

Sampling in a Quantum Population, and Applications

Niek Bouman and Serge Fehr

Centrum Wiskunde & Informatica (CWI), Amsterdam, The Netherlands

{n.j.bouman,s.fehr}@cwi.nl

Abstract

We propose a framework for analyzing classical sampling strategies for estimating the Hamming weight of a large string from a few sample positions, when applied to a multi-qubit quantum system instead. The framework shows how to interpret the result of such a strategy and how to define its accuracy when applied to a quantum system. Furthermore, we show how the accuracy of any strategy relates to its accuracy in its classical usage, which is well understood for the important examples.

We show the usefulness of our framework by using it to obtain new and simple security proofs for the following quantum-cryptographic schemes: quantum oblivious-transfer from bit-commitment, and BB84 quantum-key-distribution.

Keywords: Random sampling, quantum key distribution, quantum oblivious transfer.

1 Introduction

Sampling allows to learn some information on a large population by merely looking at a comparably small number of individuals. For instance it is possible to predict the outcome of an election with very good accuracy by analyzing a relatively small subset of all the votes. In this work, we initiate the study of sampling in a *quantum* population, where we want to be able to learn information on a large quantum state by measuring only a small part. Specifically, we investigate the quantum-version of the following classical sampling problem (and of variants thereof). Given a bit-string $\mathbf{q} = (q_1, \dots, q_n) \in \{0, 1\}^n$ of length n , the task is to estimate the Hamming weight of \mathbf{q} by sampling and looking at only a few positions within \mathbf{q} . This classical sampling problem is well understood. For instance the following particular *sampling strategy* works well: sample (with or without replacement) a linear number of positions uniformly at random, and compute an estimate for the Hamming weight of \mathbf{q} by scaling the Hamming weight of the sample accordingly; Hoeffding's bounds guarantee that the estimate is close to the real Hamming weight except with small probability. Such a sampling strategy in particular allows to *test* whether \mathbf{q} is close to the all-zero string $(0, \dots, 0)$ by looking only at a relatively small number of positions, where the test is accepted if and only if all the sample positions are zero, i.e., the estimated Hamming weight vanishes.

In the quantum version of the above sampling problem, the string \mathbf{q} is replaced by a n -qubit quantum system A . It is obvious that a sampling strategy from the classical can be *applied* to the quantum setting as well: pick a sample of qubit positions within A , measure (in the computational basis) these sample positions, and compute the estimate as dictated by the sampling strategy from the observed values (i.e., typically, scale the Hamming weight of the measured sample appropriately). However, what is a-priori

not clear, is how to formally *interpret* the computed estimate. In the special case of testing closeness to the all-zero string, one expects that if the measurement of a random sample only produces zeros then the initial state of A must have been close to the all-zero state $|0\rangle \cdots |0\rangle$. But what is the right way to measure closeness here? For instance it must allow for states of the form $|q\rangle$ where $q \in \{0, 1\}^n$ has small Hamming weight, but it must also allow for superpositions with arbitrary states that come with a very small amplitude. In the general case of a sampling strategy that, in its classical usage, aims at estimating the Hamming weight (rather than at testing closeness to the all-zero string), it is not even clear what the estimate actually estimates when the sampling strategy is applied to a n -qubit quantum system, since we cannot speak of the Hamming weight of a quantum state. Furthermore, how can we quantify in a meaningful way how accurate a sampling strategy is, and how hard is it to compute (good bounds on) the accuracy of different sampling strategies, when applied to a quantum population? Finally, a last subtlety that is inherent to the quantum setting is that the execution of a sampling strategy actually changes the state of A due to the measurements.

In this work, we present a framework that answers the above questions and allows us to fully understand how a classical sampling strategy behaves when applied to a quantum population, i.e., to a n -qubit system or, more general, to n copies of an arbitrary “atomic” system. Our framework incorporates the following. First, we specify an abstract property on the state of A (after the measurements done by the sampling strategy), with the intended meaning that this is the property one should conclude from the outcome of the sampling strategy when applied to A . We also demonstrate that this property has useful consequences: specifically, that a suitable measurement will lead to a high-entropy outcome; this is handy in particular for quantum-cryptographic purposes. Then, we define a meaningful measure, sort of a “quantum error probability” (although technically speaking it is not a probability), that tells how reliable it is to conclude the specified property from the outcome of the sampling strategy. Finally, we show that for *any* sampling strategy, the quantum error probability of the strategy, as we define it, is bounded by the square-root of its classical error probability. This means that in order to understand how well a sampling strategy performs in the quantum setting, it suffices to analyze it in the classical setting, which is typically much simpler. Furthermore, for typical sampling strategies, like when picking the sample uniformly at random, there are well-known good bounds on the classical error probability.

We demonstrate the usefulness of our framework by means of two applications. Our applications do not constitute actual new results, but they provide new and simple(r) proofs for known results, both in the area of quantum cryptography. We take this as strong indication for the usefulness of the framework, and that the framework is likely to prove valuable in other applications as well.

The first application is to quantum oblivious transfer (QOT). It is well known that QOT is not possible from scratch; however, one can build a secure QOT scheme when given a bit-commitment (BC) primitive “for free”.¹ Like QOT, also QBC is impossible from scratch; nevertheless, the implication from BC to QOT is interesting from a theoretical point of view, since the corresponding implication does not hold in the classical setting. The existence of a QOT scheme based on a BC was suggested by Bennett *et al.* in 1991 [BBCS92];² however, no security proof was provided. Mayers and Salvail proved security of the QOT scheme against a restricted adversary that only performs *individual* measurements [MS94], and finally, in 1995, Yao gave a security proof against a general adversary, which is allowed to do fully *coherent* measurements [Yao95]. However, from today’s perspective, Yao’s proof is still not fully satisfactory:

¹We use BC and OT as short-hands of the respective abstract primitives, bit commitment and oblivious transfer, and we write QBC and QOT for potential schemes implementing the respective primitives in the quantum setting.

²At that time, QBC was thought to be possible, and thus the QOT scheme was claimed to be implementable from scratch.

it is very technical, without intuition and hard to follow, and it measures the adversary’s information in terms of “accessible information”, which has proven to be a too weak information measure.

Here, we show how our framework for analyzing sampling strategies in the quantum setting leads to a conceptually very simple and easy-to-understand security proof for QOT from BC. The proof essentially works as follows: When considering a purified version of the QOT scheme, the commit-and-open phase of the QOT scheme can be viewed as executing a specific sampling strategy. From the framework, it then follows that some crucial piece of information has high entropy from the adversary’s point of view. The proof is then concluded by applying the privacy amplification theorem. In recent work of the second author [DFL⁺09], it is shown that the same kind of analysis is not restricted to QOT but actually applies to a large class of two-party quantum-cryptographic schemes which are based on a commit-and-open phase.

The second application we discuss is to quantum key-distribution (QKD). Also here, our framework allows for a simple and easy-to-understand security proof, namely for the BB84 QKD scheme.³ Similar to our proof for QOT, we can view the checking phase of the BB84 scheme as executing a specific sampling strategy (although here some additional non-trivial observation needs to be made). From the framework, we can then conclude that the raw key has high entropy from the adversary’s point of view, and again privacy amplification finishes the job.

As for QOT, also QKD schemes initially came without security proofs, and proving QKD schemes rigorously secure turned out to be an extremely challenging and subtle task. Nowadays, though, the security of QKD schemes is better understood, and we know of various ways of proving, say, BB84 secure, ranging from Shor and Preskill’s proof based on quantum error-correcting codes to Renner’s approach using a quantum De Finetti theorem which allows to reduce security against general attacks to security against the much weaker class of so-called collective attacks. As such, our proof may safely be viewed as “yet another BB84 QKD proof”. Nevertheless, when compared to other proofs, it has some nice features: It provides an explicit and easy-to-compute expression for the security of the scheme (in contrast to most proofs in the literature which merely provide an asymptotic analysis), it does not require any “symmetrization of the qubits” (e.g. by applying a random permutation) from the protocol, and it is technically not very involved (e.g. compared to the proofs involving Renner’s quantum De Finetti theorem). Furthermore, it gives immediately a *direct* security proof, rather than a reduction to the security against collective attacks.

2 Notation, Terminology, and Some Tools

Strings and Hamming Weight. Throughout the paper, \mathcal{A} denotes some fixed finite alphabet with $0 \in \mathcal{A}$. It is safe to think of \mathcal{A} as $\{0, 1\}$, but our claims also hold for larger alphabets. For a string $\mathbf{q} = (q_1, \dots, q_n) \in \mathcal{A}^n$ of arbitrary length $n \geq 0$, the *Hamming weight* of \mathbf{q} is defined as the number of non-zero entries in \mathbf{q} : $\text{wt}(\mathbf{q}) := |\{i \in [n] : q_i \neq 0\}|$, where we use $[n]$ as short hand for $\{1, \dots, n\}$. We also use the notion of the *relative Hamming weight* of \mathbf{q} , defined as $\omega(\mathbf{q}) := \text{wt}(\mathbf{q})/n$. By convention, the relative Hamming weight of the empty string \perp is set to $\omega(\perp) := 0$. For a string $\mathbf{q} = (q_1, \dots, q_n) \in \mathcal{A}^n$ and a subset $J \subset [n]$, we write $\mathbf{q}_J := (q_i)_{i \in J}$ for the restriction of \mathbf{q} to the positions $i \in J$.

³Actually, we prove security for an entanglement-based version of BB84, which was first proposed by Ekert, and which implies security for the original BB84 scheme.

Random Variables and Hoeffding’s Inequalities. Formally, a *random variable* is a function $X : \Omega \rightarrow \mathcal{X}$ with the sample space Ω of a probability space (Ω, \Pr) as domain, and some arbitrary finite set \mathcal{X} as range. The *distribution* of X , which we denote as P_X , is given by $P_X(x) = \Pr[X = x] = \Pr[\{\omega \in \Omega : X(\omega) = x\}]$. The joint distribution of two (or more) random variables X and Y is denoted by P_{XY} , i.e., $P_{XY}(x, y) = \Pr[X = x \wedge Y = y]$. Usually, we leave the probability space (Ω, \Pr) implicit, and understand random variables to be defined by their joint distribution, or by some “experiment” that uniquely determines their joint distribution. Random variables X and Y are *independent* if $P_{XY} = P_X P_Y$ (in the sense that $P_{XY}(x, y) = P_X(x)P_Y(y) \forall x \in \mathcal{X}, y \in \mathcal{Y}$).

We will make extensive use of Hoeffding’s inequalities for random sampling with and without replacement, as developed in [Hoe63]. The following theorem summarizes these inequalities, tailored to our needs.⁴

Theorem 1 (Hoeffding). *Let $\mathbf{b} \in \{0, 1\}^n$ be a bit string with relative Hamming weight $\mu = \omega(\mathbf{b})$. Let the random variables X_1, X_2, \dots, X_k be obtained by sampling k random entries from \mathbf{b} with replacement, i.e., the X_i ’s are independent and $P_{X_i}(1) = \mu$. Furthermore, let the random variables Y_1, Y_2, \dots, Y_k be obtained by sampling k random entries from \mathbf{b} without replacement. Then, for any $\delta > 0$, the random variables $\bar{X} := \frac{1}{k} \sum_i X_i$ and $\bar{Y} := \frac{1}{k} \sum_i Y_i$ satisfy*

$$\Pr[|\bar{Y} - \mu| \geq \delta] \leq \Pr[|\bar{X} - \mu| \geq \delta] \leq 2 \exp(-2\delta^2 k).$$

For the case of sampling without replacement, a slightly sharper bound was found by Serfling [Ser74]:

$$\Pr[|\bar{Y} - \mu| \geq \delta] \leq 2 \exp\left(-\frac{2\delta^2 kn}{n-k+1}\right).$$

Quantum Systems and States. We assume the reader to be familiar with the basic concepts of quantum information theory; we merely fix some terminology and notation here. A quantum system A is associated with a complex Hilbert space, $\mathcal{H} = \mathbb{C}^d$, its *state space*. The *state* of A is given, in the case of a *pure* state, by a norm-1 state vector $|\varphi\rangle \in \mathcal{H}$, respectively, in the case of a *mixed* state, by a trace-1 positive-semi-definite operator/matrix $\rho : \mathcal{H} \rightarrow \mathcal{H}$. In order to simplify language, we are sometimes a bit sloppy in distinguishing between a quantum system, its state, and the state vector or density matrix describing the state. By default, we write \mathcal{H}_A for the state space of system A , and ρ_A (respectively $|\varphi_A\rangle$ in case of a pure state) for the state of A .

The state space of a *bipartite* quantum system AB , consisting of two (or more) subsystems, is given by $\mathcal{H}_{AB} = \mathcal{H}_A \otimes \mathcal{H}_B$. If the state of AB is given by ρ_{AB} then the state of subsystem A , when treated as a stand-alone system, is given by the *partial trace* $\rho_A = \text{tr}_B(\rho_{AB})$, and correspondingly for B . *Measuring* a system A in basis $\{|i\rangle\}_{i \in I}$, where $\{|i\rangle\}_{i \in I}$ is an orthonormal basis of \mathcal{H}_A , means applying the measurement described by the projectors $\{|i\rangle\langle i|\}_{i \in I}$, such that outcome $i \in I$ is observed with probability $p_i = \text{tr}(|i\rangle\langle i|\rho_A)$ (respectively $p_i = |\langle i|\varphi_A\rangle|^2$ in case of a pure state). If A is a subsystem of a bipartite system AB , then it means applying the measurement described by the projectors $\{|i\rangle\langle i| \otimes \mathbb{I}_B\}_{i \in I}$, where \mathbb{I}_B is the identity operator on \mathcal{H}_B .

A *qubit* is a quantum system A with state space $\mathcal{H}_A = \mathbb{C}^2$. The *computational basis* $\{|0\rangle, |1\rangle\}$ (for a qubit) is given by $|0\rangle = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$ and $|1\rangle = \begin{pmatrix} 0 \\ 1 \end{pmatrix}$, and the *Hadamard basis* by $H\{|0\rangle, |1\rangle\} = \{H|0\rangle, H|1\rangle\}$, where H denotes the 2-dimensional *Hadamard matrix* $H = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$. The state space of an n -qubit

⁴Interestingly, the inequality with respect to random sampling *without* replacement does not seem to be very commonly known.

system $A = A_1 \cdots A_n$ is given by $\mathcal{H}_A = (\mathbb{C}^2)^{\otimes n} = \mathbb{C}^2 \otimes \cdots \otimes \mathbb{C}^2$. For $\mathbf{x} = (x_1, \dots, x_n)$ and $\boldsymbol{\theta} = (\theta_1, \dots, \theta_n)$ in $\{0, 1\}^n$, we write $|\mathbf{x}\rangle$ for $|\mathbf{x}\rangle = |x_1\rangle \cdots |x_n\rangle$ and $H^\boldsymbol{\theta}$ for $H^\boldsymbol{\theta} = H^{\theta_1} \otimes \cdots \otimes H^{\theta_n}$, and thus $H^\boldsymbol{\theta}|\mathbf{x}\rangle$ for $H^\boldsymbol{\theta}|\mathbf{x}\rangle = H^{\theta_1}|x_1\rangle \cdots H^{\theta_n}|x_n\rangle$. Finally, we write $\{|0\rangle, |1\rangle\}^{\otimes n} = \{|\mathbf{x}\rangle : \mathbf{x} \in \{0, 1\}^n\}$ for the computational basis on an n -qubit system, and $H^\boldsymbol{\theta}\{|0\rangle, |1\rangle\}^{\otimes n} = \{H^\boldsymbol{\theta}|\mathbf{x}\rangle : \mathbf{x} \in \{0, 1\}^n\} = H^{\theta_1}\{|0\rangle, |1\rangle\} \otimes \cdots \otimes H^{\theta_n}\{|0\rangle, |1\rangle\}$ for the basis that is made up of the computational basis on the subsystems A_i with $\theta_i = 0$ and of the Hadamard basis on the subsystems A_i with $\theta_i = 1$. In order to simplify notation, we will sometimes abuse terminology and speak of the basis $\boldsymbol{\theta}$ when we actually mean $H^\boldsymbol{\theta}\{|0\rangle, |1\rangle\}^{\otimes n}$.

