A Tight High-Order Entropic Quantum Uncertainty Relation With Applications

Ivan B. Damgård[1,⋆], Serge Fehr[2,⋆⋆], Renato Renner[3,⋆⋆⋆,†], Louis Salvail[1,‡], and Christian Schaffner[2,†]

[1] Basic Research in Computer Science (BRICS), funded by the Danish National Research Foundation, Department of Computer Science, University of Aarhus, Denmark, {ivan|salvail}@brics.dk.

[2] Center for Mathematics and Computer Science (CWI), Amsterdam, Netherlands, {fehr|c.schaffner}@cwi.nl

[3] Cambridge University, UK, r.renner@damtp.cam.ac.uk

Abstract. We derive a new entropic quantum uncertainty relation involving min-entropy. The relation is tight and can be applied in various quantum-cryptographic settings.

Protocols for quantum 1-out-of-2 Oblivious Transfer and quantum Bit Commitment are presented and the uncertainty relation is used to prove the security of these protocols in the bounded-quantum-storage model according to new strong security definitions.

As another application, we consider the realistic setting of Quantum Key Distribution (QKD) against quantum-memory-bounded eavesdroppers. The uncertainty relation allows to prove the security of QKD protocols in this setting while tolerating considerably higher error rates compared to the standard model with unbounded adversaries. For instance, for the six-state protocol with one-way communication, a bit-flip error rate of up to 17% can be tolerated (compared to 13% in the standard model).

Our uncertainty relation also yields a lower bound on the min-entropy key uncertainty against known-plaintext attacks when quantum ciphers are composed. Previously, the key uncertainty of these ciphers was only known with respect to Shannon entropy.

1 Introduction

A problem often encountered in quantum cryptography is the following: through some interaction between the players, a quantum state ρ is generated and then measured by one of the players (call her Alice in the following). Assuming Alice is honest, we want to know how unpredictable her measurement outcome

⋆ FICS, Foundations in Cryptography and Security, funded by the Danish Natural Sciences Research Council.

⋆⋆ Supported by the Dutch Organization for Scientific Research (NWO).

⋆⋆⋆ Supported by HP Labs Bristol.

† Supported by the European project SECOQC.

‡ QUSEP, Quantum Security in Practice, funded by the Danish Natural Science Research Council.

is to the adversary. Once a lower bound on the adversary's uncertainty about Alice's measurement outcome is established, it is usually easy to prove the desired security property of the protocol. Many existing constructions in quantum cryptography have been proved secure following this paradigm.

Typically, Alice does not make her measurement in a fixed basis, but chooses at random among a set of different bases. These bases are usually chosen to be pairwise *mutually unbiased*, meaning that if ρ is such that the measurement outcome in one basis is fixed then this implies that the uncertainty about the outcome of the measurement in the other basis is maximal. In this way, one hopes to keep the adversary's uncertainty high, even if ρ is (partially) under the adversary's control.

An inequality that lower bounds the adversary's uncertainty in such a scenario is called an *uncertainty relation*. There exist uncertainty relations for different measures of uncertainty, but cryptographic applications typically require the adversary's min-entropy to be bounded from below.

In this paper, we introduce a new general and tight entropic uncertainty relation. Since the relation is expressed in terms of high-order entropy (i.e. min-entropy), it is applicable to a large class of natural protocols in quantum cryptography. In particular, the new relation can be applied in situations where an n-qubit state ρ has each of its qubits measured in a random and independent basis sampled uniformly from a fixed set $\mathcal{B}$ of bases. $\mathcal{B}$ does not necessarily have to be mutually unbiased, but we assume a lower bound h (i.e. an *average entropic uncertainty bound*) on the average Shannon entropy of the distribution P_ϑ, obtained by measuring an arbitrary 1-qubit state in basis $\vartheta \in \mathcal{B}$, meaning that $\frac{1}{|\mathcal{B}|}\sum_\vartheta \mathrm{H}(P_\vartheta) \geq h$.

Uncertainty Relation (informal): *Let $\mathcal{B}$ be a set of bases with an average entropic uncertainty bound h as above. Let P_θ denote the probability distribution defined by measuring an arbitrary n-qubit state ρ in basis $\theta \in \mathcal{B}^n$. For a $\theta \in_R \mathcal{B}^n$ chosen uniformly at random, it holds except with negligible probability that*

$$\mathrm{H}_\infty(P_\theta) \gtrsim nh \ . \tag{1}$$

Observe that (1) cannot be improved significantly since the min-entropy of a distribution is at most equal to the Shannon entropy. Our uncertainty relation is therefore asymptotically tight when the bound h is tight.

Any lower bound on the Shannon entropy associated to a set of measurements $\mathcal{B}$ can be used in (1). In the special case where the set of bases is $\mathcal{B} = \{+, \times\}$ (i.e. the two BB84 bases), h is known precisely using Maassen and Uffink's entropic relation, see inequality (2) below. We get $h = \frac{1}{2}$ and (1) results in $\mathrm{H}_\infty(P_\theta) \gtrsim \frac{n}{2}$. Uncertainty relations for the BB84 coding scheme [3] are useful since this coding is widely used in quantum cryptography. Its resilience to imperfect quantum channels, sources, and detectors is an important advantage in practice.

We now discuss applications of our high-order uncertainty relation to important scenarios in cryptography: two-party cryptography, quantum key distribution and quantum encryption.

Application I: Two-Party Cryptography in the Bounded-Quantum-Storage Model. Entropic uncertainty relations are powerful tools for the security analysis of cryptographic protocols in the bounded-quantum-storage model. In this model, the adversary is unbounded in every respect, except that at a certain time, his quantum memory is reduced to a certain size (by performing some measurement). In [13], an uncertainty relation involving min-entropy was shown and used in the analysis of protocols for Rabin oblivious transfer (ROT) and bit commitment. This uncertainty relation only applies in the case when n qubits are all measured in one out of two mutually unbiased bases.

A major difference between our result (1) and the one from [13] is that while both relations bound the min-entropy conditioned on an event, this event happens in our case with probability essentially 1 (on average) whereas the corresponding event from [13] only happens with probability about 1/2. In Sect. 4, we prove the following:

1-2 OT in the Bounded-Quantum-Storage Model: *There exists a non-interactive protocol for 1-out-of-2 oblivious transfer (1-2 OT) of ℓ-bit messages, secure against adversaries with quantum memory size at most $n/4 - 2\ell$. Here, n is the number of qubits transmitted in the protocol and ℓ can be a constant fraction of n. Honest players need no quantum memory.*

Since all flavors of OT are known to be equivalent under classical information-theoretic reductions, and a ROT protocol is already known from [13], the above result may seem insignificant. This is not the case, however, for several reasons: First, although it may in principle be possible to obtain a protocol for *1-2 OT* from the ROT protocol of [13] using the standard black-box reduction, the fact that we need to call the ROT primitive many times would force the bound on the adversary's memory to be *sub*linear (in the number of transmitted qubits). Second, the techniques used in [13] do not seem applicable to *1-2 OT*, unless via the inefficient generic reduction to ROT. And, third, we prove security according to a stronger definition than the one used in [13], namely a quantum version of a recent classical definition for information theoretic *1-2 OT* [10]. The definition ensures that all (dishonest) players' inputs are well defined (and can be extracted when formalized appropriately). In particular, this implies security under sequential composition whereas composability of the protocol from [13] was not proven.

Furthermore, our techniques for *1-2 OT* imply almost directly a non-interactive bit commitment scheme (in the bounded-quantum-storage model) satisfying a composable security definition. As an immediate consequence, we obtain secure *string* commitment schemes. This improves over the bit commitment construction of [13], respectively its analysis, which does *not* guarantee composability and thus does *not* necessarily allow for string commitments. This application can be found in Sect. 5.

