

# Incorporating Device Characterization into Security Proofs

Ernest Y.-Z. Tan<sup>1</sup> and Shlok Nahar<sup>1\*</sup>

*Institute for Quantum Computing and Department of Physics and Astronomy, [University of Waterloo](#), Waterloo, Ontario N2L 3G1, Canada*



(Received 25 November 2025; revised 22 February 2026; accepted 25 March 2026; published 29 May 2026)

Typical security proofs for quantum key distribution (QKD) rely on having some model for the devices, with the security guarantees implicitly relying on the values of various parameters of the model, such as dark count rates or detector efficiencies. Hence to deploy QKD in practice, we must establish how to certify or characterize the model parameters of a manufacturer's QKD devices. We present a rigorous framework for analyzing such procedures, laying out concrete requirements for both the security proofs and the certification or characterization procedures. In doing so, we describe various forms of conclusions that can and cannot be validly drawn from such procedures, addressing some potential misconceptions. We also discuss connections to composable security frameworks and some technical aspects that remain to be resolved in that direction.

DOI: [10.1103/PRX.7.020342](https://doi.org/10.1103/PRX.7.020342)

## I. INTRODUCTION

Quantum cryptography is a promising prospect for applications of quantum technology, allowing for the realization of cryptographic tasks that would be classically impossible. For instance, a prominent form of quantum cryptography is quantum key distribution (QKD), which distributes secret keys with information-theoretic security between two honest parties Alice and Bob, even in the presence of an eavesdropper Eve. Accordingly, security proofs have been developed for a wide variety of protocols in quantum cryptography, aiming to ensure that they can indeed be safely deployed. For this work, we shall focus our discussion mainly on QKD as a central example, though we note that the overall principles should broadly generalize to many other classes of protocols.

Importantly, such security proofs are only valid for particular classes of models for the device used to implement the corresponding protocol. (For QKD protocols in particular, it might be more accurate to refer to a *pair* of devices held by Alice and Bob, but for brevity in this work, we shall often simply refer to it as a single device.) Given this, we should account for the fact that real-life devices would usually have various imperfections as compared to some idealized “perfect” behavior. If these imperfections are sufficiently extreme, lying outside the class of

device models covered by the security proof, they might even cause a protocol performed with such a device to be completely insecure. As an example, consider a decoy-state QKD protocol with an imperfectly phase-randomized laser as described in Ref. [1, Eq. (5)], focusing on the high-loss regime. If the phase randomization is perfect or sufficiently close to perfect, the protocol would securely yield nonzero key rate even in the high-loss regime. However, in the “maximally imperfect” phase-coherent case, Alice sends out linearly independent states, thus rendering the protocol completely insecure against an unambiguous state discrimination attack at high loss.

Hence it is important to introduce some process by which we could try to prevent the deployment of QKD devices that do not satisfy the model assumptions required for the security proofs to hold. More specifically, we could aim to achieve this via some notion of a **certification procedure**, which either approves the device for use in QKD protocols or otherwise rejects the device [2]. In order for this certification procedure to be practical, we would want to phrase it in terms of a tractable number of **parameters** that capture “all relevant device properties” for a QKD security proof. (We discuss various considerations for choosing these parameters in the main body of this work.) Informally, we expect that the certification procedure should bound the values of these parameters, and then accept or reject the device based on whether these bounds seem “good enough” for us to prove the resulting QKD protocols are secure. While this has some similarities to the notions of “parameter estimation” or “acceptance testing” in QKD protocols [3], the conceptual distinction here is that it involves the device that will be used by the honest parties Alice and Bob, rather than the operations

\*Contact author: [sanahar@uwaterloo.ca](mailto:sanahar@uwaterloo.ca)

*Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](#) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.*

that Eve performs on states transmitted during the QKD protocol.

### A. Fundamental obstacles in formalizing certification properties

If a certification procedure could *guarantee* a bound on some parameter with “no probabilistic uncertainty” (put another way, if it gives a bound that holds with probability 1), then there would be no issues with simply saying that the certification procedure only approves a device if the bound on that parameter falls within some range for which we can prove security of the resulting QKD protocols. Unfortunately, demanding such a guarantee is usually too stringent a requirement when considering most types of device parameters, especially in the context of quantum devices that have some “inherent” random behavior.

More realistically, one would usually only be able to make the weaker statement of constructing a **confidence interval** for some parameter, which contains the true value of the parameter with high probability (we defer a formal discussion of confidence intervals to the main body of this work). In such circumstances, it is still reasonable to suppose that the certification procedure approves the device whenever the confidence interval falls within some suitable range of values. Critically, however, here there is some small but nonzero probability that the confidence interval fails to correctly capture the true parameter value, and a full security framework should rigorously account for this fact in some form.

It is tempting to suggest that in the presence of such probabilistic uncertainty, one could simply claim that *conditioned* on the certification procedure accepting, with high probability the QKD protocols using the accepted device will be secure. Unfortunately, such a claim is fundamentally impossible in general—a very similar issue arises regarding the standard security criterion for QKD, when discussing the states conditioned on the protocol accepting; see Ref. [4, Sec. 4.3]. To see this, consider the case where the device undergoing the certification procedure is genuinely insecure (in the sense that any QKD keys it produces can be perfectly guessed by Eve), for instance the aforementioned example of a laser emitting fully phase-coherent states. Due to the probabilistic nature of the certification procedure, there will be a nonzero probability of approving such a device—for a well-designed certification procedure, this probability would presumably be extremely small, but it is not literally zero. Importantly, the device is *still* fundamentally insecure even when conditioned on the certification procedure approving; it is not valid to claim that it is “secure with high probability conditioned on approval.”

Put another way, such a claim would be a standard mistake regarding conflation of conditional probabilities: even given that the probability of approval conditioned on the

device being insecure is small, this does *not* imply that the probability of the device being insecure conditioned on approval is small. Connecting these probabilities rigorously would require an explicit choice of a **Bayesian prior** regarding the device behavior, which does not feature within the usual frameworks of quantum cryptography.

### B. Our contributions

In light of the above, a more careful treatment of probabilistic uncertainty in the certification procedure is necessary. This is the main purpose of this work: we lay out a framework via which this issue can be properly addressed, and establish rigorous statements regarding the final states produced when a QKD device is subject to such a certification procedure followed by multiple QKD protocol instances (in the event of approval). In the process, we obtain clear guidelines for what is required from both the certification procedures and the QKD security proofs, which we informally summarize as follows:

- (1) The certification procedure should construct confidence intervals for all parameters “relevant to” a security proof, and reject the device if any of these confidence intervals are not contained within some predesignated ranges.
- (2) The security proofs should have some “robustness” to imperfections, in the sense that they must ensure security holds for all devices with parameters lying within some nontrivial ranges, rather than only idealized devices with some “perfect” values for those parameters.

We give a more detailed and rigorous description of these criteria in the main body of this work, particularly Proposition 1 and Secs. III and IV. As long as these criteria are satisfied, we can obtain useful conclusions about the states produced by the QKD protocol instances, though we re-emphasize that due to the fundamental obstacles mentioned above, it is *not* valid to suppose they yield statements of the form (for instance) “conditioned on the certification procedure approving, the probability that Eve can guess the final keys is small.”

This paper is structured as follows. In Sec. II, we lay out our framework for describing certification procedures and QKD protocols, and present our main result (Proposition 1) regarding rigorous statements that can be made for the states produced at the end of such a process. In Sec. III, we discuss what this implies in terms of what is required in the certification procedure, and conversely, in Sec. IV we discuss what is required from the QKD security proofs. In Sec. V, we provide a pedagogical example that illustrates the application of our work to characterizing dark count rates in a QKD detection setup. Finally, in Sec. VI we summarize our findings. Some technical details and variants are deferred to the appendixes, such as the possibilities

of “adaptive” protocol choices based on characterizing device parameters (Appendix B) and connections to the framework of Abstract Cryptography (Appendix C).

## II. PROPOSED FRAMEWORK

We will use  $d(\rho, \sigma) := \frac{1}{2} \|\rho - \sigma\|_1$  to denote the **trace distance** between any two states  $\rho, \sigma$ , as described in Ref. [5]. All states should be understood to be normalized states, including when conditioned on some event. In this work, to denote a quantum channel  $\mathcal{E}$  from register  $Q$  to register  $Q'$ ; we will often use the abbreviated notation  $\mathcal{E} : Q \rightarrow Q'$  rather than writing out the formal statement of it being a completely positive trace-preserving map [5] from states on  $Q$  to states on  $Q'$ . We will also leave tensor products with identity channels implicit; e.g., for a channel  $\mathcal{E} : Q \rightarrow Q'$ , we often use the compact notation  $\mathcal{E}[\rho_{QR}] := (\mathcal{E} \otimes \mathcal{I}_R)[\rho_{QR}]$ , where  $\mathcal{I}_R$  is an identity channel on  $R$ .

### A. A prelude: Existing QKD security definition

Currently, many QKD security proofs are centered around showing that a particular trace-distance security definition holds, as discussed in Refs. [6–8]. This security definition can be formulated in a number of equivalent ways; here we shall present a formulation convenient for our work.

Specifically, first suppose at the end of a QKD protocol, let  $K^A$  and  $K^B$  denote the final keys obtained by Alice and Bob, respectively, and let  $E^{\text{final}}$  denote all registers that may potentially be available to an adversary in the future. We shall allow the length of the final keys  $K^A K^B$  to be picked during the protocol, i.e., we allow for what have been called **variable-length** or **adaptive-length** protocols [6–8]. While many existing QKD proofs have instead focused on **fixed-length** protocols that simply make an accept/abort decision and output keys of fixed length when they accept, these can be viewed as just a special case of variable-length protocols by viewing the abort outcome as a key of zero length.

Given the above register notation, we moreover need a formalism to describe possible attacks by Eve. For the purposes of this work, we shall take a fairly “minimalist” perspective that should apply to a broad range of formalisms—we only require that for any particular attack  $\mathcal{A}$ , there exists a corresponding quantum channel  $\mathcal{E}^{\mathcal{A}} : Q^{\text{initial}} \rightarrow K^A K^B E^{\text{final}}$  describing how that attack transforms any state on some initial register  $Q^{\text{initial}}$ , describing “all relevant registers” at the start of the protocol, to some final state on  $K^A K^B E^{\text{final}}$ , where  $E^{\text{final}}$  denotes all registers Eve holds at the end of the protocol. We emphasize we do not place any restrictions on the “type of mathematical object”  $\mathcal{A}$  is formalized as, leaving it up to the relevant use case (it could for instance be a channel, a quantum comb [9], a causal box [10], or perhaps some even more general

structure), other than this fairly minimal requirement. In particular, we are not necessarily *identifying* the attacks  $\mathcal{A}$  with the channels  $\mathcal{E}^{\mathcal{A}}$ , in that, e.g., we do not require the mapping  $\mathcal{A} \rightarrow \mathcal{E}^{\mathcal{A}}$  to be bijective—this allows  $\mathcal{A}$  to potentially be an element of some much “larger” set or class than the set of quantum channels. Note that the channel  $\mathcal{E}^{\mathcal{A}}$  is intended to describe “all physical processes” during the protocol, including actions by honest parties such as state preparations and measurements for prepare-and-measure protocols.

With this in mind, the standard trace-distance security definition can be phrased as saying that a QKD protocol is  $\epsilon^{\text{QKD}}$ -secure (also sometimes called  $\epsilon^{\text{QKD}}$ -sound [8]) if for any attack  $\mathcal{A}$  and any initial state  $\sigma$ , we have

$$d(\mathcal{E}^{\mathcal{A}}[\sigma_{Q^{\text{initial}}}], \mathcal{K} \circ \mathcal{E}^{\mathcal{A}}[\sigma_{Q^{\text{initial}}}] ) \leq \epsilon^{\text{QKD}}, \quad (1)$$

where  $\mathcal{K}$  is a channel that reads the length  $\ell$  of the keys on registers  $K^A K^B$  and replaces them with perfectly correlated uniform keys of *that length*, i.e., the state

$$\omega_{K^A K^B}^{\ell} := \sum_{k \in \{0,1\}^{\ell}} 2^{-\ell} |kk\rangle\langle kk|_{K^A K^B}. \quad (2)$$

The above security definition is sometimes presented with an additional purifying register than the channel  $\mathcal{E}^{\mathcal{A}}$  does not act on, viewing this additional purification as being held by Eve. For the purposes of this work, however, we shall view any such purification as implicitly being included in the domain of the channel  $\mathcal{E}^{\mathcal{A}}$ , simply requiring that  $\mathcal{E}^{\mathcal{A}}$  acts as the identity on that register.

