Sifting attacks in finite-size quantum key distribution

2016 New J. Phys. 18 053001

(http://iopscience.iop.org/1367-2630/18/5/053001)

New Journal of Physics
The open access journal at the forefront of physics

Deutsche Physikalische Gesellschaft DPG
IOP Institute of Physics

Published in partnership with: Deutsche Physikalische Gesellschaft and the Institute of Physics

OPEN ACCESS

RECEIVED
2 September 2015

REVISED
28 January 2016

ACCEPTED FOR PUBLICATION
6 April 2016

PUBLISHED
29 April 2016

PAPER

Sifting attacks in finite-size quantum key distribution

Corsin Pfister[1,2,4], **Norbert Lütkenhaus**[3], **Stephanie Wehner**[1,2] **and Patrick J Coles**[3]

[1] QuTech, Delft University of Technology, Lorentzweg 1, 2628 CJ Delft, The Netherlands
[2] Centre for Quantum Technologies, 3 Science Drive 2, 117543, Singapore
[3] Institute for Quantum Computing and Department of Physics and Astronomy, University of Waterloo, N2L3G1 Waterloo, Ontario, Canada
[4] Author to whom any correspondence should be addressed.

E-mail: mail@corsinpfister.com

Keywords: quantum key distribution, security loophole, quantum information

Abstract

A central assumption in quantum key distribution (QKD) is that Eve has no knowledge about which rounds will be used for parameter estimation or key distillation. Here we show that this assumption is violated for *iterative sifting*, a sifting procedure that has been employed in some (but not all) of the recently suggested QKD protocols in order to increase their efficiency. We show that iterative sifting leads to two security issues: (1) some rounds are more likely to be key rounds than others, (2) the public communication of past measurement choices changes this bias round by round. We analyze these two previously unnoticed problems, present eavesdropping strategies that exploit them, and find that the two problems are independent. We discuss some sifting protocols in the literature that are immune to these problems. While some of these would be inefficient replacements for iterative sifting, we find that the sifting subroutine of an *asymptotically* secure protocol suggested by Lo *et al* (2005 *J. Cryptol.* **18** 133–65), which we call LCA sifting, has an efficiency on par with that of iterative sifting. One of our main results is to show that LCA sifting can be adapted to achieve secure sifting in the *finite*-key regime. More precisely, we combine LCA sifting with a certain parameter estimation protocol, and we prove the finite-key security of this combination. Hence we propose that LCA sifting should replace iterative sifting in future QKD implementations. More generally, we present two formal criteria for a sifting protocol that guarantee its finite-key security. Our criteria may guide the design of future protocols and inspire a more rigorous QKD analysis, which has neglected sifting-related attacks so far.

1. Introduction

Quantum key distribution (QKD) allows for unconditionally secure communication between two parties (Alice and Bob). A recent breakthrough in the theory of QKD is the treatment of finite-key scenarios, pioneered by Renner and collaborators (see [1], for example). This has made QKD theory practically relevant, since the asymptotic regime associated with infinitely many exchanged quantum signals is an insufficient description of actual experiments. In practice, Alice and Bob have limited time, which in turn limits the number of photons they can exchange. For example, in satellite-based QKD [2] where, say, Bob is on the satellite and Alice is on the ground, the time allotted for exchanging quantum signals corresponds to the time for the satellite to pass overhead Alice's laboratory on the ground. Even if such considerations would not play a role, the necessity of error correction forces the consideration of finite-size QKD because error correcting codes operate on blocks of fixed finite length.

Finite-key analysis attempts to rigorously establish the security of finite-size keys extracted from finite raw data. A systematic framework for such analysis was developed by Tomamichel *et al* [3] involving the smooth entropy formalism. This framework was later extended to a decoy-state protocol by Lim *et al* [4]. An alternative framework was developed by Hayashi and collaborators [5, 6]. Other extensions of the finite-key framework include the treatment of device-independency by Tomamichel *et al* [7], Curty *et al* [8] and Lim *et al* [9], and

continuous-variable protocols by Furrer *et al* [10] and Leverrier [11]. The framework used in the aforementioned works, relying on some fairly technical results[5], represents the current state-of-the-art in the level of mathematical rigor for QKD security proofs. These theoretical advances have led to experimental implementations [12–14] with finite-key analysis.

For practical reasons, it is important to consider not only a protocol's security but also its efficiency. Ideally a protocol should use as little quantum communication as possible, for a given length of the final secret key. For example, it was noted by Lo *et al* [15] that—in the asymptotic regime—protocols with biased basis-choice probabilities can dramatically decrease the necessary amount of quantum communication per bit of the raw key. This is because a bias increases the probability that Alice and Bob measure in the same basis. As a consequence, when Alice and Bob perform the *sifting* step of the protocol, where they discard the outcomes of all measurements that have been made in different bases, they lose less data (see figure 2 and the discussion in section 5).

Some authors have adapted this bias in the basis choice in finite-key protocols and combined it with another measure to further decrease the amount of data that is lost through sifting. In the resulting sifting scheme, which we call *iterative sifting*, Alice and Bob announce previous basis choices while the quantum communication is still in process, and they terminate the quantum communication as soon as they have collected sufficiently many measurement outcomes in identical bases. This way, less quantum communication takes place, while at the same time they always make sure that they collect enough data. The implicit assumption here is that the knowledge of previous basis choices, but not of upcoming ones, does not help a potential eavesdropper.

As we show in this article, this assumption is wrong. Iterative sifting breaks the security proofs that have been presented for these protocols. This sifting scheme was part of theoretical protocols [3, 4, 8, 9] and has found experimental implementations [12]. Therefore, some (but not all) of the recently suggested protocols in QKD have serious security flaws.

1.1. Summary of the results

The issue with iterative sifting that we point out is as follows. Typical QKD protocols involve randomly choosing some rounds to be used for parameter estimation (PE) (i.e. testing for the presence of an eavesdropper Eve) and other rounds for key generation (KG). Naturally, if Eve knows ahead of time whether a round will be used for PE, i.e., if Eve knows which rounds will form the *sample* for testing for an eavesdropper's presence, then she can adjust her attack appropriately and the protocol is insecure. Hence a central assumption in the QKD security analysis is that Eve has no knowledge about the sample. We show that this assumption is violated for iterative sifting.

To be more precise, the iterative sifting scheme has two problems which, to our knowledge, have been neither addressed nor noted in the literature:

- *Non-uniform sampling*: The sampling probability, due to which the key bits and the encoding basis are chosen, is not uniform[6]. In other words, there is an *a priori* bias: Eve knows ahead of time that some rounds are more likely to end up in the sample than others.
- *Basis information leak*: Alice and Bob's public communication about their previous basis choices (which, in iterative sifting, happens before the quantum communication is over) allows Eve to update her knowledge about which of the upcoming (qu)bits end up in the sample. As a consequence, the quantum information that passes the channel thereafter can be correlated to this knowledge of Eve.

It is conceivable that these two problems become smaller as the size of the exchanged data increases. This would remain to be shown. More importantly, however, the protocols in question are designed to be secure for finite key lengths. In the light of these two problems, the analysis in the literature does currently not account for these finite-size effects. This is not a purely theoretical objection but a practically very relevant issue, as we present some eavesdropping attacks that exploit the problems.

As we discuss in section 5, the basis information leak can trivially be avoided by fixing the number of rounds in advance, and only announcing the basis choices after all quantum communication has taken place. We examine some sifting protocols from the literature with this property. In contrast to protocols that use iterative sifting, they often use fresh uniform randomness for the choice of the sample, and therefore are trivially

[5] These results include the uncertainty principle for smooth entropies and the operational meanings of these entropies.

[6] In general, the sampling probability (which decides over which of the bits are chosen as test bits) is distinguished from the probability distribution which decides in which basis the information is encrypted. In the literature, however, iterative sifting is combined with parameter estimation in a way such that bits measured in the X-basis are raw key bits, and bits measured in the Z-basis are used for parameter estimation. We will discuss this in more detail in the second half of section 2.

sampling uniformly. This means that they are secure with respect to our concerns. However, we find that there is room for improvement over these protocols regarding efficiency aspects.

Concretely, we note that one aspect that makes iterative sifting very efficient is the PE protocol that is used with it: after sifting, it simply uses the Z-bits as the sample for PE and the X-bits for raw key, which is why we call it the *single-basis parameter estimation* (*SBPE*). This is efficient because the sample choice requires no aditional randomness and no authenticated communication. While SBPE is insecure when used in conjunction with iterative sifting, it turns out to be secure when used with a sifting subroutine of a protocol suggested by Lo, Chau and Ardehali (LCA), which we call *LCA sifting*. The combination of LCA sifting and SBPE is essentially as efficient as iterative sifting. It has trivially no basis information leak and, as we prove, samples uniformly (see proposition 2). We therefore suggest this combination in future QKD protocols.

More generally, we find clear and explicit mathematical criteria that are sufficient for a sifting protocol to be secure in combination with SBPE. In contrast, current literature on QKD does not state such assumptions explicitly, but rather uses them implicitly.

In our formulation, they take the form of two equations

$$P_\Theta(\vartheta) = P_\Theta(\vartheta') \quad \forall\, \vartheta, \vartheta' \in \{0, 1\}_k^l \quad \text{and} \tag{1}$$

$$\rho_{A^l B^l \Theta^l} = \rho_{A^l B^l} \otimes \rho_{\Theta^l}. \tag{2}$$

Here, equation (1) expresses the absence of non-uniform sampling, i.e., that the probability $P_\Theta(\vartheta)$ for a partitioning ϑ of the total rounds into sample rounds and KG rounds is independent of ϑ. Equation (2) expresses the absence of basis information leak, which is formally expressed by stating that the classical communication Θ^l associated with the sifting process is uncorrelated (i.e., in a tensor product state) with Alice's and Bob's quantum systems $A^l B^l$. (The precise details of these two equations will be explained in section 6.) We find that the two problems are in fact independent. Hence, security from one of the two problems does not imply security from the other. The two formal criteria can be used to check whether a candidate protocol is subject to the two problems or not.

1.2. Outline of the paper

We introduce the iterative sifting protocol in section 2, where we also explain our conventions and notation. We give a detailed description of the two problems with iterative sifting in section 3. We show how these problems can be exploited in section 4 by presenting some intercept-resend attack strategies.

In section 5, we discuss some sifting protocols that are immune to these problems. We study how ideas of existing protocols can be combined to get new secure protocols that are more efficient. As a result, we suggest the aforementioned combination of LCA sifting and SBPE, and prove its security.

In section 6, we give a more general answer to the question of how the two problems can be avoided by presenting formal mathematical criteria that a sifting protocol needs to satisfy in order to avoid the problems. We conclude with a summary in section 7.

2. Iterative sifting and PE

A typical QKD protocol consists of the following subroutines [3]:

(i) Preparation, distribution, measurement and sifting, which we collectively refer to as 'sifting'.

(ii) Parameter estimation.

(iii) Error correction.

(iv) Privacy amplification.

What we discuss in this paper refers to the subroutines (i) and (ii), whereas subroutines (iii) and (iv) are not of our concern. Even though the word sifting usually only refers to the process of discarding part of the data acquired in the measurements, we refer to the preparation, distribution, measurement and sifting together as 'sifting', because they are intertwined in iterative sifting.

Our focus in this article is on a particular sifting scheme that we call iterative sifting. It has been formulated in slightly different ways in the literature, where the differences lie mostly in the choice of the wording and in whether it is realized as a prepare-and-measure protocol [3, 4, 8, 12] or as an entanglement-based protocol [9]. These details are irrelevant for the problems that we describe. Another difference is that some of the above-mentioned references take into consideration that sometimes, a measurement may not take place (no-detection event) or may have an inconclusive outcome. This is done by adding a third symbol $\varnothing$ to the set of possible outcomes, turning the otherwise dichotomic measurements into trichotomic ones with symbols $\{0, 1, \varnothing\}$. We

choose not to do so, because the problems that we describe arise independently of whether no-detection events or inconclusive measurements take place. Incorporating them would not solve the problems that we address but rather complicate things and distract from the main issues that we want to point out.

The essence of the iterative sifting protocol is shown in protocol 1. There, and in the rest of the paper, we use the notation

$$[r] := \{1, 2, \ldots, r\} \quad \text{for all } r \in \mathbb{N}_+. \tag{3}$$

Our formulation of this protocol is close to the one described in [3], with the main difference that we choose an entanglement-based protocol instead of a prepare-and-measure protocol. This will have the advantage that the formal criteria in section 6 are easier to formulate, but a prepare-and-measure based protocol would otherwise be equally valid to demonstrate our points.

In the protocol, Alice iteratively prepares qubit pairs in a maximally entangled state (step 1) and sends one half of the pair to Bob (step 2)[7]. Then, Alice and Bob each measure their qubit with respect to a basis $a_i, b_i \in \{0, 1\}$, respectively, where 0 stands for the X-basis and 1 stands for the Z-basis (steps 3 and 4). Thereby, Alice and Bob make their basis choice independently, where for each of them, 0 (X) is chosen with probability p_x, and 1 (Z) with probability p_z. These probabilities p_x and p_z are parameters of the protocol. The important and problematic parts of the protocol are step 5 and the subsequent check of the termination condition (TC): after *each* measurement, Alice and Bob communicate their basis choice over an authenticated classical channel. With this information at hand, they then check whether the TC is satisfied: if for at least n of the qubit pairs they had so far, they both measured in the X-basis, and for at least k of them, they both measured in the Z-basis, the TC is satisfied and they enter the *final phase* of the protocol by continuing with Step 6. These *quota* n and k are parameters of the protocol. If the condition is not met, they repeat the steps 1–5 (which we call the *loop phase* of the protocol) until they meet this condition. Because of this iteration, whose TC depends on the history[8] of the protocol run up to that point, we call it the iterative sifting protocol. Its number of rounds is a random variable that we denote by M. We denote possible values of M by m (see the TC and step 6).

Protocol 1. The iterative sifting protocol.

Iterative sifting

Parameters: $n, k \in \mathbb{N}_+$; $p_x, p_z \in [0, 1]$ with $p_x + p_z = 1$.
Output: For $l = n + k$, the outputs are:
Alice: l-bit string $(s_i)_{i=1}^l \in \{0, 1\}^l$ (sifted outcomes),
Bob: l-bit string $(t_i)_{i=1}^l \in \{0, 1\}^l$ (sifted outcomes),
public: l-bit string $(\vartheta_i)_{i=1}^l \in \{0, 1\}^l$ with $\sum_i \vartheta_i = k$ (basis choices, sifted), where 0 means X-basis and 1 means Z-basis.
Number of rounds: Random variable M, determined by reaching the termination condition (TC) after step 5.

The protocol

Loop phase: Steps 1–5 are iterated roundwise (round index $r = 1, 2, \ldots$) until the TC after step 5 is reached. Starting with round $r = 1$, Alice and Bob do:

Step 1: (Preparation): Alice prepares a qubit pair in a maximally entangled state.

Step 2: (Channel use): Alice uses the quantum channel to send half of the qubit pair to Bob.

Step 3: (Random bit generation): Alice and Bob each (independently) generate a random classical bit a_r and b_r, respectively, where 0 is generated with probability p_x and 1 with probability p_z.

Step 4: (Measurement): Alice measures her share in the X-basis (if $a_r = 0$) or in the Z-basis (if $a_r = 1$), and stores the outcome in a classical bit y_r. Likewise, Bob measures his share in the X-basis (if $b_r = 0$) or in the Z-basis (if $b_r = 1$), and stores the outcome in a classical bit y'_r.

Step 5: (Interim report): Alice and Bob communicate their basis choice a_r and b_r over a public authenticated channel. Then they determine the sets

$$u(r) := \{j \in [r] | a_j = b_j = 0\},$$
$$v(r) := \{j \in [r] | a_j = b_j = 1\}$$

TC: If the condition ($|u(r)| \geqslant n$ and $|v(r)| \geqslant k$) is reached, Alice and Bob set $m := r$ and proceed with step 6. Otherwise, they increment r by one and repeat from step 1.

[7] Choosing a maximally entangled state as the state that Alice prepares maximizes the probability that the correlation test in the PE (after sifting) is passed, i.e. the maximally entangled state maximizes the robustness of the protocol. However, for the security of the protocol, which is the concern of the present article, the choice of the state that Alice prepares is irrelevant.

[8] By the *history* of a protocol run, we mean the record of everything that happened during the run of the protocol. In the case of iterative sifting, this means the random bits a_r, b_r, the measurement outcomes y_r, y'_r etc.

(Continued.)

Final phase: The following steps are performed only once:

Step 6: (Random discarding): Alice and Bob choose a subset $u \subseteq u(m)$ of size n at random, i.e. each subset of size k is equally likely to be chosen. Analogously, they choose a subset $v \subseteq v(m)$ of size k at random. Then they discard the bits a_r, b_r, y_r and y'_r for which $r \notin u \cup v$.

Step 7: (Order-preserving relabeling): Let r_i be the ith element of $u \cup v$. Then Alice determines $(s_i)_{i=1}^l \in \{0, 1\}^l$, Bob determines $(t_i)_{i=1}^l \in \{0, 1\}^l$ and together they determine $(\vartheta_i)_{i=1}^l \in \{0, 1\}^l$, where for every $i \in [l]$,

$$s_i = y_{r_i}, \quad t_i = y'_{r_i}, \quad \vartheta_i = a_{r_i} (= b_{r_i}).$$

Step 8: (Output): Alice [Bob] locally outputs $(s_i)_{i=1}^l$ $[(t_i)_{i=1}^l]$, and they publicly output $(\vartheta_i)_{i=1}^l$.

After the loop phase of the protocol, in which the whole data is generated, Alice and Bob enter the final phase of the protocol, in which this data is processed. This processing consists of discarding data of rounds in which Alice and Bob measured in different bases, as well as randomly discarding a surplus of data for rounds where both measured in the same basis, where a 'surplus' refers to having more than $n(k)$ rounds in which both measured in the X (Z) basis, respectively. This discarding of surplus is done to simplify the analysis of the protocol, which is easier if the number of bits where both measured in the X (Z) basis is fixed to a number n (k). Since after the loop phase, Alice and Bob can end up with more bits measured in this same basis, they throw away surplus at random. Finally, after throwing away the surplus, Alice and Bob locally output the remaining bit strings $(s_i)_{i=1}^l$ and $(t_i)_{i=1}^l$ of measurement outcomes and publicly output the remaining bit string $(\vartheta_i)_{i=1}^l$ of basis choices.

Iterative sifting is problematic, but to fully understand why, one needs to see how the output of the iterative sifting protocol is processed in the subsequent subroutine (ii), the PE, where Alice and Bob check for the presence of an eavesdropper. Protocols that use iterative sifting use a particular protocol for PE. To make clear what we are talking about, we have written it out in protocol 2.

Alice and Bob start the protocol with the strings $(s_i)_{i=1}^l$, $(t_i)_{i=1}^l$ and $(\vartheta_i)_{i=1}^l$ that they got from sifting. Then, in a first step, they communicate the *test bits*. The test bits are those bits s_i, t_i that resulted from measurements in the Z-basis, i.e. the bits s_i, t_i with i such that $\vartheta_i = 1$. Then, they determine the fraction of the test bits that are different for Alice and Bob, i.e. they determine the *test bit error rate*. If it is higher than a certain protocol parameter $q_{\text{tol}} \in [0, 1]$, they abort. Otherwise, they locally output the *raw keys*, which are the bits s_i, t_i that result from measurements in the X-basis, i.e. those s_i, t_i with i for which $\vartheta_i = 0$.

It is important to emphasize that if the output of iterative sifting serves as the input of the PE protocol as in protocol 2, then the bits that result from measurements in the X-basis are used for the raw key, and the bits that result from measurements in the Z-basis are used for PE (i.e. they form the *sample* for the PE). Hence, the sample is determined by the basis choice; no additional randomness is injected to choose the sample.