Application II: Quantum Key Distribution. We also apply our uncertainty relation to quantum key distribution (QKD) settings. QKD is the art of distributing a secret key between two distant parties, Alice and Bob, using only a completely insecure quantum channel and authentic classical communication. QKD protocols typically provide information-theoretic security, i.e., even an adversary with unlimited resources cannot get any information about the key. A major difficulty when implementing QKD schemes is that they require a low-noise quantum channel. The tolerated noise level depends on the actual protocol and on the desired security of the key. Because the quality of the channel typically decreases with its length, the maximum tolerated noise level is an important parameter limiting the maximum distance between Alice and Bob.

We consider a model in which the adversary has a limited amount of quantum memory to store the information she intercepts during the protocol execution. In this model, we show that the maximum tolerated noise level is larger than in the standard scenario where the adversary has unlimited resources. For *one-way QKD protocols* which are protocols where error-correction is performed non-interactively (i.e., a single classical message is sent from one party to the other), we show the following result:

QKD Against Quantum-Memory-Bounded Eavesdroppers: *Let $\mathcal{B}$ be a set of orthonormal bases of $\mathcal{H}_2$ with average entropic uncertainty bound h. Then, a* one-way QKD-protocol *produces a secure key against eavesdroppers whose quantum-memory size is sublinear in the length of the raw key at a positive rate as long as the bit-flip probability p of the quantum channel fulfills* $\mathrm{H}_{\mathrm{bin}}(p) < h$ *where* $\mathrm{H}_{\mathrm{bin}}(\cdot)$ *denotes the binary Shannon-entropy function.*

Although this result does not allow us to improve (i.e. compared to unbounded adversaries) the maximum error-rate for the BB84 protocol (the four-state protocol), the six-state protocol can be shown secure against adversaries with memory bound sublinear in the secret-key length as long as the bit-flip error-rate is less than 17%. This improves over the maximal error-rate of 13% for the same protocol against unbounded adversaries. We also show that the generalization of the six-state protocols to more bases (not necessarily mutually unbiased) can be shown secure (against memory-bounded adversaries) for a maximal error-rate up to 20% provided the number of bases is large enough.

The quantum-memory-bounded eavesdropper model studied here is not comparable to other restrictions on adversaries considered in the literature (e.g. *individual attacks*, where the eavesdropper is assumed to apply independent measurements to each qubit sent over the quantum channel [18, 23]). In fact, these assumptions are generally artificial and their purpose is to simplify security proofs rather than to relax the conditions on the quality of the communication channel from which secure key can be generated. We believe that the quantum-memory-bounded eavesdropper model is more realistic.

Application III: Key-Uncertainty of Quantum Ciphers. In [15], symmetric quantum ciphers encrypting classical messages with classical secret-keys are considered. It is shown that under known-plaintext attacks, the Shannon uncertainty of the secret-key can be much higher for some quantum ciphers than for any classical one. The Shannon secret-key uncertainty $\mathrm{H}(K|C, M)$ of classical ciphers C encrypting messages M of size m with keys K of size $k > m$ is always such that $\mathrm{H}(K|C,M) \leq k - m$. In the quantum case, the Shannon secret-key uncertainty is defined as the minimum residual uncertainty about key K given the best measurement (POVM) $P_M(C)$ applied to quantum cipher C given plaintext M. Examples of quantum ciphers are provided with $k = m + 1$ such that $\mathrm{H}(K|P_M(C)) = m/2+1$ and with $k = 2m$ such that $\mathrm{H}(K|P_M(C)) \geq 2m-1$. All ciphers in [15] have their keys consisting of two parts. The first part chooses one basis out a set $\mathcal{B}$ of bases while the other part is used as a classical one-time-pad. The message is first encrypted with the one-time-pad before being rotated in the basis indicated by the key. In this case, Theorem 4 in [15] states that the Shannon secret-key uncertainty adds up under repetitions with independent and random keys[1]: if $\mathrm{H}(K|P_M(C)) \geq h$ then n repetitions with independent keys satisfy $\mathrm{H}(K_1,\ldots,K_n|P_{M_1,\ldots,M_n}(C_1,\ldots,C_n)) \geq nh$. Our uncertainty relation allows to obtain a stronger result. The analysis in [15] shows that these quantum ciphers with Shannon secret-key uncertainty h satisfy the condition of our uncertainty relation. As result we obtain a lower bound on the min-entropy key uncertainty given the outcome of any quantum measurement applied to all ciphers and given all plaintexts. When $\mathrm{H}(K|P_M(C)) \geq h$ our uncertainty relation tells us that $\mathrm{H}_\infty(K_1,\ldots,K_n|P_{M_1,\ldots,M_n}(C_1,\ldots,C_n)) \gtrsim nh$. Notice that unlike the two previous applications, this time the result holds unconditionally. Details of this application will be provided in the full version.

History and Related Work. The history of uncertainty relations starts with Heisenberg who showed that the outcomes of two non-commuting observables A and B applied to any state ρ are not easy to predict simultaneously. However, Heisenberg only speaks about the variance of the measurement results. Because his result had several shortcomings (as pointed out in [19, 16]), more general forms of uncertainty relations were proposed by Bialynicki-Birula and Mycielski [7] and by Deutsch [16]. The new relations were called *entropic uncertainty relations*, because they are expressed using Shannon entropy instead of the statistical variance and, hence, are purely information theoretic statements. For instance, Deutsch's uncertainty relation [16] states that $\mathrm{H}(P)+\mathrm{H}(Q) \geq -2\log\frac{1+c}{2}$, where P, Q are random variables representing the measurement results and c is the maximum inner product norm between any eigenvectors of A and B. First conjectured by Kraus [21], Maassen and Uffink [24] improved Deutsch's relation to the optimal

$$\mathrm{H}(P) + \mathrm{H}(Q) \geq -2\log c \ . \tag{2}$$

Although a bound on Shannon entropy can be helpful in some cases, it is usually not good enough in cryptographic applications. The main tool to reduce the adversary's information—privacy amplification [5, 20, 4, 27, 25]—only works if a bound on the adversary's min-entropy (in fact collision entropy) is known. Unfortunately, knowing the Shannon entropy of a distribution does in general not allow to bound its higher order Rényi entropies.

[1] The proof of Theorem 4 in [15] is incorrect but can easily be fixed without changing the statement.

An entropic uncertainty relation involving Rényi entropy of order 2 (i.e. *collision entropy*) was introduced by Larsen [22, 30]. Larsen's relation quantifies precisely the collision entropy for the set $\{A_i\}_{i=1}^{d+1}$ of *all* maximally non-commuting observables, where d is the dimension of the Hilbert space. Its use is therefore restricted to quantum coding schemes that take advantage of *all* $d+1$ observables, i.e. to schemes that are difficult to implement in practice. Uncertainty relations in terms of Rényi entropy have also been studied in a different context by Bialynicki-Birula [6].

2 Preliminaries

2.1 Notation and Terminology

For any positive integer d, $\mathcal{H}_d$ stands for the complex Hilbert space of dimension d and $\mathcal{P}(\mathcal{H})$ for the set of density operators, i.e., positive semi-definite trace-1 matrices, acting on $\mathcal{H}$. The pair $\{|0\rangle, |1\rangle\}$ denotes the computational or rectilinear or "+" basis for the 2-dimensional Hilbert space $\mathcal{H}_2$. The diagonal or "×" basis is defined as $\{|0\rangle_\times, |1\rangle_\times\}$ where $|0\rangle_\times = (|0\rangle + |1\rangle)/\sqrt{2}$ and $|1\rangle_\times = (|0\rangle - |1\rangle)/\sqrt{2}$. The circular or "$\oslash$" basis consists of vectors $(|0\rangle + i|1\rangle)/\sqrt{2}$ and $(|0\rangle - i|1\rangle)/\sqrt{2}$. Measuring a qubit in the $+$-basis (resp. $\times$-basis) means applying the measurement described by projectors $|0\rangle\langle 0|$ and $|1\rangle\langle 1|$ (resp. projectors $|0\rangle_\times\langle 0|_\times$ and $|1\rangle_\times\langle 1|_\times$). When the context requires it, we write $|0\rangle_+$ and $|1\rangle_+$ instead of $|0\rangle$ and $|1\rangle$, respectively. If we want to choose the $+$ or $\times$-basis according to the bit $b \in \{0,1\}$, we write $[+,\times]_b$.