We measure closeness of two states ρ and σ by their *trace distance*: $\Delta(\rho, \sigma) := \frac{1}{2}\text{tr}|\rho - \sigma|$, where for any square matrix M , $|M|$ denotes the positive-semi-definite square-root of $M^\dagger M$. For *pure* states $|\varphi\rangle$ and $|\psi\rangle$, the trace distance of the corresponding density matrices coincides with $\Delta(|\varphi\rangle\langle\varphi|, |\psi\rangle\langle\psi|) = \sqrt{1 - |\langle\varphi|\psi\rangle|^2}$. If the states of two systems A and B are ϵ -close, i.e. $\Delta(\rho_A, \rho_B) \leq \epsilon$, then A and B cannot be distinguished with advantage greater than ϵ ; in other words, A behaves exactly like B , except with probability ϵ .

Classical and Hybrid Systems (and States). Subsystem X of a bipartite quantum system XE is called *classical*, if the state of XE is given by a density matrix of the form

$$\rho_{XE} = \sum_{x \in \mathcal{X}} P_X(x) |x\rangle\langle x| \otimes \rho_E^x,$$

where \mathcal{X} is a finite set of cardinality $|\mathcal{X}| = \dim(\mathcal{H}_X)$, $P_X : \mathcal{X} \rightarrow [0, 1]$ is a probability distribution, $\{|x\rangle\}_{x \in \mathcal{X}}$ is some fixed orthonormal basis of \mathcal{H}_X , and ρ_E^x is a density matrix on \mathcal{H}_E for every $x \in \mathcal{X}$. Such a state, called *hybrid* or *cq-* (for *classical-quantum*) state, can equivalently be understood as consisting of a *random variable* X with distribution P_X , taking on values in \mathcal{X} , and a system E that is in state ρ_E^x exactly when X takes on the value x . This formalism naturally extends to two (or more) classical systems X, Y etc.

If the state of XE satisfies $\rho_{XE} = \rho_X \otimes \rho_E$, where $\rho_X = \text{tr}_E(\rho_{XE}) = \sum_x P_X(x) |x\rangle\langle x|$ and $\rho_E = \text{tr}_X(\rho_{XE}) = \sum_x P_X(x) \rho_E^x$, then X is *independent* of E , and thus no information on X can be obtained from system E . Moreover, if $\rho_{XE} = \frac{1}{|\mathcal{X}|} \mathbb{I}_X \otimes \rho_E$, where \mathbb{I}_X denotes the identity on \mathcal{H}_X , then X is *random-and-independent* of E . This is what is aimed for in quantum cryptography, when X represents a classical cryptographic key and E the adversary's potential quantum information on X .

It is not too hard to see that for two hybrid states ρ_{XE} and $\rho_{XE'}$ with the same (distribution of) X , the trace distance between ρ_{XE} and $\rho_{XE'}$ can be computed as $\Delta(\rho_{XE}, \rho_{XE'}) = \sum_x P_X(x) \Delta(\rho_E^x, \rho_{E'}^x)$.

Min-Entropy and Privacy Amplification. We make use of Renner's notion of the *conditional min-entropy* $H_{\min}(\rho_{XE}|E)$ of a system X conditioned on another system E [Ren05]. Although the notion makes sense for arbitrary states, we restrict to hybrid states ρ_{XE} with classical X . If the hybrid state ρ_{XE} is clear from the context, we may write $H_{\min}(X|E)$ instead of $H_{\min}(\rho_{XE}|E)$. The formal definition, given by $H_{\min}(\rho_{XE}|E) := \sup_{\sigma_E} \max\{h \in \mathbb{R} : 2^{-h} \cdot \mathbb{I}_X \otimes \sigma_E - \rho_{XE} \geq 0\}$ where the supremum is over all density matrices σ_E on \mathcal{H}_E , is not very relevant to us; we merely rely on some elementary properties. For instance, the *chain rule* guarantees that $H_{\min}(X|YE) \geq H_{\min}(XY|E) - \log(|\mathcal{Y}|) \geq H_{\min}(X|E) - \log(|\mathcal{Y}|)$ for classical X and Y with respective ranges \mathcal{X} and \mathcal{Y} , where here and throughout the article \log denotes the *binary* logarithm, whereas \ln denotes the *natural* logarithm. Furthermore, it holds that if E' is obtained from E by measuring (part of) E , then $H_{\min}(X|E') \geq H_{\min}(X|E)$.

Finally, we make use of Renner’s privacy amplification theorem [RK05, Ren05], as given below. Recall that a function $g : \mathcal{R} \times \mathcal{X} \rightarrow \{0, 1\}^\ell$ is called a *universal* (hash) function, if for the random variable R , uniformly distributed over \mathcal{R} , and for any distinct $x, y \in \mathcal{X}$: $\Pr[g(R, x) = g(R, y)] \leq 2^{-\ell}$.

Theorem 2 (Privacy amplification). *Let ρ_{XE} be a hybrid state with classical X . Let $g : \mathcal{R} \times \mathcal{X} \rightarrow \{0, 1\}^\ell$ be a universal hash function, and let R be uniformly distributed over \mathcal{R} , independent of X and E . Then $K = g(R, X)$ satisfies*

$$\Delta(\rho_{KRE}, \frac{1}{|\mathcal{K}|} \mathbb{I}_K \otimes \rho_{RE}) \leq \frac{1}{2} \cdot 2^{-\frac{1}{2}(\text{H}_{\min}(X|E) - \ell)}.$$

Informally, Theorem 2 states that if X contains sufficiently more than ℓ bits of entropy when given E , then ℓ nearly random-and-independent bits can be extracted from X .

3 Sampling in a Classical Population

As a warm-up, and in order to study some useful examples and introduce some convenient notation, we start with the classical sampling problem, which is rather well-understood.

3.1 Sampling Strategies

Let $\mathbf{q} = (q_1, \dots, q_n) \in \mathcal{A}^n$ be a string of given length n . We consider the problem of estimating the relative Hamming weight $\omega(\mathbf{q})$ by only looking at a substring \mathbf{q}_t of \mathbf{q} , for a small subset $t \subset [n]$.⁵ Actually, we are interested in the equivalent problem of estimating the relative Hamming weight $\omega(\mathbf{q}_{\bar{t}})$ of the *remaining* string $\mathbf{q}_{\bar{t}}$, where \bar{t} is the complement $\bar{t} = [n] \setminus t$ of t .⁶ A canonical way to do so would be to sample a uniformly random subset (say, of a certain small size) of positions, and compute the relative Hamming weight of the sample as estimate. Very generally, we allow any strategy that picks a subset $t \subset [n]$ according to some probability distribution and computes the estimate for $\omega(\mathbf{q}_{\bar{t}})$ as some (possibly randomized) function of t and \mathbf{q}_t , i.e., as $f(t, \mathbf{q}_t, s)$ for a *seed* s that is sampled according to some probability distribution. This motivates the following formal definition.

Definition 1 (Sampling strategy). *A sampling strategy Ψ consists of a triple (P_T, P_S, f) , where P_T is a distribution over the subsets of $[n]$, P_S is a (independent) distribution over a finite set \mathcal{S} , and f is a function*

$$f : \{(t, v) : t \subset [n], v \in \mathcal{A}^{|t|}\} \times \mathcal{S} \rightarrow \mathbb{R}.$$

We stress that a sampling strategy Ψ , as defined here, specifies how to choose the sample subset as well as how to compute the estimate from the sample (thus a more appropriate but lengthy name would be a “sample-and-estimate strategy”).

Remark 1. *By definition, the choice of the seed s is specified to be independent of t , i.e., $P_{TS} = P_T P_S$. Sometimes, however, it is convenient to allow s to depend on t . We can actually do so without*

⁵More generally, we may consider the problem of estimating the Hamming distance of \mathbf{q} to some arbitrary reference string \mathbf{q}_0 ; but this can obviously be done simply by estimating the Hamming weight of $\mathbf{q}' = \mathbf{q} - \mathbf{q}_0$.

⁶The reason for this, as will become clear later, is that in our applications, the sampled positions within \mathbf{q} will be *discarded*, and thus we will be interested merely in the remaining positions.

contradicting Definition 1. Namely, to comply with the independence requirement, we would simply choose a (typically huge) “container” seed that contains a seed for every possible choice of t , each one chosen with the corresponding distribution, and it is then part of f 's task, when given t , to select the seed that is actually needed out of the container seed.⁷

A sampling strategy Ψ can obviously also be used to *test* if \mathbf{q} (or actually $\mathbf{q}_{\bar{t}}$) is close to the all-zero string $0 \cdots 0$: compute the estimate for $\omega(\mathbf{q}_{\bar{t}})$ as dictated by Ψ , and *accept* if the estimate vanishes and else *reject*.

We briefly discuss five example sampling strategies. The examples should illustrate the generality of the definition, and some of the examples will be used later on; however, the reader is free to skip (some of) them. We start with the canonical example mentioned in the beginning.

Example 1 (Random sampling *without* replacement).

In random sampling without replacement, k *distinct* indices i_1, \dots, i_k within $[n]$ are chosen uniformly at random, where k is some parameter, and the relative Hamming weight of $\mathbf{q}_{\{i_1, \dots, i_k\}}$ is used as estimate for $\omega(\mathbf{q}_{\bar{t}})$. Formally, this sampling strategy is given by $\Psi = (P_T, P_S, f)$ where $P_T(t) = 1/\binom{n}{k}$ if $|t| = k$ and else $P_T(t) = 0$, $\mathcal{S} = \{\perp\}$ and thus $P_S(\perp) = 1$, and $f(t, \mathbf{q}_t, \perp) = f(t, \mathbf{q}_t) = \omega(\mathbf{q}_t)$. \diamond

With the second example, we show that also sampling with replacement is captured by our definition.

Example 2 (Random sampling *with* replacement).

In random sampling with replacement, k indices i_1, \dots, i_k are chosen independently uniformly at random within $[n]$, where k is some parameter, and the relative Hamming weight of the string $(q_{i_1}, \dots, q_{i_k})$ is used as estimate for $\omega(\mathbf{q}_{\bar{t}})$. Note that here i_ℓ may coincide with $i_{\ell'}$ for $\ell \neq \ell'$, in which case $(q_{i_1}, \dots, q_{i_k})$ is not equal to $\mathbf{q}_{\{i_1, \dots, i_k\}}$. To make this fit into Definition 1, we set t to be $\{i_1, \dots, i_k\}$, and we let $f(t, \mathbf{q}_t, s)$ be given by $\omega(q_{j_1}, \dots, q_{j_k})$, where j_1, \dots, j_k is determined by the seed s among all possibilities with $\{j_1, \dots, j_k\} = t$. It is cumbersome and of no importance to us to determine the correct distributions P_T and P_S for t and s , respectively; it is sufficient to realize that random sampling with replacement *is* captured by Definition 1. \diamond

Next, we sample by picking a uniformly random subset (without restricting its size).

Example 3 (Uniformly random subset sampling).

The sample set t is chosen as a uniformly random subset of $[n]$, and the estimate is computed as the relative Hamming weight of the sample \mathbf{q}_t . Formally, $P_T(t) = 1/2^n$ for any $t \subseteq [n]$, and $\mathcal{S} = \{\perp\}$ and $f(t, \mathbf{q}_t, \perp) = f(t, \mathbf{q}_t) = \omega(\mathbf{q}_t)$. \diamond

As a fourth example, we consider a somewhat unnatural and in some sense non-optimal sampling strategy. This example, though, will be of use in our analysis of quantum oblivious transfer in Section 5.

Example 4 (Random sampling without replacement, using only part of the sample).

This example can be viewed as a composition of Example 1 and 3. Namely, t is chosen as a random subset of fixed size k , as in Example 1, so that $P_T(t) = 1/\binom{n}{k}$ for $t \subseteq [n]$ with $|t| = k$. But now, only part of the sample \mathbf{q}_t is used to compute the estimate. Namely, the estimate is computed as

$$f(t, \mathbf{q}_t, s) = \omega(\mathbf{q}_s).$$

⁷Alternatively, we could simply drop the independence requirement in Definition 1; however, we feel it is conceptually easier to think of the seed as being independently chosen.

where the seed s is chosen as a uniformly random subset s of t ; i.e., $P_S(s) = 1/2^t$ for any $s \subseteq t$. Recall from Remark 1 that the choice of s is allowed to depend on t . We would like to point out that when we use Example 4 in Section 5, it is useful that the restriction to the subset s is part of the evaluation of f , rather than part of the selection of the sample subset t . \diamond

In the fifth example we consider another somewhat unnatural sampling strategy, which though will be useful for the QKD proof in Section 6.

Example 5 (Pairwise one-out-of-two sampling, using only part of the sample).

For this example, it is convenient to consider the index set from which the subset t is chosen, to be of the form $[n] \times \{0, 1\}$. Namely, we consider the string $\mathbf{q} \in \mathcal{A}^{2n}$ to be indexed by *pairs* of indices, $\mathbf{q} = (q_{ij})$, where $i \in [n]$ and $j \in \{0, 1\}$; in other words, we consider \mathbf{q} to consist of n pairs (q_{i0}, q_{i1}) . The subset $t \subset [n] \times \{0, 1\}$ is chosen as $t = \{(1, j_1), \dots, (n, j_n)\}$ where every j_k is picked independently at random in $\{0, 1\}$. In other words, t selects one element from each pair (q_{i0}, q_{i1}) . Furthermore, the estimate for $\omega(\mathbf{q}_{\bar{t}})$ is computed from \mathbf{q}_t as $f(t, \mathbf{q}_t, s) = \omega(\mathbf{q}_s)$ where the seed s is a random subset $s \subset t$ of size k . \diamond

Example 6 (Pairwise *biased* one-out-of-two sampling, using only part of the sample).