Protocol 2. The single-basis parameter estimation (SBPE) protocol.

Single-basis parameter estimation (SBPE)

Protocol parameters: n, $k \in \mathbb{N}_+$, p_x, $p_z \in [0, 1]$ with $p_x + p_z = 1$ and $q_{\text{tol}} \in [0, 1]$.

Input: For $l = n + k$, the inputs are:

Alice: l-bit string $(s_i)_{i=1}^l \in \{0, 1\}^l$ (measurement outcomes, sifted),

Bob: l-bit string $(t_i)_{i=1}^l \in \{0, 1\}^l$ (measurement outcomes, sifted),

public: l-bit string $(\vartheta_i)_{i=1}^l \in \{0, 1\}^l$ with $\sum_i \vartheta_i = k$ (basis choices, sifted), where 0 means X-basis and 1 means Z-basis.

Output: Either no output (if the protocol aborts in step 2) or:

Alice: n-bit string $(x_j)_{j=1}^n \in \{0, 1\}^n$ (raw key),

Bob: n-bit string $(x'_j)_{j=1}^n \in \{0, 1\}^n$ (raw key).

The protocol

Step 1: (Test bit communication): Alice and Bob communicate their test bits, i.e. the bits s_i and t_i with i for which $\vartheta_i = 1$, over a public authenticated channel.

Step 2: (Correlation test): Alice and Bob determine the *test bit error rate*

$$\lambda_{\text{test}} := \frac{1}{k} \sum_{i=1}^{l} \vartheta_i (s_i \oplus t_i),$$

where $\oplus$ denotes addition modulo 2, and do the *correlation test*: if $\lambda_{\text{test}} \leqslant q_{\text{tol}}$, they continue the protocol and move on to step 3. If $\lambda_{\text{test}} > q_{\text{tol}}$, they abort.

(Continued.)

Step 3: (Raw key output): Let i_j be the jth element of $\{i \in [l] | \vartheta_i = 0\}$. Then Alice outputs the n-bit string $(x_j)_{j=1}^n$ and Bob outputs the n-bit string $(x'_j)_{j=1}^n$, where

$$x_j = s_{i_j}, \quad x'_j = t_{i_j}.$$

This is not necessarily a problem by itself. However, as we will show in section 3.1, in iterative sifting, some rounds are more likely to end up in the sample than other rounds. This leads to non-uniform sampling, which is a problem since uniform sampling is one of the assumptions that enter the analysis of the PE. This seems to be unnoticed so far, as we found that protocols in the literature that use iterative sifting as a subroutine use SBPE as a subroutine for PE (or something equivalent) [3, 4, 8, 9, 12]. In contrast, the LCA sifting protocol that we discuss in section 5 *does* sample uniformly, even if bits from X-measurements are used for the raw key and Z-measurements are used for paremeter estimation, without injecting additional randomness.

We will discuss randomness injection for the sample choice in more detail in section 5. The idea behind the PE is the following: if the correlation test passes, then the likelihood that Eve knows much about the raw key is sufficiently low. The exact statement of this is subtle, and involves more details than are necessary for our purposes. We refer to [3] for more details. Here, what is important is that this estimate of Eve's knowledge is done via estimating another probability that we call the *tail probability* $p_{\text{tail}}(\mu)$ which, for $\mu \in [0, 1]$, is given by

$$p_{\text{tail}}(\mu) = P[\Lambda_{\text{key}} \geqslant \Lambda_{\text{test}} + \mu | \Lambda_{\text{test}} \leqslant q_{\text{tol}}]. \tag{4}$$

Here, Λ_{test} is the random variable of the test bit error rate λ_{test} determined in the PE protocol

$$\lambda_{\text{test}} := \frac{1}{k} \sum_{i=1}^{l} \vartheta_i (s_i \oplus t_i). \tag{5}$$

The random variable Λ_{key} is the random variable of a quantity that is not actually measured: it is the random variable of the error rate on the raw key bits *if they had been measured in the Z-basis*. Since in the actual protocol, the raw key bits have been measured in the X-basis, the random variable Λ_{key} is the result of a *Gedankenexperiment* rather than an actually measured quantity. We will define Λ_{key} formally in section 6.

The usual analysis, as in [3], aims at proving that

$$p_{\text{tail}}(\mu) \leqslant \frac{\exp\left(-2\frac{kn}{l}\frac{k}{k+1}\mu^2\right)}{p_{\text{pass}}}, \tag{6}$$

where

$$p_{\text{pass}} = P[\Lambda_{\text{test}} \leqslant q_{\text{tol}}]. \tag{7}$$

Inequality (6) is turned into an inequality about the eavesdropper's knowledge about the raw key using an uncertainty relation for smooth entropies [3, 16].

2.1. Notation and terminology

In the following sections, we will have a closer look at the probabilities of certain outputs of the iterative sifting protocol in protocol 1. For example, in section 3.1 we will consider the probability that iterative sifting with parameters $n = 1$, $k = 2$ outputs the string $\vartheta = (\vartheta_i)_{i=1}^3 = (1, 1, 0)$. Since the output of the protocol is probabilistic, the output string becomes a random variable. We denote random variables by capital letters and their values by lower case letters. For example, the random variable for the output string ϑ is denoted by Θ, and the probability of the output string to have a certain value ϑ is $P[\Theta = \vartheta]$. For strings in $\vartheta = (\vartheta_i)_{i=1}^l \in \{0, 1\}^l$, we write $(\vartheta_i)_{i=1}^l = \vartheta_1\vartheta_2 \ldots \vartheta_l$ instead of $(\vartheta_i)_{i=1}^l = (\vartheta_1, \vartheta_2, \ldots, \vartheta_l)$, i.e. we omit the brackets and commas. For example, we write $110 \in \{0, 1\}^3$ instead of $(1, 1, 0) \in \{0, 1\}^3$, so the probability that we calculate in section 3.1 is $P[\Theta = 110]$. Other random variables that we consider include the random variable A_1 (B_1) of Alice's (Bob's) first basis choice a_1 (b_1) or the random variable M of the number m of total rounds performed in the loop phase of the iterative sifting protocol.

To simplify the calculations, it is convenient to introduce the following terminology. For a round r in the loop phase of the iterative sifting protocol, r is an X-agreement if $a_r = b_r = 0$, r is a Z-agreement if $a_r = b_r = 1$ and r is a disagreement if $a_r \neq b_r$. We sometimes say that r is an agreement if it is an X- or a Z-agreement.

For calculations with random variables like Θ, A_1, B_1 or M, the sample space of the relevant underlying probability space is the set of all possible histories of the iterative sifting protocol. This set is hard to model, as it contains not only all possible strings $(a_r)_r$, $(b_r)_r$, $(y_r)_r$ and $(y'_r)_r$ of the loop phase (which can be arbitrarily long) but also a record of the choice of the subsets u and v in the random discarding during the final phase. It is,

however, not necessary for our calculations to have the underlying sample space explicitly written out. In order to avoid unnecessarily complicating things, we therefore only deal with the relevant events, random variables and their probability mass functions directly, assuming that the reader understands what probability space they are meant to be defined on. In contrast, the LCA sifting protocol which we discuss in section 5, has a simpler set of histories, and we will derive a probability space model for it in appendix C.

We often write expressions in terms of probability mass functions instead of in terms of probability weights of events, e.g. we write

$$P_\Theta(\vartheta) \coloneqq P[\Theta = \vartheta]. \tag{8}$$

3. The problems

3.1. Non-uniform sampling

To show that iterative sifting leads to non-uniform sampling, we calculate the sampling probabilities for some example parameters $k,\ n \in \mathbb{N}_+$ as functions of the probabilities p_x and p_z. By a sampling probability, we mean the probability that some subset of k of the $l = n + k$ bits is used as a sample for the PE, i.e. the sampling probabilities are $P_\Theta(\vartheta)$ for $\vartheta \in \{0,\ 1\}_k^l$, where

$$\{0,\ 1\}_k^l \coloneqq \left\{ (\vartheta_i)_{i=1}^l \in \{0,\ 1\}^l \;\middle|\; \sum_{i=1}^l \vartheta_i = k \right\} \tag{9}$$

is the set of all l-bit strings with Hamming weight k. We say that sampling is uniform if $P_\Theta(\vartheta)$ is the same for all $\vartheta \in \{0,\ 1\}_k^l$, and non-uniform otherwise. While non-uniform sampling already arises in the case of the smallest possible parameters $k = n = 1$, the results are even more interesting in cases where $k \neq n$. Let us consider iterative sifting (protocol 1) with $n = 1,\ k = 2$ and arbitrary $p_x,\ p_z \in [0,\ 1]$. Let Θ denote the random variable of the string $\vartheta = (\vartheta_i)_{i=1}^3 = \vartheta_1\vartheta_2\vartheta_3$ of sifted basis choices which is generated by the protocol. The possible values of Θ are 110, 101 and 011. The probabilities of these strings are given as follows (see appendix A for a proof).

Proposition 1. *For the iterative sifting protocol as in protocol 1 with $n = 1$ and $k = 2$, it holds that*

$$P_\Theta(110) = g_z^2, \quad \text{where} \quad g_z = \frac{p_z^2}{p_z^2 + p_x^2}. \tag{10}$$

For the other two possible values of Θ, it holds that

$$P_\Theta(011) = P_\Theta(101) = \frac{1 - g_z^2}{2}. \tag{11}$$

Hence, different samples have different probabilities, in general. In order for the sampling probability P_Θ to be uniform, in the case where $n = 1$ and $k = 2$, we need to have $P_\Theta(\vartheta) = 1/3$ for $\vartheta = 011,\ 101,\ 110$. This holds if and only if $g_z = g_z^*$, where $g_z^* = 1/\sqrt{3}$, which in turn is equivalent to $p_z = p_z^*$, where

$$p_z^* = \frac{(3 + 2\sqrt{3})(1 + \sqrt{\sqrt{3} - 1})}{\sqrt{3}} \approx 0.539. \tag{12}$$

This is bad news for iterative sifting: it means that iterative sifting leads to non-uniform sampling for all values of p_z except $p_z = p_z^*$. Interestingly, the value of p_z^* does not seem to be a probability that has been considered in the QKD literature. In particular, p_z^* corresponds to neither the symmetric case $p_z = 1/2$ nor to a certain asymmetric probability which has been suggested to be chosen in order to maximize the key rate [3].

The value g_z can be interpreted as the probability that in a certain round of the loop phase, Alice and Bob have a Z-agreement, given that they have an agreement in that round (this conditional is why the p_z^2 is renormalized with the factor $1/(p_z^2 + p_x^2)$). Hence, g_z^2 is the probability that Alice and Bob's first two basis agreements are Z-agreements. Therefore, $P_\Theta(110) = g_z^2$ is what one would intuitively expect: to end up with $\Theta = 110$, the first two basis agreements need to be Z-agreements, and conversely, whenever the first two basis agreements are Z-agreements, Alice and Bob end up with $\Theta = 110$.

More generally, it turns out that for $n = 1$ and for $k \in \mathbb{N}_+$ arbitrary, the iterative sifting protocol leads to

$$P_\Theta(1\ \ldots\ 10) = g_z^k, \tag{13}$$

$$P_\Theta(\vartheta) = \frac{1 - g_z^k}{k} \quad \text{for all other } \vartheta \in \{0,\ 1\}_k^l. \tag{14}$$

This is a uniform probability distribution if and only if $g_z = g_z^*$, where

$$g_z^* = \left(\frac{1}{k+1}\right)^{1/k}, \tag{15}$$

which is true iff $p_z = p_z^*$, where

$$p_z^* = \frac{g_z^* - \sqrt{g_z^*(1-g_z^*)}}{2g_z^* - 1}. \tag{16}$$

Hence, we conclude that iterative sifting does not lead to uniformly random sampling, unless p_x and p_z are chosen in a very particular way. This particular choice does not seem to correspond to anything that has been considered in the literature so far.

3.2. Basis information leak

In iterative sifting, information about Alice's and Bob's basis choices reaches Eve in every round of the loop phase. In step 5 of round r, Alice and Bob communicate their basis choice a_r, b_r of that round. They do so because they want to condition their upcoming action on the strings $a_1 \ldots a_r$ and $b_1 \ldots b_r$: if they have enough basis agreements, they quit the loop phase; otherwise they keep looping.

What seems to have remained unnoticed in the literature is that Eve can also condition her actions on $a_1 \ldots a_r$ and $b_1 \ldots b_r$. This means that if there is a round $r+1$, Eve can correlate the state of the qubit that Alice sends to Bob in round $r+1$ with $a_1 \ldots a_r$ and $b_1 \ldots b_r$. Hence, the state of the qubit that Bob measures is correlated with the classical register that keeps the information about the basis choice. Note that the basis information leak tells Eve how close Alice and Bob are to meeting their quotas for each basis. Eve can tailor her attack on future rounds based on this information. For example, if Alice and Bob have already met their Z-quota, but not their X-quota, then Eve can measure in the X-basis, knowing that, if Alice and Bob happen to both measure in the Z-basis, the round may be discarded anyway.

We want to emphasize that the basis information leak is not resolved by injecting additional randomness for the choice of the sample. As we will discuss in section 5, such additional randomness can ensure that the sampling is uniform, but it does not help against the basis information leak. Randomness injection for the sample is effectively equivalent to performing a random permutation on the qubits [17]. This does not remove the correlation between the classical basis information register and the qubits.

We will see more concretely how the basis information leak is a problem when we present an eavesdropping attack in section 4.1 and when we treat the problem more formally in section 6.

4. Eavesdropping attacks

A detailed analysis of the effect of non-uniform sampling and basis information leak on the key rate is beyond the scope of the present paper. It would involve developing a new security analysis for a whole protocol involving iterative sifting. Instead of attempting to find a modified analysis for iterative sifting, we will discuss alternative protocols in section 5.

However, to give an intuitive idea of the effect, we will calculate another figure of merit: the error rate for an intercept-resend attack. We devise a strategy for Eve to attack the iterative sifting protocol during its loop phase and calculate the expected value of the error rate

$$E = \frac{1}{l}\sum_{i=1}^{l} S_i \oplus T_i \tag{17}$$

that results from this attack. Here, $\oplus$ denotes addition modulo 2 and S_i and T_i are the random variables of the bits s_i and t_i, respectively, which are generated by the protocol. One would typically expect an error rate no lower than 25% for an intercept-resend attack [18], which is why our results below are alarming.

4.1. Attack on non-uniform sampling

Let us first consider an attack on non-uniform sampling, i.e., on the fact that not every possible value of Θ is equally likely. It will be a particular kind of intercept-resend attack, i.e. Eve intercepts all the qubits that Alice sends to Bob during the loop phase, measures them in some basis and afterwards, prepares another qubit in the eigenstate associated with her outcome and sends it to Bob. Then we will show that the attack strategy leads to an error rate below 25%.

For the error rate calculation, we assume that the X- and Z-basis is the same for Alice, Bob and Eve, and that they are mutually unbiased. This way, if Alice and Bob measure in the same basis, but Eve measures in the other basis, then Eve introduces an error probability of $1/2$ on this qubit. Moreover, for simplicity, we make this

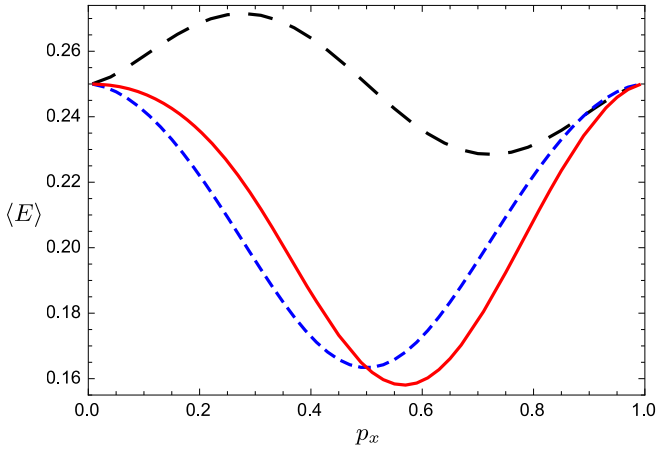

Figure 1. The error rate for three different eavesdropping attacks iterative sifting: (1) attack on non-uniform sampling (long-dashed, black curve), (2) attack on basis-information leak (short-dashed, blue curve), (3) attack on both problems (solid, red curve).

calculation for the easiest possible choice of parameters. Consider the iterative sifting iterative sifting protocol (protocol 1) with the parameters $k = n = 1$. From equations (15) and (16), we get that the sampling probabilities in this case are

$$P_\Theta(01) = \frac{p_x^2}{p_x^2 + p_z^2}, \quad P_\Theta(10) = \frac{p_z^2}{p_x^2 + p_z^2}. \tag{18}$$

These sampling probabilities are uniform for the symmetric case $p_x = p_z$, but are non-uniform for all other values. In the following, we assume $p_x > 1/2$, which makes the sample $\Theta = 01$ more likely than the sample $\Theta = 10$. We choose the following attack: in the first round of the loop phase, she attacks in the X-basis, and in all the other rounds, she attacks in the Z-basis. We choose the attack this way because we know that the first non-discarded basis agreement is more likely to be an X-agreement, whereas the second one is more likely to be a Z-agreement[9].

We calculate the expected error rate for this attack in appendix B.1. The black curve in figure 1 shows $\langle E \rangle$ as a function of p_x for this attack. Notice that $\langle E \rangle$ falls below 25% for $1/2 < p_x < 1$, and reaches a minimum of $\langle E \rangle \approx 22.8\%$ for $p_x \approx 0.73$.

The concerned reader might worry that the 25% error rate associated with the intercept-resend attack was derived under the assumption of equal weighting for the two bases X and Z, whereas it seems here that we choose unequal weightings. However, for the protocol under consideration, the *a priori* probability distribution $\{p_x, p_z\}$ is not the relevant quantity. Rather, the fact that $n = k$ in our example ensures that the X and Z bases enter in with equal weighting.

4.2. Attack on basis information leak

We now give an eavesdropping strategy that exploits the basis information leak. It is an adaptive strategy, in which Eve's action in round $r + 1$ depends on the past communication of the strings $a_1 \ldots a_r$ and $b_1 \ldots b_r$. Again, we consider the simple case of $n = k = 1$. To make sure our attack is really exploiting the basis information leak and not the non-uniform sampling, we set $p_x = p_z = 1/2$. In this case, from equation (18), the sampling is uniform:

$$P_\Theta(01) = P_\Theta(10) = \frac{1}{2}. \tag{19}$$

Before we define Eve's strategy, we want to give some intuition. Suppose that during the protocol, Eve learns that Alice and Bob just had their first basis agreement. If this first agreement is a Z-agreement, say, what does this mean for Eve? She knows that the protocol will now remain in the loop phase until they end up with an X-agreement. Suppose that she now decides that she will measure all the remaining qubits in the X-basis. Then, if the next basis agreement of Alice and Bob is an X-agreement, Eve knows the raw key bit perfectly, and her measurement on that bit did not introduce an error. If the next basis agreement is a Z-agreement, she may

[9] The attentive reader may point out that this attack could be improved by making Eve's basis choice dependent on the communication between Alice and Bob. This is correct, but we intentionally design the attack such that Eve ignores Alice and Bob's communication. That allows one to see the effect of non-uniform sampling alone and to compare it to attacks on basis information leak alone, see sections 4.2 and 4.3.