The behavior of a (mixed) quantum state in a register E is fully described by its density matrix ρ_E. We often consider cases where a quantum state may depend on some classical random variable X, in that the state is described by the density matrix ρ_E^x if and only if $X = x$. For an observer who has access to the state but not X, the behavior of the state is determined by the density matrix $\rho_E := \sum_x P_X(x)\rho_E^x$, whereas the joint state, consisting of the classical X and the quantum register E is described by the density matrix $\rho_{XE} := \sum_x P_X(x)|x\rangle\langle x| \otimes \rho_E^x$, where we understand $\{|x\rangle\}_{x\in\mathcal{X}}$ to be the standard (orthonormal) basis of $\mathcal{H}_{|\mathcal{X}|}$. Joint states with such *c*lassical and *q*uantum parts are called *cq-states*. We also write $\rho_X := \sum_x P_X(x)|x\rangle\langle x|$ for the quantum representation of the classical random variable X. This notation extends naturally to quantum states that depend on several classical random variables (i.e. to ccq-states, cccq-states etc.). Given a cq-state ρ_{XE} as above, by saying that there exists a random variable Y such that ρ_{XYE} satisfies some condition, we mean that ρ_{XE} can be understood as $\rho_{XE} = \mathrm{tr}_Y(\rho_{XYE})$ for some ccq-state ρ_{XYE} and that ρ_{XYE} satisfies the required condition.[2]

We would like to point out that $\rho_{XE} = \rho_X \otimes \rho_E$ holds if and only if the quantum part is independent of X (in that $\rho_E^x = \rho_E$ for any x), where the latter

[2] The quantum version is similar to the case of distributions of classical random variables where given X, the existence of a certain Y is understood that there exists a joint distribution P_{XY} with $\sum_y P_{XY}(\cdot, y) = P_X$.

in particular implies that no information on X can be learned by observing only ρ_E. Similarly, X is uniformly random and independent of the quantum state in register E if and only if $\rho_{XE} = \frac{1}{|\mathcal{X}|}\mathbb{1} \otimes \rho_E$, where $\frac{1}{|\mathcal{X}|}\mathbb{1}$ is the density matrix of the fully mixed state of suitable dimension. Finally, if two states like ρ_{XE} and $\rho_X \otimes \rho_E$ are ε-close in terms of their trace distance $\delta(\rho,\sigma) = \frac{1}{2}\operatorname{tr}(|\rho-\sigma|)$, which we write as $\rho_{XE} \approx_\varepsilon \rho_X \otimes \rho_E$, then the real system ρ_{XE} "behaves" as the ideal system $\rho_X \otimes \rho_E$ except with probability ε in that for any evolution of the system no observer can distinguish the real from the ideal one with advantage greater than ε [27].

2.2 Smooth Rényi Entropy

We briefly recall the notion of (conditional) *smooth* min-entropy [25, 28]. For more details, we refer to the aforementioned literature. Let X be a random variable over alphabet $\mathcal{X}$ with distribution P_X. The standard notion of min-entropy is given by $\mathrm{H}_\infty(X) = -\log\big(\max_x P_X(x)\big)$ and that of max-entropy by $\mathrm{H}_0(X) = \log\big|\{x \in \mathcal{X} : P_X(x) > 0\}\big|$. More general, for any event $\mathcal{E}$ (defined by $P_{\mathcal{E}|X}(x)$ for all $x \in \mathcal{X}$) $\mathrm{H}_\infty(X\mathcal{E})$ may be defined similarly simply by replacing P_X by $P_{X\mathcal{E}}$. Note that the "distribution" $P_{X\mathcal{E}}$ is not normalized; $\mathrm{H}_\infty(X\mathcal{E})$ is still well defined, though. For an arbitrary $\varepsilon \geq 0$, the smooth version $\mathrm{H}^\varepsilon_\infty(X)$ is defined as follows. $\mathrm{H}^\varepsilon_\infty(X)$ is the *maximum* of the standard min-entropy $\mathrm{H}_\infty(X\mathcal{E})$, where the maximum is taken over all events $\mathcal{E}$ with $\Pr(\mathcal{E}) \geq 1-\varepsilon$. Informally, this can be understood that if $\mathrm{H}^\varepsilon_\infty(X) = r$ then the standard min-entropy of X equals r as well, except with probability ε. As ε can be interpreted as an error probability, we typically require ε to be negligible in the security parameter n.

For random variables X and Y, the *conditional* smooth min-entropy $\mathrm{H}^\varepsilon_\infty(X \,|\, Y)$ is defined as $\mathrm{H}^\varepsilon_\infty(X \,|\, Y) = \max_{\mathcal{E}} \min_y \mathrm{H}_\infty(X\mathcal{E} \,|\, Y = y)$, where the quantification over $\mathcal{E}$ is over all events $\mathcal{E}$ (defined by $P_{\mathcal{E}|XY}$) with $\Pr(\mathcal{E}) \geq 1-\varepsilon$. In Sect. 6, we work with smooth min-entropy conditioned on a quantum state. We refer the reader to [25] for the definition of this quantum version. We will make use of the following chain rule for smooth min-entropy [28], which in spirit was already shown in [8].

Lemma 1. $\mathrm{H}^{\varepsilon+\varepsilon'}_\infty(X \,|\, Y) > \mathrm{H}^\varepsilon_\infty(XY) - \mathrm{H}_0(Y) - \log\big(\frac{1}{\varepsilon'}\big)$ *for all* $\varepsilon, \varepsilon' > 0$.

2.3 Azuma's Inequality

In the following and throughout the paper, the expected value of a real random variable R is denoted by $\mathbb{E}[R]$. Similarly, $\mathbb{E}[R|\mathcal{E}]$ and $\mathbb{E}[R|S]$ denote the conditional expectation of R conditioned on an event $\mathcal{E}$ respectively random variable S.

Definition 1. *A list of real random variables* $R_1, \ldots, R_n$ *is called a* martingale difference sequence *if* $\mathbb{E}[R_i \,|\, R_1, \ldots, R_{i-1}] = 0$ *with probability 1 for every* $1 \leq i \leq n$, *i.e., if* $\mathbb{E}[R_i \,|\, R_1 = r_1, \ldots, R_{i-1} = r_{i-1}] = 0$ *for every* $1 \leq i \leq n$ *and* $r_1, \ldots, r_{i-1} \in \mathbb{R}$.

The following lemma follows directly from Azuma's inequality [2, 1].

Lemma 2. *Let $R_1, \ldots, R_n$ be a martingale difference sequence such that $|R_i| \leq c$ for every $1 \leq i \leq n$. Then,* $\Pr\big[\sum_i R_i \geq \lambda n\big] \leq \exp\big(-\frac{\lambda^2 n}{2c^2}\big)$ *for any $\lambda > 0$.*

3 The Uncertainty Relation

We start with a classical tool which itself might be of independent interest.

Theorem 1. *Let $Z_1, \ldots, Z_n$ be n (not necessarily independent) random variables over alphabet $\mathcal{Z}$, and let $h \geq 0$ be such that*

$$\mathrm{H}(Z_i \,|\, Z_1 = z_1, \ldots, Z_{i-1} = z_{i-1}) \geq h \tag{3}$$

for all $1 \leq i \leq n$ and $z_1, \ldots, z_{i-1} \in \mathcal{Z}$. Then for any $0 < \lambda < \frac{1}{2}$

$$\mathrm{H}_\infty^\varepsilon(Z_1, \ldots, Z_n) \geq (h - 2\lambda)n \ ,$$

where $\varepsilon = \exp\big(-\frac{\lambda^2 n}{32 \log(|\mathcal{Z}|/\lambda)^2}\big)$.

