Sifting problems in finite-size quantum key distribution Extended Abstract

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

[1] *QuTech, Delft University of Technology, Lorentzweg 1, 2628 CJ Delft, 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*

GENERAL BACKGROUND

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 often an insufficient description of actual experiments. In practice, Alice and Bob have limited time and/or dim light sources that limit the number of photons they can exchange. For example, in satellite-based QKD 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.

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. [2] involving the smooth entropy formalism. This framework was later extended to a decoy-state protocol by Lim et al. [3]. An alternative framework was developed by Hayashi and collaborators [4, 5]. Other extensions of the finite-key framework include the treatment of device-independent protocols by Curty et al. [6] and Lim et al. [7], and continuous-variable protocols by Furrer et al. [8] and Leverrier [9]. The framework used in the aforementioned works, relying on some fairly technical results,[1] 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 with finite-key analysis. Some examples of such experiments include Refs. [10–12].

One can argue that all real QKD implementations are of finite size. It is therefore of utmost importance to fully understand the subtleties of finite-size effects. We point out that this understanding is required not only for theoretical satisfaction but also for the correct calculation of practically relevant quantities such as the key rate or the noise tolerance of a QKD protocol.

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

SIFTING PROBLEMS

General situation

In this work, we address some finite-size effects that seem to have remained unnoticed so far. We analyze some subtleties that occur in the sifting[2] stage of QKD. Namely, we point out a discrepancy between the assumptions that enter into security proofs and the actual sifting protocol used in several works on finite-key analysis. Our calculations suggest that this discrepancy might be exploited by the eavesdropper, and hence if not fixed, might lead to a security vulnerability.

At the same time, we see no reason to believe that this discrepancy cannot be eliminated. We believe that once this gap is closed, the theoretical framework developed in the aforementioned works is valid and contains the relevant ideas.

In the meantime, our work suggests that there is an open area of research that needs to be addressed: how to link sifting protocols to the statistical assumptions needed for security proofs.

The particular sifting scheme that we call into question, which we call *iterative sifting*, involves an iterative procedure where Alice and Bob communicate their measurement choices after each completed round until a quota is met. This iterative sifting procedure was part of theoretical protocols [2, 3, 6, 7] and has found experimental implementation [10].

Ultimately the issue 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 ahead of time which of the rounds form the *sample* for testing for an eavesdropper's presence, then she can adjust her attack appropriately and security is compromised. Hence a central assumption in the QKD security analysis is that Eve has no knowledge about the sample.

[2] Sifting is the process where Alice and Bob announce their measurement bases and discard some data, e.g., if the measurement bases disagree.

Contribution: Identified problems

We show that this assumption, that Eve has no knowledge about the sample, 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:

1. *Non-uniform sampling*: The sampling probability, due to which the sampled bits and the encoding basis are chosen, is not uniform.[3] 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.

2. *Basis information leak*: Alice and Bob's public communication about their previous basis choices (which happens before the quantum communication is over) allows Eve to update her knowledge about which of the upcoming bits end up in the sample. As a consequence, the quantum information that passes through the channel thereafter will 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 and quantified. More importantly, however, the protocols in question are designed to be secure for all finite key lengths. In the light of these two problems, the security analysis of iterative sifting in the literature is incomplete. This is not a purely theoretical objection but a practically very relevant issue.

To illustrate these two problems and to show their potential harm to the security of a QKD protocol, we present some simple intercept-resend attack strategies on iterative sifting and calculate the expected error rates they cause:

1. We show that non-uniform sampling can be exploited using a fixed non-adaptive strategy, with error rates as low as $\approx 22.8\%$.

2. Eve can exploit the basis information leak using an adaptive attack strategy that conditions the basis used for the attack on the classical data that Alice and Bob exchanged previously. In this case, error rates as low as $\approx 16.3\%$ are possible.

1 & 2. An attack strategy simultaneously exploiting both problems achieves expected error rates of $\approx 15.8\%$.

These numbers are well-below the typical 25% value for intercept-resend attacks [13]. We also show that the two identified problems are in fact two independent problems that appear concurrently in iterated sifting but which can occur separately in general.

Contribution: Suggested solution

There are two ways to address these problems. One way is to investigate whether the security analysis of the QKD protocols involving iterative sifting can be modified such that the two problems are addressed and the security of the protocol is proved, but perhaps with more pessimistic key rates than originally thought. In other words, one could leave the protocol unchanged but change its analysis. The second way is to find clear mathematical formulations of the assumptions in the analysis that are violated by these two problems. This allows one to investigate candidate protocols that could substitute for iterative sifting without violating these assumptions. In other words, the second way is to leave the assumptions of the security analysis unchanged but to change the protocol which is supposed to satisfy these assumptions.

We choose to address the problems in the second way. We present a candidate for a protocol that avoids the two problems, and support our candidate by some example calculations. More generally, we present two formal mathematical criteria that can be checked for a protocol in order to guarantee that the usual security analysis employed in the literature is valid. These criteria are expressed in terms of a state $\rho_{A^N B^N \Theta^N}$ which is associated with a sifting protocol. It is the quantum-quantum-classical state that the protocol would output if all the measurements of the protocol were skipped and the qubit systems kept in their state instead. Then A^N and B^N are Alice's and Bob's sifted qubit systems, and Θ^N is a public register containing their basis choice and sample information. Informally stated, the two criteria are:

1. The sampling probability encoded in the register Θ^N is uniform over all samples of the appropriate size. If this is guaranteed, then there is no non-uniform sampling.

2. Alice's and Bob's quantum information is uncorrelated to the basis choice and sample information, i.e. the according systems are in a product state, $\rho_{A^N B^N \Theta^N} = \rho_{A^N B^N} \otimes \rho_{\Theta^N}$. If this is guaranteed, then there is no basis information leak.

We show that these two criteria are sufficient in the sense that the relevant statistical inequality needed for the parameter estimation analysis can be derived from these two assumptions. Although the state $\rho_{A^N B^N \Theta^N}$ is only fully known if Eve's attack strategy is known, the two criteria can be checked independently of Eve's actions if the sifting protocol is properly modeled.

[3] In general, the sampling probability (which decides 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 iterative sifting, however, the two coincide.

CONCLUSION

In conclusion, we have shown that there is a gap between the physical protocol used for finite-key sifting and the security analysis in the literature and, furthermore, that this gap can be exploited by an eavesdropper. We believe that this observation will lead eventually to a more complete framework for finite-key analysis. In the sense that a framework is needed to address the issues we raised, our work opens up an area for research.

We advocate that in future QKD protocols and their analysis, the sifting process is no longer only assumed to satisfy the necessary conditions for parameter estimation but that the sifting is explicitly modeled and proved to satisfy these conditions. The sufficient formal criteria that we found can serve as a guiding light in this undertaking.

[1] Renato Renner. Security of Quantum Key Distribution. (16242), December 2005.

[2] Marco Tomamichel, Charles Ci Wen Lim, Nicolas Gisin, and Renato Renner. Tight finite-key analysis for quantum cryptography. *Nat. Commun.*, 3(634), 2012.

[3] Charles Ci Wen Lim, Marcos Curty, Nino Walenta, Feihu Xu, and Hugo Zbinden. Concise security bounds for practical decoy-state quantum key distribution. *Phys. Rev. A*, 89(2):022307, 2014.

[4] Masahito Hayashi and Toyohiro Tsurumaru. Simple and Tight Security Analysis of the Bennett-Brassard 1984 Protocol with Finite Key Lengths. 093014:9, 2011.

[5] Masahito Hayashi and Ryota Nakayama. Security analysis of the decoy method with the Bennett-Brassard 1984 protocol for finite key lengths. *New Journal of Physics*, 16, 2014.

[6] Marcos Curty, Feihu Xu, Wei Cui, Charles Ci Wen Lim, Kiyoshi Tamaki, and Hoi-Kwong Lo. Finite-key analysis for measurement-device-independent quantum key distribution. *Nat. Commun.*, 5, 2014.

[7] Charles Ci Wen Lim, Christopher Portmann, Marco Tomamichel, Renato Renner, and Nicolas Gisin. Device-independent quantum key distribution with local bell test. *Phys. Rev. X*, 3(3):031006, 2013.

[8] Fabian Furrer, Torsten Franz, Mario Berta, Anthony Leverrier, Volkher B. Scholz, Marco Tomamichel, and Reinhard F. Werner. Continuous Variable Quantum Key Distribution: Finite-Key Analysis of Composable Security against Coherent Attacks. pages 1–10, December 2011.

[9] Anthony Leverrier. Composable security proof for continuous-variable quantum key distribution with coherent states. August 2014.

[10] Davide Bacco, Matteo Canale, Nicola Laurenti, Giuseppe Vallone, and Paolo Villoresi. Experimental quantum key distribution with finite-key security analysis for noisy channels. *Nat. Commun.*, 4, 2013.

[11] Feihu Xu, Shihan Sajeed, Sarah Kaiser, Zhiyuan Tang, Li Qian, Vadim Makarov, and Hoi-Kwong Lo. Experimental quantum key distribution with source flaws and tight finite-key analysis. pages 1–11, August 2014.

[12] Boris Korzh, Charles Ci Wen Lim, Raphael Houlmann, Nicolas Gisin, Ming Jun Li, Daniel Nolan, Bruno Sanguinetti, Rob Thew, and Hugo Zbinden. Provably Secure and Practical Quantum Key Distribution over 307 km of Optical Fibre. *Nature Photonics*, 9(3):7, 2014.

[13] Valerio Scarani, Helle Bechmann-Pasquinucci, Nicolas J. Cerf, Miloslav Dušek, Norbert Lütkenhaus, and Momtchil Peev. The security of practical quantum key distribution. *Rev. Mod. Phys.*, 81(3):1301–1350, 2009.

Sifting problems in finite-size quantum key distribution
Technical details

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

[1] *QuTech, Delft University of Technology, Lorentzweg 1, 2628 CJ Delft, 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*

A central assumption in quantum key distribution (QKD) security analysis is that Eve has no knowledge about which rounds will be used for parameter estimation or for key distillation. However, we show that this assumption is violated for a common implementation of QKD that we call *iterative sifting*, which involves an iterative procedure of communicating measurement choices of completed rounds until a quota is met. We show that iterated sifting leads to two problems. Firstly, some rounds are more likely to be key rounds than others. This bias allows Eve to choose an attack strategy that attempts to read out the key when this is less likely to be detected. Secondly, the public communication of past measurement choices changes this bias round by round, which allows Eve to gradually update her attack strategy to further reduce the likelihood of being detected.