In this example we consider a similar situation as in Example 5, except that we now construct t by sampling every j_k according to the *Bernoulli* distribution $(p, 1-p)$. Consequently, we compute the estimate for $\omega(\mathbf{q}_{\bar{t}})$ slightly differently, but we will make this clear in Appendix A.6. \diamond

3.2 The Error Probability

After having introduced the general notion of a sampling strategy, we next want to define a measure that captures for a given sampling strategy how well it performs, i.e., with what probability the estimate, $f(t, \mathbf{q}_t, s)$, is how close to the real value, $\omega(\mathbf{q}_{\bar{t}})$. For the definition, it will be convenient to introduce the following notation. For a given sampling strategy $\Psi = (P_T, P_S, f)$, consider arbitrary but fixed choices for the subset $t \subset [n]$ and the seed $s \in \mathcal{S}$ with $P_T(t) > 0$ and $P_S(s) > 0$. Furthermore, fix an arbitrary $\delta > 0$. Define $B_{t,s}^\delta(\Psi) \subseteq \mathcal{A}^n$ as

$$B_{t,s}^\delta(\Psi) := \{\mathbf{b} \in \mathcal{A}^n : |\omega(\mathbf{b}_{\bar{t}}) - f(t, \mathbf{b}_t, s)| < \delta\},$$

i.e., as the set of all strings \mathbf{q} for which the estimate is δ -close to the real value, assuming that subset t and seed s have been used. To simplify notation, if Ψ is clear from the context, we simply write $B_{t,s}^\delta$ instead of $B_{t,s}^\delta(\Psi)$. By replacing the specific values t and s by the corresponding (independent) random variables T and S , with distributions P_T and P_S , respectively, we obtain the *random variable* $B_{T,S}^\delta$, whose range consists of subsets of \mathcal{A}^n . By means of this random variable, we now define the *error probability* of a sampling strategy as follows.

Definition 2 (Error probability). *The (classical) error probability of a sampling strategy $\Psi = (P_T, P_S, f)$ is defined as the following value, parametrized by $0 < \delta < 1$:*

$$\varepsilon_{\text{class}}^\delta(\Psi) := \max_{\mathbf{q} \in \mathcal{A}^n} \Pr \left[\mathbf{q} \notin B_{T,S}^\delta(\Psi) \right].$$

By definition of the error probability, it is guaranteed that for any string $\mathbf{q} \in \mathcal{A}^n$, the estimated value is δ -close to the real value except with probability at most $\varepsilon_{\text{class}}^\delta(\Psi)$. When used as a sampling strategy to

test closeness to the all-zero string, $\varepsilon_{\text{class}}^\delta(\Psi)$ determines the probability of accepting even though \mathbf{q}_i is “not close” to the all-zero string, in the sense that its relative Hamming weight exceeds δ . Whenever Ψ is clear from the context, we will write $\varepsilon_{\text{class}}^\delta$ instead of $\varepsilon_{\text{class}}^\delta(\Psi)$.

In Appendix A, we analyze the error probabilities for the sampling strategies considered in Examples 1 to 5, excluding Example 2, and we show them all to be exponentially small by applying Hoeffding’s inequality in a suitable way.

4 Sampling in a Quantum Population

We now want to study the behavior of a sampling strategy when applied to a quantum population. More specifically, let $A = A_1 \cdots A_n$ be an n -partite quantum system, where the state space of each system A_i equals $\mathcal{H}_{A_i} = \mathbb{C}^d$ with $d = |\mathcal{A}|$, and let $\{|a\rangle\}_{a \in \mathcal{A}}$ be a fixed orthonormal basis of \mathbb{C}^d . We allow A to be entangled with some additional system E with arbitrary finite-dimensional state-space \mathcal{H}_E . We may assume the joint state of AE to be pure, and as such be given by a state vector $|\varphi_{AE}\rangle \in \mathcal{H}_A \otimes \mathcal{H}_E$; if not, then it can be purified by increasing the dimension of \mathcal{H}_E .

Similar to the classical sampling problem of testing closeness to the all-zero string, we can consider here the problem of testing if the state of A is close to the all-zero *reference state* $|\varphi_A^\circ\rangle = |0\rangle \cdots |0\rangle$ by looking at, which here means *measuring*, only a few of the subsystems of A . More generally, we will be interested in the sampling problem of estimating the “Hamming weight of the state of A ”, although it is not clear at the moment what this should mean. Actually, like in the classical case, we are interested in testing closeness to the all-zero state, respectively estimating the Hamming weight, of the *remaining subsystems* of A .

It is obvious that a sampling strategy $\Psi = (P_T, P_S, f)$ can be applied in a straightforward way to the setting at hand: sample t according to P_T , measure the subsystems A_i with $i \in t$ in basis $\{|a\rangle\}_{a \in \mathcal{A}}$ to observe $\mathbf{q}_t \in \mathcal{A}^{|t|}$, and compute the estimate as $f(t, \mathbf{q}_t, s)$ for s chosen according to P_S (respectively, for testing closeness to the all-zero state, accept or reject depending on the value of the estimate). However, it is a-priori *not* clear, how to interpret the outcome. Measuring a random subset of the subsystems of A and observing 0 all the time indeed seems to suggest that the original state of A , and thus the remaining subsystems, must be in some sense close to the all-zero state; but what is the right way to formalize this? In the case of a general sampling strategy for estimating the (relative) Hamming weight, what does the estimate actually estimate? And, do all strategies that perform well in the classical setting also perform well in the quantum setting?

We give in this section a rigorous analysis of sampling strategies when applied to a n -partite quantum system A . Our analysis completely answers above concerns. Later in the paper, we demonstrate the usefulness of our analysis of sampling strategies for studying and analyzing quantum-cryptographic schemes.

4.1 Analyzing Sampling Strategies in the Quantum Setting

We start by suggesting the property on the remaining subsystems of A that one should expect to be able to conclude from the outcome of a sampling strategy. A somewhat natural approach is as follows.

Definition 3. For system AE , and similarly for any subsystem of A , we say that the state $|\varphi_{AE}\rangle$ of AE has relative Hamming weight β within A if it is of the form $|\varphi_{AE}\rangle = |\mathbf{b}\rangle|\varphi_E\rangle$ with $\mathbf{b} \in \mathcal{A}^n$ and $\omega(\mathbf{b}) = \beta$.

Now, given the outcome $f(t, \mathbf{q}_t, s)$ of a sampling strategy when applied to A , we want to be able to conclude that, up to a small error, the state of the remaining subsystem $A_{\bar{t}}E$ is a *superposition* of states with relative Hamming weight close to $f(t, \mathbf{q}_t, s)$ within $A_{\bar{t}}$. To analyze this, we extend some of the notions introduced in the classical setting. Recall the definition of $B_{t,s}^\delta$, consisting of all strings $\mathbf{b} \in \mathcal{A}^n$ with $|\omega(\mathbf{b}_{\bar{t}}) - f(t, \mathbf{b}_t, s)| < \delta$. By slightly abusing notation, we extend this notion to the quantum setting and write

$$\text{span}(B_{t,s}^\delta) := \text{span}(\{|\mathbf{b}\rangle : \mathbf{b} \in B_{t,s}^\delta\}) = \text{span}(\{|\mathbf{b}\rangle : |\omega(\mathbf{b}_{\bar{t}}) - f(t, \mathbf{b}_t, s)| < \delta\}).$$

Note that if the state $|\varphi_{AE}\rangle$ of AE happens to be in $\text{span}(B_{t,s}^\delta) \otimes \mathcal{H}_E$ for some t and s , and if exactly these t and s are chosen when applying the sampling strategy to A , then *with certainty* the state of $A_{\bar{t}}E$ (after the measurement) is in a superposition of states with relative Hamming weight δ -close to $f(t, \mathbf{q}_t, s)$ within $A_{\bar{t}}$, regardless of the measurement outcome \mathbf{q}_t .

Next, we want to extend the notion of error probability (Definition 2) to the quantum setting. The following approach turns out to be fruitful. We consider the *hybrid* system $TSAE$, consisting of the classical random variables T and S with distribution $P_{TS} = P_T P_S$, describing the choices of t and s , respectively, and of the actual quantum systems A and E . The state of $TSAE$ is given by

$$\rho_{TSAE} = \sum_{t,s} P_{TS}(t, s) |t, s\rangle\langle t, s| \otimes |\varphi_{AE}\rangle\langle\varphi_{AE}|.$$

Note that TS is independent of AE : $\rho_{TSAE} = \rho_{TS} \otimes \rho_{AE}$; indeed, in a sampling strategy t and s are chosen independently of the state of AE . We compare this *real* state of $TSAE$ with an *ideal* state which is of the form

$$\tilde{\rho}_{TSAE} = \sum_{t,s} P_{TS}(t, s) |t, s\rangle\langle t, s| \otimes |\tilde{\varphi}_{AE}^{ts}\rangle\langle\tilde{\varphi}_{AE}^{ts}| \quad \text{with} \quad |\tilde{\varphi}_{AE}^{ts}\rangle \in \text{span}(B_{t,s}^\delta) \otimes \mathcal{H}_E \quad \forall t, s \quad (1)$$

for some given $\delta > 0$. Thus, T and S have the same distribution as in the real state, but here we allow AE to depend on T and S , and for each particular choice t and s for T and S , respectively, we require the state of AE to be in $\text{span}(B_{t,s}^\delta) \otimes \mathcal{H}_E$. Thus, in an “ideal world” where the state of the hybrid system $TSAE$ is given by $\tilde{\rho}_{TSAE}$, it holds *with certainty* that the state $|\psi_{A_{\bar{t}}E}\rangle$ of $A_{\bar{t}}E$, after having measured A_t and having observed \mathbf{q}_t , is in a superposition of states with relative Hamming weight δ -close to $\beta := f(t, \mathbf{q}_t, s)$ within $A_{\bar{t}}$. We now define the quantum error probability of a sampling strategy by looking at how far away the closest ideal state $\tilde{\rho}_{TSAE}$ is from the real state ρ_{TSAE} .

Definition 4 (Quantum error probability). *The quantum error probability of a sampling strategy $\Psi = (P_T, P_S, f)$ is defined as the following value, parametrized by $0 < \delta < 1$:*

$$\varepsilon_{\text{quant}}^\delta(\Psi) = \max_{\mathcal{H}_E} \max_{|\varphi_{AE}\rangle} \min_{\tilde{\rho}_{TSAE}} \Delta(\rho_{TSAE}, \tilde{\rho}_{TSAE}),$$

where the first max is over all finite-dimensional state spaces \mathcal{H}_E , the second max is over all state vectors $|\varphi_{AE}\rangle \in \mathcal{H}_A \otimes \mathcal{H}_E$, and the min is over all ideal states $\tilde{\rho}_{TSAE}$ as in (1).⁸

⁸It is not too hard to see, in particular after having gained some more insight via the proof of Theorem 3 below, that these min and max exist.

As with $B_{t,s}^\delta$ and $\varepsilon_{\text{class}}^\delta$, we simply write $\varepsilon_{\text{quant}}^\delta$ when Ψ is clear from the context. We stress the meaningfulness of the definition: it guarantees that on average over the choice of t and s , the state of $A_{\bar{t}}E$ is $\varepsilon_{\text{quant}}^\delta$ -close to a superposition of states with Hamming weight δ -close to $f(t, \mathbf{q}_t, s)$ within $A_{\bar{t}}$, and as such it *behaves* like a superposition of such states, except with probability $\varepsilon_{\text{quant}}^\delta$. We will argue below and demonstrate in the subsequent sections that being (close to) a superposition of states with given approximate (relative) Hamming weight has some useful consequences.

Remark 2. *Similarly to footnote 5, also here the results of the section immediately generalize from the all-zero reference state $|0\rangle \cdots |0\rangle$ to an arbitrary reference state $|\varphi_A^\circ\rangle$ of the form $|\varphi_A^\circ\rangle = U_1|0\rangle \otimes \cdots \otimes U_n|0\rangle$ for unitary operators U_i acting on \mathbb{C}^d . Indeed, the generalization follows simply by a suitable change of basis, defined by the U_i 's. Or, in the special case where $\mathcal{A} = \{0, 1\}$ and*

$$|\varphi_A^\circ\rangle = H^{\hat{\theta}}|\hat{\mathbf{x}}\rangle = H^{\hat{\theta}_1}|\hat{x}_1\rangle \otimes \cdots \otimes H^{\hat{\theta}_n}|\hat{x}_n\rangle$$

for a fixed reference basis $\hat{\theta} \in \{0, 1\}^n$ and a fixed reference string $\hat{\mathbf{x}} \in \{0, 1\}^n$, we can, alternatively, replace in the definitions and results the computational by the Hadamard basis whenever $\hat{\theta}_i = 1$, and speak of the (relative) Hamming distance to $\hat{\mathbf{x}}$ rather than of the (relative) Hamming weight.

4.2 The Quantum vs. the Classical Error Probability

It remains to discuss how difficult it is to actually *compute* the quantum error probability for given sampling strategies, and how the *quantum* error probability $\varepsilon_{\text{quant}}^\delta$ relates to the corresponding *classical* error probability $\varepsilon_{\text{class}}^\delta$. To this end, we show the following simple relationship between $\varepsilon_{\text{quant}}^\delta$ and $\varepsilon_{\text{class}}^\delta$.

Theorem 3. *For any sampling strategy Ψ and for any $\delta > 0$:*

$$\varepsilon_{\text{quant}}^\delta(\Psi) \leq \sqrt{\varepsilon_{\text{class}}^\delta(\Psi)}.$$

As a consequence of this theorem, it suffices to analyze a sampling strategy in the classical setting, which is much easier, in order to understand how it behaves in the quantum setting. In particular, sampling strategies that are known to behave well in the classical setting, like examples 1 to 5, are also automatically guaranteed to behave well in the quantum setting. We will use this in the application sections.

Our bound on $\varepsilon_{\text{quant}}^\delta$ is in general tight. Indeed, in Appendix C we show tightness for an explicit class of sampling strategies, which e.g. includes Example 1 and Example 5. Here, we just mention the tightness result.