If the Z_i's are *independent* and have Shannon-entropy at least h, it is known (see [28]) that the smooth min-entropy of $Z_1, \ldots, Z_n$ is at least nh for large enough n. Informally, Theorem 1 guarantees that when the independence-condition is relaxed to a lower bound on the Shannon entropy of Z_i *given any previous history*, then we still have min-entropy of (almost) nh except with negligible probability ε.

Proof (sketch). The idea is to use Lemma 2 for cleverly chosen R_i's. For any i we write $Z^i := (Z_1, \ldots, Z_i)$ (with Z^0 being the "empty symbol"), and similarly for other sequences. We want to show that $\Pr\big[P_{Z^n}(Z^n) \geq 2^{-(h-2\lambda)n}\big] \leq \varepsilon$. By the definition of smooth min-entropy, this then implies the claim. Note that $P_{Z^n}(Z^n) \geq 2^{-(h-2\lambda)n}$ is equivalent to

$$\sum_{i=1}^{n} \Big(\log\big(P_{Z_i \,|\, Z^{i-1}}(Z_i \,|\, Z^{i-1})\big) + h \Big) \geq 2\lambda n \ .$$

We set $S_i := \log P_{Z_i|Z^{i-1}}(Z_i \,|\, Z^{i-1})$. For such a sequence of real-valued random variables $S_1, \ldots, S_n$, it is easy to verify that $R_1, \ldots, R_n$ where $R_i := S_i - \mathbb{E}[S_i \,|\, S^{i-1}]$ forms a martingale difference sequence. If the $|R_i|$ were bounded by c, we could use Lemma 2 to conclude that

$$\Pr\left[\sum_{i=1}^{n} \Big(S_i - \mathbb{E}\big[S_i \,|\, S^{i-1}\big] \Big) \geq \lambda n \right] \leq \exp\left(-\frac{\lambda^2 n}{2c^2} \right) \ .$$

As by assumption $\mathbb{E}[S_i \,|\, S^{i-1}] \leq -h$, this would give us a bound similar to what we want to show. In order to enforce a bound on $|R_i|$, S_i needs to be truncated whenever $P_{Z_i \,|\, Z^{i-1}}(Z_i \,|\, Z^{i-1})$ is smaller than some $\delta > 0$. It is then a subtle and technically involved matter of choosing δ and ε appropriately in order to finish the proof, as shown in the full version of the paper [12]. $\square$

We now state and prove the new entropic uncertainty relation in its most general form. A special case will then be introduced (Corollary 1) and used in the security analysis of all protocols we consider in the following.

Definition 2. *Let $\mathcal{B}$ be a finite set of orthonormal bases in the d-dimensional Hilbert space $\mathcal{H}_d$. We call $h \geq 0$ an* average entropic uncertainty bound *for $\mathcal{B}$ if every state in $\mathcal{H}_d$ satisfies $\frac{1}{|\mathcal{B}|}\sum_{\vartheta\in\mathcal{B}} \mathrm{H}(P_\vartheta) \geq h$, where P_ϑ is the distribution obtained by measuring the state in basis ϑ.*

Note that by the convexity of the Shannon entropy H, a lower bound for all *pure* states in $\mathcal{H}_d$ suffices to imply the bound for all (possibly mixed) states.

Theorem 2. *Let $\mathcal{B}$ be a set of orthonormal bases in $\mathcal{H}_d$ with an average entropic uncertainty bound h, and let $\rho \in \mathcal{P}(\mathcal{H}_d^{\otimes n})$ be an arbitrary quantum state. Let $\Theta = (\Theta_1, \ldots, \Theta_n)$ be uniformly distributed over $\mathcal{B}^n$ and let $X = (X_1, \ldots, X_n)$ be the outcome when measuring ρ in basis Θ, distributed over $\{0, \ldots, d-1\}^n$. Then for any $0 < \lambda < \frac{1}{2}$ and $\lambda' > 0$,*

$$\mathrm{H}_\infty^{\varepsilon+\varepsilon'}(X \mid \Theta) \geq (h - 2\lambda - \lambda')\, n$$

with $\varepsilon = \exp\Big(-\frac{\lambda^2 n}{32(\log(|\mathcal{B}|\cdot d/\lambda))^2}\Big)$ *and* $\varepsilon' = 2^{-\lambda' n}$.

Proof. Define $Z_i := (X_i, \Theta_i)$ and $Z^i := (Z_1, \ldots, Z_i)$. Let z^{i-1} be arbitrary in $(\{0, \ldots, d-1\} \times \mathcal{B})^{i-1}$. Then

$$\mathrm{H}(Z_i \mid Z^{i-1} = z^{i-1}) = \mathrm{H}(X_i \mid \Theta_i, Z^{i-1} = z^{i-1}) + \mathrm{H}(\Theta_i \mid Z^{i-1} = z^{i-1}) \geq h + \log|\mathcal{B}|\,,$$

where the inequality follows from the fact that Θ_i is chosen uniformly at random and from the definition of h. Note that h lower bounds the average entropy for any system in $\mathcal{H}_d$, and thus in particular for the i-th subsystem of ρ, with all previous d-dimensional subsystems measured. We use the chain rule for smooth min-entropy (Lemma 1) and Theorem 1 to conclude that,

$$\mathrm{H}_\infty^{\varepsilon+\varepsilon'}(X \mid \Theta) > \mathrm{H}_\infty^{\varepsilon}(Z) - \mathrm{H}_0(\Theta) - \log\left(\tfrac{1}{\varepsilon'}\right) \geq (h - 2\lambda)n - \lambda' n \ ,$$

for ε and ε' as claimed. □

For the special case where $\mathcal{B} = \{+, \times\}$ is the set of BB84 bases, we can use the uncertainty relation of Maassen and Uffink [24] (see (2) with $c = 1/\sqrt{2}$), which, using our terminology, states that $\mathcal{B}$ has average entropic uncertainty bound $h = \frac{1}{2}$. Theorem 2 then immediately gives the following corollary.

Corollary 1. *Let $\rho \in \mathcal{P}(\mathcal{H}_2^{\otimes n})$ be an arbitrary n-qubit quantum state. Let Θ be uniformly distributed over $\{+, \times\}^n$, and let X be the outcome when measuring ρ in basis Θ. Then for any $0 < \lambda < \frac{1}{2}$ and $\lambda' > 0$,*

$$\mathrm{H}_\infty^{\varepsilon+\varepsilon'}(X \mid \Theta) \geq \left(\tfrac{1}{2} - 2\lambda - \lambda'\right) n$$

where $\varepsilon = \exp\big(-\frac{\lambda^2 n}{32(2-\log(\lambda))^2}\big)$ *and* $\varepsilon' = 2^{-\lambda' n}$.

Maassen and Uffink's relation being optimal means there exists a quantum state ρ—namely the product state of eigenstates of the subsystems, e.g. $\rho = |0\rangle\langle 0|^{\otimes n}$—for which $\mathrm{H}(X \mid \Theta) = \frac{n}{2}$. On the other hand, we have shown that $(\frac{1}{2} - \lambda)n \leq \mathrm{H}_\infty^\varepsilon(X \mid \Theta)$ for $\lambda > 0$ arbitrarily close to 0. For the product state ρ, the X_i's are independent and we know from [28] that in this case $\mathrm{H}_\infty^\varepsilon(X \mid \Theta)$ approaches $\mathrm{H}(X \mid \Theta) = \frac{n}{2}$. It follows that the relation cannot be significantly improved even when considering Rényi entropy of lower order than min-entropy (but higher than Shannon entropy).