We analyze these two previously unnoticed problems, present attack strategies that exploit them, calculate their error rates and find that the two problems are independent. We propose a simple correction to recover the security of the protocol whereby Alice and Bob announce their basis choices only after the quantum communication is completed. More generally, we present two formal mathematical criteria for a sifting protocol that ensure that these two problems are avoided and that the parameter estimation leads to correct conclusions. These may be guiding lights in the design of future protocols and in the development of a rigorous formal QKD analysis, which has neglected sifting-related problems so far.

CONTENTS

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 often an insufficient description of actual experiments. In practice, Alice and Bob have limited time and/or dim light sources that limit the number of photons they can exchange. For example, in satellite-based QKD 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.

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. [2] involving the smooth entropy formalism. This framework was later extended to a decoy-state protocol by Lim et al. [3]. An alternative framework was developed by Hayashi and collaborators [4, 5]. Other extensions of the finite-key framework include the treatment of device-independent protocols by Curty et al. [6] and Lim et al. [7], and continuous-variable protocols by Furrer et al. [8] and Leverrier [9]. The framework used in the aforementioned works, relying on some fairly technical results,[1] 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 with finite-key analysis. Some examples of such experiments include Refs. [10–12].

One can argue that all real QKD implementations are of finite size. It is therefore of utmost importance to fully understand the subtleties of finite-size effects. In this work, we point out a subtlety that occurs in the sifting[2] stage of QKD, which was not recognized previously. Namely, we point out a discrepancy between the assumptions that enter the theoretical treatment of the sifted data and the actual sifting protocol used in several works on finite-key analysis. Our calculations suggest that this discrepancy might be exploited by the eavesdropper, and hence if not fixed, might lead to a security vulnerability. At the same time, we emphasize that we see no reason to believe that this discrepancy cannot be eliminated. We believe that once this gap is closed, the theoretical framework developed in the aforementioned works is valid and contains all the relevant ideas.

The particular sifting scheme that we call into question, which we call *iterative sifting*, involves an iterative procedure where Alice and Bob communicate their measurement choices after each completed round until a quota is met. This iterative sifting procedure was part of theoretical protocols [2, 3, 6, 7] and has found experimental implementations [10].

Ultimately the issue 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 ahead of time which of the rounds form the *sample* for testing for an eavesdropper's presence, then she can adjust her attack appropriately and security is compromised. 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. On the other hand, we propose an alternative sifting scheme that may fix the problem.

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.[3] 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 happens before the quantum communication is over) allows Eve to update her knowledge about which of the upcoming bits end up in the sample. As a consequence, the quantum information that passes the channel thereafter will 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, their analysis does currently not account for these finite-size effects. This is not a purely theoretical objection but a practically very relevant issue.

---

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

[2] Sifting is the process where Alice and Bob announce their measurement bases and discard some data, e.g., if the measurement bases disagree.

[3] 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 iterative sifting, however, the two coincide.

There are two ways to address these problems. The first way would be to investigate whether the security analysis of the QKD protocols involving iterated sifting can be modified such that the two problems are addressed and the security of the protocol is proved. In other words, the first way would be to leave the protocol unchanged but to change the assumptions on which the security analysis is based. The second way is to find clear mathematical formulations of the assumptions in the analysis that are violated by these two problems. This allows to investigate candidate protocols that could substitute iterative sifting without violating these assumptions. In other words, the second way is to leave the assumptions of the security analysis unchanged but to change the protocol which is supposed to satisfy these assumptions.

We choose to address the problem in the second way. We find that the two mentioned problems are in fact two independent problems that appear concurrently in iterated sifting but which can occur separately in general. We show how these problems open up the possibility of successful attacks by an eavesdropper. We present a candidate for a protocol that avoids the two problems, and support our candidate by some example calculations. More generally, we present two formal mathematical criteria that can be checked for a protocol in order to guarantee that the usual security analysis employed in the literature is valid. These criteria are expressed in terms of a state $\rho_{A^N B^N \Theta^N}$ which is associated with a sifting protocol. It is the quantum-quantum-classical state that the protocol would output if all the measurements of the protocol were skipped and the qubit systems kept in their state instead. Then A^N and B^N are Alice's and Bob's sifted qubit systems, and Θ^N is a public register containing their basis choice and sample information. Informally stated, the two criteria are:

- The sampling probability encoded in the register Θ^N is uniform over all samples of the appropriate size. If this is guaranteed, then there is no non-uniform sampling.
- Alice's and Bob's quantum information is uncorrelated to the basis choice and sample information, i.e. the according systems are in a product state, $\rho_{A^N B^N \Theta^N} = \rho_{A^N B^N} \otimes \rho_{\Theta^N}$. If this is guaranteed, then there is no basis information leak.

We show that these two criteria are sufficient in the sense that the relevant statistical inequality of the parameter estimation analysis can be derived from these two assumptions. Although the state $\rho_{A^N B^N \Theta^N}$ is only fully known if Eve's attack strategy is known, the two criteria can be checked independently of Eve's actions if the sifting protocol is properly modeled. We therefore advocate that in future protocol designs, the sifting process is modeled explicitly. The two criteria can then help as a guideline.

I. WHAT EXACTLY IS THE PROBLEMATIC SCHEME?

A typical QKD protocol consists of the following subroutines [2]:

(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 entirely to the subroutine (i) and (ii), whereas subroutines (iii) and (iv) are not of our concern. We refer to subroutine (i) collectively as "sifting". 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 the way in which these processes are linked in our protocol of interest make it hard to separate the sifting process from the other three processes.

Our focus in this article is on a particular protocol for (i), which we call *iterated sifting*. This protocol is problematic, but to fully understand why, one needs to see how the output of the iterated sifting protocol is processed in the subsequent subroutine (ii), the parameter estimation. The typical analysis of the parameter estimation is false if the parameter estimation is preceded by iterative sifting. This is why we will occasionally refer to the parameter estimation in Sections II and III and discuss it in more detail in Section IV.

A. Description of the iterative sifting protocol

In this section, we will have a closer look at the sifting protocol in question, which we call the *iterative sifting protocol*. It has found slightly different formulations 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 [2, 3, 6, 10] or as an entanglement-based protocol [7]. These details are irrelevant for the problems that we describe. In order to fix notation and to have a clearly formulated version of the iterated sifting protocol at hand, we write down a version of the protocol in Protocol 1. Its formulation and notation is close to the one described in [2], 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 analysis in Section IV is easier to formulate, but a prepare-and-measure based protocol would otherwise be equally valid to demonstrate our points.

Iterative Sifting

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

Output: For $N = n + k$, the outputs are:

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

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

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

Number of rounds: Stochastic number M, determined by reaching the termination condition (TC) after Step 5 of the loop.

The protocol

Loop phase: Steps 1 to 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 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.

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

$$\mathcal{X}(r) = \{j \in \{1, \ldots, r\} \mid a_j = b_j = 0\}\,,$$
$$\mathcal{Z}(r) = \{j \in \{1, \ldots, r\} \mid a_j = b_j = 1\}$$

TC: If the condition ($\mathcal{X}(r) \geq n$ and $\mathcal{Z}(r) \geq k$) is reached, Alice and Bob set $m := r$ and proceed with Step 6. Otherwise, they increment r by one and repeat from Step 1.

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

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

Step 7: (Order-preserving relabeling): Let r_i be the i-th element of $\mathcal{X} \cup \mathcal{Z}$. Then Alice determines $(s_i)_{i=1}^N \in \{0,1\}^N$, Bob determines $(t_i)_{i=1}^N \in \{0,1\}^N$ and together they determine $(c_i)_{i=1}^N \in \{0,1\}^N$, where for every $i \in \{1, \ldots, N\}$,

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

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

Protocol 1: The iterative sifting protocol.

In the protocol, Alice and Bob iteratively prepares qubit pairs in a maximally entangled state (Step 1) and sends one half of the pair to Bob (Step 2). Then, Alice and Bob each measure their half of the 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, 0 (X) is chosen with probability p_x, and 1 (Z) is chosen with probability p_z, which are parameters of the protocol. The important and problematic part 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 termination condition 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 termination condition 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 to 5 (which we call the *loop phase* of the protocol) until they meet this condition. Because of this iteration, whose termination condition depends on the course of the protocol up to that point, we call it the *iterative sifting protocol*.

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.

As we will see in Section II, the problematic components of the protocol are the interim report (Step 5) and the termination condition (TC). The introduction of a termination condition leads to the problem of non-uniform sampling (Section II B), and the communication in Step 5 leads to the problem that we call basis information leak (Section II C).

Without going into details at this point, we point out that the problem of non-uniform sampling, which comes along with the introduction of a termination condition, arises because of the following convention that is made in the context of iterated sifting.

Convention: *QKD schemes that use the iterated sifting protocol (Protocol 1) as a subroutine for sifting make the following convention for the parameter estimation subroutine: Bits s_i, t_i that result from measuring in the X-basis ($\vartheta_i = 0$) are raw key bits, whereas bits s_i, t_i that result from measuring in the Z-basis ($\vartheta_i = 1$) are test bits.*

This convention means that no additional randomness is injected in the parameter estimation step to decide on which bits are used for the correlation test. The randomness for this *sampling* comes entirely from the probability distribution p_x, p_z (Step 3), together with the termination condition (TC) and the random discarding (Step 6). It is the combination of this convention and the termination condition which is problematic.

What are potential reasons for choosing iterative sifting as the sifting subroutine of a QKD protocol? What speaks for iterated sifting is the fact that the loop phase never runs longer than necessary to achieve the quota n for the raw key bits and the quota k for the test bits. This makes iterated sifting a promising candidate for a sifting routine that is efficient in the sense that it delivers many raw key and test bits per total number of uses of the quantum channel. In fact, schemes that use iterated sifting usually set the probabilities p_x and p_z, which we have left as free parameters, to values that optimize this ratio, namely [2, 7, 10]

$$p_x = \frac{1}{1+\sqrt{\frac{k}{n}}}, \quad p_z = 1 - p_x \,. \tag{1}$$

This efficiency, in turn, means that iterated sifting is a promising candidate for a subroutine of a QKD protocol with a high expected *key rate*, i.e. a protocol with a high number of secret key bits that are generated per use of the quantum channel. The expected key rate being the most important efficiency parameter of a QKD protocol may explain the popularity of iterated sifting in the literature.