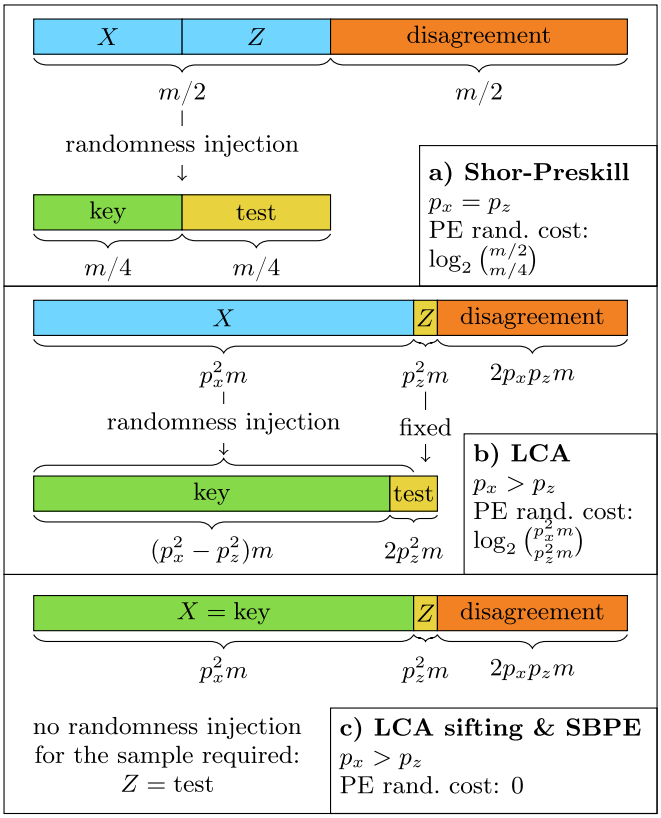

Figure 2. Comparison of the expected sifting efficiencies. (a) In the protocol of Shor and Preskill [21], only about a quarter of the measurement results end up in the raw key. Moreover, a relatively large amount of randomness needs to be injected for the sample choice, which in turn increases the length of pre-shared secret key that Alice and Bob use for authenticated communication. (b) The protocol by Lo *et al* [15] allows for a bias, $p_x > p_z$. This way, the expected fraction of bits with basis disagreements shrinks from one half to $2p_x p_z$. The proportions drawn in this figure correspond to $p_x = 0.8$. However, it still requires randomness injection for the choice of the sample. (c) If, instead, LCA sifting and SBPE are used, as we suggest, then no randomness injection is required for the choice of the sample. Moreover, less bits are consumed for parameter estimation in the finite-key regime, resulting in a longer raw key.

introduce an error on that test bit. However, there will be a chance that Alice and Bob discard this test bit, because they have a total of two (or more, in the end) Z-agreements, and the protocol forces them to discard all Z-agreements except $k = 1$ of them. Hence, learning that the first basis agreement was a Z-agreement brings Eve into an favorable position: she knows that attacking in the X-basis for the rest of the loop phase will necessarily tell her the raw key bit, while she has quite some chance to remain undetected.

This intuition inspires the following intercept-resend attack. Before the first round of the loop phase, Eve flips a fair coin. Let F be the random variable of the coin flip outcome and let 0 and 1 be its possible values. If $F = 0$, then in the first round, Eve attacks in the X-basis, and if $F = 1$, she attacks in the Z-basis. In the subsequent rounds, she keeps attacking in that basis until Alice and Bob first reached a basis agreement. If it is an X-agreement (equivalent to $\Theta = 01$), Eve attacks in the Z-basis in all remaining rounds, and if it is a Z-agreement (equivalent to $\Theta = 10$), she attacks in the X-basis in all remaining rounds[10].

We calculate the expected error rate for this attack in the appendix B.2. We find that

$$\langle E \rangle = \frac{2 - \ln 2}{8} \approx 16.3\%. \tag{20}$$

Hence, the basis information leak allows Eve to go far below the typical expected error rate of 25% for intercept-resend attacks [19]. The blue curve in figure 1 shows, more generally, $\langle E \rangle$ as a function of p_x, for this attack.

4.3. Independence of the two problems

Are non-uniform sampling and basis information leak really two different problems, or is one a consequence of the other? We will argue now that the two problems are in fact independent. To this end, we describe a protocol that suffers from non-uniform sampling but not from basis information leak, and another protocol that suffers from basis information leak but not from non-uniform sampling.

[10] We let Eve flip a coin in order to make the attack symmetric between X and Z. This allows for a more meaningful comparison with the attack on non-uniform sampling, as this attack here does not exploit non-uniform sampling even if $p_x \neq 1/2$, see section 4.1 and 4.3.

We have already seen an instance of a protocol that suffers from basis information leak but not from non-uniform sampling: in section 4.2, we looked at the iterative sifting protocol with $n = k = 1$ and $p_x = p_z = 1/2$, in which case the sampling is uniform. Hence, there was no exploitation of non-uniform sampling, but the attack strategy exploited basis information leak.

What about the other way round? Can non-uniform sampling occur without basis information leak? A closer look at the attack on non-uniform sampling presented in section 4.1 hints that this is possible: the attack strategy works, even though it completely ignores the communication between Alice and Bob, so it did not make any use of the basis information leak due to this communication.

A more dramatic example shows clearly that non-uniform sampling can occur without basis information leak. To this end, we forget about iterative sifting for a moment and look at a different protocol. Consider a sifting-protocol in which Alice and Bob agree in advance that they will measure the first $n = 100$ qubits in the X-basis, and that they will measure the second $k = 100$ qubits in the Z-basis, without any communication during the protocol. Of course, there is no hope for this protocol to be useful for QKD, but it serves well to demonstrate our point. It leads to a very dramatic form of non-uniform sampling, because $P_\Theta(0 \ldots 01 \ldots 1) = 1$ and $P_\Theta(\vartheta) = 0$ for all other $\vartheta \in \{0, 1\}^l_k$. If Eve attacks the first 100 rounds in X and the second 100 rounds in Z, then she knows the raw key perfectly, without introducing any error. At the same time, there is no communication between Alice and Bob during the protocol, so no information about the basis choice is *leaked during the protocol*. Instead, Eve (who is always assumed to know the protocol) already had this information before the first round.

Hence, we conclude that the problems of non-uniform sampling and basis information leak are independent. They just happen to occur simultaneously for iterative sifting, but they can occur separately in general. We will see the independence of the two problems more formally in section 6.

4.4. Attack on both problems

Since the two problems are independent, it is interesting to devise an attack that exploits both of them. Let us again consider $k = n = 1$ and suppose $p_x > 1/2$ to ensure that we have non-uniform sampling. Suppose Eve begins in the same way as in the attack on non-uniform sampling, measuring in the X-basis. However, as in the attack on the basis-information leak, she makes her attack adaptive by following the rule that she switches to the Z-basis when Alice and Bob announce that they had an X-agreement. If Alice and Bob announce a Z-agreement, Eve keeps attacking in the X-basis.

We give an expression for the error rate induced by this attack in appendix B.3. The red curve in figure 1 shows a plot of this error rate as a function of p_x. As one can see, the error rate attains its minimum of $\langle E \rangle \approx 15.8\%$ for $p_x \approx 0.57$. Hence, this combined attack on both problems performs much better than the one on non-uniform sampling alone (with a minimal error rate of $\sim$22.8%) and even better than the attack on the basis information leak alone (with a minimal error rate of $\sim$16.3%).

5. Solutions to the problems

How can these problems be avoided? Roughly speaking, we can say that protocols with iterative sifting are characterized by three properties that make it efficient: (1) asymmetric basis choice probabilities and quota, $p_x > p_z$ and $n > k$, (2) SBPE (protocol 2), (3) communication in step 5 of the loop phase. As we have seen, it is the communication which causes the basis information leak.

An obvious fix to this problem is to take this communication out of the loop phase and to postpone it to the final phase, when all the quantum communication is over. Then there is no classical communication during the loop phase, and hence, there cannot be a TC that depends on classical communication. Instead, the number of rounds in the loop phase is set to a fixed number $m \in \mathbb{N}_+$. This number m then becomes a parameter of the protocol.

Fixing the number of rounds introduces a new issue: there is no guarantee that the quotas for X- and Z-agreements will be met after m rounds. In order to perform the PE, however, the quotas n and k must be met. Otherwise, Inequality (6) is not applicable, because the number of X- and Z-agreements in the loop phase are random numbers that can be below n and k, respectively. Thus, unless one wants to introduce a new tail probability analysis as well, there is a strictly positive probability that Alice and Bob have to abort the sifting protocol because they have too many basis disagreements. If the sifting scheme is modified in this way, it no longer involves any communication about the basis choices during its loop phase. Thus, it is trivially true that there is no basis information leak.

Many protocols in the QKD literature have such a fixed number m of rounds (which is often denoted by N instead) and an according abort event. It seems that before iterative sifting was introduced, the sifting procedure was either not clearly written out in the protocols, or it had such a fixed round number. For example, in the original

BB84 paper [20], the sifting scheme is not written out in enough detail to say whether this is the case, but the protocol for which Shor and Preskill showed asymptotic security uses a fixed number of rounds [21]. In addition, they use symmetric basis choice probabilities and quota, i.e. $p_x = p_z = 1/2$ and $k = n$. Alice sends $4n + \delta$ qubits to Bob (where δ is a positive but small overhead) without any intermediate classical communication. Afterwards, they compare their bases and check whether they have at least n X-agreements and at least n Z-agreements. If not, they abort, otherwise they choose n X-agreements and n Z-agreements and discard the rest.

With the remaining $2n$ bits, they continue with PE. However, instead of performing SBPE, they choose n bits at random (i.e. with fresh randomness) for PE and use the rest for the raw key. Hence, this protocol shares none of the three properties with iterative sifting that we listed above.

This scheme trivially has no basis information leak. In addition, it trivially samples uniformly, as the whole sample is chosen with fresh randomness that is injected for that purpose. Thus, it is secure with respect to the concerns raised in this article. However, it is unnecessarily inefficient: speaking in expectation values, half of the bits are discarded because they were determined in different bases, and another quarter of the bits is used for PE, leaving only a quarter of the original bits for the raw key, see figure 2(a).

A similar protocol has recently been suggested by Tomamichel and Leverrier with a complete proof of its security, modeling all its subroutines [22]. They also use symmetric basis choice probabilities $p_x = p_z$ and randomness injection for the sample choice. However, they do not use half of the sifted bits for PE but less. Their protocol also samples uniformly, because additional randomness is injected for the choice of the sample.

To increase the efficiency, LCA suggested to use asymmetric basis choice probabilities and quota, i.e. $p_x > 0$ and $k \neq n$. As shown in figure 2(b), this decreases the number of expected disagreements from a value of $m/2$ to a value of $2p_x p_z m$. This is great for efficiency: for larger block lengths, relatively smaller samples are required to gain the same confidence that Alice's and Bob's bits are correlated[11]. In the limit where $m \to \infty$, the probability p_x can be chosen to be arbitrarily close to one, and the fraction of data lost due to basis disagreements converges to zero. We call this protocol *LCA sifting*. It shares property (1) with iterative sifting.

As for the protocol of Shor–Preskill, LCA did not consider SBPE. Their PE also requires some randomness injection for the choice of the sample: the Z-agreements form one half of the sample, and the other half is chosen at random from the X-agreements. Then, not just one but two error rates are determined, namely on the X-part and the Z-part of the sample separately. Only if *both* error rates are below a fixed error tolerance, they continue the protocol using the rest as the raw key (for details, see their article [15]). The LCA protocol trivially has no basis information leak. In addition, it turns out that it also samples uniformly. This is in fact non-trivial, and to our knowledge, it was not proved in the literature. We fill this gap: the uniform sampling property of the LCA protocol turns out to be a corollary of proposition 2 below. Thus, the LCA protocol could be used as a secure replacement for iterative sifting.

On the one hand, we suggest using the sifting part of LCA protocol. To be clear about the details of the sifting scheme, we have written it out in our notation in protocol 3. On the other hand, we find that the PE part of the LCA protocol is unnecessarily complicated and inefficient: it needs randomness injection for part of the sample choice, and it requires the estimation of two instead of one error rate. What if, instead, LCA sifting is followed by SBPE, i.e., only the error rate on the Z-agreements is determined? The critical question is whether this would still lead to uniform sampling. As the following propositin shows, this is indeed the case.

Proposition 2. *The combination of LCA sifting (protocol 3) and SBPE (protocol 2) samples uniformly. In other words, the LCA sifting protocol satisfies*

$$P_\Theta(\vartheta) = P_\Theta(\vartheta') \quad \forall\, \vartheta, \vartheta' \in \{0, 1\}^l_k. \tag{21}$$

In constrast to protocols that use randomness injection for the sample choice, the uniform sampling property is non-trivial to prove for LCA sifting with SBPE. We prove proposition 2 in appendix C (see the corollary of proposition 8). This shows that the combination of LCA sifting and SBPE is secure and can therefore be used to replace iterative sifting[12]. For protocols that use these subroutines, the abort probability p_{abort} of the sifting step is important because it affects the key rate of the QKD protocol. We calculate p_{abort} in appendix C as well (proposition 8).

This is good news for efficiency, as no randomness injection is required for the choice of the sample. Since this random sample choice would need to be communicated between Alice and Bob in an authenticated way, this also uses up less secret key from the initial key pool (see [23] for a discussion of the key cost of classical

[11] This can be seen from inequality (6), for example.

[12] This also establishes uniform sampling for the whole LCA protocol (with the PE protocol with randomness injection instead of SBPE). This is because the PE protocol of LCA can now be seen as a two-stage random sampling without replacement, where in both stages, the sampling probabilities are uniform. This leads to overall uniform sampling.

postprocessing). One can see in figure 2 that in the finite-key regime, this also leads to a larger raw key. Together with proposition 3, which we will discuss in section 6, this also establishes security of the protocol in the finite-key regime. In contrast, the original work of LCA [15] only establishes asymptotic security.

Suggestion. Use LCA sifting (protocol 3) and SBPE (protocol 2).

Let us briefly remark about the efficiency LCA sifting in comparison to that of iterative sifting. They differ in that LCA sifting has no communication during the loop phase, see property (3) above. The question is whether this necessarily means that the efficiency is strongly reduced in comparison with iterative sifting.

Protocol 3. The Lo–Chau–Ardehali (LCA) sifting protocol.

LCA sifting

Protocol parameters: $n,\ k,\ m \in \mathbb{N}_+$ with $m \geqslant n + k \in \mathbb{N}_+$ and $p_x,\ p_z \in [0,\ 1]$ with $p_x + p_z = 1$.
Output: For $l = n + k$, the outputs are:
Alice: l-bit string $(s_i)_{i=1}^l \in \{0,\ 1\}^l$ (measurement outcomes, sifted) or $s = \perp$ (if the protocol aborts),
Bob: l-bit string $(t_i)_{i=1}^l \in \{0,\ 1\}^l$ (measurement outcomes, sifted) or $t = \perp$ (if the protocol aborts),
public: l-bit string $(\vartheta_i)_{i=1}^l \in \{0,\ 1\}^l$ with $\sum_i \vartheta_i = k$ (basis choices, sifted), where 0 means X-basis and 1 means Z-basis, or $\vartheta = \perp$ (if the protocol aborts).
Number of rounds: Fixed number m (protocol parameter)

The protocol

Loop phase: Steps 1–4 are repeated m times (round index $r = 1,\ldots,m$). Starting with round $r = 1$, Alice and Bob do the following:

Step 1: (Preparation): Alice prepares a qubit pair in a maximally entangled state.

Step 2: (Channel use): Alice uses the quantum channel to send one share of the qubit pair to Bob.

Step 3: (Random bit generation): Alice and Bob each (independently) generate a random classical bit a_r and b_r, respectively, where 0 is generated with probability p_x and 1 is generated with probability p_z.

Step 4: (Measurement): Alice measures her share in the X-basis (if $a_r = 0$) or in the Z-basis (if $a_r = 1$), and stores the outcome in a classical bit y_r. Likewise, Bob measures his share in the X-basis (if $b_r = 0$) or in the Z-basis (if $b_r = 1$), and stores the outcome in a classical bit y'_r.

Final phase: The following steps are performed in a single run:

Step 5: (Quota Check): Alice and Bob determine the sets

$$u(m) = \{r \in [m] | a_r = b_r = 0\},$$
$$v(m) = \{r \in [m] | a_r = b_r = 1\}.$$

They check whether the quota condition ($u(m) \geqslant n$ and $v(m) \geqslant k$) holds. If it holds, they proceed with Step 6. Otherwise, they abort.

Step 6: (Random Discarding): Alice and Bob choose a subset $u \subseteq u(m)$ of size k at random, i.e. each subset of size k is equally likely to be chosen. Analogously, they choose a subset $v \subseteq v(m)$ of size k at random. Then they discard the bits a_r, b_r, y_r and y'_r for which $r \notin u \cup v$.

Step 7: (Order-preserving relabeling): Let r_i be the ith element of $u \cup v$. Then Alice determines $(s_i)_{i=1}^l \in \{0,\ 1\}^l$, Bob determines $(t_i)_{i=1}^l \in \{0,\ 1\}^l$ and together they determine $(\vartheta_i)_{i=1}^l \in \{0,\ 1\}^l$, where for every $i \in [l]$,

$$s_i = y_{r_i}, \quad t_i = y'_{r_i}, \quad \vartheta_i = a_{r_i} (= b_{r_i}).$$

Step 8: (Output): Alice locally outputs $(s_i)_{i=1}^l$, Bob locally outputs $(t_i)_{i=1}^l$ and they publicly output $(\vartheta_i)_{i=1}^l$.

We define the efficiency η of a sifting protocol as

$$\eta = \frac{R}{M}, \tag{22}$$

where R is the random variable of the number of rounds that are kept after sifting and M is the random variable of the total number of rounds performed in the loop phase of the protocol. We explain this in more detail in appendix D. A plot of the expected efficiency for iterative sifting and for LCA sifting is shown in figure 3 for the special case of symmetric probabilities $p_x = p_z$ and identical quota $n = k$ (this special case is computationally much easier to calculate; for other choices, the computation becomes very hard). We find that iterative sifting is

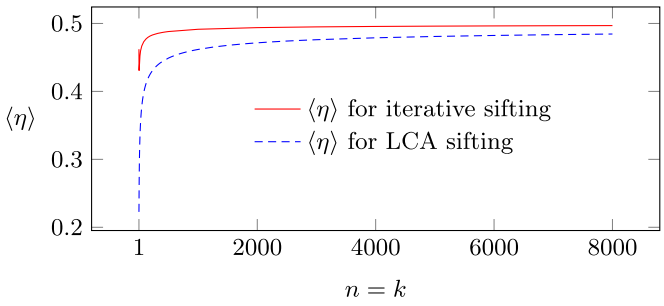

Figure 3. Efficiency comparison of the two sifting protocols. The plots show lower bounds on the expected efficiencies for symmetric probabilities $p_x = p_z = 1/2$ and for identical quotas $n = k$. The solid red curve shows a lower bound on the expected value of the efficiency for the iterative sifting protocol as a function of $n = k$. For the LCA sifting protocol, an optimization over the additional parameter m has been made for each value of $n = k$.

more efficient, as expected, but the difference between the two efficiencies becomes insignificant for practically relevant quota sizes n and k.

6. Formal criteria for good sifting

In section 3, we have seen that iterative sifting leads to problems. In section 5, we showed that these problems can be avoided by using LCA sifting (protocol 3) and SBPE (protocol 2). In this section, we give a more complete answer to the question of how these problems can be avoided by presenting two simple formal criteria that are sufficient for a sifting protocol to lead to a correct PE. More precisely, we describe two formal properties of the state produced by a sifting protocol which guarantee that if the protocol is followed by SBPE (protocol 2), then inequality (6) holds. As indicated in the introduction, the two properties take the form of equalities, see equations (1) and (2). We prove the sufficiency of these two criteria by deriving (6) from them in proposition 3 below.