Another tight corollary is obtained if we consider the set of measurements $\mathcal{B} = \{+, \times, \oslash\}$. In [29], Sánchez-Ruiz has shown that for this $\mathcal{B}$ the average entropic uncertainty bound $h = \frac{2}{3}$ is optimal. It implies that $\mathrm{H}_\infty^\varepsilon(X \mid \Theta) \approx \mathrm{H}(X \mid \Theta) = \frac{2n}{3}$ for negligible ε. In the full version [12], we compute the average uncertainty bound for the set of *all bases* of a d-dimensional Hilbert space.

4 Application: Oblivious Transfer

4.1 Privacy Amplification and a Min-Entropy-Splitting Lemma

Recall, a class $\mathcal{F}$ of hash functions from, say, $\{0,1\}^n$ to $\{0,1\}^\ell$ is called *two-universal* [9, 31] if $\Pr[F(x) = F(x')] \leq 1/2^\ell$ for any distinct $x, x' \in \{0,1\}^n$ and for F uniformly distributed over $\mathcal{F}$.

Theorem 3 (Privacy Amplification [27, 25]). *Let $\varepsilon \geq 0$. Let $\rho_{XU\mathsf{E}}$ be a ccq-state, where X takes values in $\{0,1\}^n$, U in the finite domain $\mathcal{U}$ and register E contains q qubits. Let F be the random and independent choice of a member of a two-universal class of hash functions $\mathcal{F}$ from $\{0,1\}^n$ into $\{0,1\}^\ell$. Then,*

$$\delta\big(\rho_{F(X)FU\mathsf{E}}, \tfrac{1}{2^\ell}\mathbb{1} \otimes \rho_{FU\mathsf{E}}\big) \leq \frac{1}{2}\, 2^{-\frac{1}{2}\big(H_\infty^\varepsilon(X|U) - q - \ell\big)} + 2\varepsilon \ . \tag{4}$$

The theorem stated here is slightly different from the version given in [27, 25] in that the classical and the quantum parts of the adversary's knowledge are treated differently. A derivation of the above theorem starting from the result in [25] can be found in the full version [12].

A second tool we need is the following Min-Entropy-Splitting Lemma. Note that if the joint entropy of two random variables X_0 and X_1 is large, then one is tempted to conclude that at least one of X_0 and X_1 must still have large entropy, e.g. half of the original entropy. Whereas this is indeed true for Shannon entropy, it is in general not true for min-entropy. The following lemma, though, which appeared in a preliminary version of [33], shows that it *is* true in a randomized sense. For completeness, the proof can be found in the full version [12].

Lemma 3 (Min-Entropy-Splitting Lemma). *Let $\varepsilon \geq 0$, and let X_0, X_1 be random variables (over possibly different alphabets) with $\mathrm{H}_\infty^\varepsilon(X_0 X_1) \geq \alpha$. Then, there exists a binary random variable C over $\{0,1\}$ such that $\mathrm{H}_\infty^\varepsilon(X_{1-C} C) \geq \alpha/2$.*

The corollary below follows rather straightforwardly by noting that (for normalized as well as non-normalized distributions) $\mathrm{H}_\infty(X_0X_1 \,|\, Z) \geq \alpha$ holds exactly if $\mathrm{H}_\infty(X_0X_1 \,|\, Z=z) \geq \alpha$ for all z, applying the Min-Entropy-Splitting Lemma, and then using the Chain Rule, Lemma 1.

Corollary 2. *Let $\varepsilon \geq 0$, and let X_0, X_1 and Z be random variables such that $\mathrm{H}_\infty^\varepsilon(X_0X_1 \,|\, Z) \geq \alpha$. Then, there exists a binary random variable C over $\{0,1\}$ such that $\mathrm{H}_\infty^{\varepsilon+\varepsilon'}(X_{1-C} \,|\, ZC) \geq \alpha/2 - 1 - \log(1/\varepsilon')$ for any $\varepsilon' > 0$.*

4.2 The Definition

In *1-2 OT^ℓ*, the sender Alice sends two ℓ-bit strings S_0, S_1 to the receiver Bob in such a way that Bob can choose which string to receive, but does not learn anything about the other. On the other hand, Alice does not get to know which string Bob has chosen. The common way to build *1-2 OT^ℓ* is by constructing a protocol for (Sender-)Randomized *1-2 OT^ℓ*, which then can easily be converted into an ordinary *1-2 OT^ℓ* (see, e.g., [14]). *Rand 1-2 OT^ℓ* essentially coincides with ordinary *1-2 OT^ℓ*, except that the two strings S_0 and S_1 are not *input* by the sender but generated uniformly at random during the protocol and *output* to the sender.

For the formal definition of the security requirements of a quantum protocol for *Rand 1-2 OT^ℓ*, let us fix the following notation: Let C denote the binary random variable describing receiver R's choice bit, let S_0, S_1 denote the ℓ-bit long random variables describing sender S's output strings, and let Y denote the ℓ-bit long random variable describing R's output string (supposed to be S_C). Furthermore, for a fixed candidate protocol for *Rand 1-2 OT^ℓ*, and for a fixed input distribution for C, the overall quantum state in case of a dishonest sender $\tilde{\mathsf{S}}$ is given by the ccq-state $\rho_{CY\tilde{\mathsf{S}}}$. Analogously, in the case of a dishonest receiver $\tilde{\mathsf{R}}$, we have the ccq-state $\rho_{S_0S_1\tilde{\mathsf{R}}}$.

Definition 3 (*Rand 1-2 OT^ℓ*). *An ε-secure* Rand 1-2 OT^ℓ *is a quantum protocol between S and R, with R having input $C \in \{0,1\}$ while S has no input, such that for any distribution of C, if S and R follow the protocol, then S gets output $S_0, S_1 \in \{0,1\}^\ell$ and R gets $Y = S_C$ except with probability ε, and the following two properties hold:*

ε-Receiver-security: *If R is honest, then for any $\tilde{\mathsf{S}}$, there exist random variables S_0', S_1' such that $\Pr\big[Y = S_C'\big] \geq 1 - \varepsilon$ and $\delta\big(\rho_{CS_0'S_1'\tilde{\mathsf{S}}}, \rho_C \otimes \rho_{S_0'S_1'\tilde{\mathsf{S}}}\big) \leq \varepsilon$.*

ε-Sender-security: *If S is honest, then for any $\tilde{\mathsf{R}}$, there exists a binary random variable D such that $\delta\big(\rho_{S_{1-D}S_DD\tilde{\mathsf{R}}}, \frac{1}{|2^\ell|}\mathbb{1} \otimes \rho_{S_DD\tilde{\mathsf{R}}}\big) \leq \varepsilon$.*

If any of the above holds for $\varepsilon = 0$, then the corresponding property is said to hold perfectly. *If one of the properties only holds with respect to a restricted class $\mathfrak{S}$ of $\tilde{\mathsf{S}}$'s respectively $\mathfrak{R}$ of $\tilde{\mathsf{R}}$'s, then this property is said to hold and the protocol is said to be secure* against *$\mathfrak{S}$ respectively $\mathfrak{R}$.*

Receiver-security, as defined here, implies that whatever a dishonest sender does is as good as the following: generate the ccq-state $\rho_{S'_0 S'_1 \tilde{\mathsf{S}}}$ independently of C, let R know S'_C, and output $\rho_{\tilde{\mathsf{S}}}$. On the other hand, sender-security implies that whatever a dishonest receiver does is as good as the following: generate the ccq-state $\rho_{S_D D \tilde{\mathsf{R}}}$, let S know S_D and an independent uniformly distributed S_{1-D}, and output $\rho_{\tilde{\mathsf{R}}}$. In other words, a protocol satisfying Definition 3 is a secure implementation of the natural *Rand 1-2 OT^ℓ* ideal functionality, except that it allows a dishonest sender to influence the distribution of S_0 and S_1, and the dishonest receiver to influence the distribution of the string of his choice. This is in particular good enough for constructing a standard *1-2 OT^ℓ* in the straightforward way.