B. Remarks on probability calculations for iterative sifting

In Section II, we will calculate the probabilities of certain outputs of Protocol 1. For example, in Section II B 1, we will calculate the probability that Protocol 1 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 $\vartheta = (\vartheta_i)_{i=1}^N \in \{0,1\}^N$, we write $(\vartheta_i)_{i=1}^N = \vartheta_1\vartheta_2\ldots,\vartheta_N$ instead of $(\vartheta_i)_{i=1}^N = (\vartheta_1, \vartheta_2, \ldots, \vartheta_N)$, 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 II B 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 protocol.

To simplify the calculations, it is convenient to introduce the following terminology. For a round r in the loop phase of Protocol 1, we say that 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 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 disagreements in the loop phase, respectively.

For calculations with random variables like Θ, A_1, B_1, M, N_x, N_z or N_d, the sample space of the relevant underlying probability space is the set of all possible courses of Protocol 1. 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 $\mathcal{X}$ and $\mathcal{Z}$ 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.

We write events as logical statements of the random variables, e.g. $\Theta = 110 \wedge N_x \geq 2$ 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 $P[\Theta = 110 \wedge N_x \geq 2]$. 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, e.g. we write

$$P_{\Theta N_x N_z N_d}(\vartheta, n_x, n_z, n_d) := P[\Theta = \vartheta, N_x = n_x, N_z = n_z, N_d = n_d]\,. \tag{2}$$

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

$$P[A_1 \neq B_1] = P_{A_1 B_1}(0,1) + P_{A_1 B_1}(1,0)\,. \tag{3}$$

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

II. WHAT ARE THE PROBLEMS?

A. The origin of the problems: The link between sifting and parameter estimation

As mentioned at the beginning of Section I, the first two subroutines of a QKD protocol are the sifting and the subsequent parameter estimation. As we will describe in more detail in Section IV, protocols that use iterated sifting as their sifting subroutine (and hence are problematic) use a protocol for parameter estimation which uses the bits that result from measurements in the Z-basis for the correlation test and which outputs the bits that come from measurements in the X-basis as raw keys (c.f. the Convention in Section I). To make clear what we are talking about, we have written out such a protocol in Protocol 2.

Protocol 2 basically does the following. Alice and Bob start with the strings $(s_i)_{i=1}^N$, $(t_i)_{i=1}^N$ and $(\vartheta_i)_{i=1}^N$ 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 size of the fraction of the test bits that are anticorrelated between 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 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 $\vartheta_i = 0$.

The idea behind the parameter estimation 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 [2] 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* $p_{\text{tail}}(\mu)$, where for $\mu \in [0,1]$, $p_{\text{tail}}(\mu)$ is given by

$$p_{\text{tail}}(\mu) = P[\Lambda_{\text{key}} \geq \Lambda_{\text{test}} + \mu \mid \Lambda_{\text{test}} \leq q_{\text{tol}}]\,. \tag{4}$$

Here, Λ_{test} is the random variable of the test bit error rate λ_{test} determined in Protocol 2. The random variable Λ_{key} is less easy to understand with the background we have developed this far. We will give a precise formal definition of this random variable in Section IV. At this point here, we give an informal intuition on Λ_{key} instead.

Consider Protocol 1, but with the following modification: Instead of measuring in the X-basis when $a_r = 0$ ($b_r = 0$), Alice and Bob measure in the Z-basis, no matter which value the bit a_r (b_r) has. This way, the measurement choice bits a_r and b_r are the same as in Protocol 1, but the measurement outcome bits are different. Instead of calling them s_r and t_r, let them be denoted by z_r and z'_r, since they are all measured in the Z-basis. For this modified protocol, let Λ_{key} be the random variable of the quantity

$$\lambda_{\text{key}} = \frac{1}{n}\sum_{i=1}^{N}(1-\vartheta_i)(z_i \oplus z'_i)\,. \tag{5}$$

Parameter Estimation

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 $N = n + k$, the inputs are:

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

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

public: N-bit string $(\vartheta_i)_{i=1}^N \in \{0,1\}^N$ 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 $\vartheta_i = 1$.

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

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

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

Step 3: (Raw key output): Let i_j be the j-th element of $\{i \in \{1, \ldots, N\} \mid \vartheta_i = 0\}$. 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

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

Protocol 2: The parameter estimation protocol.

Note that while the bits z_i and z'_i for i with $\vartheta_i = 0$ are different from the bits s_i, t_i of Protocol 1, the bits z_i and z'_i for i with $\vartheta_i = 1$ are the same as the bits s_i and t_i in Protocol 1. Hence, the quantity

$$\lambda_{\text{test}} := \frac{1}{k} \sum_{i=1}^{N} \vartheta_i (s_i \oplus t_i) \tag{6}$$

is the same as the quantity λ_{key} as

$$\lambda_{\text{test}} := \frac{1}{k} \sum_{i=1}^{N} \vartheta_i (z_i \oplus z'_i) \,. \tag{7}$$

Therefore, assume that we run Protocol 1 and then determine λ_{test} as part of the parameter estimation, Protocol 2. Then, the validity of (8) tells us the following: *Given* that correlation test is passed ($\Lambda_{\text{test}} \leq q_{\text{tol}}$), we know what the probability *would* be that $\Lambda_{\text{key}} \geq \Lambda_{\text{test}} + \mu$ *if we had run the modified protocol* (in which everything is measured in Z). The relevance of this knowledge about the random variable Λ_{key} of this hypothetical protocol is that it allows to make statements about $H_{\min}^{\epsilon}(X|E)$, where X is the raw key output of the parameter estimation, Protocol 2, for the output of Protocol 1 (which has *actually been performed*) as its input. The connection between (8) and $H_{\min}^{\epsilon}(X|E)$ is made by virtue of an uncertainty relation for smooth entropies [2, 13], and a statement about the extractability of a secure key from the raw key can then be made by employing the *Quantum Leftover Hash Lemma* [1, 2, 14]. Hence, the validity of (8), which connects parameters determined in the actual protocol with parameters of a hypothetical protocol, means that passing the correlation test really tells us something about the extractability of a secure key from the raw key output of the parameter estimation.

Specifically, the tail probability analysis aims at proving that [2]

$$p_{\text{tail}}(\mu) \leq \frac{\exp\left(-2\frac{kn}{N}\frac{k}{k+1}\mu^2\right)}{p_{\text{pass}}} \,, \tag{8}$$

where

$$p_{\text{pass}} = P[\Lambda_{\text{test}} \geq q_{\text{tol}}] \tag{9}$$

is the probability that the correlation test passes. The derivation of (8) is done using a large deviation bound due to Serfling [15]. For this bound to be correctly applied, two assumptions must be satisfied:

(A1) The sampling probability is uniform,

(A2) The basis choice register is uncorrelated to Alice's and Bob's qubits before measuring.

Assumption (A1) can be expressed as

$$P_\Theta(\vartheta) = P_\Theta(\vartheta') \quad \text{for all } \vartheta, \vartheta' \in \{0,1\}^N \text{ with } \sum_{i=1}^N \vartheta_i = \sum_{i=1}^N \vartheta'_i = k\,. \tag{10}$$

It states that the probability of a subset of the qubits to be test bits (i.e. to be the *sample*) is equally high for all subsets of size k. We can state this assumption by saying that the *sampling probability*, i.e. the probability of picking a subset of the qubits as a sample, is uniform over the set

$$\{0,1\}^N_k := \left\{ (\vartheta_i)_{i=1}^N \in \{0,1\}^N \,\middle|\, \sum_{i=1}^N \vartheta_i = k \right\}. \tag{11}$$

Therefore, we call it the *uniform sampling assumption*. We point out that this is an assumption about the input of Protocol 2, i.e. an assumption about the output of the sifting protocol. Assumption (A2) will be easy to be stated once we have established the framework in which we will do the formal tail probability analysis in Section IV. We will not go into the details of this assumption here. We will give an intuitive picture of (A2) in Section II C and treat it formally in Section IV, where we derive (8) from Assumptions (A1) and (A2). Again, however, what is important about (A2) is that it is not an assumption about the parameter estimation protocol but about the sifting protocol. Hence, although violations of the two assumptions happen in the sifting subroutine of QKD protocols, their effects lie in the failure of the estimate of p_{tail} via the parameter estimation subroutine.

The two problems with iterative sifting lie in its violation of these two assumptions: Protocol 1 violates both (A1) and (A2). For easy reference, we give these two violations some name. We say that a violation of Assumption (A1) leads to the problem of *non-uniform sampling*, and that the violation of Assumption (A2) leads to the problem of *basis information leak*.

B. Non-uniform sampling

1. *Iterative sifting leads to 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. 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 $k \neq n$. Hence, we consider the iterative sifting protocol (Protocol 1) with $n = 1$, $k = 2$ and arbitrary $p_x, p_z \in [0,1]$ with $p_x + p_z = 1$. 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. This is a random variable with possible values 110, 101 and 011.

We are now going to calculate the probability $P_\Theta(110)$ that the protocol outputs a sifted basis choice string which indicates that the first two sifted qubits have been measured in the Z basis and the last sifted qubit has been measured in the X-basis. For this, we would like to remind the reader of the conventions that we made in Section I B, where we explained the meaning of the random variables N_x, N_z and N_d and our conventions for the notation.