Proposition 1. *There exist natural sampling strategies for which the inequality in Theorem 3 is an equality.*

Proof of Theorem 3. We need to show that for any $|\varphi_{AE}\rangle \in \mathcal{H}_A \otimes \mathcal{H}_E$, with arbitrary \mathcal{H}_E , there exists a suitable ideal state $\tilde{\rho}_{TSAE}$ with $\Delta(\rho_{TSAE}, \tilde{\rho}_{TSAE}) \leq (\varepsilon_{\text{class}}^\delta)^{1/2}$. We construct $\tilde{\rho}_{TSAE}$ as in (1), where the $|\tilde{\varphi}_{AE}^{ts}\rangle$'s are defined by the following decomposition.

$$|\varphi_{AE}\rangle = \langle \tilde{\varphi}_{AE}^{ts} | \varphi_{AE} \rangle |\tilde{\varphi}_{AE}^{ts}\rangle + \langle \tilde{\varphi}_{AE}^{ts\perp} | \varphi_{AE} \rangle |\tilde{\varphi}_{AE}^{ts\perp}\rangle,$$

with $|\tilde{\varphi}_{AE}^{ts}\rangle \in \text{span}(B_{t,s}^\delta) \otimes \mathcal{H}_E$, $|\tilde{\varphi}_{AE}^{ts\perp}\rangle \in \text{span}(B_{t,s}^\delta)^\perp \otimes \mathcal{H}_E$ and $|\langle \tilde{\varphi}_{AE}^{ts} | \varphi_{AE} \rangle|^2 + |\langle \tilde{\varphi}_{AE}^{ts\perp} | \varphi_{AE} \rangle|^2 = 1$. In other words, $|\tilde{\varphi}_{AE}^{ts}\rangle$ is obtained as the re-normalized projection of $|\varphi_{AE}\rangle$ into $\text{span}(B_{t,s}^\delta) \otimes \mathcal{H}_E$. Note

that $|\langle \tilde{\varphi}_{AE}^{ts\perp} | \varphi_{AE} \rangle|^2$ equals the probability $\Pr[\mathbf{Q} \notin B_{t,s}^\delta]$, where the random variable \mathbf{Q} is obtained by measuring subsystem A of $|\varphi_{AE}\rangle$ in basis $\{|a\rangle\}_{a \in \mathcal{A}}^{\otimes n}$. Furthermore,

$$\sum_{t,s} P_{TS}(t,s) |\langle \tilde{\varphi}_{AE}^{ts\perp} | \varphi_{AE} \rangle|^2 = \sum_{t,s} P_{TS}(t,s) \Pr[\mathbf{Q} \notin B_{t,s}^\delta] = \Pr[\mathbf{Q} \notin B_{T,S}^\delta] = \sum_{\mathbf{q}} P_{\mathbf{Q}}(\mathbf{q}) \Pr[\mathbf{q} \notin B_{T,S}^\delta],$$

where by definition of $\varepsilon_{\text{class}}^\delta$, the latter is upper bounded by $\varepsilon_{\text{class}}^\delta$. From elementary properties of the trace distance, and using Jensen's inequality, we can now conclude that

$$\begin{aligned} \Delta(\rho_{TSAE}, \tilde{\rho}_{TSAE}) &= \sum_{t,s} P_{TS}(t,s) \Delta(|\varphi_{AE}\rangle\langle\varphi_{AE}|, |\tilde{\varphi}_{AE}^{ts}\rangle\langle\tilde{\varphi}_{AE}^{ts}|) = \sum_{t,s} P_{TS}(t,s) \sqrt{1 - |\langle \tilde{\varphi}_{AE}^{ts} | \varphi_{AE} \rangle|^2} \\ &= \sum_{t,s} P_{TS}(t,s) |\langle \tilde{\varphi}_{AE}^{ts\perp} | \varphi_{AE} \rangle| \leq \sqrt{\sum_{t,s} P_{TS}(t,s) |\langle \tilde{\varphi}_{AE}^{ts\perp} | \varphi_{AE} \rangle|^2} \leq \sqrt{\varepsilon_{\text{class}}^\delta}, \end{aligned}$$

which was to be shown. \square

As a side remark, we point out that the particular ideal state $\tilde{\rho}_{TSAE}$ constructed in the proof minimizes the distance to ρ_{TSAE} ; this follows from the so-called Hilbert projection theorem.

4.3 Superpositions with a Small Number of Terms

We give here some argument why being (close to) a superposition of states with a given approximate Hamming weight may be a useful property in the analyses of quantum-cryptographic schemes. For simplicity, and since this will be the case in our applications, we now restrict to the binary case where $\mathcal{A} = \{0, 1\}$. Our argument is based on the following lemma, which follows immediately from Lemma 3.1.13 in [Ren05]; for completeness, we give a direct proof of Lemma 1 in Appendix B. Informally, it states that measuring (part of) a *superposition* of a small number of orthogonal states produces a similar amount of uncertainty as when measuring the *mixture* of these orthogonal states.

Lemma 1. *Let A and E be arbitrary quantum systems, let $\{|i\rangle\}_{i \in I}$ and $\{|w\rangle\}_{w \in \mathcal{W}}$ be orthonormal bases of \mathcal{H}_A , and let $|\varphi_{AE}\rangle$ and ρ_{AE}^{mix} be of the form*

$$|\varphi_{AE}\rangle = \sum_{i \in J} \alpha_i |i\rangle |\varphi_E^i\rangle \in \mathcal{H}_A \otimes \mathcal{H}_E \quad \text{and} \quad \rho_{AE}^{\text{mix}} = \sum_{i \in J} |\alpha_i|^2 |i\rangle\langle i| \otimes |\varphi_E^i\rangle\langle \varphi_E^i|$$

for some subset $J \subseteq I$. Furthermore, let ρ_{WE} and ρ_{WE}^{mix} describe the hybrid systems obtained by measuring subsystem A of $|\varphi_{AE}\rangle$ and ρ_{AE}^{mix} , respectively, in basis $\{|w\rangle\}_{w \in \mathcal{W}}$ to observe outcome W . Then,

$$H_{\min}(\rho_{WE}|E) \geq H_{\min}(\rho_{WE}^{\text{mix}}|E) - \log |J|.$$

We apply Lemma 1 to an n -qubit system A where $|\varphi_{AE}\rangle$ is a superposition of states with relative Hamming weight δ -close to β within A :⁹

$$|\varphi_{AE}\rangle = \sum_{\substack{\mathbf{b} \in \{0,1\}^n \\ |\omega(\mathbf{b}) - \beta| \leq \delta}} |\mathbf{b}\rangle |\varphi_E^{\mathbf{b}}\rangle.$$

⁹System A considered here corresponds to the subsystem $A_{\bar{t}}$ in the previous section, after having measured A_t of the ideal state.

It is well known that $|\{\mathbf{b} \in \{0, 1\}^n : |\omega(\mathbf{b}) - \beta| \leq \delta\}| \leq |\{\mathbf{b} \in \{0, 1\}^n : \omega(\mathbf{b}) \leq \beta + \delta\}| \leq 2^{h(\beta+\delta)n}$ for $\beta + \delta \leq \frac{1}{2}$, where the function $h : [0, 1] \rightarrow [0, 1]$ is the *binary entropy function*, defined as $h(p) = -(p \log(p) + (1 - p) \log(1 - p))$ for $0 < p < 1$ and as 0 for $p = 0$ or 1 .¹⁰

Since measuring qubits within a state $|\mathbf{b}\rangle$ in the *Hadamard* basis produces uniformly random bits, we can conclude the following.

Corollary 1. *Let A be an n -qubit system, let the state $|\varphi_{AE}\rangle$ of AE be a superposition of states with relative Hamming weight δ -close to β within A , where $\delta + \beta \leq \frac{1}{2}$, and let the random variable \mathbf{X} be obtained by measuring A in basis $H^\theta\{|0\rangle, |1\rangle\}^{\otimes n}$ for $\theta \in \{0, 1\}^n$. Then*

$$H_{\min}(\mathbf{X}|E) \geq \text{wt}(\theta) - h(\beta + \delta)n.$$

Consider now the following quantum-cryptographic setting. Bob prepares and hands over to Alice an n -qubit quantum system A , which ought to be in state $|\varphi_A^\circ\rangle = |0\rangle \cdots |0\rangle$. However, since Bob might be dishonest, the state of A could be anything, even entangled with some system E controlled by Bob. Our results now imply the following: Alice can apply a suitable sampling strategy to convince herself that the joint state of the remaining subsystem of A and of E is (close to) a superposition of states with bounded relative Hamming weight. From Corollary 1, we can then conclude that with respect to the min-entropy of the measurement outcome, the state of A behaves similarly to the case where Bob honestly prepares A to be in state $|\varphi_A^\circ\rangle$. By Remark 2, i.e., by doing a suitable change of basis, the same holds if $|\varphi_A^\circ\rangle = H^{\hat{\theta}}|\hat{\mathbf{x}}\rangle$ for arbitrary fixed $\hat{\theta}, \hat{\mathbf{x}} \in \{0, 1\}^n$, where $\text{wt}(\theta)$ is replaced by the Hamming distance between θ and $\hat{\theta}$. We will make use of this in the applications in the upcoming sections.

5 Application I: Quantum Oblivious Transfer (QOT)

5.1 The Bennett *et al.* QOT Scheme

In a (one-out-of-two) *oblivious transfer*, OT for short, Alice sends two messages, $\mathbf{m}_0, \mathbf{m}_1 \in \{0, 1\}^\ell$ to Bob. Bob may choose to receive one of the two message, \mathbf{m}_c . The security requirements demand that Bob learns no information on the other message, \mathbf{m}_{1-c} , while at the same time Alice remains ignorant about Bob's choice bit c .

Back in 1992, Bennett *et al.* proposed a quantum scheme for OT, i.e., a QOT scheme [BBCS92]. The scheme makes use of a *bit commitment* (BC), which at that point in time was believed to be implementable with unconditional security by a quantum scheme. Bennett *et al.*, however, merely claimed security of their scheme without providing any proof. In 1994, Mayers and Salvail proved the QOT scheme secure against a limited class of attacks [MS94], and, subsequently, Yao presented a full security proof without limiting the adversary's capabilities [Yao95]. However, Yao's proof is lengthy and very technical, and thus hard to understand. Furthermore, security is phrased and proven in terms of *accessible information*, of which we now know that it is a too weak information measure to guarantee security as required.

Here we show how our sampling-strategy framework naturally leads to a new security proof for Bennett *et al.*'s QOT scheme. The new proof is simple and conceptually easy-to-understand, and security is expressed and proven by means of a security definition that is currently accepted to be "the right one".

¹⁰There exists a corresponding upper bound for the cardinality of a q -ary Hamming ball (with arbitrary q), expressed in terms of the so-called q -ary entropy function; we do not elaborate on this here, since we now focus on the binary case.

Furthermore, it allows for an explicit bound on the imperfection of the scheme for any set of parameters (number of transmitted qubits, length of messages etc.), rather than merely providing an asymptotic security claim. Nowadays, we of course know that BC (as well as QOT) cannot be implemented with unconditional security by means of a quantum scheme: QBC is impossible [May97, LC97]. As such QOT cannot be instantiated from scratch. Nevertheless, the existence of a QOT scheme based on a (hypothetical) BC is still an interesting result, since in the non-quantum world, a BC alone does *not* allow to implement OT.

Below, we describe Bennett *et al.*'s QOT scheme (with some minor modifications), which we denote as QOT. Actually, QOT corresponds to the *randomized* OT used within Bennett *et al.*'s QOT scheme, where the messages m_0 and m_1 , called k_0 and k_1 in QOT, are not *input* by Alice (her input is empty: \perp) but randomly produced during the course of the scheme and then *output* to Alice. The desired non-randomized OT is then obtained simply by one-time-pad encrypting Alice's input messages m_0 and m_1 with the keys k_0 and k_1 , respectively. Security of the non-randomized OT follows immediately from the security of the randomized OT by the properties of the one-time-pad.

QOT is parametrized by parameters $n, k, \ell \in \mathbb{N}$, where n is the number of qubits communicated, ℓ the bit-length of the messages/keys k_0, k_1 , and k is the size of the "test set" t , which we require to be at most $n/2$. QOT makes use of a universal hash function $g : \mathcal{R} \times \{0, 1\}^n \rightarrow \{0, 1\}^\ell$. For $x' \in \{0, 1\}^{n'}$ with $n' < n$, we define $g(r, x')$ as $g(r, x)$ where $x \in \{0, 1\}^n$ is obtained from x' by padding it with sufficiently many 0's. Furthermore, the scheme makes use of a BC, which we assume to be an ideal BC functionality. Alternatively, at the cost of losing unconditional security against dishonest Alice, we may use a BC implementation that is unconditionally binding and computationally hiding.¹¹ Finally, for simplicity, we assume a *noise-free* quantum channel. For the more realistic setting of noisy quantum communication, an error-correcting code can be applied in a similar fashion as in the original scheme; this will not significantly affect our proof. In the upcoming protocol¹² descriptions, we make use of our convention to speak about a basis θ (or $\hat{\theta}$) in $\{0, 1\}^n$ when we actually mean $H^\theta\{|0\rangle, |1\rangle\}^{\otimes n}$ (respectively $H^{\hat{\theta}}\{|0\rangle, |1\rangle\}^{\otimes n}$). Protocol QOT is shown below.

Protocol QOT($\perp; c$)

1. (*Preparation*) Alice chooses $x, \theta \in \{0, 1\}^n$ at random and sends the n qubits $H^\theta|x\rangle$ to Bob. Bob selects $\hat{\theta} \in \{0, 1\}^n$ at random and measures the received qubits in basis $\hat{\theta}$, obtaining $\hat{x} \in \{0, 1\}^n$.
 2. (*Commitment*) Bob commits bit-wise to $\hat{\theta}$ and \hat{x} . Alice samples a random subset $t \subset [n]$ of cardinality k and asks Bob to open the commitments to $\hat{\theta}_i$ and \hat{x}_i for all $i \in t$. Alice verifies the openings and that $\hat{x}_i = x_i$ whenever $\hat{\theta}_i = \theta_i$, and she aborts in case of an inconsistency.
 3. (*Set partitioning*) Alice sends θ to Bob. Bob partitions \bar{t} into the subsets $I_c = \{i \in \bar{t} : \theta_i = \hat{\theta}_i\}$ and $I_{1-c} = \{i \in \bar{t} : \theta_i \neq \hat{\theta}_i\}$ and sends I_0 and I_1 to Alice.
 4. (*Key extraction*) Alice chooses and sends to Bob a random $r \in \mathcal{R}$, and computes $k_0 := g(r, x_{I_0})$ and $k_1 := g(r, x_{I_1})$. Bob computes $\hat{k}_c = g(r, \hat{x}_{I_c})$.
-

¹¹In case of an unconditionally hiding and computationally binding BC scheme, our techniques do not apply directly; how to handle this case is shown in [DFL⁺09].

¹²A *protocol* is an interactive algorithm between two (or in general more) entities, whereas a *scheme* in general may consist of several protocols (like for BC); since the cryptographic tasks considered in this article (QOT and QKD) ask for a single protocol, we use the terms *protocol* and *scheme* interchangeably.

It is trivial to see that for honest Alice and Bob: $\hat{k}_c = k_c$. Furthermore, security against dishonest Alice, who is trying to learn information on c , is easy to see and not the issue here: in case of a perfect BC functionality, Alice learns no information on c no matter what she does; in case of a computationally hiding BC implementation, all information she obtains on c is “hidden within the commitments”, and thus computational security follows from the computational hiding property.

Proving security against dishonest Bob is much more subtle, and is the goal of this section. Clearly, if Bob indeed measures the qubits in the preparation phase with respect to some choice $\hat{\theta}$, then security is easy to see: no matter how he partitions \bar{t} into I_0 and I_1 , on at least one of x_{I_0} and x_{I_1} he has some lower bounded uncertainty, and privacy amplification finishes the job. The intuition is now that the commitment phase forces Bob to essentially measure all qubits with respect to some choice $\hat{\theta}$, as otherwise he will get caught. However, proving this rigorously is non-trivial.

5.2 The Security Proof

For our proof of security against dishonest Bob, we first introduce a slightly modified version of the protocol, QOT*, given below. QOT* is only of proof-technical interest because it asks Alice to perform some actions that she could not do in practice. However, her actions are well-defined, and it follows from standard arguments that Bob’s view of QOT is exactly the same as of QOT*. It thus suffices to prove security (against dishonest Bob) for QOT*.

QOT* is obtained from QOT by means of the following two modifications. First, for every $i \in [n]$, instead of sending $H^{\theta_i}|x_i\rangle$, Alice prepares an EPR pair A_iB_i of which she sends B_i to Bob and measures A_i , at some later point in the protocol, in basis θ_i to obtain x_i . By elementary properties of EPR pairs, and since actions on different subsystems commute, this does not affect Bob’s view of the protocol. Second, Alice measures her qubits A_t within the test subset t in Bob’s basis $\hat{\theta}_t$ (rather than in θ_t) to obtain x_t , but she still only verifies correctness of Bob’s \hat{x}_i ’s with $i \in t$ for which $\hat{\theta}_i = \theta_i$. Note that by assumption on the BC, the string $\hat{\theta}$ to which Bob can open his commitments is uniquely determined at this point, and thus Alice’s action is well-defined, although not doable in real life. This modification only influences Alice’s bits x_i for which $i \in t$ and $\hat{\theta}_i \neq \theta_i$; however, since these bits are not used in the protocol, it has no effect on Bob’s view.