We would like to point out the importance of requiring the existence of S'_0 and S'_1 in the formulation of receiver-security in a quantum setting: requiring only that the sender learns no information on C, as is sufficient in the classical setting (see e.g. [10]), does not prevent a dishonest sender from obtaining S_0, S_1 by a suitable measurement *after* the execution of the protocol in such a way that he can choose $S_0 \oplus S_1$ at will, and S_C is the string the receiver has obtained in the protocol. This would for instance make the straightforward construction of a bit commitment[3] based on *1-2 OT* insecure.

4.3 The Protocol

We introduce a quantum protocol for *Rand 1-2 OT^ℓ* that will be shown perfectly receiver-secure against any sender and ε-sender-secure against any quantum-memory-bounded receiver for a negligible ε. The first two steps of the protocol are identical to Wiesner's "conjugate coding" protocol [32] from circa 1970 for *"transmitting two messages either but not both of which may be received"*.

The simple protocol is described in Fig. 1, where for $x \in \{0,1\}^n$ and $I \subseteq \{1,\ldots,n\}$ we define $x|_I$ to be the restriction of x to the bits x_i with $i \in I$. The sender S sends random BB84 states to the receiver R, who measures all received qubits according to his choice bit C. S then picks randomly two functions from a fixed two-universal class of hash functions $\mathcal{F}$ from $\{0,1\}^n$ to $\{0,1\}^\ell$, where ℓ is to be determined later, and applies them to the bits encoded in the $+$ respectively the bits encoded in $\times$-basis to obtain the output strings S_0 and S_1. Note that we may apply a function $f \in \mathcal{F}$ to a n'-bit string with $n' < n$ by padding it with zeros (which does not decrease its entropy). S announces the encoding bases and the hash functions to the receiver who then can compute S_C. Intuitively, a dishonest receiver who cannot store all the qubits until the right bases are announced, will measure some qubits in the wrong basis and thus cannot learn both strings simultaneously.

We would like to stress that although protocol description and analysis are designed for an ideal setting with perfect noiseless quantum communication and

[3] The committer sends two random bits of parity equal to the bit he wants to commit to, the verifier chooses to receive at random one of those bits.

RAND 1-2 QOT $^{\ell}$: Let c be R's choice bit.
1. S picks $x \in_R \{0,1\}^n$ and $\theta \in_R \{+,\times\}^n$, and sends $|x_1\rangle_{\theta_1},\ldots,|x_n\rangle_{\theta_n}$ to R.
2. R measures all qubits in basis $[+,\times]_c$. Let $x' \in \{0,1\}^n$ be the result.
3. S picks two hash functions $f_0, f_1 \in_R \mathcal{F}$, announces θ and f_0, f_1 to R, and outputs $s_0 := f_0(x|_{I_0})$ and $s_1 := f_1(x|_{I_1})$ where $I_b := \{i : \theta_i = [+,\times]_b\}$.
4. R outputs $s_c = f_c(x'|_{I_c})$.

Fig. 1. Quantum Protocol for *Rand 1-2 OT* $^{\ell}$.

with perfect sources and detectors, all our results can easily be extended to a more realistic noisy setting along the same lines as in [13].

It is clear by the non-interactivity of RAND 1-2 QOT $^{\ell}$ that a dishonest sender cannot learn anything about the receiver's choice bit. The proof of receiver-security according to Definition 3 can be found in the full version [12]; the idea, though, simply is to have a dishonest $\tilde{\mathsf{S}}$ execute the protocol with a receiver that has *unbounded quantum memory* and that way can compute S_0' and S_1'.

Proposition 1. RAND 1-2 QOT $^{\ell}$ *is perfectly receiver-secure.*

4.4 Security Against Memory-Bounded Dishonest Receivers

We model dishonest receivers in RAND 1-2 QOT $^{\ell}$ under the assumption that the maximum size of their quantum storage is bounded. Such adversaries are only required to have bounded quantum storage when Step 3 in RAND 1-2 QOT $^{\ell}$ is reached; before and after that, the adversary can store and carry out arbitrary quantum computations involving any number of qubits. Let $\mathfrak{R}_q$ denote the set of all possible quantum dishonest receivers $\tilde{\mathsf{R}}$ in RAND 1-2 QOT $^{\ell}$ which have quantum memory of size at most q when Step 3 is reached. We stress once more that apart from the restriction on the size of the quantum memory available to the adversary, no other assumption is made. In particular, the adversary is not assumed to be computationally bounded and the size of his classical memory is not restricted.

Theorem 4. RAND 1-2 QOT $^{\ell}$ *is ε-secure against $\mathfrak{R}_q$ for a negligible (in n) ε if $n/4 - 2\ell - q \in \Omega(n)$.*

For improved readability, we merely give a sketch of the proof; the formal proof that takes care of all the ε's is given in the full version [12].

Proof (sketch). It remains to show sender-security. Let X be the random variable that describes the sender's choice of x, where we understand the distribution of X to be conditioned on the classical information that $\tilde{\mathsf{R}}$ obtained by measuring all but γn qubits. A standard purification argument, that was also used in [13], shows that the same X can be obtained by measuring a quantum state in basis $\theta \in_R \{+,\times\}^n$, described by the random variable Θ: for each qubit $|x_i\rangle_{\theta_i}$ the sender S is instructed to send to R, S instead prepares an EPR pair

$|\Phi\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$ and sends one part to R while keeping the other, and when Step 3 is reached, S measures her qubits.

The uncertainty relation, Theorem 1, implies that the smooth min-entropy of X given Θ is approximately $n/2$. Let now X_0 and X_1 be the two substrings of X consisting of the bits encoded in the basis $+$ or $\times$, respectively. Then the Min-Entropy-Splitting Lemma, respectively Corollary 2, implies the existence of a binary D such that X_{1-D} has approximately $n/4$ bits of smooth min-entropy given Θ and D. From the random and independent choice of the hash functions F_0, F_1 and from the Chain Rule, Lemma 1, it follows that X_{1-D} has still about $n/4-\ell$ bits of smooth min-entropy when conditioning on Θ, D, F_D and $F_D(X_D)$. The Privacy Amplification Theorem 3, then guarantees that $S_{1-D} = F_{1-D}(X_{1-D})$ is close to random, given $\Theta, D, F_D, S_D, F_{1-D}$ and $\tilde{\mathsf{R}}$'s quantum state of size q, if $n/4 - 2\ell - q$ is positive and linear in n. □

We note that by adapting recent and more advanced techniques [33] to the quantum case, the security of Rand 1-2 QOT$^{\ell}$ can be proven against $\mathfrak{R}_q$ if $n/4 - \ell - q \in \Omega(n)$.

5 Application: Quantum Bit Commitment

The binding criterion for classical commitments usually requires that after the committing phase and for any dishonest committer, there exists a bit $d \in \{0,1\}$ that can only be opened with negligible probability. In the quantum world, the binding property cannot be defined the same way. If the commitment is unconditionally concealing, the committer can place himself in superposition of committing to 0 and 1 and only later make a measurement that fixes the choice. For this reason, the previous standard approach (see e.g. [17]) was to use a weaker binding condition only requiring that the probabilities p_0 and p_1 (to successfully open $b = 0$ and $b = 1$ respectively), satisfy $p_0 + p_1 \lesssim 1$. The bit commitment scheme proposed in [13] was shown to be binding in this weak sense.

We first argue that this weak notion is not really satisfactory. For instance, it does not capture the expected behavior of a commitment scheme by allowing a dishonest committer who can open the commitment with probability 1/2 to any value, and with probability 1/2 is unable to open it at all (depending on some event occurring during the opening). Another shortcoming of this notion is that committing bit by bit does not yield a secure string commitment—the argument that one is tempted to use requires independence of the p_b's between the different executions, which in general does not hold. We now argue that this notion is *unnecessarily* weak, even when taking into account a committer committing in superposition. We propose the following definition.