*Remark 1.* Informally, the state  $\mathcal{K} \circ \mathcal{E}^{\mathcal{A}}[\sigma_{Q^{\text{initial}}}]$  in Eq. (1) is a somewhat “ideal” state, and the security criterion states that the real state produced at the end of the protocol is  $\epsilon^{\text{QKD}}$ -close to this “ideal” state. To avoid possible misconceptions, we briefly emphasize that this “ideal” state is *not* of the form  $\sum_{k \in \{0,1\}^{\ell}} 2^{-\ell} |kk\rangle\langle kk|_{K^A K^B} \otimes \rho_{E^{\text{final}}}^{(\ell)}$  for any specific length  $\ell$  (and some state  $\rho_{E^{\text{final}}}^{(\ell)}$ ), but rather it is a mixture of such states, with the weights in the mixture being generally *unknown* (as they depend on Eve’s arbitrary attack). In particular, this means that for fixed-length protocols, it would *not* be correct to say that the real state conditioned on accepting is  $\epsilon^{\text{QKD}}$ -close to a state of that form, since the security definition only involves a mixture of the states conditioned on accepting and aborting. We emphasize however that as discussed in Ref. [8], in the context of the standard QKD assumptions, this definition still suffices to yield composable security.

We now turn to the question of how to slightly modify this security definition in the context of device imperfections and certification.

### B. Security with respect to device models

In our above discussion of the security definition, in order to obtain a channel  $\mathcal{E}^{\mathcal{A}}$  from the attack  $\mathcal{A}$ , we in fact implicitly needed to rely on having some model for the QKD device used in the protocol. This works well enough under a “basic” QKD security framework where the device (or at least all its “relevant” properties) is assumed to be exactly characterized, but in the presence of device imperfections or inexact characterization, it is helpful to make this dependence on the device model more explicit. Hence we shall now slightly rephrase the security definition in a manner more suited for this context, which we shall use for this work.

Consider a QKD device that is about to undergo a certification procedure, intended to certify it for use in up to  $N_{\max}$  instances of some QKD protocol (after which it is discarded or sent for recertification) [11]. We shall suppose that the device behavior across all those  $N_{\max}$  instances is described by some **device model**, which we shall denote as  $\mathcal{M}$ . The device model  $\mathcal{M}$  is intended to fully capture the behavior of the device throughout those  $N_{\max}$  protocol instances, in the following sense. Suppose that at the end of each instance, we let  $K_j^A K_j^B$  denote the keys produced by that instance, and we let  $E_j$  denote *all other registers* that might possibly be accessible to Eve in any form in the future (there are some subtleties in interpreting this, which we defer to Remark 2 later). Similar to the earlier discussion, let us use  $\mathcal{A}$  to denote some arbitrary attack by Eve across all  $N_{\max}$  instances [12]. Once again, we do not impose restrictions on the exact “types of mathematical objects” used to formalize the device model  $\mathcal{M}$  and the attack  $\mathcal{A}$ . Our central requirement is only the following:

*Condition 1.* Given any model  $\mathcal{M}$  and any attack  $\mathcal{A}$ , there exists a corresponding sequence of channels  $\mathcal{E}_j^{\mathcal{A}, \mathcal{M}} : Q_{j-1} \rightarrow K_j^A K_j^B E_j$ , describing the physical transformations induced by the  $j$ th QKD instance from states on some “initial” register  $Q_{j-1}$  to states on  $K_j^A K_j^B E_j$ .

Under this framework, we introduce the following definition of what it means for a sequence of QKD protocol instances to be secure for a given device model  $\mathcal{M}$ :

*Definition 1.* Following the above description, consider a device for use in up to  $N_{\max}$  instances of a QKD protocol. For each  $j \in \{1, 2, \dots, N_{\max}\}$ , we say that the  $j$ th QKD protocol instance is  $\epsilon_j^{\text{QKD}}$ -**secure under device model  $\mathcal{M}$**  if for any attack  $\mathcal{A}$  and any initial states  $\sigma^{(j)}$ , we have

$$d\left(\mathcal{E}_j^{\mathcal{A}, \mathcal{M}}\left[\sigma_{Q_{j-1}}^{(j)}\right], \mathcal{K}_j \circ \mathcal{E}_j^{\mathcal{A}, \mathcal{M}}\left[\sigma_{Q_{j-1}}^{(j)}\right]\right) \leq \epsilon_j^{\text{QKD}}, \quad (3)$$

where  $\mathcal{K}_j$  is a channel that reads the length  $\ell$  of the keys on registers  $K_j^A K_j^B$  and replaces them with perfectly correlated

uniform keys of *that length*, i.e., the state

$$\omega_{K_j^A K_j^B}^\ell := \sum_{k \in \{0,1\}^\ell} 2^{-\ell} |kk\rangle\langle kk|_{K_j^A K_j^B}. \quad (4)$$

Observe that in a context where the device is used sequentially for multiple instances of a QKD protocol, we can write any state generated after  $n$  such instances in the following form (as long as certain technicalities hold regarding memory registers; see Remark 2):

$$\mathcal{E}_n^{\mathcal{A}, \mathcal{M}} \circ \mathcal{F}_{n-1} \circ \dots \circ \mathcal{E}_2^{\mathcal{A}, \mathcal{M}} \circ \mathcal{F}_1 \circ \mathcal{E}_1^{\mathcal{A}, \mathcal{M}} \left[ \sigma_{Q_0}^{(0)} \right], \quad (5)$$

for some initial state  $\sigma^{(0)}$  and some channels  $\mathcal{F}_j : K_j^A K_j^B E_j \rightarrow Q_j$  that represent arbitrary “physical processes” taking place between the completion of the  $j$ th protocol instance and the following instance. In the subsequent discussions, we will be discussing properties of states with the form in Eq. (5), with the understanding that they are meant to describe the state at the end of  $n$  QKD protocol instances.

We highlight that in this formalism we have chosen, we allow each of the  $\mathcal{F}_j$  channels to potentially act nontrivially on the keys  $K_j^A K_j^B$  produced by the preceding channel, not just the “side-information” register  $E_j$ . This is because these channels are intended to capture “all physical processes” that take place in each instance, and in many practical applications of QKD, one would presumably want to make use of the keys produced in each instance before the next instance has fully completed. This could potentially leak information about those keys to an adversary—as an example, if  $K_j^A K_j^B$  are used in a one-time-pad protocol, then the ciphertext can be nontrivially correlated to  $K_j^A K_j^B$  (this being essentially the reason one-time-pads cannot be used to encrypt more than one message). Another particularly critical example would be using part of the keys produced by one instance to authenticate the subsequent instance, as discussed in Ref. [13]. By including these keys in the domain of  $\mathcal{F}_j$ , our formalism is able to capture such scenarios as well; in particular, we do not require that the keys remain “unused” until the devices have reached the end of their certification lifespan (which would clearly be impractical).

*Remark 2.* Some care is needed in interpreting the registers  $E_j$  in the above definition. Specifically, we first emphasize it is known that if the behavior of the device in some QKD protocol instance depends nontrivially on some “memory” of the raw key retained from previous instances, then this fundamentally breaks the composable security of QKD [14]. While that work was focused on device-independent QKD, the same issue arises for any other form of QKD (including measurement-device-independent QKD) in the absence of assumptions ruling out such forms

of memory effects in the devices, whether they be for state preparation or measurement. Therefore, we should consider how this is to be reflected in the above formalism, denoting the memory retained in a QKD device after each protocol instance as  $R_j$ .

One option is to view  $R_j$  as being implicitly part of the  $E_j$  register, but in that case showing that the bound Eq. (3) holds would require proving that the memory is *also* basically “decoupled” from the secret key—with this, we would have no interpretational difficulties with the modeling and theoretical framework; however, proving that the memory is indeed decoupled from the key would involve invoking suitable physical assumptions. Alternatively, we could exclude  $R_j$  from the  $E_j$  register, but note that in the above formalism this is tantamount to saying that *no “allowed” physical process* gives Eve access to  $R_j$  in any form, in that none of the subsequent  $\mathcal{F}_j$  and  $\mathcal{E}_j^{A,\mathcal{M}}$  channels can act on  $R_j$  in any way—for practical purposes, this might be operationally equivalent to supposing the devices do not in fact retain memory between protocol instances.

While these options might appear to be somewhat restrictive conditions, we highlight that the existence of the [14] attack demonstrates that QKD would *genuinely* become insecure in the absence of assumptions ruling out “nontrivial” memory effects, and hence some assumptions of this nature are in fact necessary (unless countermeasures such as Refs. [15–17] are applied, but we leave for future work an analysis of how they interact with a certification framework).

### C. Main result

With this formalism, we now turn to the main question of how to analyze certification procedures in these terms. We shall suppose that given some device with model  $\mathcal{M}$ , the certification procedure either **approves** the device for use in up to  $N_{\max}$  QKD protocol instances, in which case it outputs a classical register  $F = 1$ , or it **rejects** the device, in which case it outputs  $F = 0$  (with the probabilities of each outcome being some function of  $\mathcal{M}$ ). In that case, the relevant state to consider before the start of the first QKD protocol instance can be written in the form

$$\sigma_{Q_0 F} = \Pr[F = 0] \sigma_{Q_0|F=0} \otimes |0\rangle\langle 0|_F + \Pr[F = 1] \sigma_{Q_0|F=1} \otimes |1\rangle\langle 1|_F, \quad (6)$$

where  $Q_0$  is the input register to the first channel  $\mathcal{E}_1^{A,\mathcal{M}}$  discussed in Condition 1, and  $\sigma_{Q_0|F=0}$  and  $\sigma_{Q_0|F=1}$ , respectively, denote the states on this register conditioned on  $F = 0$  or  $F = 1$ .

We highlight that in order to be able to draw any nontrivial conclusions in a certification framework, it is necessary to implicitly restrict the “universe of possible device models” in some sense. For instance, if we were to consider the hypothetical possibility of a device that behaves *exactly*

*identically in every way* to some “ideal” honest device during the certification procedure, only to later behave in some malicious fashion when used in the actual QKD protocols, then it would be impossible for a certification procedure to make nontrivial conclusions about this device (since by hypothesis, none of its malicious behavior can be “observed” during certification). Therefore, it is necessary to rule out at least some forms of device behavior purely *by assumption*, rather than relying on a certification procedure to do so. We shall formalize this in this work by supposing that all device models  $\mathcal{M}$  lie within some universal set (or class) denoted as  $\mathcal{U}_{\text{models}}$ , which is chosen to be broad enough to capture all “reasonable” device behaviors, but still sufficiently specific to allow proving nontrivial statements such as the criteria in Proposition 1 below.

*Remark 3.* As a simple example, we could for instance restrict  $\mathcal{U}_{\text{models}}$  to only consist of device models such that the state preparation device and the measurement device behave in an independent and identically distributed (IID) manner every time they are used, even if we do not exactly know the state prepared or measurement performed in a single usage. Note that this is a milder assumption than assuming a restriction to “IID collective attacks” (see, e.g., Ref. [8]) in QKD security proofs; we return to this point at the end of Sec. III, when discussing confidence intervals. Qualitatively, this rules out the above hypothetical of the device performing very differently during the certification procedure compared to the actual protocol, and allows us to obtain nontrivial conclusions from Proposition 1 later. However, we emphasize that this is just a starting example, and one could consider more general choices of  $\mathcal{U}_{\text{models}}$  if desired, as we elaborate on at the end of Sec. III.

Since a device model  $\mathcal{M}$  could in general be a very complicated object (as it should describe the entire behavior of the device over all  $N_{\max}$  protocol instances), we shall introduce an additional concept to simplify the subsequent discussion and the practical certification procedure. Namely, we shall suppose that each device model  $\mathcal{M}$  is associated to a list of  $N_{\text{param}}$  **parameters**  $(\mu_1, \mu_2, \dots, \mu_{N_{\text{param}}}) =: \boldsymbol{\mu}$ , which informally speaking summarizes some “main” properties relevant to a security analysis. As a basic example, one of these parameters could for instance be an upper bound [18] on the dark count rate of the detectors, where the bound is to hold over all  $N_{\max}$  protocol instances. Again, this is just a simple example, which is not meant to be definitive—for now we retain flexibility in what can be considered a parameter of the device model, returning to this question later after presenting a notion of what we require from this concept.

*Remark 4.* Note that the simplification in terms of parameters is not strictly mathematically necessary for our

subsequent results, which could in principle be rephrased directly in terms of the device models  $\mathcal{M}$  themselves. However, for practical purposes it is much easier to discuss the certification procedure in terms of a concise list of parameters  $\mu$ .

With this, we present our main technical statement regarding device certification. We highlight that it formally holds for any choice of the universe of models  $\mathcal{U}_{\text{models}}$ ; however, an “overly broad” choice of  $\mathcal{U}_{\text{models}}$  might make it impossible for the criteria in this proposition to be satisfied, as discussed above, making this statement vacuous in those circumstances. Also as discussed above, the channels  $\mathcal{F}_j$  in the statement should be understood to represent any arbitrary “physical processes” taking place between the QKD instances.

*Proposition 1.* Following the above description, consider a device for use in up to  $N_{\text{max}}$  instances of a QKD protocol, and let  $S_{\text{robust}}$  be some set of parameter values  $\mu$ . Suppose the following criteria are satisfied:

- (1) For every device model  $\mathcal{M} \in \mathcal{U}_{\text{models}}$  with parameters  $\mu \notin S_{\text{robust}}$ , the certification procedure rejects with probability at least  $1 - \epsilon^{\text{cert}}$ , i.e.,  $\Pr[F = 1] \leq \epsilon^{\text{cert}}$ .
- (2) For every device model  $\mathcal{M} \in \mathcal{U}_{\text{models}}$  with parameters  $\mu \in S_{\text{robust}}$ , the  $j$ th QKD protocol is  $\epsilon_j^{\text{QKD}}$ -secure under device model  $\mathcal{M}$  (Definition 1).