Proposition 1: *For Protocol 1 with $n = 1$, $k = 2$, it holds that*

$$P_\Theta(110) = g_z^2\,, \quad \textit{where} \quad g_z = \frac{p_z^2}{p_z^2 + p_x^2}\,. \tag{12}$$

Proof. We have that (explanations for these equalities follow below)

$$P_\Theta(110) = \sum_{n_x=1}^{N}\sum_{n_z=2}^{N}\sum_{n_d=0}^{N} P_{\Theta N_x N_z N_d}(110, n_x, n_z, n_d) \tag{13}$$

$$= \sum_{n_z=2}^{N}\sum_{n_d=0}^{N} P_{\Theta N_x N_z N_d}(110, 1, n_z, n_d) \tag{14}$$

$$= \sum_{n_z=2}^{N}\sum_{n_d=0}^{N} p_x^2 (p_z^2)^k (2p_x p_z)^d \binom{n_z+n_d}{n_d} \tag{15}$$

$$= g_z^2\,, \quad \text{where } g_z = \frac{p_z^2}{p_z^2+p_x^2}\,. \tag{16}$$

Equation (13) is just stating that P_Θ is the marginal of $P_{\Theta N_x N_z N_d}$. Equation (14) follows from

$$P_{\Theta N_x N_z N_d}(110, n_x, n_z, n_d) = 0 \quad \text{for } n_x \geq 2\,. \tag{17}$$

One can see (17) as follows: If $N_x \geq 2$, then necessarily $N_z = 2$, because $N_x > n \wedge N_z > z$ is impossible in Protocol 1. This means that during the random discarding, no Z-agreement gets discarded. Moreover, if $N_x \geq 2$, 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 $N_x \geq 2$, so $\Theta = 110$ is impossible in that case. To see why equation (15) holds, note that the event

$$\Theta = 110 \wedge N_x = 1 \wedge N_z = n_z \wedge N_d = n_d \tag{18}$$

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 because conversely, 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

$$P_{\Theta N_x N_z N_d}(110, 1, n_z, n_d) = \binom{n_z+n_d}{n_d} p_x^2 (p_z^2)^{n_z} (2p_x p_z)^{n_d}\,. \tag{19}$$

This explains equation (15). Finally, equation (16) is just an evaluation of the expression in the line above. This completes the proof. □

For the other two possible values of Θ, it turns out that

$$P_\Theta(011) = P_\Theta(101) = \frac{1-g_z^2}{2}\,. \tag{20}$$

As explained in Section II A, the security analysis of QKD protocols usually assumes that the sampling probability P_Θ is uniform over $\{0,1\}_k^N$, Equations (10) and (11). For this to be true 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

$$g_z^* = \frac{1}{\sqrt{3}}\,, \tag{21}$$

which in turn is equivalent to $p_z = p_z^*$, where

$$p_z^* = \frac{\left(3+2\sqrt{3}\right)\left(1+\sqrt{\sqrt{3}-1}\right)}{\sqrt{3}} \approx 0.539\,. \tag{22}$$

This is bad news for iterative sifting: It means that iterative sifting leads to non-uniform sampling for all values of p_z expect $p_z = p_z^*$. Interestingly, the value of p_z^* does not seem to be a probability that has been considered in the context of QKD and sifting. In particular, p_z^* corresponds to neither the symmetric case $p_z = 1/2$ nor to the probability that has been suggested to be chosen in order to maximize the key rate, c.f. Equation (1).

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 do not have a disagreement (which 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, Protocol 1 leads to

$$P_\Theta(1\ldots10) = g_z^k\,, \quad P_\Theta(\vartheta) = \frac{1-g_z^k}{k} \quad \text{for all other } \vartheta \in \{0,1\}_k^N\,. \tag{23}$$

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

$$g_z^* = \left(\frac{1}{k+1}\right)^{1/k}\,, \tag{24}$$

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

$$p_z^* = \frac{g_z^* - \sqrt{g_z^*(1-g_z^*)}}{2g_z^* - 1}\,. \tag{25}$$

This, in general, corresponds to neither the symmetric case $p_z = p_x$ nor to the probabilities that have been suggested to be chosen in order to maximize the key rate, c.f. Equation (1).

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. Hence, the security analysis of QKD protocols that use iterative sifting as a subroutine needs to be modified, or the iterative sifting needs to be replaced by a sifting scheme that avoids this problem.

2. An attack on non-uniform sampling

How severe is the non-uniform sampling caused by iterated sifting? A decisive answer would require a proper analysis of a full QKD protocol that contains iterated sifting as its sifting subroutine. This analysis would need to answer the question whether a secure key can still be extracted, and if so, what would be the decrease in the key rate compared to uniform sampling. Such an analysis is beyond the scope of the present paper.

To give an idea of the effects of non-uniform sampling nonetheless, we are going to calculate another figure of merit. We devise a strategy for Eve to attack Protocol 1 during its loop phase and calculate the expected value of the *error rate*

$$E = \frac{1}{N}\sum_{i=1}^{N} S_i \oplus T_i \tag{26}$$

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 Protocol 1. We will devise the attack strategy such that Eve particularly makes use of the fact that not every 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 0.25. This is alarming, since intercept-resend attacks ought to introduce an error rate of at least 0.25 if the sifting scheme is secure [16].

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, i.e. they correspond to anticommuting observables. 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, to keep things easy, we want to make this calculation for the easiest possible choice of parameters. Consider Protocol 1 with the parameters $k = n = 1$. From equations (24) and (25), we get that the sampling probabilities in this case are

$$P_\Theta(01) = \frac{p_x^2}{p_x^2+p_z^2}\,, \quad P_\Theta(10) = \frac{p_z^2}{p_x^2+p_z^2}\,. \tag{27}$$

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 $p_z > 1/2$, which makes the sample $\Theta = 10$ more likely than the sample $\Theta = 01$. 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 basis agreement is more likely to be an X-agreement, whereas the second basis agreement is more likely to be a Z-agreement. We will calculate now that the above attack introduces an expected error rate below 0.25 for $1/2 < p_x < 1$, and reaches a minimum of $\langle E\rangle \approx 0.228$ for $p_x \approx 0.73$.

The calculation goes as follows. We first make a split:

$$\langle E\rangle = \sum_{\vartheta} P[\Theta = \vartheta]\,\langle E|\Theta = \vartheta\rangle \tag{28}$$

$$= \underbrace{P[\Theta = 01]\,\langle E|\Theta = 01\rangle}_{\triangle_x} + \underbrace{P[\Theta = 10]\,\langle E|\Theta = 10\rangle}_{\triangle_z}\,. \tag{29}$$

We have that

$$\begin{aligned}\triangle_x = \sum_{n_x=1}^{\infty} \Big(& P[\Theta = 01 \wedge N_x = n_x \wedge A_1 = B_1 = 0]\,\langle E|\Theta = 01 \wedge N_x = n_x \wedge A_1 = B_1 = 0\rangle \\ & + P[\Theta = 01 \wedge N_x = n_x \wedge A_1 \neq B_1]\,\langle E|\Theta = 01 \wedge N_x = n_x \wedge A_1 \neq B_1\rangle \\ & + \underbrace{P[\Theta = 01 \wedge N_x = n_x \wedge A_1 = B_1 = 1]}_{0}\,\langle E|\Theta = 01 \wedge N_x = n_x \wedge A_1 = B_1 = 1\rangle \Big)\end{aligned} \tag{30}$$

$$\begin{aligned}= \sum_{n_x=1}^{\infty} \Bigg(& \sum_{n_d=0}^{\infty} (p_x^2)^{n_x} p_z^2 (2p_x p_z)^{n_d} \binom{n_x + n_d - 1}{n_d} \frac{1}{4} \left(1 - \frac{1}{n_x}\right) \\ & + \sum_{n_d} (p_x^2)^{n_x} p_z^2 (2p_x p_z)^{n_d} \binom{n_x + n_d - 1}{n_x} \frac{1}{4} \Bigg)\end{aligned} \tag{31}$$

$$= \frac{1}{4} \sum_{n_x=1}^{\infty} \sum_{n_d=0}^{\infty} (p_x^2)^{n_x} p_z^2 (2p_x p_z)^{n_d} \left(\binom{n_x + n_d - 1}{n_d} (1 - \frac{1}{n_x}) + \binom{n_x + n_d - 1}{n_x} \right)\,. \tag{32}$$

In a similar way, we get:

$$\triangle_z = \frac{1}{4} \sum_{n_z=1}^{\infty} \sum_{n_d=0}^{\infty} p_x^2 (p_z^2)^{n_z} (2p_x p_z)^{n_d} \left(\binom{n_z + n_d - 1}{n_d} + \binom{n_z + n_d - 1}{n_d} \left(1 + \frac{1}{n_x}\right) \right) \tag{33}$$

Equations (29), (32) and (33) taken together result in

$$\begin{aligned}\langle E\rangle = \sum_{n_d=0}^{\infty} (2p_x p_z)^{n_d} \Bigg(& \sum_{n_x=1}^{\infty} (p_x^2)^{n_x} p_x^2 \left(\binom{n_x + n_d - 1}{n_d} \left(1 - \frac{1}{n_x}\right) + \binom{n_x + n_d - 1}{n_x} \right) \\ & + \sum_{n_z=1}^{\infty} p_x^2 (p_z^2)^{n_z} \left(\binom{n_z + n_d - 1}{n_z} + \binom{n_z + n_d - 1}{n_d} \left(1 + \frac{1}{n_z}\right) \right) \Bigg)\,.\end{aligned} \tag{34}$$

Figure 1 shows a plot of $\langle E\rangle$ as in (34) as a function of p_x. As one can see, $\langle E\rangle$ achieves a minimum of $\langle E\rangle \approx 0.228$ for $p_x \approx 0.73$.

C. Basis information leak

We will now turn our attention to the second problem with iterative sifting. We mentioned in Section II A that the conclusions of the parameter estimation are only valid if (A1) the sampling probability is uniform and (A2) the basis choice register is uncorrelated to Alice's and Bob's qubits before measuring. We have seen in Section II B that iterative sifting violates (A1) and that this leads to an advantage for Eve. In this section, we will see in Section II C 1 that iterative sifting leads to a violation of (A2). (We will explain this more formally in Section IV.) We say that this violation of (A2) leads to a *basis information leak*, because it allows Eve to know which basis agreements Alice and Bob already had during the course of the protocol. We show in Section II C 2 that this allows for an adaptive intercept-resend attack that introduces an error rate of only ≈ 0.163.

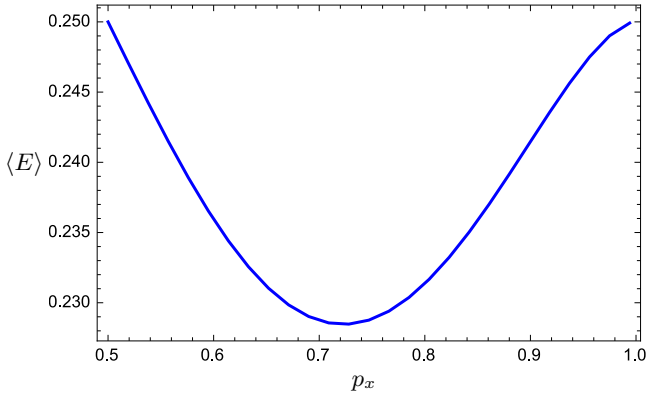

FIG. 1. The error rate induced by Eve's attack on non-uniform sampling.