Protocol QOT*($\perp; c$)

1. (*Preparation*) Alice prepares n EPR pairs of the form $(|0\rangle|0\rangle + |1\rangle|1\rangle)/\sqrt{2}$, and sends one qubit of each pair to Bob, who proceeds as in the original scheme QOT to obtain $\hat{\theta}$ and \hat{x} . Alice chooses a random $\theta \in \{0, 1\}^n$, but she does not measure her qubits yet.
 2. (*Commitment*) Bob commits to $\hat{\theta}$ and \hat{x} , and Alice chooses a random subset $t \subset [n]$ of cardinality k , as in QOT. Next, Alice measures her qubits that are indexed by t in Bob’s basis $\hat{\theta}_t$ to obtain x_t . Then, Alice sends t to Bob and they proceed as in QOT, meaning that Alice verifies that $\hat{x}_i = x_i$ for $i \in t$ only when $\hat{\theta}_i = \theta_i$.
 3. (*Set partitioning*) As in QOT. Additionally, Alice measures her qubits corresponding to I_0 in basis θ_{I_0} to obtain x_{I_0} and her qubits corresponding to I_1 in basis θ_{I_1} to obtain x_{I_1} .
 4. (*Key extraction*) Exactly as in the original scheme QOT.
-

Our proof for the security of QOT*, and thus of QOT, against dishonest Bob follows quite easily from

our treatment of sampling strategies from Section 4. The proof is given below, after the formal security statement in Theorem 4. We would like to point out that our security guarantee implies the security definition proposed and studied in [FS08] for (randomized) OT, which in particular implies sequential *composability* when used as a sub-routine in a classical outer protocol.

Theorem 4 (Security of QOT). *Consider an execution of QOT (respectively QOT*) between honest Alice and dishonest Bob. Let \mathbf{K}_0 and \mathbf{K}_1 be the keys in $\{0, 1\}^\ell$ output by Alice. Then, there exists a bit c so that \mathbf{K}_{1-c} is close to random-and-independent of Bob's view (given \mathbf{K}_c) in that for any $\epsilon, \delta > 0$:*

$$\begin{aligned} & \Delta(\rho_{\mathbf{K}_{1-c}\mathbf{K}_cE}, \frac{1}{2^\ell} \mathbb{I} \otimes \rho_{\mathbf{K}_cE}) \\ & \leq \frac{1}{2} \cdot 2^{-\frac{1}{2} \left(\left(\frac{1}{4} - \frac{\epsilon}{2} - h(\delta) \right) (n-k) - \ell \right)} + \sqrt{6} \exp(-\delta^2 k / 100) + 2 \exp(-2\epsilon^2 (n-k)), \end{aligned}$$

where E denotes the quantum state output by Bob, and \mathbb{I} the identity operator on \mathbb{C}^{2^ℓ} .

On a high level, the proof is as follows. Alice's checking procedure can be understood as applying a sampling strategy to the qubits she holds. From this we obtain that (except with a small error) the joint state she shares with Bob is a superposition of states with small relative Hamming weight within her subsystem $A_{\bar{t}}$. This implies that the joint state is a superposition of states with small relative Hamming weight also within $A_{I_{1-c}}$, where $c \in \{0, 1\}$ is chosen such that $\theta_i \neq \hat{\theta}_i$ for approximately half (or more) of the indices i in I_{1-c} . It then follows from Corollary 1 that $\mathbf{x}_{I_{1-c}}$, obtained by measuring $A_{I_{1-c}}$ in basis $\boldsymbol{\theta}_{I_{1-c}}$, has high min-entropy, so that privacy amplification concludes the proof. The formal proof, which takes care of the details and keeps track of the error term, is given below.

Proof. We consider the state

$$|\varphi_{AE_o}\rangle \in \mathcal{H}_{A_1} \otimes \cdots \otimes \mathcal{H}_{A_n} \otimes \mathcal{H}_{E_o},$$

shared between Alice and Bob, after Bob has committed to $\hat{\boldsymbol{\theta}}$ and $\hat{\mathbf{x}}$, but before Alice chooses the test subset t . $|\varphi_{AE_o}\rangle$ is obtained from the n EPR-pairs by an arbitrary quantum operation (possibly involving measurements), applied only to Bob's part. Without loss of generality, we may assume that, given the commitments, the joint state is indeed pure. Furthermore, we consider the strings $\hat{\boldsymbol{\theta}}$ and $\hat{\mathbf{x}}$, to which Bob has committed. By the unconditional binding property, these are uniquely determined. For concreteness, and in order to have the notation fit nicely with Section 4, we assume $\hat{\boldsymbol{\theta}} = \hat{\mathbf{x}} = (0, \dots, 0) \in \{0, 1\}^n$; however, by Remark 2, the very same reasoning works for any $\hat{\boldsymbol{\theta}}$ and $\hat{\mathbf{x}}$.

The crucial observation now is that Alice's checking procedure within the commitment phase of QOT* can be understood as applying a sampling strategy to the state $|\varphi_{AE_o}\rangle$ in order to test closeness of A to the all-zero state $|0\rangle \cdots |0\rangle$. Indeed, Alice chooses a random subset $t \subset [n]$ of cardinality k , measures A_t (in the computational basis) to obtain \mathbf{x}_t , and decides whether to accept or reject based on \mathbf{x}_t ; specifically, she takes a random subset $s \subseteq t$, given by $s = \{i \in t : \theta_i = \hat{\theta}_i\}$, and accepts if and only if $x_s = 0$ for all $i \in s$. This is precisely the sampling strategy Ψ studied in Example 4, adapted to test closeness to $|0\rangle \cdots |0\rangle$ by accepting if and only if $f(t, \mathbf{x}_t, s) = 0$. Note that, by the random choices of the θ_i 's, s is indeed a random subset of t .

Thus, we can conclude that at the end of the commitment phase, for any fixed $\delta > 0$, the joint state of $A_{\bar{t}}E_o$ has collapsed to a state $|\psi_{A_{\bar{t}}E_o}\rangle$ that is (on average over Alice's choice of t and s) $\epsilon_{\text{quant}}^\delta$ -close to being a superposition of states with relative Hamming weight at most δ within $A_{\bar{t}}$ (or else Alice has aborted). We proceed by assuming that the state $|\psi_{A_{\bar{t}}E_o}\rangle$ equals a superposition of states with small

relative Hamming weight, and we book-keep the error $\varepsilon_{\text{quant}}^\delta$.¹³ Recall that by Theorem 3 and Example 4 (and its analysis in Appendix A.4),

$$\varepsilon_{\text{quant}}^\delta \leq \sqrt{\varepsilon_{\text{class}}^\delta} \leq \sqrt{6} \exp(-k\delta^2/100).$$

By the random choices of the θ_i 's, it follows from Hoeffding's inequality (Theorem 1) that the Hamming weight of $\theta_{\bar{t}}$ is lower bounded by $\text{wt}(\theta_{\bar{t}}) \geq (\frac{1}{2} - \epsilon)(n - k)$ except with probability at most $2 \exp(-2\epsilon^2(n - k))$.¹⁴ In the sequel, we assume that the bound holds, and we book-keep the error. It follows that regardless of how Bob divides \bar{t} into I_0 and I_1 , there exists $c \in \{0, 1\}$ such that $\text{wt}(\theta_{I_{1-c}}) \geq \frac{1}{2}(\frac{1}{2} - \epsilon)(n - k)$ (if Bob is honest, then c coincides with his input bit).

By re-arranging Alice's qubits, we write the state $|\psi_{A_{\bar{t}}E_0}\rangle$ as $|\psi_{A^{1-c}A^cE_0}\rangle$, where $A^0 := A_{I_0}$ and $A^1 := A_{I_1}$. Since $|\psi_{A_{\bar{t}}E_0}\rangle$ is a superposition of states with Hamming weight at most $(n - k)\delta$ within $A_{\bar{t}}$, it is easy to see that $|\psi_{A^{1-c}A^cE_0}\rangle$ is a superposition of states with Hamming weight at most $(n - k)\delta$ within A^{1-c} . Let the random variables \mathbf{X}_{1-c} and \mathbf{X}_c describe the outcome of measuring A^{1-c} and A^c in bases $\theta_{I_{1-c}}$ and θ_{I_c} , respectively, and let $\rho_{\mathbf{X}_{1-c}\mathbf{X}_cE_0}$ be the corresponding hybrid state. We may think of $\rho_{\mathbf{X}_{1-c}\mathbf{X}_cE_0}$ being obtained by *first* measuring A^{1-c} , resulting in a hybrid state $\rho_{\mathbf{X}_{1-c}A^cE_0}$, and *then* measuring A^c ; indeed, the order in which these measurements take place have no effect on the final state.

We can now apply Corollary 1 to the hybrid state $\rho_{\mathbf{X}_{1-c}A^cE_0}$ obtained from measuring subsystem A^{1-c} within $|\psi_{A^{1-c}A^cE_0}\rangle$ and conclude that

$$H_{\min}(\mathbf{X}_{1-c}|A^cE_0) \geq \text{wt}(\theta_{I_{1-c}}) - h(\delta) \log(|I_{1-c}|) \geq \left(\frac{1}{4} - \frac{\epsilon}{2} - h(\delta)\right)(n - k).$$

By a basic property of the min-entropy ("measuring only destroys information"), it follows that the same bound in particular holds for $H_{\min}(\mathbf{X}_{1-c}|\mathbf{X}_cE_0)$. Applying privacy amplification (Theorem 2), incorporating the error-probabilities (expressed in terms of trace distance) obtained along the proof, and noting that Bob's processing of his information to obtain his final quantum state E does not increase the trace-distance, concludes the proof. \square

6 Application II: Quantum Key Distribution (QKD)

In quantum key distribution (QKD), Alice and Bob want to agree on a secret key in the presence of an adversary Eve. Alice and Bob are assumed to be able to communicate over a quantum channel and over an authenticated classical channel.¹⁵ Eve may eavesdrop the classical channel (but not insert or modify messages), and she has full control over the quantum channel. The first and still most prominent QKD scheme is the famous BB84 QKD scheme due to Bennett and Brassard [BB84].

In this section, we show how our sampling-strategy framework leads to a simple security proof for the BB84 QKD scheme. Proving QKD schemes rigorously secure is a highly non-trivial task, and as such our new proof nicely demonstrates the power of the sampling-strategy framework. Furthermore, our new proof has some nice features. For instance, it allows us to explicitly state (a bound on) the error

¹³It now follows immediately from Corollary 1 that $H_{\min}(X_0X_1|E_0)$ is "large", where X_0 collects the bits obtained by measuring A_{I_0} in basis θ_{I_0} , and correspondingly for X_1 . However, in the end we need that $H_{\min}(X_{1-c}|X_cE_0)$ is "large" for some c , which does *not* follow from the former. Because of that, we need to make a small detour.

¹⁴Actually, for the one-sided bound, we could save the factor two in front of the exp.

¹⁵If the classical channel between Alice and Bob is not authentic, then authenticity of the communication can still be achieved by information-theoretic authentication techniques, at the cost of requiring Alice and Bob to initially share a short secret key.

probability of the QKD scheme for any given choices of the parameters. Additionally, our proof does not seem to take unnecessary detours or to make use of “loose bounds”, and therefore we feel that the bound on the error probability we obtain is rather tight (although we have no formal argument to support this).

Our proof strategy can also be applied to other QKD schemes that are based on the BB84 encoding. For example, Lo *et al.*'s QKD scheme¹⁶ [LCA05] can be proven secure by following exactly our proof, except that one needs to analyze a slightly different sampling strategy, namely the one from Example 6. On the other hand, it is yet unknown whether our framework can be used to prove e.g. the six-state QKD protocol [Bru98] secure.

Actually, the QKD scheme we analyze is the entanglement-based version of the BB84 scheme (as initially suggested by Ekert [Eke91]). However, it is very well known and not too hard to show that security of the entanglement-based version implies security of the original BB84 QKD scheme.

The entanglement-based QKD scheme, QKD, is parametrized by the total number n of qubits sent in the protocol and the number k of qubits used to estimate the error rate of the quantum channel (where we require $k \leq n/2$). Additional parameters, which are determined during the course of the protocol, are the observed error rate β and the number $\ell \in \mathbb{N} \cup \{0\}$ of extracted key bits. QKD makes use of a universal hash function $g : \mathcal{R} \times \{0, 1\}^{n-k} \rightarrow \{0, 1\}^\ell$ and a linear binary error correcting code of length $n - k$ that allows to correct up to a β' -fraction of errors (except maybe with negligible probability) for some $\beta' > \beta$. The choice of how much β' exceeds β is a trade-off between keeping the probability that Alice and Bob end up with different keys small and increasing the size of the extractable key. We will write m for the bit size of the syndrome of this error-correcting code. Protocol QKD can be found below.

Protocol QKD

1. (*Qubit distribution*) Alice prepares n EPR pairs of the form $(|0\rangle|0\rangle + |1\rangle|1\rangle)/\sqrt{2}$, and sends one qubit of each pair to Bob, who confirms the receipt of the qubits. Then, Alice picks random $\theta \in \{0, 1\}^n$ and sends it to Bob, and Alice and Bob measure their respective qubits in basis θ to obtain \mathbf{x} on Alice's side respectively \mathbf{y} on Bob's side.
 2. (*Error estimation*) Alice chooses a random subset $s \subset [n]$ of size k and sends it to Bob. Then, Alice and Bob exchange \mathbf{x}_s and \mathbf{y}_s and compute $\beta := \omega(\mathbf{x}_s \oplus \mathbf{y}_s)$.
 3. (*Error correction*) Alice sends the syndrome syn of $\mathbf{x}_{\bar{s}}$ to Bob with respect to a suitable linear error correcting code (as described above). Bob uses syn to correct the errors in $\mathbf{y}_{\bar{s}}$ and obtains $\hat{\mathbf{x}}_{\bar{s}}$. Let m be the bit-size of syn .
 4. (*Key distillation*) Alice chooses a random seed r for a universal hash function g with range $\{0, 1\}^\ell$, where ℓ satisfies $\ell < (1 - h(\beta))n - k - m$ (or $\ell = 0$ if the right-hand side is not positive), and sends it to Bob. Then, Alice and Bob compute $\mathbf{k} := g(r, \mathbf{x}_{\bar{s}})$ and $\hat{\mathbf{k}} := g(r, \hat{\mathbf{x}}_{\bar{s}})$, respectively.
-

It is not hard to see that $\mathbf{k} = \hat{\mathbf{k}}$ except with negligible probability (in n). Furthermore, if no Eve interacts with the quantum communication in the qubit distribution phase then $\mathbf{x} = \mathbf{y}$ in case of a noise-free quantum channel, or more generally, $\omega(\mathbf{x} - \mathbf{y}) \approx \phi$ in case the quantum channel is noisy and introduces an error probability $0 \leq \phi < \frac{1}{2}$. It follows that $\beta \approx \phi$, so that using an error correcting code that approaches the Shannon bound, Alice and Bob can extract close to $(1 - 2h(\phi))(n - k)$ bits of

¹⁶In this scheme, Alice and Bob bias the choice of the bases so that they measure a bigger fraction of the qubits in the same basis.

secret key, which is positive for ϕ smaller than approximately 11%. The difficult part is to prove security against an active adversary Eve. We first state the formal security claim.