Then for any  $n \in \{1, 2, \dots, N_{\text{max}}\}$ , any state  $\sigma_{Q_0 F}$  [as in Eq. (6)] at the start of the first protocol instance, any attack  $\mathcal{A}$  and device model  $\mathcal{M} \in \mathcal{U}_{\text{models}}$ , and any sequence of channels  $\mathcal{F}_j : K_j^A K_j^B E_j \rightarrow Q_j$  for  $j \in \{1, 2, \dots, n-1\}$ , we have

$$\begin{aligned} & \Pr[F = 1] d\left(\bar{\mathcal{E}}_n^{A, \mathcal{M}} \circ \dots \bar{\mathcal{E}}_1^{A, \mathcal{M}} [\sigma_{Q_0|F=1}], \right. \\ & \quad \left. \mathcal{K}_n \circ \bar{\mathcal{E}}_n^{A, \mathcal{M}} \circ \dots \mathcal{K}_1 \circ \bar{\mathcal{E}}_1^{A, \mathcal{M}} [\sigma_{Q_0|F=1}]\right) \leq \epsilon_n^{\text{total}}, \\ & \text{with } \epsilon_n^{\text{total}} = \epsilon^{\text{cert}} + \sum_{j=1}^n \epsilon_j^{\text{QKD}}, \end{aligned} \quad (7)$$

where for notational compactness we write  $\bar{\mathcal{E}}_j^{A, \mathcal{M}} := \mathcal{E}_j^{A, \mathcal{M}} \circ \mathcal{F}_{j-1}$  (defining  $\mathcal{F}_0$  to be the identity channel on  $Q_0$ ), and  $\mathcal{K}_j$  is a channel that reads the length  $\ell$  of the keys on registers  $K_j^A K_j^B$  and replaces them with perfectly correlated uniform keys of that length, i.e., the state  $\omega^\ell$  in Eq. (4).

The proof is fairly straightforward, hence we defer it to Appendix A [19]. For now we simply discuss its qualitative implications. In particular, observe that it allows us to more precisely specify some goals we should aim for

in security proofs and a certification procedure, via the two criteria. Specifically, the second criterion tells us we should aim to construct security proofs that have some “robustness” to device imperfections, in the sense that the QKD protocols are secure for all devices with model parameters in some choice of set  $S_{\text{robust}}$ . From the other direction, the first criterion tells us the certification procedure should aim to reject all devices with model parameters outside  $S_{\text{robust}}$  with high probability. This is the main proposal of our work, and in Secs. III and IV, we explore in greater detail how these goals could be achieved in practice.

Before proceeding, we briefly discuss how to interpret the trace-distance bound obtained from the above result. In particular, observe that the state  $\bar{\mathcal{E}}_n^{A, \mathcal{M}} \circ \dots \bar{\mathcal{E}}_1^{A, \mathcal{M}} [\sigma_{Q_0|F=1}]$  in the above formula is basically the real state produced after  $n$  protocol instances for a given attack and device model (conditioned on the certification procedure approving), while the state  $\mathcal{K}_n \circ \bar{\mathcal{E}}_n^{A, \mathcal{M}} \circ \dots \mathcal{K}_1 \circ \bar{\mathcal{E}}_1^{A, \mathcal{M}} [\sigma_{Q_0|F=1}]$  is essentially a more “ideal” version in which the keys produced at the end of each protocol instance have been immediately replaced by perfect keys of that length. Hence our above result is an upper bound on the trace distance between these states, weighted by the prefactor of  $\Pr[F = 1]$ . Note that the possibility of Eve performing different actions depending on the certification outcome can be implicitly incorporated into the framework we describe here, simply by defining the channels  $\mathcal{E}_n^{A, \mathcal{M}}$  to act on a classical register storing that outcome as well.

However, we emphasize (similar to Remark 1) that it would *not* be correct to view the above bound as saying that the trace distance between the states conditioned on the certification procedure approving (i.e.,  $\bar{\mathcal{E}}_n^{A, \mathcal{M}} \circ \dots \bar{\mathcal{E}}_1^{A, \mathcal{M}} [\sigma_{Q_0|F=1}]$  and  $\mathcal{K}_n \circ \bar{\mathcal{E}}_n^{A, \mathcal{M}} \circ \dots \mathcal{K}_1 \circ \bar{\mathcal{E}}_1^{A, \mathcal{M}} [\sigma_{Q_0|F=1}]$ ) is small. (An analogous point also applies for conditioning on some or all of the QKD protocol instances accepting, as already discussed in Remark 1.) This is due to the prefactor of  $\Pr[F = 1]$  in Eq. (7), which allows the trace distance to potentially be large if  $\Pr[F = 1]$  is small. Still, we can recover some intuitive properties from this bound as follows. Given a state  $\sigma_{Q_0 F}$  as in Eq. (6) at the start of the first protocol instance, we can write the final state at the end of all the instances in the form

$$\begin{aligned} \rho_{K_n^A K_n^B E_n F} &= \Pr[F = 0] \rho_{K_n^A K_n^B E_n|F=0} \otimes |0\rangle\langle 0|_F \\ &+ \Pr[F = 1] \bar{\mathcal{E}}_n^{A, \mathcal{M}} \circ \dots \bar{\mathcal{E}}_1^{A, \mathcal{M}} [\sigma_{Q_0|F=1}] \\ &\otimes |1\rangle\langle 1|_F, \end{aligned} \quad (8)$$

i.e., a mixture of the cases conditioned on the certification procedure approving or rejecting. If we now define another

state  $\tilde{\rho}$  via

$$\begin{aligned} \tilde{\rho}_{K_n^A K_n^B E_n F} &= \Pr[F = 0] \rho_{K_n^A K_n^B E_n | F=0} \otimes |0\rangle\langle 0|_F \\ &\quad + \Pr[F = 1] \mathcal{K}_n \circ \bar{\mathcal{E}}_n^{A, \mathcal{M}} \circ \dots \mathcal{K}_1 \circ \bar{\mathcal{E}}_1^{A, \mathcal{M}} \\ &\quad [\sigma_{Q_0 | F=1}] \otimes |1\rangle\langle 1|_F, \end{aligned} \quad (9)$$

then the bound Eq. (7) is equivalent to

$$d(\rho_{K_n^A K_n^B E_n F}, \tilde{\rho}_{K_n^A K_n^B E_n F}) \leq \epsilon_n^{\text{total}}, \quad (10)$$

i.e., these states are genuinely  $\epsilon_n^{\text{total}}$ -close in trace distance.

We stress that the states  $\rho_{K_n^A K_n^B E_n F}$  and  $\tilde{\rho}_{K_n^A K_n^B E_n F}$  are *mixtures* of the states conditioned on the certification procedure approving or rejecting, not just the states conditioned on approving. However, in many contexts  $\rho_{K_n^A K_n^B E_n F}$  is indeed the final state of interest, since the actual state that would be physically produced is indeed this mixture. As for  $\tilde{\rho}_{K_n^A K_n^B E_n F}$ , while it is not as “ideal” a state as one might initially have hoped for, it should still possess many useful properties if we presume that the component  $\rho_{K_n^A K_n^B E_n | F=0}$  conditioned on the certification procedure rejecting is essentially “trivial” in some sense (for instance, the device is never used for any QKD protocols and all the registers  $K_j^A K_j^B$  are just set to some trivial “null” values). In particular, for readers already familiar with the Abstract Cryptography (AC) framework [8,20], we discuss in Appendix C how this analysis suffices to show we have a composable secure construction of a particular resource (appropriately modified from the usual resource constructed by QKD), under certain assumptions.

For now, we simply note that due to the fundamental properties of trace distance, the trace-distance bound implies for instance that if we were to replace the actual state  $\rho_{K_n^A K_n^B E_n F}$  with  $\tilde{\rho}_{K_n^A K_n^B E_n F}$  in any other security analysis, then the probability of any event cannot change by more than  $\epsilon_n^{\text{total}}$ , which would already be a useful operational result for many security statements (as long as one keeps in mind that  $\tilde{\rho}_{K_n^A K_n^B E_n F}$  has the structure in Eq. (9) for some “unknown” probabilities  $\Pr[F = 0]$ ,  $\Pr[F = 1]$ ).

One might question whether the  $\Pr[F = 1]$  prefactor in the Proposition 1 bound is truly necessary. We highlight however that it is straightforward to construct counterexamples (even with simple IID devices) such that the final state conditioned on the certification procedure approving satisfies

$$\begin{aligned} d(\bar{\mathcal{E}}_n^{A, \mathcal{M}} \circ \dots \bar{\mathcal{E}}_1^{A, \mathcal{M}} [\sigma_{Q_0 | F=1}], \mathcal{K}_n \circ \bar{\mathcal{E}}_n^{A, \mathcal{M}} \\ \circ \dots \mathcal{K}_1 \circ \bar{\mathcal{E}}_1^{A, \mathcal{M}} [\sigma_{Q_0 | F=1}]) \approx 1, \end{aligned} \quad (11)$$

i.e., it is very far from the “ideal” state. (This does not contradict Proposition 1, precisely because the bound in

it has the  $\Pr[F = 1]$  prefactor.) Specifically, we simply need to again consider a “genuinely insecure” device as described in the introduction, such as the example of a fully phase-coherent laser being used for QKD protocols in a high-loss regime. As discussed in the introduction, Eve would always be able to perfectly learn the QKD keys generated using such a device, even conditioned on the certification approving—hence the trace distance between the states conditioned on approving, as in Eq. (11), would indeed be close to 1. Therefore, any attempt to bound the trace distance involving the state conditioned on the certification procedure accepting would necessarily need to account for  $\Pr[F = 1]$  in some form, even if it might not be strictly in the form presented in Eq. (7).

### III. CONSIDERATIONS FOR THE CERTIFICATION PROCEDURE

In this section, we outline some examples of how a certification procedure could be designed to reject “bad” devices, in the sense of those with model parameters outside of some set  $S_{\text{robust}}$  for which we can prove the QKD protocols are secure (see Sec. IV). We again highlight that the following examples are not necessarily definitive, in the sense that Proposition 1 is not concerned with the details of exactly how this is achieved—it only requires that such devices are rejected with high probability. Still, we believe that the scenarios considered here are reasonably reflective of procedures that could be applied in practice; at the end of this section, we also discuss some natural extensions.

To begin, we briefly present the basic definition of a confidence interval:

*Definition 2.* Consider a probability distribution described by some fixed value  $\mu$ , usually referred to as a **parameter**. If we have two random variables  $X^{\text{low}}$  and  $X^{\text{upp}}$  such that the interval  $\mathcal{I}^{\text{conf}} := [X^{\text{low}}, X^{\text{upp}}]$  contains  $\mu$  with probability at least  $1 - \epsilon$ , then we say  $\mathcal{I}^{\text{conf}}$  is a **confidence interval** for  $\mu$  at confidence level  $\epsilon$ .

We emphasize a standard conceptual point regarding confidence intervals: in the above definition,  $\mu$  is to be understood as a *fixed* value, with no “randomness” involved. Therefore, when one says that  $\mathcal{I}^{\text{conf}}$  contains  $\mu$  with some probability, the “random quantity” is *not* the value  $\mu$  but rather the interval  $\mathcal{I}^{\text{conf}}$ , and it would be a misconception to speak of (for instance) the probability distribution of  $\mu$ . As a default example, suppose we have  $n$  random variables  $X_1, X_2, \dots, X_n$  following an IID distribution  $P_{X_1 X_2 \dots X_n} = P_{X_1} P_{X_2} \dots P_{X_n}$ , where the expected value for any individual  $X_j$  is some value  $\mu$ . A common method to construct a confidence interval for  $\mu$  is to compute the “sample mean”  $X_{\text{avg}} := (1/n) \sum_j X_j$ , and then

set  $\mathcal{I}^{\text{conf}} = [X_{\text{avg}} - \delta, X_{\text{avg}} + \delta]$  for some suitably chosen value  $\delta$  (depending on  $n$ ,  $\epsilon$  and some other parameters, such as the range of the random variables). Notice that in this process, the value of  $\mu$  is a *fixed* quantity, with the “random quantities” being  $\{X_j\}$ ,  $X_{\text{avg}}$ , and thus  $\mathcal{I}^{\text{conf}}$ .