In order to state the two criteria and the random variable Λ_{key} in (6) formally, we need to define a certain kind of quantum state $\rho_{A^l B^l \Theta^l}$ associated with a sifting protocol. To explain what this state is, we explain what the state $\rho_{A^l B^l \Theta^l}$ is like for LCA sifting. It is a state that is best described in a variation of the protocol. Suppose that Alice and Bob run the protocol, but they skip the measurement in every round. Instead, they keep each qubit system in their lab without modifying its state. With current technology, this is practically impossible, but since $\rho_{A^l B^l \Theta^l}$ is a purely mathematical construct, we do not worry about the technical feasibility. Notice that Alice and Bob still make basis choices, compare them and discard rounds—they just do not actually perform the measurements. Let us compare the output of this modified protocol with the output of the original protocol:

	Original protocol	Modified protocol
Alice:	l bits $s = (s_i)_{i=1}^l$	l-qubit state ρ_{A^l}
Bob:	l bits $t = (t_i)_{i=1}^l$	l-qubit state ρ_{B^l}
Public:	l bits $\vartheta = (\vartheta_i)_{i=1}^l$	l bits $\vartheta = (\vartheta_i)_{i=1}^l$

Hence, if we model the classical bit string ϑ as the state of a classical register Θ^l, we can say that the output of the modified protocol is a quantum–quantum-classical state $\rho_{A^l B^l \Theta^l}$. More generally, the state $\rho_{A^l B^l \Theta^l}$ associated with a sifting protocol is its output state in the case where all the measurements are skipped.

This state still carries all the probabilistic information of the original protocol. To see this, let $\mathbb{X} = \{\mathbb{X}_0, \mathbb{X}_1\}$ and $\mathbb{Z} = \{\mathbb{Z}_0, \mathbb{Z}_1\}$ be the POVMs describing Alice's X- and Z-measurement, let $\mathbb{X}' = \{\mathbb{X}'_0, \mathbb{X}'_1\}$ and $\mathbb{Z}' = \{\mathbb{Z}'_0, \mathbb{Z}'_1\}$ be the POVMs describing Bob's X- and Z-measurement, and let $\mathbb{M} = \{\mathbb{M}_0, \mathbb{M}_1\}$ be the projective measurement on Θ with respect to which the state of the register Θ is diagonal. Define the operators

$$\begin{aligned} &\mathbb{O}_0 = \mathbb{X}_0, \mathbb{O}_1 = \mathbb{X}_1, \mathbb{O}_2 = \mathbb{Z}_0, \mathbb{O}_3 = \mathbb{Z}_1, \\ &\mathbb{O}'_0 = \mathbb{X}'_0, \mathbb{O}'_1 = \mathbb{X}'_1, \mathbb{O}'_2 = \mathbb{Z}'_0, \mathbb{O}'_3 = \mathbb{Z}'_1. \end{aligned} \tag{23}$$

Then, the probability distribution over the output of the protocol is

$$P_{ST\Theta}(s, t, \vartheta) = \operatorname{tr}(\Pi(s, t, \vartheta)\rho_{(AB\Theta)^l}), \tag{24}$$

where $\rho_{(AB\Theta)^l}$ is the same state as $\rho_{A^lB^l\Theta^l}$, but with the registers reordered in the obvious way, and where

$$\Pi(s,\, t,\, \vartheta) = \bigotimes_{i=1}^{l} (\mathbb{O}_{2\vartheta_i+s_i} \otimes \mathbb{O}'_{2\vartheta_i+t_i} \otimes \mathbb{M}_{\vartheta_i}). \tag{25}$$

With the state $\rho_{A^lB^l\Theta^l}$ associated with a sifting protocol at hand, it is easy to define the random variable Λ_{key} associated with the protocol. The relevant probability space is the discrete probability space $(\Omega_{ZZ'\Theta},\, P_{ZZ'\Theta})$, where $\Omega_{ZZ'\Theta}$ is the sample space

$$\Omega_{ZZ'\Theta} = \{0,\ 1\}^l \times \{0,\ 1\}^l \times \{0,\, 1\}_k^l \tag{26}$$

and where $P_{ZZ'\Theta}$ is the probability mass function

$$\begin{aligned} P_{ZZ'\Theta}:\ \ \Omega_{ZZ'\Theta} &\to [0,\, 1] \\ (z,\, z',\, \vartheta) &\mapsto \operatorname{tr}\left(\left(\bigotimes_{i=1}^{l} \mathbb{Z}_{z_i}\right) \otimes \left(\bigotimes_{i=1}^{l} \mathbb{Z}'_{z'_i}\right) \otimes \left(\bigotimes_{i=1}^{l} \mathbb{M}_{\vartheta_i}\right) \rho_{A^lB^l\Theta^l}\right). \end{aligned} \tag{27}$$

The probability mass function $P_{ZZ'\Theta}$ corresponds to a *Gedankenexperiment* in which Alice and Bob measure *all* qubits in the Z-basis.

Now we are able to formally say what the random variable Λ_{key} of a sifting protocol is. Let $\rho_{A^lB^l\Theta^l}$ be the state associated with the sifting protocol, let $(\Omega_{ZZ'\Theta},\, P_{ZZ'\Theta})$ be the probability space as in equations (26) and (27). Then Λ_{key} is the random variable

$$\begin{aligned} \Lambda_{\text{key}}:\ \ \Omega_{ZZ'\Theta} &\to [0,\, 1] \\ (z,\, z',\, \vartheta) &\mapsto \frac{1}{n}\sum_{i=1}^{N} (1 - \vartheta_i)(z \oplus z'), \end{aligned} \tag{28}$$

which is the *key bit error rate*. Analogously, we have the *test bit error rate*

$$\begin{aligned} \Lambda_{\text{test}}:\ \ \Omega_{ZZ'\Theta} &\to [0,\, 1] \\ (z,\, z',\, \vartheta) &\mapsto \frac{1}{k}\sum_{i=1}^{l} \vartheta_i (z \oplus z'). \end{aligned} \tag{29}$$

This allows us to formally define the tail probability p_{tail}. We define it via the same formula as in (4), which we repeat here for the reader's convenience:

$$p_{\text{tail}}(\mu) = P[\Lambda_{\text{key}} \geqslant \Lambda_{\text{test}} + \mu | \Lambda_{\text{test}} \leqslant q_{\text{tol}}]. \tag{4}$$

The difference is that now, we have formally defined all the components of the equality. The following proposition states the tail probability bound in a formal way.

Proposition 3 (**Tail probability estimate**). *Let $\rho_{A^lB^l\Theta^l}$ be a density-operator of a system $A^lB^l\Theta^l$ where* A *and* B *are qubit systems and Θ is a classical system, let $\{\mathbb{Z}_0,\, \mathbb{Z}_1\}$ and $\{\mathbb{Z}'_0,\, \mathbb{Z}'_1\}$ be POVMs on the quantum systems* A *and* B*, respectively, let $\{\mathbb{M}_0,\, \mathbb{M}_1\}$ be the read-out measurement of the classical system Θ, let $\Lambda_{\text{key}},\ \Lambda_{\text{test}}$ be random variables on the discrete probability space $(\Omega_{ZZ'\Theta},\, P_{ZZ'\Theta})$ as defined in equations (26)–(29) and let p_{tail} be as in equation (4). Let $\rho_{A^lB^l}$ and ρ_{Θ^l} denote the according reduced states of $\rho_{A^lB^l\Theta^l}$ and P_Θ denote the according marginal of $P_{ZZ'\Theta}$. If the two conditions*

$$P_\Theta(\vartheta) = P_\Theta(\vartheta') \quad \forall\, \vartheta,\, \vartheta' \in \{0,\, 1\}_k^l \quad \text{and} \tag{1}$$

$$\rho_{A^lB^l\Theta^l} = \rho_{A^lB^l} \otimes \rho_{\Theta^l} \tag{2}$$

hold, then

$$p_{\text{tail}}(\mu) \leqslant \frac{\exp\left(-2\frac{kn}{l}\frac{k}{k+1}\mu^2\right)}{p_{\text{pass}}}, \tag{6}$$

where

$$p_{\text{pass}} = P[\Lambda_{\text{test}} \leqslant q_{\text{tol}}]. \tag{7}$$

We prove proposition 3 in appendix E. The formulation of proposition 3 allows us to see the formal requirements on a sifting protocol to lead to a correct PE when followed by SBPE: condition (1) is exactly the statement that the sampling probability does not depend on the sample, i.e. the protocol leads to uniform sampling. There is one thing that we want to point out here: while it is sufficient for the sampling probabilities to be the inverse of the number of possible samples, i.e.

$$P_\Theta(\vartheta) = \frac{1}{|\{0,1\}^l_k|} = \binom{l}{k}^{-1} \quad \forall\, \vartheta \in \{0,1\}^l_k, \tag{30}$$

condition (1) is strictly weaker. In the case where there is a non-zero probability that the protocol aborts during the sifting phase (as it is the case for LCA sifting), the sampling probabilities do not add up to 1 but rather to $1 - p_{\text{abort}}$, where p_{abort} is the probability that the protocol aborts during the sifting phase.

Condition (2) is the formal statement of what it means for a protocol that the basis choice register is uncorrelated with Alice's and Bob's qubits before measuring. Proposition 3 states that if these two conditions are satisfied, then the correlation test of the SBPE protocol leads to the right conclusion. Hence, these are the two conditions that a sifting protocol needs to satisfy in order to be a good sifting protocol.

We point out that the digression to a classical probability space, equations (26)–(29) and (4), is a mere change of notation. However, the fact that it is possible to express proposition 3 in terms of a classical probability space shows that this part of a QKD security analysis is purely classical.

7. Conclusion

In recent years QKD has emerged as a commercial technology, with the prospect of global QKD networks on the horizon [19]. All QKD implementations have finite size, and yet only recently has finite-key analysis approached mathematical rigor [3–6, 8–11]. In this work, we showed that further modifications of the protocols and/or their analysis are needed to make finite-key analysis rigorous.

We pointed out that *sifting*—a stage of QKD that is often overlooked with respect to security analysis—is actually crucial for security. A carelessly designed sifting subroutine can jeopardize the security of an otherwise reliable protocol. We found that *iterative sifting*, a sifting protocol that has both been proposed theoretically [3, 4, 8, 9] and been implemented experimentally [12], violates two assumptions in the typical security analysis. We showed how the violation of these assumptions can be exploited by an eavesdropper, leading to intercept-resend attacks with unexpectedly low error rates (see figure 1).

We presented an alternative scheme, LCA sifting and SBPE, and proved that it solves the two problems. We derived an expression for its abort probability and therefore provided everything that is needed for its future use as a subroutine. We argued that this scheme is more economical and efficient than some other other previously proposed protocols, as it does not require an additional random seed for the sample and at the same time allows for asymmetric basis choice probabilities. As we explained, the latter allows for a significantly higher sifting efficiency [15].

We gave the precise mathematical form of the two assumptions that are needed for secure sifting in equations (1) and (2). In doing so, we have provided a guide for the construction of future protocols: when designing a sifting protocol, one just needs to check these two conditions in order to make sure that the usual analysis of the PE based on inequality (6) is correct and the protocol is secure. This may require a mathematical model for the state $\rho_{A^lB^l\Theta^l}$ or for the probabilities of the output strings $(\vartheta_i)_{i=1}^l$, $(s_i)_{i=1}^l$ and $(t_i)_{i=1}^l$ generated by the sifting protocol. Such models are rarely provided in the literature. In the case of iterative sifting, the absence of such a model to check the desired properties has led to a wrong security analysis.

This points to a deeper problem in QKD security analysis: there is often a gap between the physical protocols that are written down as instructions for Alice and Bob and the mathematics of the security proof. This is not a purely pedantic issue, but rather a very practical one which can be exploited by eavesdroppers. In the future, we advocate that each step in the physical QKD protocol be explicitly mathematically modeled. In particular, we emphasize that sifting protocols must be proved to (rather than assumed to) satisfy the desired assumptions of the analysis. We believe our work will ultimately inspire more complete security proofs of finite-size QKD.

Acknowledgments

We would like to thank Marco Tomamichel, David Elkouss and Jędrzej Kaniewski for insightful discussions. CP and SW are funded by Singapore's MOE Tier 3A Grant MOE2012-T3-1-009, and STW, Netherlands. PJC and NL are supported by Industry Canada, Sandia National Laboratories, NSERC Discovery Grant, and Ontario Research Fund (ORF).

Appendix

Conventions

We make some notational conventions for the appendix (in addition to the ones we made in section 2). For the iterative sifting protocol as in protocol 1, we denote by N_x the random variable of the number of X-agreements,

and analogously, N_z and N_d are the random variables of the number of Z-agreements and disagreements in the loop phase, respectively. We write events as logical statements of the random variables, e.g. $\Theta = 110 \wedge N_x \geqslant 2$ is the event in which the protocol runs with more than two X-agreements and produces the output $\vartheta = 110$, and its probability is given by $P[\Theta = 110 \wedge N_x \geqslant 2]$. In cases where all involved random variables have fixed values, we occasionally write expressions in terms of probability mass functions instead of in terms of probability weights of events (as we have done it in the main article), e.g. we write

$$P_{\Theta N_x N_z N_d}(\vartheta, n_x, n_z, n_d) \coloneqq P[\Theta = \vartheta, N_x = n_x, N_z = n_z, N_d = n_d]. \tag{31}$$

In cases with inequalities, it is however shorter to use the event notation, e.g.

$$P[A_1 \neq B_1] = P_{A_1 B_1}(0, 1) + P_{A_1 B_1}(1, 0). \tag{32}$$

We will use whatever notation we find more appropriate in each case.

Appendix A. Sampling probability calculation for iterative sifting

In this appendix, we prove proposition 1, i.e. we calculate the sampling probabilities $P_\Theta(\vartheta)$ for iterative sifting with $n = 1$ and $k = 2$ and find $P_\Theta(110) = g_z^2$ and $P_\Theta(101) = P_\Theta(011) = (1 - g_z^2)/2$, where $g_z = p_z^2/(p_z^2 + p_x^2)$.

Proof of proposition 1. We first write out the sequence of equalities that lead to the claim. We explain each equality below. The sequence of equalities looks as follows:

$$P_\Theta(110) = \sum_{n_x=1}^{\infty} \sum_{n_z=2}^{\infty} \sum_{n_d=0}^{\infty} P_{\Theta N_x N_z N_d}(110, n_x, n_z, n_d), \tag{A1}$$

$$= \sum_{n_z=2}^{\infty} \sum_{n_d=0}^{\infty} P_{\Theta N_x N_z N_d}(110, 1, n_z, n_d), \tag{A2}$$

$$= \sum_{n_z=2}^{\infty} \sum_{n_d=0}^{\infty} p_x^2 (p_z^2)^k (2p_x p_z)^d \binom{n_z + n_d}{n_d}, \tag{A3}$$

$$= g_z^2, \quad \text{where } g_z = \frac{p_z^2}{p_z^2 + p_x^2}. \tag{A4}$$

Equation (A1) is just stating that P_Θ is the marginal of $P_{\Theta N_x N_z N_d}$. The ranges of the sums can be explained as follows. The iterative sifting protocol always runs until there have been at least n X-agreements and at least k Z-agreements. Therefore

$$P_{\Theta N_x N_z N_d}(\theta, n_x, n_z, n_d) = 0 \quad \text{if } n_x < n \text{ or } n_z < k. \tag{A5}$$

In our case, $n = 1$ and $k = 2$, hence the limits of the sums.

Equation (A2) follows from

$$P_{\Theta N_x N_z N_d}(110, n_x, n_z, n_d) = 0 \quad \text{for } n_x \geqslant 2. \tag{A6}$$

One can see (A6) as follows: if $N_x \geqslant 2$, then necessarily $N_z = 2$, because $N_x > n \wedge N_z > k$ is impossible in iterative sifting (the loop phase of the protocol is terminated as soon as both quota are met). This means that during the random discarding, no Z-agreement gets discarded. Moreover, if $N_x \geqslant 2$, then the last round of the loop phase must be a Z-agreement. Since this Z-agreement is not discarded, we have that Θ must necessarily end in a 1 if $N_x \geqslant 2$, so $\Theta = 110$ is impossible in that case.

To see why equation (A3) holds, note that the event

$$\Theta = 110 \wedge N_x = 1 \wedge N_z = n_z \wedge N_d = n_d \tag{A7}$$

consists of all runs of the protocol in which one X-agreement, n_z Z-agreements and n_d disagreements occurred, and where the X-agreement was the last round of the loop phase. This is because in every such run, one necessarily ends up with $\Theta = 110$, and if $\Theta = 110$, then the last round of the loop phase must be an X-agreement. There are $\binom{n_z + n_d}{n_d}$ such runs, and each of them has the probability $p_x^2 (p_z^2)^{n_z} (2p_x p_z)^{n_d}$, and therefore

$$P_{\Theta N_x N_z N_d}(110, 1, n_z, n_d) = \binom{n_z + n_d}{n_d} p_x^2 (p_z^2)^{n_z} (2p_x p_z)^{n_d}. \tag{A8}$$

This explains equation (A3). Finally, equation (A4) is just an evaluation of the expression in the line above. This shows $P_\Theta(110) = g_z^2$.

It remains to be shown that $P_\Theta(101) = P_\Theta(011) = (1 - g_z^2)/2$. In analogy to the above, it holds that

$$P_\Theta(101) = \sum_{n_x=1}^{\infty}\sum_{n_z=2n_d=0}^{\infty}\sum P_{\Theta N_x N_z N_d}(101, n_x, n_z, n_d), \tag{A9}$$

$$= \sum_{n_x=2n_d=0}^{\infty}\sum P_{\Theta N_x N_z N_d}(101, n_x, 2, n_d). \tag{A10}$$

Equation (A9) is, in analogy to equation (A1), stating that P_Θ is the marginal of $P_{\Theta N_x N_z N_d}$, and the same argumentation for the limits of the sums applies. Equation (A10) is explained by a similar reasoning as for equation (A2): it follows from

$$P_{\Theta N_x N_z N_d}(101, n_x, n_z, n_d) = 0 \quad \text{for } n_z \geqslant 3. \tag{A11}$$

For equation (A11), note that if $N_z \geqslant 3$, then $N_x = 1$ because $N_x > n \wedge N_z > k$ is impossible in iterative sifting. Thus, no X-agreement gets discarded. Moreover, if $N_z \geqslant 3$, then the last round of the loop phase must be an X-agreement. Since this X-agreement is not discarded, Θ necessarily ends in a 0 if $N_z \geqslant 3$, so $\Theta = 101$ is impossible in this case.

Analogously, it holds that

$$P_\Theta(011) = \sum_{n_x=1}^{\infty}\sum_{n_z=2n_d=0}^{\infty}\sum P_{\Theta N_x N_z N_d}(011, n_x, n_z, n_d), \tag{A12}$$

$$= \sum_{n_x=2n_d=0}^{\infty}\sum P_{\Theta N_x N_z N_d}(011, n_x, 2, n_d). \tag{A13}$$

The next step is to realize that for every $n_x \geqslant 2$ and for every $n_d \in \{0, 1, 2, \ldots\}$, it holds that

$$P_{\Theta N_x N_z N_d}(101, n_x, 2, n_d) = P_{\Theta N_x N_z N_d}(011, n_x, 2, n_d). \tag{A14}$$

This is because the event

$$(\Theta = 101, N_x = n_x, N_z = 2, N_d = n_d) \tag{A15}$$

and the event

$$(\Theta = 011, N_x = n_x, N_z = 2, N_d = n_d) \tag{A16}$$

consist of equally many histories of the protocol, and each of these histories has the same probability. Equations (A10), (A13) and (A14) imply $P_\Theta(101) = P_\Theta(011)$. Since $P_\Theta(011) + P_\Theta(101) + P_\Theta(110) = 1$ and $P(110) = g_z^2$, it holds that $P_\Theta(011) = P_\Theta(101) = (1 - g_z^2)/2$ as claimed. □

Appendix B. Error rate calculations for the attacks on iterative sifting

B.1. Attack that exploits non-uniform sampling

Here, we calculate the expected error rate for the attack on iterative sifting which exploits non-uniform sampling, as explained in section 4.1. We first recall the relevant conventions that we made in the main article. The iterative sifting protocol is described in protocol 1. Eve performs an intercept-resend attack during the loop phase of the protocol. In the first round, she attacks in the X-basis, and in all the other rounds of the loop phase, she attacks in the Z-basis. We defined the error rate in equation (17) in the main article, namely

$$E = \frac{1}{l}\sum_{i=1}^{l} S_i \oplus T_i. \tag{B1}$$

Moreover, recall that we assume that the X- and Z-basis is the same for Alice, Bob and Eve, and that they are mutually unbiased. This way, if Alice and Bob measure in the same basis, but Eve measures in the other basis, then Eve introduces an error probability of 1/2 on this qubit.