1. *Iterative sifting leads to a basis information leak*

In iterative sifting, Protocol 1, 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 $a_1 \dots a_r$ and $b_1 \dots b_r$: if they show enough basis agreements, they quit the loop phase; otherwise they keep repeating. What seems to have remained unnoticed in the literature: Eve can also condition her actions on $a_1 \dots a_r$ and $b_1 \dots b_r$. The way we have chosen the quantum communication in Protocol 1, 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 $a_1 \dots a_r$ and $b_1 \dots b_r$. Hence, the state of the qubit that Bob measures is correlated with the classical register that keeps the information about the basis choice. We will state this more formally in Section IV.

2. *An attack on basis information leak*

We will now design an attack strategy for Eve that makes use of the basis information leak. It is an adaptive strategy, in which Eve's action in round $r+1$ depend on the past communication of the strings $a_1 \dots a_r$ and $b_1 \dots b_r$. Again, we make things as simple as possible and consider Protocol 1 with the smallest possible parameters $n = k = 1$. 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, the sampling is uniform,

$$P_\Theta(01) = P_\Theta(10) = \frac{1}{2} \tag{35}$$

(c.f. Equation (27) in Section II B 2).

Before we define Eve's strategy, we want to give some intuition. Suppose that in the course of the protocol, Eve learns that Alice and Bob just had their first basis agreement. If this first agreement is an 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, Alice introduces an error on that test bit. However, there will be a chance that Alice and Bob discard this test bit, because they have a surplus 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 will necessarily tell her the raw key bit, while she has quite some chance to remain undetected. Of course, this intuition does not prove anything, but it is the guideline for designing our attack on basis information leak. For this attack, we can then calculate the error rate, which has more evidential value.

Consider 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 (is equivalent

to $\Theta = 01$), Eve attacks in the Z-basis in all remaining rounds, and if it is a Z-agreement (which is equivalent to $\Theta = 10$), she attacks in the X-basis in all remaining rounds.

As in Section II B 2, let $\langle E \rangle$ be the expected value of the error rate as defined in Equation (26). Again, we assume that the X- and Z-basis are the same for Alice, Bob and Eve and that they correspond to anticommuting observables. Then we have that (explanations below)

$$\begin{aligned}
\langle E \rangle &= P_F(0)\,\langle E|F=0\rangle + P_F(1)\,\langle E|F=1\rangle && (36)\\
&= \langle E|F=0\rangle && (37)\\
&= \underbrace{P_\Theta(01)}_{1/2}\,\langle E|F=0\wedge\Theta=01\rangle + \underbrace{P_\Theta(10)}_{1/2}\,\underbrace{\langle E|F=0\wedge\Theta=10\rangle}_{1/4}\,. && (38)
\end{aligned}$$

Equality (36) is just a decomposition of $\langle E \rangle$ into conditional expectations. Equality (37) 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 be calculated in Equation (38) 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:

$$\begin{aligned}
\langle E|F=0\wedge\Theta=01\rangle &= \sum_{n_x=1}^{\infty} \langle E|F=0\wedge\Theta=01\wedge N_x=n_x\rangle\,\underbrace{P_{N_x|\Theta F}(n_x|01,0)}_{P_{N_x|\Theta}(n_x|01)} && (39)\\
&= \sum_{n_x=1}^{\infty} \underbrace{\langle E|F=0\wedge\Theta=01\wedge N_x=n_x\rangle}_{\frac{n_x-1}{4n_x}}\ \underbrace{P_{N_x\Theta}(n_x,01)}_{\sum_{n_d=0}^{\infty}(p_x^2)^{n_x}p_z^2(2p_xp_z)^{n_d}\binom{n_x+n_d}{n_d}}\ \underbrace{\frac{1}{P_\Theta(01)}}_{2} && (40)\\
&= \sum_{n_x=1}^{\infty} \frac{n_x-1}{2n_x} \sum_{n_d=0}^{\infty} (p_x^2)^{n_x} p_z^2 (2p_xp_z)^{n_d} \binom{n_x+n_d}{n_d} && (41)\\
&= \frac{1}{4}(1-\ln 2)\,, && (42)
\end{aligned}$$

where ln denotes the logarithm to base e. Therefore,

$$\begin{aligned}
\langle E \rangle &= \frac{1}{2}\frac{1}{4}(1-\ln 2) + \frac{1}{2}\frac{1}{4} && (43)\\
&= \frac{2-\ln 2}{8} && (44)\\
&\approx 0.163\,. && (45)
\end{aligned}$$

Hence, the basis information leak allows Eve to go far below the critical expected error rate of 0.25 that one would expect for secure sifting [16].

D. Relation between the two problems

1. *Independence of the two problems*

Are non-uniform sampling and basis information leak really two different problems, or is one of them a consequence of the other? The attack strategy in Section II C 2 shows that basis information leak can occur without the problem of non-uniform sifting, because the attack strategy works even though sampling is uniform. Can non-uniform sifting occur without basis information leak? A closer look at the attack strategy that we devised in Section II B 2 hints that this is possible: The attack strategy works, even though it completely ignores the communication of 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. 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. This leads to a very dramatic form of non-uniform sampling, because $P_\Theta(0\ldots01\ldots1) = 1$ and $P_\Theta(\vartheta) = 0$ for all other $\vartheta \in \{0,1\}_k^N$. 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 indeed two independent problems. They just happen to occur simultaneously for iterated sifting, but they can occur separately in general.

2. An attack that exploits both problems

If the two problems are independent, can we devise an attack strategy for Eve that exploits both of them? This is indeed the case. We again choose the parameters $k = n = 1$ for Protocol 1. Moreover, we assume that $p_x > 1$ to make sure that we have non-uniform sampling. Let Eve's attack be as follows. She begins in the same way as in the attack on non-uniform sampling and attacks 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.

It turns out that for this attack strategy, the expected value of the error rate is

$$\langle E\rangle = \sum_{n_z=1}^{\infty}\sum_{n_d=0}^{\infty} p_x^2 p_z^{2n_z}(2p_xp_z)^{n_d}\binom{n_z+n_d}{n_d}\frac{1}{4} + \sum_{n_x=1}^{\infty}\sum_{n_d=0}^{\infty} p_x^{2n_x} p_z^{2}(2p_xp_z)^{n_d}\binom{n_x+n_d}{n_d}\frac{n_x-1}{4n_x}\,. \tag{46}$$

A plot of (46) is shown in Figure 2 as a function of p_x. As one can see, the expected error rate has a minimum of $\langle E\rangle \approx 0.158$ 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 ≈ 0.228, see Section II B 2) and even better than the attack on the basis information leak alone (with a minimal expected error rate of ≈ 0.163, see Section II C 2).

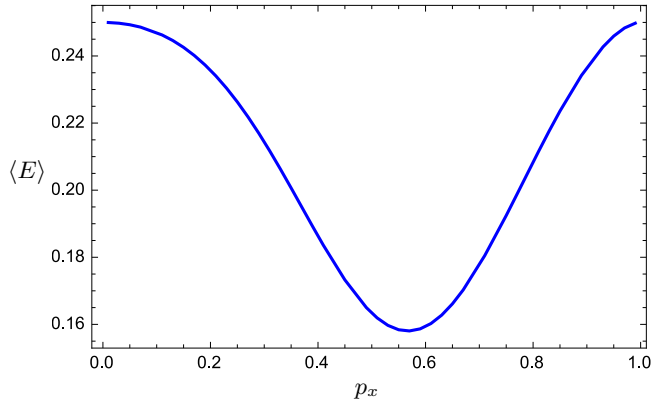

FIG. 2. The error rate induced by Eve's attack on both problems.

III. HOW CAN THESE PROBLEMS BE FIXED? A CANDIDATE PROTOCOL

We have seen that iterative sifting leads to the problems of non-uniform sampling and to a basis information leak. What does this mean for the analysis of QKD protocols? There are two ways to deal with this problem:

(W1) Modify the analysis of the QKD protocols that use iterated sifting as their sifting subroutine,

(W2) Devise new protocols that use sifting subroutines that do not exhibit non-uniform sampling or basis information leak.

Reasons for choosing way (W1), i.e. for sticking to iterated sifting, could be the hope that iterated sifting leads to a high key rates (c.f. the discussion at the end of Section I A). In that case, one would need to find a way to show that it is possible to extract a secure from the raw key without relying on Inequality (8), or find a new way to prove it.

Choosing option (W2) is attractive because the tail probability analysis to arrive at Inequality (8) would not need to be modified. This does not mean that the QKD security analysis can remain unchanged altogether. The new

protocol may require including new security parameters that do not arise for iterated sifting. (We will see below that for the protocol that we propose, this is the case.)

We choose to deal with the problem in way (W2). In Section III A, we propose a protocol that may replace iterated sifting as a sifting subroutine. Unlike the iterated sifting protocol, its loop phase has a fixed number of rounds after which it enters the final phase, without a termination condition that depends on previous rounds. For this reason, we call it the *fixed round number protocol*. We give some evidence on why this may solve the problems in Section III B.

A. Description of the fixed round number protocol

Our approach to designing a sifting protocol that does not exhibit the problems described above is to identify the elements of Protocol 1 that cause the problems and to modify them such that the problems may disappear. As we have seen, it is the communication in Step 5 during the loop phase of Protocol 1 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 termination condition that depends on classical communication. The condition needs to be fixed from the beginning, and it can be nothing but to set the number M of rounds to a fixed number. This number then becomes a parameter of the protocol.

This introduces a new problem: There is no guarantee that after the fixed number $M \geq n + k$, the number N_x of X-agreements is at least n and the number N_z of Z-agreements is at least k. In order to perform the parameter estimation according to Protocol 2, however, these are the minimal numbers of basis agreements that Alice and Bob need to have; otherwise, Inequality (8) is meaningless. Thus, unless one wants to introduce a new tail probability analysis as well (which would make choosing way (W2) somewhat pointless), there is a strictly positive probability that Alice and Bob have to abort the sifting protocol because they have too many basis disagreements.

The protocol that results from the described changes is written out in Protocol 3. The protocol still involves random discarding (Step 6). This is necessary to guarantee that the tail probability analysis that leads to Inequality (8) still goes through. If this would not be done, then the sample size and the total number of bits would be stochastic, which would require a more sophisticated analysis. However, as we will see in Section III B, this makes the computation of the sample probabilities (which is required to prove uniform sampling) more difficult.