Critically, however, the standard valid interpretation of confidence intervals is fully compatible with the framework suggested by Proposition 1, as we now describe. As a crude starting demonstration, consider a very simplistic scenario where  $\mathcal{U}_{\text{models}}$  consists only of devices that behave in an IID manner every time they are used, and the only parameter of importance is, say, the dark count rate  $\mu_{\text{dark}}$  in the photon detection device (which is constant across all uses, by the IID restriction). Suppose furthermore that the set  $S_{\text{robust}}$  is simply some (fixed) interval  $[\mu_{\text{dark}}^{\text{low}}, \mu_{\text{dark}}^{\text{upp}}]$ , i.e., we have established security proofs such that QKD protocols performed with this device are secure whenever the device has dark count rate  $\mu_{\text{dark}}$  within the interval  $[\mu_{\text{dark}}^{\text{low}}, \mu_{\text{dark}}^{\text{upp}}]$ . In that case, if we design the certification procedure such that it constructs a confidence interval  $\mathcal{I}_{\text{dark}}^{\text{conf}}$  for  $\mu_{\text{dark}}$  at confidence level  $\epsilon^{\text{cert}}$ , and set the approval condition to

$$\text{Approve if and only if } \mathcal{I}_{\text{dark}}^{\text{conf}} \subseteq [\mu_{\text{dark}}^{\text{low}}, \mu_{\text{dark}}^{\text{upp}}], \quad (12)$$

then it follows directly from the definition of confidence intervals that indeed, this certification procedure fulfills the criterion in Proposition 1 that it rejects with probability at least  $1 - \epsilon^{\text{cert}}$  whenever  $\mu_{\text{dark}} \notin [\mu_{\text{dark}}^{\text{low}}, \mu_{\text{dark}}^{\text{upp}}]$ . Specifically, simply recall that from Definition 2 we have that  $\mathcal{I}_{\text{dark}}^{\text{conf}}$  contains  $\mu_{\text{dark}}$  with probability at least  $1 - \epsilon^{\text{cert}}$ ; therefore, whenever  $\mu_{\text{dark}} \notin [\mu_{\text{dark}}^{\text{low}}, \mu_{\text{dark}}^{\text{upp}}]$ , with that probability  $\mathcal{I}_{\text{dark}}^{\text{conf}}$  contains the value  $\mu_{\text{dark}}$  outside  $[\mu_{\text{dark}}^{\text{low}}, \mu_{\text{dark}}^{\text{upp}}]$ , causing the certification to reject. Also notice that as one would intuitively expect, it is advantageous to construct  $\mathcal{I}_{\text{dark}}^{\text{conf}}$  as small as possible, since it makes the certification procedure more likely to approve “good” devices with  $\mu_{\text{dark}} \in [\mu_{\text{dark}}^{\text{low}}, \mu_{\text{dark}}^{\text{upp}}]$ , while still retaining the desired security properties (as long as  $\mathcal{I}_{\text{dark}}^{\text{conf}}$  still satisfies Definition 2). Finally, we acknowledge that a natural question is whether one could relax the approval condition to  $\mathcal{I}_{\text{dark}}^{\text{conf}}$  having nonempty overlap with  $[\mu_{\text{dark}}^{\text{low}}, \mu_{\text{dark}}^{\text{upp}}]$ ; however, this is not strong enough to show the desired property.

More realistically, we would want to consider multiple parameters  $(\mu_1, \mu_2, \dots, \mu_{N_{\text{param}}}) =: \boldsymbol{\mu}$  as described in the previous section. Suppose for simplicity that  $S_{\text{robust}}$  is simply defined by “componentwise” intervals, i.e., it has the following form (for some fixed values  $\mu_k^{\text{low}}, \mu_k^{\text{upp}}$ ):

$$S_{\text{robust}} = \{\boldsymbol{\mu} \in \mathbb{R}^{N_{\text{param}}} \mid \mu_k \in [\mu_k^{\text{low}}, \mu_k^{\text{upp}}] \text{ for all } j\}. \quad (13)$$

In that case, a certification procedure satisfying the criterion in Proposition 1 would be to construct for each  $k$  a confidence interval  $\mathcal{I}_k^{\text{conf}}$  for  $\mu_k$  at confidence level  $\epsilon^{\text{cert}}$ ,

and set the approval condition to

$$\text{Approve if and only if } \mathcal{I}_k^{\text{conf}} \subseteq [\mu_k^{\text{low}}, \mu_k^{\text{upp}}] \text{ for all } k. \quad (14)$$

Note that we do not require the confidence interval constructions to be independent in any way; we only require that each confidence interval individually satisfies Definition 2. Again, to see that this has the desired property, simply observe that if  $\boldsymbol{\mu} \notin S_{\text{robust}}$ , that means at least one  $\mu_k$  lies outside the corresponding interval  $[\mu_k^{\text{low}}, \mu_k^{\text{upp}}]$ . Therefore, with probability at least  $1 - \epsilon^{\text{cert}}$  the corresponding confidence interval  $\mathcal{I}_k^{\text{conf}}$  contains this value outside  $[\mu_k^{\text{low}}, \mu_k^{\text{upp}}]$ , forcing the certification to reject (without any independence assumptions across the confidence intervals). Again, one might ask whether the approval condition could be relaxed to having  $\mathcal{I}_k^{\text{conf}} \subseteq [\mu_k^{\text{low}}, \mu_k^{\text{upp}}]$  for only some  $k$  rather than all, or to only requiring nonempty overlap between  $\mathcal{I}_k^{\text{conf}}$  and  $[\mu_k^{\text{low}}, \mu_k^{\text{upp}}]$ —however, these versions are again not strong enough to show the desired property.

We believe that the above example with multiple parameters should conceptually suffice to match certification procedures that can be performed in practice (though it does clarify precisely what should be set as the approval condition). In particular, we note that the above structure is completely indifferent to the exact method used to construct the confidence intervals, as long as they genuinely satisfy the foundational definition (Definition 2), and hence is very flexible in that regard. One might even wish to consider the more general concept of **confidence regions**, which are subsets of  $\mathbb{R}^{N_{\text{param}}}$  that are not necessarily defined by “componentwise” intervals (see Definition B1 for more details); similarly, one could consider choices of  $S_{\text{robust}}$  that are not of the form Eq. (13). Our above analysis generalizes straightforwardly by requiring the confidence region to be completely contained in  $S_{\text{robust}}$ .

We close this section with some discussions of points to note in practice.

**Assumptions in constructing confidence intervals.** It is worth highlighting that in order to construct a confidence interval in the first place, one usually needs to make some sort of assumption on the device behavior during the certification procedure, such as an IID or independence property, or perhaps a martingale property (for instance, to apply Azuma’s inequality). However, it seems reasonable to suggest that since this is a consideration of the *devices* manufactured for use by the honest parties, one could take a less “adversarial” perspective on their behavior as compared to a QKD security proof analyzing Eve’s *attack* on the quantum states being transmitted, and thus reasonably restrict  $\mathcal{U}_{\text{models}}$  such that the devices have the required independence properties. (In other words, this is weaker than an assumption that Eve is restricted to “IID collective attacks” [8].) Furthermore, in any case it seems

unlikely that one would be able to draw nontrivial conclusions from certification without at least some assumptions of this nature.

**Nonconstant behavior.** While the above discussion is very general regarding what can be considered a “parameter,” it may be worth discussing some specific scenarios more closely. For instance, it is natural to consider the possibility of a QKD device where its performance “degrades” over time in some fashion. One might wonder whether this would require us to include an entire *continuum* of parameters in our formalism, describing its behavior at every point in time. However, this would be a rather impractical way of capturing this within our formalism. Rather, we believe a more reasonable approach would be to assume the universe of models  $\mathcal{U}_{\text{models}}$  is restricted such that for any device model  $\mathcal{M}$  within this universe, this “degrading behavior” over its entire lifespan can be described in terms of some *single* parameter or a small number of parameters.

For example, single-photon detectors used in satellite QKD are expected to have time-varying dark count rates due to radiation damage and changes in temperature (see, e.g., Ref. [21]). This time-variation can be studied and considered to be a part of the universe of models  $\mathcal{U}_{\text{models}}$ —for instance, we could capture this by restricting  $\mathcal{U}_{\text{models}}$  to be such that for any model  $\mathcal{M} \in \mathcal{U}_{\text{models}}$ , the dark count rate in the  $j$ th QKD instance can be lower and upper bounded by some functions  $f^{\text{low}}(j, \mu_{\text{temp}})$  and  $f^{\text{upp}}(j, \mu_{\text{temp}})$ , respectively, where  $\mu_{\text{temp}}$  is the temperature around the detector, which one can aim to construct a confidence interval for during certification. This is mainly a basic example to capture the spirit of the idea; practically, one should give a more sophisticated description of  $\mathcal{U}_{\text{models}}$  and how the models within it can be described in terms of some parameters. Note that there is freedom to “trade-off” against the security proofs here, in the sense that if the device behavior is expected to degrade over time as described by the model  $\mathcal{M} \in \mathcal{U}_{\text{models}}$ , then it may be possible to compensate for this in the security proofs for the QKD protocols (i.e., criterion 1 in Proposition 1) by reducing the key lengths in later QKD instances, to ensure that the  $j$ th instance is indeed  $\epsilon_j^{\text{QKD}}$ -secure for models with  $\mu \in S_{\text{robust}}$ .

Another consideration is whether some aspects of the device behavior might deviate dramatically from the honest behavior when subject to some form of adversarial attack, such as a Trojan-horse attack [22]. In that case, the relevant parameter to assign to it within our formalism would be for instance a *bound* on how far it can deviate from the honest behavior even when subjected to such attacks; for instance, by using an optical isolator [22], one might be able to bound the maximum extent of deviation (under “reasonable” assumptions [23] about the intensity of light the adversary can send, and so on, which would be captured within our framework by restricting the universe  $\mathcal{U}_{\text{models}}$ ). The security proofs would then have to ensure

that as long as this parameter lies within  $S_{\text{robust}}$ , the QKD protocol instances are still secure.

Overall, this highlights the need to “embed” assumptions into  $\mathcal{U}_{\text{models}}$  somewhat carefully, so that we can still describe the models with a tractable number of parameters for certification, while still being able to prove security for models with  $\mu \in S_{\text{robust}}$ . This is something that is pragmatically necessary, as the space of “all conceivable device models” would be unreasonably large, and impossible to design a practical certification procedure for. Practically, this would mean choosing  $\mathcal{U}_{\text{models}}$  according to the accepted best practices at a given time, such that they indeed capture all device behaviors that are “realistically” known to occur (even in the presence of adversarial attacks during the QKD protocol), while still being tractable to analyze. It is also important to ensure that the security proofs for the QKD protocols remain “up-to-date” with the choice of  $\mathcal{U}_{\text{models}}$ , so that they indeed ensure the required security properties for models with  $\mu \in S_{\text{robust}}$ . In the next section, we turn to discussing this point.

**Practical challenges in metrology.** A significant practical challenge remains in the *metrological* quality of the certification procedure. In a realistic deployment, the certification is performed using reference devices (e.g., single-photon detectors, laser sources, etc.) that are themselves subject to uncertainties and imperfections. Further, in order to accurately capture errors such as misalignment, a realistic certification procedure would involve a complete certification of the entire detector or source module rather than the individual parameters that comprise of it. Rigorously developing a certification procedure that accounts for these imperfections is a challenging task that we leave to future work. Our framework highlights that the security of a deployed system relies as much on the quality of these certification procedures as it does on the security proof techniques.

#### IV. CONSIDERATIONS FOR THE SECURITY PROOFS

An important point arising from the preceding discussions is that in order for this framework to be practical, we do genuinely need QKD security proofs that have “non-trivial robustness” to imperfections, in the sense that we can show criterion 2 in Proposition 1 holds for some choice of  $S_{\text{robust}}$  that does not consist only of a single point. This is because from the preceding discussion, as long as the certification procedure can only construct confidence intervals of nonzero width, it would never approve the device if  $S_{\text{robust}}$  consists only of a single point (or is otherwise “too small” for reasonably sized confidence intervals to be contained within it). Hence it highlights the importance of constructing QKD security proofs in such a way that they have some robustness to device imperfections, as is a subject of much ongoing work [24–32].

Furthermore, given the preceding discussion of the flexibility in choosing parameters, we see that it is important for security proofs to remain “in touch” with the practical reality of what can be measured during a certification procedure. One should choose the parameters and the universe  $\mathcal{U}_{\text{models}}$  in such a way that there is compatibility between the certification procedures and the security proofs, while still “realistically” reflecting the actual behavior of the devices. It is likely that this will need to involve an ongoing process of updating the models and choices of parameters, as better understanding is gained of device characteristics and potential flaws.

## V. PEDAGOGICAL EXAMPLE: DETECTOR DARK COUNTS

To demonstrate the application of our framework more concretely, we consider the characterization of a receiver’s device for the BB84 protocol. We focus specifically on the characterization of the detectors’ dark count rates, illustrating the application of our work.

### A. The physical model

We consider a standard passive basis-choice detection setup consisting of a basis-choice beam splitter leading to two measurement arms (one for the Z basis, one for the X basis), utilizing a total of four single-photon detectors. We assume that the detection setup couples perfectly to a single spatiotemporal input mode, and all detectors have unit efficiencies. Further, we assume that the repetition rate of the protocol is fixed. This repetition rate allows us to define the probability of a dark count of each detector in a single protocol round. We *do not* assume that this dark count probability is perfectly known. Finally, we assume that the detector behavior is IID in each round, with no memory effects (again, we emphasize that this is *not* restricting an attacker to collective attacks: it only involves the device held by the receiver).