Definition 4. *An unconditionally secure commitment scheme is called* binding*, if for every (dishonest) committer there exists a classical binary random variable D whose distribution* cannot *be influenced by the (dishonest) committer after the commit phase and with the property that the committer's probability to successfully open the commitment to $1 - D$ is negligible.*

Note that this definition still allows a committer to commit to a superposition and otherwise honestly follow the protocol. D is then simply defined to be the outcome when the register that carries the superposition is measured. On the other hand, the definition captures exactly what one expects from a commitment scheme, except that the bit, to which the committer can open the commitment, is not fixed right after the commit phase. However, once committed, the dishonest committer *cannot influence* its distribution anymore, and thus this is not of any help to him, because he can always pretend not to know that bit.

It is also clear that with this stronger notion of the binding property, the obvious reduction from a string to a bit commitment scheme by committing bit-wise can be proven secure: the i-th execution of the bit commitment scheme guarantees a random variable D_i such that the committer cannot open the i-th bit commitment to $1 - D_i$, and thus there exists a random variable S, namely $S = (D_1, D_2, \ldots)$, such that the committer cannot open the list of commitments to any other string than S.

We show in the following that the quantum bit-commitment scheme COMM from [13] fulfills the stronger notion of binding from Definition 4 above. Let $\mathfrak{C}_q$ denote the set of all possible quantum dishonest committers $\tilde{\mathsf{C}}$ in COMM which have quantum memory of size at most q at the start of the opening phase. Then the following holds.

Theorem 5. *The quantum bit-commitment scheme* COMM *from [13] is binding according to Definition 4 against* $\mathfrak{C}_q$ *if* $n/4 - q \in \Omega(n)$.

Proof (Sketch). By considering a purified version of the scheme and using the uncertainty relation, one can argue that X has (smooth) min-entropy about $n/2$ given Θ. The Min-Entropy-Splitting Lemma implies that there exists D such that X_{1-D} has smooth min-entropy about $n/4$ given Θ and D. Privacy amplification implies that $F(X_{1-D})$ is close to random given Θ, D, F and $\tilde{\mathsf{C}}$'s quantum register of size q, where F is a two-universal one-bit-output hash function, which in particular implies that $\tilde{\mathsf{C}}$ cannot guess X_{1-D}. □

6 Application: Quantum Key Distribution

Let $\mathcal{B}$ be a set of orthonormal bases on a Hilbert space $\mathcal{H}$, and assume that the basis vectors of each basis $\vartheta \in \mathcal{B}$ are parametrized by the elements of some fixed set $\mathcal{X}$. We then consider QKD protocols consisting of the steps described in Fig. 2. Note that the quantum channel is only used in the preparation step. Afterwards, the communication between Alice and Bob is only classical (over an authentic channel).

As shown in [25] (Lemma 6.4.1), the length ℓ of the secret key that can be generated in the privacy amplification step of the protocol described above is given by[4]

$$\ell \approx \mathrm{H}_\infty^\varepsilon(X \mid E) - \mathrm{H}_0(C) \ ,$$

[4] The approximation in this and the following equations holds up to some small additive value which depends logarithmically on the desired security ε of the final key.

One-Way QKD: let $N \in \mathbb{N}$ be arbitrary

1. *Preparation:* For $i = 1 \dots N$, Alice chooses at random a basis $\vartheta_i \in \mathcal{B}$ and a random element $X_i \in \mathcal{X}$. She encodes X_i into the state of a quantum system (e.g., a photon) according to the basis ϑ_i and sends this system to Bob. Bob measures each of the states he receives according to a randomly chosen basis ϑ'_i and stores the outcome Y_i of this measurement.
2. *Sifting:* Alice and Bob publicly announce their choices of bases and keep their data at position i only if $\vartheta_i = \vartheta'_i$. In the following, we denote by X and Y the concatenation of the remaining data X_i and Y_i, respectively. X and Y are sometimes called the *sifted raw key.*
3. *Error correction:* Alice computes some error correction information C depending on X and sends C to Bob. Bob computes a guess $\hat{X}$ for Alice's string X, using C and Y.
4. *Privacy amplification:* Alice chooses at random a function f from a two-universal family of hash functions and announces f to Bob. Alice and Bob then compute the final key by applying f to their strings X and $\hat{X}$, respectively.

Fig. 2. General form for *one-way* QKD protocols.

where E denotes the (quantum) system containing all the information Eve might have gained during the preparation step of the protocol and where $\mathrm{H}_0(C)$ is the number of error correction bits sent from Alice to Bob. Note that this formula can be seen as a generalization of the well known expression by Csiszár and Körner for classical key agreement [11].

Let us now assume that Eve's system E can be decomposed into a classical part Z and a purely quantum part E'. Then, using the chain rule (Lemma 3.2.9 in [25]), we find

$$\ell \approx \mathrm{H}^{\varepsilon}_{\infty}(X \mid ZE') - \mathrm{H}_0(C) \gtrsim \mathrm{H}^{\varepsilon}_{\infty}(X \mid Z) - \mathrm{H}_0(E') - \mathrm{H}_0(C) \ .$$

Because, during the preparation step, Eve does not know the encoding bases which are chosen at random from the set $\mathcal{B}$, we can apply our uncertainty relation (Theorem 2) to get a lower bound for the min-entropy of X conditioned on Eve's classical information Z, i.e., $\mathrm{H}^{\varepsilon}_{\infty}(X \mid Z) \geq Mh$, where M denotes the length of the sifted raw key X and h is the average entropic uncertainty bound for $\mathcal{B}$. Let q be the bound on the size of Eve's quantum memory E'. Moreover, let e be the average amount of error correction information that Alice has to send to Bob per symbol of the sifted raw key X. Then

$$\ell \gtrsim M(h - e) - q \ .$$

Hence, if the memory bound only grows sublinearly in the length M of the sifted raw key, then the *key rate*, i.e., the number of key bits generated per bit of the sifted raw key, is lower bounded by

$$\mathrm{rate} \geq h - e \ .$$

The Binary-Channel Setting. For a binary channel (where $\mathcal{H}$ has dimension two), the average amount of error correction information e is given by the binary Shannon entropy[5] $\mathrm{H}_{\mathrm{bin}}(p) = -\big(p\log(p) + (1-p)\log(1-p)\big)$, where p is the bit-flip probability of the quantum channel (for classical bits encoded according to some orthonormal basis as described above). The achievable key rate of a QKD protocol using a binary quantum channel is thus given by $\mathrm{rate}_{\mathrm{binary}} \geq h - \mathrm{H}_{\mathrm{bin}}(p)$. Summing up, we have derived the following theorem.

Theorem 6. *Let $\mathcal{B}$ be a set of orthonormal bases of $\mathcal{H}_2$ with average entropic uncertainty bound h. Then, a* one-way *QKD-protocol as in Fig. 2 produces a secure key against eavesdroppers whose quantum-memory size is sublinear in the length of the raw key (i.e., sublinear in the number of qubits sent from Alice to Bob) at a positive rate as long as the bit-flip probability p fulfills* $\mathrm{H}_{\mathrm{bin}}(p) < h$.

For the BB84 protocol, we have $h = \frac{1}{2}$ and $\mathrm{H}_{\mathrm{bin}}(p) < \frac{1}{2}$ is satisfied as long as $p \leq 11\%$. This bound coincides with the known bound for security against an unbounded adversary. So, the memory-bound does not give an advantage here.[6]

The situation is different for the six-state protocol where $h = \frac{2}{3}$. In this case, security against memory-bounded adversaries is guaranteed (i.e. $\mathrm{H}_{\mathrm{bin}}(p) < \frac{2}{3}$) as long as $p \leq 17\%$. If one requires security against an unbounded adversary, the threshold for the same protocol lies below 13%, and even the best known QKD protocol on binary channels with one-way classical post-processing can only tolerate noise up to roughly 14.1% [26]. It has also been shown that, in the unbounded model, no such protocol can tolerate an error rate of more than 16.3%.