Note that we cannot expect that Eve has (nearly) no information on \mathbf{K} , i.e. that $\Delta(\rho_{\mathbf{K}E}, \frac{1}{|\mathcal{K}|} \mathbb{I}_{\mathbf{K}} \otimes \rho_E)$ is small, since the bit-length ℓ of \mathbf{K} is not fixed but depends on the course of the protocol, and Eve can influence and thus obtain information on ℓ (and thus on \mathbf{K}). Theorem 5 though guarantees that the bit-length ℓ is the *only* information Eve learns on \mathbf{K} , in other words, \mathbf{K} is essentially random-and-independent of E when given ℓ .

Theorem 5 (Security of QKD). *Consider an execution of QKD in the presence of an adversary Eve. Let \mathbf{K} be the key obtained by Alice, and let E be Eve's quantum system at the end of the protocol. Let $\tilde{\mathbf{K}}$ be chosen uniformly at random of the same bit-length as \mathbf{K} . Then, for any δ with $\beta + \delta \leq \frac{1}{2}$:*

$$\Delta(\rho_{\mathbf{K}E}, \rho_{\tilde{\mathbf{K}}E}) \leq \frac{1}{2} \cdot 2^{-\frac{1}{2}((1-h(\beta+\delta))^{n-k-m-\ell})} + 2 \exp(-\frac{1}{6}\delta^2 k).$$

From an application point of view, the following question is of interest. Given the parameters n and k , and given a course of the protocol with observed error rate β and where an error-correcting code with syndrome length m was used, what is the maximal size ℓ of the extractable key \mathbf{K} if we want $\Delta(\rho_{\mathbf{K}E}, \rho_{\tilde{\mathbf{K}}E}) \leq \epsilon$ for a given ϵ ? From the bound in Theorem 5, it follows that for every choice of δ (with $\beta + \delta \leq \frac{1}{2}$), one can easily compute a possible value for ℓ simply by solving for ℓ . In order to compute the optimal value, one needs to maximize ℓ over the choice of δ .

The formal proof of Theorem 5 is given below. Informally, the argument goes as follows. The error estimation phase can be understood as applying a sampling strategy. From this, we can conclude that the state from which the raw key, $\mathbf{x}_{\bar{s}}$, is obtained, is a superposition of states with bounded Hamming weight, so that Corollary 1 guarantees a certain amount of min-entropy within $\mathbf{x}_{\bar{s}}$. Privacy amplification then finishes the proof.

To indeed be able to model the error estimation procedure as a sampling strategy, we will need to consider a modified but *equivalent* way for Alice and Bob to jointly obtain \mathbf{x}_s and \mathbf{y}_s from the initial joint state, which will allow them to obtain the XOR-sum $\mathbf{x}_s \oplus \mathbf{y}_s$, and thus to compute β , *before* they measure the remaining part of the state, whose outcome then determines $\mathbf{x}_{\bar{s}}$. This modification is based on the so-called CNOT operation, U_{CNOT} , acting on $\mathbb{C}^2 \otimes \mathbb{C}^2$, and its properties that

$$U_{\text{CNOT}}(|b\rangle|c\rangle) = |b\rangle|b \oplus c\rangle \quad \text{and} \quad U_{\text{CNOT}}(H|b\rangle H|c\rangle) = H|b \oplus c\rangle H|c\rangle, \quad (2)$$

where the first holds by definition of U_{CNOT} , and the second is straightforward to verify.

Proof. Throughout the proof, we use capital letters, Θ , \mathbf{X} etc. for the *random variables* representing the corresponding choices of θ , \mathbf{x} etc. in protocol QKD. Let the state, shared by Alice, Bob and Eve right after the quantum communication in the qubit distribution phase, be denoted by $|\psi_{ABE_{\circ}}\rangle$; ¹⁷ without loss of generality, we may indeed assume the shared state to be pure. For every $i \in [n]$, Alice and Bob then measure the respective qubits A_i and B_i from $|\psi_{ABE_{\circ}}\rangle$ in basis Θ_i , obtaining X_i and Y_i . This results in the hybrid state $\rho_{\Theta \mathbf{X} \mathbf{Y} E_{\circ}}$. For the proof, it will be convenient to introduce the additional random variables $\mathbf{W} = (W_1, \dots, W_n)$ and $\mathbf{Z} = (Z_1, \dots, Z_n)$, defined by

$$Z_i := X_i \oplus Y_i \quad \text{and} \quad W_i := \begin{cases} X_i & \text{if } \Theta_i = 0 \\ Y_i & \text{if } \Theta_i = 1 \end{cases}. \quad (3)$$

¹⁷Note that E_{\circ} represents Eve's quantum state just after the quantum communication stage, whereas E represents Eve's entire state of knowledge at the end of the protocol (i.e., the quantum information and all classical information gathered during execution of QKD).

Note that, when given Θ , the random variables \mathbf{W} and \mathbf{Z} are uniquely determined by \mathbf{X} and \mathbf{Y} *and vice versa*, and thus we may equivalently analyze the hybrid state $\rho_{\Theta \mathbf{W} \mathbf{Z} E_o}$.

For the analysis, we will consider a slightly *different* experiment for Alice and Bob to obtain the very *same* state $\rho_{\Theta \mathbf{W} \mathbf{Z} E_o}$; the advantage of the modified experiment is that it can be understood as a sampling strategy. The modified experiment is as follows. First, the CNOT transformation is applied to every qubit pair $A_i B_i$ within $|\psi_{ABE_o}\rangle$ for $i \in [n]$, such that the state $|\varphi_{ABE_o}\rangle = (U_{\text{CNOT}}^{\otimes n} \otimes \mathbb{I}_{E_o})|\psi_{ABE_o}\rangle$ is obtained. Next, Θ is chosen at random as in the original scheme, and for every $i \in [n]$ the qubit pair $A_i B_i$ of the transformed state is measured as in the original scheme depending on Θ_i ; however, if $\Theta_i = 0$ then the resulting bits are denoted by W_i and Z_i , respectively, and if $\Theta_i = 1$ then they are denoted by Z_i and W_i , respectively, such that which bit is assigned to which variable depends on Θ_i . This is illustrated in Figure 1 (left and middle), where light and dark colored ovals represent measurements in the computational and Hadamard basis, respectively. It now follows immediately from the properties (2) of the CNOT transformation and from the relation (3) between \mathbf{X} , \mathbf{Y} and \mathbf{W} , \mathbf{Z} that the state $\rho_{\Theta \mathbf{W} \mathbf{Z} E_o}$ (or, equivalently, $\rho_{\Theta \mathbf{X} \mathbf{Y} E_o}$) obtained in this modified experiment is exactly the same as in the original.




Figure 1: Original and modified experiments for obtaining the same state $\rho_{\Theta \mathbf{W} \mathbf{Z} E_o}$.

An additional modification we may do without influencing the final state is to *delay* some of the measurements: we assume that first the qubits are measured that lead to the Z_i 's, and only at some later point, namely after the *error estimation* phase, the qubits leading to the W_i 's are measured (as illustrated in Figure 1, right). This can be done since the relative Hamming weight of $X_S \oplus Y_S$ for a random subset $S \subset [n]$ (of size k) can be computed given \mathbf{Z} alone.

The crucial observation is now that this modified experiment can be viewed as a particular sampling strategy Ψ , as a matter of fact as the sampling strategy discussed in Example 5, being applied to systems A and B of the state $|\varphi_{ABE_o}\rangle$. Indeed: first, a subset of the $2n$ qubit positions is selected according to some probability distribution, namely of each pair $A_i B_i$ one qubit is selected at random (determined by Θ_i). Then, the selected qubits are measured to obtain the bit string $\mathbf{Z} = (Z_1, \dots, Z_n)$. And, finally, a value β is computed as a (randomized) function of \mathbf{Z} : $\beta = \omega(\mathbf{Z}_S)$ for a random $S \subset [n]$ of size k . We point out that here the reference basis (as explained in Remark 2) is not the computational basis for all qubits, but is the Hadamard basis on the qubits in system A and the computational basis in system B ; however, as discussed in Remark 2, we may still apply the results from Section 4 (appropriately adapted).

It thus follows that for any fixed $\delta > 0$, the remaining state, from which \mathbf{W} is then obtained, is (on average over Θ and S) $\varepsilon_{\text{quant}}^\delta$ -close to a state which is (for any possible values for Θ , \mathbf{Z} and S) a superposition of states with relative Hamming weight in a δ -neighborhood of β . Note that the latter has to be understood with respect to the fixed reference basis (i.e., the Hadamard basis on A and the computational basis on B). In the following, we assume that the remaining state *equals* such a superposition, but we

remember the error

$$\varepsilon_{\text{quant}}^{\delta} \leq \sqrt{\varepsilon_{\text{class}}^{\delta}} \leq 2 \exp\left(-\frac{1}{6}\delta^2 k\right).$$

where the bound on $\varepsilon_{\text{class}}^{\delta}$ is derived in Appendix A.5.

Recall that \mathbf{W} is now obtained by measuring the remaining qubits; however, the basis used is opposite to the reference basis, namely the computational basis on the qubits A_i and the Hadamard basis on the qubits B_i . Hence, by Corollary 1 (and the subsequent discussion) we get a lower bound on the min-entropy of \mathbf{W} :

$$H_{\min}(\mathbf{W}|\Theta\mathbf{Z}SE_o) \geq (1 - h(\beta + \delta))n.$$

Since \mathbf{W} is uniquely determined by \mathbf{X} (and vice versa) when given Θ and \mathbf{Z} , the same lower bound also holds for $H_{\min}(\mathbf{X}|\Theta\mathbf{Z}SE_o)$. Note that in QKD, the k qubit-pairs that are used for estimating β are not used anymore in the key distillation phase, so we are actually interested in the min-entropy of $\mathbf{X}_{\bar{S}}$. Additionally, we should take into account that Alice sends an m -bit syndrome SYN during the error correction phase. Hence, by using the chain rule, we obtain

$$H_{\min}(\mathbf{X}_{\bar{S}}|\Theta\mathbf{Z}\mathbf{X}_S SYN E_o) \geq (1 - h(\beta + \delta))n - k - m.^{18}$$

Finally, we apply privacy amplification (Theorem 2) which concludes the proof. \square

7 Conclusion

We have shown a framework for predicting some property (namely the approximate Hamming weight, appropriately defined) of a population of quantum states, by measuring a small sample subset. The framework allows for new and simple security proofs for important quantum cryptographic protocols: the Bennett *et al.* QOT and the BB84 QKD scheme. We find it particularly interesting that with our framework, the protocols for QOT and QKD can be proven secure by means of very similar techniques, even though they implement fundamentally different cryptographic primitives, and are intuitively secure due to very different reasons (namely in QOT the commitments force Bob to measure the communicated qubits, whereas in QKD Eve disturbs the communicated qubits when trying to observe them).

References

- [BB84] Charles H. Bennett and Gilles 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.
- [BBCS92] Charles H. Bennett, Gilles Brassard, Claude Crépeau, and Marie-Hélène Skubiszewska. Practical quantum oblivious transfer. In *CRYPTO '91: Proceedings of the 11th Annual International Cryptology Conference on Advances in Cryptology*, pages 351–366, London, UK, 1992. Springer-Verlag.

¹⁸Probably, it is possible to prove the lower bound: $(1 - h(\beta + \delta))(n - k) - m$ using a different sampling strategy. However, for that case the error probability of the related classical sampling strategy becomes harder to analyze. We have chosen for the current proof strategy and bound for the sake of simplicity.

- [Bru98] Dagmar Bruss. Optimal eavesdropping in quantum cryptography with six states. *Physical Review Letters*, 81:3018, 1998. <http://arxiv.org/abs/quant-ph/9805019>.
- [DFL⁺09] Ivan Damgård, Serge Fehr, Carolin Lunemann, Louis Salvail, and Christian Schaffner. Improving the security of quantum protocols, 2009. <http://arxiv.org/abs/0902.3918>.
- [Eke91] Artur K. Ekert. Quantum cryptography based on bell’s theorem. *Physical Review Letter*, 67(6):661–663, August 1991.
- [FS08] Serge Fehr and Christian Schaffner. Composing quantum protocols in a classical environment. <http://arxiv.org/abs/0804.1059>, 2008.
- [Hoe63] W. Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963.
- [LC97] H.-K. Lo and H. F. Chau. Is quantum bit commitment really possible? *Physical Review Letters*, 78:3410–3413, April 1997.
- [LCA05] H.-K. Lo, H. F. Chau, and M. Ardehali. Efficient quantum key distribution scheme and a proof of its unconditional security. *J. Cryptol.*, 18(2):133–165, 2005.
- [May97] Dominic Mayers. Unconditionally secure quantum bit commitment is impossible. *Physical Review Letters*, 78(17):3414–3417, April 1997.
- [MS94] Dominic Mayers and Louis Salvail. Quantum oblivious transfer is secure against all individual measurements. In *Proceedings of the Third Workshop on Physics and Computation — PhysComp ’94*, pages 69–77. IEEE Computer Society Press, 1994.
- [Ren05] Renato Renner. *Security of Quantum Key Distribution*. PhD thesis, ETH Zürich (Switzerland), September 2005. <http://arxiv.org/abs/quant-ph/0512258>.
- [RK05] Renato Renner and Robert 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.
- [Ser74] R. J. Serfling. Probability inequalities for the sum in sampling without replacement. *The Annals of Statistics*, 2(1):39–48, 1974.
- [Yao95] Andrew Chi-chih Yao. Security of quantum protocols against coherent measurements. In *Proceedings of 26th Annual ACM Symposium on the Theory of Computing*, pages 67–75, 1995.

A Error Probabilities of the Example Sampling Strategies

A.1 Example 1 — Random sampling *without* replacement

It follows immediately from Theorem 1 that the estimate is δ -close to the relative Hamming weight $\omega(\mathbf{q})$ of \mathbf{q} except with probability at most $2 \exp(-2\delta^2 k)$. However, we want to analyze closeness of the

estimate to $\omega(\mathbf{q}_{\bar{T}})$ (still treating T as a random variable). This can be derived easily as follows. We can write $\omega(\mathbf{q}) = \alpha\omega(\mathbf{q}_T) + (1 - \alpha)\omega(\mathbf{q}_{\bar{T}})$, where $\alpha := k/n$, and thus can see that

$$\omega(\mathbf{q}_{\bar{T}}) - \omega(\mathbf{q}_T) = \frac{1}{1 - \alpha} \left(\omega(\mathbf{q}) - \alpha\omega(\mathbf{q}_T) \right) - \omega(\mathbf{q}_T) = \frac{1}{1 - \alpha} \left(\omega(\mathbf{q}) - \omega(\mathbf{q}_T) \right)$$

so that

$$\begin{aligned} \varepsilon_{\text{class}}^\delta &= \max_{\mathbf{q}} \Pr \left[\mathbf{q} \notin B_{T,S}^\delta \right] = \max_{\mathbf{q}} \Pr \left[|\omega(\mathbf{q}_{\bar{T}}) - \omega(\mathbf{q}_T)| \geq \delta \right] \\ &= \max_{\mathbf{q}} \Pr \left[|\omega(\mathbf{q}) - \omega(\mathbf{q}_T)| \geq (1 - \alpha)\delta \right] \leq 2 \exp(-2(1 - \alpha)^2 \delta^2 k). \end{aligned} \quad (4)$$

Under assumption of $k \leq n/2$, we obtain a simple bound for the latter expression,

$$\varepsilon_{\text{class}}^\delta \leq 2 \exp(-2(1 - \alpha)^2 \delta^2 k) \leq 2 \exp(-\frac{1}{2} \delta^2 k). \quad (5)$$

We obtain the following bound if we use the bound from [Ser74]:

$$\begin{aligned} \varepsilon_{\text{class}}^\delta &= \max_{\mathbf{q}} \Pr \left[|\omega(\mathbf{q}) - \omega(\mathbf{q}_T)| \geq (1 - \alpha)\delta \right] \\ &\leq 2 \exp\left(-\frac{2(1 - \alpha)^2 \delta^2 k n}{n - k + 1}\right) = 2 \exp\left(-\frac{2k(n - k)^2 \delta^2}{n(n - k + 1)}\right) \leq 2 \exp\left(-\frac{\delta^2 k n}{n + 2}\right). \end{aligned}$$

for $k \leq n/2$, because $-\frac{2k(n - k)^2 \delta^2}{n(n - k + 1)}$ is convex in k , and $-\frac{\delta^2 k n}{n + 2}$ is linear in k and equality holds at $k = 0$ and $k = n/2$, hence it is a tight linear upper bound.