In short, we assume all devices behave perfectly as per specification, except for the dark count probabilities which must be characterized. This defines our universe of models  $\mathcal{U}_{\text{models}}$ . We associate with each device model  $\mathcal{M}$  a single parameter  $\mu_{\text{dark}}$ , which represents the maximum dark count probability of all detectors. While we could define the dark count probability in each detector to be an independent parameter in our device model, we find this simpler, single parameter description is sufficient for security proofs [29,33].

### B. The certification procedure

Let us define the set  $S_{\text{robust}}$  as the subset of all parameter values where all four detectors have a dark count probability at most  $\mu_{\text{dark}}^{\text{upp}}$  (for some value  $\mu_{\text{dark}}^{\text{upp}}$  that is fixed before

the certification procedure takes place):

$$S_{\text{robust}} = \{ \mu_{\text{dark}} \in [0, 1] \mid \mu_{\text{dark}} \leq \mu_{\text{dark}}^{\text{upp}} \}. \quad (15)$$

*Remark 5.* We emphasize a basic but crucial point: the value of  $\mu_{\text{dark}}^{\text{upp}}$  in this example (and thus  $S_{\text{robust}}$ ) is to be fixed *before* the certification procedure begins, and not adjusted afterward. The choice of this value does not affect the security guarantees (as long as the subsequent QKD protocols are implemented with the key lengths prescribed in Sec. VC later); however, in order for the certification procedure to be “practically useful,” we need it to also accept with reasonably high probability whenever the device behaves according to the expected “correct” behavior. In particular, this means that we would usually need to choose  $\mu_{\text{dark}}^{\text{upp}}$  to be somewhat higher than the expected “correct” value—we stress however that the corresponding “margin of tolerance” is a *distinct* concept from the value  $\delta$  we describe below (which cannot be arbitrarily chosen without affecting the security guarantees), and should not be confused with or merged into it.

The certification procedure needs to be designed such that if  $\mu_{\text{dark}} \notin S_{\text{robust}}$ , the certification procedure rejects with probability at least  $1 - \epsilon^{\text{cert}}$  as described in criterion 1 of Proposition 1. Such a procedure can be implemented by formally integrating finite-size statistical analysis into existing methods [34] for characterizing dark count probabilities of individual detectors. We sketch out the main steps in such a procedure:

- (1) **Data collection:** Each detector is individually isolated from any input light. The detectors are each operated for  $n_{\text{char}}$  time bins. We record the resulting vector of click counts  $\mathbf{k} = (k_1, k_2, k_3, k_4)$  for each of the four detectors, and define a corresponding vector of click frequencies

$$\mathbf{f} = (f_1, f_2, f_3, f_4) := \left( \frac{k_1}{n_{\text{char}}}, \frac{k_2}{n_{\text{char}}}, \frac{k_3}{n_{\text{char}}}, \frac{k_4}{n_{\text{char}}} \right). \quad (16)$$

More generally, one could use a different number of time bins for each detector, but it is not hard to modify the analysis accordingly if that is the case.

- (2) **Statistical bounds:** We can treat the click count  $k_j$  of each detector as a realization of a binomial random variable—note this is because we have restricted the detector behavior to be IID within our universe of models. For such variables, one can apply a standard IID concentration inequality (e.g., the Hoeffding bound [35, Theorem 2], or other tail bounds on the binomial distribution) to calculate a finite-size penalty term  $\delta$ , with the following property: the probability that the random variable  $f_j + \delta$

lies below the true dark count rate of that detector is at most  $\epsilon^{\text{cert}}$ . Note that  $\delta$  will depend on the number of time-bins  $n_{\text{char}}$  and the desired  $\epsilon^{\text{cert}}$  value. With this, we straightforwardly have the fact that the probability that the random variable

$$\max_j \{f_j + \delta\} \quad (17)$$

lies below the true maximum dark count rate over all detectors is also at most  $\epsilon^{\text{cert}}$  [36]. (In the language of confidence intervals,  $[0, \max_j \{f_j + \delta\}]$  is a confidence interval for the maximum dark count rate over all detectors, where we have set the left endpoint to the “trivial” extremal value of 0 for a dark count rate.) Using this, we can explicitly set the approval condition: the certification procedure approves the device if and only if the observed frequencies satisfy

$$\max_j \{f_j + \delta\} < \mu_{\text{dark}}^{\text{upp}}. \quad (18)$$

From the preceding discussion, we have that if the true dark count rate of any detector is larger than  $\mu_{\text{dark}}^{\text{upp}}$ , then the probability of the frequencies  $f_j$  fluctuating low enough to satisfy Eq. (18) is at most  $\epsilon^{\text{cert}}$ .

*Remark 6.* The certification procedure presented here is straightforward only because of the large number of simplifying assumptions made in the physical model. More realistic device models with afterpulsing, nonunit efficiencies, dead times, etc., would require a more involved certification procedure, which we leave for future work.

In the above example, we framed the approval condition in terms of a single random variable  $\max_j \{f_j + \delta\}$ , rather than four separate conditions that each of the random variables  $f_j + \delta$  is at most  $\mu_{\text{dark}}^{\text{upp}}$ . While the latter is logically equivalent, we chose to present the former for several reasons. First, this directly matches the framing for single-parameter confidence intervals in Sec. III; second, it demonstrates constructing confidence intervals for parameters that do not literally correspond exactly to the expected value of some IID process (in this case, the *maximum* dark count rate over the four detectors). Such a perspective may be useful in other contexts, for instance when considering nonconstant behavior as mentioned in Sec. III.

### C. Robust QKD security proof

For the QKD security proof, we can appeal to existing works [29,33] that can prove security with an upper bound

on the dark count rate. Thus, for every device model  $\mathcal{M} \in \mathcal{U}_{\text{models}}$  with parameter  $\mu_{\text{dark}} < \mu_{\text{dark}}^{\text{upp}}$ , the  $j$ th QKD protocol can be proved to be  $\epsilon_j^{\text{QKD}}$ -secure under device model  $\mathcal{M}$  (Definition 1) by a direct application of Refs. [29,33] (for a suitable choice of the final key length, which depends on the value  $\mu_{\text{dark}}^{\text{upp}}$ ). This is exactly criterion 2 defined in Proposition 1.

### D. Application of the framework

This procedure directly satisfies the requirements of Proposition 1. The approval condition in the previous subsection is constructed specifically such that if the maximum dark count rate over all detectors  $\mu_{\text{dark}}$  were greater than  $\mu_{\text{dark}}^{\text{upp}}$  (i.e., the parameter was outside  $S_{\text{robust}}$ ), then the probability of the certification procedure approving the device is upper bounded by  $\epsilon^{\text{cert}}$ . This corresponds exactly to criterion 1 defined in Proposition 1. Similarly, as described in Sec. VC, a device model  $\mathcal{M} \in \mathcal{U}_{\text{models}}$  with parameter  $\mu_{\text{dark}} \in S_{\text{robust}}$  can be shown to satisfy criterion 2 defined in Proposition 1.

As a result, we can conclude from Proposition 1 that the joint process involving characterization followed by  $n$  QKD protocols (with the keys potentially being used between each QKD protocol) is “ $\epsilon_n^{\text{total}}$ -secure” in the sense formalized in Eq. (7).

## VI. CONCLUSION

Overall, in this work we have formalized a framework that allows us to draw rigorous conclusions regarding QKD device certification, as long as the certification procedure and QKD security proofs satisfy some criteria, via Proposition 1. In particular, this establishes those criteria as a useful point to focus on when considering those aspects, providing guiding principles for future considerations of device certification procedures.

Of course, a valid consideration is the question of whether there can be potential alternatives to those criteria. One natural possibility is to observe that it may be too restrictive to focus only on certification procedures that make a “binary” decision of whether to approve or reject the devices. More flexibly, we should consider the possibility that after the certification procedure constructs confidence intervals for the parameters  $\mu$ , these confidence intervals are used to make “adaptive” choices of the QKD protocols performed using the device, for instance by having the protocols output keys of shorter length if the confidence intervals suggest that the device imperfections are “larger.” This could perhaps be more accurately referred to as a “characterization” procedure, rather than “certification.” (This distinction is similar in spirit to the difference between fixed-length and variable-length QKD protocols [8].) Using the ideas presented in this work, we

can indeed obtain one possible rigorous formalization of this idea, which we present in Appendix B; we defer a more detailed analysis to future work.

Finally, we reiterate that while Proposition 1 is presented in terms of a trace-distance bound, we believe it is possible to connect this to the framework of Abstract Cryptography [8,20] under some technical assumptions, hence yielding a composable notion of security. While it is often said that the standard QKD trace-distance security criterion Eq. (1) “implies composable security,” this property is not an immediate consequence of an arbitrary trace-distance security definition, but rather it must be proven within a suitable formalization of the notion of “composable security.” Hence when we additionally introduce a certification procedure as in this work, it is necessary to re-evaluate whether an analogous claim is valid. We outline some relevant points regarding this in Appendix C, though a fully formal claim would require addressing certain technicalities regarding how to model QKD devices as “resources” within the Abstract Cryptography framework.

### ACKNOWLEDGMENTS

We thank Jan Krause, Norbert Lütkenhaus and Jerome Wiesemann for helpful discussions. This work was performed at the Institute for Quantum Computing, at the University of Waterloo, which is supported by Innovation, Science, and Economic Development Canada. This work was supported by NSERC under the Discovery Grants Program, Grant No. 341495. This research was supported by Alliance QUINT.

### DATA AVAILABILITY

There are no publicly available research data or software supporting this manuscript. Requests for further information or data should be sent to the authors.

### APPENDIX A: PROOF OF MAIN PROPOSITION

For clarity, we first restate the claim of the main proposition:

*Proposition A1.* Following the above description, consider a device for use in up to  $N_{\max}$  instances of a QKD protocol, and let  $S_{\text{robust}}$  be some set of parameter values  $\mu$ . Suppose the following criteria are satisfied:

- (1) For every device model  $\mathcal{M} \in \mathcal{U}_{\text{models}}$  with parameters  $\mu \notin S_{\text{robust}}$ , the certification procedure rejects with probability at least  $1 - \epsilon^{\text{cert}}$ , i.e.,  $\Pr[F = 1] \leq \epsilon^{\text{cert}}$ .
- (2) For every device model  $\mathcal{M} \in \mathcal{U}_{\text{models}}$  with parameters  $\mu \in S_{\text{robust}}$ , the  $j$ th QKD protocol is  $\epsilon_j^{\text{QKD}}$ -secure under device model  $\mathcal{M}$  (Definition 1).

Then for any  $n \in \{1, 2, \dots, N_{\max}\}$ , any state  $\sigma_{Q_0 F}$  [as in Eq. (6)] at the start of the first protocol instance, any attack

$\mathcal{A}$  and device model  $\mathcal{M} \in \mathcal{U}_{\text{models}}$ , and any sequence of channels  $\mathcal{F}_j : K_j^A K_j^B E_j \rightarrow Q_j$  for  $j \in \{1, 2, \dots, n-1\}$ , we have

$$\begin{aligned} \Pr[F = 1] d\left(\bar{\mathcal{E}}_n^{A, \mathcal{M}} \circ \dots \bar{\mathcal{E}}_1^{A, \mathcal{M}}[\sigma_{Q_0|F=1}], \mathcal{K}_n \circ \bar{\mathcal{E}}_n^{A, \mathcal{M}} \right. \\ \left. \circ \dots \mathcal{K}_1 \circ \bar{\mathcal{E}}_1^{A, \mathcal{M}}[\sigma_{Q_0|F=1}]\right) \leq \epsilon_n^{\text{total}}, \end{aligned}$$

(7a)

$$\text{with } \epsilon_n^{\text{total}} = \epsilon^{\text{cert}} + \sum_{j=1}^n \epsilon_j^{\text{QKD}},$$

where for notational compactness we write  $\bar{\mathcal{E}}_j^{A, \mathcal{M}} := \mathcal{E}_j^{A, \mathcal{M}} \circ \mathcal{F}_{j-1}$  (defining  $\mathcal{F}_0$  to be the identity channel on  $Q_0$ ), and  $\mathcal{K}_j$  is a channel that reads the length  $\ell$  of the keys on registers  $K_j^A K_j^B$  and replaces them with perfectly correlated uniform keys of that length, i.e., the state  $\omega^\ell$  in Eq. (4).

*Proof.* We prove the proposition by considering two exhaustive cases. First consider the case of a model  $\mathcal{M} \in \mathcal{U}_{\text{models}}$  with parameters  $\mu \notin S_{\text{robust}}$ . Then criterion 1 states that  $\Pr[F = 1] \leq \epsilon^{\text{cert}}$ . Since the trace distance between two quantum states is upper bounded by 1, we get

$$\begin{aligned} \Pr[F = 1] d\left(\bar{\mathcal{E}}_n^{A, \mathcal{M}} \circ \dots \bar{\mathcal{E}}_1^{A, \mathcal{M}}[\sigma_{Q_0|F=1}], \mathcal{K}_n \circ \bar{\mathcal{E}}_n^{A, \mathcal{M}} \right. \\ \left. \circ \dots \mathcal{K}_1 \circ \bar{\mathcal{E}}_1^{A, \mathcal{M}}[\sigma_{Q_0|F=1}]\right) \leq \epsilon^{\text{cert}} \leq \epsilon_n^{\text{total}}, \end{aligned} \quad (\text{A1})$$

as required.