B. On the prospect of the fixed round number protocol to solve the problems

Since the fixed round number protocol, Protocol 3, does not involve any communication about the basis choices during its loop phase, it is trivially true that there is no basis information leak. How about uniform sampling? The fact that Protocol 3 involves random discarding makes it difficult to find closed-form expressions for the sampling probabilities $P_\Theta(\vartheta)$ as functions of the parameters k, n, p_x, p_z and M. Instead of attempting this, we give an example calculation. We show that Protocol 3 leads to uniform sampling in the case where $n = k = 1$ and where $p_x, p_z = 1 - p_x \in [0, 1]$ and $M \in \mathbb{N}_+, M \geq n + k$ are arbitrary. Hence, consider Protocol 3, with $n = k = 1$ and with $M \in \mathbb{N}_+$, $M \geq 2$. We are now going to show that for these parameters, $P_\Theta(01) = P_\Theta(10)$.

We have that

$$P_\Theta(10) = \sum_{n_x=0}^{M} \sum_{n_z=0}^{M-n_x} P_{\Theta N_x N_z}(10, n_x, n_z)\,. \tag{47}$$

The difference to the case of Protocol 1 is that now, the sums do not go to infinity but are subject to the fixed number M of total rounds. Moreover, N_d is now fully determined by N_x and N_z, namely $N_d = M - N_x - N_z$. We can rewrite (47) as

$$P_\Theta(10) = \sum_{n_x=0}^{M} \sum_{n_z=0}^{M-n_x} P_{\Theta|N_x N_z}(10|n_x, n_z) P_{N_x N_z}(n_x, n_z)\,, \tag{48}$$

where $P_{\Theta|N_x N_z}(10|n_x, n_z)$ is the probability that Alice end up with $\Theta 10$, *given* that they had n_x X-agreements and n_z Z-agreements during the loop phase of the protocol. The key observation for proving uniform sampling is that

$$P_{\Theta|N_x N_z}(01|n_x n_z) = P_{\Theta|N_x N_z}(10|n_x, n_z)\,. \tag{49}$$

Fixed Round Number Sifting

Protocol Parameters: $n, k \in \mathbb{N}_+$, $p_x, p_z \in [0,1]$ with $p_x + p_z = 1$ and $M \in \mathbb{N}_+$ with $M \geq n + k$.

Output: Either no output (if the protocol aborts in Step 5') or (for $N = n + k$)

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

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

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

Number of rounds: Fixed number M (protocol parameter)

The protocol

Loop phase: Steps 1 to 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

$$\mathcal{X}(M) = \{r \in \{1, \ldots, M\} \mid a_r = b_r = 0\}\,,$$
$$\mathcal{Z}(M) = \{r \in \{1, \ldots, M\} \mid a_r = b_r = 1\}$$

They check whether the quota condition ($\mathcal{X}(M) \geq n$ and $\mathcal{Z}(M) \geq k$) holds. If it holds, they proceed with Step 6. Otherwise, they abort.

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

Step 7: (Order-preserving relabeling): Let r_i be the i-th element of $\mathcal{X} \cup \mathcal{Z}$. Then Alice determines $(s_i)_{i=1}^N \in \{0,1\}^N$, Bob determines $(t_i)_{i=1}^N \in \{0,1\}^N$ and together they determine $(\vartheta_i)_{i=1}^N \in \{0,1\}^N$, where for every $i \in \{1, \ldots, N\}$,

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

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

Protocol 3: The fixed round number sifting protocol.

Equations (48) and (49) imply that

$$P_\Theta(10) = \sum_{n_x=0}^{M} \sum_{n_z=0}^{M-n_x} P_{\Theta|N_x N_z}(01|n_x, n_z) P_{N_x N_z}(n_x, n_z) \tag{50}$$
$$= P_\Theta(01)\,. \tag{51}$$

C. An efficiency comparison with iterative sifting

We define the efficiency η of a sifting protocol as

$$\eta = \frac{R}{M}\,, \tag{52}$$

where R is the number of rounds that are kept after sifting and M is the total number of rounds performed in the loop phase of the protocol. The efficiency η depends on the particular course of the protocol: different runs of the

protocol may have different efficiencies. Therefore, η is a random variable. 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 fixed round number protocol, where M_F is fixed but but R_F is a random variable. (Note that the fixed round number protocol may abort, in which case $R_F = 0$).

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:

$$\langle \eta_I \rangle = \left\langle \frac{R_I}{M_I} \right\rangle \tag{53}$$

$$= (n+k) \left\langle \frac{1}{M_I} \right\rangle \tag{54}$$

$$= (n+k) \sum_{m=n+k}^{\infty} \frac{1}{m} P_{M_I}(m) \tag{55}$$

$$= (n+k) \sum_{m=n+k}^{\infty} \frac{1}{m} \sum_{n_d=0}^{m-n-k} P_{M_I N_d}(m, n_d) \tag{56}$$

$$= (n+k) \sum_{m=n+k}^{\infty} \frac{1}{m} \sum_{n_d=0}^{m-n-k} \left(P_{M_I N_d A_m B_m}(m, n_d, 0, 0) + P_{M_I N_d A_m B_m}(m, n_d, 1, 1) \right) \tag{57}$$

$$= (n+k) \sum_{m=n+k}^{\infty} \frac{1}{m} \sum_{n_d=0}^{m-n-k} \left((p_x^2)^n (p_z^2)^{m-n-d} (2p_x p_z)^{n_d} \binom{m-1}{n_d} \binom{m-n_d-1}{n-1} + \right. \tag{58}$$

$$\left. (p_x^2)^{m-k-d} (p_z^2)^k (2p_x p_z)^{n_d} \binom{m-1}{n_d} \binom{m-n_d-1}{k-1} \right) \tag{59}$$

$$= (n+k) \sum_{m=n+k}^{\infty} \frac{1}{m} (2p_x p_z)^{n_d} \binom{m-1}{n_d} \left((p_x^2)^n (p_z^2)^{m-n-d} \binom{m-n_d-1}{n-1} + (p_x^2)^{m-k-d} (p_z^2)^k \binom{m-n_d-1}{k-1} \right). \tag{60}$$

For the case of the fixed round number protocol, we have:

$$\langle \eta_F \rangle = \frac{R_F}{M_F} \tag{61}$$

$$= \frac{1}{M} \langle R_F \rangle \tag{62}$$

$$= \frac{1}{M} (n+k) P[N_x \geq n \wedge N_z \geq k] \tag{63}$$

$$= \frac{1}{M} (n+k) \sum_{n_d=0}^{M-n-k} P[N_x \geq n \wedge N_z \geq k \wedge N_d = d] \tag{64}$$

$$= \frac{n+k}{M} \sum_{n_d=0}^{M-n-k} \sum_{n_z=k}^{M-n_d-n} P[N_x \geq n \wedge N_z = n_z \wedge N_d = n_d] \tag{65}$$

$$= \frac{n+k}{M} \sum_{n_d=0}^{M-n-k} \sum_{n_z=k}^{M-n_d-n} (p_x^2)^{M-n_z-n_d} (p_x^2)^{n_z} (2p_x p_z)^{n_d} \binom{M}{n_d} \binom{M-n_d}{n_z}. \tag{66}$$

The two values are plotted in Figure 3 for symmetric probabilities as a function of n. The number k is set to be identical to n in these plots. Since the fixed round number protocol depends on an additional parameter M, a choice had to be made in order to plot a value. We chose to optimize over M for each value of n.

IV. FORMAL CRITERIA FOR FUTURE SIFTING PROTOCOLS TO BE SECURE

In Section II, we have seen that iterative sifting leads to problems. In Section III, we have given a partial answer to the question how these problems can be avoided by presenting a candidate protocol and making some example

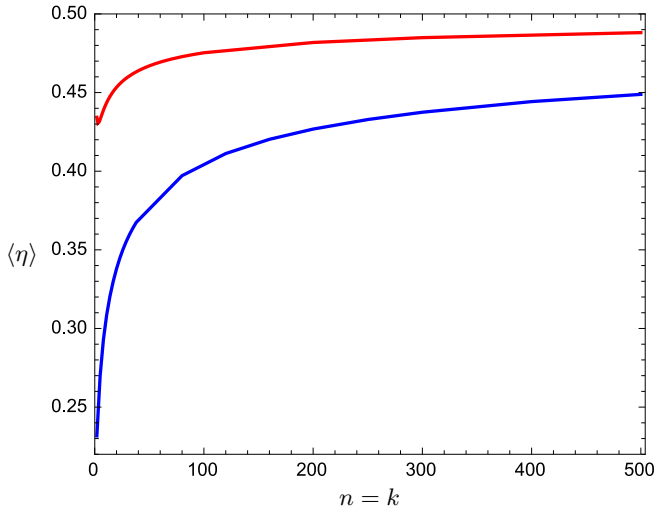

FIG. 3. Efficiency comparison of the two sifting protocols. The plots are calculated for symmetric probabilities $p_x = p_z = 1/2$. The red curve shows the expected value of the efficiency for the iterated sifting protocol as a function of the raw key length n, which is kept at the same value as the number k of test bits. For the fixed round number protocol, the additional parameter M has been optimized for each value of $n = k$.

calculations. In this section, we give a more complete answer to the question by presenting two simple formal criteria that are sufficient for a sifting protocol to lead to a correct parameter estimation. More precisely, we describe two formal properties of the state produced by a sifting protocol which guarantee that if the protocol is followed by the parameter estimation protocol (Protocol 2), then Inequality (8) is true. We prove the sufficiency of these two criteria by deriving (8) from these two criteria.