A.2 Example 2 — Random sampling *with* replacement

Computing the error probability for Example 2 actually turns out to be tricky. Although, as in Example 1 above, Theorem 1 applies and guarantees that the estimate is likely to be close to $\omega(\mathbf{q})$, showing that the estimate is likely to be close to $\omega(\mathbf{q}_{\bar{T}})$ seems to be non-trivial here. Since we make no further use of this example sampling strategy, we refrain from analyzing its error probability.

A.3 Example 3 — Uniformly random subset sampling

Note that for any fixed choice $k = |t|$, t is obtained as in random sampling without replacement. Because t is sampled uniformly at random, the expectation of k is given by $E[k] = n/2$. Hence, by making use of Hoeffding's inequality, we can say that for $0 < \beta < \frac{1}{2}$, $\Pr\left[\left|\frac{k}{n} - \frac{1}{2}\right| \geq \beta\right] \leq 2 \exp(-2\beta^2 n)$.

Informally, the idea is to start off with an upper bound on $\varepsilon_{\text{class}}^\delta$ obtained in Appendix A.1 (the case of sampling without replacement), and transform it into an upper bound that holds under the assumption that $k \in [(\frac{1}{2} - \beta)n, (\frac{1}{2} + \beta)n]$. Note that we cannot use the simple bound (5) from Appendix A.1, because that result was obtained under the assumption that $k \leq n/2$, and here this assumption does not hold. Instead, we use bound (4) from Appendix A.1,

$$\varepsilon_{\text{class}}^\delta \leq 2 \exp\left(-2\left(1 - \frac{k}{n}\right)^2 \delta^2 k\right) \quad (6)$$

which *does* hold for all $k \in \{0, \dots, n\}$.

To get an upper bound for (6), we replace the first occurrence of k in that expression (in the numerator of the fraction) by an upper bound for k , and the second occurrence of k by a lower bound for k .

The upper and lower bound for k are simply given by the (appropriate) boundary points of the interval $[(\frac{1}{2} - \beta)n, (\frac{1}{2} + \beta)n]$. I.e.,

$$2 \exp\left(-2n\delta^2\left(1 - \frac{(\frac{1}{2} + \beta)n}{n}\right)^2\left(\frac{1}{2} - \beta\right)\right) = 2 \exp\left(-2n\delta^2\left(\frac{1}{2} - \beta\right)^3\right)$$

To compute $\varepsilon_{\text{class}}^\delta$, we use a union bound to combine the upper bound above, which holds under assumption that k lies inside the previously defined interval, with the upper bound on the probability that k does *not* lie in this interval,

$$\varepsilon_{\text{class}}^\delta \leq 2 \exp\left(-2n\delta^2\left(\frac{1}{2} - \beta\right)^3\right) + 2 \exp(-2\beta^2n).$$

Setting $\beta = \delta/4$ in the expression above yields $-n\delta^2(2 - \delta)^3/32$ for the exponent of the first summand, and $-n\delta^2/8$ for the exponent of the second summand. Because $0 < \delta < 1$ (Definition 2), a suitable upper bound for both exponents is $-n\delta^2/32$.¹⁹ This gives the following simpler bound,

$$\varepsilon_{\text{class}}^\delta \leq 4 \exp(-n\delta^2/32).$$

A.4 Example 4 — Random sampling without replacement, using only part of the sample

From Appendix A.1, we know that $\Pr[|\omega(\mathbf{q}_{\bar{T}}) - \omega(\mathbf{q}_T)| \geq \xi] \leq 2 \exp(-\frac{1}{2}\xi^2k)$, for $k < n/2$. Additionally, the selection of the seed s and the computation of $f(t, \mathbf{q}_t, s)$ can be viewed as applying uniformly random subset sampling to \mathbf{q}_t . Hence, it follows from Appendix A.3 that $\max_{\mathbf{q}} \Pr[|\omega(\mathbf{q}_T) - \omega(\mathbf{q}_S)| \geq \gamma] \leq 4 \exp(-k\gamma^2/32)$. Setting $\delta = \xi + \gamma$, and using triangle inequality and union bound, we obtain

$$\begin{aligned} \varepsilon_{\text{class}}^\delta &= \max_{\mathbf{q}} \Pr[|\omega(\mathbf{q}_S) - \omega(\mathbf{q}_{\bar{T}})| \geq \delta] \\ &\leq \min_{0 < \xi < \delta} \left[2 \exp(-\frac{1}{2}\xi^2k) + 4 \exp(-k(\delta - \xi)^2/32) \right] \\ &\leq 6 \exp(-k\delta^2/50), \end{aligned}$$

where the last inequality follows from setting $\xi = \delta/5$ such that the two exponents coincide.

A.5 Example 5 — Pairwise one-out-of-two sampling, using only part of the sample

For $\mathcal{A} = \{0, 1\}$, a bound on the error probability $\varepsilon_{\text{class}}^\delta$ is obtained as follows. Let \mathbf{q} be arbitrary, indexed as discussed earlier. First, we show that $\omega(\mathbf{q}_{\bar{T}})$ is likely to be close to $\omega(\mathbf{q}_T)$. For this, consider the pairs (q_{i0}, q_{i1}) for which $q_{i0} \neq q_{i1}$. Let there be ℓ such pairs (where obviously $\ell \leq n$.) We denote the restrictions of \mathbf{q}_T and $\mathbf{q}_{\bar{T}}$ to these indices i with $q_{i0} \neq q_{i1}$ by $\tilde{\mathbf{q}}_T$ and $\tilde{\mathbf{q}}_{\bar{T}}$, respectively. It is easy to see that $\text{wt}(\tilde{\mathbf{q}}_T) + \text{wt}(\tilde{\mathbf{q}}_{\bar{T}}) = \ell$. It follows that for any $\epsilon > 0$ we have

$$\begin{aligned} \Pr[|\omega(\mathbf{q}_{\bar{T}}) - \omega(\mathbf{q}_T)| \geq \epsilon] &= \Pr[|\text{wt}(\mathbf{q}_T) - \text{wt}(\mathbf{q}_{\bar{T}})| \geq n\epsilon] \\ &= \Pr[|\text{wt}(\tilde{\mathbf{q}}_T) - \text{wt}(\tilde{\mathbf{q}}_{\bar{T}})| \geq n\epsilon] = \Pr[|2\text{wt}(\tilde{\mathbf{q}}_T) - \ell| \geq n\epsilon] \\ &\leq 2 \exp\left(-2\left(\frac{n\epsilon}{2\ell}\right)^2\ell\right) = 2 \exp\left(-\frac{n\epsilon^2}{2} \cdot \frac{n}{\ell}\right) \leq 2 \exp\left(-\frac{1}{2}\epsilon^2n\right), \end{aligned}$$

¹⁹Note that our goal is to find a short and simple expression, rather than finding the tightest bound.

where the third equality follows from replacing $\text{wt}(\tilde{\mathbf{q}}_{\bar{T}})$ by $\ell - \text{wt}(\tilde{\mathbf{q}}_T)$, and the first inequality follows from Hoeffding's inequality (as each entry of $\text{wt}(\tilde{\mathbf{q}}_T)$ is 0 with independent probability $\frac{1}{2}$).

Furthermore, for any $\gamma > 0$ we have the following relation involving \mathbf{q}_S :

$$\Pr[|\omega(\mathbf{q}_T) - \omega(\mathbf{q}_S)| \geq \gamma] \leq 2 \exp(-2k\gamma^2),$$

which follows from directly applying Hoeffding's inequality. Applying the union bound and letting $\delta = \epsilon + \gamma$, we obtain

$$\begin{aligned} \varepsilon_{\text{class}}^\delta &= \max_{\mathbf{q}} \Pr[|\omega(\mathbf{q}_{\bar{T}}) - \omega(\mathbf{q}_S)| \geq \delta] < 2 \min_{\epsilon \in (0, \delta)} [\exp(-\frac{1}{2}\epsilon^2 n) + \exp(-2k(\delta - \epsilon)^2)] \\ &\leq 4 \exp\left(-\frac{2kn\delta^2}{(2\sqrt{k} + \sqrt{n})^2}\right) \leq 4 \exp(-\frac{1}{3}\delta^2 k), \end{aligned}$$

where the last line follows from choosing ϵ such that the two exponents coincide, and from doing some simplifications while assuming $k \leq n/2$.

A.6 Example 6 — Pairwise *biased* one-out-of-two sampling, using only part of the sample

It will be convenient to define the index set t as the union of two subsets, $t_0 \subset [n] \times \{0\}$ and $t_1 \subset [n] \times \{1\}$. Note that the complements of these subsets should now be understood as $\bar{t}_0 = ([n] \times \{0\}) \setminus t_0$ and $\bar{t}_1 = ([n] \times \{1\}) \setminus t_1$. Let t_0 and t_1 be constructed as follows. We first sample a set $\tilde{t} \subset [n]$; for each element of $[n]$, we include it in \tilde{t} with probability p . Then, $t_0 := \tilde{t} \times \{0\}$ and $t_1 := ([n] \setminus \tilde{t}) \times \{1\}$. Like t , the seed s is also defined as the union of two randomly chosen sets, $s = s_0 \cup s_1$, where $s_0 \subset t_0$ and $s_1 \subset t_1$.²⁰ These sets have fixed size; for a parameter $k \in \mathbb{N}$, $|s_0| = \frac{k}{2}$ and $|s_1| = \frac{k}{2}$. Now, the estimate for $\omega(\mathbf{q}_{\bar{t}})$ is computed as $f(t, \mathbf{q}_t, s) = \frac{1}{n}(|\bar{t}_0| \omega(\mathbf{q}_{s_0}) + |\bar{t}_1| \omega(\mathbf{q}_{s_1}))$.

We need to show that $\omega(\mathbf{q}_{\bar{T}})$ is likely to be close to $\omega(\mathbf{q}_S)$. Because we compute an estimate for $\omega(\mathbf{q}_{\bar{T}})$ as a function of $\omega(\mathbf{q}_{S_0})$ and $\omega(\mathbf{q}_{S_1})$, we will first show that (with high probability) $\omega(\mathbf{q}_{T_0}) \approx \omega(\mathbf{q}_{S_0})$ and $\omega(\mathbf{q}_{T_1}) \approx \omega(\mathbf{q}_{S_1})$. Then, we argue that $\omega(\mathbf{q}_{\bar{T}_0}) \approx \omega(\mathbf{q}_{T_0})$ and $\omega(\mathbf{q}_{\bar{T}_1}) \approx \omega(\mathbf{q}_{T_1})$, from which we can also conclude (using the union bound) that $\omega(\mathbf{q}_{\bar{T}_0}) \approx \omega(\mathbf{q}_{S_0})$ and $\omega(\mathbf{q}_{\bar{T}_1}) \approx \omega(\mathbf{q}_{S_1})$. Finally, we apply the union bound again and combine the two bounds to obtain an upper bound for $\Pr[|\omega(\mathbf{q}_{\bar{T}}) - \frac{1}{n}(|\bar{T}_0| \omega(\mathbf{q}_{S_0}) + |\bar{T}_1| \omega(\mathbf{q}_{S_1}))| \geq \delta]$.

The first step in the proof follows directly from Hoeffding's inequality,

$$\Pr[|\omega(\mathbf{q}_{T_0}) - \omega(\mathbf{q}_{S_0})| \geq \gamma] \leq 2 \exp(-2|S_0|\gamma^2) = 2 \exp(-k\gamma^2), \quad \text{for any } \gamma > 0.$$

Trivially, this bound also applies to the relation between $\omega(\mathbf{q}_{T_1})$ and $\omega(\mathbf{q}_{S_1})$, if we substitute appropriately. The second step, showing that $\omega(\bar{T}_0)$ (respectively $\omega(\bar{T}_1)$) is likely to be close to $\omega(T_0)$ (resp. $\omega(T_1)$), is slightly more involved. Namely, although the sum of the sizes of T_0 and T_1 is constant (to be precise, $|T_0| + |T_1| = n$), their individual sizes are random. In Example 3 (see also Appendix A.3), we have already encountered a similar, though not identical, situation, i.e., Example 3 considers uniformly random one-out-of-two sampling whereas here we analyze one-out-of-two sampling according to a Bernoulli $(p, 1-p)$ distribution. Nonetheless, it is straightforward to generalize the proof of Appendix A.3 to this (more general) case.

Let $X := |T_0|$. The expectation of X is given by $E[X] = np$. Let \mathcal{E} be the event that $X \in [(p - \beta)n, (p + \beta)n]$, for $\beta > 0$. From Hoeffding's inequality, we know that $\Pr[\bar{\mathcal{E}}] = \Pr[|\frac{X}{n} - p| \geq$

²⁰Again, Remark 1 applies.