Next consider the case of a model  $\mathcal{M} \in \mathcal{U}_{\text{models}}$  with parameters  $\mu \in S_{\text{robust}}$ . Observe that criterion 2 implies that for all  $j$  and any attack  $\mathcal{A}$ ,

$$\sup_{\sigma_{Q_{j-1}}^{(j)}} d\left(\bar{\mathcal{E}}_j^{A, \mathcal{M}}[\sigma_{Q_{j-1}}^{(j)}], \mathcal{K}_j \circ \bar{\mathcal{E}}_j^{A, \mathcal{M}}[\sigma_{Q_{j-1}}^{(j)}]\right) \leq \epsilon_j^{\text{QKD}}, \quad (\text{A2})$$

with the supremum being taken over all quantum states  $\sigma_{Q_{j-1}}^{(j)}$ . We use this to first show that for the  $n=2$  case, we have  $d(\bar{\mathcal{E}}_2^{A, \mathcal{M}} \circ \bar{\mathcal{E}}_1^{A, \mathcal{M}}[\sigma_{Q_0|F=1}], \mathcal{K}_2 \circ \bar{\mathcal{E}}_2^{A, \mathcal{M}} \circ \mathcal{K}_1 \circ \bar{\mathcal{E}}_1^{A, \mathcal{M}}[\sigma_{Q_0|F=1}]) \leq \epsilon_1^{\text{QKD}} + \epsilon_2^{\text{QKD}}$ , as follows:

$$\begin{aligned}
& d\left(\bar{\mathcal{E}}_2^{A,\mathcal{M}} \circ \bar{\mathcal{E}}_1^{A,\mathcal{M}} [\sigma_{Q_0|F=1}], \mathcal{K}_2 \circ \bar{\mathcal{E}}_2^{A,\mathcal{M}} \circ \mathcal{K}_1 \circ \bar{\mathcal{E}}_1^{A,\mathcal{M}} [\sigma_{Q_0|F=1}]\right) \\
&= \frac{1}{2} \left\| \left( \bar{\mathcal{E}}_2^{A,\mathcal{M}} \circ \bar{\mathcal{E}}_1^{A,\mathcal{M}} - \mathcal{K}_2 \circ \bar{\mathcal{E}}_2^{A,\mathcal{M}} \circ \mathcal{K}_1 \circ \bar{\mathcal{E}}_1^{A,\mathcal{M}} \right) [\sigma_{Q_0|F=1}] \right\|_1 \\
&= \frac{1}{2} \left\| \left( \bar{\mathcal{E}}_2^{A,\mathcal{M}} \circ \bar{\mathcal{E}}_1^{A,\mathcal{M}} - \bar{\mathcal{E}}_2^{A,\mathcal{M}} \circ \mathcal{K}_1 \circ \bar{\mathcal{E}}_1^{A,\mathcal{M}} + \bar{\mathcal{E}}_2^{A,\mathcal{M}} \circ \mathcal{K}_1 \circ \bar{\mathcal{E}}_1^{A,\mathcal{M}} - \mathcal{K}_2 \circ \bar{\mathcal{E}}_2^{A,\mathcal{M}} \circ \mathcal{K}_1 \circ \bar{\mathcal{E}}_1^{A,\mathcal{M}} \right) [\sigma_{Q_0|F=1}] \right\|_1 \\
&\leq \frac{1}{2} \left\| \bar{\mathcal{E}}_2^{A,\mathcal{M}} \circ \left( \bar{\mathcal{E}}_1^{A,\mathcal{M}} - \mathcal{K}_1 \circ \bar{\mathcal{E}}_1^{A,\mathcal{M}} \right) [\sigma_{Q_0|F=1}] \right\|_1 + \frac{1}{2} \left\| \left( \bar{\mathcal{E}}_2^{A,\mathcal{M}} - \mathcal{K}_2 \circ \bar{\mathcal{E}}_2^{A,\mathcal{M}} \right) \circ \bar{\mathcal{E}}_1^{A,\mathcal{M}} [\sigma_{Q_0|F=1}] \right\|_1, \quad (\text{A3})
\end{aligned}$$

$$\leq \frac{1}{2} \left\| \left( \bar{\mathcal{E}}_1^{A,\mathcal{M}} - \mathcal{K}_1 \circ \bar{\mathcal{E}}_1^{A,\mathcal{M}} \right) [\sigma_{Q_0|F=1}] \right\|_1 + \frac{1}{2} \left\| \left( \bar{\mathcal{E}}_2^{A,\mathcal{M}} - \mathcal{K}_2 \circ \bar{\mathcal{E}}_2^{A,\mathcal{M}} \right) \circ \bar{\mathcal{E}}_1^{A,\mathcal{M}} [\sigma_{Q_0|F=1}] \right\|_1, \quad (\text{A4})$$

$$\begin{aligned}
&\leq \sup_{\sigma_{Q_0}^{(1)}} \frac{1}{2} \left\| \left( \bar{\mathcal{E}}_1^{A,\mathcal{M}} - \mathcal{K}_1 \circ \bar{\mathcal{E}}_1^{A,\mathcal{M}} \right) [\sigma_{Q_0}^{(1)}] \right\|_1 + \sup_{\sigma_{Q_1}^{(2)}} \frac{1}{2} \left\| \left( \bar{\mathcal{E}}_2^{A,\mathcal{M}} - \mathcal{K}_2 \circ \bar{\mathcal{E}}_2^{A,\mathcal{M}} \right) [\sigma_{Q_1}^{(2)}] \right\|_1 \\
&\leq \epsilon_1^{\text{QKD}} + \epsilon_2^{\text{QKD}}, \quad (\text{A5})
\end{aligned}$$

where Eq. (A3) follows from the triangle inequality of the 1-norm, Eq. (A4) follows from the fact that the 1-norm is nonincreasing under the action of channels, and Eq. (A5) follows from Eq. (A2). The general case

$$\begin{aligned}
& d\left(\bar{\mathcal{E}}_n^{A,\mathcal{M}} \circ \dots \circ \bar{\mathcal{E}}_1^{A,\mathcal{M}} [\sigma_{Q_0|F=1}], \mathcal{K}_n \circ \bar{\mathcal{E}}_n^{A,\mathcal{M}} \right. \\
&\quad \left. \circ \dots \circ \mathcal{K}_1 \circ \bar{\mathcal{E}}_1^{A,\mathcal{M}} [\sigma_{Q_0|F=1}]\right) \leq \sum_{j=1}^n \epsilon_j^{\text{QKD}} \leq \epsilon_n^{\text{total}}
\end{aligned}$$

follows by iterating the same argument  $n$  times, yielding the claimed result.

Note that the above proof also works identically if instead we set  $\epsilon_n^{\text{total}} = \max \left\{ \epsilon^{\text{cert}}, \sum_{j=1}^n \epsilon_j^{\text{QKD}} \right\}$ , as claimed. ■

## APPENDIX B: ADAPTIVE PROTOCOLS BASED ON “CHARACTERIZATION” RATHER THAN “CERTIFICATION”

The results in this appendix are more succinctly described by introducing the concept of confidence regions, which include confidence intervals as a special case:

*Definition B1.* Consider a probability distribution described by some tuple of parameters  $\mu$ . A **confidence region** for  $\mu$  at confidence level  $\epsilon$  is a (set-valued) random variable  $S^{\text{conf}}$  such that  $S^{\text{conf}}$  contains  $\mu$  with probability at least  $1 - \epsilon$ .

With this in mind, we can more rigorously describe the process of how a characterization procedure should “estimate” the device parameters using confidence regions, followed by performing QKD protocol instances that explicitly depend on the obtained confidence regions.

*Proposition B1.* Consider a device which undergoes a characterization procedure followed by up to  $N_{\text{max}}$  QKD protocol instances (which can depend on the output of the characterization procedure), such that the following criteria are satisfied:

- (1) For any device model  $\mathcal{M} \in \mathcal{U}_{\text{models}}$ , with parameters  $\mu$ , the characterization procedure outputs a confidence region for  $\mu$  at confidence level  $\epsilon^{\text{cert}}$ . We shall denote this confidence region as  $\Lambda$ , and in the subsequent statements we shall also freely treat it as a classical register, via the standard identification between random variables and classical registers [37].
- (2) Conditioned on the value  $\lambda$  stored in  $\Lambda$ , the QKD protocol instances satisfy the following criterion: for any device model  $\mathcal{M}$  with parameters  $\mu \in \lambda$ , the  $j$ th QKD protocol instance is  $\epsilon_j^{\text{QKD}}$ -secure under device model  $\mathcal{M}$  (Definition 1). For full clarity, in the subsequent expressions we will instead write the channels in that definition as  $\mathcal{E}_j^{A,\mathcal{M},\lambda}$ , explicitly denoting their dependence on  $\lambda$  (since the protocols being performed can depend on  $\lambda$ ).

Then for any  $n \in \{1, 2, \dots, N_{\text{max}}\}$ , any state  $\sigma_{Q_0\Lambda}$  at the start of the first protocol instance (after the characterization procedure), any attack  $\mathcal{A}$  and device model  $\mathcal{M} \in \mathcal{U}_{\text{models}}$ ,

and any sequence of channels  $\mathcal{F}_j : K_j^A K_j^B E_j \rightarrow Q_j$  for  $j \in \{1, 2, \dots, n-1\}$ , we have

$$\sum_{\lambda} \Pr[\Lambda = \lambda] d\left(\bar{\mathcal{E}}_n^{A, \mathcal{M}, \lambda} \circ \dots \circ \bar{\mathcal{E}}_1^{A, \mathcal{M}, \lambda} [\sigma_{Q_0|\lambda}], \mathcal{K}_n \circ \bar{\mathcal{E}}_n^{A, \mathcal{M}, \lambda} \circ \dots \circ \mathcal{K}_1 \circ \bar{\mathcal{E}}_1^{A, \mathcal{M}, \lambda} [\sigma_{Q_0|\lambda}]\right) \leq \epsilon_n^{\text{total}},$$

with  $\epsilon_n^{\text{total}} = \epsilon^{\text{cert}} + \sum_{j=1}^n \epsilon_j^{\text{QKD}}, \quad (\text{B1})$

where for notational compactness we write  $\bar{\mathcal{E}}_j^{A, \mathcal{M}, \lambda} := \mathcal{E}_j^{A, \mathcal{M}, \lambda} \circ \mathcal{F}_{j-1}$  (defining  $\mathcal{F}_0$  to be the identity channel on  $Q_0$ ), and  $\mathcal{K}_j$  is a channel that reads the length  $\ell$  of the keys on registers  $K_j^A K_j^B$  and replaces them with perfectly correlated uniform keys of that length, i.e., the state  $\omega^\ell$  in Eq. (4).

Note that the above description implicitly incorporates as a special case the possibility that some outcome values from the characterization procedure lead to the devices being “rejected” for use in QKD, simply by having the subsequent “QKD protocols” all output zero-length keys (which trivially fulfill the  $\epsilon_j^{\text{QKD}}$ -security criterion). Also, similar to the discussion in Sec. II C, the bound Eq. (B1) can be equivalently rewritten as

$$d\left(\rho_{K_n^A K_n^B E_n \Lambda}, \tilde{\rho}_{K_n^A K_n^B E_n \Lambda}\right) \leq \epsilon_n^{\text{total}}, \quad (\text{B2})$$

where

$$\rho_{K_n^A K_n^B E_n \Lambda} = \sum_{\lambda} \Pr[\Lambda = \lambda] \bar{\mathcal{E}}_n^{A, \mathcal{M}, \lambda} \circ \dots \circ \bar{\mathcal{E}}_1^{A, \mathcal{M}, \lambda} [\sigma_{Q_0|\lambda}] \otimes |\lambda\rangle\langle\lambda|, \quad (\text{B3})$$

$$\tilde{\rho}_{K_n^A K_n^B E_n \Lambda} = \sum_{\lambda} \Pr[\Lambda = \lambda] \mathcal{K}_n \circ \bar{\mathcal{E}}_n^{A, \mathcal{M}, \lambda} \circ \dots \circ \mathcal{K}_1 \circ \bar{\mathcal{E}}_1^{A, \mathcal{M}, \lambda} [\sigma_{Q_0|\lambda}] \otimes |\lambda\rangle\langle\lambda|. \quad (\text{B4})$$

Viewed this way, Proposition B1 is essentially stating that the *mixture* of states produced by averaging over the outcomes of the characterization procedure (i.e.,  $\rho_{K_n^A K_n^B E_n \Lambda}$ ) is close in trace distance to a somewhat more “ideal” state  $\tilde{\rho}_{K_n^A K_n^B E_n \Lambda}$ , which is similar to  $\rho_{K_n^A K_n^B E_n \Lambda}$  except that the key registers  $K_j^A K_j^B$  are replaced by ideal keys of the same length immediately after each protocol instance.