In order to state the two criteria and the random variable Λ_{key} in (8) formally, we need to define a certain kind of quantum state $\rho_{A^N B^N \Theta^N}$ associated with a sifting protocol. To explain what this state is, we explain what the state $\rho_{A^N B^N \Theta^N}$ is like for Protocol 3. It is a state that is best described in a variation of Protocol 3. Suppose that Alice and Bob run Protocol 3, but they skip the measurement (Step 4) in every round. Instead, they keep each qubit system in their lab without modifying its state. This may be practically difficult, but since $\rho_{A^N B^N \Theta^N}$ is a purely mathematical construct, we do not worry about the technical feasibility. Let us compare the output of this modified protocol with the output of the original Protocol 3:

	original Protocol 3	Protocol 3 without measurements
Alice:	classical bit string $s = (s_i)_{i=1}^N$	N-qubit state ρ_{A^N}
Bob:	classical bit string $t = (t_i)_{i=1}^N$	N-qubit state ρ_{B^N}
public:	classical bit string $\vartheta = (\vartheta_i)_{i=1}^N$	classical bit string $\vartheta = (\vartheta_i)_{i=1}^N$

Hence, if we model the classical bit string ϑ as the state of a classical register Θ^N, we can say that the output of the modified Protocol 3 is a $\mathrm{Q}^N\mathrm{Q}^N\mathrm{C}^N$-state $\rho_{A^N B^N \Theta^N}$. This is the state $\rho_{A^N B^N \Theta^N}$ associated with Protocol 3 that we mean. More generally, the state $\rho_{A^N B^N \Theta^N}$ 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

$$\begin{aligned} &\mathbb{O}_0 = \mathbb{X}_0\,,\ \mathbb{O}_1 = \mathbb{X}_1\,,\ \mathbb{O}_2 = \mathbb{Z}_0\,,\ \mathbb{O}_3 = \mathbb{Z}_1\,, \\ &\mathbb{O}'_0 = \mathbb{X}'_0\,,\ \mathbb{O}'_1 = \mathbb{X}'_1\,,\ \mathbb{O}'_2 = \mathbb{Z}'_0\,,\ \mathbb{O}'_3 = \mathbb{Z}'_1\,. \end{aligned} \tag{67}$$

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

$$P_{ST\Theta}(s,t,\vartheta) = \mathrm{tr}(\Pi(s,t,\vartheta)\rho_{(AB\Theta)^N})\,, \tag{68}$$

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

$$\Pi(s,t,\vartheta) = \bigotimes_{i=1}^{N} \left(\mathbb{O}_{2\vartheta_i + s_i} \otimes \mathbb{O}'_{2\vartheta_i + s_i} \otimes \mathbb{M}_{\vartheta_i} \right) . \tag{69}$$

With the state $\rho_{A^N B^N \Theta^N}$ 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

$$\Omega_{ZZ'\Theta} = \{0,1\}^N \times \{0,1\}^N \times \{0,1\}_k^N \tag{70}$$

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

$$\begin{aligned} P_{ZZ'\Theta} : \quad \Omega_{ZZ'\Theta} &\to [0,1] \\ (z,z',\vartheta) &\mapsto \operatorname{tr}\left(\left(\bigotimes_{i=1}^N \mathbb{Z}_{z_i} \right) \otimes \left(\bigotimes_{i=1}^N \mathbb{Z}'_{z'_i} \right) \otimes \left(\bigotimes_{i=1}^N \mathbb{M}_{\vartheta_i} \right) \rho_{(A^N B^N \Theta)^N} \right) . \end{aligned} \tag{71}$$

Now we are able to formally say what the random variable Λ_{key} of a sifting protocol is. Let $\rho_{A^N B^N \Theta^N}$ be the state associated with the sifting protocol, let $(\Omega_{ZZ'\Theta}, P_{ZZ'\Theta})$ be the probability space as in Equations (70) and (71). Then Λ_{key} is the random variable

$$\begin{aligned} \Lambda_{\mathrm{key}} : \quad \Omega_{ZZ'\Theta} &\to [0,1] \\ (z,z',\vartheta) &\mapsto \frac{1}{n} \sum_{i=1}^{N} (1-\vartheta_i)(z \oplus z') , \end{aligned} \tag{72}$$

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

$$\begin{aligned} \Lambda_{\mathrm{test}} : \quad \Omega_{ZZ'\Theta} &\to [0,1] \\ (z,z',\vartheta) &\mapsto \frac{1}{k} \sum_{i=1}^{N} \vartheta_i (z \oplus z') . \end{aligned} \tag{73}$$

This allows us to define the *tail probability*

$$p_{\mathrm{tail}}(\mu) = P[\Lambda_{\mathrm{key}} \geq \Lambda_{\mathrm{test}} + \mu \mid \Lambda_{\mathrm{test}} \leq q_{\mathrm{tol}}] . \tag{74}$$

The following proposition makes the formal statement of the tail probability analysis.

Proposition 2 (Tail probability estimate)**:** *Let $\rho_{A^N B^N \Theta^N}$ be a $Q^N Q^N C^N$-state, let $\{\mathbb{Z}_0, \mathbb{Z}_1\}$ and $\{\mathbb{Z}'_0, \mathbb{Z}'_1\}$ be POVMs on the quantum systems A and B, let $\{\mathbb{M}_0, \mathbb{M}_1\}$ be the read-out measurement of the classical system Θ, let Λ_{key}, Λ_{test} be random variables on the discrete probability space $(\Omega_{ZZ'\Theta}, P_{ZZ'\Theta})$ as defined in Equations (70) to (73) and let p_{tail} be as in Equation (74). Let $\rho_{A^N B^N}$ and ρ_{Θ^N} denote the according reduced states of $\rho_{A^N B^N \Theta^N}$ and P_Θ denote the according marginal of $P_{ZZ'\Theta}$. If the two conditions*

$$P_\Theta(\vartheta) = \binom{N}{k}^{-1} \quad \forall\, \vartheta \in \{0,1\}_k^N \quad \textit{and} \tag{75}$$

$$\rho_{A^N B^N \Theta^N} = \rho_{A^N B^N} \otimes \rho_{\Theta^N} \tag{76}$$

hold, then

$$p_{\mathrm{tail}}(\mu) \leq \frac{\exp\left(-2 \frac{kn}{N} \frac{k}{k+1} \mu^2 \right)}{p_{\mathrm{pass}}} , \tag{77}$$

where

$$p_{\mathrm{pass}} = P[\Lambda_{\mathrm{test}} \geq q_{\mathrm{tol}}] . \tag{78}$$

The formulation of Proposition 2 allows us to see the formal requirements on a sifting protocol to lead to a correct parameter estimation: Condition (75) is exactly the statement that the protocol leads to uniform sampling. The condition reads the same as before, the only difference being that now we have formulated the probability distribution P_Θ in terms of the state associated with the protocol. Condition (76) 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 2 states that if these two conditions are satisfied, then the correlation test of the parameter estimation 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 whole digression to a classical probability space, Equations (70) to (74), is a mere change of notation. However, the fact that it is possible to express Proposition 2 in terms of a classical probability space shows that this part of a QKD security analysis is purely classical.

We now turn to the proof.

Proof of Proposition 2. According to Baye's Theorem, we have that

$$p_{\text{tail}} = P[\Lambda_{\text{key}} \geq \Lambda_{\text{test}} + \mu \mid \Lambda_{\text{test}} \leq q_{\text{tol}}] \tag{79}$$

$$= \frac{P[\Lambda_{\text{test}} \leq q_{\text{tol}} \mid \Lambda_{\text{key}} \geq \Lambda_{\text{test}} + \mu] P[\Lambda_{\text{key}} \geq \Lambda_{\text{test}} + \mu]}{P[\Lambda_{\text{test}} \leq q_{\text{tol}}]} \tag{80}$$

$$\leq \frac{P[\Lambda_{\text{key}} \geq \Lambda_{\text{test}} + \mu]}{p_{\text{pass}}} \,. \tag{81}$$

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

$$\begin{aligned} \Lambda_{\text{tot}} : \quad \Omega_{ZZ'\Theta} &\to [0,1] \\ (z, z', \vartheta) &\mapsto \frac{1}{N} \sum_{i=1}^{N} z \oplus z' \,. \end{aligned} \tag{82}$$

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

$$\Lambda_{\text{key}}(z, z, \vartheta) \geq \Lambda_{\text{test}}(z, z, \vartheta) + \mu \tag{83}$$

$$\iff \frac{1}{n} \sum_{i=1}^{N} (1 - \vartheta_i)(z_i \oplus z'_i) \geq \frac{1}{k} \sum_{i=1}^{N} \vartheta_i (z_i \oplus z'_i) + \mu \tag{84}$$

$$\iff \frac{1}{n} \sum_{i=1}^{N} (1 - \vartheta_i)(z_i \oplus z'_i) + \frac{1}{k} \sum_{i=1}^{N} (1 - \vartheta_i)(z_i \oplus z'_i) \geq \frac{1}{k} \sum_{i=1}^{N} \vartheta_i (z_i \oplus z'_i) + \frac{1}{k} \sum_{i=1}^{N} (1 - \vartheta_i)(z_i \oplus z'_i) + \mu \tag{85}$$

$$\iff \left(\frac{1}{n} + \frac{1}{k} \right) \sum_{i=1}^{N} (1 - \vartheta_i)(z_i \oplus z'_i) \geq \frac{1}{k} \sum_{i=1}^{N} (z_i \oplus z'_i) + \mu \tag{86}$$

$$\iff \frac{k}{N} \left(\frac{1}{n} + \frac{1}{k} \right) \sum_{i=1}^{N} (1 - \vartheta_i)(z_i \oplus z'_i) \geq \frac{k}{N} \frac{1}{k} \sum_{i=1}^{N} (z_i \oplus z'_i) + \frac{k}{N} \mu \tag{87}$$

$$\iff \Lambda_{\text{key}}(z, z, \vartheta) \geq \Lambda_{\text{tot}}(z, z, \vartheta) + \frac{k}{N} \mu \,. \tag{88}$$

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},

$$\Sigma_{\text{key}} = n \Lambda_{\text{key}} \,, \quad \Sigma_{\text{test}} = k \Lambda_{\text{test}} \,, \quad \Sigma_{\text{tot}} = N \Lambda_{\text{tot}} \,. \tag{89}$$

This gives us

$$\Lambda_{\text{key}} \geq \Lambda_{\text{tot}} + \frac{k}{N} \mu \iff \Sigma_{\text{key}} \geq n \left(\frac{\Sigma_{\text{tot}}}{N} + \frac{N - n}{N} \mu \right) \tag{90}$$

Therefore,

$$P[\Lambda_{\text{key}} \geq \Lambda_{\text{test}} + \mu] = P\left[\Sigma_{\text{key}} \geq n \left(\tfrac{\Sigma_{\text{tot}}}{N} + \tfrac{N-n}{N} \mu \right) \right] \tag{91}$$

and hence

$$p_{\text{tail}} \leq \frac{P\left[\Sigma_{\text{key}} \geq n\left(\frac{\Sigma_{\text{tot}}}{N} + \frac{N-n}{N}\mu\right)\right]}{p_{\text{pass}}} \tag{92}$$

$$= \frac{\sum_{\sigma_{\text{tot}}} P[\Sigma_{\text{tot}} = \sigma_{\text{tot}}] P\left[\Sigma_{\text{key}} \geq n\left(\frac{\sigma_{\text{tot}}}{N} + \frac{N-n}{N}\mu\right) \,\middle|\, \Sigma_{\text{tot}} = \sigma_{\text{tot}}\right]}{p_{\text{pass}}} \tag{93}$$

$$= \frac{\sum_{\sigma_{\text{tot}}} P[\Sigma_{\text{tot}} = \sigma_{\text{tot}}] \sum_j P\left[\Sigma_{\text{key}} = j \,|\, \Sigma_{\text{tot}} = \sigma_{\text{tot}}\right]}{p_{\text{pass}}}, \tag{94}$$