$\beta] \leq 2 \exp(-2\beta^2 n)$. Like in Appendix A.3, we find an upper bound that holds conditioned on the event \mathcal{E} , by substituting the boundary points of the interval used to define \mathcal{E} in (6),

$$\begin{aligned} \Pr[|\omega(\mathbf{q}_{T_0}) - \omega(\mathbf{q}_{\bar{T}_0})| \geq \delta \mid \mathcal{E}] &\leq -2(p - \beta)n \left(1 - \frac{(p + \beta)n}{n}\right)^2 \\ &= 2 \exp(-2n\delta^2(1 - p - \beta)^2(p - \beta)). \end{aligned}$$

Next, we apply the union bound to show that for $0 < \epsilon < \gamma$

$$\Pr[|\omega(\mathbf{q}_{\bar{T}_0}) - \omega(\mathbf{q}_{S_0})| \geq \gamma \mid \mathcal{E}] \leq 2 \exp(-2n\epsilon^2(1 - p - \beta)^2(p - \beta)) + 2 \exp(-k(\gamma - \epsilon)^2)$$

By substituting p by $1 - p$ in the expression above, we also obtain

$$\Pr[|\omega(\mathbf{q}_{\bar{T}_1}) - \omega(\mathbf{q}_{S_1})| \geq \gamma \mid \mathcal{E}] \leq 2 \exp(-2n\epsilon^2(p - \beta)^2(1 - p - \beta)) + 2 \exp(-k(\gamma - \epsilon)^2)$$

Finally, we combine the two bounds and we get rid of the conditioning on \mathcal{E} by adding $\Pr[\bar{\mathcal{E}}]$. For any $\delta > 0$ and $0 < \epsilon < \delta$, we may write

$$\begin{aligned} \epsilon_{\text{class}}^\delta &= \max_{\mathbf{q}} \Pr[|\omega(\mathbf{q}_{\bar{T}}) - \frac{1}{n}(|\bar{T}_0| \omega(\mathbf{q}_{S_0}) + |\bar{T}_1| \omega(\mathbf{q}_{S_1}))| \geq \delta] \\ &= \max_{\mathbf{q}} \Pr[|\text{wt}(\mathbf{q}_{\bar{T}}) - |\bar{T}_0| \omega(\mathbf{q}_{S_0}) + |\bar{T}_1| \omega(\mathbf{q}_{S_1})| \geq n\delta] \\ &= \max_{\mathbf{q}} \Pr[|\text{wt}(\mathbf{q}_{\bar{T}}) - |\bar{T}_0| \omega(\mathbf{q}_{S_0}) + |\bar{T}_1| \omega(\mathbf{q}_{S_1})| \geq (|\bar{T}_0|\delta + |\bar{T}_1|\delta)] \\ &\leq \max_{\mathbf{q}} \Pr[|\omega(\mathbf{q}_{\bar{T}_0}) - \omega(\mathbf{q}_{S_0})| \geq \delta] + \Pr[|\omega(\mathbf{q}_{\bar{T}_1}) - \omega(\mathbf{q}_{S_1})| \geq \delta] \\ &\leq 2 \exp(-2n\epsilon^2(1 - p - \beta)^2(p - \beta)) + 2 \exp(-2n\epsilon^2(p - \beta)^2(1 - p - \beta)) + \dots \\ &\quad + 4 \exp(-k(\delta - \epsilon)^2) + 2 \exp(-2\beta^2 n) \end{aligned}$$

B Proof of Lemma 1

Proof. We will show that $|J|\rho_{WE}^{\text{mix}} \geq \rho_{WE}$, to be understood in that $|J|\rho_{WE}^{\text{mix}} - \rho_{WE}$ is positive semi-definite. With this shown, it then follows that for any density matrix σ_E and for any non-negative $h \in \mathbb{R}$

$$2^{-(h - \log |J|)} \cdot \mathbb{I}_W \otimes \sigma_E - \rho_{WE} \geq 2^{-h} |J| \cdot \mathbb{I}_W \otimes \sigma_E - |J| \rho_{WE}^{\text{mix}} = |J| (2^{-h} \cdot \mathbb{I}_W \otimes \sigma_E - \rho_{WE}^{\text{mix}})$$

so that if the right-hand side is positive semi-definite then so is the left-hand side. The claimed bound $H_{\min}(\rho_{WE}|E) \geq H_{\min}(\rho_{WE}^{\text{mix}}|E) - \log |J|$ then follows by the definition of the min-entropy.

Writing out the measurements explicitly yields

$$\rho_{WE} = \sum_{w \in \mathcal{W}} (|w\rangle\langle w| \otimes \mathbb{I}_E) |\varphi_{AE}\rangle\langle\varphi_{AE}| (|w\rangle\langle w| \otimes \mathbb{I}_E) = \sum_{w \in \mathcal{W}} \sum_{i, j \in J} \alpha_i \bar{\alpha}_j |w\rangle\langle w| |i\rangle\langle j| |w\rangle\langle w| \otimes |\varphi_E^i\rangle\langle\varphi_E^j|$$

and

$$\rho_{WE}^{\text{mix}} = \sum_{i \in J} |\alpha_i|^2 \sum_{w \in \mathcal{W}} |w\rangle\langle w| \otimes |\varphi_E^i\rangle\langle\varphi_E^i|.$$

We want to show that $\langle \xi | (|J| \rho_{WE}^{\text{mix}} - \rho_{WE}) | \xi \rangle \geq 0$ for all $|\xi\rangle \in \mathcal{H}_W \otimes \mathcal{H}_E$. We first consider $|\xi\rangle$ of the special form $|\xi\rangle = |v\rangle |\psi_E\rangle$ with $v \in \mathcal{W}$, and compute/bound $\langle \xi | \rho_{WE} | \xi \rangle$ and $\langle \xi | \rho_{WE}^{\text{mix}} | \xi \rangle$ as

$$\begin{aligned} \langle \xi | \rho_{WE} | \xi \rangle &= \sum_{i,j \in J} \alpha_i \bar{\alpha}_j \langle v|i\rangle \langle j|v\rangle \langle \psi_E | \varphi_E^i \rangle \langle \varphi_E^j | \psi_E \rangle = \left(\sum_{i \in J} \alpha_i \langle v|i\rangle \langle \psi_E | \varphi_E^i \rangle \right) \left(\sum_{j \in J} \bar{\alpha}_j \langle j|v\rangle \langle \varphi_E^j | \psi_E \rangle \right) \\ &= \left| \sum_{i \in J} \alpha_i \langle v|i\rangle \langle \psi_E | \varphi_E^i \rangle \right|^2, \end{aligned}$$

and

$$\langle \xi | \rho_{WE}^{\text{mix}} | \xi \rangle = \sum_{i \in J} |\alpha_i|^2 \langle v|i\rangle^2 |\langle \psi_E | \varphi_E^i \rangle|^2 \geq \frac{1}{|J|} \left| \sum_{i \in J} \alpha_i \langle v|i\rangle \langle \psi_E | \varphi_E^i \rangle \right|^2 = \frac{1}{|J|} \langle \xi | \rho_{WE} | \xi \rangle,$$

where the inequality follows from Cauchy-Schwarz inequality. The claim, $\langle \xi | (|J| \rho_{WE}^{\text{mix}} - \rho_{WE}) | \xi \rangle \geq 0$, for an *arbitrary* $|\xi\rangle = \sum_{w \in \mathcal{W}} \beta_w |w\rangle |\psi_E^w\rangle \in \mathcal{H}_W \otimes \mathcal{H}_E$ now follows by linearity, and by noting that $\langle v, \psi_E | \rho_{WE} | v', \psi_E' \rangle = 0 = \langle v, \psi_E | \rho_{WE}^{\text{mix}} | v', \psi_E' \rangle$ for all *distinct* $v, v' \in \mathcal{W}$, so that all “cross-products” vanish. \square

C The Tightness of Theorem 3

We show here that in general the inequality from Theorem 3 is tight. Specifically, we specify a natural class of sampling strategies for which Theorem 3 is an equality. Informally, this class consists of sampling strategies that behave in exactly the same way if the randomized choices T and S are replaced by *fixed* choices t_\circ and s_\circ , and instead the coordinates of \mathbf{q} are shuffled by means of a uniformly random permutation (chosen from a subgroup of all permutations). The formal definition is given below, but let us point out already here that Example 1 as well as the QKD sampling strategy discussed in Example 5 belong to this class. Indeed, for Example 1, instead of choosing a random subset T of size k one can equivalently choose a fixed subset and randomly permute the positions of \mathbf{q} . And, similarly for Example 5, instead of choosing left or right from each pair (q_{i0}, q_{i1}) at random and then choosing a random subset of size k of the selected q_{ij} 's, one can equivalently fix these choices and swap each pair (q_{i0}, q_{i1}) with probability $\frac{1}{2}$ and apply a random permutation to the first index.

Let S_n denote the symmetric group of degree n , i.e. the group of permutations on $[n]$. For any $\pi \in S_n$ and $\mathbf{q} = (q_1, \dots, q_n) \in \mathcal{A}^n$, we write $\pi \mathbf{q}$ to express that π permutes the *positions* of the elements of \mathbf{q} , i.e., $\pi \mathbf{q} = (q_{\pi^{-1}(1)}, \dots, q_{\pi^{-1}(n)})$. If \mathcal{V} is a set of strings $\mathbf{q} \in \mathcal{A}^n$, then $\pi \mathcal{V}$ means that the permutation π acts element-wise on \mathcal{V} .

Definition 5 (*G*-Symmetry of a sampling strategy). *Let Ψ be a sampling strategy, let G be a subgroup of S_n , where n is the size of the population to which Ψ is applied, and let Π be a random permutation, uniformly distributed over G . We call Ψ *G*-symmetric, if there exist $t_\circ \subset [n]$ and $s_\circ \in \mathcal{S}$ such that*

$$(\omega(\mathbf{q}_{\bar{T}}), f(T, \mathbf{q}_T, S)) \sim (\omega((\Pi \mathbf{q})_{\bar{t}_\circ}), f(t_\circ, (\Pi \mathbf{q})_{t_\circ}, s_\circ))$$

where “ \sim ” means that the pairs have the same probability distribution.

A direct consequence of this definition is the following relation, which we will apply later in this section.

$$\begin{aligned} B_{T,S}^\delta &= \{\mathbf{q} \in \{0,1\}^n : |\omega(\mathbf{q}_{\bar{T}}) - f(T, \mathbf{q}_T, S)| < \delta\} \\ &\sim \{\mathbf{q} \in \{0,1\}^n : |\omega((\Pi\mathbf{q})_{\bar{t}_o}) - f(t_o, (\Pi\mathbf{q})_{t_o}, s_o)| < \delta\} = \Pi^{-1}B_{t_o, s_o}^\delta. \end{aligned}$$

We can now rephrase Proposition 1 and prove it.

Proposition 1 (Rephrased). *For any G -symmetric sampling strategy Ψ_G^{sym} and any $\delta > 0$:*

$$\varepsilon_{\text{quant}}^\delta(\Psi_G^{\text{sym}}) = \sqrt{\varepsilon_{\text{class}}^\delta(\Psi_G^{\text{sym}})}$$

Proof. We need to show that there exists a system E and a state $|\varphi_{AE}\rangle$ such that $\Delta(\rho_{TSAE}, \tilde{\rho}_{TSAE})^2 = \varepsilon_{\text{class}}^\delta$ for $\tilde{\rho}_{TSAE}$ that minimizes the left hand side. As pointed out after the proof of Theorem 3, the particular construction of $\tilde{\rho}_{TSAE}$ used in the proof of Theorem 3 does minimize $\Delta(\rho_{TSAE}, \tilde{\rho}_{TSAE})$. Hence, it suffices to show that there exists a system E and a state $|\varphi_{AE}\rangle$ (that depends on G) such that

$$\Delta(\rho_{TSAE}, \tilde{\rho}_{TSAE})^2 \stackrel{(7)}{=} \left[\sum_{t,s} P_{TS}(t,s) |\langle \varphi_{AE} | \tilde{\varphi}_{AE}^{\perp ts} \rangle| \right]^2 \stackrel{(8)}{=} \sum_{t,s} P_{TS}(t,s) |\langle \varphi_{AE} | \tilde{\varphi}_{AE}^{\perp ts} \rangle|^2 \stackrel{(9)}{=} \varepsilon_{\text{class}}^\delta.$$

where $\tilde{\rho}_{TSAE}$ and $|\tilde{\varphi}_{AE}^{\perp ts}\rangle$ are constructed as in the proof of Theorem 3. The derivation of equality (7) can be found in the proof of Theorem 3. The outline of the remaining part of the proof is as follows; we first present a candidate for $|\varphi_{AE}\rangle$ and then we show that equalities (8) and (9) do indeed hold for this state.

We choose E to be empty. Furthermore, we define

$$|\varphi_{AE}\rangle := \frac{1}{\sqrt{|G|}} \sum_{\pi \in G} |\pi \mathbf{q}^*\rangle.$$

where \mathbf{q}^* is such that $\Pr[\mathbf{q}^* \notin B_{T,S}^\delta] = \varepsilon_{\text{class}}^\delta$. It follows from the projection construction for $\tilde{\rho}_{TSAE}$ that

$$|\tilde{\varphi}_{AE}^{\perp ts}\rangle = \frac{1}{\sqrt{|H_{t,s}|}} \sum_{\pi \in H_{t,s}} |\pi \mathbf{q}^*\rangle,$$

where $H_{t,s} \subseteq G$, i.e. $H_{t,s} := \{\pi \in G : \pi \mathbf{q}^* \notin B_{t,s}^\delta\}$.

To prove equality (8), we need to show that the inner product $|\langle \varphi_{AE} | \tilde{\varphi}_{AE}^{\perp ts} \rangle|$ is independent of t and s . Because $|\varphi_{AE}\rangle$ is a uniform superposition over permutations of \mathbf{q}^* and $|\tilde{\varphi}_{AE}^{\perp ts}\rangle$ is a renormalized projection of $|\varphi_{AE}\rangle$, we can easily compute this inner product, $|\langle \varphi_{AE} | \tilde{\varphi}_{AE}^{\perp ts} \rangle| = |H_{t,s}| / \sqrt{|G| \cdot |H_{t,s}|} = \sqrt{|H_{t,s}| / |G|}$. It suffices to show that $|H_{t,s}|$ is independent of (t, s) . It follows from the G -symmetry that there exists a π such that $B_{t,s}^\delta = \pi B_{t_o, s_o}^\delta$. Furthermore, let Π be a random permutation, uniformly distributed over G . By definition of $H_{t,s}$ and because Π is *uniformly* distributed over G , we may write

$$|H_{t,s}| = |G| \cdot \Pr[\Pi \mathbf{q}^* \notin B_{t,s}^\delta] = |G| \cdot \Pr[\mathbf{q}^* \notin \Pi^{-1} B_{t_o, s_o}^\delta] = |G| \cdot \Pr[\mathbf{q}^* \notin \Pi^{-1} B_{t_o, s_o}^\delta], \quad (10)$$

where the last expression is clearly independent of (t, s) .

Now, let us focus on equality (9). We derived in the proof of Theorem 3 that $\sum_{t,s} P_{TS}(t,s) |\langle \varphi_{AE} | \tilde{\varphi}_{AE}^{\perp ts} \rangle|^2 = \sum_{\mathbf{q}} P_{\mathbf{Q}}(\mathbf{q}) \Pr[\mathbf{q} \notin B_{T,S}^\delta]$, where the random variable \mathbf{Q} is obtained by measuring subsystem A of $|\varphi_{AE}\rangle$.

By definition of $|\varphi_{AE}\rangle$, $P_{\mathbf{Q}}(\mathbf{q}) > 0$ only for \mathbf{q} of the form $\pi\mathbf{q}^*$ for some $\pi \in G$. Hence, to prove equality (9), we have to show that for any $\pi \in G$, $\Pr[\pi\mathbf{q}^* \notin B_{T,S}^\delta] = \varepsilon_{\text{class}}^\delta$. This follows directly from the G -symmetry,

$$\Pr[\pi\mathbf{q}^* \notin B_{T,S}^\delta] = \Pr[\pi\mathbf{q}^* \notin \Pi^{-1}B_{t_o,s_o}^\delta] = \Pr[\mathbf{q}^* \notin \pi^{-1}\Pi^{-1}B_{t_o,s_o}^\delta] = \Pr[\mathbf{q}^* \notin \Pi^{-1}B_{T,S}^\delta] = \Pr[\mathbf{q}^* \notin B_{T,S}^\delta]. \quad (11)$$

Finally, note that (10) and (11) rely on the group structure of G . □