We now present the proof of Proposition B1, which uses essentially the same ideas as the variable-length security proof in Ref. [38].

*Proof.* Consider any device model  $\mathcal{M} \in \mathcal{U}_{\text{models}}$ , with corresponding parameters  $\mu$ . Partition the summation in

Eq. (B1) into summations over the  $\lambda$  values such that  $\mu \notin \lambda$ , and the  $\lambda$  values such that  $\mu \in \lambda$ . For the former, we have

$$\sum_{\lambda \text{ s.t. } \mu \notin \lambda} \Pr[\Lambda = \lambda] d\left(\bar{\mathcal{E}}_n^{A, \mathcal{M}, \lambda} \circ \dots \circ \bar{\mathcal{E}}_1^{A, \mathcal{M}, \lambda} [\sigma_{Q_0|\lambda}], \mathcal{K}_n \circ \bar{\mathcal{E}}_n^{A, \mathcal{M}, \lambda} \circ \dots \circ \mathcal{K}_1 \circ \bar{\mathcal{E}}_1^{A, \mathcal{M}, \lambda} [\sigma_{Q_0|\lambda}]\right) \leq \sum_{\lambda \text{ s.t. } \mu \notin \lambda} \Pr[\Lambda = \lambda] \leq \epsilon^{\text{cert}}, \quad (\text{B5})$$

where the last line holds by recalling that by hypothesis,  $\Lambda$  is a confidence region for  $\mu$  at confidence level  $\epsilon^{\text{cert}}$ . For the latter, we have

$$\sum_{\lambda \text{ s.t. } \mu \in \lambda} \Pr[\Lambda = \lambda] d\left(\bar{\mathcal{E}}_n^{A, \mathcal{M}, \lambda} \circ \dots \circ \bar{\mathcal{E}}_1^{A, \mathcal{M}, \lambda} [\sigma_{Q_0|\lambda}], \mathcal{K}_n \circ \bar{\mathcal{E}}_n^{A, \mathcal{M}, \lambda} \circ \dots \circ \mathcal{K}_1 \circ \bar{\mathcal{E}}_1^{A, \mathcal{M}, \lambda} [\sigma_{Q_0|\lambda}]\right) \leq \sum_{j=1}^n \epsilon_j^{\text{QKD}}, \quad (\text{B6})$$

by the same reasoning as in Appendix A, noting that for each  $\lambda$  value in the sum, the trace-distance term is evaluated for QKD protocol instances that are  $\epsilon_j^{\text{QKD}}$ -secure under this device model  $\mathcal{M}$  by hypothesis (since  $\mu \in \lambda$ ). Adding these bounds yields the claimed final result. ■

## APPENDIX C: COMPOSABILITY UNDER THE ABSTRACT CRYPTOGRAPHY FRAMEWORK

This appendix is written assuming reader familiarity with the AC framework presented in Ref. [8]. (An earlier version of the framework can be found in Ref. [20], but there are some differences in the formalism.)

We shall discuss what can be stated in the AC framework (with respect to the version described in Ref. [8]) regarding an “overall” protocol consisting of a certification procedure that either approves or rejects a device, followed by  $N_{\text{max}}$  QKD protocol instances if it approves. For this discussion, we shall require that when the certification procedure rejects, the devices are never used for any QKD protocols.

*Remark C1.* Ideally, one might wish to instead prove some sort of composable security statement regarding the certification procedure as a “standalone” protocol, then combine it with security statements for the individual QKD protocol instances. However, it appears difficult to define a suitable notion of an ideal resource to construct from the certification procedure, and we were unable to satisfactorily resolve this question. We hence leave this possibility for future work.

We first recall that in the AC framework, one aims to make statements that a protocol constructs some resource from another resource within some  $\epsilon$ . To connect this to the formalism of device models described in Sec. II, let us impose the following condition on device models, focusing on prepare-and-measure QKD (the modifications for entanglement-based or measurement-device-independent QKD are straightforward, by changing which parties hold preparation and/or measurement devices). We note that the statement of this condition is slightly informal, as it is not entirely clear whether the resources outlined within it can be completely formalized within the existing AC framework—however, we believe that some version of such formalization should be possible, and hence the remainder of this appendix will proceed under the assumption that this can be done.

*Condition 2.* We suppose that for every device model  $\mathcal{M} \in \mathcal{U}_{\text{models}}$ , it is possible to define a corresponding pair of resources in the AC framework, namely, a state preparation device for Alice that can repeatedly accept a classical input and generate a state based on the input (and possibly some memory it stores over time), and a measurement device for Bob that can repeatedly accept a classical input and a quantum system, then outputs the outcome of some measurement on that quantum system depending on the classical input (and possibly some memory it stores over time) [39]. We additionally require these resources to have an operation that “clears” any memory in the devices.

Note that at the practical level, the “memory clearing” operation may be as simple as leaving the devices unused for some period of time, under the physical assumption that memory effects are sufficiently weak that they will dissipate given enough time. We impose this requirement simply as an approach to address the issues discussed in Remark 2, though it might be possible to explore in future work whether there are other ways to overcome the issue.

*Remark C2.* As a minor variant of Condition C1, it might be possible to instead attempt to simply *identify* the models with those resources, since at the typical level of abstraction in the AC framework, this would also suffice to capture all properties relevant to the subsequent discussion. Yet another possible variant might be to instead attempt to define a somewhat more general resource that can accept as input some specification of a device model  $\mathcal{M} \in \mathcal{U}_{\text{models}}$  from Eve, and then gives Alice and Bob various state preparation and measurement functionalities depending on this specification. However, we avoid doing so here in order to avoid needing to formalize the models  $\mathcal{M} \in \mathcal{U}_{\text{models}}$  in some fashion compatible with being inputs to a resource in the AC framework; instead, we couch our subsequent security statements in terms of a quantifier over all models  $\mathcal{M} \in \mathcal{U}_{\text{models}}$ . (This is analogous to the discussion in Ref. [8, Sec. VI.D] for device-independent QKD,

regarding whether the security statement is quantified over all device behaviors, as compared to allowing the devices to be “universal computers” that can have their behavior determined by an input from Eve.)

Having established this, we now need to specify the starting resources and the resource we aim to construct, for both honest and dishonest Eve [8]. Following the above discussion, we view these resources as being implicitly defined as functions of models  $\mathcal{M} \in \mathcal{U}_{\text{models}}$ , and we will later present our security statements in terms of a quantifier over  $\mathcal{M} \in \mathcal{U}_{\text{models}}$ . First considering dishonest Eve, we take the starting resources  $\mathcal{R}_{\text{ini}}$  to be, similarly to Ref. [8]:

- (1) An insecure quantum channel between Alice and Bob.
- (2) An authenticated classical channel between Alice and Bob.
- (3) The resource described above in Condition C1 (implicitly depending on the model  $\mathcal{M}$ ).

As for the resource we aim to construct, we define it as follows (essentially a multiple-key version of the secret key resource in Ref. [8, Fig. 18]), given an increasing sequence of times  $t_j$  for  $j \in \{1, 2, \dots, N_{\text{max}}\}$ :

- (1)  $\mathcal{R}_{\text{ideal}}$  is a resource where at each time  $t_j$ , it accepts a classical input value  $\ell_j \in \mathbb{N}$  from Eve’s interface, and then outputs perfect shared secret keys of length  $\ell_j$  to Alice and Bob’s interfaces.

With this in mind, we can define the protocol we are considering, performed using the starting resources  $\mathcal{R}_{\text{ini}}$ . In order to minimize difficulties regarding composability, we shall suppose that the devices (as defined in Condition C1) are not used for any purposes other than this protocol—it might be possible to do so if the device memories are cleared appropriately, but we shall not consider this in depth here.

- (1) First, a certification procedure is performed on  $\mathcal{R}_{\text{ini}}$ , generating either an approve or reject decision that is used to determine the next action.
- (2) If the certification procedure approves, the resources  $\mathcal{R}_{\text{ini}}$  are then used for  $N_{\text{max}}$  QKD protocol instances that output keys at the times  $t_j$ , under the constraint after each QKD protocol instance, *the device memories are cleared before they are used again*. If the certification procedure rejects, the devices are discarded, which we shall formalize by having the protocol output zero-length keys at all the times  $t_j$ .

In the AC framework as described in Ref. [8], the above protocol would be a converter applied to Alice and Bob’s interfaces in  $\mathcal{R}_{\text{ini}}$ , with the inner interface connecting to  $\mathcal{R}_{\text{ini}}$ , and the outer interface producing keys at times  $t_j$ .

(Note that the approve/reject decision from the certification step is to be interpreted as an “internal process” of the protocol and is not released at any of the interfaces.)

With this, we now state our main claim under the assumption that Condition C1 can be properly formalized, followed by an outline of how it would be proven given a suitable formalization of that condition. Since distinguishability in the AC framework is always to be considered with respect to some class of distinguishers, we phrase our following result as being implicitly in terms of the class of distinguishers. Also, as discussed in Remark C2, we simply state the result with a quantifier over all models  $\mathcal{M} \in \mathcal{U}_{\text{models}}$ , rather than attempting to build the dependence on the model  $\mathcal{M}$  “into” the resource in some fashion.

*Proposition C1.* Consider any device model  $\mathcal{M} \in \mathcal{U}_{\text{models}}$  satisfying Condition C1 and take any increasing sequence of times  $t_j$ . Let  $\mathcal{R}_{\text{ini}}$  and  $\mathcal{R}_{\text{ideal}}$  be the resources defined above, for the case of dishonest Eve, and consider the protocol defined above. Suppose furthermore that this protocol satisfies the two criteria in Proposition 1 for some class of attacks  $\mathcal{A}$ , and define

$$\epsilon^{\text{total}} := \epsilon^{\text{cert}} + \sum_{j=1}^{N_{\text{max}}} \epsilon_j^{\text{QKD}}. \quad (\text{C1})$$

Then this protocol constructs the resource  $\mathcal{R}_{\text{ideal}}$  from  $\mathcal{R}_{\text{ini}}$  within  $\epsilon^{\text{total}}$  in the sense of Ref. [8, Definition 1], as long as the class of attacks  $\mathcal{A}$  is large enough to describe all distinguisher strategies considered in that definition.

*Proof outline.* The idea is basically identical to Ref. [8, Sec. III.B.2]; here we outline the main structure. By definition, we need to construct a suitable simulator  $\sigma_E$  to attach to Eve’s interface in  $\mathcal{R}_{\text{ideal}}$ , such that  $\mathcal{R}_{\text{ideal}}\sigma_E$  is  $\epsilon^{\text{total}}$ -indistinguishable from  $\pi_{AB}\mathcal{R}_{\text{ini}}$ , where  $\pi_{AB}$  denotes the protocol described above performed by Alice and Bob. We follow a basically identical construction to Ref. [8, Sec. III.B.2] for this simulator. Specifically, we define it to simply run an “internal instance” of  $\pi_{AB}\mathcal{R}_{\text{ini}}$ , so its “outer” interface interacts accordingly with Eve. As for how it interacts on its “inner” interface (i.e., the Eve interface on  $\mathcal{R}_{\text{ideal}}$ ), we define it as follows: for each time  $t_j$ , if the certification step rejected in the “internal instance” of  $\pi_{AB}\mathcal{R}_{\text{ini}}$ , then the simulator simply sets  $\ell_j = 0$ , otherwise (if the certification step approved) the simulator sets  $\ell_j$  equal to the length of the key produced from the  $j$ th QKD protocol in the “internal instance” of  $\pi_{AB}\mathcal{R}_{\text{ini}}$ .