The calculation of $\langle E\rangle$ for this attack goes as follows. We first make a split:

$$\langle E\rangle = \sum_{\vartheta} P[\Theta = \vartheta]\langle E|\Theta = \vartheta\rangle \tag{B2}$$

$$= \underbrace{P[\Theta = 01]\langle E|\Theta = 01\rangle}_{\Delta_x} + \underbrace{P[\Theta = 10]\langle E|\Theta = 10\rangle}_{\Delta_z}. \tag{B3}$$

We have that

$$\begin{aligned}\Delta_x = \sum_{n_x=1}^{\infty} (&P[\Theta = 01 \wedge N_x = n_x \wedge A_1 = B_1 = 0]\langle E|\Theta = 01 \wedge N_x = n_x \wedge A_1 = B_1 = 0\rangle \\ &+ P[\Theta = 01 \wedge N_x = n_x \wedge A_1 \neq B_1]\langle E|\Theta = 01 \wedge N_x = n_x \wedge A_1 \neq B_1\rangle \\ &+ \underbrace{P[\Theta = 01 \wedge N_x = n_x \wedge A_1 = B_1 = 1]}_{0}\langle E|\Theta = 01 \wedge N_x = n_x \wedge A_1 = B_1 = 1\rangle).\end{aligned} \tag{B4}$$

The third summand on the right-hand side of equation (B4) vanishes because $\Theta = 01$ is impossible if Alice and Bob have a z-agreement in the first round of the loop phase. The event

$$\Theta = 01 \wedge N_x = n_x \wedge A_1 = B_1 = 0 \tag{B5}$$

consists of all histories of the protocol in which Alice and Bob have an x-agreement in the first round and n_x X-agreements in total. Infinitely many such histories are possible because an arbitrary number of disagreements is possible. We express the probability of the event (B5) as the marginal of the probability of the event

$$\Theta = 01 \wedge N_x = n_x \wedge A_1 = B_1 = 0 \wedge N_d = n_d. \tag{B6}$$

The event (B6) consists of $\binom{n_x + n_d + 1}{n_d}$ histories of the protocol, and each history has the probability $(p_x^2)^{n_x} p_z^2 (2p_x p_z)^{n_d}$. Therefore

$$P[\Theta = 01 \wedge N_x = n_x \wedge A_1 = B_1 = 0] = \sum_{n_d=0}^{\infty} P[\Theta = 01 \wedge N_x = n_x \wedge A_1 = B_1 = 0 \wedge N_d = n_d], \tag{B7}$$

$$= \sum_{n_d=0}^{\infty} (p_x^2)^{n_x} p_z^2 (2p_x p_z)^{n_d} \binom{n_x + n_d - 1}{n_d}. \tag{B8}$$

Moreover, we have that

$$\langle E|\Theta = 01 \wedge N_x = n_x \wedge A_1 = B_1 = 0\rangle = \frac{1}{4}\left(1 - \frac{1}{n_x}\right). \tag{B9}$$

The validity of (B9) can be seen as follows. On the second bit of S and T, there is no error because it comes from a round in which all parties have measured in the Z-basis. Hence, the left had side of (B9) is the probability of getting an error on the first bit of S and T, divided by the total number of bits, 2. Hence, we need to determine the error probability of the first bit. If $N_x = 1$, then the first bit comes from the first round of the loop phase, in which Alice, Bob and Eve have measured in the X-basis and hence, there is no error. However, for $N_x = n_x$, the first bit of S and T is chosen at random from one of the n_x X-agreements. In only one of these n_x rounds, Eve has measured in the X-basis, and in $n_x - 1$ rounds, she measured in the Z-basis. Hence, the probability that Eve measured in the wrong basis on the first bit of S and T is $(n_x - 1)/n_x$, and therefore the error probability of the first bit is $1/2 \cdot (n_x - 1)/n_x$. Thus

$$\langle E|\Theta = 01 \wedge N_x = n_x \wedge A_1 = B_1 = 0\rangle = \frac{1}{2} \cdot \frac{1}{2}\left(\frac{n_x}{n_x - 1}\right), \tag{B10}$$

$$= \frac{1}{4}\left(1 - \frac{1}{n_x}\right). \tag{B11}$$

Similarly, we get

$$P[\Theta = 01 \wedge N_x = n_x \wedge A_1 \neq B_1] = \sum_{n_d=0}^{\infty} (p_x^2)^{n_x} p_z^2 (2p_x p_z)^{n_d} \binom{n_x + n_d - 1}{n_x} \tag{B12}$$

and

$$\langle E|\Theta = 01 \wedge N_x = n_x \wedge A_1 \neq B_1\rangle = \frac{1}{4}. \tag{B13}$$

Taking equations (B8), (B9), (B12), (B13) together we get that

$$\Delta_x = \frac{1}{4} \sum_{n_x=1}^{\infty} \sum_{n_d=0}^{\infty} (p_x^2)^{n_x} p_z^2 (2p_x p_z)^{n_d} \left(\binom{n_x + n_d - 1}{n_d}\left(1 - \frac{1}{n_x}\right) + \binom{n_x + n_d - 1}{n_x}\right). \tag{B14}$$

In a similar way, we get

$$\Delta_z = \frac{1}{4} \sum_{n_z=1}^{\infty} \sum_{n_d=0}^{\infty} p_x^2 (p_z^2)^{n_z} (2p_x p_z)^{n_d} \left(\binom{n_z + n_d - 1}{n_d} + \binom{n_z + n_d - 1}{n_d}\left(1 + \frac{1}{n_x}\right)\right). \tag{B15}$$

Equations (B3), (B14), (B15) taken together result in

$$\langle E\rangle = \sum_{n_d=0}^{\infty}(2p_xp_z)^{n_d}\left(\sum_{n_x=1}^{\infty}(p_x^2)^{n_x}p_x^2\left(\binom{n_x+n_d-1}{n_d}\left(1-\frac{1}{n_x}\right)+\binom{n_x+n_d-1}{n_x}\right)\right.$$
$$\left.+\sum_{n_z=1}^{\infty}p_x^2(p_z^2)^{n_z}\left(\binom{n_z+n_d-1}{n_z}+\binom{n_z+n_d-1}{n_d}\left(1+\frac{1}{n_z}\right)\right)\right). \tag{B16}$$

Figure 1 in the main article shows a plot of $\langle E\rangle$ as in (B16) as a function of p_x. As one can see, $\langle E\rangle$ achieves a minimum of $\langle E\rangle \approx 22.8\%$ for $p_x \approx 0.73$.

B.2. Attack that exploits basis-information leak

Now we calculate the expected error rate of iterative sifting for the attack which exploits basis-information leak as described in section 4.2. As before, let $\langle E\rangle$ be the expected value of the error rate as defined in equation (17). Again, we assume that the X- and Z-basis are the same for Alice, Bob and Eve and that they are mutually unbiased. Recall the strategy of Eve's intercept-resend attack: before the first round of the loop phase, Eve flips a fair coin. Let F be the random variable of the coin flip outcome and let 0 and 1 be its possible values. If $F=0$, then in the first round, Eve attacks in the X basis, and if $F=1$, she attacks in the Z-basis. In the subsequent rounds, she keeps attacking in that basis until Alice and Bob first reached a basis agreement. If it is an X-agreement (equivalent to $\Theta = 01$), Eve attacks in the Z-basis in all remaining rounds, and if it is a Z-agreement (equivalent to $\Theta = 10$), she attacks in the X-basis in all remaining rounds.

The calculation of $\langle E\rangle$ goes as follows:

$$\langle E\rangle = P_F(0)\langle E|F=0\rangle + P_F(1)\langle E|F=1\rangle, \tag{B17}$$
$$= \langle E|F=0\rangle, \tag{B18}$$
$$= \underbrace{P_\Theta(01)}_{1/2}\langle E|F=0\wedge\Theta=01\rangle + \underbrace{P_\Theta(10)}_{1/2}\underbrace{\langle E|F=0\wedge\Theta=10\rangle}_{1/4}. \tag{B19}$$

Equality (B17) is just a decomposition of $\langle E\rangle$ into conditional expectations. Equality (B18) follows from the fact that the problem is symmetric under the exchange of X and Z, i.e. under the exchange of 0 and 1. The only quantity that is not trivial to calculate in equation (B19) is the expected value of the error rate, given that Eve first measures in X and that the first basis agreement is an X-agreement. It is calculated as follows:

$$\langle E|F=0\wedge\Theta=01\rangle = \sum_{n_x=1}^{\infty}\langle E|F=0\wedge\Theta=01\wedge N_x=n_x\rangle\underbrace{P_{N_x|\Theta F}(n_x|01,0)}_{P_{N_x|\Theta}(n_x|01)}, \tag{B20}$$
$$= \sum_{n_x=1}^{\infty}\underbrace{\langle E|F=0\wedge\Theta=01\wedge N_x=n_x\rangle}_{\frac{n_x-1}{4n_x}}\underbrace{P_{N_x\Theta}(n_x,01)}_{\sum_{n_d=0}^{\infty}(p_x^2)^{n_x}p_z^2(2p_xp_z)^{n_d}\binom{n_x+n_d}{n_d}}\underbrace{\frac{1}{P_\Theta(01)}}_{2}, \tag{B21}$$
$$= \sum_{n_x=1}^{\infty}\frac{n_x-1}{2n_x}\sum_{n_d=0}^{\infty}(p_x^2)^{n_x}p_z^2(2p_xp_z)^{n_d}\binom{n_x+n_d}{n_d} \tag{B22}$$
$$= \frac{1}{4}(1-\ln 2), \tag{B23}$$

where ln denotes the logarithm to base e. Therefore

$$\langle E\rangle = \frac{1}{2}\frac{1}{4}(1-\ln 2)+\frac{1}{2}\frac{1}{4}, \tag{B24}$$
$$= \frac{2-\ln 2}{8}, \tag{B25}$$
$$\approx 16.3\%. \tag{B26}$$

B.3. Attack that exploits both problems

Here we present the error rate induced by the intercept-resend attack presented in section 4.4, which exploits both non-uniform sampling and basis information leak. Let us recall the attack strategy. In the first round of the loop phase of the iterative sifting protocol, she attacks in the X-basis. She keeps doing that in subsequent rounds until Alice and Bob announce a basis-agreement. If they announce an X-agreement, Eve attacks in the Z-basis in all the following rounds. Otherwise, she keeps attacking in the X-basis.

The calculation of the error rate is similar to the calculations done in appendices B.1 and B.2. We only show the result here:

$$\langle E\rangle = \sum_{n_z=1 n_d=0}^{\infty}\sum^{\infty} p_x^2 p_z^{2n_z} (2p_x p_z)^{n_d}\binom{n_z + n_d}{n_d}\frac{1}{4} + \sum_{n_x=1 n_d=0}^{\infty}\sum^{\infty} p_x^{2n_x} p_z^2 (2p_x p_z)^{n_d}\binom{n_x + n_d}{n_d}\frac{n_x - 1}{4n_x}. \tag{B27}$$

A plot of (B27) is shown in figure 1 as a function of p_x. As one can see, the expected error rate has a minimum of $\langle E\rangle \approx 15.8\%$ for $p_x \approx 0.57$. Hence, this combined attack on both problems performs much better than the one on non-uniform sampling alone (with a minimal expected error rate of ≈22.8%, see section 4.1) and even better than the attack on the basis information leak alone (with a minimal expected error rate of ≈16.3%, see section 4.2).

Appendix C. Sampling and abort probability calculation for LCA sifting

In this appendix, we derive the general form of the probabillity distribution $P_\Theta(\vartheta)$ for LCA sifting (protocol 3) as a function of the parameters n, k, m, p_x and p_z. This achieves two goals: firstly, it turns out that the sampling probability $P_\Theta(\vartheta)$ is independent of the sample $\vartheta \in \{0,1\}^l_k$, which shows that the protocol samples uniformly. Secondly, we calculate the abort probability $p_{\text{abort}} = P_\Theta(\perp)$. This abort probability influences the key rate of potential QKD protocols that use this protocol as a subroutine, which makes p_{abort} an important performance parameter of the protocol.

We start by describing in appendix C.1 how we think that proofs of sampling probabilities should be formalized and how the general strategy of our proof looks like. In appendices C.2–C.4, we show the proofs and finally derive P_Θ.

C.1. On probabilistic models of the protocol

LCA sifting gives rise to a set Ω of histories of the protocol. This set can be modeled as the set $\Omega = \Omega_{ABYY'STUV\Theta}$ of all tuples

$$\omega = (a, b, y, y', s, t, u, v, \vartheta), \tag{C1}$$

where each entry varies over all its possible values. There are finitely many such histories, and each of them as a probability associated with it. This can be expressed more formally in the language of discrete probability theory[13] by saying that Ω forms the sample space of a discrete probability space (Ω, P), on which a probability mass function P is defined such that $P(\omega)$ is the probability of a history ω. Note that by choosing $\Omega = \Omega_{AB\ldots\Theta}$, we also include impossible combinations of $a, b, \ldots, \vartheta$. For example, a history ω as in (C1) with $u = v$ is not possible, because u stands for the X-agreements chosen for the raw key and v stands for the Z-agreements chosen for the sample, and the two cannot coincide. This is not a problem for our model, because in this case, we simply have $P(\omega) = 0$.

In this probability theory language, the strings $a, b, \ldots, \vartheta$ are values that random variables $A, B, \ldots, \Theta$ can take. Random variables are maps from the sample space Ω to a set which is called the *range* or *codomain* of the random variable. For example, the random variable A is a map

$$\begin{aligned} A : \Omega &\to \mathcal{A} \\ \omega &\mapsto A(\omega), \end{aligned} \tag{C2}$$

where $\mathcal{A}$ is the codomain of A. We denote the codomains of random variables with calligraphic letters (except for the random variable Θ, whose codomain we denote by $\text{co}(\Theta)$). According to the protocol, we have

$$\mathcal{A} = \{0,1\}^m = \{(a_i)_{i=1}^m | a_i \in \{0,1\} \forall i \in [m]\}. \tag{C3}$$

In the case where we model

$$\Omega = \Omega_{ABYY'STUV\Theta} = \mathcal{A}\times\mathcal{B}\times\mathcal{Y}\times\mathcal{Y}'\times\mathcal{S}\times\mathcal{T}\times\mathcal{U}\times\mathcal{V}\times\text{co}(\Theta), \tag{C4}$$

the random variables are simply the (set-theoretic) projections on the respective components, e.g.

$$\begin{aligned} A : \Omega = \mathcal{A}\times\mathcal{B}\times\ldots\times\text{co}(\Theta) &\to \mathcal{A}, \\ (a, b, \ldots, \vartheta) &\mapsto a. \end{aligned} \tag{C5}$$

[13] By *discrete* probability theory, we mean probability theory with a discrete sample space Ω, i.e. where Ω is finite or countably infinite.

Then, the probability $P_A(a)$ that $A = a$ is given by

$$\begin{aligned} P_A\colon \mathcal{A} &\to [0,\,1] \\ a &\mapsto \sum_{\omega\in A^{-1}(a)} P(\omega) \\ &= \sum_{(b,y,\ldots,\vartheta)} P_{AB\ldots\Theta}(a,\,b,\ldots,\vartheta), \end{aligned} \tag{C6}$$

where we have written $P = P_{AB\ldots\Theta}$. This is because in the case where $\Omega = \Omega_{AB\ldots\Theta}$, P is simply the joint probability distribution of the random variables $A,\,B,\,\ldots,\,\Theta$.

Setting $(\Omega,\,P) = (\Omega_{AB\ldots\Theta},\,P_{AB\ldots\Theta})$ is sufficient to describe the probabilities of the random variables $A,\,B,\,\ldots,\,\Theta$ and functions thereof. For our purposes, however, this description is overloaded. We do not need to incorporate all the random variables $A,\,B,\,\ldots,\,\Theta$ in Ω and P. One reason is that some of the random variables are completely determined by some of the other random variables. For example, the string s of Alice's sifted measurement outcomes is completely determined by Alice's measurement outcomes a and the subsets u and v. In the probability theory language, this is expressed as the fact that the random variable S is a function of the random variables A, U and V,

$$S \equiv S(A,\,U,\,V), \tag{C7}$$

or more precisely

$$\begin{aligned} S : \mathcal{A}\times\mathcal{U}\times\mathcal{V} &\to \quad \mathcal{S} \\ (a,\,u,\,v) &\mapsto S(a,\,u,\,v) \end{aligned} \tag{C8}$$

and its probability distribution is given by

$$P_S(s) = \sum_{(a,u,v)\in S^{-1}(s)} P_{AUV}(a,\,u,\,v), \tag{C9}$$

$$= \sum_{\omega\in(S\circ A\times U\times V)^{-1}(s)} P(\omega). \tag{C10}$$

There are more such dependencies in our list of random variables:

$$T \equiv T(B,\,U,\,V), \tag{C11}$$

$$\Theta \equiv \Theta(U,\,V). \tag{C12}$$

Hence, setting

$$(\Omega,\,P) = (\Omega_{ABYY'UV},\,P_{ABYY'UV}) \tag{C13}$$

and using the dependencies (C7), (C11) and (C12) leads to an equally powerful description, but with a smaller probability space.

For our purposes, the space (C13) is still overloaded. We are only interested in the distribution P_Θ of Θ. According to (C12), the relevant probability space is $(\Omega_{UV},\,P_{UV})$, and Θ is a random variable

$$\begin{aligned} \Theta : \Omega_{UV} = \mathcal{U}\times\mathcal{V} &\to \mathrm{co}(\Theta), \\ (u,\,v) &\mapsto \vartheta(u,\,v). \end{aligned} \tag{C14}$$

Then, P_Θ is given by

$$\begin{aligned} P_\Theta : \mathrm{co}(\Theta) &\to \quad [0,\,1] \\ \vartheta \quad &\mapsto \sum_{(u,v)\in\Theta^{-1}(\vartheta)} P_{UV}(u,\,v). \end{aligned} \tag{C15}$$

It is difficult to write down the probability mass function P_{UV} directly. Instead, we will derive the probability mass function P_{ABUV} on the sample space Ω_{ABUV}, and arrive at the probability distribution P_{UV} via marginalization of P_{ABUV}:

$$P_{UV}(u,\,v) = \sum_{(a,b)\in\mathcal{A}\times\mathcal{B}} P_{ABUV}(a,\,b,\,u,\,v). \tag{C16}$$

Hence, the relevant probability space for our proof of uniform sampling of LCA sifting is the probability space $(\Omega_{ABUV},\,P_{ABUV})$.

C.2. Formalization of $(\Omega_{ABUV},\,P_{ABUV})$

According to what we said in the last subsection, the probability space that is relevant for our proof of uniform sampling of LCA sifting is the space $(\Omega_{ABUV},\,P_{ABUV})$, which describes the probabilities of the basis choice strings a and b of Alice and Bob, as well as the choices u and v of the rounds that are used for the raw key and for PE, respectively. We are going to formalize this space in this subsection.