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

$$j = \left\lceil n\left(\frac{\sigma_{\text{tot}}}{N} + \frac{N-n}{N}\mu\right)\right\rceil, \left\lceil n\left(\frac{\sigma_{\text{tot}}}{N} + \frac{N-n}{N}\mu\right)\right\rceil + 1, \ldots, n, \tag{95}$$

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

$$h(\sigma_{\text{tot}}, N, n, j) := P\left[\Sigma_{\text{key}} = j \,|\, \Sigma_{\text{tot}} = \sigma_{\text{tot}}\right] \tag{96}$$

$$= \frac{P[\Sigma_{\text{key}} = j \wedge \Sigma_{\text{tot}} = \sigma_{\text{tot}}]}{P[\Sigma_{\text{tot}} = \sigma_{\text{tot}}]} \tag{97}$$

$$= \frac{P[\Omega_{j\sigma_{\text{tot}}}]}{P[\Omega_{\sigma_{\text{tot}}}]}, \tag{98}$$

$$= \frac{\sum_{(z,z',\vartheta)\in\Omega_{j\sigma_{\text{tot}}}} P_{ZZ'\Theta}(z,z',\vartheta)}{\sum_{(z,z',\vartheta)\in\Omega_{\sigma_{\text{tot}}}} P_{ZZ'\Theta}(z,z',\vartheta)} \tag{99}$$

where

$$\Omega_{j\sigma_{\text{tot}}} = \{(z,z',\vartheta) \in \Omega_{ZZ'\Theta} \mid \Sigma_{\text{key}}(z,z',\vartheta) = j \wedge \Sigma_{\text{tot}}(z,z',\vartheta) = \sigma_{\text{tot}}\}, \tag{100}$$

$$\Omega_{\sigma_{\text{tot}}} = \{(z,z',\vartheta) \in \Omega_{ZZ'\Theta} \mid \Sigma_{\text{tot}}(z,z',\vartheta) = \sigma_{\text{tot}}\}. \tag{101}$$

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

$$P_{ZZ'\Theta}(z,z',\vartheta) = P_{ZZ'}(z,z')P_\Theta(\vartheta) \tag{102}$$

$$= P_{ZZ'}(z,z')\binom{N}{k}^{-1}, \tag{103}$$

where $P_{ZZ'}$ and P_Θ are the according marginal distributions of $P_{ZZ'\Theta}$. Equation (102) follows form (76), and Equation (103) follows from Equation (75). This implies

$$h(\sigma_{\text{tot}}, N, n, j) = \frac{\sum_{(z,z',\vartheta)\in\Omega_{j\sigma_{\text{tot}}}} P_{ZZ'}(z,z')\binom{N}{k}^{-1}}{\sum_{(z,z',\vartheta)\in\Omega_{\sigma_{\text{tot}}}} P_{ZZ'}(z,z')\binom{N}{k}^{-1}} \tag{104}$$

$$= \frac{\sum_{(z,z',\vartheta)\in\Omega_{j\sigma_{\text{tot}}}} P_{ZZ'}(z,z')}{\sum_{(z,z',\vartheta)\in\Omega_{\sigma_{\text{tot}}}} P_{ZZ'}(z,z')} \tag{105}$$

$$= \frac{\sum_{(z,z')\in\Gamma_{\sigma_{\text{tot}}}} P_{ZZ'}(z,z')\binom{\sigma_{\text{tot}}}{j}\binom{N-\sigma_{\text{tot}}}{n-j}}{\sum_{(z,z')\in\Gamma_{\sigma_{\text{tot}}}} P_{ZZ'}(z,z')\binom{N}{n}} \tag{106}$$

$$= \binom{\sigma_{\text{tot}}}{j}\binom{N-\sigma_{\text{tot}}}{n-j}\binom{N}{n}^{-1}, \tag{107}$$

where

$$\Gamma_{\sigma_{\text{tot}}} = \left\{(z,z') \in \{0,1\}^N \times \{0,1\}^N \,\middle|\, \sum_{i=1}^N z_i \oplus z'_i = \sigma_{\text{tot}}\right\}. \tag{108}$$

Equation (107) means that $h(\sigma_{\text{tot}}, N, n, j)$ is a hypergeometric distribution. We are interested in the *tail* of this distribution,

$$H(\sigma_{\text{tot}}, N, n, l) := \sum_{j=l}^{n} h(\sigma_{\text{tot}}, N, n, j)\,, \tag{109}$$

because according to Equations (94) and (95),

$$p_{\text{tail}} \leq \frac{\sum_{\sigma_{\text{tot}}} P[\Sigma_{\text{tot}} = \sigma_{\text{tot}}] H(\sigma_{\text{tot}}, N, n, l)}{p_{\text{pass}}}\,, \tag{110}$$

where

$$l = \left\lceil n \left(\frac{\sigma_{\text{tot}}}{N} + \frac{N-n}{N}\mu \right) \right\rceil . \tag{111}$$

There are several well-known bounds on the tail of a hypergeometric distribution. For our case, *Serfling's bound* is a suitable one [15]. The appropriate special case of Serfling's bound for this case reads

$$H(\sigma_{\text{tot}}, N, n, l) \leq \exp\left(-2\frac{(N-n)n}{N} \frac{N-n}{N-n+1} \mu^2 \right) \tag{112}$$

$$= \exp\left(-2\frac{kn}{N} \frac{k}{k+1} \mu^2 \right) . \tag{113}$$

(Instead of Serfling's bound, one may use *Hoeffding's bound* [17]. 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 [18, 19], so these references are easier to understand in our context.) Inequalities (110) and (113) together imply

$$p_{\text{tail}} \leq \frac{\sum_{\sigma_{\text{tot}}} P[\Sigma_{\text{tot}} = \sigma_{\text{tot}}] H(\sigma_{\text{tot}}, N, n, l)}{p_{\text{pass}}} \tag{114}$$

$$\leq \frac{\exp\left(-2\frac{kn}{N} \frac{k}{k+1} \mu^2 \right)}{p_{\text{pass}}}\,, \tag{115}$$

which completes the proof. □

ACKNOWLEDGMENTS

We would like to thank Marco Tomamichel and David Elkouss for insightful discussions.

[1] Renato Renner. Security of Quantum Key Distribution. (16242), December 2005.
[2] Marco Tomamichel, Charles Ci Wen Lim, Nicolas Gisin, and Renato Renner. Tight finite-key analysis for quantum cryptography. *Nat. Commun.*, 3(634), 2012.
[3] Charles Ci Wen Lim, Marcos Curty, Nino Walenta, Feihu Xu, and Hugo Zbinden. Concise security bounds for practical decoy-state quantum key distribution. *Phys. Rev. A*, 89(2):022307, 2014.
[4] Masahito Hayashi and Toyohiro Tsurumaru. Simple and Tight Security Analysis of the Bennett-Brassard 1984 Protocol with Finite Key Lengths. 093014:9, 2011.
[5] Masahito Hayashi and Ryota Nakayama. Security analysis of the decoy method with the Bennett-Brassard 1984 protocol for finite key lengths. *New Journal of Physics*, 16, 2014.
[6] Marcos Curty, Feihu Xu, Wei Cui, Charles Ci Wen Lim, Kiyoshi Tamaki, and Hoi-Kwong Lo. Finite-key analysis for measurement-device-independent quantum key distribution. *Nat. Commun.*, 5, 2014.
[7] Charles Ci Wen Lim, Christopher Portmann, Marco Tomamichel, Renato Renner, and Nicolas Gisin. Device-independent quantum key distribution with local bell test. *Phys. Rev. X*, 3(3):031006, 2013.
[8] Fabian Furrer, Torsten Franz, Mario Berta, Anthony Leverrier, Volkher B. Scholz, Marco Tomamichel, and Reinhard F. Werner. Continuous Variable Quantum Key Distribution: Finite-Key Analysis of Composable Security against Coherent Attacks. pages 1–10, December 2011.

[9] Anthony Leverrier. Composable security proof for continuous-variable quantum key distribution with coherent states. August 2014.
[10] Davide Bacco, Matteo Canale, Nicola Laurenti, Giuseppe Vallone, and Paolo Villoresi. Experimental quantum key distribution with finite-key security analysis for noisy channels. *Nat. Commun.*, 4, 2013.
[11] Feihu Xu, Shihan Sajeed, Sarah Kaiser, Zhiyuan Tang, Li Qian, Vadim Makarov, and Hoi-Kwong Lo. Experimental quantum key distribution with source flaws and tight finite-key analysis. pages 1–11, August 2014.
[12] Boris Korzh, Charles Ci Wen Lim, Raphael Houlmann, Nicolas Gisin, Ming Jun Li, Daniel Nolan, Bruno Sanguinetti, Rob Thew, and Hugo Zbinden. Provably Secure and Practical Quantum Key Distribution over 307 km of Optical Fibre. *Nature Photonics*, 9(3):7, 2014.
[13] Marco Tomamichel and Renato Renner. Uncertainty relation for smooth entropies. *Phys. Rev. Lett.*, 106(11):110506, 2011.
[14] Marco Tomamichel, Christian Schaffner, Adam Smith, and Renato Renner. Leftover hashing against quantum side information. *IEEE Trans. Inf. Theory*, 57(8):5524–5535, 2011.
[15] Robert J. Serfling. Probability inequalities for the sum in sampling without replacement. *Ann. Stat.*, 2(1):39–48, 1974.
[16] Valerio Scarani, Helle Bechmann-Pasquinucci, Nicolas J. Cerf, Miloslav Dušek, Norbert Lütkenhaus, and Momtchil Peev. The security of practical quantum key distribution. *Rev. Mod. Phys.*, 81(3):1301–1350, 2009.
[17] Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *J. Am. Statist. Assoc.*, 58(301):13–30, 1963.
[18] Václav Chvátal. The tail of the hypergeometric distribution. *Discrete Math.*, 25(3):285–287, 1978.
[19] Matthew Skala. Hypergeometric tail inequalities: ending the insanity. 2013.