Observe that  $\pi_{AB}\mathcal{R}_{\text{ini}}$  is perfectly indistinguishable from  $\mathcal{R}_{\text{ideal}}\sigma_E$  up until time  $t_1$  (as also noted in Ref. [8, Sec. III.B.2]); there are absolutely no differences in their behavior up to that point. At time  $t_1$ , the first potential discrepancy occurs, since  $\pi_{AB}\mathcal{R}_{\text{ini}}$  outputs the actual key at

Alice and Bob’s interfaces, whereas  $\mathcal{R}_{\text{ideal}}\sigma_E$  outputs a “replaced” version of the key. Let us write the registers held by the distinguisher at that point as  $K_1^A K_1^B E_1$ , where  $E_1$  denotes all registers held by the distinguisher just before  $K_1^A K_1^B$  were released. With this, observe that the states held by the distinguisher at this time in the two scenarios are, respectively, the reduced states (on  $K_1^A K_1^B E_1$ ) of states with the form

$$\begin{aligned} \rho_{K_1^A K_1^B E_1 F} &= \Pr[F = 0] \rho_{K_1^A K_1^B E_1 | F=0} \otimes |0\rangle\langle 0|_F \\ &\quad + \Pr[F = 1] \mathcal{E}_1^{\mathcal{A}, \mathcal{M}} [\sigma_{Q_0 | F=1}] \otimes |1\rangle\langle 1|_F, \\ \tilde{\rho}_{K_1^A K_1^B E_1 F} &= \mathcal{K}_1 \left[ \rho_{K_1^A K_1^B E_1 F} \right] \\ &= \Pr[F = 0] \rho_{K_1^A K_1^B E_1 | F=0} \otimes |0\rangle\langle 0|_F \\ &\quad + \Pr[F = 1] \mathcal{K}_1 \circ \mathcal{E}_1^{\mathcal{A}, \mathcal{M}} [\sigma_{Q_0 | F=1}] \otimes |1\rangle\langle 1|_F, \end{aligned} \quad (\text{C2})$$

where  $\mathcal{K}_1, \mathcal{E}_1^{\mathcal{A}, \mathcal{M}}, \sigma_{Q_0 F}$  are as in the Proposition 1 notation (for some  $\mathcal{A}$  corresponding to the distinguisher’s strategy), and to obtain the last equality we recall that in the conditional state  $\rho_{K_1^A K_1^B E_1 | F=0}$ , the key registers  $K_1^A K_1^B$  are fixed to zero-length keys, and therefore  $\mathcal{K}_1$  leaves it unchanged. Furthermore, by Proposition 1 we have

$$\begin{aligned} \Pr[F = 1] d \left( \mathcal{E}_1^{\mathcal{A}, \mathcal{M}} [\sigma_{Q_0 | F=1}], \mathcal{K}_1 \circ \mathcal{E}_1^{\mathcal{A}, \mathcal{M}} [\sigma_{Q_0 | F=1}] \right) \\ \leq \epsilon^{\text{cert}} + \epsilon_1^{\text{QKD}}, \end{aligned} \quad (\text{C3})$$

and thus we have

$$d \left( \rho_{K_1^A K_1^B E_1 F}, \tilde{\rho}_{K_1^A K_1^B E_1 F} \right) \leq \epsilon^{\text{cert}} + \epsilon_1^{\text{QKD}} \leq \epsilon^{\text{total}}, \quad (\text{C4})$$

so the states held by the distinguisher at time  $t_1$  (on  $K_1^A K_1^B E_1$ ) are  $\epsilon^{\text{total}}$ -indistinguishable.

Proceeding on to later times, recall we suppose that after the first QKD protocol finishes, the devices have their memories cleared before they are used again in any form. Therefore, without loss of generality we can suppose any memory registers in the devices are traced out immediately after  $t_1$  (since they cannot have any operational impact if the devices are not used before these memories are cleared). Under this picture, observe that the registers  $K_1^A K_1^B E_1 F$  suffice to capture the *entire* state of the resources and the distinguisher at that point. As we have shown above in Eq. (C4), the states on these registers are  $\epsilon^{\text{total}}$ -indistinguishable in the two scenarios. From this, we can conclude that up until time  $t_2$  (when the keys from the second QKD protocol are produced), the distinguishing advantage remains at most  $\epsilon^{\text{total}}$ , because all further states held by the distinguisher up until that time can then be described via some completely positive, trace-preserving (CPTP) map acting on those states, which does not increase trace distance.

Then for the states at time  $t_2$ , we apply the same argument structure as above to conclude the trace distance is again at most  $\epsilon^{\text{cert}} + \epsilon_1^{\text{QKD}} + \epsilon_2^{\text{QKD}} \leq \epsilon^{\text{total}}$ . Continuing similarly for the remainder of the protocol, we conclude  $\pi_{AB}\mathcal{R}_{\text{ini}}$  is and  $\mathcal{R}_{\text{ideal}}\sigma_E$  are always  $\epsilon^{\text{total}}$ -indistinguishable, yielding the desired result. ■

*Remark C3.* The “adaptive” device-characterization version described in Appendix B would also construct the same resource within  $\epsilon^{\text{total}}$  (under the same informal conditions discussed above), via an analogous argument. Specifically, rather than having the simulator set  $\ell_j = 0$  for all  $j$  whenever the certification step aborts, we simply have it supply the length  $\ell_j$  of the key produced from the  $j$ th QKD protocol in the “internal instance” of  $\pi_{AB}\mathcal{R}_{\text{ini}}$  (in this version, these lengths will also implicitly depend on the outcome from the certification step). The rest of the proof outline carries through in the same manner.

We now turn to the case where Eve is honest, as is required for a full analysis under the AC framework. The following discussion is to be understood to apply both for the device certification procedure in Sec. II and the “adaptive” device-characterization version described in Appendix B, as the same claims apply to both. Again following Ref. [8], we take the starting resources to be the same as above, except that the insecure quantum channel outputs nothing on Eve’s interface and is taken to have some fixed “honest” noise level between Alice and Bob. Similarly, the resource we aim to construct is one that outputs nothing at Eve’s interface, and at each time  $t_j$ , it outputs a mixture of perfect shared secret keys of different lengths to Alice and Bob’s interfaces, such that the distribution of lengths matches that which would be obtained from the honest behavior of the insecure quantum channel. The analysis of this case is fairly straightforward, following the same arguments in Ref. [8, Sec. III.B.4]; we again conclude that for any  $\mathcal{M} \in \mathcal{U}_{\text{models}}$  the protocol indeed constructs the claimed resource within  $\epsilon^{\text{total}}$ . Note however that there might be many  $\mathcal{M} \in \mathcal{U}_{\text{models}}$  such that the resulting constructed resource in this case is not of practical interest (e.g., one that almost never produces any keys, because for that model  $\mathcal{M}$  the certification procedure almost always rejects). Hence if desired, we can instead focus only on a specific model  $\mathcal{M}_{\text{hon}} \in \mathcal{U}_{\text{models}}$  corresponding to some “honest” device model that yields a practically interesting final resource.

[1] S. Nahar, T. Upadhyaya, and N. Lütkenhaus, Imperfect phase randomization and generalized decoy-state quantum key distribution, *Phys. Rev. Appl.* **20**, 064031 (2023).

[2] We choose to use the terms “approve”/“reject” in order to avoid confusion with the “accept”/“abort” decisions in QKD protocols themselves.

[3] R. Renner, *Security of Quantum Key Distribution* (ETH Zürich, 2005).

[4] D. Tupkary, E. Y. Z. Tan, S. Nahar, L. Kamin, and N. Lütkenhaus, QKD security proofs for decoy-state BB84: Protocol variations, proof techniques, gaps and limitations, [arXiv:2502.10340](https://arxiv.org/abs/2502.10340) [quant-ph].

[5] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information. 10th Anniversary Edition* (Cambridge University, New York, 2010).

[6] M. Ben-Or, M. Horodecki, D. W. Leung, D. Mayers, and J. Oppenheim, *Theory of Cryptography* (Springer, Berlin, Heidelberg, 2005), pp. 386–406.

[7] M. Hayashi and T. Tsurumaru, Concise and tight security analysis of the Bennett–Brassard 1984 protocol with finite key lengths, *New J. Phys.* **14**, 093014 (2012).

[8] C. Portmann and R. Renner, Security in quantum cryptography, *Rev. Mod. Phys.* **94**, 025008 (2022).

[9] G. Chiribella, G. M. D’Ariano, and P. Perinotti, Theoretical framework for quantum networks, *Phys. Rev. A* **80**, 022339 (2009).

[10] C. Portmann, C. Matt, U. Maurer, R. Renner, and B. Tackmann, Causal boxes: Quantum information-processing systems closed under composition, *IEEE Trans. Inf. Theory* **63**, 3277 (2017).

[11] More generally, one might want to consider the possibility of some form of “mid-lifespan” certification of the QKD device. However, we leave a detailed analysis of this to be addressed in future work—for this work, we shall simply take the perspective that this could potentially be viewed as a “fresh” certification of the device for another number of QKD protocol instances.

[12] An alternative formalism could be to specify one attack  $\mathcal{A}_j$  for each instance, but that can simply be considered a special case of the formalism we describe here, for instance by setting  $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2, \dots)$ .

[13] C. Portmann, Key recycling in authentication, *IEEE Trans. Inf. Theory* **60**, 4383 (2014).

[14] J. Barrett, R. Colbeck, and A. Kent, Memory attacks on device-independent quantum cryptography, *Phys. Rev. Lett.* **110**, 010503 (2013).

[15] M. Curty and H.-K. Lo, Foiling covert channels and malicious classical post-processing units in quantum key distribution, *npj Quantum Inf.* **5**, 1 (2019).

[16] F. G. Lacerda, J. M. Renes, and R. Renner, Classical leakage resilience from fault-tolerant quantum computation, *J. Cryptol.* **32**, 1071 (2019).

[17] M. McKague and L. Sheridan, *Information Theoretic Security*, edited by C. Padró, Lecture Notes in Computer Science (Springer International Publishing, Cham, 2014), pp. 122–141.

[18] In this example we choose to specify the parameter to be an *upper bound* on the dark count rate rather than “the” dark count rate, since the latter may not be a well-defined single number if the detector behavior is not IID.

[19] In fact, from the proof it can be seen that the upper bound in Eq. (7) can be sharpened to  $\epsilon_n^{\text{total}} = \max \left\{ \epsilon^{\text{cert}}, \sum_{j=1}^n \epsilon^{\text{QKD}} \right\}$ .

- [20] U. Maurer and R. Renner, *The Second Symposium on Innovations in Computer Science, ICS 2011* (Tsinghua University, 2011), pp. 1–21.
- [21] A. Lenart, T. Islam, S. Sivasankaran, P. Neilson, B. Hiding, D. K. L. Oi, and A. Ling, Comparing a radiation damage model for avalanche photodiodes through in-situ observation of cubesat based devices, *Commun. Phys.* **8**, 118 (2025).
- [22] A. Vakhitov, V. Makarov, and D. R. Hjelme, Large pulse attack as a method of conventional optical eavesdropping in quantum cryptography, *J. Mod. Opt.* **48**, 2023 (2001).
- [23] M. Lucamarini, I. Choi, M. B. Ward, J. F. Dynes, Z. L. Yuan, and A. J. Shields, Practical security bounds against the Trojan-horse attack in quantum key distribution, *Phys. Rev. X* **5**, 031030 (2015).
- [24] A. Arqand, T. Metger, and E. Y. Z. Tan, Mutual information chain rules for security proofs robust against device imperfections, [arXiv:2407.20396](https://arxiv.org/abs/2407.20396) [quant-ph].
- [25] G. Currás-Lorenzo, S. Nahar, N. Lütkenhaus, K. Tamaki, and M. Curty, Security of quantum key distribution with imperfect phase randomisation, *Quantum Sci. Technol.* **9**, 015025 (2024).
- [26] G. Currás-Lorenzo, M. Pereira, G. Kato, M. Curty, and K. Tamaki, Security of high-speed quantum key distribution with imperfect sources, *Quantum* **3**, 525 (2025).
- [27] L. Kamin, J. Burniston, and E. Y. Z. Tan, Rényi security framework against coherent attacks applied to decoy-state QKD, [arXiv:2504.12248](https://arxiv.org/abs/2504.12248) [quant-ph].
- [28] A. Marwah and F. Dupuis, Proving security of BB84 under source correlations, [arXiv:2402.12346](https://arxiv.org/abs/2402.12346) [quant-ph].
- [29] S. Nahar, D. Tupkary, and N. Lütkenhaus, Imperfect detectors for adversarial tasks with applications to quantum key distribution, *Quantum* **10**, 2044 (2026).
- [30] X. Sixto, Á. Navarrete, M. Pereira, G. Currás-Lorenzo, K. Tamaki, and M. Curty, Quantum key distribution with imperfectly isolated devices, *Quantum Sci. Technol.* **10**, 035034 (2025).
- [31] D. Tupkary, S. Nahar, P. Sinha, and N. Lütkenhaus, Phase error rate estimation in QKD with imperfect detectors, *Quantum* **9**, 1937 (2025).
- [32] V. Zapatero, Á. Navarrete, K. Tamaki, and M. Curty, Security of quantum key distribution with intensity correlations, *Quantum* **5**, 602 (2021).
- [33] S. Nahar, A proof-technique-independent framework for detector imperfections in QKD, Ph.D. thesis, University of Waterloo, 2026.
- [34] ETSI, Quantum Key Distribution (QKD); Component characterization: characterizing optical components for QKD systems (2016).
- [35] W. Hoeffding, Probability inequalities for sums of bounded random variables, *J. Am. Stat. Assoc.* **58**, 13 (1963).
- [36] To see this: note that this event is equivalent to *all*  $f_j + \delta$  values lying below the true maximum dark count rate. Now focusing on the detector with the highest true dark count rate, this implies, in particular, that the  $f_j + \delta$  value for that detector lies below its true dark count rate, and thus this event has probability at most  $\epsilon^{\text{cert}}$ .
- [37] For the purposes of this work, we shall suppose that  $\Lambda$  can only take values in a countable set, to avoid technical issues with uncountably infinite sets. This should not be a significant restriction in practice, since real-life characterization procedures can typically only output values belonging to some countable set (for instance, finite-precision numbers of some form), even when estimating real-valued parameters.
- [38] D. Tupkary, E. Y.-Z. Tan, and N. Lütkenhaus, Security proof for variable-length quantum key distribution, *Phys. Rev. Res.* **6**, 023002 (2024).
- [39] Strictly speaking, if the certification procedure involves interacting with these devices in some more general fashion, then these resources also need to be further generalized in order to allow such operations, but the overall conclusion remains.