17] Aram W Harrow and Ashley Montanaro. Quantum computational supremacy. *Nature*, 549(7671):203, 2017.
- [LP11] Richard Lindner and Chris Peikert. Better key sizes (and attacks) for lwe-based encryption. In *Cryptographers Track at the RSA Conference*, pages 319–339. Springer, 2011.

- [Mah17a] Urmila Mahadev. Classical homomorphic encryption for quantum circuits. *Arxiv preprint arXiv:1708.02130v1*, 2017.
- [Mah17b] Urmila Mahadev. Classical homomorphic encryption for quantum circuits. *arXiv preprint arXiv:1708.02130*, 2017.
- [Mah18] Urmila Mahadev. Classical verification of quantum computations. Manuscript, 2018.
- [MP11] Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. Cryptology ePrint Archive, Report 2011/501, 2011. <http://eprint.iacr.org/2011/501>.
- [MS14] Carl A Miller and Yaoyun Shi. Universal security for randomness expansion from the spot-checking protocol. *arXiv preprint arXiv:1411.6608*, 2014.
- [MS16] Carl A Miller and Yaoyun Shi. Robust protocols for securely expanding randomness and distributing keys using untrusted quantum devices. *Journal of the ACM (JACM)*, 63(4):33, 2016.
- [MU88] Hans Maassen and Jos BM Uffink. Generalized entropic uncertainty relations. *Physical Review Letters*, 60(12):1103, 1988.
- [PAM⁺10] S. Pironio, A. Acin, S. Massar, A. Boyer De La Giroday, D. N. Matsukevich, P. Maunz, S. Olmschenk, D. Hayes, L. Luo, T. A. Manning, and et al. Random numbers certified by Bell’s theorem. *Nature*, 464(7291), 2010.
- [Pei09] Chris Peikert. Public-key cryptosystems from the worst-case shortest vector problem. In *Proceedings of the forty-first annual ACM symposium on Theory of computing*, pages 333–342. ACM, 2009.
- [PM11] S. Pironio and S. Massar. Security of practical private randomness generation. Technical report arXiv:1111.6056, 2011.
- [PRS17] Chris Peikert, Oded Regev, and Noah Stephens-Davidowitz. Pseudorandomness of ring-lwe for any ring and modulus. In Hamed Hatami, Pierre McKenzie, and Valerie King, editors, *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, pages 461–473. ACM, 2017.
- [Reg05] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *STOC*, pages 84–93. ACM, 2005.
- [TCR09] Marco Tomamichel, Roger Colbeck, and Renato Renner. A fully quantum asymptotic equipartition property. *IEEE Transactions on Information Theory*, 55(12):5840–5847, 2009.
- [Tom15] Marco Tomamichel. *Quantum Information Processing with Finite Resources: Mathematical Foundations*, volume 5. Springer, 2015.
- [VV11] Umesh Vazirani and Thomas Vidick. Certifiable quantum dice: or, true random number generation secure against quantum adversaries. In *Proceedings of the 44th symposium on Theory of Computing, STOC ’12*, pages 61–76. ACM, 2011. Also available as arXiv:1111.6054.

[Wil11] Mark M Wilde. From classical to quantum shannon theory. *arXiv preprint arXiv:1106.1445*, 2011.