We start by determining the sample space

$$\Omega_{ABUV} = \mathcal{A} \times \mathcal{B} \times \mathcal{U} \times \mathcal{V}. \tag{C17}$$

In the loop phase of the protocol, Alice and Bob generate basis choice strings $a = (a_i)_{i=1}^m \in \{0, 1\}^m$, $b = (b_i)_{i=1}^m \in \{0, 1\}^m$. This happens in every run, no matter whether Alice and Bob abort the protocol in the final phase. Hence

$$\mathcal{A} = \mathcal{B} = \{0, 1\}^m. \tag{C18}$$

In the final phase of the protocol, Alice and Bob do a quota check, in which they determine the rounds in which both measured in the X-basis (X-agreement) the rounds in which both measured in the Z-basis (Z-agreements). In the case where they had less than n X-agreements or less than k Z-agreements, they abort. In this case, Alice and Bob do not choose subsets u and v of their X- and Z-agreements, respectively. We model this by saying that in this case, $u = v = \perp$, where $\perp$ is just a symbol indicating that Alice and Bob abort. In the case where the quota check of the protocol is successful, Alice and Bob choose random subsets $u \subseteq u(m)$ of size n and $v \subseteq v(m)$ of size k. We represent these subets by bit strings $u \in \{0, 1\}_n^m$, $v \in \{0, 1\}_k^m$, where

$$\begin{aligned} \{0, 1\}_n^m &= \left\{ (u_i)_{i=1}^m \in \{0, 1\}^m \;\middle|\; \sum_{i=1}^m u_i = n \right\}, \\ \{0, 1\}_k^m &= \left\{ (v_i)_{i=1}^m \in \{0, 1\}^m \;\middle|\; \sum_{i=1}^m v_i = k \right\}. \end{aligned} \tag{C19}$$

They are to be interpreted as follows: for $u \in \{0, 1\}_n^m$ and $i \in [m]$, $u_i = 1$ means that i is contained in the subset $u \subseteq u(m)$, and $u_i = 0$ means that i is not contained, and likewise for $v \in \{0, 1\}_k^m$. The requirement that the subsets u and v have size n and k translates into the conditions that the string components sum up to n and k, respectively. Taking the two possibilities (the protocol aborts or the quota check is successful) together, we have that

$$\mathcal{U} = \{0, 1\}_n^m \bigcup \{\perp\}, \tag{C20}$$

$$\mathcal{V} = \{0, 1\}_k^m \bigcup \{\perp\}, \tag{C21}$$

and hence

$$\Omega_{ABUV} = \mathcal{A} \times \mathcal{B} \times \mathcal{U} \times \mathcal{V} = \{0, 1\}^m \times \{0, 1\}^m \times (\{0, 1\}_n^m \bigcup \{\perp\}) \times (\{0, 1\}_k^m \bigcup \{\perp\}). \tag{C22}$$

This is the sample space of the probability space $(\Omega_{ABUV}, P_{ABUV})$ that we are looking for.

Next, we determine the probability mass function P_{ABUV}. We can write

$$P_{ABUV}(a, b, u, v) = P_{AB}(a, b) P_{UV|AB}(u, v|a, b), \tag{C23}$$

where $P_{UV|AB}(u, v|a, b)$ is the probability that $U = u$ and $V = v$, conditioned on $A = a$ and $B = b$. The probability distribution $P_{AB}(a, b)$ is easily determined. Each bit a_i, b_i, $i \in [m]$ is generated independently at random and takes the value 0 with probability p_x and the value 1 with probability p_z. Hence

$$\forall (a, b) \in \mathcal{A} \times \mathcal{B}: \quad P_{AB}(a, b) = \prod_{i=1}^m p_x^{1-a_i} p_z^{a_i} p_x^{1-b_i} p_z^{b_i}, \tag{C24}$$

$$= p_x^{m-|a|} p_z^{|a|} p_x^{m-|b|} p_z^{|b|}, \tag{C25}$$

$$= p_x^{2m-|a|-|b|} p_z^{|a|+|b|}, \tag{C26}$$

where for a string $a \in \{0, 1\}^m$, we write

$$|a| := \sum_{i=1}^m a_i. \tag{C27}$$

The conditional probability distribution $P_{UV|AB}$ is a bit more tricky to write down. What is crucial for this conditional probability is whether the strings a and b have at least n X-agreements and at least k Z-agreements. We want to give this condition a formula as follows. Imagine Alice and Bob want to count their X- and Z-agreements. To do so, they can first determine the string $a \wedge b$, given by

$$a \wedge b := (a_i b_i)_{i=1}^m. \tag{C28}$$

The ith entry $a_i b_i$ of $a \wedge b$ is 1 if the corresponding bits a_i and b_i are both 1, i.e. if they had a Z-agreement, and 0 otherwise. Hence, to count their Z-agreements, they can sum up the components of $a \wedge b$:

$$\text{number of } Z\text{-agreements} = \sum_{i=1}^m a_i b_i = |a \wedge b|. \tag{C29}$$

Therefore, the condition that Alice and Bob had at least k Z-agreements can be expressed as

$$|a \wedge b| \geqslant k. \tag{C30}$$

Likewise, the condition that they had at least n X-agreements can be written as

$$|\bar{a} \wedge \bar{b}| \geqslant n, \tag{C31}$$

where for a string $a \in \{0, 1\}^m$, we write

$$\bar{a} = (1 - a_i)_{i=1}^m \in \{0, 1\}^m. \tag{C32}$$

Taken together, the quota check condition reads

$$|\bar{a} \wedge \bar{b}| \geqslant n \quad \text{and} \quad |a \wedge b| \geqslant k. \tag{C33}$$

In the case where condition (C33) is not satisfied, Alice and Bob abort, and therefore it must be that $(u, v) = (\perp, \perp)$. We can write this as

$$\begin{aligned} &\forall (a, b) \in \{0, 1\}^m \times \{0, 1\}^m \text{ such that } (|a \wedge b| < k \text{or } |\bar{a} \wedge \bar{b}| < n): \\ &\quad P_{UV|AB}(u, v|a, b) = \chi(u = v = \perp), \end{aligned} \tag{C34}$$

where χ is the *indicator* function, which evaluates to 1 if its argument is true and which evaluates to 0 if its argument is false.

For $(a, b) \in \{0, 1\}^m \times \{0, 1\}^m$ such that condition (C33) is satisfied, the conditional probability $P_{UV|AB}$ is a little more difficult to write down. In that case, both $u = \perp$ and $v = \perp$ are impossible. Moreover, only those $u \in \{0, 1\}_n^m$ are possible which are subsets of Alice and Bob's X-agreements, i.e. which satisfy

$$u_i = 1 \Longrightarrow a_i = b_i = 0 \quad \forall i \in [m]. \tag{C35}$$

Note that

$$\begin{aligned} &\forall (a, b, u) \in \{0, 1\}^m \times \{0, 1\}^m \times \{0, 1\}_n^m: \\ &\quad (u_i = 1 \Longrightarrow a_i = b_i = 0) \Longleftrightarrow |\bar{a} \wedge \bar{b} \wedge u| = n. \end{aligned} \tag{C36}$$

Hence, the condition that u is a subset of the X-agreements simply reads

$$|\bar{a} \wedge \bar{b} \wedge u| = n, \tag{C37}$$

and likewise, the condition that v is a subset of the Z-agreements reads

$$|a \wedge b \wedge v| = k. \tag{C38}$$

Hence, in the case where (C33) holds, only those $(u, v) \in \{0, 1\}_n^m \times \{0, 1\}_k^m$ are possible for which

$$|\bar{a} \wedge \bar{b} \wedge u| = n \quad \text{and} \quad |a \wedge b \wedge v| = k. \tag{C39}$$

We can combine the two conditions in a single formula:

$$\forall (a, b, u, v) \in \{0, 1\}^m \times \{0, 1\}^m \times \{0, 1\}_n^m \times \{0, 1\}_k^m: \tag{C40}$$

$$(|\bar{a} \wedge \bar{b} \wedge u| = n \text{ and } |a \wedge b \wedge v| = k) \Longleftrightarrow |\bar{a} \wedge \bar{b} \wedge u| + |a \wedge b \wedge v| = l, \tag{C41}$$

where $l \coloneqq n + k$. If this condition is satisfied, then the pair u is a subset of the X-agreements. Since the number of X-agreements is given by $|\bar{a} \wedge \bar{b}|$, we have that

$$\text{number of subsets of } X\text{-agreements of size } n = \binom{|\bar{a} \wedge \bar{b}|}{n}. \tag{C42}$$

Since Alice and Bob are discarding surplus fully at random, each such subset is equally likely, and thus, has a probability of $1 \Big/ \binom{|\bar{a} \wedge \bar{b}|}{n}$. Arguing similarly for v and noting that the choices of u and v are independent when the quota condition is passed leads to

$$\begin{aligned} &\forall (a, b) \in \{0, 1\}^m \times \{0, 1\}^m \text{ such that } |a \wedge b| \geqslant k \text{ and } |\bar{a} \wedge \bar{b}| \geqslant n: \\ &P_{UV|AB}(u, v|a, b) = \chi(u \neq \perp, v \neq \perp, |\bar{a} \wedge \bar{b} \wedge u| + |a \wedge b \wedge v| = l) \binom{|\bar{a} \wedge \bar{b}|}{n}^{-1} \binom{|a \wedge b|}{k}^{-1}. \end{aligned} \tag{C43}$$

These two cases fully determine the conditional probability, i.e. (C34) and (C43) determine $P_{UV|AB}$ for all $(a, b) \in \{0, 1\}^m \times \{0, 1\}^m$, namely:

$$P_{UV|AB}(u, v|a, b) = \begin{cases} \chi(u = v = \perp) & \text{if } |a \wedge b| < k \text{ or } |\bar{a} \wedge \bar{b}| < n \\ \chi(u \neq \perp, v \neq \perp, |\bar{a} \wedge \bar{b} \wedge u| + |a \wedge b \wedge v| = l) \binom{|\bar{a} \wedge \bar{b}|}{n}^{-1} \binom{|a \wedge b|}{k}^{-1} & \text{if } |a \wedge b| \geqslant k \text{ and } |\bar{a} \wedge \bar{b}| \geqslant n. \end{cases} \tag{C44}$$

We can write this as

$$P_{UV|AB}(u, v|a, b) = \chi(|a \wedge b| < k \text{ or } |\bar{a} \wedge \bar{b}| < n)\chi(u = v = \perp) \tag{C45}$$

$$+\chi(|a \wedge b| \geqslant k \text{ and } |\bar{a} \wedge \bar{b}| \geqslant n)\chi(u \neq \perp, v \neq \perp, |\bar{a} \wedge \bar{b} \wedge u| + |a \wedge b \wedge v| = l) \times \binom{|\bar{a} \wedge \bar{b}|}{n}^{-1} \binom{|a \wedge b|}{k}^{-1}$$

$$= \chi(|a \wedge b| < k \text{ or } |\bar{a} \wedge \bar{b}| < n)\chi(u = v = \perp) + \chi(u \neq \perp, v \neq \perp, |\bar{a} \wedge \bar{b} \wedge u| + |a \wedge b \wedge v| = l) \binom{|\bar{a} \wedge \bar{b}|}{n}^{-1} \binom{|a \wedge b|}{k}^{-1}, \tag{C46}$$

where the last equality follows form

$$u \neq \perp, v \neq \perp, |\bar{a} \wedge \bar{b} \wedge u| + |a \wedge b \wedge v| = l \Longrightarrow |a \wedge b| \geqslant k \text{ and } |\bar{a} \wedge \bar{b}| \geqslant n. \tag{C47}$$

Taking (C23), (C26) and (C46) together, we get

$$P_{ABUV}(a, b, u, v) = p_x^{2m-|a|-|b|} p_z^{|a|+|b|} (\chi(|a \wedge b| < k \text{ or } |\bar{a} \wedge \bar{b}| < n)\chi(u = v = \perp) + \chi(u \neq \perp, v \neq \perp, |\bar{a} \wedge \bar{b} \wedge u| + |a \wedge b \wedge v| = l) \binom{|\bar{a} \wedge \bar{b}|}{n}^{-1} \binom{|a \wedge b|}{k}^{-1}). \tag{C48}$$

This concludes our formalization of $(\Omega_{ABUV}, P_{ABUV})$.

Definition 4. We define the discrete probability space $(\Omega_{ABUV}, P_{ABUV})$ by equations (C22) and (C48).

C.3. Marginalization to (Ω_{UV}, P_{UV})

Definition 5. We define the probability space (Ω_{UV}, P_{UV}) by

$$\Omega_{UV} := \mathcal{U} \times \mathcal{V} = (\{0, 1\}_n^m \bigcup \{\perp\}) \times (\{0, 1\}_k^m \bigcup \{\perp\}), \tag{C49}$$

$$P_{UV}(u, v) := \sum_{a,b \in \mathcal{A} \times \mathcal{B}} P_{ABUV}(a, b, u, v). \tag{C50}$$

Proposition 6. *It holds that*

$$P_{UV}(u, v) = \chi(u = v = \perp) \left(\sum_{n_x=0}^{n-1} \sum_{n_z=0}^{m-n_x} + \sum_{n_x=n}^{m} \sum_{n_z=0}^{\min(m-n_x, k-1)} \right) \binom{m}{n_x} \binom{m - n_x}{n_z} 2^{m-n_x-n_z} p_x^{m+n_x-n_z} p_z^{m-n_x+n_z}$$

$$+ \chi(u \neq \perp, v \neq \perp, |u \wedge v| = 0) \sum_{n_x=n}^{m-k} \sum_{n_z=k}^{m-n_x} \binom{m - l}{n_x - n} \binom{m - k - n_x}{n_z - k}$$

$$\times 2^{m-n_x-n_z} p_x^{m+n_x-n_z} p_z^{m-n_x+n_z} \binom{n_x}{n}^{-1} \binom{n_z}{k}^{-1}. \tag{C51}$$

Proof. To show equation (C51), we need to show three things:

$$P_{UV}(\perp, \perp) = \left(\sum_{n_x=0}^{n-1} \sum_{n_z=0}^{m-n_x} + \sum_{n_x=n}^{m} \sum_{n_z=0}^{\min\{m-n_x, k-1\}} \right) \binom{m}{n_x} \binom{m - n_x}{n_z} 2^{m-n_x-n_z} p_x^{m+n_x-n_z} p_z^{m-n_x+n_z}, \tag{i}$$

$\forall (u, v) \in \{0, 1\}_n^m \times \{0, 1\}_k^m$:

$$P_{UV}(u, v) = \chi(|u \wedge v| = 0) \sum_{n_x=n}^{m-k} \sum_{n_z=k}^{m-n_x} \binom{m - l}{n_x - n} \binom{m - k - n_x}{n_z - k} 2^{m-n_x-n_z} p_x^{m+n_x-n_z} p_z^{m-n_z+n_z} \binom{n_x}{n} \binom{n_z}{k}, \tag{ii}$$

$$\forall (u, v) \in (\{\perp\} \times \{0, 1\}_k^m) \bigcup (\{0, 1\}_n^m \times \{\perp\}) \colon P_{UV}(u, v) = 0. \tag{iii}$$

We start with showing (i). We have that

$$P_{UV}(\perp,\perp) = \sum_{(a,b)\in\mathcal{A}\times\mathcal{B}} P_{ABUV}(a,b,\perp,\perp), \tag{C52}$$

$$= \sum_{(a,b)\in\mathcal{A}\times\mathcal{B}} p_x^{2m-|a|-|b|}\, p_z^{|a|+|b|}\, \chi(|a\wedge b| < k \text{ or } |\bar{a}\wedge\bar{b}| < n), \tag{C53}$$

$$= \sum_{(a,b)\in\Gamma_{\text{abort}}} p_x^{2m-|a|-|b|}\, p_z^{|a|+|b|}, \tag{C54}$$

where

$$\Gamma_{\text{abort}} = \{(a,b)\in\{0,1\}^m\times\{0,1\}^m \mid |a\wedge b| < k \text{ or } |\bar{a}\wedge\bar{b}| < n\}. \tag{C55}$$

We can partition Γ_{abort} as follows:

$$\Gamma_{\text{abort}} = \bigsqcup_{(n_x,n_z)\in I_{\text{abort}}} \Gamma(n_x,n_z), \tag{C56}$$

where the 'square cup' $\sqcup$ stands for disjoint union (the union of disjoint sets) and where

$$I_{\text{abort}} = \{(n_x,n_z)\in\{0,\ldots,m\}\times\{0,\ldots,m\} \mid n_x+n_z\leqslant m,\ (n_x<n \text{ or } n_z<k)\}, \tag{C57}$$

$$\Gamma(n_x,n_z) = \{(a,b)\in\{0,1\}^m\times\{0,1\}^m \mid |a\wedge b| = n_x,\ |\bar{a}\wedge\bar{b}| = n_z\}. \tag{C58}$$

Hence

$$P_{UV}(\perp,\perp) = \sum_{(n_x,n_z)\in I_{\text{abort}}} \sum_{(a,b)\in\Gamma(n_x,n_z)} p_x^{2m-|a|-|b|}\, p_z^{|a|+|b|} \tag{C59}$$

The set $\Gamma(n_x,n_z)$ consists of all $(a,b)\in\{0,1\}^m\times\{0,1\}^m$ with exactly n_x X-agreements and exactly n_z Z-agreements. For these strings

$$\forall(a,b)\in\Gamma(n_x,n_z)\colon\ p_x^{2m-|a|-|b|}\, p_z^{|a|+|b|} = p_x^{2n_x} p_z^{2n_z} (p_x p_z)^{m-n_x-n_z}, \tag{C60}$$

$$= p_x^{m+n_x-n_z} p_z^{m-n_z+n_x}, \tag{C61}$$

so equation (C59) simplifies to

$$P_{UV}(\perp,\perp) = \sum_{(n_x,n_z)\in I_{\text{abort}}} |\Gamma(n_x,n_z)|\ p_x^{m+n_x-n_z} p_z^{m-n_z+n_x}. \tag{C62}$$

The number $|\Gamma(n_x,n_z)|$ of elements of $\Gamma(n_x,n_z)$ is given by

$$|\Gamma(n_x,n_z)| = \binom{m}{n_x}\binom{m-n_x}{n_z} 2^{m-n_x-n_z}. \tag{C63}$$

This can be seen as follows: $\binom{m}{n_x}$ is the number of possible distributions of the n_x X-agreements over the m rounds, and $\binom{m-n_x}{n_x}$ is the number of possible distributions of the n_z Z-agreements over the remaining $m-n_x$ rounds. For the rounds where the strings have basis agreement, they are fully determined, but for i in the remaining $m-n_x-n_z$ rounds, we can have that either $a_i=0$ and $b_i=1$ for a basis disagreement or $a_i=1$ and $b_i=0$. Thus, there are two possibilities for every disagreement, which explains the factor $2^{m-n_x-n_z}$. Combining equations (C62) and (C63) yields

$$P_{UV}(\perp,\perp) = \sum_{(n_x,n_z)\in I_{\text{abort}}} \binom{m}{n_x}\binom{m-n_x}{n_z} 2^{m-n_x-n_z} p_x^{m+n_x-n_z} p_z^{m-n_z+n_x}, \tag{C64}$$

$$= \left(\sum_{n_x=0}^{n-1}\sum_{n_z=0}^{m-n_x} + \sum_{n_x=n}^{m}\sum_{n_z=0}^{\min(m-n_x,k-1)}\right)\binom{m}{n_x}\binom{m-n_x}{n_z} 2^{m-n_x-n_z} p_x^{m+n_x-n_z} p_z^{m-n_x+n_z}, \tag{C65}$$

where the last equation follows from splitting up I_{abort} into the two respective sets. This shows (i).