The performance of QKD protocols against quantum-memory bounded eavesdroppers can be improved further by making the choice of the encoding bases more random. For example, they might be chosen from the set of all possible orthonormal bases on a two-dimensional Hilbert space. As shown in the full version [12], the average entropic uncertainty bound is then given by $h \approx 0.72$ and $\mathrm{H}_{\mathrm{bin}}(p) < 0.72$ is satisfied if $p \lesssim 20\%$. For an unbounded adversary, the thresholds are the same as for the six-state protocol.

7 Open Problems

It is interesting to investigate whether the uncertainty relation (Theorem 2) still holds if the measurement bases $(\Theta_1, \ldots, \Theta_n)$ are randomly chosen from a relatively small subset of $\mathcal{B}^n$ (rather than from the entire set $\mathcal{B}^n$). Such an extension would reduce the amount of randomness that is needed in applications. In particular, in the context of QKD with quantum-memory-bounded eavesdroppers, it would allow for more efficient protocols that use a relatively short initial secret key in order to select the bases for the preparation and measurement of the states and, hence, avoid the sifting step.

[5] This value of e is only achieved if an optimal error-correction scheme is used. In practical implementations, the value of e might be slightly larger.

[6] Note, however, that the analysis given here might not be optimal.

Another open problem is to consider protocols using higher-dimensional quantum systems. The results mentioned in the previous paragraph show that for high-dimensional systems, the average entropic uncertainty bound converges to its theoretical maximum. The maximal tolerated channel noise might thus be higher for such protocols (depending on the noise model for higher-dimensional quantum channels).

References

1. N. Alon and J. Spencer. *The Probabilistic Method*. Series in Discrete Mathematics and Optimization. Wiley-Interscience, 2nd edition, 2000.
2. K. Azuma. Weighted sums of certain dependent random variables. *Tôhoku Mathematical Journal*, 19:357–367, 1967.
3. C. H. Bennett and G. Brassard. Quantum cryptography: Public key distribution and coin tossing. In *Proceedings of IEEE International Conference on Computers, Systems, and Signal Processing*, pages 175–179, 1984.
4. C. H. Bennett, G. Brassard, C. Crépeau, and U. M. Maurer. Generalized privacy amplification. *IEEE Transactions on Information Theory*, 41:1915–1923, Nov. 1995.
5. C. H. Bennett, G. Brassard, and J.-M. Robert. Privacy amplification by public discussion. *SIAM J. Comput.*, 17(2):210–229, 1988.
6. I. Bialynicki-Birula. Formulation of the uncertainty relations in terms of the Rényi entropies. *Physical Review A*, 74:052101, 2006.
7. I. Bialynicki-Birula and J. Mycielski. Uncertainty relations for information entropy. *Communications in Mathematical Physics*, 129(44), 1975.
8. C. Cachin. Smooth entropy and Rényi entropy. In *Advances in Cryptology—EUROCRYPT '97*, volume 1233 of *Lecture Notes in Computer Science*, pages 193–208. Springer, 1997.
9. J. L. Carter and M. N. Wegman. Universal classes of hash functions. In *9th Annual ACM Symposium on Theory of Computing (STOC)*, pages 106–112, 1977.
10. C. Crépeau, G. Savvides, C. Schaffner, and J. Wullschleger. Information-theoretic conditions for two-party secure function evaluation. In *Advances in Cryptology—EUROCRYPT '06*, volume 4004 of *Lecture Notes in Computer Science*, pages 538–554. Springer, 2006.
11. I. Csiszár and J. Körner. Broadcast channels with confidential messages. *IEEE Transactions on Information Theory*, 24(3):339–348, May 1978.
12. I. B. Damgård, S. Fehr, R. Renner, L. Salvail, and C. Schaffner. A tight high-order entropic quantum uncertainty relation with applications. 2007. Available at `http://arxiv.org/abs/quant-ph/0612014`.
13. I. B. Damgård, S. Fehr, L. Salvail, and C. Schaffner. Cryptography in the bounded quantum-storage model. In *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 449–458, 2005.
14. I. B. Damgård, S. Fehr, L. Salvail, and C. Schaffner. Oblivious transfer and linear functions. In *Advances in Cryptology—CRYPTO '06*, volume 4117 of *Lecture Notes in Computer Science*, pages 427–444. Springer, 2006.
15. I. B. Damgård, T. B. Pedersen, and L. Salvail. On the key-uncertainty of quantum ciphers and the computational security of one-way quantum transmission. In *Advances in Cryptology—EUROCRYPT '04*, volume 3027 of *Lecture Notes in Computer Science*, pages 91–108. Springer, 2004.

16. D. Deutsch. Uncertainty in quantum measurements. *Physical Review Letters*, 50(9):631–633, February 1983.
17. P. Dumais, D. Mayers, and L. Salvail. Perfectly concealing quantum bit commitment from any quantum one-way permutation. In *Advances in Cryptology—EUROCRYPT '00*, volume 1807 of *Lecture Notes in Computer Science*, pages 300–315. Springer, 2000.
18. C. A. Fuchs, N. Gisin, R. B. Griffiths, C.-S. Niu, and A. Peres. Optimal eavesdropping in quantum cryptography. I. Information bound and optimal strategy. *Physical Review A*, 56:1163 – 1172, 1997.
19. J. Hilgevood and J. Uffink. The mathematical expression of the uncertainty principle. In *Microphysical Reality and Quantum Description*. Kluwer Academic, 1988.
20. R. Impagliazzo, L. A. Levin, and M. Luby. Pseudo-random generation from one-way functions. In *21st Annual ACM Symposium on Theory of Computing (STOC)*, pages 12–24, 1989.
21. K. Kraus. Complementary observables and uncertainty relations. *Physical Review D*, 35(10):3070–3075, May 1987.
22. U. Larsen. Superspace geometry: the exact uncertainty relationship between complementary aspects. *Journal of Physics A: Mathematical and General*, 23(7):1041–1061, April 1990.
23. N. Lütkenhaus. Security against individual attacks for realistic quantum key distribution. *Physical Review A*, 61:052304, 2000.
24. H. Maassen and J. B. M. Uffink. Generalized entropic uncertainty relations. *Physical Review Letters*, 60(12):1103–1106, March 1988.
25. R. Renner. *Security of Quantum Key Distribution*. PhD thesis, ETH Zürich, 2005. `http://arxiv.org/abs/quant-ph/0512258`.
26. R. Renner, N. Gisin, and B. Kraus. An information-theoretic security proof for QKD protocols. *Phys. Rev. A*, 72(012332), July 2005.
27. R. Renner and R. König. Universally composable privacy amplification against quantum adversaries. In *Theory of Cryptography Conference (TCC)*, volume 3378 of *Lecture Notes in Computer Science*, pages 407–425. Springer, 2005.
28. R. Renner and S. Wolf. Simple and tight bounds for information reconciliation and privacy amplification. In *Advances in Cryptology—ASIACRYPT 2005*, Lecture Notes in Computer Science, pages 199–216. Springer, 2005.
29. J. Sánchez-Ruiz. Entropic uncertainty and certainty relations for complementary observables. *Physics Letters A*, 173(3):233–239, February 1993.
30. J. Sánchez-Ruiz. Improved bounds in the entropic uncertainty and certainty relations for complementary observables. *Physics Letters A*, 201(2–3):125–131, May 1995.
31. M. N. Wegman and J. L. Carter. New classes and applications of hash functions. In *20th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 175–182, 1979.
32. S. Wiesner. Conjugate coding. *SIGACT News*, 15(1):78–88, 1983. Original manuscript written circa 1970.
33. J. Wullschleger. Oblivious-Transfer amplification. In *Advances in Cryptology—EUROCRYPT '07*, Lecture Notes in Computer Science. Springer, 2007.