We proceed with showing (ii). We get from equation (C48) that

$$\forall(u,v)\in\{0,1\}_n^m\times\{0,1\}_k^m\colon \tag{C66}$$

$$P_{UV}(u,v) = \sum_{(a,b)\in\{0,1\}^m\times\{0,1\}^m} p_x^{2m-|a|-|b|}\, p_z^{|a|+|b|}\, \chi(u\neq\perp, v\neq\perp, |\bar{a}\wedge\bar{b}\wedge u| + |a\wedge b\wedge v| = l)\binom{|\bar{a}\wedge\bar{b}|}{n}^{-1}\binom{|a\wedge b|}{k}^{-1} \tag{C67}$$

$$= \sum_{(a,b)\in\Phi(u,v)} p_x^{2m-|a|-|b|} p_z^{|a|+|b|} \binom{|\bar{a}\wedge\bar{b}|}{n}^{-1} \binom{|a\wedge b|}{k}^{-1}, \tag{C68}$$

where

$$\Phi(u,v) = \{(a,b)\in\{0,1\}^m\times\{0,1\}^m \,|\, |\bar{a}\wedge\bar{b}\wedge u| + |a\wedge b\wedge v| = l\}. \tag{C69}$$

In analogy to the way we split up Γ_{abort} above, we now split up $\Phi(u,v)$:

$$\Phi(u,v) = \bigsqcup_{(n_x,n_z)\in I_{\text{pass}}} \Phi(u,v,n_x,n_z), \tag{C70}$$

where

$$I_{\text{pass}} = \{(n_x,n_z)\in\{0,\ldots,m\}\times\{0,\ldots,m\} \,|\, n_x+n_z\leqslant m,\ n_x\geqslant n,\ n_z\geqslant k\}, \tag{C71}$$

$$\begin{aligned}\Phi(u,v,n_x,n_z) = \{(a,b)\in\{0,1\}^m\times\{0,1\}^m \,|\, &|\bar{a}\wedge\bar{b}\wedge u| + |a\wedge b\wedge v| \\ &= l,\ |\bar{a}\wedge\bar{b}| = n_x,\ |a\wedge b| = n_z\}.\end{aligned} \tag{C72}$$

This gives us

$$\begin{aligned}&\forall (u,v)\in\{0,1\}_n^m\times\{0,1\}_k^m\colon P_{UV}(u,v) \\ &= \sum_{(n_x,n_z)\in I_{\text{pass}}} \sum_{(a,b)\in\Phi(u,v,n_x,n_z)} p_x^{2m-|a|-|b|} p_z^{|a|+|b|} \binom{|\bar{a}\wedge\bar{b}|}{n}^{-1} \binom{|a\wedge b|}{k}^{-1}.\end{aligned} \tag{C73}$$

Again, in analogy to our calculation of $P_{UV}(u,v)$, the sets $\Phi(u,v,n_x,n_z)$ are sets on which the summand in equation (C73) is constant. More precisely, for every $(a,b,u,v)\in\{0,1\}^m\times\{0,1\}^m\times I_{\text{pass}}$, it holds that

$$\begin{aligned}&\forall (a,b)\in\Phi(u,v,n_x,n_z)\colon \quad p_x^{2m-|a|-|b|} p_z^{|a|+|b|} \binom{|\bar{a}\wedge\bar{b}|}{n}^{-1} \binom{|a\wedge b|}{k}^{-1} \\ &= p_x^{2n_x} p_z^{2n_z} (p_x p_z)^{m-n_x-n_z} \binom{n_x}{n}\binom{n_z}{k},\end{aligned} \tag{C74}$$

$$= p_x^{m+n_x-n_z} p_z^{m-n_z+n_z} \binom{n_x}{n}\binom{n_z}{k}. \tag{C75}$$

This leads us to determining the size of $\Phi(u,v,n_x,n_z)$. In words, this set contains all pairs $(a,b)\in\{0,1\}^m\times\{0,1\}^m$ with n_x X-agreements and n_z Z-agreements such that n X-agreements are located where $u_i=1$ and k Z-agreements are located where $v_i=1$. The size of this set is

$$|\Phi(u,v,n_x,n_z)| = \chi(|u\wedge v| = 0)\binom{m-l}{n_x-n}\binom{m-k-n_x}{n_z-k}2^{m-n_x-n_z}. \tag{C76}$$

This can be seen as follows. If $|u\wedge v|\neq 0$, there cannot be any $(a,b)\in\{0,1\}^m\times\{0,1\}^m$ such that $|\bar{a}\wedge\bar{b}\wedge u| + |a\wedge b\wedge v| = l$, and hence the set must be empty in that case. This explains the factor $\chi(|u\wedge v|=0)$. For those $(u,v)\in\{0,1\}_n^m\times\{0,1\}_k^m$ for which $|u\wedge v|=0$, the strings $(a,b)\in\Phi(u,v,n_x,n_z)$ are determined on $n+k=l$ positions by u and v. On the remaining $m-l$ rounds are partitioned into n_x-n rounds of X-agreements, n_z-k Z-agreements and $m-n_x-n_z$ disagreements. There are $\binom{m-l}{n_x-n}\binom{m-k-n_x}{n_z-k}$ such partitions. Finally, on each position of the $m-n_x-n_z$ disagreements, we have the two possibilities $(a_i,b_i)=(0,1)$ and $(a_i,b_i)=(1,0)$, which contributes the factor $2^{m-n_x-n_z}$. Taking equations (C75) and (C76) together, we get

$\forall (u,v)\in\{0,1\}_n^m\times\{0,1\}_k^m$:

$$P_{UV}(u,v) = \sum_{(n_x,n_z)\in I_{\text{pass}}} \chi(|u\wedge v| = 0)\binom{m-l}{n_x-n}\binom{m-k-n_x}{n_z-k}2^{m-n_x-n_z} p_x^{m+n_x-n_z} p_z^{m-n_z+n_z} \binom{n_x}{n}\binom{n_z}{k}, \tag{C77}$$

$$= \chi(|u\wedge v| = 0)\sum_{n_x=n}^{m-k}\sum_{n_z=k}^{m-n_x}\binom{m-l}{n_x-n}\binom{m-k-n_x}{n_z-k}2^{m-n_x-n_z} p_x^{m+n_x-n_z} p_z^{m-n_z+n_z} \binom{n_x}{n}\binom{n_z}{k}. \tag{C78}$$

This shows (ii).

The remaining case (iii) is easily shown. It follows directly from (C48), because

$$\begin{aligned}&\forall (u,v)\in(\{\perp\}\times\{0,1\}_k^m)\cup(\{0,1\}_n^m\times\{\perp\})\colon \chi(u=v=\perp) \\ &= \chi(u\neq\perp, v\neq\perp, |\bar{a}\wedge\bar{b}\wedge u| + |a\wedge b\wedge v| = l) = 0.\end{aligned} \tag{C79}$$

This shows (iii) and therefore completes the proof. □

C.4. Formalization of Θ and derivation of P_Θ

We have derived the probability space (Ω_{UV}, P_{UV}) as demanded in appendix C.1. Now we are left to define the random variable

$$\begin{aligned}\Theta :\ \Omega_{UV} &\to \mathrm{co}(\Theta)\\ (u, v) &\mapsto \begin{cases} h(u, v) & \text{if } (u, v) \in \{(u, v) \in \{0, 1\}_n^m \times \{0, 1\}_k^m | |u \wedge v| = 0\},\\ \perp & \text{otherwise.}\end{cases}\end{aligned} \tag{C80}$$

and to derive an expression for

$$\begin{aligned}P_\Theta : \mathrm{co}(\Theta) &\to [0, 1]\\ \vartheta &\mapsto \sum_{(u,v)\in\Theta^{-1}(\vartheta)} P_{UV}(u, v).\end{aligned} \tag{C81}$$

The range $\mathrm{co}(\Theta)$ of Θ is given by

$$\mathrm{co}(\Theta) = \{0, 1\}_k^l \bigcup \{\perp\}, \tag{C82}$$

where an element of $\{0, 1\}_k^l$ is a sifted basis choice string as in LCA sifting and where we set $\theta = \perp$ in the case where Alice and Bob abort the protocol.

To derive the random variable Θ, assume that Alice and Bob arrived at strings $(u, v) \in \mathcal{U} \times \mathcal{V}$. How do these two strings determine the sifted basis choice string ϑ? Let us first assume the case where $(u, v) \in \{0, 1\}_n^m \times \{0, 1\}_k^m$ such that $|u \wedge v| = 0$. The relevant set of indices in this case is the set of round indices r for which $u_r = 1$ or $v_r = 1$:

$$\alpha(u, v) \coloneqq \{r \in \{0, 1\}^m | u_r = 1 \text{ or } v_r = 1\}. \tag{C83}$$

Note that $|\alpha(u, v)| = n + k = l$. For $i \in [l]$, we define

$$\alpha_i(u, v) \coloneqq \text{the ith element of } \alpha(u, v). \tag{C84}$$

With this notation at hand, we can determine ϑ from u and v as follows: for $i \in [l]$, we have that $\vartheta_i = 0$ if $u_{\alpha_i(u,v)} = 1$ and $\vartheta_i = 1$ if $v_{\alpha_i(u,v)} = 1$. (Note that for $i \in [l]$, it always holds either $u_{\alpha_i(u,v)} = 1$ or $v_{\alpha_i(u,v)} = 1$, but never both, so this is well-defined.) We can write this in terms of a helper function h as

$$\begin{aligned}h : \{(u, v) \in \{0, 1\}_n^m \times \{0, 1\}_k^m | |u \wedge v| = 0\} &\to \{0, 1\}_k^l\\ (u, v) &\mapsto (h_i(u, v))_{i=1}^l,\end{aligned} \tag{C85}$$

where

$$h_i(u, v) = \begin{cases} 0 & \text{if } u_{\alpha_i(u,v)} = 1,\\ 1 & \text{if } v_{\alpha_i(u,v)} = 1.\end{cases} \tag{C86}$$

This determines Θ for all $(u, v) \in \{0, 1\}_n^m \times \{0, 1\}_k^m$ such that $|u \wedge v| = 0$. However, since these are the only pairs (u, v) for which a sifted basis choice string $\vartheta \in \{0, 1\}_k^l$ is generated, we just let Θ send all other pairs (u,v) to $\perp$:

$$\begin{aligned}\Theta : \mathcal{U} \times \mathcal{V} &\to \mathrm{co}(\Theta)\\ (u, v) &\mapsto \begin{cases} h(u, v) & \text{if } (u, v) \in \{(u, v) \in \{0, 1\}_n^m \times \{0, 1\}_k^m | |u \wedge v| = 0\},\\ \perp & \text{otherwise.}\end{cases}\end{aligned} \tag{C87}$$

This way, pairs (u, v) are mapped to $\perp$ which cannot occur in the protocol (e.g. $(\perp, b)$ with $b \in \{0, 1\}_k^l$). This is unproblematic, because for these pairs, $P_{UV}(u, v) = 0$, so according to equation (C81), they do not contribute to P_Θ.

Definition 7. We define the sifted basis choice string random variable Θ on Ω_{UV} by equation (C87). Its associated probability mass function P_Θ is given by (C81).

We are ready to state the result.

Proposition 8. *For LCA sifting (protocol 3), we have that*

$$p_{\text{abort}} = P_\Theta(\perp) = \left(\sum_{n_x=0}^{n-1}\sum_{n_z=0}^{m-n_x} + \sum_{n_x=n}^{m}\sum_{n_z=0}^{\min(m-n_x,k-1)}\right)\binom{m}{n_x}\binom{m-n_x}{n_z} 2^{m-n_x-n_z} p_x^{m+n_x-n_z} p_z^{m-n_x+n_z}, \tag{C88}$$

$$\forall\,\vartheta\in\{0,1\}_k^l\colon P_\Theta(\vartheta)=\binom{m}{n+k}\sum_{n_x=n}^{m-k}\sum_{n_z=k}^{m-n_x}\binom{m-n-k}{n_x-n}$$
$$\times\binom{m-k-n_x}{n_z-k}2^{m-n_x-n_z}p_x^{m+n_x-n_z}p_z^{m-n_x+n_z}\binom{n_x}{n}^{-1}\binom{n_z}{k}^{-1}. \tag{C89}$$

Before we prove proposition 8, let us point out its importance. Equation (C88) is the probability that the sifting protocol aborts because Alice and Bob did not reach the quota on the X- and Z-agreements, and is therefore a performance parameter of the protocol. Equation (C89) is the sampling probability for each $\vartheta\in\{0,1\}_k^l$. Since (C89) is independent of $\vartheta\in\{0,1\}_k^l$, we get uniform sampling as a corollary of proposition 8.

Corollary. *The combination of LCA sifting (protocol 3) and SBPE (protocol 2) samples uniformly. In other words, the LCA sifting protocol satisfies*

$$P_\Theta(\vartheta)=P_\Theta(\vartheta')\quad\forall\,\vartheta,\vartheta'\in\{0,1\}_k^l. \tag{C90}$$

This proves proposition 2. It leads us to proposing the protocol as a secure alternative to the insecure iterative sifting protocol.

Now we proceed to the proof of proposition 8.

Proof of proposition 8. We first show equation (C88). By definition, it holds that

$$P_\Theta(\perp)=\sum_{(u,v)\in\Theta^{-1}(\perp)}P_{UV}(u,v), \tag{C91}$$

where

$$\Theta^{-1}(\perp)=(\{\perp\}\times\{0,1\}_k^m)\cup(\{0,1\}_n^m\cup\{\perp\})\cup\{(\perp,\perp)\}\cup\{(u,v)\in\{0,1\}_n^m\times\{0,1\}_k^m||u\wedge v|\neq 0\}. \tag{C92}$$

We know from proposition 6 that

$$\forall(u,v)\in(\{\perp\}\times\{0,1\}_k^m)\bigcup(\{0,1\}_n^m\bigcup\{\perp\})\colon\quad P_{UV}(u,v)=0. \tag{C93}$$

Since

$$\forall(a,b,u,v)\in\{0,1\}^m\times\{0,1\}^m\times\{0,1\}_n^m\times\{0,1\}_k^m\colon$$
$$|u\wedge v|\neq 0\Longrightarrow|\bar a\wedge\bar b\wedge u|+|a\wedge b\wedge v|\neq 0, \tag{C94}$$

we also have

$$\forall(u,v)\in\{(u',v')\in\{0,1\}_n^m\times\{0,1\}_k^m||u'\wedge v'|\neq 0\}\colon\quad P_{UV}(u,v)=0. \tag{C95}$$

Thus

$$P_\Theta(\perp)=P_{UV}(\perp,\perp), \tag{C96}$$
$$=\left(\sum_{n_x=0}^{n-1}\sum_{n_z=0}^{m-n_x}+\sum_{n_x=n}^{m}\sum_{n_z=0}^{\min(m-n_x,k-1)}\right)\binom{m}{n_x}\binom{m-n_x}{n_z}2^{m-n_x-n_z}p_x^{m+n_x-n_z}p_z^{m-n_x+n_z}, \tag{C97}$$

where the last equality follows form proposition 6. This shows equation (C88).

We proceed with showing equation (C89). We have that

$$\forall\,\vartheta\in\{0,1\}_k^l\colon\ P_\Theta(\vartheta)=\sum_{(u,v)\in\Theta^{-1}(\vartheta)}P_{UV}(u,v), \tag{C98}$$
$$=\sum_{(u,v)\in h^{-1}(\vartheta)}P_{UV}(u,v), \tag{C99}$$

where

$$h^{-1}(\vartheta)=\left\{(u,v)\in\{0,1\}_n^m\times\{0,1\}_k^m\ \middle|\ \begin{array}{l}|u\wedge v|=0,\\ \vartheta_i=0\Longrightarrow u_{\alpha_i(u,v)}=1,\\ \vartheta_i=1\Longrightarrow v_{\alpha_i(u,v)}=1\end{array}\right\}. \tag{C100}$$

Since $|u \wedge v| = 0$ for all $(u, v) \in h^{-1}(\vartheta)$, we know from proposition 6 that

$$\forall (u, v) \in (u, v) \in h^{-1}(\vartheta):$$
$$P_{UV}(u, v) = \sum_{n_x=n}^{m-k} \sum_{n_z=k}^{m-n_x} \binom{m-l}{n_x-n} \binom{m-k-n_x}{n_z-k} 2^{m-n_x-n_z} p_x^{m+n_x-n_z} p_z^{m-n_x+n_z} \binom{n_x}{n}^{-1} \binom{n_z}{k}^{-1}. \tag{C101}$$

Thus

$$\forall \vartheta \in \{0, 1\}_k^l:$$
$$P_\Theta(\vartheta) = |h^{-1}(\vartheta)| \sum_{n_x=n}^{m-k} \sum_{n_z=k}^{m-n_x} \binom{m-l}{n_x-n} \binom{m-k-n_x}{n_z-k} 2^{m-n_x-n_z} p_x^{m+n_x-n_z} p_z^{m-n_x+n_z} \binom{n_x}{n}^{-1} \binom{n_z}{k}^{-1}. \tag{C102}$$

For every $\vartheta \in \{0, 1\}_k^l$, the set $h^{-1}(\vartheta)$ is the set of all pairs $(u, v) \in \{0, 1\}_n^m \times \{0, 1\}_k^m$ such that the following two properties are satisfied:

- $|u \wedge v| = 0$,
- for the set $\alpha(u, v)$ as in equation (C84), it holds for every $i \in [m]$ that $u_{\alpha_i(u,v)} = 1$ if $\vartheta_i = 0$ and $v_{\alpha_i(u,v)} = 1$ if $\vartheta_i = 1$.

Now note that the only thing that matters is the question which $l = n + k$ elements of $[m]$ form the subset $\alpha_i(u, v) \subseteq [m]$: for every subset $\alpha \subseteq [m]$ of size l, there is exactly one pair (u, v) which satisfies the above two properties such that $\alpha = \alpha_i(u, v)$. Hence, counting the elements of $h^{-1}(\vartheta)$ is the same as counting the l-element subsets of $[m]$, and therefore

$$|h^{-1}(\vartheta)| = \binom{m}{n+k}. \tag{C103}$$

This reduces equation (C102) to

$$\forall \vartheta \in \{0, 1\}_k^l: P_\Theta(\vartheta) = \binom{m}{n+k} \sum_{n_x=n}^{m-k} \sum_{n_z=k}^{m-n_x} \binom{m-n-k}{n_x-n} \binom{m-k-n_x}{n_z-k} 2^{m-n_x-n_z} p_x^{m+n_x-n_z} p_z^{m-n_x+n_z} \binom{n_x}{n}^{-1} \binom{n_z}{k}^{-1}, \tag{C104}$$

which is what we wanted to show. □

Appendix D. Efficiency calculation

Here we compare the efficiencies of iterative sifting and LCA sifting. Recall from equation (22) that we define the efficiency η of a sifting protocol as

$$\eta = \frac{R}{M}, \tag{D1}$$

where R is the random variable of the number of rounds that are kept after sifting and M is the random variable of the total number of rounds performed in the loop phase of the protocol. The efficiency η depends on the particular history of the protocol: different runs of the protocol may have different efficiencies. Therefore, η is a random variable. In the following, R_I and M_I denote the random variables R and M for the iterative sifting protocol, and R_L and M_L denote the corresponding random variables for the LCA protocol. Whereas in the case of iterative sifting, the number R_I is fixed and the number M_I is a random variable, the opposite is true for the LCA sifting protocol, where $M_L = m$ is fixed but but R_L is a random variable. (Note that the LCA sifting protocol may abort, in which case $R_L = 0$).

To compare the efficiencies of the two protocols, we calculate the expected value of η in each case. We first do this for the case of iterative sifting. Recall that A_r, B_r is the random variable of Alice's and Bob's basis choice in round r, respectively, and that N_d is the number of basis disagreements. Then we have:

$$\langle \eta_I \rangle = \left\langle \frac{R_I}{M_I} \right\rangle, \tag{D2}$$
$$= (n + k) \left\langle \frac{1}{M_I} \right\rangle, \tag{D3}$$
$$= (n + k) \sum_{m=n+k}^{\infty} \frac{1}{m} P_{M_I}(m), \tag{D4}$$

$$= (n+k)\sum_{m=n+k}^{\infty}\frac{1}{m}\sum_{n_d=0}^{m-n-k} P_{M_I N_d}(m, n_d), \tag{D5}$$

$$= (n+k)\sum_{m=n+k}^{\infty}\frac{1}{m}\sum_{n_d=0}^{m-n-k} (P_{M_I N_d A_m B_m}(m, n_d, 0, 0) + P_{M_I N_d A_m B_m}(m, n_d, 1, 1)), \tag{D6}$$

$$= (n+k)\sum_{m=n+k}^{\infty}\frac{1}{m}\sum_{n_d=0}^{m-n-k}\left((p_x^2)^n (p_z^2)^{m-n-n_d}(2p_x p_z)^{n_d}\binom{m-1}{n_d}\binom{m-n_d-1}{n-1}\right.$$
$$\left. + (p_x^2)^{m-k-n_d}(p_z^2)^k (2p_x p_z)^{n_d}\binom{m-1}{n_d}\binom{m-n_d-1}{k-1}\right), \tag{D7}$$

$$= (n+k)\sum_{m=n+k}^{\infty}\frac{1}{m}\sum_{n_d=0}^{m-n-k}(2p_x p_z)^{n_d}\binom{m-1}{n_d}\left((p_x^2)^n (p_z^2)^{m-n-n_d}\binom{m-n_d-1}{n-1}\right.$$
$$\left. + (p_x^2)^{m-k-n_d}(p_z^2)^k\binom{m-n_d-1}{k-1}\right). \tag{D8}$$

For the case of the LCA sifting protocol, we have:

$$\langle \eta_L \rangle = \frac{R_L}{M_L}, \tag{D9}$$

$$= \frac{1}{m}\langle R_L \rangle, \tag{D10}$$

$$= \frac{1}{m}(n+k)P[N_x \geqslant n \wedge N_z \geqslant k], \tag{D11}$$

$$= \frac{1}{m}(n+k)\sum_{n_d=0}^{m-n-k} P[N_x \geqslant n \wedge N_z \geqslant k \wedge N_d = d], \tag{D12}$$

$$= \frac{n+k}{m}\sum_{n_d=0}^{m-n-k}\sum_{n_z=k}^{m-n_d-n} P[N_x \geqslant n \wedge N_z = n_z \wedge N_d = n_d], \tag{D13}$$

$$= \frac{n+k}{m}\sum_{n_d=0}^{m-n-k}\sum_{n_z=k}^{m-n_d-n}(p_x^2)^{m-n_z-n_d}(p_z^2)^{n_z}(2p_x p_z)^{n_d}\binom{m}{n_d}\binom{m-n_d}{n_z}. \tag{D14}$$

The calculation of the expected efficiencies (D8) and (D14) requires a lot of computational power. We wrote programs that compute numerical lower bounds on $\langle \eta_I \rangle$ and $\langle \eta_L \rangle$ for the case where the probabilities are symmetric ($p_x = p_z = 1/2$) and where the quotas coincide ($n = k$). A plot of these lower bounds is shown in figure 3. In order to plot the lower bound on $\langle \eta_L \rangle$, a choice for m had to be made for each value of $n = k$. Our program choses an m which is likely to maximize the expected efficiency for the given value of $n = k$. Note that $1/2$, being the expected fraction of basis agreements, is an upper bound on the expected efficiencies. Hence, figure 3 indicates that the difference in the expected efficiencies becomes insignificant for practically relevant values of the block length $n + k$. This means that replacing iterative sifting by LCA sifting is unlikely to have a significant effect on the key rate of a QKD protocol.

Appendix E. Proof of the sufficiency of the formal criteria

In this appendix, we prove that the two formal criteria for good sifting, (1) and (2), are sufficient for good sifting in the sense that the relevant statistical inequality, (6), follows from these two conditions. In other words, we prove proposition 3.

Proof of proposition 3. According to Bayes' theorem, we have that

$$p_{\text{tail}} = P[\Lambda_{\text{key}} \geqslant \Lambda_{\text{test}} + \mu | \Lambda_{\text{test}} \leqslant q_{\text{tol}}], \tag{E1}$$

$$= \frac{P[\Lambda_{\text{test}} \leqslant q_{\text{tol}} | \Lambda_{\text{key}} \geqslant \Lambda_{\text{test}} + \mu] P[\Lambda_{\text{key}} \geqslant \Lambda_{\text{test}} + \mu]}{P[\Lambda_{\text{test}} \leqslant q_{\text{tol}}]}, \tag{E2}$$

$$\leqslant \frac{P[\Lambda_{\text{key}} \geqslant \Lambda_{\text{test}} + \mu]}{p_{\text{pass}}}. \tag{E3}$$

We define the *total error rate* Λ_{tot} as the random variable

$$\begin{aligned}\Lambda_{\text{tot}}: \ \Omega_{ZZ'\Theta} &\to [0,1]\\ (z, z', \vartheta) &\mapsto \frac{1}{l}\sum_{i=1}^{l} z \oplus z'.\end{aligned} \tag{E4}$$

For all $(z, z', \vartheta) \in \Omega_{ZZ'\Theta}$, it holds that

$$\Lambda_{\text{key}}(z, z, \vartheta) \geqslant \Lambda_{\text{test}}(z, z, \vartheta) + \mu, \tag{E5}$$

$$\iff \frac{1}{n}\sum_{i=1}^{l}(1-\vartheta_i)(z_i \oplus z_i') \geqslant \frac{1}{k}\sum_{i=1}^{l}\vartheta_i(z_i \oplus z_i') + \mu, \tag{E6}$$

$$\iff \frac{1}{n}\sum_{i=1}^{l}(1-\vartheta_i)(z_i \oplus z_i') + \frac{1}{k}\sum_{i=1}^{l}(1-\vartheta_i)(z_i \oplus z_i') \geqslant \frac{1}{k}\sum_{i=1}^{l}\vartheta_i(z_i \oplus z_i') + \frac{1}{k}\sum_{i=1}^{l}(1-\vartheta_i)(z_i \oplus z_i') + \mu, \tag{E7}$$

$$\iff \left(\frac{1}{n}+\frac{1}{k}\right)\sum_{i=1}^{l}(1-\vartheta_i)(z_i \oplus z_i') \geqslant \frac{1}{k}\sum_{i=1}^{l}(z_i \oplus z_i') + \mu, \tag{E8}$$

$$\iff \frac{k}{l}\left(\frac{1}{n}+\frac{1}{k}\right)\sum_{i=1}^{l}(1-\vartheta_i)(z_i \oplus z_i') \geqslant \frac{k}{l}\frac{1}{k}\sum_{i=1}^{l}(z_i \oplus z_i') + \frac{k}{l}\mu, \tag{E9}$$

$$\iff \Lambda_{\text{key}}(z, z, \vartheta) \geqslant \Lambda_{\text{tot}}(z, z, \vartheta) + \frac{k}{l}\mu. \tag{E10}$$

We express the error *rates* Λ_{key}, Λ_{test} and Λ_{tot} in terms of the error *numbers* Σ_{key}, Σ_{test} and Σ_{tot},

$$\Sigma_{\text{key}} = n\Lambda_{\text{key}}, \quad \Sigma_{\text{test}} = k\Lambda_{\text{test}}, \quad \Sigma_{\text{tot}} = l\Lambda_{\text{tot}}. \tag{E11}$$

This gives us

$$\Lambda_{\text{key}} \geqslant \Lambda_{\text{tot}} + \frac{k}{l}\mu \iff \Sigma_{\text{key}} \geqslant n\left(\frac{\Sigma_{\text{tot}}}{l} + \frac{l-n}{l}\mu\right). \tag{E12}$$

Therefore

$$P[\Lambda_{\text{key}} \geqslant \Lambda_{\text{test}} + \mu] = P\left[\Sigma_{\text{key}} \geqslant n\left(\frac{\Sigma_{\text{tot}}}{l} + \frac{l-n}{l}\mu\right)\right] \tag{E13}$$

and hence

$$p_{\text{tail}} \leqslant \frac{P\left[\Sigma_{\text{key}} \geqslant n\left(\frac{\Sigma_{\text{tot}}}{l} + \frac{l-n}{l}\mu\right)\right]}{p_{\text{pass}}}, \tag{E14}$$

$$= \frac{\sum_{\sigma_{\text{tot}}} P[\Sigma_{\text{tot}} = \sigma_{\text{tot}}] P\left[\Sigma_{\text{key}} \geqslant n\left(\frac{\sigma_{\text{tot}}}{l} + \frac{l-n}{l}\mu\right) \mid \Sigma_{\text{tot}} = \sigma_{\text{tot}}\right]}{p_{\text{pass}}}, \tag{E15}$$

$$= \frac{\sum_{\sigma_{\text{tot}}} P[\Sigma_{\text{tot}} = \sigma_{\text{tot}}] \sum_j P[\Sigma_{\text{key}} = j \mid \Sigma_{\text{tot}} = \sigma_{\text{tot}}]}{p_{\text{pass}}}, \tag{E16}$$

where the sum over j ranges over all possible values of Σ_{key} that are larger or equal to the according value, i.e.

$$j = \left\lceil n\left(\frac{\sigma_{\text{tot}}}{l} + \frac{l-n}{l}\mu\right)\right\rceil, \left\lceil n\left(\frac{\sigma_{\text{tot}}}{l} + \frac{l-n}{l}\mu\right)\right\rceil + 1, \ldots, n, \tag{E17}$$

where $\lceil\cdot\rceil$ denotes the ceiling function

$$h(\sigma_{\text{tot}}, l, n, j) := P[\Sigma_{\text{key}} = j \mid \Sigma_{\text{tot}} = \sigma_{\text{tot}}], \tag{E18}$$

$$= \frac{P[\Sigma_{\text{key}} = j \wedge \Sigma_{\text{tot}} = \sigma_{\text{tot}}]}{P[\Sigma_{\text{tot}} = \sigma_{\text{tot}}]}, \tag{E19}$$

$$= \frac{P[\Omega_{j\sigma_{\text{tot}}}]}{P[\Omega_{\sigma_{\text{tot}}}]}, \tag{E20}$$

$$= \frac{\sum_{(z,z',\vartheta)\in\Omega_{j\sigma_{\text{tot}}}} P_{ZZ'\Theta}(z,\, z',\, \vartheta)}{\sum_{(z,z',\vartheta)\in\Omega_{\sigma_{\text{tot}}}} P_{ZZ'\Theta}(z,\, z',\, \vartheta)}, \tag{E21}$$

where

$$\Omega_{j\sigma_{\text{tot}}} = \{(z,\, z',\, \vartheta) \in \Omega_{ZZ'\Theta} | \Sigma_{\text{key}}(z,\, z',\, \vartheta) = j \wedge \Sigma_{\text{tot}}(z,\, z',\, \vartheta) = \sigma_{\text{tot}}\}, \tag{E22}$$

$$\Omega_{\sigma_{\text{tot}}} = \{(z,\, z',\, \vartheta) \in \Omega_{ZZ'\Theta} | \Sigma_{\text{tot}}(z,\, z',\, \vartheta) = \sigma_{\text{tot}}\}. \tag{E23}$$

It holds for all $(z,\, z',\, \vartheta) \in \Omega_{ZZ'\Theta}$ that

$$P_{ZZ'\Theta}(z,\, z',\, \vartheta) = P_{ZZ'}(z,\, z')P_{\Theta}(\vartheta) \tag{E24}$$

$$= P_{ZZ'}(z,\, z')c, \tag{E25}$$

where $P_{ZZ'}$ and P_{Θ} are the according marginal distributions of $P_{ZZ'\Theta}$. Equation (E24) follows from (2), and equation (E25) follows from equation (1). This implies

$$h(\sigma_{\text{tot}},\, l,\, n,\, j) = \frac{\sum_{(z,z',\vartheta)\in\Omega_{j\sigma_{\text{tot}}}} P_{ZZ'}(z,\, z')p}{\sum_{(z,z',\vartheta)\in\Omega_{\sigma_{\text{tot}}}} P_{ZZ'}(z,\, z')p}, \tag{E26}$$

$$= \frac{\sum_{(z,z',\vartheta)\in\Omega_{j\sigma_{\text{tot}}}} P_{ZZ'}(z,\, z')}{\sum_{(z,z',\vartheta)\in\Omega_{\sigma_{\text{tot}}}} P_{ZZ'}(z,\, z')}, \tag{E27}$$

$$= \frac{\sum_{(z,z')\in\Gamma_{\sigma_{\text{tot}}}} P_{ZZ'}(z,\, z')\binom{\sigma_{\text{tot}}}{j}\binom{l-\sigma_{\text{tot}}}{n-j}}{\sum_{(z,z')\in\Gamma_{\sigma_{\text{tot}}}} P_{ZZ'}(z,\, z')\binom{l}{n}}, \tag{E28}$$

$$= \binom{\sigma_{\text{tot}}}{j}\binom{l-\sigma_{\text{tot}}}{n-j}\binom{l}{n}^{-1}, \tag{E29}$$

where

$$\Gamma_{\sigma_{\text{tot}}} = \left\{(z,\, z') \in \{0,\ 1\}^l \times \{0,\ 1\}^l \,|\, \sum_{i=1}^{l} z_i \oplus z_i' = \sigma_{\text{tot}}\right\}. \tag{E30}$$

Equation (E29) means that $h(\sigma_{\text{tot}},\, l,\, n,\, j)$ is a hypergeometric distribution. We are interested in the *tail* of this distribution,

$$H(\sigma_{\text{tot}},\, l,\, n,\, d) := \sum_{j=d}^{n} h(\sigma_{\text{tot}},\, l,\, n,\, j), \tag{E31}$$

because according to equations (E16) and (E17),

$$p_{\text{tail}} \leqslant \frac{\sum_{\sigma_{\text{tot}}} P[\Sigma_{\text{tot}} = \sigma_{\text{tot}}] H(\sigma_{\text{tot}},\, l,\, n,\, d)}{p_{\text{pass}}}, \tag{E32}$$

where

$$d = \left\lceil n\left(\frac{\sigma_{\text{tot}}}{l} + \frac{l-n}{l}\mu\right)\right\rceil. \tag{E33}$$

There are several well-known bounds on the tail of a hypergeometric distribution [24]. For our case, *Serfling's bound* is a suitable one [25]. The appropriate special case of Serfling's bound for this case reads

$$H(\sigma_{\text{tot}},\, l,\, n,\, d) \leqslant \exp\left(-2\frac{(l-n)n}{l}\frac{l-n}{l-n+1}\mu^2\right), \tag{E34}$$

$$= \exp\left(-2\frac{kn}{l}\frac{k}{k+1}\mu^2\right). \tag{E35}$$

(Instead of Serfling's bound, one may use *Hoeffding's bound* [26]. That bound is weaker than Serfling's bound in this case, but it has the advantage that it has been formulated directly in terms of hypergeometric distributions [27, 28], so these references are easier to understand in our context.) Inequalities (E32) and (E35) together imply

$$p_{\text{tail}} \leqslant \frac{\sum_{\sigma_{\text{tot}}} P[\Sigma_{\text{tot}} = \sigma_{\text{tot}}] H(\sigma_{\text{tot}},\, l,\, n,\, d)}{p_{\text{pass}}}, \tag{E36}$$

$$\leqslant \frac{\exp\left(-2\frac{kn}{l}\frac{k}{k+1}\mu^2\right)}{p_{\text{pass}}}, \tag{E37}$$

which completes the proof. □

References

[1] Renner R 2005 Security of quantum key distribution *PhD Thesis* Zürich ETH (arXiv:quant-ph/0512258)
[2] Bonato A, Tomaello C, Da Deppo V, Naletto G and Villoresi P 2009 Feasibility of satellite quantum key distribution *New J. Phys.* **11** 045017
[3] Tomamichel M, Lim C C W, Gisin N and Renner R 2012 Tight finite-key analysis for quantum cryptography *Nat. Commun.* **3** 634
[4] Lim C C W, Curty M, Walenta N, Xu F and Zbinden H 2014 Concise security bounds for practical decoy-state quantum key distribution *Phys. Rev.* A **89** 022307
[5] Hayashi M and Tsurumaru T 2011 Concise and tight security analysis of the Bennett–Brassard 1984 protocol with finite key lengths *New J. Phys.* **9** 093014
[6] Hayashi M and Nakayama R 2014 Security analysis of the decoy method with the Bennett–Brassard 1984 protocol for finite key lengths *New J. Phys.* **16** 063009
[7] Tomamichel M, Fehr S, Kaniewski J and Wehner S 2013 A monogamy-of-entanglement game with applications to device-independent quantum cryptography *New J. Phys.* **15** 103002
[8] Curty M, Xu F, Cui W, Lim C C W, Tamaki K and Lo H-K 2014 Finite-key analysis for measurement-device-independent quantum key distribution *Nat. Commun.* **5** 3732
[9] Lim C C W, Portmann C, Tomamichel M, Renner R and Gisin N 2013 Device-independent quantum key distribution with local Bell test *Phys. Rev.* X **3** 031006
[10] Furrer F, Franz T, Berta M, Leverrier A, Scholz V B, Tomamichel M and Werner R F 2011 Continuous variable quantum key distribution: finite-key analysis of composable security against coherent attacks *Phys. Rev. Lett.* **109** 100502
[11] Leverrier A 2014 Composable security proof for continuous-variable quantum key distribution with coherent states *Phys. Rev. Lett.* **114** 070501
[12] Bacco D, Canale M, Laurenti N, Vallone G and Villoresi P 2013 Experimental quantum key distribution with finite-key security analysis for noisy channels *Nat. Commun.* **4** 2363
[13] Xu F, Sajeed S, Kaiser S, Tang Z, Qian L, Makarov V and Lo H-K 2014 Experimental quantum key distribution with source flaws *Phys. Rev.* A **92** 032305
[14] Korzh B, Lim C C W, Houlmann R, Gisin N, Li M J, Nolan D, Sanguinetti B, Thew R and Zbinden H 2014 Provably secure and practical quantum key distribution over 307 km of optical fibre *Nat. Photon.* **9** 7
[15] Lo H-K, Chau H F and Ardehali M 2005 Efficient quantum key distribution scheme and a proof of its unconditional security *J. Cryptol.* **18** 133–65
[16] Tomamichel M and Renner R 2011 Uncertainty relation for smooth entropies *Phys. Rev. Lett.* **106** 110506
[17] Renner R 2007 Symmetry of large physical systems implies independence of subsystems *Nat. Phys.* **3** 645–9
[18] Huttner B and Ekert A K 1994 Information gain in eavesdropping *J. Mod. Opt.* **41** 2455–66
[19] Scarani V, Bechmann-Pasquinucci H, Cerf N J, Dusek M, Lutkenhaus N and Peev M 2009 The security of practical quantum key distribution *Rev. Mod. Phys.* **81** 1301–50
[20] Bennett C H and Brassard G 1984 Quantum cryptography: public key distribution and coin tossing *Proc. IEEE Int. Conf. on Comp. Sys. and Signal Process* pp 175–9
[21] Shor P W and Preskill J 2000 Simple proof of security of the bb84 quantum key distribution protocol *Phys. Rev. Lett.* **85** 441–4
[22] Tomamichel M and Leverrier A 2015 A rigorous and complete proof of finite key security of quantum key distribution (arXiv:1506.08458)
[23] Fung F C-H, Ma X and Chau H 2010 Practical issues in quantum-key-distribution postprocessing *Phys. Rev.* A **81** 012318
[24] Boucheron S, Lugosi G and Massart P 2013 *Concentration Inequalities: A Nonasymptotic Theory of Independence* (Oxford: Oxford University Press)
[25] Serfling R J 1974 Probability inequalities for the sum in sampling without replacement *Ann. Stat.* **2** 39–48
[26] Hoeffding W 1963 Probability inequalities for sums of bounded random variables *J. Am. Statist. Assoc.* **58** 13–30
[27] Chvátal V 1978 The tail of the hypergeometric distribution *Discrete Math.* **25** 285–7
[28] Skala M 2013 Hypergeometric tail inequalities: ending the insanity (arXiv:1311.5939)